“Moderates”[*]

David E. Broockman[†] Benjamin E. Lauderdale[‡]

October 17, 2023

Abstract

Many Americans express a mix of conservative and liberal views across issues. Prior research indicates these voters are cross-pressured. A recent, influential article “Moderates” (Fowler et al. 2023) argues that these voters instead largely have centrist views on issues. To reach this conclusion, “Moderates” develops a method to determine which voters’ views are well-summarized by left-right ideology. “Moderates” finds that most voters’ views are, and therefore concludes that the large number of voters with centrist estimated ideologies—“moderates”—must hold centrist views on issues. We show that this method systematically overstates how many voters’ views are well-summarized by left-right ideology: it assumes voters’ views are unless they either answer questions randomly or form a single cluster with distinctive views. In simulations, we show this bias is large. The article’s core conclusion that voters who express a mix of conservative and liberal views can be inferred to support centrist policies therefore remains in doubt.

[*]We thank Jack Blumenau, Alex Coppock, Kevin DeLuca, Jamie Druckman, Elizabeth Elder, Anthony Folwer, Don Green, Chris Hanretty, Greg Huber, Dan Hopkins, Josh Kalla, Yph Lelkes, Gabe Lenz, Neil O’Brian, Eric Schickler, Sean Westwood, and seminar participants at the University of Chicago Harris School for helpful feedback. All remaining errors are our own.

[†]Associate Professor, Political Science, University of California, Berkeley. dbroockman@berkeley.edu, https://polisci.berkeley.edu/people/person/david-broockman.

[‡]Professor, Department of Political Science, University College London. b.lauderdale@ucl.ac.uk, https://benjaminlauderdale.net.

A large number of Americans express a mix of liberal and conservative views across issues. A recent, already-influential[1] article in the *American Political Science Review*, "Moderates" (Fowler et al. 2023), offers a novel interpretation of these voters' views. Prior research largely characterizes these voters as genuinely cross-pressured across issues (e.g., Ahler and Broockman 2018). "Moderates" notes that many might actually "hold centrist views" on issues, which they sometimes fail to express due to limited response options or measurement error. "Moderates" finds that (1) most ostensibly cross-pressured voters hold centrist views on issues, and (2) these centrist voters ("moderates") are "central to electoral change," being more likely to change their votes in response to candidate ideology and quality.

To reach these conclusions, "Moderates" does not measure whether voters explicitly support centrist policies; indeed, prior research shows that they rarely do (Broockman 2016). Rather, the article first develops a method to determine which voters' views on individual issues can be reliably predicted from their left-right ideology. It finds that the vast majority of voters' views are "well-summarized" by their left-right ideology, and so it is possible to infer their views on individual issues from their left-right ideology. Because voters who express different mixes of liberal and conservative views are all estimated to have moderate one-dimensional ideologies, "Moderates" therefore infers that most of these voters actually support centrist policies.

In this short paper, we show that, unfortunately, the method "Moderates" offers does not accomplish what it claims, and the article's empirical findings are therefore unreliable. The key bias in its method arises because, when categorizing voters as holding views well-explained or not-well-explained by one-dimensional ideology, the method assumes non-ideological voters are characterised by *one particular pattern* of views. This leads the method to systematically overestimate the share of voters whose views on issues can be inferred from their one-dimensional ideology: when truly non-ideological voters have a different set of views than the particular set the model assumes all non-ideologues have, the method often still classifies

[1] The article is among the most-cited articles in the *APSR* this year, and has been cited in the *New York Times*.

them as one-dimensional ideologues.

We first describe the "Moderates" method, why it yields biased results, and illustrate this with real and simulated data (with further examples in the supplemental appendix). In our main simulation, the "Moderates" model misclassified 75% of completely non-ideological voters as one-dimensional ideologues, and 44% of them "moderate" ideologues. We conclude by further elaborating the substantive implications of these contrasting findings and discussing ways forward toward the article's ambitious goals.

"Moderates"' Goals and Method

Converse (1964) coined the term "constraint" to describe the strength of correlations between Americans' views on different issues. Converse (1964) concluded that *ideological* constraint in particular was relatively rare—i.e., that most Americans' views were not structured by a single liberal-conservative dimension. Since Converse (1964), understanding to what extent American public opinion exhibits constraint in general or one-dimensional ideological constraint in particular has been a central scholarly focus (e.g., Ansolabehere, Rodden and Snyder 2008; Lauderdale, Hanretty and Vivyan 2018; Kinder and Kalmoe 2017). Scholars have studied this question for many reasons, including that, if left-right ideology predicts voters' views on issues well, voters' issue views can be inferred from their overall left-right ideology alone.

"Moderates" proposes a method to categorize voters into a group whose issue views are well-summarized by their left-right ideology (ideologues) and those whose views are not (non-ideologues). The upshot of this method is that the article can then focus on the former group when attempting to make inferences about voters' views on issues from their estimated left-right ideology. The downside is that, if the model misclassifies non-ideological voters as ideologues, it may produce inaccurate inferences about their views on issues from their estimated ideologies.

To identify ideologues, the method estimates a three mode mixture model which takes respondents' answers to binary survey questions at a single point in time as input and estimates the probability that respondents fall into one of three categories:

1. “Downsians” (i.e., one-dimensional ideologues) have responses that arise from a one-dimensional model of left-right ideology.

2. “Conversians” (i.e., non-ideologues) share a common response probability to each question.

3. “Inattentives” have a 50% probability of answering Yes to each question, as one might if one were randomly clicking in a survey.

Mathematically, the three modes of the mixture model each correspond to profiles of predicted probabilities of respondent i giving a “Yes” response to each binary issue question j. Where Λ is the cumulative logistic distribution function:

$$p\left(Y_{ij}=1|\text{“Downsian”}\right) = \Lambda\left(\beta_j(x_i-\alpha_j)\right) \tag{1}$$

$$p\left(Y_{ij}=1|\text{“Conversian”}\right) = \lambda_j \tag{2}$$

$$p\left(Y_{ij}=1|\text{“Inattentives”}\right) = 0.5 \tag{3}$$

Key Bias

“Moderates” claims that the “Conversian” category “flexibly” (p. 4, 5) captures respondents who express views that are “genuine” but “not well summarized by a single ideological dimension” (p. 1). But the model does not measure whether voters’ views are “*well*” or “*not well*” summarized by a single ideological dimension. It instead asks which of the three modes of the mixture model each respondent is *best* explained by. This is problematic because, as we will explain, it is easier for most voters’ views to be explained by the “Downsian” mode than the “Conversian” mode simply because the “Downsian” mode accommodates many subtypes of respondents but the “Conversian” mode does not allow more than one type of respondent to coexist within it.

In particular, in the model, the probability that "Conversians" answer "Yes" on any given issue j is λ_j for *all* Conversians. The model therefore assumes that "Conversians" are a homogeneous type which form a single cluster with a shared pattern of views. The "Conversian" mode of the model *can* thus effectively capture one group of voters with a particular idiosyncratic belief system—e.g., libertarians, where λ values might be high ('conservative') across all respondents of this type for economic issues but low ('liberal') for social issues.

However, contrary to what "Moderates" claims, voters of the exact type Converse imagined—voters with various patterns of beliefs not well-described by left-right ideology—cannot coexist in the model's "Conversian" category. For example, Converse (1964, p. 235) argues that voters' views towards social groups generate non-ideological constraint in their views, such as whether voters are "sympathetic to [Blacks] as a group" (see Elder and O'Brian 2022). Some voters might be highly sympathetic, affecting their views across a range of relevant issues; others may harbor animosity, and therefore hold the opposite set of positions on those issues. The "Moderates" model does not allow non-ideologues with hostility towards Blacks and non-ideologues with sympathy towards Blacks to coexist in the "Conversian" category. This makes the model's "Conversian" category an overly restrictive description of voters with "genuine" views "not well summarized by a single ideological dimension"—and certainly of Converse's (1964) conception of public opinion.

The model's restrictive conception of non-ideologues makes it vulnerable to miscategorizing a large share of true non-ideologues as moderate ideologues. Individuals whose responses do not precisely fit the model's chosen single cluster of "Conversians," the single cluster of "Inattentives," or any point on the "Downsian" dimension, are still necessarily classified as one of these by the model. The model typically describes most such respondents as "Downsians" simply because that mode of the mixture model is more flexible than the others in the response patterns it can describe: whereas the "Downsian" mode is consistent with varying profiles of responses across the range of estimated respondent ideology x_i, the "Conversian" mode is restricted to a single predicted probability profile. Voters whose responses are not well-explained by one-dimensional ideology yet do not closely resemble the single cluster of "Conversians" are thus typically categorized as "Downsians."

Understanding this Bias Graphically

A graphical illustration clarifies this source of bias. The points in Figure 1 show data from a simulated survey. In the simulation, 1,000 voters' answers to 100 survey questions are determined by voters' true latent views in one of two correlated dimensions and random measurement error.[2] One could imagine that these dimensions correspond with views on two distinct issues, such as abortion and taxes. The x- and y-axes correspond with these two underlying dimensions. The colors show the categorizations the "Moderates" method produced.

