Public Reason 15(2)-16(1): 70-87

Political Knowledge: Measurement, Elitism, and Dogmatism

Kristoffer Ahlstrom-Vij Birkbeck
Featurespace/Visa

Abstract. Political knowledge is a resource: having a lot of it means being in a position to navigate the political world, and stand a better chance of connecting your fundamental political goals with successful means. The present piece argues that standard political knowledge tests measure political knowledge, so understood, and uses counterfactual modeling to demonstrate the difference having such knowledge can make to political choice. It then takes up two of the most forceful objections to political knowledge and its measurement, by (a) rejecting the idea that knowledge scales encode elitist assumptions, and (b) arguing that, even if political knowledge breeds dogmatism, such dogmatism can be expected to serve a protective function in exactly the type of hostile epistemic environments – filled with lies, falsehoods, and misinformation – that make up the political domain.

Keywords: political knowledge; public ignorance; information effects; knowledge scales.

1. THE TRADITIONALIST AND THE REVISIONIST

It is well-established that most of us know very little when it comes to politically relevant facts (e.g., Delli Carpini and Keeter 1996; Achen and Bartels 2016; Somin 2016). But what follows from this? There are two schools of thought: the traditionalist and the revisionist.1 The traditionalist takes the straightforward line that public ignorance poses problems for democracy. The two main problems are as follows:

- **The problem with ignorant voting:** Voters might be rationally ignorant (e.g., Downs 1957). But what is individually rational can still be trumped by moral obligations in cases of aggregate harm. (My attempting to free-ride on recycling might be rational, but I still ought to recycle.) Something similar goes for ignorant voting: it is a 'collectively harmful activity' (Brennan 2016; see also Somin 2016).[1]
- **The problem with ignorant non-voting:** Political knowledge is unevenly distributed, with higher levels found among people who are white, educated, older, and male. Since those with high levels of political knowledge are also more likely to vote, political igno- rance means that those already disadvantaged remain so, and that injustices will likely be perpetuated rather than eradicated through the democratic process (e.g., Delli Carpini and Keeter 1996, 138).

The revisionist, by contrast, typically argues that, voters use varieties of heuristics and rules of thumb (e.g., party-political cues) to make up for what they do not know

1] I am here setting aside a large literature responding to concerns about public ignorance by suggesting that a large public is likely to outperform a smaller number of more competent individuals, or arguing that there are epistemic benefits of deliberative or liberal institutions. See (Ahlstrom-Vij 2020b) for an argument that these considerations fail to cast doubt upon the traditionalist line covered here.

(e.g., Popkin 1991). For example, if I don't know much about economic policy, but I know that party X is 'my' party, I can just vote whatever their position is on the matter.

Importantly, the revisionist narrative is testable, since it entails a counterfactual: *By relying on heuristics, we vote and hold political preferences in the same way we would have, had we been fully informed.* Note that 'fully informed' is a term of art here: it does not designate political omniscience, but is usually operationalised as getting a perfect score on a civic style political knowledge test (more on these in Section 2 below). Note also that this counterfactual provides a normative standard for thinking about voter competency. In particular, consider Arthur Lupia's – a prominent revisionist – who suggests that

> [...] we should evaluate a voter as competent regardless of how she reaches a conclusion, as long as it is the conclusion she would have reached had she been aware of the best available information (Lupia 2006, 226).

With this in mind, we can capture what constitutes a sufficient amount of political knowledge – sufficient for competence, as Lupia would say – as follows:

> **CHANGES:** You have a sufficient amount of political knowledge for purposes of political preference and choice if it is the case that, had you known more, this would not have changed your political preferences and choices.

Differently put, you know enough if your preferences and choices would have survived your knowing more. With that in mind, we can test the relevant counterfactual as follows, to see if the revisionist's story holds up. We first gather data on people's level of political knowledge as given by some appropriate political knowledge test, alongside their political preferences and demographics – or, alternatively, we rely on established surveys of political opinion, such as the American National Election Studies (ANES) in the U.S. and the British Election Study (BES) in the U.K. Then, we use regression analysis to identify the relationship between test scores, demographic variables, and the probability of reporting particular political preferences. Finally, we use that model to increase the knowledge variable of each to the maximum, and note for each respondent how that probability changes. Assuming that the relevant knowledge tests do in fact measure political knowledge (more on this in Section 2), the presence of any information effects – i.e., differences between our actual preferences and the preferences we would have held, had we been fully informed – is evidence that we do not know enough.

Such information effects also speak directly to the revisionist counterfactual. Henrik Os- carsson explains:

> If the popular [revisionist] view of low information rationality [through the use of heuristics] is correct – that most uninformed voters most of the time make the same voting choices as they would have had they been fully informed – we would not expect any significant information effects, and certainly not any important changes in the aggregate outcome of elections. More knowledgeable voters would not differ from less knowledgeable in political preference or behavior since the use of heuristics would be a successful compensatory strategy" (Oscarsson 2007, 304).

As it happens, Information effects have been identified in a number of geographical contexts over the past couple of decades (see Ahlstrom-Vij (2020a) for a discussion). For example, Scott Althaus (2003) finds that fully informed preferences tend to be 'more dovish and interventionist on foreign policy, less conservative on social, environmental, and equal rights issues, and more conservative on morality issues and questions about the proper limits of government activity' (2-4). Across all survey questions considered by Althaus, the average difference on the collec- tive level between surveyed and modelled (reported) preferences is 6.5 percentage points, and '[s]hifts in collective preferences [were] large enough to change what appears to be majority or plurality consensus on an issue occur quite frequently' (126), and 'correcting for information effects changes collective preferences in nearly half of governance questions' (128).

