



BIROn - Birkbeck Institutional Research Online

Ahlstrom-Vij, Kristoffer (2021) Do we live in a 'Post-Truth' era? *Political Studies* , ISSN 0032-3217.

Downloaded from: <https://eprints.bbk.ac.uk/id/eprint/44894/>

Usage Guidelines:

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

DO WE LIVE IN A “POST-TRUTH” ERA?

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

Abstract: Have we entered a “post-truth” era? The present paper attempts to answer this question by (a) offering an explication of the notion of “post-truth” from recent discussions; (b) deriving a testable implication from that explication, to the effect that we should expect to see decreasing *information effects*—i.e., differences between actual preferences and estimated, fully informed preferences—on central political issues over time; and then (c) putting the relevant narrative to the test by way of counterfactual modelling, using election year data for the period of 2004-2016 from the American National Election Studies’ (ANES) Times Series Study. The implication in question turns out to be consistent with the data: at least in a US context, we do see evidence of a decrease in information effects on key, political issues—immigration, same-sex adoption, and gun laws, in particular—in the period 2004 to 2016. This offers some novel, empirical evidence for the “post-truth” narrative.

1. What is “Post-Truth”?

It has been suggested that we have entered a “post-truth” era. But what exactly does this mean? The Oxford Dictionaries famously declared “post-truth” the word of the year in 2016, defining it as ‘relating to or denoting circumstances in which objective facts are less influential in shaping public opinion than appeals to emotion and personal belief’. Several subsequent commentators have taken this definition as their starting point. For example, Lee McIntyre (2018) comments on it by suggesting that ‘the prefix “post” is meant to indicate not so much the idea that we are “past” truth [...] but in the sense that truth has been eclipsed—that it is irrelevant’ (5). Eclipsed by what? According to Oxford Dictionaries, ‘emotion and personal belief’. As Matthew D’Ancona (2017) puts it, the ‘essence of Post-Truth culture [...] depends not upon evidence, but on feeling’ (68).¹

The most promising gloss on these claims interprets them in the context of an influential account of political preference and behaviour, on which we tend not to concern ourselves with policy details,

¹ There are other writers, besides McIntyre and D’Ancona, who pursue a “post-truth” narrative, including Davis (2017) and Ball (2017). However, Davis (2017) makes clear that he is really concerned with a long history of lies and deceit, or with bullshit in the broadest sense (far broader than in Frankfurt 2005); and Ball (2017) is not conceptually focused enough to warrant any attention beyond what will be given to McIntyre and D’Ancona.

and instead form political preferences and make political decisions primarily by latching on to commitments that jibe with our partisan identities. Consider, for example, Christopher Achen and Larry Bartels' recent and highly influential book, *Democracy for Realists* (2016). On their picture, 'political preferences and judgments that look and feel like the bases of partisanship and voting behaviour are, in reality, often *consequences* of party and group loyalties' (268). This is in keeping with a long tradition in the study of political behaviour. Already in the classic Columbia Studies, we find the suggestion that 'political preferences may be [...] considered analogous to cultural tastes [...] [having] their origin in ethnic, sectional, class, and family traditions' (Berelson *et al.* 1954: 310-311). Similarly, the subsequent and highly influential Michigan Model famously maintains that political preferences are greatly influenced by partisan loyalties formed early in life (Campbell *et al.* 1960).²

So, perhaps McIntyre, D'Ancona, and others discussing "post-truth" in similar terms have in mind the idea that, when it comes to politics, it is not that we sit down and make a calculation of the cost and benefits of the relevant policies—"the facts"—and then form our preferences. Instead, we start from the emotional attachments that flow from our group-identities—the 'emotions' in the Oxford definition or 'feelings' in D'Ancona's gloss—and then opt for the policies that come with being the type of person we are. But note two things. First, the claim would have to be that the type of tendencies identified by political scientists in the Michigan Model tradition have become significantly more *pronounced* in recent years, potentially swamping any consideration for 'the facts' in political preference formation. Second, the claim cannot be that what is true is *irrelevant* to our political attitudes, as McIntyre seems to suggest. On the contrary, *what is true about what commitments happen to jibe with what partisan identities* is highly relevant. A distinction is in order: for any candidate political option, there are *identity cues*, relating to whether people of certain identity groups go for that option; and then there is everything else that is true about that option, and specifically its empirical nature (e.g., what does the option involve, in terms of political policies and strategies?) and implications (e.g., what are the likely consequences of implementing that option?). Call the latter *the empirical substance of politics*.

This distinction is in line with the idea, put forward recently by Ian MacMullen (2020), that politics is *factual*, as opposed to *post-factual*, 'to the extent that political decisions are informed by and responsive to relevant empirical conditions' (98), if those conditions are spelled out in terms of what is referred to as the empirical substance of politics above. Moreover, note that statements about "post-truth" are often accompanied by a claim about *timing*. For example, McIntyre suggests that, '[a]lthough the Brexit vote and the US presidential election may seem inextricably tied up with post-truth, neither was the cause of it—they were the result' (2018: 15). For his part, D'Ancona claims that '2016 was the

² Similar suggestions can be found in recent work in cognitive science (e.g., Sloman and Fernbach 2017). Critics of the Michigan Model have tended to suggest that the public is able to rely on 'cues and shortcuts' and, as such, 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.

year that definitely launched the era of “Post-Truth” (2017: 7), and A. C. Grayling (in Coughlan 2017) that the relevant shift happened sometime after the financial crash in 2008. So, if these commentators are right, we entered the “post-truth” era sometime in the period between 2008 and 2016. This opens up for the following rendering of the “post-truth” claim:

(PT) Sometime in the period of 2008 to 2016, the empirical substance of politics became significantly less relevant to what political preferences we form.

This claim has two virtues. First, it is agnostic on the underlying mechanism. Consider, in particular, MacMullen’s (2020) identification of four different types of ‘post-factual’ politics: on *unconscious* post-factualism, people systematically (but unwittingly) opt for unreliable belief-forming methods in politics³; on *metaphysical* post-factualism, people deny that there are any truths of the matter in politics⁴; on *motivational* post-factualism, people don’t care enough about what is true to be guided by it in political contexts⁵; and on *epistemic* post-factualism, people are sceptical about there being any unbiased or otherwise reliable means to ever find out what is true on politically relevant matter.⁶ Note that, on each form of post-factualism, the upshot is the same: when people form their political preferences, partisan considerations—or identity cues, as I called them earlier—weigh far more heavily than the empirical substance of politics. So, if we have entered a ‘post-factual’ era in *any* of MacMullen’s senses, we should expect (PT) to be true.