Voters with conservative views on both dimensions (upper-right) or liberal views on both dimensions (lower-left) could fairly be described as one-dimensional ideologues, as their views in both domains can be predicted well from an overall left-right ideology. Voters in the upper-left and lower-right have liberal views in one domain and conservative views in another. Their views cannot be reliably predicted by a single left-right ideological dimension, and therefore meet our (and Fowler et al.'s (2023)) definition of non-ideologues.

The bias in the "Moderates" model is evident: the model only characterizes non-ideologues in the top-left quadrant as non-ideologues ("Conversians"), even though voters in the bottom-right quadrant are similarly poorly described by one dimension. Although Fowler et al. (2023, p. 1) leave vague what it means for voters' views to be "well" or "not well" summarized by one dimension, it is clear that voters in the top-left and bottom-right are similarly "not well" summarized by one dimension. Yet the latter are miscategorized as one-dimensional ideologues ("Downsians")—thereby overstating the share of one-dimensional ideologues for whom views on individual issues can be inferred from left-right ideology. Furthermore, because the non-ideologues at the bottom-right are near the middle of one-dimensional ideology, they would

[2] This simulation uses a two-dimensional data-generating process exclusively to facilitate graphical exposition. The next section presents a highly multidimensional simulation.

Figure 1: Motivating Example in Simulated Dataset

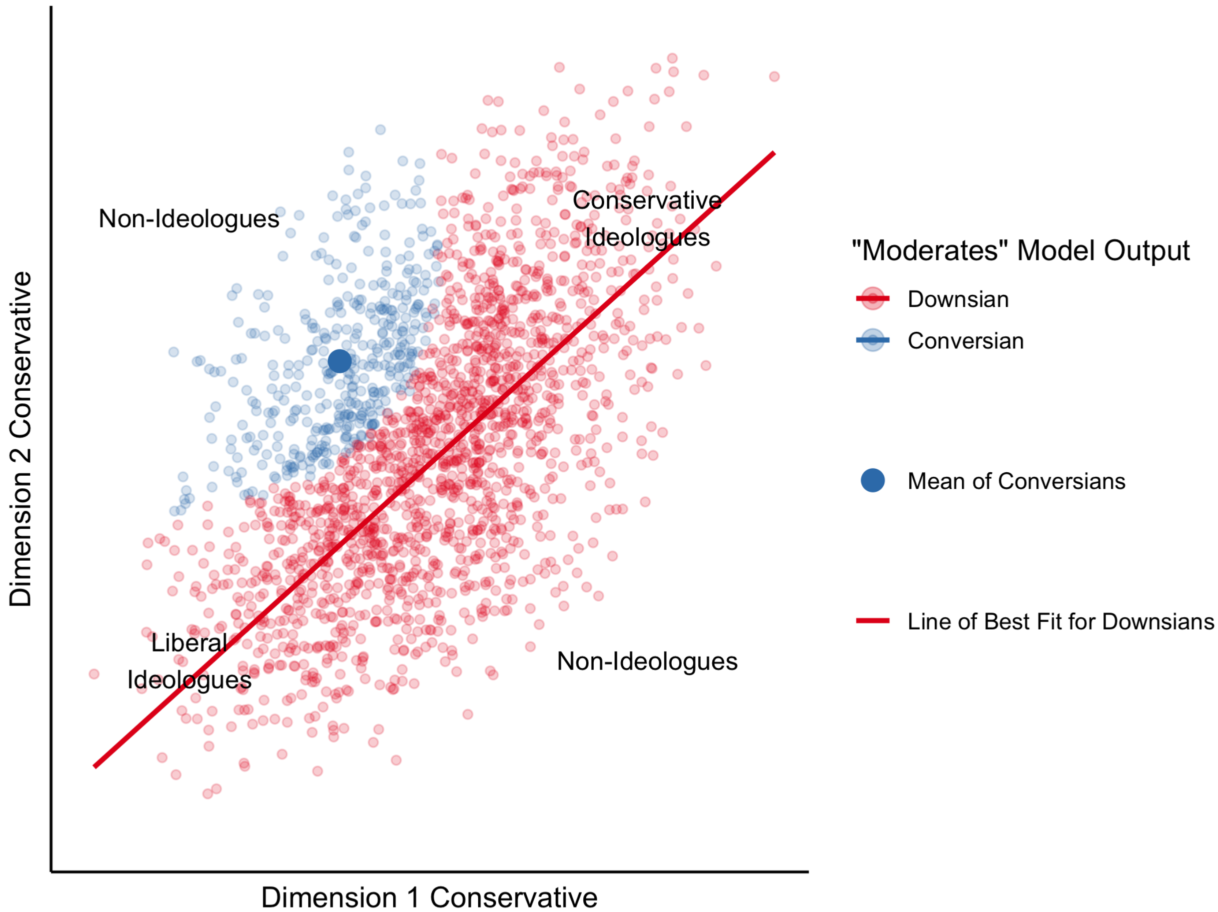

be described as "moderates" and inferred to have centrist views on issues—despite actually being conservative on one issue and liberal on another.

The "Moderates" method makes these mistakes because, within the context of multidimensional data, the method asks two questions. First, the model asks: what is the best way to place a freely movable line (the estimated ideological spectrum for Downsians) and a freely movable point (the estimated mean issue positions of Conversians) such that the distance of all the points from the movable line, the movable point, and the fixed point (Inattentives) is minimized?

Second, to categorize respondents, the method then asks whether each point is closest to the line (Downsians), to the movable point (Conversians), or to the fixed point (Inattentives). The answer is often that more points are closer to the movable line than either of the two points, as was the case in Figure 1. This is simply because a line is "bigger" than a point, i.e., it is a more flexible description of the data than the two points are, not because of anything about the structure of ideology. In other words, the main finding of "Moderates" that more voters are labeled as Downsian (close to a line) than Conversian (close to a point) is a feature of geometry not politics.

Our argument in the previous subsection can be understood in terms of this graphical exposition. In the context of the geometry of a multidimensional spatial model,[3] the model selects only *one* point to represent the views of "Conversians" (a vector of λ_js shared across all Conversian respondents)—the blue point in Figure 1. This corresponds with what we noted earlier, that in the "Moderates" model, non-ideologues cannot have multiple different profiles of genuine views other than those that arise from left-right ideology. But the model allows ideologues to have a variety of response patterns—anywhere closer to the red line than the blue point. This leads the model to estimate that there are far more ideologues than non-ideologues—as witnessed in how a substantial share of (one-dimensional) non-ideologues were categorized as ideologues in Figure 1.[4]

Figure 2 shows that the same pattern appears in the datasets "Moderates" used. The colors again depict the categorizations the "Moderates" model produced in the article's own datasets. To display the data graphically, we follow "Moderates" in using a (two-dimensional) IRT model to estimate two latent dimensions, and orient them so that larger values correspond with conservative positions. "Conversians" largely occupy a narrow cluster in every dataset. Furthermore, where this is varies politically by year: "Conversians" lean conservative in some years, moderate in others, and liberal in others. This indicates that "Conversians" are not a stable group "Moderates" discovered.[5]

[3] All possible respondent by issue matrices of data from all possible data generating processes can be redescribed in terms of multidimensional spatial models. This is the basis of principle components analysis, factor analysis, and related methods including spatial voting models.

[4] Figure OA2 shows that similar results obtain when the latent dimensions are uncorrelated.

[5] Figure OA1 shows similar results in our simulated data.

Figure 2: Motivating Example Applied to Datasets in "Moderates"