In terms of vote choice in particular, Larry Bartels (1996) finds that, on average, 'Democrats do almost two percentage points better and incumbents do almost five percentage points better than they would if all voters in presidential elections were, in fact, fully informed' (220). Andr´e Blais et al. (2009) simulate the outcome of six past Canadian elections, and see a likely difference in outcome in one, and an average information effect of 2.3 percentage points across parties and elections. Henrik Oscarsson (2007) simulates six past Swedish elections, with a likely electoral difference in outcome in two under a fully informed electorate, and an average net gain of 2.7 percentage points for right parties. 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, and two instances where the difference between actual and simulated support exceeds ten percentage points (with a social democratic party losing out, and a liberal one gaining).

Return to the revisionist counterfactual with these results in mind. In each case, these infor- mation effects suggest that we likely would have voted differently and held different preferences, had we been fully informed (but otherwise been identical across measured covariates) – contrary to the revisionist narrative.

2. WHAT POLITICAL KNOWLEDGE IS AND HOW TO MEASURE IT

So, the presence of information effects is bad news for the revisionist. However, any estimate of informed political preferences will only be as good as the conception and measure of political knowledge that it relies on. What, then, is political knowledge? Delli Carpini and Keeter offer a helpful gloss on its domain:

> [...] a general familiarity with (1) the rules of the game (the institutions and pro- cesses of elections and governance); (2) the substance of politics (the major domestic and international issues of the day, current social and economic conditions, key political initiates, and so forth); and (3) people and parties (the promises, performances, and attributes of candidates, public officials, and the political parties) is critical to the maintenance of a healthy democracy (Delli Carpini and Keeter 1996, 14).

This makes clear that political knowledge is a resource: having a lot of it means being in a position to navigate the political world, and to stand a better chance of connecting your fundamental political goals with successful means – e.g., by voting for the candidate or party that is best place to realise those goals.[2]

How, then, do we measure political knowledge? The standard way is to ask a number of factual questions about what government is and does – along the lines of Delli Carpini and Keeter's three-part gloss above – and then add up the number of correct answers for a 'knowledge score.' A couple of points are in order:

- Why think these tests measure anything? This is fairly easy to establish on formal grounds, since we in any given instance can test for internal consistency (the extent to which items co-vary) and unidimensionality (whether the items likely measure a single latent trait).
- Why think they measure political knowledge specifically? If political knowledge is a resource, these scales should find more of it among the privileged and less among the disadvantaged, which turns out to be the case (e.g., Althaus 2003, 135; see also Delli Carpini and Keeter 1996, 177).
- What topics do the items need to cover? Since people tend to be generalists – i.e., those who know a lot in one political area tend to know a lot in other areas – we 'need not be overly concerned with the mix of specific topics covered by individual items' (Delli Carpini and Keeter 1996, 174).

- Do the scales need to be long? Tests with as few as five items tend to perform as well as longer ones (Delli Carpini and Keeter 1996 and 1993).

To get a sense of how this works in practice, as it relates to the type of information effects considered in the previous section, consider the six, true/false knowledge items in Table 1, from the March 2019 (wave 15, N = 30,842) panel survey of the British Election Study.

Item
1. Polling stations close at 10.00pm on election day.
2. No-one may stand for parliament unless they pay a deposit.
3. Only taxpayers are allowed to vote in a general election.
4. The Liberal Democrats favour a system of proportional representation.
5. MPs from different parties are on parliamentary committees.
6. The number of members of parliament is about 100.

Table 1: The six general knowledge items from Fieldhouse et al. 2020.

2] See also Zaller 1992 on how the politically informed tend to be the most ideologically coherent, a point that we will nd reason to return to below.

Looking at the response patterns in this particular data set, four of the six items (2, 4, 5, and 6) make for a scale that both has a high internal consistency and is likely unidimensional, i.e., measuring a single, latent feature, as given by a factor analysis.[3]

Moreover, looking at the proportion of respondents answering all four question correctly, we see the patterns we should expect to see, if political knowledge is a resource: in particular, Figure 1 shows that those with higher education, higher income, of older age, and who report being more likely to vote are over-represented among the politicall

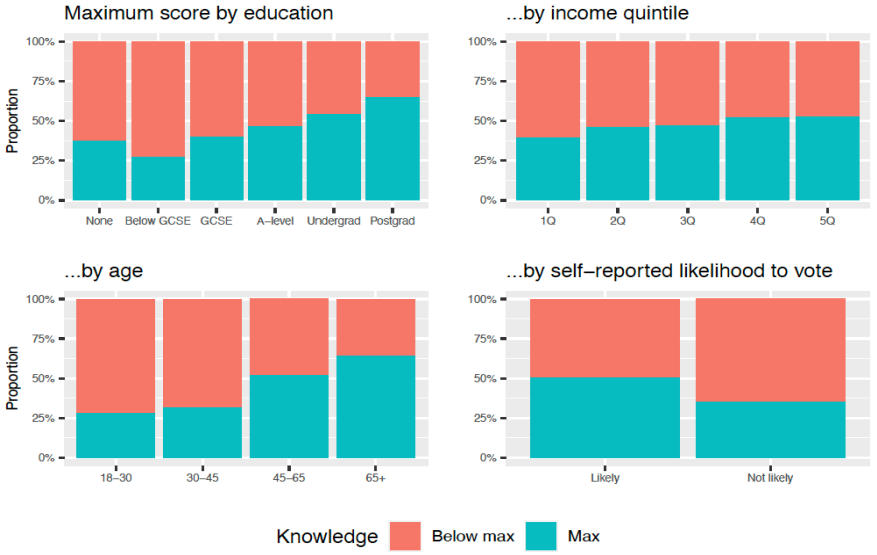

Figure 1: Proportion of respondents in Fieldhouse et al. (2020) with a maximum knowledge score by education, income, age, and self-reported likelihood to vote in the general election.

