



BIROn - Birkbeck Institutional Research Online

Ahlstrom-Vij, Kristoffer (2022) The case for modelled democracy. *Episteme* 19 (1), pp. 89-110. ISSN 1742-3600.

Downloaded from: <https://eprints.bbk.ac.uk/id/eprint/30773/>

Usage Guidelines:

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

or alternatively

THE CASE FOR MODELLED DEMOCRACY

Kristoffer Ahlstrom-Vij

Department of Philosophy

Birkbeck College, University of London

k.ahlstrom-vij@bbk.ac.uk

Abstract: The fact that most of us are ignorant on politically relevant matters presents a problem for democracy. In light of this, some have suggested that we should impose *epistemic* constraints on democratic participation, and specifically that the franchise be restricted along competency lines—a suggestion that in turn runs the risk of violating a long-standing condition on political legitimacy to the effect that legitimate political arrangements cannot be open to reasonable objections. The present paper therefore outlines a way to solve the problem of public ignorance without restricting the franchise. The proposal involves filtering the electoral input of a universal franchise through a statistical model that simulates what the public’s political preferences would have been, had they been informed on politically relevant matters. The result is *modelled democracy*. A case is made that such democracy both solves the problem of public ignorance and satisfies the aforementioned condition on legitimacy.

1. Introduction

The fact that most of us are ignorant on politically relevant matters is well established (see, e.g., Delli Carpini and Keeter 1996). Indeed, as pointed out by Jeffrey Friedman (1998), it is ‘one of the strongest findings that have been produced by any social science—possibly *the* strongest’ (397). Some have argued that we cannot infer from widespread public ignorance that there is a problem for democracy, because a large public is likely to outperform a smaller number of more competent individuals (e.g., Condorcet 1785/1785; Converse 1990; Hong and Page 2004; Landmore 2013); or because of the epistemic benefits of deliberative or liberal institutions (e.g. Buchanan 2004; Gutmann and Thompson 2004; Talisse 2005). While a full treatment of these possibilities would take us too far afield, I have already argued elsewhere (Ahlstrom-Vij

forthcoming) that they fail to throw doubt upon the idea that widespread public ignorance has a detrimental effect on the quality of democratic decision-making.

This, moreover, reinforces a long-standing tradition stretching all the way back to Plato's *Republic*, looking to impose epistemic constraints on democratic participation. John Stuart Mill (2008/1861) famously defended such constraints in the form of plural votes for the educated. However, David Estlund (2008) argues that Mill's proposal violates the *qualified acceptability criterion* (QAC), a necessary condition on political legitimacy that holds that, if a political arrangement is open to objections that are not unreasonable, it is not politically legitimate. According to Estlund, it is not unreasonable to maintain that the educated will not be better at ruling than the uneducated, since the educated might be coextensive with people sharing some common set of biases on account of being wealthy, white or privileged. This is the *demographic objection* against epistocracy, or rule by the knowledgeable.

More recently, Jason Brennan (2011) has defended an epistocratic system that would 'restrict electoral power to citizens who can demonstrate competence' (701) with reference to a *competence principle*, protecting 'a citizen's right not to be subject to high stakes decisions made by incompetent and morally unreasonable people' (702). But Brennan grants that his proposal violates the QAC, a 'defining feature of liberal political philosophy' (714). A majority of political philosophers would agree with Brennan here (although see Mulligan 2015). This makes for an unfortunate impasse, and motivates my question: *Is it possible to solve the problem of widespread public ignorance, without restricting the franchise?* I will argue that it is. My proposal builds on Scott Althaus's (2003) work on 'fully informed preferences,' and involves filtering the universal franchise through a statistical model that simulates what the public's preferences would have been, had they been informed on politically relevant matters.¹ This is *modelled democracy*. I will make the case that such democracy both meets the challenge posed by widespread public ignorance and satisfies the QAC.

¹ Brennan (2016: 220-222) briefly mentions a similar approach and refers to it as 'government by simulated oracle.' However, not only does Brennan not spell out the approach in question; he also categorizes it as a form of epistocracy which, for reasons that will become clear in what follows, is a mistake.

That said, it should be acknowledged at the very outset that the proposal is radical, controversial, and incomplete. It is radical since its implementation would require substantial changes to current political arrangements, and controversial in relying on philosophical theories—for example, about the relationship between choices and informed preferences—that are not universally accepted. And it is incomplete on account of how its radical nature and controversial assumptions warrants a more extensive treatment than can be offered in a single paper. But the purpose of the following is not to establish that we ought to implement modelled democracy. It is simply to offer a convincing enough case for it to be clear that the proposal warrants further investigation. In that respect, we can think of this as a pilot study for a more complete case for modelled democracy.

2. Modelling Fully Informed Preferences

As already noted, the proposal to be defended involves keeping the franchise universal, but filtering the electoral input through a statistical model that simulates what the public's preferences would have been, had they been informed on politically relevant matters. There are a number of ways to estimate what the distribution of informed preferences among the public would be, save for through simulation. Let's motivate a simulation approach by considering some problems for the two most prominent alternatives.

2.1. Two Alternatives

One way to try to isolate an informed public is by using filtering questions in opinion surveys (e.g., Bishop, Oldendick, and Tuchfarber 1983). The relevant questions are challenging enough to significantly increase the number of 'Don't Know' responses, which are then ignored, leaving the remaining responses as a proxy for a more informed response set. However, as pointed out by Althaus (2003), '[t]hose who tend to be filtered away will come predominantly from certain demographic groups whose subjective needs, wants, and values may not be adequately represented by the remaining opinion givers' (100). So, whether or not the opinion givers surviving the filter are informed, they are most likely not *representative*.

Another approach proceeds experimentally, by immersing a representative sample in high-quality information, and then observing any changes in their reported preferences. This is the idea behind James

Fishkin's (e.g., 1999) deliberative polling. Here the problem is the opposite of that facing the filtering approach: whether or not the subjects involved in deliberative polling are representative, they are not necessarily more *informed* in the sense we are after here. As Althaus (2003: 100-101) notes, being informed is not simply a matter of having been immersed in high-quality information, but also of having attained specific cognitive styles and strategies, and acquired a rich set of political information that has implications for how new information is subsequently processed. These traits and resources take more than a couple of days to acquire and are as such not necessarily present in contexts of deliberative polling.²

2.2. The Simulation Approach

This brings us to Althaus's simulation approach, which—if successful—serves to model an informed *and* representative public. The simulation proceeds by asking people about their political preferences and collecting standard demographic information, while also testing their political knowledge through a 'civic style' knowledge test. In a US context, the test items might include questions about what roles are held by people named—e.g., the President, the Speaker of the House, Supreme Court Justices, foreign leaders—or about the constitutional powers of different federal branches, the holders of the majority in both houses of Congress, changes in the federal budget deficit, etc. People are considered politically informed to the extent that they answer the questions correctly, and *fully* informed if they answer all of them correctly.

This provides the material for fitting a logistic regression model, estimating the probability of reporting some particular policy preference as a function of a person's test score and demographics. Using that model, we can then, for each person in the data set, increase the value of her test score variable to the maximum one while leaving everything else as is, and note the models' (revised) estimate of the probability

² A similar problem presents itself for López-Guerra's (2014) proposal of an 'enfranchisement lottery', on which, prior to an election, a set of 'pre-voters' are randomly selected through a lottery, but only franchised if prepared to 'gather in relatively small groups to participate in a competence-building process carefully designed to optimize their knowledge about the alternatives on the ballot' (4). While not offering a lot of details regarding the relevant process, López-Guerra refers to Fishkin's work, and suggests that 'the enfranchisement lottery can easily be regarded as part of the family of innovations known as *mini-publics*: small, descriptively representative bodies of citizens who are asked to ponder and decide on some matter of public concern' (27; see also 36).

of her reporting some particular policy preference. Assuming a robust enough relationship between reported and actual preferences, this serves to model, for each respondent, what her political preferences on the matters surveyed *would* have been, had she been fully informed, as operationalized by the knowledge items.

Note that the model does *not* look at the policy preferences of the most informed and then takes it that those are the ones everyone would have had, had they been fully informed. Rather, it factors in differences in preference between more or less informed people sharing a demographic profile, and then estimates on the basis of that the probability that any given person of that profile would have reported some particular policy preference, had she been fully informed *while maintaining that profile*.