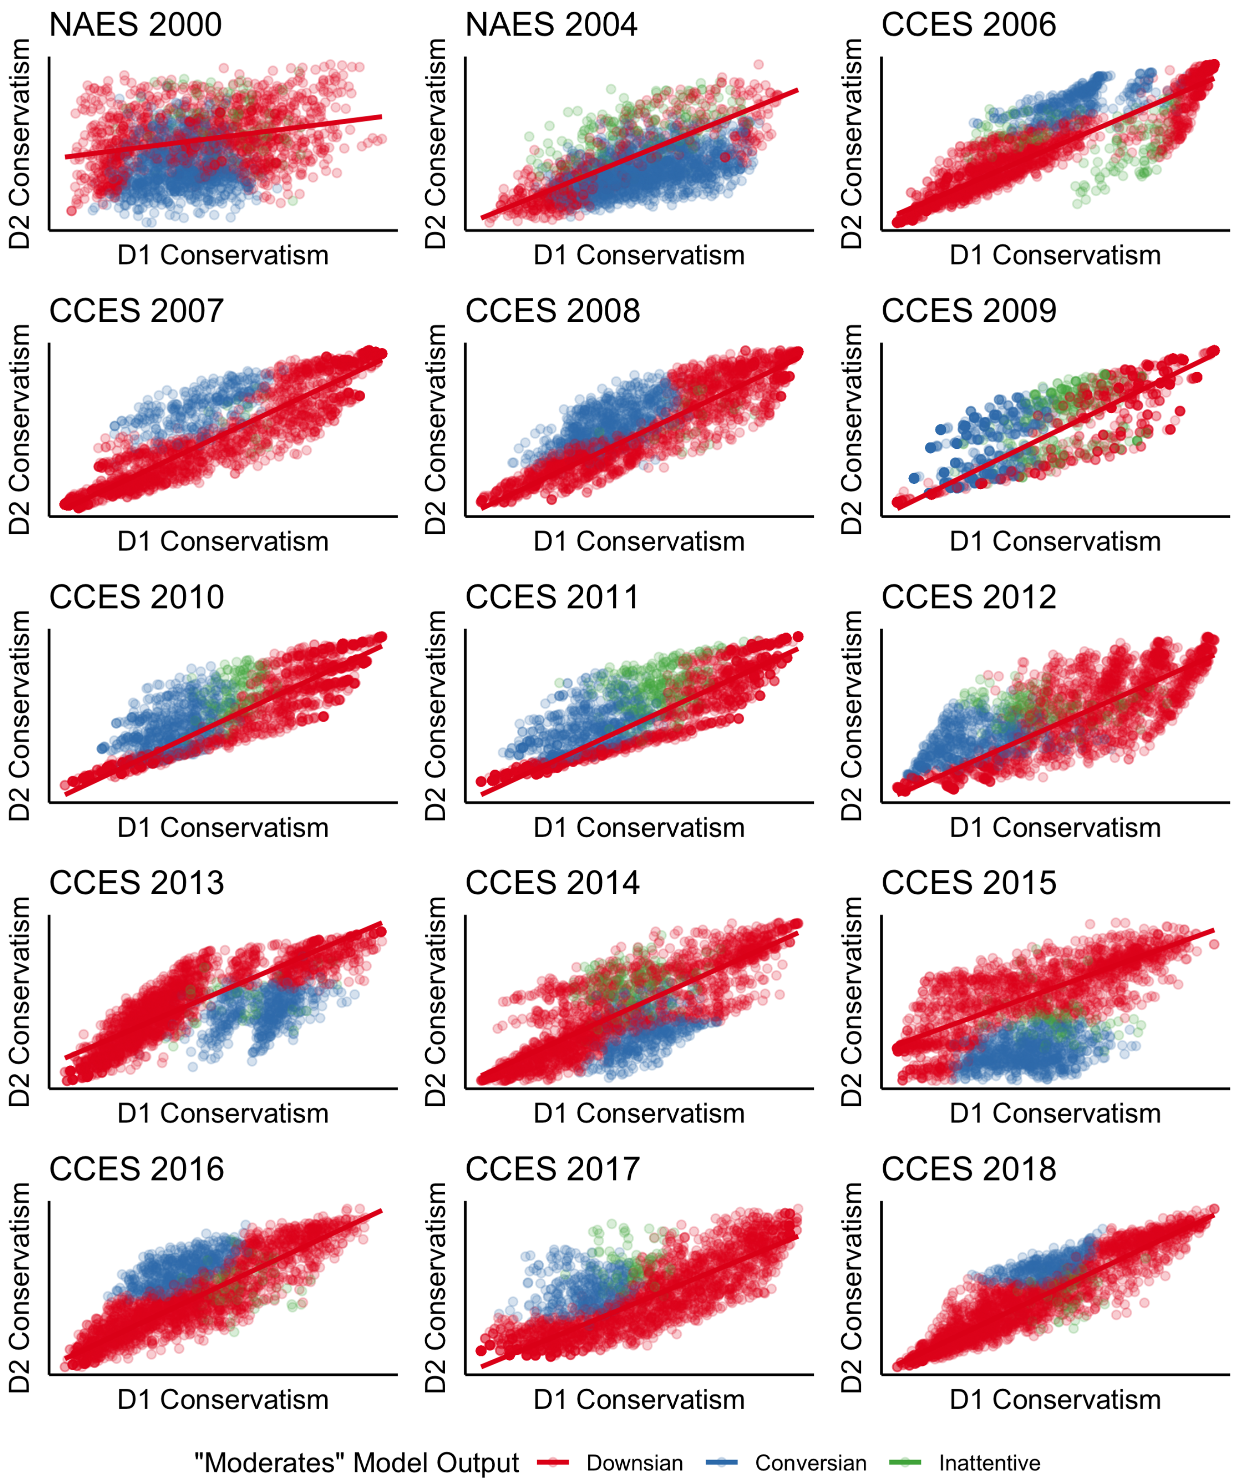

Notes: Appendix A2 discusses details.

In this section we considered two-dimensional examples for the purpose of straightforward graphical exposition.[6] In the next section we evaluate this bias in simulations of highly multidimensional voters, and show that it is even more severe.

Estimates in Simulated Data

We next present simulations showing that the "Moderates" model indicates there are more one-dimensional ideologues ("Downsians") than non-ideologues ("Conversians") in data from a variety of plausible data-generating processes, including from complete non-ideologues.[7] In these simulations, respondents have true latent views in each of 15 different issue domains. We simulate several data generating processes (DGPs) that produce these true latent views. 5,000 respondents from each DGP answer a simulated survey. The survey asks 5 issue questions in each of the 15 issue domains,[8] for a total of 75 questions. We use a Rasch model to probabilistically translate respondents' latent views in the 15 issue domains into observed binary responses to the 75 questions.[9] Finally, we apply the "Moderates" model to the observed binary responses. See Appendix B for details.

Our main simulation relies on three DGPs with varying ideological constraint, as ideological constraint varies across real voters (Kinder and Kalmoe 2017). In particular, we simulate:

- A "Non-Ideologues" DGP. Respondents in this data-generating process exhibit no ideological constraint; their views in each issue domain are entirely uncorrelated with those in all other issue domains.

- A "Some Ideological Constraint" DGP where ideology determines half of the variation in

[6]Appendix D in "Moderates" finds that a two-dimensional version of the model classifies even more voters as ideologues. The authors interpret this as indicating that the one-dimensional model potentially underestimates the share of ideologues. We interpret this as that their two-dimensional mixture model's Downsian submodel is even more flexible relative to the Conversian submodel than the one-dimensional version.

[7]"Moderates" presents simulation studies, but these only show that the model recovers the data-generating parameters when the data are generated from the model itself, not from alternative models of non-ideological voters.

[8]This creates a form of non-ideological constraint, as Converse (1964) theorized.

[9]This represents the presence of measurement error when latent views are translated into observed responses and/or genuine opinion variation within issue domains.

respondents' views in every issue domain.

- A "One-Dimensional Ideologues" DGP, where respondents' cross-domain ideology solely determines respondents' views in every issue domain.

If it functions properly, the "Moderates" model should identify the "Non-Ideologues" DGP's voters as "Conversians," as the article claims that the "Conversian" category captures respondents who express views that are "genuine" but "not well summarized by a single ideological dimension" (p. 1). By contrast, the latent views held by the "One-Dimensional Ideologues" group are entirely determined by one-dimensional ideology, matching the definition of "Downsians" in "Moderates." By construction, the share of individuals' true latent issue views which are moderate[10] are identical in every DGP.

We combine 5,000 responses from each of these three DGPs into one dataset, and then estimate the "Moderates" model in this dataset. Table 1 shows the results.

Table 1: Simulation 1 Results

Ground Truth in DGP		**"Moderates" Model: Share Categorized As...**			
True Type	Share of True Latent Preferences Moderate	Ideologues ("Downsians")	Moderate Ideologues ("Downsian Moderates")	Non-Ideologues ("Conversian")	Inattentive
Non-Ideologues (100% "Conversian")	33%	75%	44%	25%	0%
Some Ideological Constraint (Semi-"Downsian")	33%	88%	24%	12%	0%
One-Dimensional Ideologues (100% "Downsian")	33%	94%	18%	6%	0%

Table 1's first row shows that the "Moderates" model miscategorized *three-quarters* of non-

[10]Moderate latent issue preferences are defined as those in the middle third of each issue's distribution, mirroring Fowler et al.'s (2023) definition of moderate latent ideology.

ideological respondents as ideologues ("Downsians"). These respondents are simulated to match Fowler et al.'s (2023) verbal description of the "Conversian" type, yet are usually categorized as "Downsian."

In addition, the "Moderates" model disproportionately miscategorizes non-ideologues as *moderate* ideologues ("moderates"). In our simulations, non-ideological, semi-ideological, and purely ideological respondents' true latent views are similarly moderate. However, the model classified about twice as many non-ideologues as moderate ideologues. This result is consistent with Broockman (2016): respondents with *any* mix of liberal or conservative issue views are disproportionately likely to have moderate estimated ideologies, even if their views on those issues are not moderate.

In Appendix B.3, we present a second simulation which includes several more DGPs, including pure partisans and inattentive respondents. Of greatest interest is a "Libertarian" DGP we introduce, motivated by our earlier discussion. As we noted, the model's "Conversian" category does not formalize a diverse variety of voters with genuine views unconstrained by one-dimensional ideology, it rather formalizes a single cluster of voters with a distinctive pattern of views, such as libertarians. In Appendix B.3 we confirm this by showing that, when introducing a single cluster of voters with distinctive views ("Libertarians"), the model labels them all as "Conversian" non-ideologues—but almost no other voters as such, including almost none of the true non-ideologues. Figures OA3 and OA4 illustrate this visually.