So, it looks like we are measuring something, and that that something moreover is plausibly thought of as political knowledge. But the revisionist will want to know whether what we are measuring makes a difference to political preferences and choice: some might have more knowledge than others, but perhaps those with less do just fine.

To evaluate this, remember CHANGES: if knowing more would have changed my pref- erences/choices, I do not (presently) know enough. As we have seen, such (counterfactual) differences are measured by information effects. For illustration, Figure 2 shows what happens if we go through the three steps in Section 2 with regards to respondents' vote in the UK's 2016 EU referendum, from the above BES data set.[4]

Given the fairly substantial information effect (4.3 percentage points), when compared to the size of typical information effects (see Section 1), it looks like a

3] The standardized (Chronbach's) alpha for the scale is 0.772. The factor analysis was performed using fa.parallel in R's psych package (Revelle 2018).

4] For full information about the modeling involved, see the Appendix in Section 7.

sizeable proportion in the sample did not have enough political knowledge – and that the knowledge measured does make a difference.[5]

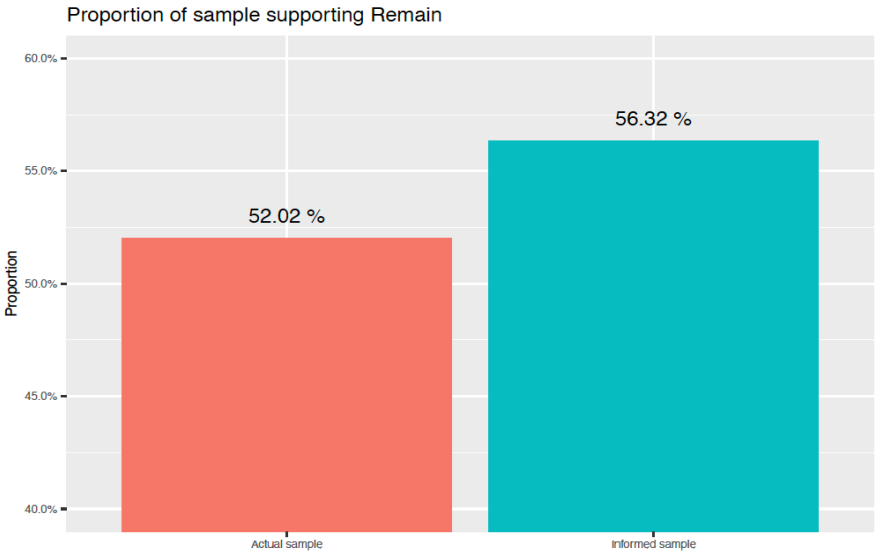

Figure 2: Proportion of respondents in Fieldhouse et al. (2020) supporting 'Remain' as opposed to 'Leave' in the 2016 EU referendum in the actual sample (left) and as estimated by a logistic model for a fully informed sample (right), holding constant gender, age, income, ethnicity, and education.

3. ARE KNOWLEDGE SCALES ELITIST?

We are now in a position to evaluate one of the most forceful, revisionist lines on knowledge scales: Lupia's charge that these scales are driven by elitism. He explains:

> Most political-knowledge questions are not derived from a replicable or transparent logic about how their answers bear on a voter's ability to make decisions of a particular quality in the voting booth. Instead, the questions are generated by a worldview that is shared by a select set of academics, journalists, and politicos, but few others. [...] The elitist move is to assume that these questions have a similar value to citizens whose societal responsibilities can be very different than their own (Lupia 2006, 219).

5] Someone might object that the knowledge tests involved simply measure a form of technical political knowl-edge, as opposed to what we might call situated political knowledge, such as knowing what it is like to go hungry for days on end, or to be on the receiving end of racism. But note that, by controlling for standard demographic variables (e.g., income, gender, ethnicity, etc.) with which these forms of knowledge likely travel, the type of model relied on here preserves such situated knowledge in estimating informed preferences. Specifically, it can be thought of as modeling what you would have preferred, had you had all of the situated knowledge you currently have (e.g., on account of a direct experience with hardship and injustice), and also been informed in the sense of having technical political knowledge (e.g., by being highly informed on how the political world operates, and how to navigate it).

I will suggest that Lupia goes wrong on two points. First, he confuses testing with educating for political knowledge. Second, he disregards evidence that what is tested for is not mere political trivia. Let us consider these points in turn.

3.1 Confusing testing with education

Lupia asks us to imagine that we have identified some set of facts, A-Z, and we want to determine whether knowing these facts is necessary for competent voting (task t).