As for its second virtue, (PT) is empirically testable. Specifically, if (PT) is true, we should expect the following to be the case:

(TI) Information effects—i.e., differences between actual preferences and (estimated) fully informed ones—have diminished over time.

The notion of an information effect comes out of political scientists’ desire to understand the role of political knowledge in political attitude and preference formation (see, e.g., Althaus 2003; see also Ahlstrom-Vij *forthcoming*, Caplan 2007, and Delli Carpini and Keeter 1996). By using large-scale survey data, containing both established measures of people’s degrees of political knowledge and their demographics and socioeconomic features, we can employ statistical models to estimate how they likely *would* have responded, had they been *fully informed* (more on this in Section 2). Information effects, in

³ According to MacMullen, such post-factualism might arise for any number of reasons, ranging from a variety of well-known cognitive biases to more specific forms of *motivated reasoning* (e.g., Kahan 2016).

⁴ See, e.g., Calcutt (2016) for a diagnosis along these lines, as well as chapter 6 of McIntyre (2018).

⁵ See, e.g., Goodin and Spiekermann (2018: 365-67) on ‘epistemic agnosticism.’

⁶ As noted by MacMullen, this type of ‘post-factualism’ has not received a lot of attention in the literature.

turn, measure the distance between people’s actual (reported) preferences or attitudes, and the (estimated) preferences or attitudes they likely would have had, under full information.

Relating this back to (TI), the idea is that, since (PT) holds that the empirical substance of politics—and thereby also what we *know* about the relevant substantive issues in politics, as operationalised by way of aforementioned measures of political knowledge—is decreasingly a factor in political preference formation, we should expect information effects to have diminished over time. That is, we should expect to see that, as we move from 2008 to 2016 in particular, it is decreasingly the case that, had people known *more* about the empirical substance of politics, they would have had *different* political preferences.

To sum up, (PT) is a reasonable explication of the “post-truth” narrative, as it figures in recent discussions, and (T1) a testable implication of that narrative, so explicated. Of course, some might be sceptical that there is in fact a well-defined and coherent conception of “post-truth” underlying the relevant discussions. Perhaps those discussions are simply partisan attacks on political opponents, masquerading as a good-faith analysis of genuine social phenomena.⁷ That might be so—I find it difficult to evaluate. All the more reason, then, to understand what I am offering as an explication. It certainly captures a lot of what these discussants at least *say* that they are concerned with, and is as such if nothing else a notion that is *available* to them. It also has the dialectical advantage that it ascribes to the relevant discussants a coherent and empirically testable view held in good faith.

With that being said, the next section turns to the data and methodology used for testing if (TI) is consistent with relevant political attitudes data.

2. Survey Data and Methodology

A lot of discourse surrounding “post-truth” focuses on the US, and on the lead-up to Donald Trump’s election victory in the 2016 Presidential election in particular. For that reason, the following relies on cross-sectional data collected in the election years of 2004, 2008, 2012, and 2016, as part of the American National Election Studies’ (ANES) Times Series Study, surveying U.S. eligible voters on a range of matters, including electoral participation, voting behaviour, and public opinion.⁸ ANES is the gold standard for political attitudes and electoral data in the US. In addition, it is a particularly suitable data set for our purposes, since it has historically formed the basis for exactly the type of knowledge scales and counterfactual modelling that we will be engaged in below.

For purposes of testing (TI), available survey items were therefore reviewed in order to identify a set that (a) was featured in all of the four election years; (b) covered a wide enough range of key political issues or preferences for any subsequent analysis to support a general enough conclusion about the “post-truth” narrative; and yet (c) was small enough for the number of subsequent models not to be

⁷ I am grateful to an anonymous reviewer for this journal for raising this point.

⁸ See Appendix for further details.

too large, given that each item would require four models (one for each election year). In the end, six (binary) dependent variables were constructed. The first variable concerned *whether the respondent supports a Republican as opposed to a Democrat for president*, and as such spoke directly to matter of political preference, and by extension also to political choice.⁹ The second and third variable concerned *whether same-sex couples should be able to adopt*, and *whether less government is better than more government*, which in turn can be expected to speak to the respondent's social and economic ideology (Feldman and Johnston 2014), respectively. The fourth variable concerned *whether the respondent wants to see immigration levels increased*, thereby tapping into matters of race and diversity (Hainmuller and Hiscox 2017). Finally, the fifth and sixth variables concerned *whether the government should make it more difficult to buy a gun*, and *whether the respondent favours the death penalty for murderers*, which in turn is likely associated with respondent's level of authoritarianism (Altemeyer 1981; Farnen and Meloen 2000).

The surveys also include a number of items tapping into political knowledge, building on work by Michael Delli Carpini and Scott Keeter. The two main upshots of their work are as follows: First, so long as scales are made up of items within the broad categories of what government is and does, and of political leaders and parties, they do not need to be long to be diagnostic, and as such correlate not only with (independent) interviewer ratings of people's degree of informedness, but also with a variety of political behaviours that we have independent reasons to believe to be related to a person's degree of political knowledge (e.g., Delli Carpini and Keeter 1993: 1198-99; see also Althaus 2003). Second, since people tend to be *generalists*—if someone knows (or does not know) a lot about one area of politics, they will tend (not) to know a lot about other areas—'researchers developing national or general political knowledge scales need not be overly concerned with the mix of specific topics covered by individual items' (Delli Carpini and Keeter 1996: 174).¹⁰

For our purposes, there were four, relevant knowledge items asked in each of the four years in question: *What party held the majority in the House of Representatives before the election?* *Who is the House Speaker?* *Who is the Vice President?* and *Who is the Chief Justice of the Supreme Court?* These items were used to fit an Item Response Theory (IRT) model, an established way to measure underlying

⁹ This variable was constructed by combining two items: whether the respondent voted for a Republican as opposed to a Democrat in the last election or, if the respondent did not vote, whether they preferred a Republican to a Democrat president.

¹⁰ Following Lupia (2006) it might be objected that informed political choice does not require knowing the answer to the *specific* type of items typically used in these scales. However, as discussed, the point is that knowing the answer to small sets of such specific questions is *diagnostic* of (and as such good evidence for) having (or lacking) a *large* stock of knowledge about things that *are* necessary for informed political choice. For more in defence of knowledge scales, see Ahlstrom-Vij (manuscript).

(continuous) traits, on the basis of a set of response patterns.¹¹ Using that model, each respondent in the data set was then assigned a numerical value, representing their degree of political informedness. This formed the basis for the knowledge variable used below.