Appendix B.4 presents a qualitatively different simulation inspired by Converse (1964) where non-ideological voters' views are informed by their views towards various groups. We find that the "Moderates" model similarly misclassifies most of these voters as Downsians.

In summary, our simulations confirm three key points regarding the "Moderates" model:

- The "Conversian" category is well-suited to identify a group of respondents who share *the same* views, but does not broadly capture respondents with non-ideological views.

- The model typically miscategorizes other non-ideological voters who do not fall into this cluster as ideologues, dramatically overestimating the share of voters who are ideologues.
- The model disproportionately miscategorizes non-ideological cross-pressured voters as *moderate* ideologues, i.e., as “moderates,” because they hold a mix of liberal and conservative views.

In Appendix C, we generate simulated datasets that have the same sample sizes and pairwise correlation structures as the datasets “Moderates” analyzes. These simulated datasets contain only a single category of respondents, but the “Moderates” model still categorizes about the same proportion of respondents in them as “Downsians” as in the original datasets. This implies that the proportion of “Downsians” that “Moderates” estimated in the original datasets cannot reliably indicate how many respondents truly belong to different underlying types, as it returns the same answer in datasets where all respondents are of the same type. The proportion of each type the model estimates instead appears to depend largely on the pairwise correlation structures of issue positions, which were preserved in these simulations.

Contrary Evidence

We have shown that, because the model “Moderates” uses is unreliable, the evidence it adduces leaves us unclear how many voters are one-dimensional ideologues and, therefore, to what extent we can infer that voters have centrist views on issues if their left-right ideology is estimated to be moderate. However, other studies offer insight, and reach different conclusions.

First, with respect to how many voters are one-dimensional ideologues, standard factor analytic approaches find that only approximately 30% of the variation in expressed opinions can be explained by a single dimension (Broockman 2016). Using panel data, Lauderdale, Hanretty and Vivyan (2018) estimate that left-right ideology represents only about 25% of the persistent variation in public opinion. Freeder, Lenz and Turney (2019) find that knowledge of party positions largely explains whether people hold stable ideological views, and similarly conclude that only about 20-40% do so. Finally, Ahler and Broockman (2018) find that only 30% of voters prefer candidates whose issue positions convey a broad ideology matching their own, rather than candidates which match their idiosyncratic issue views. While the quantities of interest are defined differently, these studies all converge on estimates of the prevalence of left-right ideologues or ideology far lower than the 73% in "Moderates."

Second, with respect to how many voters support moderate policy, Broockman (2016) finds that only 18% of voters express support for policy more moderate than the two parties on the typical issue, and that voters with a mix of liberal and conservative views are no more likely to support moderate policies than consistently liberal or consistently conservative voters. By contrast, "Moderates" argues that voters with a mix of liberal and conservative views largely all support moderate policies.

These debates have important substantive implications. They center on the question of how to understand the large number of voters who express a mix of liberal and conservative views on surveys. Understanding these voters is crucial because, as scholars have long known (Lazarsfeld, Berelson and Gaudet 1944) and "Moderates" re-confirms, they are disproportionately likely to be 'floating' or 'swing' voters. If "Moderates" is correct, because these voters are mostly ideologues, that they have moderate estimated ideologies indicates that they agree with each other about what they want from government: moderate policy. But other studies generally find that this is in fact a highly heterogeneous group who want different things from government than each other—many of which are not moderate—and therefore whom it is neither straightforward nor obviously desirable for politicians to satisfy (Ahler and Broockman 2018; Broockman 2016; Lauderdale, Hanretty and Vivyan 2018).

Moving Forward

"Moderates" outlines an ambitious goal: categorizing whether individual voters are one-dimensional ideologues. The paper's creative reorientation of previous literature's focus on characterizing groups or issues to individuals holds tremendous promise. Unfortunately, this goal cannot be accomplished with tweaks to the “Moderates” model because different data are required to answer this research question, and an entirely different model capable of handling that data.

In particular, it is no accident that the contrasting studies reviewed above all rely on panel data. Panel data can determine whether expressed views which are extreme or inconsistent with left-right ideology are likely to be genuine or due to measurement error. It does so by measuring whether respondents register these same views again weeks or months later. By contrast, lacking panel data, “Moderates” is forced to rely on a restrictive modeling assumption to distinguish between measurement error and genuine non-ideological views: namely, that only one cluster of non-ideological views exists, and that non-ideological responses from respondents not in this cluster reflect measurement error.

In addition, “Moderates” relies largely on binary data from existing surveys where the two response options are each supported by one of the two major parties. Answers to such questions cannot speak to whether voters would prefer a missing moderate option between the options presented—again forcing “Moderates” to rely on a restrictive assumption: that a mix of liberal and conservative responses *across* questions can be interpreted as support for a missing centrist option on those questions.

While the data and methodology in “Moderates” therefore cannot support the article’s substantive claims, its novel ambition to distinguish between one-dimensional ideologues and non-ideologues is nevertheless well worth pursuing with richer datasets and different methods.

References

Ahler, Douglas and David Broockman. 2018. “The delegate paradox: Why polarized politicians can represent citizens best.” *Journal of Politics* 80(4):1117–1133.

Ansolabehere, Stephen, Jonathan Rodden and James Snyder. 2008. “The strength of issues: Using multiple measures to gauge preference stability, ideological constraint, and issue voting.”

American Political Science Review 102(2):215–232.

Broockman, David. 2016. "Approaches to Studying Policy Representation." *Legislative Studies Quarterly* 41(1):181–215.

Converse, Philip E. 1964. The nature of belief systems in mass publics. In *Ideology and its discontents*, ed. David Apter. New York: Glencoe Free Press.

Elder, Elizabeth and Neil O'Brian. 2022. "Social groups as the source of political belief systems." *American Political Science Review* 116(4):1407–1424.

Fowler, Anthony, Seth Hill, Jeffrey Lewis, Chris Tausanovitch, Lynn Vavreck and Christopher Warshaw. 2023. "Moderates." *American Political Science Review* 117(2):643–660.

Freeder, Sean, Gabriel Lenz and Shad Turney. 2019. "The importance of knowing "what goes with what": Reinterpreting the evidence on policy attitude stability." *Journal of Politics* 81(1):274–290.

Kinder, Donald and Nathan Kalmoe. 2017. *Neither liberal nor conservative: Ideological innocence in the American public*. Chicago.

Lauderdale, Benjamin, Chris Hanretty and Nick Vivyan. 2018. "Decomposing public opinion variation into ideology, idiosyncrasy, and instability." *Journal of Politics* 80(2):707–712.

Lazarsfeld, Paul, Bernard Berelson and Hazel Gaudet. 1944. *The people's choice*. Columbia.

Online Appendix

Contents

A Visualizing Fowler et al.'s (2023) "Conversians' '

A.1 Figure 1

In this section we show how we simulated two-dimensional data for Figure 1. We use the following code:

```
draw.2d.model.data <- function() {
  n.respondents <- 2500
  n.issue.domains <- 2
  n.issue.questions.per.domain <- 50

  # generate two correlated dimensions
  data.2d <- data.frame(true.type = rep("Two-Dimensional", n.respondents))
  generate.cor <- rnorm(n.respondents)
  data.2d$latent.d1 <- .75 * rnorm(n.respondents) + .75 * generate.cor
  data.2d$latent.d2 <- .75 * rnorm(n.respondents) + .75 * generate.cor

  # generate preferences within each domain
  for (j in 1:n.issue.domains) {
    if (j == 1) {
      latent.issue.domain.views <- rnorm(n.respondents) * .5 + data.2d$latent.d1
    } else {
      latent.issue.domain.views <- rnorm(n.respondents) * .5 + data.2d$latent.d2 * .75
    }
    for (k in 1:n.issue.questions.per.domain) {
      # Rasch model translates latent issue views into binary outcome
      b <- rnorm(1) # difficulty parameter
      data.2d[, paste0("issueopinion_", j, ".issue", k)] <- as.numeric(
        exp(latent.issue.domain.views - b) / (1 + exp(latent.issue.domain.views - b)) >
        ↪ runif(length(latent.issue.domain.views))
      )
    }
```

```
  }

  mod <- data.2d %>%
    select(starts_with("issueopinion_")) %>%
    mirt(data = ., model = 2)
  scores <- fscores(mod)
  data.2d$d1 <- scores[, 1]
  data.2d$d2 <- scores[, 2]

  return(data.2d)
}
```