> To answer [this] question about necessity, we must ask whether there is a different set of facts, perhaps even a subset of facts A-Z, that also allows the voter to accomplish t? If the answer is no, then knowledge of every fact from A to Z is necessary for the voter to choose competently. In such a case, we can assess the voter's political competence reliably by quizzing her about A through Z – as with the political-knowledge tests on which so many analysts base their judgments of voter incompetence. If we find her deficient in her knowledge of even one of these facts, we can accurately judge her incompetent at task t. If we want to increase her competence at this task, moreover, we know that a precondition of success is providing her with the knowledge of all such facts about which the quiz reveals her ignorance (Lupia 2006, 222).

Lupia is right in this: if there is some set of facts necessary for performing a task, someone can be deemed incompetent at that task if they turn out not to know some of those facts. And we can make her more competent by imparting those facts in her. Where he goes wrong is in assuming that this is how political knowledge tests work. He continues:

> The problem with this approach to assessing voter competence is that it is validity depends on establishing that facts A-Z are necessary for competence: i.e., that knowledge of no subset of these facts, or alternative set of facts, would suffice for task t to be accomplished. However, if facts other than the full set A-Z are sufficient for citizens to accomplish t, then knowing A-Z cannot be a necessary condition for competence at t. Thus, merely demonstrating that a voter does not know these facts may reveal little or nothing about her competence in the voting booth (Lupia 2006, 222).

Lupia's point here makes sense if we are looking to educate for political knowledge: if there are things everyone needs to know (necessary), these must be identified and imparted. If we succeed, the person will have become more competent; and if we fail, she will not. But knowledge tests are not educational tools. They are – naturally – constructed to test people. This is done in terms of what is diagnostic of political knowledge, and what is diagnostic might be neither necessary nor sufficient for such knowledge. Table 2 illustrates this point, using an analogy with medical diagnostics.

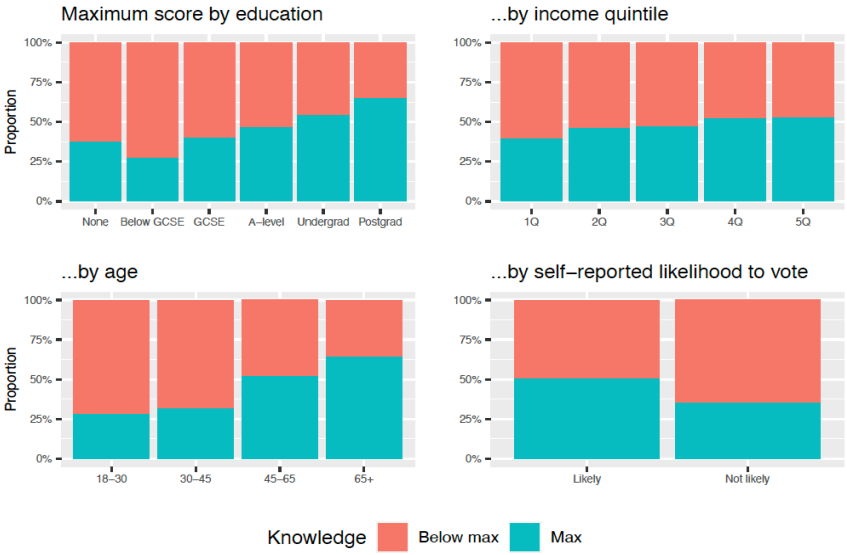

Table 2: Medical diagnostics and political knowledge cases compared.

So, there are things that might be diagnostic of some target property, even if neither neces- sary nor sufficient for having the relevant target property. Knowing that (say) there are more than 100 MPs in the Commons, or that the Liberal Democrats favour a system of proportional representation, is neither necessary nor sufficient for being a competent UK voter. Still, these are things that, if you keep up with politics and are politically informed, you will likely know these things; and if you do not, then you likely won't.

3.2 Testing for non-trivial knowledge

Lupia might grant the above point about testing versus educating, but respond that what typical knowledge items are diagnostic of is not political knowledge, but political trivia. As he writes: 'what benefit does a randomly selected citizen draw from knowing something like the name of the Chief Justice [of the Supreme Court]?' (218).

Part of the answer to this question was provided in the previous section: an item on a knowledge scale does not earn its keep for being necessary for political competency, but for being diagnostic of such competency. Moreover, two further things should be pointed out here As already noted, short knowledge scales of the kind discussed earlier correlate well with longer scales (e.g., Delli Carpini and Keeter 1996 and 1993). Beyond that, you would be hard pressed to maintain that the collection of items usually found on such longer scales are mere trivia, as opposed to substantive facts relevant to sensible political action. By way of illustrating this point, see Table 3 for the twenty-item scale from Delli Carpini and Keeter 1996, which is representative of subsequent work on political knowledge (e.g., in Althaus 2003; see also Bennett 2003).

People	Party	Civics
1. Quale	9. Relative ideological location of the two parties.	15. Times a president can be elected
2. Gorbachev	10. Party with most seats in the House	16. Whose responsibility is judicial review
3. Thatcher	11. Relative location of parties on defense spending	17. Whose responsibility is it to nominate federal judges
4. Name one candidate (and his/her party) for US House	12. Party with most seats in the Senate	18. What are the first ten amendments called
5. Mandela	13. Relative location of parties on federal spending	19. What majority is needed to override presidential veto
6. Foley	14. Relative location of parties on aid to blacks	20. How long is a senator's term
7. Rehnquist		
8. Mitchell		

Table 3: The twenty items from Delli Carpini and Keeter's 1996 (302) long knowledge scale, drawn from the 1990-91 American National Election Studies survey.