Finally, by aggregating those probabilities for each policy preference question, and comparing them to the actual proportion of survey respondents indicating (say) support for the relevant policies, we get a measure of *collective information effects*, or the aggregate differences between informed preferences and the preferences actually reported. Althaus 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’ (23-4).³ Across all survey questions, the average difference on the collective level between surveyed and modelled (reported) preferences is 6.5 percentage points, while the average individual-level effect is 12.8 percentage points (122). Still, ‘[s]hifts in collective preferences 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).⁴

³ Comparable effects are found in Bartels (1996) and Caplan (2007).

⁴ Can shifts of the relevant magnitude make an electoral difference? The short answer is ‘yes,’ but I will discuss the matter at some length in Section 3.2.

2.3. Two Worries

Two particular worries are relevant here. The first one is that, depending on what variables we include in the model, we are making different assumptions about what makes us who we are for purposes of identifying our political selves. As noted above, what is included in Althaus's models is people's demographic profile (broadly construed), which includes education, income, age, partisanship, race, gender, marital status, occupation, religious affiliation, union membership, homeowner status, parental status, financial status, and geographic region, as well as the type of community resided in (2003: 105). The best way to motivate this type of demographic model is by noting that it meets two desiderata:

First, we'll want to avoid a minimal model that holds *nothing* constant. Such a model would in effect simulate the average information effect for all, and arguably say nothing about what *you* (or me, or anyone else) would have been like, had you been fully informed. By including demographic variables that can be expected to influence attitudes and preferences, as well as some ideological (in a broad sense) variables such as partisanship, religious affiliation, and union membership, that ensure the constancy of central parts of your political self and broader value commitments, we avoid the problem with the minimal model.

Second, we don't want to control for *all* of the things that form part of your political self—including your attitudes—since that would rule out by fiat any information effects. For the same reason, we will also want to exclude some variables that are *highly predictive* of political preferences. This is counterintuitive, but an example might help. Variables tapping into authoritarian dispositions are highly predictive of political preferences (e.g., Hawkins *et. al* 2018). Including such variables when modelling fully informed preferences would constrain any information effects in inappropriate ways, by assuming that the degree to which someone is informed is irrelevant to the manners in which any authoritarian dispositions are expressed or manifested in attitudes. That is an implausible assumption. As Karen Stenner (2005) makes clear, the factors that predispose someone to authoritarianism are likely innate, but since 'the environment can influence the extent to which a potential predisposition develops or remains dormant' (146), the degree to which someone becomes informed early on in their political development is relevant to whether or not such activation will take place to begin with.

For these reasons, when modelling what is in effect the *counterfactual* state of someone's fully informed self, we want to include variables that are predictive of attitudes and preferences, but not ones that will inappropriately constrain any information effects.⁵ The type of demographic model under consideration balances these desiderata well. Of course, questions might still be raised about whether some particular, relevant variables are left out of such a model, and we will return to this matter when we discuss model robustness in Section 3.1, and unknown bias in Section 8.

Turning to the second worry, what reason do we have for believing that being 'fully informed,' in the sense at work in the relevant models, corresponds to being informed in some philosophically interesting sense? Three considerations are relevant here:

First, the tests involved typically concern several matters, such as the holders of key political positions at home and abroad, and the basics about the political platforms of the main political parties, all of which would seem crucial to making anything resembling informed political choices. Moreover, to the extent that the tests concern more esoteric matters, knowing the answer to such questions is likely correlated with knowing a great many *other* things that are directly relevant to making wise political choices. For example, someone who knows the names of several foreign leaders, or of obscure legislative procedures, will likely know these things, not as a matter of isolated pieces of trivia, but on account of knowing a great deal about the political world.

Second, the claim that tests of this particular kind are diagnostic of political knowledge has some empirical support. These tests build on indices developed by Michael Delli Carpini and Scott Keeter (1996 and 1993), whose results suggest that scales made up of as few as five items from the American National Election Studies (ANES) surveys within aforementioned categories tend to 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 will be related to a person's degree of political knowledge, such as degree of political participation and degree of opinionation.

⁵ As is clear from the type of 'elite cue' theory (e.g., in Zaller 1992 and 2012, and Berinsky 2009) to be discussed in Section 7, ideological and party identities are certainly predictive of political preferences and attitudes. However, they are *less* predictive than factors like authoritarianism (see, e.g., Hawkins *et al.* 2018).

Third, the claim that these tests are diagnostic of political knowledge finds further support in the fact that, on Althaus’s models, individual-level information effects turn out to be consistently higher for groups that we have independent reason to believe will typically be less politically informed, such as high-school dropouts, people who are poor, and people in rural areas; and consistently lower for groups we have good reason to believe will typically be more politically informed—including in the richer sense identified in relation to Fishkin’s deliberative polls—such as the college educated, the affluent, and people in cities or suburbs (Althaus 2003: 135; see also Delli Carpini and Keeter 1996: 177).⁶

These three considerations offer good reason to believe that operationalizing what it is to be informed in terms of the relevant type of tests will enable the models in question to track the phenomenon we are after, namely *informed* preferences. Of course, talking in terms of informed preferences—or even *fully* informed preferences—in this context is not meant to suggest that such preferences are *ideal* in any strong sense. Exactly what we can expect from the epistemic character of our fully informed selves will be discussed in section 7. For now, it suffices to note that, while our fully informed selves are certainly epistemically better off than we are, they are not meant to be epistemically flawless.

3. Modelled Democracy

We can now put the relevant models to philosophical work. What I will be proposing is that the relevant type of modelling enables us to assign a probability to that the respondents would have reported and more-over—assuming that such reports tend to be sincere—had some particular political preference, had they

⁶ Someone might object that the relevant tests simply measure a form of *technical* political knowledge, 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. (Many thanks to Michael Garnett for raising this matter.) But note that, by controlling for a wide range of variables (e.g., income, race, etc.) with which these forms of knowledge likely travel, the models *preserve* such situated knowledge in estimating informed preferences. Specifically, they can be thought of as modelling 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).

been fully informed. Let us consider a brief but concrete example to illustrate how this would be done in practice.

3.1. An Illustration: Modelling an Informed EU Referendum

The British Election Study's 2017 face-to-face survey (Fieldhouse *et al.* 2018) collected demographic information on 2,194 people in the UK, asked six questions tapping into people's degree of political knowledge, and surveyed respondents on a number of political issues.⁷ One such issue was how they voted in the UK's 2016 EU referendum. Imagine for a moment that those surveyed made up the entirety of the voting population in the UK, and that the survey was administered in the place of the referendum held the year before. Using the data collected, we could then fit a model to estimate for each voter the probability that they would have reported a preference for Remain or for Leave, had they been fully informed (as operationalized by the knowledge items) but otherwise no different along all demographic variables. Then, we could calculate the means of the assigned probabilities of support for Leave and Remain, respectively, to generate a breakdown of informed support between two options. For the respondents in the BES data set, this makes for an information effect whereby support for Remain increases from 49.4% to 54.5%—a difference of 5.1 percentage points that is non-trivial in the context of the actual referendum result (48.1% for Remain). In the imagined scenario, it would then be the proportion of informed support that would determine the electoral outcome.

This is, of course, a simplistic example. Still, it should hopefully give a sense of how the relevant type of modelling could be applied in an electoral context. In more realistic settings, the range of possible electoral options could be extended from two, be it referenda options, political candidates, or parties. Moreover, the model used here assumes that everyone who is fully informed votes. This is not a completely unreasonable assumption, given the correlation between education and the likelihood of voting (e.g., Field 2015: 6). But if one wanted not to make that assumption, one could include abstention as an additional response option, and then fit a multinomial model that estimates not only informed votes but also informed abstention. As it happens, fitting such a model on the present data set leaves the information effect in the direction

⁷ See Appendix for more information on the data set and subsequent models.

of Remain intact, with a Remain proportion of 55%, and the predicted proportion of voters increasing from 87% (the proportion of reported voters in the sample) to 93%.

Another consideration is the robustness of the model. In particular, do slight changes in how the model is specified—i.e., in what aspects of the participants are controlled for—make for substantially different outputs? There is only one way to find out: to compare the outputs of several models. To that end, consider two further models: a *purely* demographic model, identical to the original binomial model above, save for excluding the quasi-attitudinal variables of party identity and union membership; and an *authoritarianism* model, which includes all variables from the binomial model, as well as respondent's score on the authoritarianism scale (see Appendix). The simulated proportions of support for Remain on the two models comes out to 55.7% on the demographic model, making for an information effect of 6.3 percentage points, and 51.9% on the authoritarianism model, for an information effect of 2.5 percentage points. The thing to note here, of course, is that the upshot of the three models is the same: a non-trivial information effect towards Remain, averaging 4.7 percentage points across the three models.