We then add the estimates from the Fowler et al. (2023) model and plot them alongside 2D IRT estimates of the two underlying dimensions from the mirt package. As shown and discussed in the main text, plotting these according to their first and second dimensions illustrates that the Moderates model categorizes respondents closer to a *single point* point it selects than the main one-dimensional line it estimates as Conversian.

```
make.fig1 <- function(sim.number) {
  data.2d.withests <- draw.2d.model.data() %>%
    add.moderates.estimates()

  data.2d.withests.downsians <- filter(data.2d.withests, modpaper_category == 'Downsian')

  conversians.mean <- data.2d.withests %>%
    filter(modpaper_category == 'Conversian') %>%
    summarize(d1 = mean(d1), d2 = mean(d2))

  text <- data.frame(
    text = c('Conservative\nIdeologues', 'Liberal\nIdeologues', 'Non-Ideologues',
     ↪  'Non-Ideologues'),
    d1 = c(2.75, -2.75, 2.75, -2.75),
    d2 = c(2.25, -2.25, -2.25, 2.25)
  )

  g <- data.2d.withests %>%
    sample_n(size = 2000) %>% # make graph readable
    mutate(modpaper_category = factor(modpaper_category, ordered = T, levels =
                                        c('Downsian', 'Conversian', 'Inattentive'))) %>%
    ggplot(aes(x = d1, y = d2)) +
    geom_point(aes(color = modpaper_category), alpha = .2) +
    geom_point(data = conversians.mean, aes(x = d1, y = d2, color = 'Conversian',
                                            fill = 'Mean of Conversians'), size = 4) +
    geom_smooth(data = data.2d.withests.downsians,
                aes(x = d1, y = d2, color = 'Downsian',
                    linetype = "Line of Best Fit for Downsians"),
                se = FALSE,
                formula = y ~ x,
                method = 'lm') +
    scale_color_brewer('"Moderates" Model Output', palette = "Set1") +
    scale_linetype('') + scale_fill_brewer('') +
    ylab('Dimension 2 Conservative') + xlab('Dimension 1 Conservative') +
    guides(color = guide_legend(order = 1, override.aes = list(alpha = .3)),
           fill = guide_legend(override.aes = list(color = "#377EB8"), order = 2),
           linetype = guide_legend(override.aes = list(color = '#E41A1C'), order = 3)) +
```

```r
    theme_classic() +
    theme(axis.text.x = element_blank(), axis.ticks.x = element_blank(),
          axis.text.y = element_blank(), axis.ticks.y = element_blank()) +
    geom_text(data = text, aes(label = text), size = 3.75) +
    expand_limits(x = c(-4, 4), y = c(-4, 4))
  return(g)
}

# select two examples where conversians are above the line and two examples where
↪  conversians are below the line
gs <- alply(1:8, 1, make.fig1, .parallel = TRUE)
above <- c()
below <- c()
for(i in 1:length(gs)) {
  if( mean(gs[[i]]$data$d2[gs[[i]]$data$conversian]) > 0 ) {
    above <- c(above, i)
  } else {
    below <- c(below, i)
  }
}

g <- ggpubr::ggarrange(gs[[above[1]]] + ggtitle('Simulation 1') + theme(legend.position =
↪  'none'),
                       gs[[below[1]]] + ggtitle('Simulation 2') + theme(legend.position =
↪  'none'),
                       gs[[below[2]]] + ggtitle('Simulation 3') + theme(legend.position =
↪  'none'),
                       gs[[above[2]]] + ggtitle('Simulation 4') + theme(legend.position =
↪  'none'),
                       ncol=2, nrow=2, legend="none")
```

Furthermore, Figure OA1 shows that where "Conversians" are placed in this space varies across random simulations, as Conversians are not a meaningful group.

Figure OA1: Motivating Example: Multiple Simulations

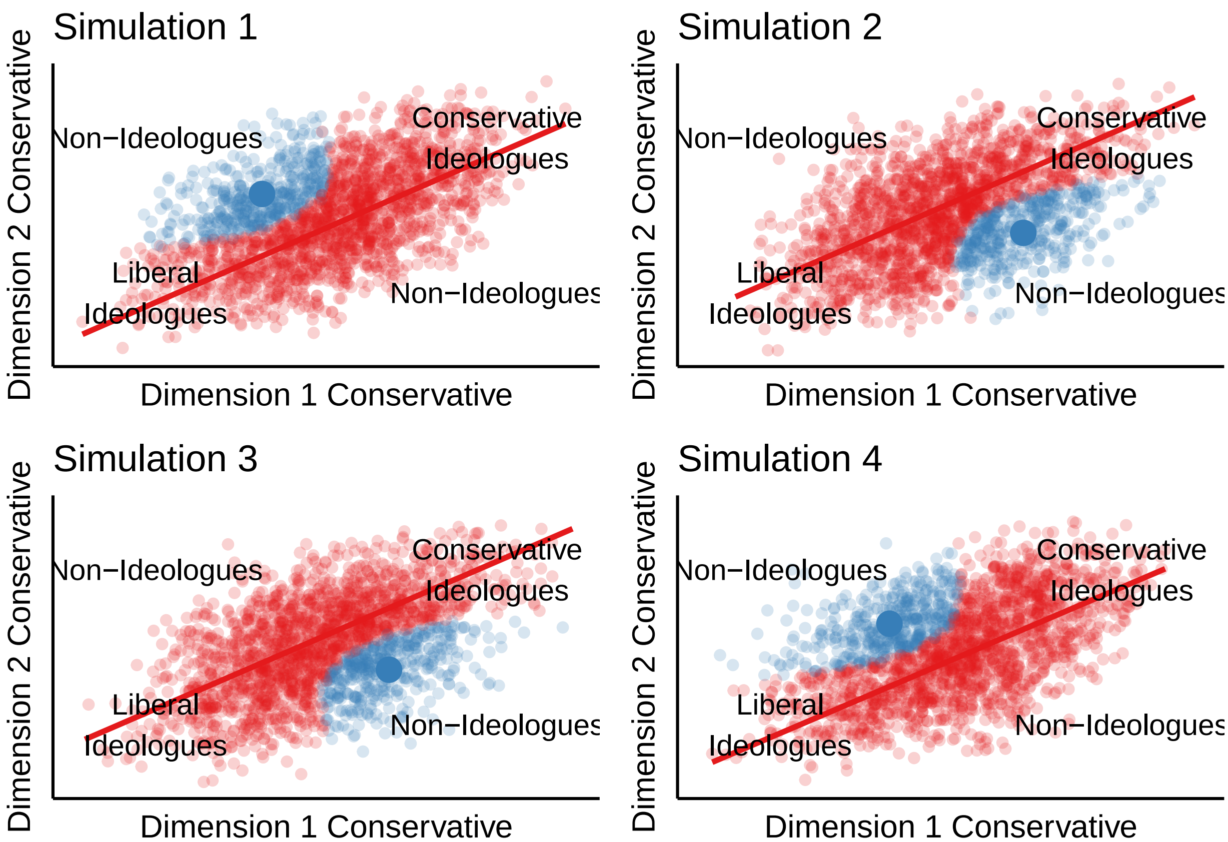

A.1.1 Version of Figure 1 With Uncorrelated Dimensions

In this section we show that the results of Figure 1 look similar even if the two underlying dimensions are not correlated.

As before, we first simulate the data.

```
draw.2d.model.data <- function() {
  n.respondents <- 2500
  n.issue.domains <- 2
  n.issue.questions.per.domain <- 50

  # generate two correlated dimensions
  data.2d <- data.frame(true.type = rep("Two-Dimensional", n.respondents))
  data.2d$latent.d1 <- rnorm(n.respondents)
  data.2d$latent.d2 <- rnorm(n.respondents)

  # generate preferences within each domain
  for (j in 1:n.issue.domains) {
    if (j == 1) {
      latent.issue.domain.views <- rnorm(n.respondents) * .5 + data.2d$latent.d1
    } else {
      latent.issue.domain.views <- rnorm(n.respondents) * .5 + data.2d$latent.d2 * .75
    }
    for (k in 1:n.issue.questions.per.domain) {
      # Rasch model translates latent issue views into binary outcome
      b <- rnorm(1) # difficulty parameter
      data.2d[, paste0("issueopinion_", j, ".issue", k)] <- as.numeric(
        exp(latent.issue.domain.views - b) / (1 + exp(latent.issue.domain.views - b)) >
         ↪  runif(length(latent.issue.domain.views))
      )
    }
  }

  mod <- data.2d %>%
    select(starts_with("issueopinion_")) %>%
    mirt(data = ., model = 2)
  scores <- fscores(mod)
  data.2d$d1 <- scores[, 1]
  data.2d$d2 <- scores[, 2]

  return(data.2d)
}
```

As before, we then add the estimates from the Fowler et al. (2023) model and plot them alongside 2D IRT estimates of the two underlying dimensions from the mirt package. As shown and discussed in the main text, plotting these according to their first and second dimensions illustrates that the Moderates model categorizes respondents closer to a *single point* point it selects than the main one-dimensional line it estimates as Conversian.

```
data.2d.withests <- draw.2d.model.data() %>%
  add.moderates.estimates()
```