4. DOES POLITICAL KNOWLEDGE BREED DOGMATISM?

A more fundamental objection to the utility of political knowledge and its measurement takes the line that such scales measure political knowledge, but that any benefits of political knowledge will be canceled out by such knowledge also making us dogmatic.

This line has been pursued recently by Michael Hannon (2021). Hannon starts by making the observation that the politically knowledgeable tend to be more partisan, in the sense of having the most stable and worked-out ideological positions (e.g., Zaller 1992). That, in itself, is neither surprising nor necessarily worrying. What is worrying, according to Hannon, is that the politically knowledge – and thereby most partisan – also are more susceptible to motivated reasoning than the rest of us. Motivated reasoning involves pursuing and processing information in ways that systematically favor your preferred viewpoints, and often in ways that serve to 'protect' beliefs and convictions that are central to your group identity. For example, when it comes to politically motivated reasoning in particular, it involves 'the formation of beliefs that maintain a person's status in [an] affinity group united by shared values' (Kahan 2016, 2).

Why think that the politically knowledgeable are more susceptible to politcally motivated reasoning than the rest of us? Hannon's primary piece of evidence involve a study by Taber and Lodge (2006). Taber and Lodge measure political knowledge using a "17-item general political knowledge scale" (760), and contend that "the politically knowledgeable, because they possess greater ammunition with which to counterargue incongruent facts, figures, and arguments, will be more susceptible

to motivated bias than will unsophisticates" (757). Taber and Lodge clearly have a particular causal story in mind: political knowledge is the mechanism that gives rise to increased levels of motivated reasoning, on account of the politically knowledgeable having "greater ammunition" with which to shoot down any challenges coming from outside their partisan fortress. Hannon is more careful on this point, suggesting instead that "partisanship is the common cause of both acquiring political knowledge and motivated reasoning" (8). This matters for whether we can expect making people more politically informed will mean also making them more partisan and/or dogmatic, in systematically rejecting viewpoitns coming from "the other side," so to speak.

For purposes of argument, let us assume Taber and Lodge's causal story, whereby the politically knowledgeble are dogmatic because they are knowledgeable. As far as knowledge scales are concerned, this raises questions, not about whether such scales are diagnostic, but whether they are measuring anything that we should have any interest in promoting. After all, while it might sound desirable to raise the level of political knowledge in society, it is not clear that we should want to do that if it will invariably lead to increased levels of dogmatism.

But this, in turn, raises the more fundamental question whether dogmatism is invariably a bad thing, from an epistemic point of view. Dogmatism – roughly, the unwillingness to engage with the perspectives and viewpoints of those with whom one disagrees – is a belief-forming disposition, and the epistemic status of such dispositions are a function of the environments in which they are employed. For example, just like blindly deferring to others might be a virtue if one is surrounded by knowledgeable sources (Ahlstrom-Vij 2015), dogmatically refusing to engage with others might be a virtue if one is surrounded by ignoramuses (Battaly 2018).

In fact, we can investigate in a more systematic fashion the implications of being highly knowledgeable, and as such (we are assuming) also highly dogmatic, versus being less knowl- edgeable using a simulation study. Specifically, let us set things up as follows:

- Assume some set of agents – say, 500 – and that each agent has a partisan affiliation, called "green" and "blue." These are assigned randomly, with a probability of 50%.
- As part of the simulation, each agent encounters 100 pieces of information in sequence. Some of these are true, and some of them are false. We code true ones as '1' and false ones as '-1,' and then quantify what an agent has learned at the end of the process by taking the sum of what they've taken on board. So, if you encounter 100 pieces of false information, you will come out at -100, designating having become maximally misled, relative to the body of information. If you encounter 100 true ones, you'll come out at 100, designating that you've come out maximally informed, relative to the body of information.

- Each piece of information is also "tagged" with a partisan cue. We can think of this as the source of the information being a representative of either "green" or "blue."

- Assume also that each agent either has a high level of knowledge knowledge or a low level of knowledge. These labels are assigned to agents randomly, with a probability of 50%. Crucially, the only difference between these two types of agents is that low knowledge agents take on board all information, while high knowledge agents only take on board those "tagged" with their own partisan label.

In other words, what we are simulating is two types of agents being faced with 100 pieces of information in sequence, where each piece of information is "tagged" with a partisan affiliation, and the politically knowledgeable agents using their increased sophistication to "shoot down" and as such reject all information that is affiliated with the other party.

Under these assumptions, what will be the effect of that dogmatism on where high and low knowledge agents will find themselves, in regards to their sum of informedness, as a result of having encountered the 100 pieces of information? This will partly be a function of the base rate of true and false information in the body of evidence as a whole (i.e., of the base rate). So let us run our simulation for 101 such base rates, ranging from 0% true (i.e., 100% false) up to 100% true. Figure 3 shows what happens.

Figure 3: Change in level of informedness as a function of the base rate of truth in the environment for high and low knowleddge agents, after having encountered 100 pieces of information.

Along the x-axis, we have the base rate of truth. On the y-axis we have the mean change in the level of informedness that the respective group of agents – high and low knowledge – has at the end of encountering all of the information, given a base rate. In other words, What we are modelling is the diachronic fact of agents encountering information, and then seeing where they end up. And where do they end up? Look at low knowledge agents first. Here the line pretty much follows the diagonal. That makes

sense: since low knowledge agents take on board all information they encounter, the change in their level of informedness at the end of the process should track the base rate.