Indeed, averaging outputs across an ensemble of models is one promising strategy for navigating disagreements about what exact model ought to be implemented. After all, while model selection in the context of modelled democracy is partly an empirical matter, in so far as it is an empirical question what factors are highly predictive of political preferences, yet such should be excluded (for the reasons given in Section 2.3), this still leaves room for philosophical discussions—and disagreements—about how to specify the model. And this is where robustness checks come into the picture, not only because they enable us to gauge the extent to which it matters which model we select (sometimes it might not), but also because it enables us to proceed by way of the average output of an ensemble of models. On such an ensemble approach, the outcome would then be given by the average output—either across a binary option (e.g.,

Leave and Remain, as in the case considered a moment ago), or across multiple options (e.g., more than two political parties)—of the range of reasonable models.^{8,9}

Whether involving a single or an ensemble of model, let us refer to an electoral process that simulates the public's informed preferences in this manner as *modelled democracy*. Government set up on the basis of such preferences is likely to be less affected by widespread public ignorance than are current democratic arrangements. As already noted earlier, this is *not* to suggest that what is being modelled is what we would have been like, had we been epistemically flawless. But as Brennan (2016: 207) notes, ideal epistocracy—

⁸ Again, in Section 2.3, I offered some reason for excluding some highly predictive variables, such as those associated with an authoritarian disposition. However, while (naturally) convincing in my view, that line of reasoning does not necessarily render it *unreasonable* to include a model with an authoritarianism variable in an ensemble of models, in the manner that is being done here.

⁹ Of course, another decision that has to be made is what specific knowledge items to use, and someone might be concerned that what knowledge is politically relevant thereby needs to be known upfront. Three points are in order here. First, in light of the literature discussed in 2.3, we do have a good sense upfront of what *types* of knowledge will be relevant, including knowledge of what government is and does, and of political leaders, parties, and alignments. Second, as for *specific* knowledge items, Delli Carpini and Keeter's (1993) results comparing different knowledge scales moreover suggest that scales made up of different items within aforementioned categories tend to correlate to roughly the same (substantial) degree with things that we (as also noted in 2.3) have independent reason to believe will track political knowledge, which means that knowledge estimates are unlikely to be wildly different under different item selections. This is, moreover, consistent with the fact that, when it comes to political knowledge, 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—and it is in light of this that Delli Carpini and Keeter (1996) conclude that 'researchers developing national or general political knowledge scales need not be overly concerned with the mix of specific topics covered by individual items' (174). Third, it is a fairly long-standing practice for large attitudes surveys such as ANES and BES to include a number of knowledge items in their surveys, on account of which there are by now decades worth of survey data to inform our decisions about what particular items to include. So, to sum up, while selecting appropriate knowledge items is certainly not a trivial matter, we have a fairly robust foundation to stand on when it comes to making such decision in a reasonable and empirically informed manner. (Many thanks to an anonymous reviewer for this journal for raising this question.)

involving all-wise and all-benevolent rulers—is no more an option than is ideal democracy. So, the challenge for anyone looking to deal with public ignorance without surrendering democracy is to identify something that constitutes a strong bet for avoiding to the greatest degree feasible the problems posed for democracy by public ignorance, without giving up on the anti-epistocratic idea that we, the people, should rule ourselves. While Section 7 will say more about exactly what our informed selves will tend to be like, if not all-wise (or all-benevolent), my contention will be that modelled democracy fits that bill.

3.2. Would it Make any Difference?

Of course, going beyond the above EU referendum example, we would want to know what difference modelled democracy can be expected to make to electoral politics. Looking at information effects in past US Presidential elections, 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). But he also finds that, in the six presidential elections considered, there is no difference in the electoral outcome: whoever actually won the election would have won, even if voters were fully informed. Should we infer that modelling informed preferences is politically irrelevant? No. Consider three types of electorally relevant information effects we can expect to see.

First, information effects can change the electoral outcome, but they will tend to be masked in cases of a small number of electoral options, as in the US context. As Kasper Hansen (2009), points out, ‘the existence of more parties allows the impact of political knowledge to push the voters to several other parties; not just one, as in the American case. In other words, the aggregated effect of knowledge will, *ceteris paribus*, be greater in a multiparty system than in a two-party system’ (542). This expectation is born out in a number of studies. For example, André Blais and colleagues (2008) simulate the outcome of six past Canadian elections, involving three to four parties, with fully informed voters, and see a likely difference in outcome in one. Henrik Oscarsson (2007) simulates six past Swedish elections, involving eight main parties, and sees a likely difference in outcome in two of them.

Second, whether or not information effects have direct consequences for electoral outcomes, they can have substantial implications for party political choice. Yosef Bhatti (2010) models three European Parliament elections—in Denmark, Finland, and Sweden—involving five to seven main parties, and finds

several cases in which the differences between actual and simulated support are in the double digits. Similarly, Hansen (2009) models two Danish elections involving six main parties, and finds no change in what coalition government would have come to power, but a substantial change in the power distribution internal to the party blocks. For example, in one of the two elections, doubling the degree to which the voters were informed would have almost doubled the level of support for the Conservatives.

Third, information effects will likely have longer-term implications for what electoral menu is presented to the electorate to begin with. This gets to an important limitation in existing simulation studies: they don't consider the aggregate effects of informed voting over time.¹⁰ They model how informed people would have voted, given the actual electoral menu on offer at the time, not how they would have voted, *had that menu potentially been different on account of informed voting at past elections*. In particular, it would be very surprising if non-trivial and, in some cases, double-digit swings in support for parties within party blocks or coalitions would not substantially change subsequent party-political strategy and positioning. Remember, on modelled democracy, political parties and candidates would be catering, not to the *actual* political preferences of the people, but to their *modelled* preferences, which would exhibit the kind of deviations from actual preferences that we see in simulation studies. Such deviations in one election would likely have significant implications for the electoral menu available to the voters at future elections.¹¹

This point about the longer-term implications of modelled democracy also helps us see why information effects are not irrelevant even in cases where they might to some degree or other be masked by a

¹⁰ The only exception I am aware of is Toka (2008), who suggests that 'an informed electorate has a beneficial impact on the level of corruption control over multiple electoral cycles' (38).

¹¹ In light of this, would modelled democracy facilitate or impede political movements or party realignments over time? That would depend primarily on the degree to which such movements and alignments are driven by political ignorance. Movements based in high levels of political knowledge will tend to be catered to; those based in ignorance will receive less political attention. Additionally, realignments will be driven partly by the ability of political actors to predict the outcomes of models employed in future elections. To the extent that those actors are able to make accurate predictions, realignments driven by high degrees of political knowledge will be facilitated; to the extent those predictions aren't accurate, realignment will to some extent be an artefact of majority-seeking parties and coalitions playing catch-up after elections. (Many thanks to an anonymous reviewer for this journal for raising this question.)

two-party arrangement (or two stable party blocks). For example, Althaus (2003: 123) finds an average, individual-level information effect of 15.7 percentage points on fiscal survey questions; 12.3 on gay rights; 14.1 on welfare; 13.8 on health care; 12.1 on equality; 13.9 on environmental matters; and 14.6 on abortion. For all of these topics, what position one takes on the spectrum has implications for one's politics; for some of them (abortion being the most obvious example), what position one takes comes close to *defining* one's politics, at least in a US context. Consequently, even if shifts of aforementioned magnitude only register as within-party shifts in a two-party system like the US, this will have non-trivial, longer-term implications for within-party positioning and, by implication, also for candidate choice.

3.3. Modelling Local Elections

The discussion so far has focused on national and (in the case of EU Parliamentary elections) supranational elections, in line with the broader literature on information effects in electoral contexts. That focus can be explained in part by the fact that such elections tend to garner more interest than local elections. But part of the explanation might also be an assumption that levels of public ignorance are lower at a local level, where people presumably have a more direct and personal experience with the relevant political issues. However, recent research has thrown doubt on that assumption. For example, both Shaker's (2012) study in Philadelphia and Rapeli's (2014) study in Finland suggest that levels of political knowledge at the local level and national level are more or less identical, while Binder and colleagues' (2016) study in Florida finds evidence of people being *less* informed about local than national politics.