Figure OA2: Motivating Example With Uncorrelated Dimensions

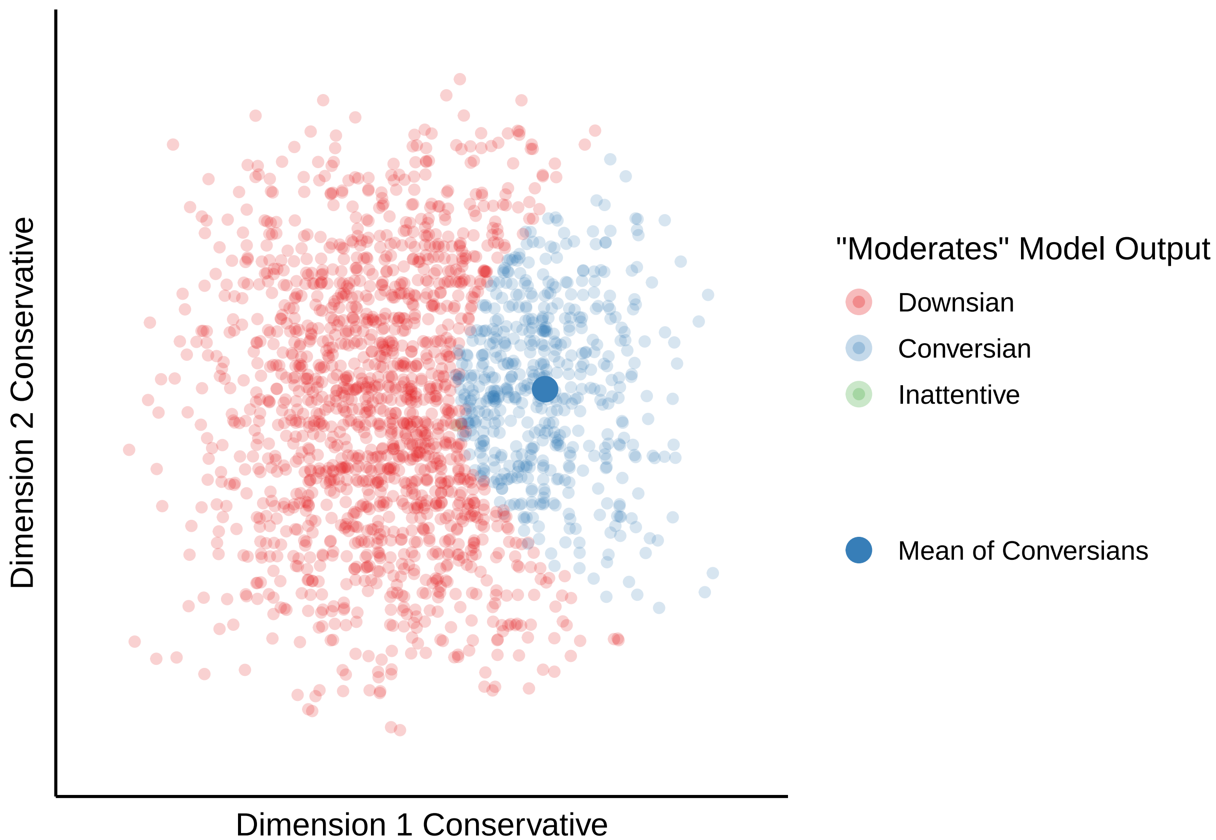

A.2 Figure 2

In this section we show how we generated Figure 2, which made a version of the same plot but in the data used in Fowler et al. (2023). Note that the 2D IRT model from the mirt package is used strictly for the purposes of visualization and does not affect the estimates the “Moderates” model produces, which we draw directly from the article’s replication data.

```
# Make plots
football.plot.moderates.data <- function(df) {
  source <- df %>%
    pull(source) %>%
    unique() %>%
    str_replace("_", " ")

  issues.mat <- df %>%
    select(cces2006_minimumwage:cces2017_buyamerican) %>%
    select_if(~ any(!is.na(.))) %>% # keep columns with any data, drop columns with all
     ↪  NA
    as.matrix()

  # 2d irt estimates from mirt package
  mod <- mirt(data = issues.mat, model = 2)
  scores <- fscores(mod)
  mirt.d1 <- scores[, 1]
  mirt.d2 <- scores[, 2]

  df <- df %>%
    mutate(
      d1 = mirt.d1,
      d2 = mirt.d2
    )

  # orient dimensions so larger values = more conservative
  if (cor(as.numeric(df$pid3), df$d1, use = "complete") < 0) df$d1 <- -1 * df$d1
  if (cor(as.numeric(df$pid3), df$d2, use = "complete") < 0) df$d2 <- -1 * df$d2

  df <- df %>%
    dplyr::rename(
      downs.prob = w1,
      conversian.prob = w2,
      inattentive.prob = w3
    ) %>%
    mutate(
      downsian = downs.prob > inattentive.prob & downs.prob > conversian.prob,
      conversian = conversian.prob > inattentive.prob & conversian.prob > downs.prob,
      inattentive = inattentive.prob > conversian.prob & inattentive.prob > downs.prob,
      modpaper_category = case_when(
        downsian ~ "Downsian",
        conversian ~ "Conversian",
        inattentive ~ "Inattentive",
        TRUE ~ NA_character_
      )
    )

  df.downsians <- filter(df, downsian)
```

```r
  g <- df %>%
    sample_n(2500) %>% # make graphs readable
    mutate(modpaper_category = factor(modpaper_category,
      ordered = T, levels =
        c("Downsian", "Conversian", "Inattentive")
    )) %>%
    ggplot(aes(x = d1, y = d2)) +
    geom_point(aes(color = modpaper_category), alpha = .2) +
    geom_smooth(
      data = df.downsians,
      aes(
        x = d1, y = d2, color = "Downsian",
        linetype = "Line of Best Fit"
      ),
      formula = y ~ x,
      se = FALSE, method = "lm"
    ) +
    scale_color_brewer('"Moderates" Model Output', palette = "Set1") +
    scale_linetype("") +
    scale_fill_manual("", values = "black") +
    ggtitle(source) +
    xlab("D1 Conservatism") +
    ylab("D2 Conservatism") +
    guides(
      color = guide_legend(order = 1),
      fill = FALSE, # guide_legend(override.aes = list(color = "#377EB8"), order = 2),
      linetype = FALSE
    ) + # guide_legend(override.aes = list(color = '#E41A1C'), order = 3)) +
    theme_classic() +
    theme(
      axis.text.x = element_blank(), axis.ticks.x = element_blank(),
      axis.text.y = element_blank(), axis.ticks.y = element_blank(),
      legend.position = "bottom"
    )

  return(g)
}

data <- read.dta(paste0(wd, "/Moderates Paper and Replication Archive/Moderates
 ↪ Replication Archive/bigsurveys_recoded4.dta"))
gs <- dlply(data, .(source), football.plot.moderates.data, .parallel = TRUE)
```

B Simulations

In this Appendix section we describe details of the simulations described in Table 1 (simulation 1) and the extended simulations reported in Table OA1 (simulation 2).

B.1 Global Parameters

```
set.seed(95821)
n.respondents <- 5000
n.issue.domains <- 15
n.issue.questions.per.domain <- 5
n.questions.total <- n.issue.domains * n.issue.questions.per.domain
```

Both simulations contain data from one or multiple DGPs. Each DGP has 5000 respondents. There are 15 issue domains represented on the survey. They are asked 5 questions in each issue domain, for a total of 75 issue questions.

B.2 Main Text Simulation (Simulation 1)

B.2.1 DGPs

Below we show the code used to generate the simulated survey responses in each DGP. DGPs 1-3 are used in the paper's main text. The remaining DGPs are used in a second simulation below.

The simulation code for each DGP below simulates the latent views of respondent i on issue j, denoted $\theta_{i,j}$, for each question for each respondent in the DGP. Note that these θ parameters are latent parameters representing respondent's true latent preferences specific to each respondent-issue. Further details about how these are translated into binary responses to a simulated survey instrument are given in the following subsection.