But now look at the high knowledge agents. They do better in "hostile" epistemic envi- ronments, for the simple reason that they reject some information in a situation where lots of information is false. And they do worse in "friendly" environments for the same reason (i.e., they miss out on true information). So the mere fact of dogmatism is epistemically beneficial in hostile epistemic environments, even if the dogmatism itself is epistemically neutral. After all, agents are not rejecting messages because of their epistemic merit, but simply because they are coming from the "wrong" party. Specifically, high knowledge agents lose less knowledge than do low knowledge agents in hostile epistemic environments, and gain less in friendly ones.

This, of course, is compatible with high knowledge agents still knowing more across both types of environments after that loss/gain, simply on account of knowing more to begin with. What their dogmatism does in hostile enviromnet is to protect their relative level of epistemic superiority vis-a-vis low knowledge agents. Heather Battaly makes this same observation when reflecting on why dogmatism can in some cases qualify as an epistemic virtue:

> So, what is a knowledge-possessing agent to do when she finds herself in an epistemically hostile environment? My proposal is that there are epistemic reasons for her to be closed-minded – to be unwilling to engage seriously with relevant intellectual options that conflict with what she already knows. [...] Why should she be closed-minded? Because, in an epistemically hostile environment, closed-mindedness is an effects-virtue. When a knowledge-possessing agent is stuck in an epistemically hos- tile environment, surrounded by falsehoods, incompetent sources, and diversions, closed-mindedness about options that conflict with what she knows will minimize the production of bad epistemic effects for her (Battaly 2018, 39).

This offers an intriguing response to Hannon's challenge to the benefits of political knowl- edge. Because even if the politically knowledgeable are maximally dogmatic, as in the above simulation, and even if it is their knowledge that makes them dogmatic (as per Taber and Lodge's suggestion), this does not cast doubt on the epistemic merits of being politically knowl- edgeable. On the contrary, such dogmatism serves to protect the knowledgeable in hostile environments – and most political epistemic environments are, after all, highly likely to be ex- actly that, on account of being marked by a high proportion of falsehoods, lies, and misinforma- tion. Consequently, being politically knowledgeable remains preferable to not being politically knowledgeable even under those assumptions.[6]

6] Thinking back to the problem about ignorant non-voting perpetuating unjust patterns in society, it might be objected that the protection offered the politically knowledgeable by their dogmatism will be to the detriment of the already marginalized. (Many thanks to Fabienne Peter for making this point.) However, note two things. First, there is some evidence to suggest that the more knowledgeable tend to be more progressive on political matters (e.g., Althaus 2003), suggesting that the knowledgeable generally make political decisions that would benefit not just them, but also the marginalized. Second, even if that

5. CONCLUSION

Political knowledge can be thought of as a resource: having a lot of it means being in a position to navigate the political world, and stand a better chance of connecting your fundamental political goals with successful means. The present piece argued that standard political knowledge tests measure political knowledge, so understood, and used counterfactual modeling to demonstrate the difference having such knowledge can make to political choice. It then took up to two of the most forceful objections to political knowledge and its measurement, by rejecting the idea that knowledge scales encode elitist assumptions, and arguing that, even if political knowledge breeds dogmatism, such dogmatism can be expected to serve a protective function in exactly the type of hostile epistemic environments – filled with lies, falsehoods, and misinformation – that make up the political domain.

ahlstromvij@gmail.com

REFERENCES

Achen, Christopher, and Larry M. Bartels. 2016. *Democracy for Realists: Why Elections Do Not Produce Responsive Government*. Princeton: Princeton University Press.

Ahlstrom-Vij, Kristoffer. 2015. The Social Virtue of Blind Deference. *Philosophy and Phenomenological Research* 91(3): 545-82. https://doi.org/https://doi.org/10.1111/phpr.12158.

———. 2020a. The Case for Modelled Democracy. *Episteme*: 1-22. https://doi.org/10.1017/epi.2020.10.

———. 2020b. The Epistemic Benefits of Democracy: A Critical Assessment. In T*he Rout- ledge Handbook of Social Epistemology*, edited by Fricker Miranda, Graham Peter, Henderson David, and Nikolaj Pedersen. Routledge.

Althaus, Scott L. 2003. *Collective Preferences in Democratic Politics: Opinion Surveys and the Will of the People*. Cambridge University Press. https://doi.org/10.1017/CBO9780511610042.

Bartels, Larry M. 1996. Uninformed Votes: Information Effects in Presidential Elections. *American Journal of Political Science* 40(1): 194-230. http://www.jstor.org/stable/2111700.

Battaly, Heather. 2018. Can Closed-Mindedness Be an Intellectual Virtue?. *Royal Institute of Philosophy Supplement* 84: 23-45. https://doi.org/10.1017/S135824611800053X.

Bennett, Stephen Earl. 2003. Is the Public's Ignorance of Politics Trivial?. *Critical Review* 5(3-4): 307-37. https://doi.org/10.1080/08913810308443585.

Bhatti, Yosef. 2010. What Would Happen If We Were Better Informed? Simulating In- creased Knowledge in European Parliament (EP) Elections. *Representation* 46(4): 391-410. https://doi.org/10.1080/00344893.2010.518078.