This raises the questions whether we should implement modelled democracy at the local level as well. Of course, any statistical model needs a sufficiently large sample to work with in estimating the relevant parameters. Would the samples be large enough at a local level? As noted, the type of model usually employed in measuring information effects is a logistic one: a binomial logistic model for binary response variables (e.g., Leave and Remain; Republican and Democrat), and a multinomial logistic model for categorical response variables with several categories or levels (e.g., representing more than two political parties or options). Sample size requirements for logistic regression is a complex area, but there are also some fairly common rules of thumb. For example, Long (1997) suggests that samples of less than 100 observations should be avoided, and that samples of more than 500 observations should be adequate for most contexts.

These values, however, should be raised in cases where we are looking to estimate a large number of parameters. To that end, Hosmer and Lemeshow (2000) suggest a rule of thumb of 10 observations per independent variable to be included in the model. For multi-level, nominal variables, this translates to needing at least 10 observations for the least common level of the relevant variable.

For context, return to our EU referendum model in Section 3.1 above. That model estimates an intercept and fourteen independent variables. Several of these variables (e.g., party identity, education, and work status) have multiple levels. In total, the model therefore ends up estimating 67 parameters. Going by the minimum set by Hosmer and Lemeshow, we would therefore want at least 670 observations. (As noted in the Appendix, the model was fitted on a sample of 1,467 complete observations.) That is, of course, a minimum, and we would ideally want much more than that. Still, it should provide some helpful context in thinking about the possibility of modelled democracy at the local level. By way of example, the average size of a US county is about 100,000. Even if the voting population is some fraction of that, it seems likely that aforementioned rules of thumb would be satisfied fairly comfortably. In the case of elections at an even more local level, involving towns or (in the UK) local councils, it would depend on the size, of course. For any reasonably complex model, such as the EU referendum model above, we would likely not want to attempt to model any election with a voting population of less than 1,000 people.

3.4. The Plan for What Follows

Having made clear how modelled democracy would operate, at what electoral levels it could be implemented, and the type of impact it would likely have on electoral politics, the remainder of the paper will consider whether such democracy satisfies the QAC. Remember, we are looking for a way to solve the problems posed for democracy by the fact of widespread public ignorance, without restricting the franchise, and while satisfying the QAC. As we have seen, modelled democracy seems well placed to address the problem of public ignorance, and it doesn't do this by restricting the franchise—but can it ever be politically legitimate? If there are reasonable objections to it, then the answer is 'no'. So, let us consider what would have to be considered the most plausible candidates for reasonable objections.

4. Modelled Democracy and the Franchise

The first challenge to modelled democracy is the suggestion that it would not constitute a rethinking of the franchise so much as an *abolishment* thereof, and consequently doesn't deserve the label 'democracy' after all. There are two ways to push this worry.

The first way is by maintaining that votes do not express preferences (informed or otherwise) to begin with. While some do indeed take votes to express preferences (e.g., Caplan 2007 and Arrow 1977), others take votes to instead express *beliefs*, including beliefs about what is just (Rawls 1971) or in the common good (e.g., Estlund 1990; Rousseau 2008/1762). This, however, is not a particularly serious worry. Thinking back to the characterization of the relevant modelling above, there is nothing that prevents us from asking people about what options out of those featured in the relevant survey they consider to correspond to what is just or in the common good, as opposed to about their preferences. We could then either have the model generate a distribution of informed beliefs, or we could transform the output back to preferences on account of how beliefs about what is just or in the common good would arguably have to be accompanied by preferences to the effect that what is just or in the common good should be implemented. I will in what follows stick to the preference formulation, while asking the reader to keep in mind the flexibility offered by the point just made.

The second way to push the worry that modelled democracy is not a form of democracy is by suggesting that it does away with voting, and that voting is a necessary component of democracy. We find this idea in Estlund (1990), who takes democracy at 'its core to be rule by the people through voting' (397). Thomas Christiano (2013) makes a similar claim when suggesting that democracy 'is a process of collective decision-making in which all or nearly all the members of the permanent adult population of the society have effective rights to participate as equals in a crucial stage of the decision-making. In the case of representative democracy this implies that all or nearly all the adult population participates in free and fair elections' (1250). If voting is *conceptually* necessary for democracy, and the input provided by people in modelled democracy doesn't amount to voting, then modelled democracy is not a form of democracy after all.

In the next section, we shall consider the (quite ambitious) nature of the type of electoral engagement that modelled democracy would involve. But let us assume, if only for the sake of argument, that such engagement would not qualify as voting. Would this be a problem? Arguably, voting is a *means* and not

something we (should) value in and of itself. More specifically, rule *by* the people (that is, popular rule), be it through voting or any other means, is good because it makes it more likely that there will be rule *for* the people, and in accordance with their preferences in particular.¹² As we will see in Section 6, for purposes of political choice, our real preferences consist in our informed preferences, and modelled democracy does a better job than does traditional forms of democracy in making government sensitive to those preferences. For that reason, the appropriate response to any objection to the effect that modelled democracy doesn't qualify as a form of democracy on account of not involving a type of activity qualifying as voting: so much the worse for 'democracy,' so understood.

5. Non-Electoral Participation and Electoral Engagement

In the previous section, we considered the suggestion that modelled democracy does away with voting. We found reason to reject that suggestion: either modelled democracy does not do away with voting; or, if it does, then the notion of voting that is invoked is restrictive to a questionable degree. We will now consider the concern that modelled democracy not so much does away with as it *ignores* other features of democracy. In particular, democracies involve varieties of political participation that have nothing to do with elections, such as protests or other forms of non-electoral, political activism. The defence of modelled democracy offered here doesn't say anything about such forms of participation. Is that a problem?

It is not clear why it would be. For purposes of this paper, modelled democracy is a position about electoral reform, and specifically about rethinking the process by way of which representatives are selected and other electoral outcomes settled. In that respect, it concerns itself with a particular democratic *activity*—namely voting—not with democratic *society* in the wider sense. This is not to deny that we engage in any number of other activities outside of the polling booth that might rightly be considered democratic activities. It is also not to deny that someone else might want to take the project of modelled democracy beyond the

¹² We find this idea in Mill's (2008/1861) argument for representative government, in terms of how 'the rights and interests of every or any person are only secure from being disregarded, when the person interested is himself able, and habitually disposed, to stand up for them' (245).

bounds of electoral politics, on account of public ignorance potentially presenting problems for non-electoral aspects of democratic politics as well. It is simply to maintain that electoral politics forms the bounds of the present investigation. How so?

An analogy might be helpful here. Some people want the UK to move from a first-past-the-post (FPTP) voting system to a proportional representation (PR) system. On the latter, if $n\%$ support party P , then P will get roughly $n\%$ of the seats in Parliament—something that is not generally the case under FPTP. Party-political motivations aside, what is driving defenders of PR is likely the idea that popular representation within powerful structures is a good thing in general. The desire for such representation might of course be felt, not just in relation to Parliament, but also in connection with (say) the judiciary, civil service, and so forth. Indeed, there is no principled reason why a supporter of PR wouldn't be interested in reforming these other domains as well. Still, since political reform is difficult and fraught with risks of unintended consequences, we don't require that defenders of PR take a firm view on the merits of all reforms pushing for proportional representation. And the same goes for modelled democracy: whether someone else might want to pursue applying the framework of modelled democracy outside of electoral politics, doing so forms no part of the particular type of electoral reform defended as part of the present investigation.

At the same time, even if the type of modelled democracy defended here is not *about* non-electoral politics, it will have to account for the fact that electoral and non-electoral, political participation do not represent isolated systems; on the contrary, they often interact. Consider an example. Say there is an active protest movement in society. In response, an electoral event is arranged—an election or a referendum, say—speaking to the claims of that particular movement. Whether that event ends up working out in favour of the movement or not (perhaps there is a majority against their specific demands), there should be a fairly strong case that the movement has been heard, and an expectation that it will acknowledge this. Non-electoral activism has been channelled through, and to that extent been resolved by, an electoral process.

Now consider how this particular case would play out on modelled democracy. The electoral event in question would involve modelling the informed preferences of the relevant electorate, including the members of the protest movement. Assume that there turns out to be an informed majority against the demands of the movement. Will its members still feel that they have been heard, in the manner that we just suggested

we can expect under current electoral arrangements? Someone might object that they will not. This generates a worry that people will feel more removed from the electoral outcome on modelled than on traditional democracy, which would have implications for people's willingness to accept the outcomes.