Additionally, a number of demographic and socioeconomic variables were included in the models as control variables. A word is in order on the choices made here, since the models constructed were *counterfactual* models—estimating what responses participants *would* have made, had the knowledge variable taken on a particular value (more on this in Section 3)—and it is well-known that care needs to be taken when controlling for variables in such models.¹² We need to control for any *confounders* that can be expected to have an effect on both someone’s degree of political knowledge *and* their political attitudes or preferences. Existing evidence suggests *gender* falls in this category, as women tend to score lower on political knowledge tests (e.g., vanHeerde-Hudson 2020), and gender has a modest effect on politically relevant behaviours (Plutzer 2020). Level of *education*, too, likely affects both your level of knowledge (e.g., Hebbelstrup and Rasmussen 2016), and relevant political behaviours (e.g., Plutzer 2020), and the same likely goes for *income* (more so in the US than in Europe; see, e.g., Vowles 2020) and *union membership* (Macdonald 2019). Moreover, to reduce the overall noise in the models, we also do well to control for any variables that can be expected to have an effect on someone’s political preferences, but not necessarily on their degree of knowledge, such as *race and ethnicity* (e.g., through a ‘shared faith’; see Dawson 1994), *religion* and *social class* (e.g., Evans and Northmore-Ball 2020), *marital status* (e.g., Denver 2008), and *age* (e.g., Plutzer 2020). All of these variables were therefore controlled for in the models.

What about *partisanship*, as measured by party identification? This is arguably the most prominent variable in political-scientific modelling, but there are two reasons for excluding it here. First, partisanship is likely affected by political knowledge, and specifically by knowledge of party and candidate positioning (Brader and Tucker 2018). This would make it a *mediator* in the language of counterfactual modelling.¹³ Controlling for a mediator—here: a node located on a direct or indirect pathway *between* political knowledge and political preference—will mean either blocking or (otherwise) mismeasuring the relevant effect. Second, even if partisanship is *not* a mediator, controlling for it in this context is likely *unnecessary*. After all, socialisation is centred around group-identity considerations relating to religion, ethnicity, class, gender, and the like—all helping shape our conceptions of who we are, and consequently also what positions ‘people like us’ take in politics (again, see Campbell *et al.* 1960; see also Green *et al.* 2002). Consequently, in so far as we control for such group variables—in

¹¹ See DeMars (2010) for an accessible introduction to IRT modelling, de Ayala (2009) for a comprehensive treatment, and the Appendix for more details on the particular model used here.


¹² Pearl (2000) is the central text here. See also Morgan and Winship (2015) on counterfactual modelling in the social sciences, and Keele (2015) for an overview of such modelling in political science in particular.

¹³ Rohrer (2018) provides a highly accessible discussion of mediators and related causal concepts.

the manner done here for gender, race and ethnicity, religion, and class—the determinants of partisanship would already be controlled for in virtue of aforementioned group-identity variables. Still, by way of robustness check, all models were also fitted with the partisanship variable included, so that the reader is able to evaluate the impact of this particular model choice on the results.¹⁴

By way of summing up the causal assumptions made for purposes of modelling, consider the graph offered in Figure 1. Note, in particular, the status of partisanship (‘Party’ in the graph) as a mediator for knowledge (‘Know.’). In the event that partisanship is an unnecessary control, the edge between knowledge and partisanship should be removed.¹⁵ As always with causal graphs of this kind, they are not presented because they necessarily offer a complete account of the mechanisms involved.¹⁶ Their primary purpose is to make explicit to the reader what assumptions are being made for purposes of modelling. If the reader disagrees, they will have the benefit of knowing exactly where the relevant disagreements lie, and what in their view needs to be done in order to improve on the relevant models.

FIGURE 1. CAUSAL GRAPH



¹⁴ As can be seen in Section 3 below, the models also controlling for partisanship yield virtually identical results to those that do not, suggesting perhaps that partisanship is an unnecessary as opposed to harmful control variable.

¹⁵ Note that, even if partisanship is not a mediator, it remains a ‘collider,’ a causal node with more than one incoming edge. Controlling for a collider can introduce spurious, non-causal relationships by opening up a non-causal ‘back-door path,’ in this case from knowledge to Y through gender and partisanship. That path can in turn be blocked by also controlling for gender. See Rohrer (2018) for a helpful discussion of colliders.

¹⁶ Indeed, once we control for certain variables, some simplifications become irrelevant. For example, income is likely causally affected by education (the more educated you are, the more you tend to earn), but there is no edge between them in Figure 1. However, since we are already controlling for both education and income, any such further relationship between the two are irrelevant for modelling purposes. Similar points apply for other simplifications in the graph, e.g., in relation to any casual connection between age and marriage; gender and income; and class and race, to name but three.

Note: A directed acyclical graph (DAG) encoding the assumed relationship between knowledge and political attitudes/preferences ('Y' in the graph), as well as the other causal determinants discussed in the text.

Using these model specifications, a logistic model was then fitted for each of the six dependent variables, and for each of the four election years in the period of 2004 to 2016, using the knowledge variable together with the aforementioned demographic and socioeconomic variables as independent variables, for a total of twenty-four models.¹⁷ As noted by Althaus (2003: 323), using logistic models has the benefit that we avoid the implausible assumption that the relationship between knowledge and political preferences is linear. By also binarizing our knowledge variable, we both avoid the further assumption that knowledge stands in a linear relationship with the logit of the (binary) dependent variables, and proceed in line with standard practice of 'doubly robust' estimation for counterfactual inference, which typically looks to approximate the situation we would have found ourselves in, had our data been the result of a randomized experimental design with a single treatment (e.g., Morgan and Winship 2015).

To that end, the knowledge variable was recoded as a binary variable, with all observations representing someone in the 75th percentile on the knowledge variable coded as 1, and everyone else as 0, thereby separating those who are for our purposes considered 'fully informed' (i.e., 'treated,' if we think on the model of an experimental design) from those who are not. (Why not operationalize 'fully informed' as achieving the *maximum* score on the scale? Primarily to avoid the charge that the bar for full information is set at an unreasonably high level, since a fairly small proportion, 12%, of respondents in the data set achieved the highest knowledge score.) The 'double robustness' of the resulting estimation owes to how effects are estimated in a context where we have both controlled for (assumed) confounds (as discussed above), and also taken steps to make up for the fact that the data have *not* come about as a result of randomized assignment. In the present case, this second layer of 'robustness' was achieved by using so-called 'propensity scores' as weights in the models.