B.2.1.1 DGP 1: Non-Ideologues Voters simulated in this DGP have genuine views in each policy domain but these are uncorrelated with their views in other policy domains.

```
data.non.ideologues <- data.frame(true.type = rep("No Ideological Constraint",
↪  n.respondents))
for (j in 1:n.issue.domains) {
  latent.issue.domain.views <- rnorm(n.respondents)
  for (k in 1:n.issue.questions.per.domain) {
    data.non.ideologues[, paste0("latentissueopinion_domain", j, ".issue", k)] <-
↪  latent.issue.domain.views
  }
}
```

B.2.1.2 DGP 2: Some Ideological Constraint Voters simulated in this DGP's views in each policy domain are equally influenced by genuine views specific to each policy domain and a cross-domain ideology. This generates the presence of some "ideological constraint" in their beliefs.

```
# Note: parameter 1/sqrt(2) chosen because it satisfies the system of equations x^2 + x^2
↪  = 1^2 (so that latent issue opinions have the same SD in these DGPs as in
↪  data.non.ideologues from DGP 1).

data.some.constraint <- data.frame(true.type = rep("Some Ideological Constraint",
↪  n.respondents))
latent.ideology <- rnorm(n.respondents)
for (j in 1:n.issue.domains) {
  latent.issue.domain.views <- 1 / sqrt(2) * rnorm(n.respondents) + 1 / sqrt(2) *
↪  latent.ideology
```

```
  for (k in 1:n.issue.questions.per.domain) {
    data.some.constraint[, paste0("latentissueopinion_domain", j, ".issue", k)] <-
 ↪  latent.issue.domain.views
  }
}
```

B.2.1.3 DGP 3: One-Dimensional Ideologues This simulation models one-dimensional ideologues. Ideology is perfectly predictive of their latent views in every issue domain.

```
data.ideologues <- data.frame(true.type = rep("Pure Ideologues", n.respondents))
ideology <- rnorm(n.respondents)
for (j in 1:n.issue.domains) {
  for (k in 1:n.issue.questions.per.domain) {
    data.ideologues[, paste0("latentissueopinion_domain", j, ".issue", k)] <- ideology
  }
}
```

B.2.2 Converting θs to responses (for non-inattentive types)

We next probabilistically convert the respondent i's $\theta_{i,j}$ parameters on each issue j to binary issue responses to each issue question using a 1 parameter IRT model (Rasch model). In this model the questions' difficulty parameters (b) are universal across DGPs and respondents. The formula for the probability that respondent i responds Yes to question j is thus

$$P(X = 1|\theta_{i,j}, b_j) = \frac{e^{\theta_{i,j}-b_j}}{1 + e^{\theta_{i,j}-b_j}}$$

.

To convert the θ parameters for each respondent on each question to responses, we (1) first calculate the probability the respondent would answer Yes to the question using a Rasch model and then (2) to represent measurement error in the observed responses, record the answer as Yes with that probability.

This is implemented as follows:

```
# draw question difficulty parameters
bs <- rnorm(n.questions.total)

# function to implement Rasch model
convert.column.to.binary <- function(theta.vec, b) {
  as.numeric(
    exp(theta.vec - b) / (1 + exp(theta.vec - b)) > runif(length(theta.vec))
  )
}
```

We apply this function to all the DGPs above (except the Inattentives DGP, where response probabilities are simply 0.5). To increase realism, half the issues are reverse coded as well.

```
reverse.coded.issue.cols <- sample(1:n.questions.total, n.questions.total / 2) + 1 +
 ↪  n.questions.total

implement.rasch.model <- function(df) {
  df.names <- names(df)

  # implement rasch model
  for (i in 2:ncol(df)) {
    df[, paste0("issueopinion_binary_", df.names[i])] <-
```

```
    convert.column.to.binary(df[, i], bs[i - 1]) # first column of data is the
     ↪ true.type label
  }

  # reverse code some items for realism
  for (i in reverse.coded.issue.cols) df[, i] <- 1 - df[, i]

  return(df)
}

df.list <- list(
  data.non.ideologues,
  data.some.constraint,
  data.ideologues
)
result.list <- lapply(df.list, implement.rasch.model)
```

B.2.3 Simulation Results

Finally, we generate the results shown in Table 1 in the main text.

```
# Add moderates estimates
sim1.results <- do.call(rbind.data.frame, result.list) %>%
  add.moderates.estimates()

# Helper functions to print results
is.in.middle.tercile <- function(x) {
  qtiles <- quantile(x, probs = c(1 / 3, 2 / 3), na.rm = T)
  return(x >= qtiles[1] & x <= qtiles[2])
}

calculate.moderate.issue.share <- function(df) {
  issue.dvs <- names(df)[str_starts(names(df), "latentissueopinion_domain")]
  for (idv in issue.dvs) df[, paste0(idv, "_moderate")] <- is.in.middle.tercile(df[,
   ↪ idv])
  share.issue.views.moderate <- rowMeans(df[, names(df)[str_ends(names(df),
 ↪ "_moderate")]], na.rm = T)
  return(share.issue.views.moderate)
}

print.results <- function(df, caption = "") {
  df$share.issue.views.moderate <- calculate.moderate.issue.share(df)

  df %>%
    group_by(true.type) %>%
    dplyr::summarize(
      `Share of True Latent Preferences Which Are Moderate` =
       ↪ mean(share.issue.views.moderate),
      Downsian = mean(downsian),
      `Downsian 'Moderate'` = mean(downsian.moderate, na.rm = T),
      Conversian = mean(conversian),
      Inattentive = mean(inattentive)
    ) %>%
```

```r
    dplyr::rename(`True Type` = true.type) %>%
    mutate_at(vars(!`True Type`), scales::percent, accuracy = 1) %>%
    kable(caption = caption, position = "H", format = "latex", booktabs = T) %>%
    add_header_above(header = c(
      "Ground Truth in DGP" = 2,
      "Moderates Paper: Share Categorized As..." = 4
    )) %>%
    column_spec(1, bold = T) %>%
    column_spec(2, width = "7em") %>%
    column_spec(4, width = "5em")
}

table1 <- sim1.results %>%
  print.results(caption = "Simulation 1 Results")
```

B.3 Simulation 2

As noted in the main text, in the Online Appendix we also conduct a second simulation which adds several additional DGPs.

B.3.1 Additional DGPs

We add the following new DGPs in this simulation.

B.3.1.1 DGP 4: Left Pure Partisans

Left pure partisans have a θ of 2 on every item.

```r
data.leftpartisans <- data.frame(true.type = rep("Left Pure Partisan", n.respondents))
for (j in 1:n.issue.domains) {
  for (k in 1:n.issue.questions.per.domain) {
    data.leftpartisans[, paste0("latentissueopinion_domain", j, ".issue", k)] <- 2
  }
}
```

B.3.1.2 DGP 5: Right Pure Partisans

Right pure partisans have a θ of -2 on every item.

```r
data.rightpartisans <- data.frame(true.type = rep("Right Pure Partisan", n.respondents))
for (j in 1:n.issue.domains) {
  for (k in 1:n.issue.questions.per.domain) {
    data.rightpartisans[, paste0("latentissueopinion_domain", j, ".issue", k)] <- -2
  }
}
```

B.3.1.3 DGP 6: Libertarians

This DGP models a group of "off-dimensional" voters who all share a *common* system of correlated beliefs across questions. A motivating example case would be libertarians. We model them has share a common set of views in each of the 15 issue domains which are correlated within those domains.

```r
data.libertarians <- data.frame(true.type = rep("Libertarians", n.respondents))
shared.latent.issue.views <- rnorm(n.issue.domains, sd = 2.5)
for (j in 1:n.issue.domains) {
  for (k in 1:n.issue.questions.per.domain) {
    data.libertarians[, paste0("latentissueopinion_domain", j, ".issue", k)] <-
 ↪  shared.latent.issue.views[j] # All respondents share the same latent views in issue
 ↪  domain j
```

```
  }
}
```

B.3.1.4 DPG 7: Inattentives Finally, we also create a DGP for inattentive respondents. Their response probabilities for every question are 0.5; they are unaffected by the difficulty parameter of the question.

```
data.inattentive <- data.frame(true.type = rep("Inattentive", n.respondents))
for (j in 1:n.issue.domains) {
  for (k in 1:n.issue.questions.per.domain) {
    data.inattentive[, paste0("issueopinion_binary_latentissueopinion_domain", j,
↪   ".issue", k)] <- rbinom(n.respondents, 1, .5)
  }
}
```

B.3.2 Results

The results are as follows:

```
df.list <- list(
  # Previously used DGPs
  data.non.ideologues,
  data.some.constraint,
  data.ideologues,
  # New DGPs (except inattentives, added below)
  data.leftpartisans,
  data.libertarians,
  data.rightpartisans
)
result.list <- lapply(df.list, implement.rasch.model)

sim2.results <- do.call(rbind.data.frame, result.list) %>%
  plyr::rbind.fill(data.inattentive) %>% # Inattentive type does not have Rasch model
↪  applied as probability for all questions is 0.5.
  add.moderates.estimates()

sim2.results %>%
  print.results(caption = "Simulation 2 Results")
```

Ground Truth in DGP		Moderates Paper: Share Categorized As...			
True Type	Share of True Latent Preferences Which Are Moderate	Downsian	Downsian 'Moderate'	Conversian	Inattentive
Inattentive	NA	2%	2%	0%	98%
Left Pure Partisan	13%	100%	0%	0%	0%
Libertarians	67%	0%	0%	100%	0%
No Ideological Constraint	52%	97%	81%	1%	2%
Pure Ideologues	52%	100%	35%	0%	0%
Right Pure Partisan	20%	100%	0%	0%	0%
Some Ideological Constraint	52%	99%	48%	0%	1%

Table OA1: Simulation 2 Results

B.3.3 Who Are Conversians?