This is a legitimate worry. Political representatives make laws, and there is ample evidence that people follow the law to the extent that they feel that their voices have been heard in the making of those laws (e.g., Tyler 2006; see also Tyler 2011). At the same time, the objection is exploiting the fact that more needs to be said about what modelled democracy would look like from the point of view of the individual voter—and, in particular, that such democracy would be both ambitious and a potentially quite engaging exercise. Specifically, it would take the form of the type of in-depth, face-to-face polling conducted by, among others, the American National Election Studies and the British Election Study. So, rather than casting a ballot in the traditional way, people would be interviewed for purposes of teasing out a number of political preferences, in addition to answering questions about political events, persons, and procedures. Such interviews have the potential of engaging individual voters at a far deeper level than is happening currently, in the case of traditional, ballot-box voting. After all, save for in the context of exit polls, we are rarely asked about the preferences underlying our votes, and even in those contexts are we rarely asked in any great detail about the wide range of policy, political, and moral positions that might form the foundation for those choices. In that respect, modelled democracy will involve a level and scale of public consultation that is unheard of—and voters will take note.

For these reasons, it is not clear why modelled democracy could not successfully channel non-electoral participation, and more generally produce outcomes that people will accept. As voters, you and I will be offered regular opportunities to engage directly with representatives of those institutions in a conversation about what we believe and why we believe it. In addition, we will receive direct feedback on the extent of our knowledge about the political world, and the role that this knowledge (or absence thereof) might or might not play in political choice in the aggregate—which in turn will form part of a conversation about the very motivation for modelled democracy. As such, electoral institutions will need to engage with and—literally—listen to voters on a scale and to a degree of depth that has never been done before, and it is hard to imagine that this will be anything but engaging for individual voters. And herein lies the response to the current worry. Aforementioned research on the role of a feeling of being heard in people subsequently

being willing to comply and cooperate suggests that, while people do not require that their input be reflected in the ultimate decision, they must feel that they have been listened to (e.g., Tyler 2006: 127). Given the type of extensive engagement with individual voters that modelled democracy would involve, it is hard to imagine that they will not.

6. Informed Preferences and Popular Rule

Someone might grant that voting is not conceptually necessary for democracy, and that modelled democracy will be both ambitious and engaging, while denying that rule in accordance with the preferences we would have had, had we been informed, is rule in accordance with *our* preferences. And if so, modelled democracy is not a form of popular rule—not a form of democracy—after all.

The most promising way to flesh out this worry is by assuming a particular notion of preference, on which preferences are closely related to choices. The relevant notion is popular in some parts of economics, where choice is sometimes understood as *revealed* preference (e.g., Samuelson 1938). This notion makes sense of the objection because it enables us to say that, unless government is set up in a manner that is a fairly straightforward function of the choices we in fact make—say, by voting—then it is not set up on the basis of *our* preferences, as revealed by us through those choices.

We can start to see why this notion is problematic, however, by reflecting on choices based on false belief. Say that someone is presented with two economic policies, identical in all respects save for the fact that one of them involves a significant reduction in immigration. If the person in question is mistaken about the actual impact of immigration on the economy—which tends to be positive—and opts for the immigration reducing policy on the grounds that this (they believe) will be beneficial for the economy, it is not clear that they *preferred* that policy. They certainly *chose* it, and we can explain why they did (they believed it to promote something that they care about, namely economic welfare). But in order to see how little that choice might have to do with what they actually preferred, imagine that the person came to realise their mistake about the economic effects of immigration. This would affect their choice behaviour. If they were offered the two policies again, they would now choose differently. More to the point, it would seem appropriate to say that having their mistakes about the relative merits of the policies corrected thereby would bring their choice behaviour in line with what they really preferred all along. Indeed, were we now to ask

them about the choice they made the first time around ‘Was that *really* what you wanted?’ they could reasonably be expected to answer in the negative. The fact that their mistaken beliefs led them to *choose* the economically detrimental policy should not lead us to believe that they ever *preferred* it.¹³

To avoid any confusion, the claim being made here is not that *welfare* is a function of (the satisfaction of) informed preferences (e.g., Brandt 1979; Railton 2003), or that what you have *reason* to do is determined by what you would do if you were fully informed, or by what a fully informed version of yourself would advise you to do (e.g., Smith 1995). Rather, the claim is that your *real* preferences are to be identified with your *informed* preferences.¹⁴ While not the standard view of preferences, it is not without its defenders, and its plausibility is particularly clear in cases where false belief is rife.¹⁵ As we have already seen, this is the case in the political realm, and in the type of case introduced a moment ago.

However, someone might object that, in the case imagined, it is not that the agent’s original choice failed to be in accordance with their real preferences, but rather that their (real) preferences changed. We may call into question this alternative reading through what I will call the *argument from regret*. Return to the agent who opts for an anti-immigration policy on economic grounds. Let’s assume that the agent comes to realize their mistake. It would be appropriate for them now to feel regret. But such regret would not make

¹³ Goodin (1995) makes a similar point when suggesting that, under circumstances of ignorance, ‘we can serve a person’s “real” preferences only by censoring the misleading indication of his preferences that is revealed in his choices’ (137).

¹⁴ Why not then model *all* of our preferences—including about whom to marry, or what career to pursue—to ensure that we live in accordance with our true preferences, even outside of the political domain? As we shall see when we consider the charge of paternalism in a moment, because unlike voting, which is a paradigmatically other-regarding type of act where your actions have implications for others and as such are held to more demanding moral standards even within a broadly liberal framework, choices about how to live your own life fall on the self-regarding end of the spectrum, and deserve greater protection from interference.

¹⁵ Harsanyi (1997) seems to defend such a view, and Goodin (1995) either defends or comes close to defending it. See also Delli Carpini and Keeter (1996), who suggest that ‘the real interest of an individual—and by extension of a group and of the polity as a whole—are reflected in the choices one would make if he or she were fully informed about the consequences’ (5).

sense had their original expression been in accordance with whatever is appropriately termed their *real* preferences on the matter. After all, on the alternative reading under consideration, the relevant regret simply comes down to this: the agent regrets that their less informed self supported a policy they now consider a bad one. But that can't be it. If it were, there would be no difference between how the agent feels about their own past support, and how they feel about some *other* person who might have supported that same (in the agent's view) bad policy. In both cases, they find it *lamentable* that a bad policy has received support. But in the case of their own past self, there is more going on—and that is what the alternative reading under consideration cannot account for. They can't *regret* the choices of others, of course, but that is not the point. The point is that the agent will feel *more* troubled by their *own* past support than about that of some other person. If the perceived badness were a mere function of the badness of supporting a bad policy, this feeling wouldn't make sense. But in the case of the agent's own past self, they didn't just lend their support to a bad policy, which might have bad consequences for others; on account of their false belief about the economic effects of immigration, they also let *themselves* down in a very specific sense: they acted contrary to what they actually preferred. That is why it is appropriate for the agent to feel regret, and for them moreover to answer the question 'Was that *really* what you wanted?' in the negative.

It might be objected that, even if we accept that your real preferences are the preferences you would have had, had you been informed, there is something *paternalistic* about in effect telling people that their (real) preferences are different from what they think (and say) they are. There are two separate matters raised by this objection that should be addressed. One concerns the evidential bar for political interventions more generally. The other concerns the relatively higher bar that applies in the context of paternalistic as opposed to non-paternalistic interventions. As we will find reason to return to the former in Section 8, I will presently focus on the second matter. Here, it should be noted that the force of the anti-paternalist line is strongest in cases of self-regarding conduct. Specifically, what seems particularly problematic is someone telling us what we should do, when the actions concerned affect only us (or, more realistically: when the likely effect on others is negligible). That is what's driving intuitions in classic cases from the paternalist

literature, involving seat belt or helmet use or recreational drug use.¹⁶ And to the extent that interventions concern self-regarding conduct, the evidential bar is thereby also higher than in cases involving other-regarding conduct. But note that voting is *not* a type of self-regarding conduct. Indeed, the very motivation for worrying about the political implications of public ignorance is that how people vote has a profound impact on all of us, in shaping the very structures that regulate how we are to live together. This calls into question not only the claim that identifying people's real preferences with the preferences they would have had, had they been informed, is at all *paternalistic*, in the sense of constituting an interference motivated by a desire to make the actor herself better off, but also—and more importantly for present purposes—that the relevant evidential asymmetry applies in the case of modelled democracy.