Blais, Andre, Elisabeth Gidengil, Patrick Fournier, and Neil Nevitte. 2009. Information, Visibility and Elections: Why Electoral Outcomes Differ When Voters Are Better

---

were not true, the appropriate response would be, not to want to see a reduction in knowledge among those who have it, but an increase among those who don't.

Informed. *European Journal of Political Research* 48(2): 256-80. https://doi.org/https://doi.org/ 10.1111/j.1475-6765.2008.00835.x.

Brennan, Jason. 2016. *Against Democracy*. Princeton: Princeton University Press.

Delli Carpini, Michael X., and Scott Keeter. 1993. Measuring Political Knowledge: Putting First Things First. *American Journal of Political Science* 37(4): 1179-1206. http://www.jstor.org/stable/2111549.

———. 1996. *What Americans Know about Politics and Why It Matters*. New Haven: Yale University Press.

Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York: Harper.

Fox, John, and Sanford Weisberg. 2011. *An R Companion to Applied Regression*. Second edition. Thousand Oaks CA: Sage. http://socserv.socsci.mcmaster.ca/jfox/Books/Companion.

Hannon, M. (2022). Are knowledgeable voters better voters? *Politics, Philosophy & Economics* 21(1), 29-54. https://doi.org/10.1177/1470594X211065080

Kahan, Dan M. 2016. The Politically Motivated Reasoning Paradigm, Part 1: What Politically Motivated Reasoning Is and How to Measure It. In *Emerging Trends in the Social and Behavioral Sciences*, 1-16. American Cancer Society. https://doi.org/https://doi.org/10.1002/9781118900772.etrds0417.

Lupia, Arthur. 2006. How Elitism Undermines the Study of Voter Competence. *Critical Review* 18(1-3): 217-32. https://doi.org/10.1080/08913810608443658.

Morgan, Stephen, and Christoper Winship. 2015. *Counterfactuals and Causal Inference : Methods and Principles for Social Research*. Cambridge University Press.

Oscarsson, Henrik. 2007. A Matter of Fact? Knowledge Effects on the Vote in Swedish General Elections, 1985-2002. *Scandinavian Political Studies* 30(3): 301-22. https://doi.org/https://doi.org/10.1111/j.1467-9477.2007.00182.x.

Popkin, Samuel. 1991. T*he Reasoning Voter : Communication and Persuasion in Presidential Campaigns*. Chicago: University of Chicago Press.

R Core Team. 2018. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. https://www.R-project.org/.

Revelle, William. 2018. *Psych: Procedures for Psychological, Psychometric, and Personality Research*. Evanston, Illinois: Northwestern University. https://CRAN.R-project.org/package=psych.

Somin, Ilya. 2016. *Democracy and Political Ignorance: Why Smaller Government Is Smarter*. Stanford, California: Stanford Law Books, an imprint of Stanford University Press.

Taber, Charles S., and Milton Lodge. 2006. Motivated Skepticism in the Evaluation of Political Beliefs. *American Journal of Political Science* 50(3): 755-69. https://doi.org/https://doi.org/10.1111/j.1540-5907.2006.00214.x.

Zaller, John. 1992. *The Nature and Origins of Mass Opinion*. Cambridge England New York, NY, USA: Cambridge University Press.

7. APPENDIX

7.1 Models and diagnostics

The modelling performed in Section 2 used 'doubly robust' estimation for counterfactual infer- ence (e.g., Morgan and Winship (2015)). The 'double robustness' owes to how counterfactual effects – in our case, what position someone would take, had they been fully informed – are es- timated in a context where we have both controlled for (assumed) confounds, and improved the balance between the two groups (here: those with a maximum score on the knowledge test and everyone else) to make up for the fact that they have not come about as a result of randomized assignment. In the present case, this second layer of 'robustness' was achieved through so-called 'propensity scores.'

In our case, propensity scores measure the probability (or propensity) that an observation will be found in the 'informed' category (i.e., attaining a maximum score on the knowledge test), as a function of someone's demographic features. The idea is to then use these scores to remove any correlation between these features and the 'informed' category, to justify a counterfactual inference. Why? It is helpful to think of this on the model of a randomized 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 will (for example) be the case that 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.

Specifically, we can use the inverse of those scores as weights (such that an observation with a low propensity is weighted heavily, and vice versa) in fitting the model. This improves the balance between the two groups, and – given an appropriately chosen set of covariates when calculating the scores – recreates a situation that would have been expected in a randomized experiment, thereby allowing greater confidence in inferring a counterfactual.

In the present case, the propensity scores were calculated by way of logistic regression, using glm in R's stats package (R Core Team (2018)). By way of illustration, consider Figure 4, where the left-hand panel shows the balance (or rather: lack thereof) for the income variable prior to applying the propensity weights – it can be seen that informed participants (teal bars, designated here as 1 or 'treated') are over-represented among the wealthy (higher quintiles), and underrepresented among the less wealthy (lower quintiles) – and the right-hand panel the balance achieved once the weights had been applied.