In our case, propensity scores measure the probability (i.e., propensity) that an observation will be found in the 'fully informed' category, as a function of someone's demographic features, and the demographic features discussed above in particular. The idea is to then use these scores to remove any correlation between these features and the 'informed' category, to justify a counterfactual inference. To see why, return to the paradigm of a randomized experimental design, where the random allocation of participants to a treatment and a control group means that no feature of the participant is predictive of being found in the treatment as opposed to in the control. Whether female or male, rich or poor (etc.), you are equally likely to end up in one group as opposed to in the other, provided assignment is truly random. In the case of observational data, by contrast, this might not be the case. In the case at hand, it

¹⁷ As already noted, a separate set of twenty-four models also controlling for partisanship were fitted as well, for a total of forty-eight. See Appendix for further details on all models, including diagnostics.

might (for example) be that some features of the observations—e.g., their level of education, their income, or what have you—are predictive of ending up in the ‘informed’ category.

By weighting our regression models with propensity score, we counteract correlations like these.¹⁸ Specifically, since propensity scores measure the probability of ending up in the ‘treatment’ category, given a set of covariates—in our case, the probability that you would be ‘informed,’ given your age, level of income, level of education, and so forth for all measured covariates—we can use the *inverse* of those scores as weights (such that an observation with a low propensity is weighted heavily) in fitting the model. Given an appropriately chosen set of covariates when calculating the scores, this recreates a situation that would have been expected in a randomized experiment, thereby allowing greater confidence in any counterfactual inference.¹⁹

Against this background, a set of propensity scores were therefore calculated for each of the four election years, measuring the probability that an observation would be found in the ‘informed’ category on the basis of its demographic features. These scores were then used as weights in fitting the corresponding models (6 models x 4 election years).²⁰ These are the models used in evaluating (TI) in the next section.

3. Putting the “Post-Truth” Claim to the Test

On the testable implication introduced above, i.e., (TI), we should expect to see information effects—i.e., differences between actual (reported) preferences and estimated fully informed ones—having diminished over time, if (PT) is true. Using the counterfactual models introduced in the previous section, we can see if that implication is consistent with the data by measuring the *difference* between (a) the aggregate proportion of affirmative answers in relation to each of the six preference variables, and (b) the models’ estimates of what that proportion *would* have been, had each person been fully informed.

The (actual) aggregate proportion was estimated by determining the proportion of support in the ANES data set for each topic and year, after having weighted each observation using the survey weights included in the data set, to approximate representativeness. The aggregate proportion of informed support was calculated by setting the (binary) knowledge variable for each respondent to 1, representing

¹⁸ Propensity score *weighting*, as used here, should not be confused with propensity score *matching*. The latter is a preprocessing technique which involves matching each ‘treated’ observation with a ‘control’ observation with the same propensity score, and then discarding all unmatched observations. Propensity score matching has been shown to actually worsen balance under some circumstances (King and Nielsen 2019), and is on that account not a recommended method for preprocessing data in the context of counterfactual inference.


¹⁹ See Morgan and Winship (2015) for a more comprehensive discussion of counterfactual inference on the basis of observational data, including through propensity score weighting.

²⁰ See the Appendix for more details on the propensity scores and weights.

their becoming informed, while otherwise remaining just the way they are, across the variables controlled for. We can think of this on the model of an imaginary ‘knowledge pill,’ that renders each respondent informed, but otherwise leaves them exactly the way they are (across measured covariates). To illustrate, imagine that—prior to the ‘knowledge pill’—a respondent reports support for the death penalty. Then, we ‘administer the pill’ (i.e., set the value of their knowledge variable to 1), and subsequently ask the models to estimate how likely the respondent now would be to report supporting the death penalty. By doing this for each respondent (for each topic and election year), aggregating all probabilities, and finally using the same survey weights as above, we can calculate the ‘informed’ proportion of support on each topic for each year.

Figure 2 shows us the *difference* we see on each of the six topics between the actual proportion of support and the estimated ‘informed’ proportion of support. The solid lines give the estimates on the models that do *not* control for partisanship, while the dashed lines in the corresponding colours give the estimates for the models that do control for partisanship, by way of robustness check.

FIGURE 2. INFORMATION EFFECTS OVER TIME



Note: Solid lines give the difference between actual and informed proportion of affirmative answer on six preference variables, with survey weights applied to approximate representativeness. Dashed lines in the corresponding colour give the estimates resulting from also controlling for partisanship. Solid points designate significant estimates for the knowledge coefficient. (See Appendix for further model details.)

Consider the top left panel of Figure 2, by way of illustration. In 2004, a Republican president would have seen about one fifth of a percentage point more public support, given a fully informed electorate; in 2008, about 2.5 percentage points more support; in 2012, about half a percentage point more support;

and in 2016, about a tenth of a percentage point less. Looking at Figure 2 in the context of (TI), what we should expect to see is these graphs converging towards 0—the dashed, horizontal line, representing the absence of any information effect—as we move from 2004 to 2016. The top left panel could be read in that way, given the substantial drop from 2008 to 2012, both years for which the knowledge coefficient estimate in the corresponding models is significant. That said, the p -value for 2008 is only marginally below 0.05, so that estimate should be interpreted with some caution. Moreover, the robustness check offered by the models also controlling for partisanship (dashed red line) provides further reason not to draw any substantial conclusions about this particular variable.

By contrast, we see a fairly clear trend towards no information effect over time when turning to the graphs regarding wanting more immigrants, wanting stricter gun laws, and wanting same-sex couples to be able to adopt. On wanting more immigrants, the knowledge coefficient estimate is significant for each of the four years. In the case of gun laws, it is significant for 2004 and 2016, between which we see a substantial drop, and the same goes for 2004 and 2012 on wanting same-sex couples to be able to adopt. As for support for the death penalty and less government, we see a trend towards a smaller information effect over the three years for which the coefficient estimate is significant (i.e., 2008-2016) but in a less clear-cut fashion than in the other cases. Moreover, in all of these cases the relevant estimates seem robust in the face of model choice, with the estimates coming out more or less identical whether we control for partisanship or not.