So, the argument from regret remains: the preferences that matter for political choice are those you would have had, had you been informed. If that argument is on the right track, then the relationship between preferences and choice behaviour will be less straightforward than what's assumed by the objection we started out with. What a person really prefers is what they would have preferred, had they been informed. In particular, if becoming informed but in all other respects remaining just the way you are changes your *reported* political preferences in some systematic way, that is good evidence that what would have been reported had the person been more informed is a better approximation of their real preferences, on account of having survived being confronted with relevant information. Moreover, the existence of information effects on political preferences provides good evidence of the antecedent: *had* we become politically informed but in all other respects remained just the way we are, our reported preferences *would* likely change in systematic ways—and, we might now say, thereby also have been brought (more) in line with our *real* preferences. But if that is right, then rule in accordance with the preferences we would have had, had we been informed, *is* rule in accordance with our preferences, contrary to the objection under consideration.¹⁷

¹⁶ This is not to deny that there are 'mixed' forms of paternalism (e.g., Feinberg 1986 and Grill 2007), where the motivation for paternalistic intervention is not *solely* the good of the persons interfered with, but merely to suggest that the anti-paternalist intuitions are strongest in non-mixed cases.

¹⁷ This line of reasoning also helps respond to a separate worry about modelled democracy: if I know that my electoral choices will be filtered through a model that estimates my informed preferences, what incentives do I have to inform myself? It should be noted in response that, if there is anything to the literature on (individual) *rational ignorance* (e.g.,

7. Partisanship, Polarization, and the Division of Ideological Labour

Another candidate objection denies that modelled democracy models *informed* preferences. In particular, on one of the most prominent theories of public opinion, such opinion is driven primarily by the cues offered by political elites. By far the most famous defence of this view is that of John Zaller (1992).¹⁸ On the objection under consideration, the challenge this type of view presents for modelled democracy is that it suggests that the ‘most informed,’ as picked out by models like those considered above, are simply the most *partisan*, the ones with the most stable ideological positions. If that is right, what modelled democracy is revealing is not so much your more informed as your more partisan self.

However, this objection reads too much into Zaller. What Zaller finds is that the more politically informed¹⁹ are the most ideologically coherent ones, in being the most adept at discriminating between information provided by elite communications, and rejecting what doesn’t fit with their ideological stance. But this doesn’t entail that there is anything *epistemically improper* going on. It *could* be that elite communications are highly predictive of public opinion because the elite is indoctrinating the rest of us—but that can’t be established by looking at mere correlations. In fact, Zaller explicitly rejects the idea that what we are dealing with is *elite domination*, involving ‘a situation in which *elites induce citizens to hold opinions that they would not hold if aware of the best available information and analysis*’ (313). Rather, he claims, ‘citizens pick and choose, *on the basis of their predispositions*, from the menu of elite-sponsored alternatives to which they have been

Downs 1957), it is not clear that we have much by way of incentive to inform ourselves under *current* electoral arrangements. (This, of course, is compatible with it being *collectively* rational for us to counteract public ignorance.) By contrast, a central conversation to have with the electorate in the context of modelled democracy is the upshot of the argument in this section: your *real* preferences are your *informed* preferences. This highlights a novel incentive for informing yourself: in so far as you would like to think and act with reference to your true preferences, you have good reason to inform yourself. (Many thanks to an anonymous reviewer for this journal for raising this question.)

¹⁸ Similar views have been defended in more recent work by Berinsky (2009) and Lenz (2012).

¹⁹ Zaller talks in terms of political *awareness*, but measures such awareness in the same way that Althaus (2003) does the degree to which people are politically informed, i.e., by way of ‘simple tests of neutral factual information about politics’ (Zaller 1992: 21).

exposed' (287), where predispositions correspond to the interests, experiences, and values—including our most fundamental values—that regulate our acceptance or rejection of political communications (22-3).

What Zaller ends up offering is an account of a particular division of ideological labour. Ideology is 'a mechanism by which ordinary citizens make contact with specialists who are knowledgeable on controversial issues and who share the citizens' predispositions' (327). So, to the extent that we take Zaller to show that modelled democracy models our most partisan self, 'partisan' need to be stripped of its pejorative force. The informed members of the public are simply the particularly discerning people who are most successful in exploiting the division of ideological labour in forming a coherent view about appropriate political options, in light of their fundamental values.

This conclusion moreover stands in the face of more recent defences of the type of 'elite cue' theory Zaller offers. Berinsky (2009) argues that people can form opinions in the absence of clear political messages from elite sources, and also that ethnic attachments that cut across partisan affiliations play a role. This, too, counts against an elite domination narrative, since it suggests that people resist elite cues when the messages involved go against ethnic attachments. Zaller's (2012) adjustments of his theory follows similar lines, by suggesting that his original theory ignored or obscured 'the part of public opinion that is not anchored in ideology and is little influenced by elites' (571). So, the picture from earlier remains: what we have is not a structure of indoctrination, but a division of ideological labour, where informed people tend to rely on elite communications from people sharing their fundamental values in forming coherent beliefs about politics, except when other attachments (e.g., to ethnic groups) trump considerations about ideology.

Of course, it follows from this that, as the degree to which people are informed increases, so does *polarization* along partisan lines. Indeed, while everyone can disagree deeply with others on political matters, only the informed (supposedly, together with the elites) can disagree in the clear-cut and systematic manner that is along distinctly partisan lines (see, e.g., Lauderdale 2013). But if so, what's so good about being more informed? Some degree of polarization is inevitable in politics, on account of how the positions we take in all likelihood ultimately turn on fundamental questions about value—including what is, arguably, the most basic question in politics: *Who deserves what?*—about which people can reasonably disagree. Moreover, even if there were a way to model away such disagreements, it is not clear that this would be desirable. Remember, the aim of modelled democracy is to say something about what we would prefer in politics, were we more

politically informed. Were the model to discard the fundamental commitments regarding who deserves what for a significant proportion of the people, for purposes of modelling a non-polarized society, it is not clear that what's subsequently said about that proportion can be said to be about *them*. That is why it is not a desideratum of modelled democracy to model away polarization, or disagreements generally. Rather, by modelling our informed selves, polarization and all, we are able to model what it is for us to disagree with others while holding positions that—on Zaller's picture—at least constitute the package of commitments that best reflects our fundamental values.

8. Modelled Democracy and Unknown Bias

A final objection is the following: the way we get to 'our' informed preferences on modelled democracy is by extrapolating from reported preferences, while controlling for what are in effect potentially biasing factors like wealth, race, sex, and so forth. But, the objection goes, we might fail to control for *unknown* bias. Focusing on the type of modelling that takes education as the proxy for being informed (e.g., in Caplan 2007), Estlund puts the objection thus:

It would not be unreasonable to worry that there might well yet be other factors that travel with higher education that skew the sample enough to outweigh the beneficial effects of the extra education. I don't say whether or not there is strong reason to have those doubts, but only that it would not be unreasonable to have them given the history of these matters. So whether or not I have the doubts, some reasonable citizens could (Estlund 2006).

Estlund's objection raises both a specific and a general worry. The specific worry that models relying on education proxies might have failed to control for some biasing factors isn't directly relevant to the type of modelling we have been concerned with, which instead relies on survey items testing directly for political knowledge. But the general worry that for *any* proxy used there is a risk of such bias is well taken. It is also one that has a response.

To see why, consider that we have *good reason* to believe that potentially biasing factors have been controlled for by Althaus and others conducting similar types of modelling. How are we to spell out

Estlund's worry in light of this? In either of two ways. On the one hand, it could refer to some *specific* variable or type of variable as being left out. Say for example that factors like wealth or gender were *not* included in the demographic profile. Then the informed preferences that the model would ascribe to us would arguably not be the ones we would have had, had we been informed, for the simple reason that we are relevantly different from the people whose preferences are used to generate 'our' informed preferences. But that is not the case the type of models we have been looking at. And if some further variable was quoted as part of a similar type of worry, the appropriate response would be to construct better models.