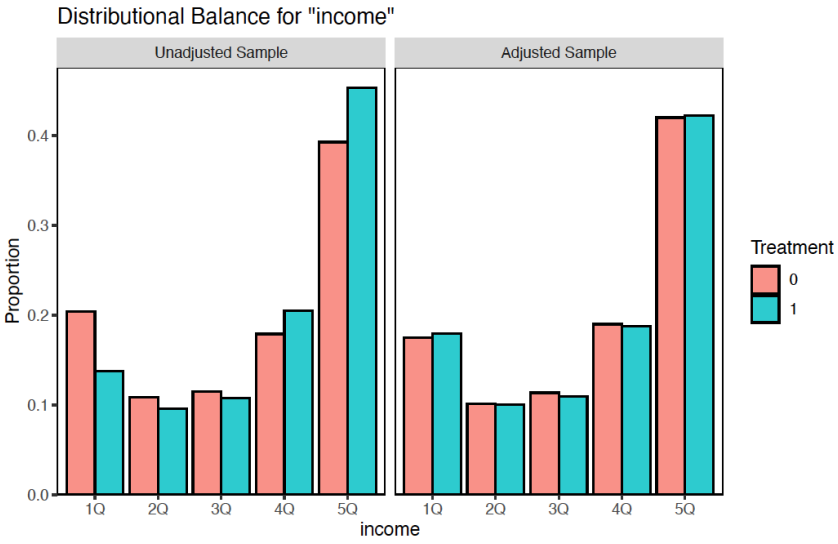

Figure 4: Balance plots for income quintiles, before and after applying propensity weights/

Using the propensity scores as weights, a logistic regression model was then fitted, again using glm. Table 4 contains details on the coefficients and coefficient values of the logistic model used, designated 'Model 1' in that table. The highest variance inflation factor value for that model was 1.187266, as given by vif in the car package (Fox and Weisberg 2011), providing no evidence of multicollinearity. There also was no indication of extreme values as measured by Cook's distance, but 431 observations had a standardized residual greater than 3. In light of these potentially influential observations, the model was refitted without those observations.

This is the model designated 'Model 2' in Table 4. If we simulate the informed outcome of the 2016 referendum on that model, we get an (increased) information effect of 6.54 percentage points (58.56% for Remain).

	Model 1	Model 2
Intercept	0.80 (0.08)***	1.11 (0.09)***
Knowledge	0.43 (0.02)***	0.63 (0.02)***
Gender (Male)	−0.23 (0.02)***	−0.43 (0.02)***
Income (2Q)	0.10 (0.04)*	0.20 (0.05)***
Income (3Q)	0.14 (0.04)***	0.22 (0.05)***
Income (4Q)	0.13 (0.04)***	0.25 (0.04)***
Income (5Q)	0.37 (0.03)***	0.50 (0.04)***
Ethnicity (White)‖	−0.25 (0.06)***	−0.35 (0.07)***
Education (No education)	−0.55 (0.05)***	−0.69 (0.06)***
Education (Below GCSE)	−0.37 (0.06)***	−0.46 (0.07)***
Education (GCSE)	−0.47 (0.03)***	−0.51 (0.04)***
Education (Undergraduate)	0.44 (0.03)***	0.45 (0.03)***
Education (Postgraduate)	0.96 (0.04)***	1.34 (0.05)***
Age (30-45)	−0.51 (0.05)***	−0.61 (0.06)***
Age (45-65)	−1.10 (0.05)***	−1.43 (0.05)***
Age (65 and over)	−1.30 (0.05)***	−1.77 (0.05)***
AIC	44319.24	38027.94
BIC	44444.52	38152.85
Log Likelihood	−22143.62	−18997.97
Deviance	46431.36	40177.62
Num. obs.	18591	18160

***$p < 0.001$; **$p < 0.01$; *$p < 0.05$.

Table 4: Logistic models

7.2. Simulation

The simulation in Section Section 4 relied on the following R script:

```
set.seed(100)
n <- 500
info_pieces <- 100

# generate results dataframe
results1 <- data.frame(base_rate = (0:100)/100, high = NA, low = NA)

# high knowledge agents reject any information tagged with
# the wrong partisan affiliation low knowledge ones take it
# all on board
for (k in 1:101) {
    # generate n agents with level of knowledge and partisan
```

```
# affiliation
agents <- data.frame(knowledge = sample(c("high", "low"),
    n, replace = TRUE), party = sample(c("blue", "green"),
    n, replace = TRUE), success = NA)

# generate some information, tagged with partisan affiliation
information <- data.frame(info = sample(c(-1, 1), info_pieces,
replace = TRUE, prob = c(1 - results1$base rate[k], results1$base rate[k]
party cue = sample(c("blue", "green"), info pieces, replace = TRUE))

for (i in 1:n) {
    temp <- NA
    for (j in 1:info_pieces) {
        if (agents$knowledge[i] == "high") {
            ifelse(information$party_cue[j] == agents$party[i],
              temp <- append(temp, information$info[j]),
              NA)
        }
        if (agents$knowledge[i] == "low") {
            temp <- append(temp, information$info[j])
        }
    }
    agents$success[i] <- sum(temp, na.rm = TRUE)
}
results1$high[k] <- mean(agents$success[agents$knowledge ==
    "high"])
results1$low[k] <- mean(agents$success[agents$knowledge ==
    "low"])
print(k)
```

ACKNOWLEDGEMENTS

Many thanks to the participants at the 2020 ECPR Summer School in Political Epistemol- ogy, the 2020 Leeds Epistemic Normativity conference, and the Epistemic Circumstances of Democracy conference at the University of Rijeka in 2021 for valuable comments. I'm particularly grateful to Michael Hannon, Sorin Baiasu, Jakub Szczepan´ski, Fabienne Peter, Carline Kiljnman, Darren Bradley, and Mona Simionescu for discussion.