4. An Alternative Explanation?

If we have entered a “post-truth” era, then we should expect decreasing levels of information effects over recent years. The previous section found evidence of such a decrease on some key political preferences. However, this only offers good evidence of the “post-truth” narrative in the absence of any competing explanation of why we should see such a decrease. It might be objected that one such explanation is that people have actually come to make use of facts to an increasing extent over the relevant period.²¹ The idea would be that, as people make greater use of facts, presumably on account of having become more informed over time, it also becomes the case that, had they known more, they would have held roughly the same beliefs, as per (TI), simply because they already know a fair amount. This is an intriguing alternative reading of the results from the previous section, and it has the following going for it: at least in the sample used here, the proportion of ‘fully informed’ respondents—again, defined as being in the 75th percentile on the knowledge variable, calculated across the four election years—increases from 11% in 2004, to 13% in 2008, to 35% in 2012, and to 47% in 2016.

However, it is one thing to know (more) things, and another to make use of the facts known in preference formation. Once we keep this distinction in mind, we will see on closer inspection that this

²¹ I am grateful to an anonymous reviewer for this journal for raising this objection.

alternative explanation is not compatible with—and as such, does not explain—the results in the previous section, nor thereby with (TI), i.e., the claim that information effects have diminished over time. To see why, we do well to consider this alternative explanation alongside (PT), i.e., the idea that the empirical substance of politics in recent years became significantly less relevant to what political preferences we form. In particular, imagine a simplified version of the results from the previous section, in terms of two points in time, t_1 and t_2 . At t_1 we have some non-trivial information effect, while at t_2 we have no such effect. That means that, at t_1 , it is generally the case that people would have held different preferences from the ones they actually hold, had they been informed. By contrast, this is not generally the case in t_2 —hence, the absence of any information effect. Why this difference between the two points in time? On (PT), it is explained with reference to how people take into account the empirical substance of politics in forming their beliefs at t_1 but not at t_2 . At t_1 that substance is a factor in people's preference formation, which is why they would in many cases have held different preferences, had they know more about that substance; at t_2 they do not factor in that substance, which is why they would not have held different preferences, had they known more.

Let us now try to explain the same pattern in information effects with reference to the alternative reading. We are to imagine that, at t_1 , people do not make use of facts, which includes facts about the empirical substance of politics. Should we expect to see any information effect under these circumstances? No. Keep in mind the distinction between knowing facts and making use of them, in the sense of their playing a role in preference formation. This distinction is important when we consider the counterfactual that we are investigating by way of information effects: had people known more about the substance of politics, would they have held different preferences? If people do *not* make use of facts at t_1 , the answer is 'no': as people do not make use of the relevant facts, knowing more of them would not have made a difference. This, of course, is contrary to the pattern of information effects we are looking to explain, whereby we have a non-trivial information effect at t_1 .

Turning to t_2 , we are now to imagine people having started to use facts. Should we expect any information effect at this point? There are two possibilities. Either people have become so informed in the period between t_1 and t_2 that there is not much more to learn, which would mean that the (marginal) information effect on preferences of whatever further knowledge they might have attained is small. Or—more plausibly, perhaps—people at t_2 are still far from politically omniscient, in which case we should see a non-trivial information effect at that point, owing to how people who are now assumed to be making use of facts in preference formation would have held different preferences, had they known (even) more—again, contrary to the pattern of information effects to be explained, whereby we see an absence of any information effect at t_2 .

To sum up, the alternative reading in terms of how people go from not using facts to using facts does not seem well-placed to explain the relevant pattern of information effects. Bringing matters back from our simplified example to the results in the previous section, if people have gone from not using facts to any significant extent in 2004 to using them to a great(er) extent in 2016, then we should expect

to see either of two things. Either (a) we should see no real difference in information effects over this time-period, as people go from not using facts (meaning that knowing more would not have made a difference) to starting to use such facts, and in the process coming to learn so much about the empirical substance of politics that the marginal effect on preferences of knowing more is vanishingly small (meaning—again—that knowing more would not have made a difference); or (b) we should see an *increase* in information effects over time, as people go from not using facts to starting to use facts, but not to the point of approximating political omniscience, and it thereby becoming the case that, had they known more, they would likely have held different preferences.

Thinking back to how we defined ‘fully informed’ earlier, and in particular to the healthy proportion of respondents in the data set that came out ‘fully informed’ on that definition, we can safely assume that the range of political knowledge measured by our scale does not encompass political omniscience. That means that (b) is the more likely possibility of the two, and that the alternative reading thereby predicts the opposite pattern of information effects from the one we saw in the previous section: that is, rather than predicting that we should see a *decrease* in information effects over time, it predicts that we should see an *increase* in those effects. In light of that, we reject this alternative reading. Consequently, the claim that (TI) offers evidence for (PT) still stands.

5. Conclusion, Limitations, and Avenues for Future Work

The analysis provided in Section 3 suggests that (TI) is consistent with the data, at least on some key political issues in a US context. Looking at the period of 2004 to 2016 in particular, we see evidence of a decrease in information effects on party preferences, particularly on three political topics that typically garner a lot of attention in political discourse, and likely tap into matters of race and diversity, social ideology, and authoritarianism, namely: immigration, same-sex adoption and gun laws. This is in line with what it was suggested we should see, if the specific “post-truth” claim identified in Section 1 is correct: at least on these issues, and over the period of time covered by these models, what you know about the substance of politics does seem to be decreasingly relevant to what preferences you form. This offers some novel, empirical evidence for the “post-truth” narrative.

One limitation of the present investigation is, however, the relatively short time-frame of twelve years (2004 to 2016). In the present case, that limitation is an artefact of needing a set of knowledge items and dependent variables with sufficient data available for each of the years, to minimise the need for imputation.²² Future work would benefit from looking at recent discussions about a “post truth” within a historical context. Needless to say, there is no reason to make romantic assumptions about a political golden age, when truth and truthfulness reigned supreme. As Hannah Arendt noted decades

²² For example, for the present ANES data set, 57% of responses are missing for the knowledge items for the year 2000; for 1996, 100% of responses are missing for the dependent variables concerning gun laws, same-sex adoption, and less government; for 1992, 100% of responses are missing for the variable concerning gun laws.

back, '[n]o one has ever doubted that truth and politics are on rather bad terms with each other, and no one, as far as I know, has ever counted truthfulness among the political virtues' (1967: 49). But equally, we cannot rule out that the relationship between politics and knowledge might vary over time, whether along a clear trend-line, or through the influence of truth waxing or waning from one period to the next.

Applying such a wider historical lens would also help us understand whether our current situation is in any way unique. While authors pushing a "post-truth" narrative typically suggest that it is, I take no view on the matter. If the analysis presented in this paper is correct, *something* seems to have happened over the relevant 12-year period that fits the "post-truth" narrative. But is this the first time we have seen such a reduction in the influence of the empirical substance of politics? If it is, then there seems to be something unique about it. If not, is it part of some longer-term trend, where the period we have looked at in this paper simply forms the tail end of it? And if so, when did the relevant decline in the influence of the substance of politics begin? And from what point did the decline start? Did the substance of politics use to be highly influential, or has it never been particularly influential, and simply become even less so?