Hence, the second way to spell out Estlund's worry proceeds without mentioning any *specific* variable, or even any particular *type* of variable. After all, as it stands, the worry is simply that there *might* be some (unspecified) unknown source of bias. But consider what that possibility would have to entail in order for modelled democracy to be in trouble. In order for our reasons for implementing modelled democracy *not* to be undercut, even the mere possibility of bias would need to be ruled out. This is to require that we be *certain* that nothing has been left out. For that requirement not to be *ad hoc*, it would have to be grounded in some more general principle about political reform, to the effect that, even if we have good reason to believe that some political system would work in the intended way, it is still reasonable to object to it. But in that case, reasonable people are going to be able to object to a lot more than to modelled democracy. Consider any given electoral system, *E*. If of the kind implemented in most developed countries, we have good reason to believe *E* to reveal the popular will (setting aside philosophical concerns about failing to track our *informed* preferences specifically). But can we be *certain* of it? No. Electoral systems might fail to track the popular will for any number of reasons, including technical problems, low voter turn-out, partisan re-districting, or different forms of voter suppression (whether intentional or not). So, there *might* be some (unspecified) unknown factor skewing the output of *E* away from the popular will. In fact, if we can be certain of *anything*, it is that there will be at least some non-trivial discrepancy between *E*'s output and the popular will, for aforementioned reasons or otherwise.

So, if we take the level of evidential support necessary for the legitimacy of a political system to be that of certainty, we can run the following argument: if there are reasonable objections to modelled democracy, the same goes for any given electoral system. Denying that there are *any* legitimate electoral systems

seems a high price to pay for being able to reject the idea of modelled democracy in principle, as in: irrespective of any particular questions one might have about specific models failing to control for specific factors. We should certainly be open to the idea that the models involved in modelled democracy have failed to control for some biasing factors. But insisting that we must be *certain* that this is not the case for modelled democracy to stand a chance of being legitimate sets the bar too high, as can be seen by the fact that it disqualifies *all* electoral systems—a Pyrrhic victory if there ever were one for the critic of modelled democracy.²⁰

9. Conclusion

The fact that most of us are ignorant on politically relevant matters is a problem for democracy. Some have suggested that the franchise therefore needs to be restricted along competency lines—a suggestion that in turn runs the risk of violating a long-standing condition on political legitimacy to the effect that legitimate political arrangements cannot be open to reasonable objections. In light of this, the present paper outlined a way to solve the problem of public ignorance without restricting the franchise. The proposal put forward involved filtering the electoral input of a universal franchise through a statistical model that simulates what the public's preferences would have been, had they been informed on politically relevant matters. The result is *modelled democracy*. A case was made that such democracy not only meets the challenge posed by widespread public ignorance without restricting the franchise, but also satisfies aforementioned condition on legitimacy. As noted at the outset, more needs to be said in order to establish that we ought to implement modelled democracy anytime soon. Still, the present investigation can be thought of as a pilot study that at the very least makes clear that the proposal warrants further investigation.

²⁰ My argument here bears some affinity to that offered by Mulligan (2015), who makes a similar observation about how the evidential standards sometimes invoked for 'reasonable' in discussions about public reason make reasonable objections too easy to generate.

References

- Ahlstrom-Vij, K. (forthcoming). 'The Epistemic Benefits of Democracy: A Critical Assessment,' in M. Fricker, P. Graham, D. Henderson, N. Pedersen, and J. Wyatt (Eds.), *The Routledge Handbook of Social Epistemology*, New York, NY: Routledge.
- Althaus, S. (2003). *Collective Preferences in Democratic Politics*, Cambridge: Cambridge University Press.
- Arrow, K. (1977). 'Current Developments in the Theory of Social Choice,' *Social Research* 44: 607-22.
- Bartels, L. (1996). 'Uninformed Voters: Information Effects in Presidential Elections,' *American Journal of Political Science* 40(1): 194–230.
- Berinsky, A. J. (2009). *In Time of War: Understanding Public Opinion from World War II to Iraq*, Chicago: Chicago University Press.
- Bhatti, Y. (2010). 'What Would Happen if We Were Better Informed? Simulating Increased Knowledge in European Parliament (EP) Elections?' *Representation* 46(4): 391-410.
- Binder, M., M. Childers, A. Hopkins, and C. Hampsey. (2016). 'In Voters' Minds, Are All Politics Really Local? Comparing Voters Knowledge of National and Local Politics,' State Policies and Policy Conference, Dallas, May 19–21. Accessed October 18, 2019. <https://www.utdallas.edu/sppc/pdf/Binder%20Childers%20Hopkins%20and%20Hampsey.pdf>.
- Bishop, G., Oldendick, R., and Tuchfarber, A. (1983). 'Effects of Filter Questions in Public Opinion Surveys,' *Public Opinion Quarterly* 47: 528-546.
- Blais, A., Gidengil, E., Fournier, P., and Nevitte, N. (2008). 'Information, Visibility and Elections: Why Electoral Outcomes Differ When Voters are Better Informed,' *European Journal of Political Research* 48: 256-280.
- Brandt, R. (1979). *A Theory of the Good and the Right*, Cambridge, MA: Harvard University Press.
- Brennan, J. (2011). 'The Right to a Competent Electorate,' *the Philosophical Quarterly* 61(245): 700-724.
- Brennan, J. (2016). *Against Democracy*, Princeton, NJ: Princeton University Press.
- Buchanan, A. (2004). 'Political Liberalism and Social Epistemology,' *Philosophy and Public Affairs* 31(2): 95-130.
- Caplan, B. (2007). *The Myth of the Rational Voter*, Princeton, NJ: Princeton University Press.
- Chalmers, P. (2012). 'A Multidimensional Item Response Theory Package for the R Environment,' *Journal of Statistical Software* 48(6): 1-29.
- Christiano, T. (2013). 'Democracy,' in *The International Encyclopedia of Ethics*, edited by H. LaFollette, Oxford: Wiley-Blackwell: 1250-1263.
- Condorcet, M. J. A. N. C. (1976). 'Essay on the Application of Mathematics to the Theory of Decision-making,' in K. M. Baker (Ed.), *Condorcet: Selected Writings*, Indiana: Bobbs-Merrill. Originally published in 1785.

- Converse, P. (1990). 'Popular Representation and the Distribution of Information,' in J. Ferejohn and J. Kuklinski (Eds.), *Information and Democratic Process*, Urbana, IL: University of Illinois Press: 369-88.
- 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.
- Downs, A. (1957). *An Economic Theory of Democracy*, New York, NY: Harper.
- Estlund, D. (1990). 'Democracy Without Preference,' *the Philosophical Review* 99(3): 397-423.
- Estlund, D. (2006). 'The Expert/Boss Fallacy,' in *Majority Fools? Rationality and the Limits of Democracy*, available at cato-unbound.org/issues/november-2006/majority-fools-irrationality-limits-democracy.
- Estlund, D. (2008). *Democratic Authority: A Philosophical Framework*, Princeton, NJ: Princeton University Press.
- Feinberg, J. (1986). *The Moral Limits of Criminal Law, Volume Three: Harm to Self*, New York and Oxford: Oxford University Press.
- Field, T. (2015). 'Who Votes? Congressional Elections and the American Electorate: 1978–2014' United States Census Bureau; available at census.gov/content/dam/Census/library/publications/2015/demo/p20-577.pdf.
- Fieldhouse, E., J. Green, G. Evans, H. Schmitt, C. van der Eijk, J. Mellon, and C. Prosser (2018). *British Election Study, 2017: Face-to-Face Post-Election Survey* [data collection]. <http://doi.org/10.5255/UKDA-SN-8418-1>.
- Fishkin, J. (1999). 'Toward a Deliberative Democracy: Experimenting with an Ideal,' in *Citizen Competence and Democratic Institutions*, edited by S. Elkin and K. E. Soltan, University Park, PA: Penn State University Press: 279–290.
- Friedman, J. (1998). 'Introduction: Public Ignorance and Democratic Theory,' *Critical Review* 12(4): 397-411.
- Goodin, R. E. (1995). *Utilitarianism as Public Policy*. Cambridge: Cambridge University Press.
- Grill, K. (2007). 'The Normative Core of Paternalism,' *Res Publica* 13: 441-58.
- Gutmann, A, and D. Thompson (1996). *Democracy and Disagreement*, Belknap Press.
- Hansen, K. P. (2009). 'Changing Patterns in the Impact of Information on Party Choice in a Multiparty System,' *International Journal of Public Opinion Research* 21(4): 525-546.
- Harsanyi, J. C. (1997). 'Utilities, Preferences, and Substantive Goods,' *Social Choice and Welfare* 14: 129-145.
- Hawkins, S., Yudkin, D., Juan-Torres, M., and Dixon, T. (2018). *Hidden Tribes: A Study of American's Polarized Landscape*, New York, NY: More in Common.
- Hong, L., and S. Page (2004) 'Groups of Diverse Problem Solvers Can Outperform Groups of High-Ability Problem Solvers,' *Proceedings of the National Academy of Sciences of the United States of America* 101(46): 16385-9.
- Hosmer, D. W., and Lemeshow, S. (2000). *Applied Logistic Regression*, 2nd Ed. New York, NY: John Wiley & Sons, Inc.