Needless to say, these questions go beyond the scope of the present paper. Still, the results we have arrived at thereby serve up several potentially fruitful avenues for further investigation. In light of the above, such investigations would likely benefit from engaging with different data sets, and also from engaging with a wider set of methodologies, including (but not restricted to) historical ones.²³

References

- Achen, C. and L. Bartels (2016). *Democracy for Realists: Why Elections do not Produce Responsive Government*, Princeton, NJ: Princeton University Press.
- Ahlstrom-Vij, K. (forthcoming). 'The Case for Modelled Democracy,' forthcoming *Episteme*.
- Ahlstrom-Vij, K. (manuscript). 'Political Knowledge: Measurement, Elitism, and Dogmatism,' unpublished manuscript.
- Altemeyer, B. (1981). *Right-wing Authoritarianism*, Winnipeg, Canada: University of Manitoba Press.
- Althaus, S. (2003). *Collective Preferences in Democratic Politics*, Cambridge: Cambridge University Press.
- Arendt, H. (1967). 'Truth and Politics,' *the New Yorker*, February 25 issue: 49-88.
- Ball, J. (2017). *Post-Truth: How Bullshit Conquered the World*, Biteback Publishing.
- Bartels, L. (1996). 'Uninformed Voters: Information Effects in Presidential Elections,' *American Journal of Political Science* 40(1): 194–230.
- Berelson, B. R., P. F. Lazarsfeld, and W. N. McPhee (1954). *Voting: A Study of Opinion Formation in a Presidential Campaign*, Chicago: University of Chicago Press.

²³ I am grateful to two anonymous reviewers and the editors of this journal for valuable comments on a previous version of this paper, as well as to Anna Alexandrova, Helen Curry, Steve John, and the rest of the audience at the History and Philosophy of Science seminar at the University of Cambridge for valuable discussion of the paper. Thanks also to Michael Hannon, Klemens Kappel, and Andrew Spear for valuable comments on early drafts of this paper.

- 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.
- Calcutt, A. (2016) 'The Truth about Post-truth Politics,' *Newsweek*, Nov 21, 2016.
- Campbell, A, P. E. Converse, W. E. Miller, and D. E. Stokes (1960). *The American Voter*, New York: John Wiley & Sons.
- Caplan, B. (2007). *The Myth of the Rational Voter*, Princeton, NJ: Princeton University Press.
- Coughlan, S. (2017). 'What Does Post-Truth Mean to a Philosopher?' BBC News, January 12; available at <https://www.bbc.co.uk/news/education-38557838>.
- D'Ancona, M. (2017). *Post Truth: The New War on Truth and How to Fight Back*, London: Ebury Press.
- Davis, E. (2017). *Post-Truth: Why We Have Reached Peak Bullshit and What We Can Do About It*, Little, Brown.
- Dawson, M. C. (1994). *Behind the Mule: Race and Class in African-American Politics*, Princeton, NJ: Princeton University Press.
- de Ayala, R. J. (2009). *The Theory and Practice of Item Response Theory*, New York, NY: Guilford Press.
- Delli Carpini, M. X., and S. Keeter (1993). 'Measuring Political Knowledge: Putting First Things First,' *American Journal of Political Science* 37(4): 1179-1206.
- Delli Carpini, M. X., and S. Keeter (1996). *What Americans Know about Politics and Why It Matters*, New Haven, CT: Yale University Press.
- DeMars, C. (2010). *Item Response Theory*, Oxford: Oxford 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.
- Evans, G. and K. Northmore-Ball (2020). 'Long-Term Factors: Class and Religious Cleavages,' in Fisher et. al. (2020): 123-135.
- Farnen, R. F., and J. D. Meloen (2000). *Democracy, Authoritarianism and Education: A CrossNational Empirical Survey*, Basingstoke: Macmillan Press.
- Feldman, S., and C. Johnston (2014). 'Understanding the Determinants of Political Ideology: Implications of Structural Complexity,' *Political Psychology* 35(3): 337-358.
- 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.
- Frankfurt, H. (2005). *On Bullshit*, Princeton, NJ: Princeton University Press.
- Goodin, R, and K. Spiekermann (2018). *An Epistemic Theory of Democracy*, Oxford: Oxford University Press.
- 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.
- Greifer, N. (2020). 'cobalt: Covariate Balance Tables and Plots,' R package version 4.0.0; <https://CRAN.R-project.org/package=cobalt>.
- Hainmueller, J., and M. J. Hiscox. (2007). 'Educated Preferences: Explaining Attitudes Towards Immigration in Europe,' *International Organization* 61(2): 399-442.
- Harrell, F. E, C. Dupont, and collaborators (2019). 'Hmisc: Harrell Miscellaneous,' R package version 4.2-0; <https://CRAN.R-project.org/package=Hmisc>.

- Hebbelstrup, S. and R. Rasmussen (2016). 'Education or Personality Traits and Intelligence as Determinants of Political Knowledge?' *Political Studies* 64(4): 1036-1054.
- Kahan, D. (2016). 'The Politically Motivated Reasoning Paradigm, Part 1: What Politically Motivated Reasoning Is and How to Measure It,' *Emerging Trends in the Social and Behavioral Sciences*, 1-16.
- Keele, L. (2015). 'The Statistics of Causal Inference: A View from Political Methodology,' *Political Analysis* 23: 313-335.
- Key, V.O. (1966). *The Responsible Electorate: Rationality in Presidential Voting 1936-1960*, Cambridge, MA: Belknap Press.
- King, G. and R. Nielsen (2019). 'Why Propensity Scores Should Not Be Used for Matching,' *Political Analysis* 27(4): 435-454.
- Lupia, A. (2006). 'How Elitism Undermines the Study of Voter Competence,' *Critical Review* 18(3): 217-232.
- Macdonald, D. (2019). 'How Labor Unions Increase Political Knowledge: Evidence from the United States,' *Political Behavior*; <https://doi.org/10.1007/s11109-019-09548-7>.
- MacMullen, I. (2020). 'Survey Article: What is "Post-factual" Politics?' *Journal of Political Philosophy* 28(1): 97-116.
- McIntyre, L. (2018). *Post-Truth*, Cambridge, MA: MIT Press.
- Morgan, S. and C. Winship (2015). *Counterfactuals and Causal Inference*, Cambridge: Cambridge University Press.
- Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*, Cambridge: Cambridge University Press.^[11]
- 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). '*ltm*: An R package for Latent Variable Modelling and Item Response Theory Analyses,' *Journal of Statistical Software* 17(5): 1-25.
- Rohrer, J. (2018). 'Thinking Clearly About Correlations and Causation: Graphical Causal Models for Observational Data,' *Advances in Methods and Practices in Psychological Science* 1(1): 27-42.
- Sloman, S. and P. Fernbach (2017). *The Knowledge Illusion*. New York: Riverhead Books.
- vanHeerde-Hudson, J. (2020). 'Political Knowledge: Measurement, Misinformation and Turnout,' in Fisher *et. al.* (2020): 369-382.
- Vowles, J. (2020). 'The Big Picture: Turnout at the Macro-Level,' in Fisher *et. al.* (2020): 57-68.
- Yen, W. M. (1993). 'Scaling Performance Assessments: Strategies for Managing Local Item Dependence,' *Journal of Educational Measurement* 30: 187-213.


Appendix

The data set used in the paper was drawn from the [ANES Time Series Cumulative Data File](#). Questions that have been asked in three or more Time Series studies since 1948 are eligible for inclusion in the data file, with variables recoded for comparability across years. Sample sizes for the years relevant to this study were 1,212 for 2004, 2,322 for 2008, 5,914 for 2012, and 4,270 for 2016. 4.9% of values were missing across the variables used (see Section 2 in the main body of the text). These were imputed with multiple imputation, using *aregImpute* in R's (R Core Team 2017) *Hmisc* package (Harrell et. al 2019).

The knowledge items given in Section 2 were used to fit a two-parameter item response theory (IRT) model in R, using *ltm* (Rizopoulos 2006), to estimate the latent ability of respondents. The discrimination values for each of the four items were good (i.e., > 1), with the difficulty values spread out nicely across the range. The guessing parameters of a corresponding three-parameter model came out very low, so the simpler, two-parameter model was used to calculate an ability score for each respondent. The resulting model had reasonable properties: tests suggested unidimensionality, local independence (by Yen's Q3; Yen 1993), and good model fit (evaluated through a plot of observed versus expected values).

As noted in the body of the text, the counterfactual models were constructed using 'doubly robust' estimation for counterfactual inference. Propensity scores were estimated using boosted logistic regression, as implemented in R's *twang* package (Ridgeway et al. 2020), and evaluated by way of the diagnostic features in that package as well as visually using *cobalt* (Greifer 2020) to confirm improved balance between the two groups. By way of illustration, consider Figure 3. The left-hand panel shows the balance (or rather: lack thereof) for the income variable in the 2016 dataset, prior to applying the propensity weights. Note that informed participants (teal bars, designated here as 1 or 'treated') are overrepresented among the wealthy, and underrepresented among the less wealthy. The right-hand panel shows the improved balance achieved once the weights have been applied.

FIGURE 3. BALANCE PLOTS FOR INCOME (2016)



Note: Balance plots for income in the 2016 data set, before and after applying propensity weights. The five groupings in each correspond to people whose income is in the 0-16th percentile (1), 17-33rd percentile (2), 34-67th percentile (3), 68-95th percentile (4), and 96-100th percentile (5).

Using these propensity scores as weights, a logistic regression model was then fitted (using *glm* in *R*'s *stats* package) for each of the six attitudinal variables and four years. Diagnostics for each model are provided in Table 1 (partisanship not controlled for) and Table 2 (partisanship controlled for).

TABLE 1. MODEL DETAILS AND DIAGNOSTICS

<i>Model</i>	<i>Knowledge coefficient</i>	<i>p-value</i>	<i>Box-Tidwell test on age variable</i>	<i>t-test: splines vs. no splines</i>	<i>McFadden</i>	<i>Highest VIF</i>	<i>Cook's > 1</i>	<i>Std. res. > 3</i>	<i>Coef. diff. w/o infl. obs.</i>
President (2004)	-0.0120	0.9138	0.0096	0.9216	0.1566	1.2744	0	30	0.0301†
President (2008)	0.1595	0.0430*	0.0075	0.9805	0.2431	1.1137	1	54	-0.0035
President (2012)	0.1126	0.0101*	0.4746	-	0.2026	1.2026	0	95	-0.0027
President (2016)	-0.0507	0.3019	0.1872	-	0.1522	1.2155	0	65	0.0032
Immigration (2004)	0.5570	0.0003*	0.0223	0.9894	0.1399	1.2811	1	20	-0.0140
Immigration (2008)	0.3722	0.0001*	0.0780	-	0.0826	1.1247	0	37	-0.0173
Immigration (2012)	0.3503	0.0000*	0.0000	0.9759	0.0531	1.2015	0	162	0.0228
Immigration (2016)	0.2349	0.0002*	0.0000	0.8728	0.0754	1.2070	0	133	0.0423
Guns (2004)	0.4702	0.0000*	0.0813	-	0.1381	1.2930	1	24	-0.0533
Guns (2008)	0.0326	0.6390	0.1841	-	0.1154	1.1169	0	69	0.0164†
Guns (2012)	0.0252	0.5227	0.0002	0.9688	0.0676	1.2207	0	113	0.0111
Guns (2016)	0.1281	0.0061*	0.1410	-	0.0739	1.2268	0	45	-0.0084
Adoption (2004)	0.9578	0.0000*	0.1089	-	0.1640	1.2948	0	23	-0.0473
Adoption (2008)	-0.0135	0.8472	0.9178	-	0.1235	1.1080	0	71	0.0259†
Adoption (2012)	0.0850	0.0372*	0.1994	-	0.0799	1.1977	0	105	0.0026
Adoption (2016)	0.0841	0.1087	0.3904	-	0.0930	1.2047	0	93	0.0045
Death penalty (2004)	-0.2036	0.0849	0.9243	-	0.1441	1.3181	0	31	0.0167
Death penalty (2008)	-0.2002	0.0049*	0.1895	-	0.0948	1.1282	0	74	-0.0125
Death penalty (2012)	-0.1907	0.0000*	0.0000	0.9600	0.0541	1.2190	0	135	-0.0057
Death penalty (2016)	-0.2548	0.0000*	0.0001	0.8629	0.0761	1.2219	0	76	0.0061
Less govt. (2004)	0.1092	0.3258	0.6722	-	0.1580	1.2769	0	28	0.0181
Less govt. (2008)	0.4298	0.0000*	0.0002	0.9352	0.1474	1.1147	0	67	-0.0096
Less govt. (2012)	0.2948	0.0000*	0.0004	0.9858	0.1262	1.2082	0	112	-0.0044
Less govt. (2016)	0.3054	0.0000*	0.3742	-	0.0790	1.2133	0	41	0.0021