- Landemore, H. (2013). *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*, Princeton and Oxford: Princeton University Press.
- Lauderdale, B. (2013). 'Does Inattention to Political Debate Explain the Polarization Gap Between the U.S. Congress and Public?' *Public Opinion Quarterly* 77: 2-23.
- Lenz, G. (2012). *Follow the Leader? How Voters Respond to Politicians, Policies, and Performance*, Chicago: Chicago University Press.
- Long, J. S. (1997). *Regression Models for Categorical and Limited Dependent Variables*, Thousand Oaks, CA: SAGE Publications, Inc.
- López-Guerra, C. (2014). *Democracy and Disenfranchisement*, New York: Oxford University Press.
- Mill, J. S. (2008). 'Considerations on Representative Government,' in J. Gray (Ed.), *On Liberty and Other Essays*, Oxford: Oxford University Press: 205-470. Originally published in 1861.
- Mulligan, T. (2015). 'On the Compatibility of Epistocracy and Public Reason,' *Social Theory and Practice* 41(3): 458-476.
- Oscarsson, H. (2007). 'A Matter of Fact? Knowledge Effects on the Vote in Swedish General Elections, 1985–2002,' *Scandinavian Political Studies* 30(3): 301-322.
- R Core Team (2017). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Railton, P. (2003). 'Facts and Values,' in *Facts, Values and Norms*, Cambridge: Cambridge University Press.
- Rapeli, L. (2014). 'Comparing Local, National and EU Knowledge: The Ignorant Public Reassessed,' *Scandinavian Political Studies* 37(4): 428-446.
- Rawls, J. (1971). *A Theory of Justice*, Cambridge, MA: Harvard University Press.
- Rizopoulos, D. (2006). 'ltm: An R package for Latent Variable Modelling and Item Response Theory Analyses,' *Journal of Statistical Software* 17(5): 1-25.
- Rousseau, J.-J. (2008). *The Social Contract*, translated by C. Betts, Oxford: Oxford University Press. Originally published in 1762.
- Samuelson, P. A. (1938). 'A Note on the Pure Theory of Consumers' Behaviour,' *Economica* 5(17): 61-71.
- Shaker, L. (2012). 'Local Political Knowledge and Assessments of Citizen Competence,' *Public Opinion Quarterly* 76(3): 525-537.
- Smith, M. (1995). 'Internal Reasons,' *Philosophy and Phenomenological Research* 55(1): 109-131.
- Stenner, K. (2005). *The Authoritarian Dynamic*. Cambridge, UK: Cambridge University Press
- Talisie, R. (2005). *Democracy After Liberalism*, New York, NY: Routledge.
- Toka, G. (2008). 'Citizen Information, Election Outcomes and Good Governance,' *Electoral Studies* 27: 31-44.

Tyler, T. R. (2006). *Why People Obey the Law*, 2nd Ed. Princeton, NJ: Princeton University Press.

Tyler, T. R. (2011). *Why People Cooperate: The Role of Social Motivations*, Princeton, NJ: Princeton University Press.

Venables, W. N. and B. D. Ripley (2002). *Modern Applied Statistics with S*. Fourth Edition. Springer, New York. ISBN 0-387-95457-0

Yen, W. M. (1993). 'Scaling Performance Assessments: Strategies for Managing Local Item Dependence,' *Journal of Educational Measurement* 30: 187-213.

Zaller, J. (1992). *The Nature and Origins of Mass Opinion*, Cambridge: Cambridge University Press.

Zaller, J. (2012). 'What *Nature and Origins* Leaves Out,' *Critical Review* 24(4): 569-642.

Appendix: Methodology for Simulation

A1. Data Set

The British Election Study data set (Fieldhouse *et al.* 2018) used for the simulation in Section 3.1 is from a face-to-face survey using an address-based random probability sample of eligible voters living in 468 wards in 234 Parliamentary Constituencies in England, Scotland, and Wales, for a total sample of 2,194 respondents. The full data set can be accessed at <https://www.britishelectionstudy.com/data-object/2017-face-to-face/>. Out of the 2,194 respondents, 902 reported voting Remain, 893 reported voting Leave, 348 reported not having voted, and 51 responded ‘Don’t know’. The ‘Don’t know’ respondents were deleted from the data set. Reported referendum votes were coded as 1 for Remain and 0 for Leave, with those reported as not having voted being coded as NA (i.e., no value). Once all incomplete responses across the variables in the models in A2-5 below were excluded, 1,467 observations remained.

A2. Ability and Authoritarianism Scales

The data set included six knowledge items. Those items were used to build a two-parameter item response theory (IRT) model in R (R Core Team 2017), using `ltm` (Rizopoulos 2006), in order to estimate the latent ability of respondents. The discrimination values for two of the six items suggested that they were too easy (i.e., not discriminating enough), and were therefore excluded. 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 Q_3 ; Yen 1993), and good model fit (evaluated through a plot of observed versus expected values).

Additionally, an authoritarianism scale (see A5 below) was built using items asking about young people not having enough respect for traditional British values, death penalty being appropriate for some crimes, and people breaking the law deserving stiffer sentences. The scale was built using a general partial credit model from R’s `mirt` package (Chalmers 2012). Each item loaded well onto the factor, and the item trace lines showed fairly well-ordered response categories. This model, too, had reasonable properties: a scree plot suggested unidimensionality, empirical plots for each item showed a good fit, and the maximum Q_3 value came out to -0.39, against an expected value of -0.5 for a scale with three items (Yen 1993).

A3. Logistic Model of Informed Referendum Vote

A logistic model was fitted using `glm` (in R’s `stats` package) with (reported) EU referendum vote as the dependent variable, and the ability score and the demographic variables (party identification, union membership, income, living situation, gender, religion, age, ethnic group, education, work status, marital status, and geographical region) as the

independent variables. The model suggested an increase in the logged odds towards Remain as the person becomes more informed ($p = 0.0007$), holding constant the remaining variables. The variance-inflation factors did not indicate any collinearity. Cook's distance values and standardised residuals didn't suggest any influential values or outliers. The McFadden value (0.20) suggested a good fit. The actual proportion of support for remain in the sample (49.4%) was calculated and compared to the simulated proportion of support for Remain, once the ability score for each participant had been increased to the maximum one (0.83), which came out to 54.5%, making for an information effect of 5.1 percentage points.

A4. Multinomial Model to Accommodate Informed Non-voting

To test whether the effect remained, even if every informed person was not assumed to vote, reported referendum votes (Remain and Leave) and non-voting were then re-coded as three nominal factors. By fitting a multinomial model with `nnet` (Venables and Ripley 2002), the probability of each respondent voting Remain, voting Leave, or not voting, if fully informed, was estimated with the same variables as in A3 above, giving a Remain proportion of 55%. The proportion of voters was predicted to increase from 87 (the proportion of reported voters in the sample) to 93%.

A5. Robustness Check

In order to also determine how sensitive the information effect in A3 was to changes in model specification, two further models were constructed: a purely demographic model, which was identical to the model in A3, save for excluding the quasi-attitudinal variables of party identity and union membership; and an authoritarianism model, which included all variables from the model in A3, as well as the respondent's score on the authoritarianism scale (see A2). For both of these models, the variance-inflation factors did not indicate any collinearity, and the Cook's distance values and standardised residuals didn't suggest any influential values or outliers. The McFadden values were 0.14 for the demographic model and 0.24 for the authoritarianism one. The latter being higher than the former (and also higher than the model in A2) is to be expected, given that authoritarianism is highly predictive in this context. The simulated proportions of support for Remain on the two models were 55.7% on the demographic model, for an information effect of 6.3 percentage points, and 51.9% on the authoritarianism model, for an information effect of 2.5 percentage points. Across all three models, we thereby saw a non-trivial information effect towards Remain, averaging 4.7 percentage points.