Note: The **2nd column** gives the coefficient for the knowledge variable, and the **3rd column** its *p*-value, marked with an asterisk when below 0.05. The **4th column** gives the *p*-value of a Box-Tidwell test on the age variable (the only continuous variable used), performed by including the product of the variable and its natural logarithm as an additional predictor, and seeing if it comes out significant, suggesting a non-linear relationship between the predictor and the logit of the outcome. In cases where that was the case (*i*) a spline version of the model was fitted (using *R*'s *splines* package); (*ii*) using that model, a prediction was made on each observation in the data set; and (*iii*) the predicted (fitted) values were then compared to the predicted values on the original model using a *t*-test, the *p*-value of which is reported in the **5th column**. The **6th column** gives the McFadden value for the model. The **7th, 8th, and 9th columns** give the highest value across all variables of a VIF (variance inflation factor) test for multicollinearity (values substantially higher than 1 suggest potential multicollinearity); the number of instances where the Cook's distance value of any observation exceeded 1 (signifying potential outliers); and the number of instances where the standardized residual of any observation exceeded 3 (signifying potentially influential observations). In cases where there was more than one instance with a standardized residual greater than 3, the model was re-fitted without those observations to measure the difference in the estimated value for the knowledge coefficient, compared to the original model. That difference is reported in the **10th column**. Instances where the absolute size of the difference is equal to or greater than half of the original coefficient estimate are marked with a dagger symbol. These differences do not alter the analysis offered in Section 3.

TABLE 2. MODEL DETAILS AND DIAGNOSTICS (PARTISANSHIP INCLUDED)

<i>Model</i>	<i>Knowledge coefficient</i>	<i>p-value</i>	<i>Box-Tidwell test on age variable</i>	<i>t-test: splines vs. no splines</i>	<i>McFadden</i>	<i>Highest VIF</i>	<i>Cook's > 1</i>	<i>Std. res. > 3</i>	<i>Coef. diff. w/o infl. obs.</i>
President (2004)	0.1173	0.4852	0.0082	0.9850	0.5396	1.3177	1	19	0.0076
President (2008)	0.0215	0.8439	0.5142	-	0.5426	1.1565	0	40	-0.0075
President (2012)	0.134	0.0406	0.7607	-	0.5782	1.2116	0	136	-0.0027
President (2016)	0.0195	0.7884	0.4687	-	0.5411	1.2345	0	105	0.0083
Immigration (2004)	0.4718	0.0027*	0.0172	0.9729	0.1355	1.2706	0	18	-0.0084
Immigration (2008)	0.3323	0.0003*	0.1000	-	0.0822	1.1211	0	41	-0.0176
Immigration (2012)	0.3708	0.0000*	0.0000	0.9973	0.0704	1.2062	0	163	0.0246
Immigration (2016)	0.2182	0.0006*	0.0000	0.8739	0.1175	1.2058	0	112	0.0308
Guns (2004)	0.4773	0.0001*	0.0902	-	0.1733	1.2910	0	21	-0.0411
Guns (2008)	0.0678	0.3367	0.2075	-	0.1298	1.1150	0	67	0.0152
Guns (2012)	0.0482	0.2442	0.0000	0.9664	0.1320	1.2185	0	134	0.0070
Guns (2016)	0.1166	0.0219*	0.0288	0.9424	0.1852	1.2325	0	82	-0.0129
Adoption (2004)	0.93	0.0000*	0.0482	0.9830	0.1882	1.2961	0	19	-0.0495
Adoption (2008)	0.0327	0.6532	0.1231	-	0.1674	1.1161	0	62	0.0168 [†]
Adoption (2012)	0.1161	0.0063*	0.2251	-	0.1363	1.1988	0	113	-0.0013
Adoption (2016)	0.0757	0.1625	0.7782	-	0.1440	1.2023	0	81	0.0019
Death penalty (2004)	-0.1388	0.2661	0.6352	-	0.1972	1.3159	0	27	0.0074
Death penalty (2008)	-0.2523	0.0004*	0.0524	-	0.1068	1.1267	0	69	-0.0199
Death penalty (2012)	-0.1931	0.0000*	0.0000	0.9602	0.0842	1.2209	0	147	-0.0071
Death penalty (2016)	-0.268	0.0000*	0.0000	0.8621	0.1243	1.2219	0	88	0.0113
Less govt. (2004)	0.1601	0.1788	0.5356	-	0.2091	1.2806	0	20	0.0035
Less govt. (2008)	0.4326	0.0000*	0.0009	0.9504	0.2421	1.1187	0	53	-0.0181
Less govt. (2012)	0.3466	0.0000*	0.0001	0.9946	0.2517	1.2121	0	124	-0.0044
Less govt. (2016)	0.3837	0.0000*	0.5844	-	0.1785	1.2126	0	78	0.0026

Note: The **2nd column** gives the coefficient for the knowledge variable, and the **3rd column** its *p*-value, marked with an asterisk when below 0.05. The **4th column** gives the *p*-value of a Box-Tidwell test on the age variable (the only continuous variable used), performed by including the product of the variable and its natural logarithm as an additional predictor, and seeing if it comes out significant, suggesting a non-linear relationship between the predictor and the logit of the outcome. In cases where that was the case (i) a spline version of the model was fitted (using *R*'s *splines* package); (ii) using that model, a prediction was made on each observation in the data set; and (iii) the predicted (fitted) values were then compared to the predicted values on the original model using a *t*-test, the *p*-value of which is reported in the **5th column**. The **6th column** gives the McFadden value for the model. The **7th, 8th, and 9th columns** give the highest value across all variables of a VIF (variance inflation factor) test for multicollinearity (values substantially higher than 1 suggest potential multicollinearity); the number of instances where the Cook's distance value of any observation exceeded 1 (signifying potential outliers); and the number of instances where the standardized residual of any observation exceeded 3 (signifying potentially influential observations). In cases where there was more than one instance with a standardized residual greater than 3, the model was re-fitted without those observations to measure the difference in the estimated value for the knowledge coefficient, compared to the original model. That difference is reported in the **10th column**. Instances where the absolute size of the difference is equal to or greater than half of the original coefficient estimate are marked with a dagger symbol. These differences do not alter the analysis offered in Section 3.