By plotting the estimated first and second dimenions from Simulation 2 as estimated using a 2D IRT model, it can be seen visually that Conversians form a single cluster of "those who are far from the first dimension in a similar way." (Note that the 2D IRT model from the mirt package is used strictly for the purposes of visualization and does not affect the estimates the "Moderates" model produces.)

```r
# 2d irt estimates from mirt package
irt2d.mod <- sim2.results %>%
  select(starts_with("issueopinion_binary_")) %>%
  mirt(data = ., model = 2)
scores <- fscores(irt2d.mod)

sim2.results$d1 <- scores[, 1]
sim2.results$d2 <- scores[, 2]
```

```r
sim2.results %>%
  sample_frac(.1) %>%
  ggplot(aes(x = d1, y = d2, color = modpaper_category)) +
  xlab("First Dimension") +
  ylab("Second Dimension") +
  scale_color_brewer('"Moderates" Model Output', palette = "Set1") +
  geom_point() +
  theme_bw()
```

The below graph splits this by voters' true underlying types, shown on the top of the facets:

```r
sim2.results %>%
  mutate(true.type = str_wrap(true.type, 15)) %>%
  sample_frac(.1) %>%
  ggplot(aes(x = d1, y = d2, color = modpaper_category)) +
  geom_point(alpha = .6) +
  xlab("First Dimension") +
  ylab("Second Dimension") +
  scale_color_brewer('"Moderates" Model Output', palette = "Set1") +
  facet_wrap(~true.type) +
  theme_bw()
```

Figure OA3: Moderates Paper's Model When Applied to Simulation 2, Visualized

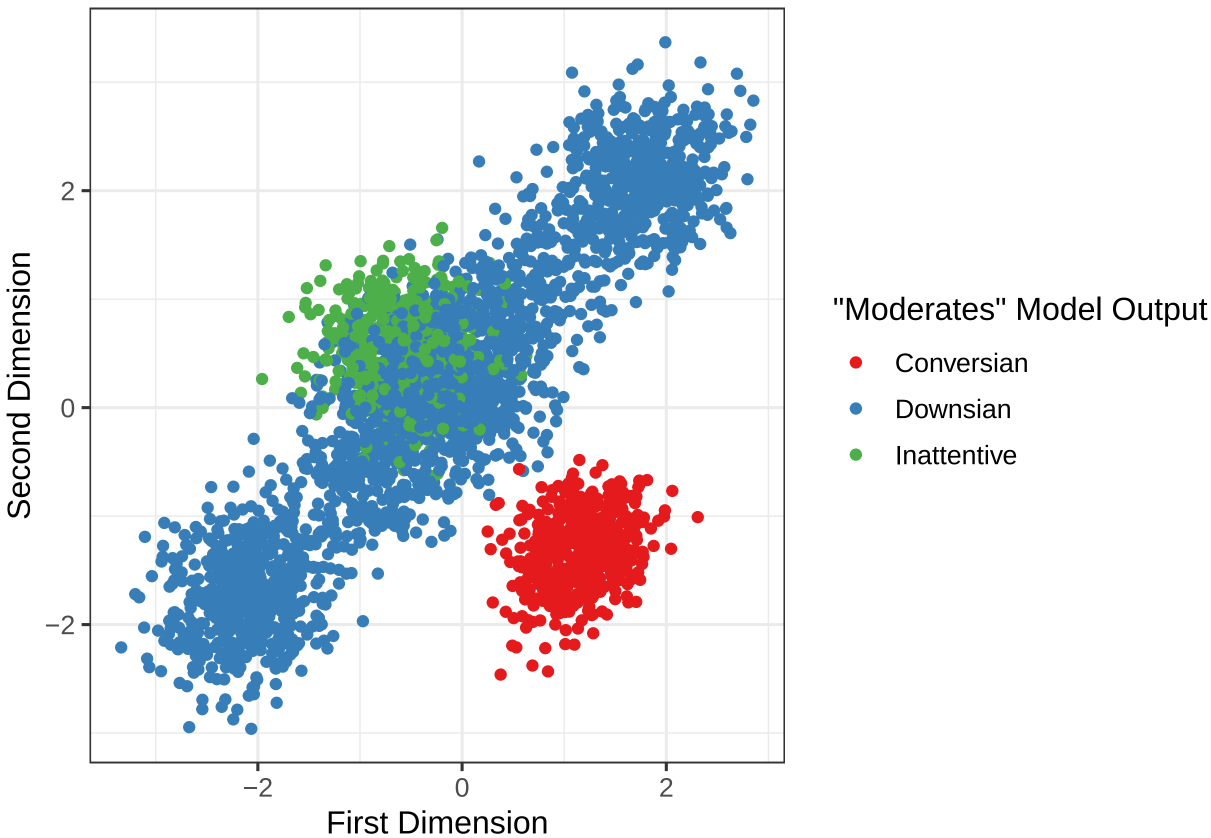

Figure OA4: Moderates Paper's Model When Applied to Simulation 2, Visualized, Split by True Type

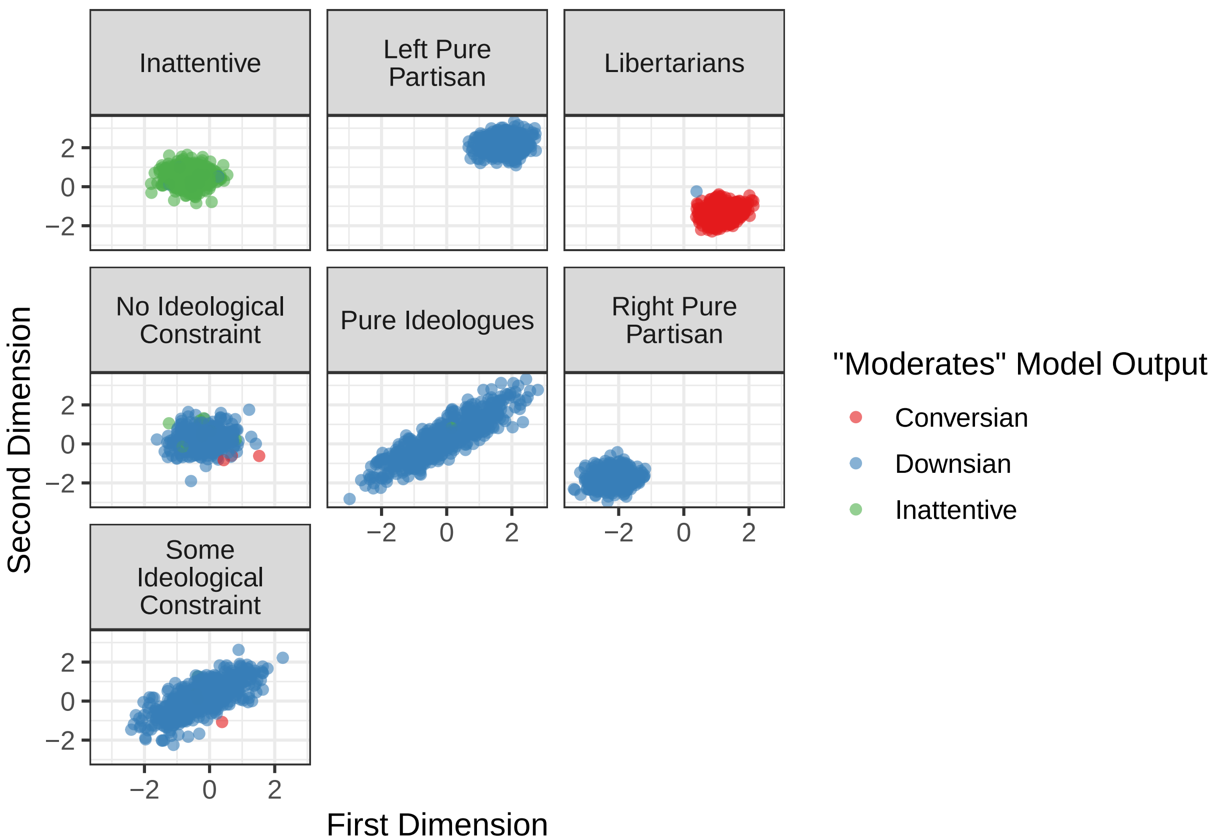

B.4 Simulation 3

B.4.1 DGP 8: Non-1D Constrained by Views Towards Social Groups

Our third simulation pursues a different approach than the first two. In this simulation, the source of the correlations between non-ideological respondents' views across issues is *not* that each question is a function of one of many possible uncorrelated issue domains (as in Simulations 1 and 2). Instead, drawing on Converse (1964), in this DGP there are instead 25 social groups towards which respondents may have attitudes. For each of the 25 social groups, there are a random 1000 of the 5000 respondents whose views toward the social group affect their attitudes on a random subset of 20 of the issue questions. For example, suppose one of the social groups is unions, that most people don't think about unions when they answer survey questions, but for a minority of respondents, their attitudes about unions inform their views on some of the survey questions (and there are 24 additional such groups). This produces correlations across questions (i.e., constraint across questions) for reasons unrelated to one-dimensional ideology. This allows us to test whether the bias in the "Moderates" model we discuss only arises when there are multiple groups of multiple issue questions which are each are determined by views in a single latent issue dimension within each group of questions (as in the prior simulation), or whether it also arises when a single question can be affected by multiple latent dimensions in public attitudes (e.g., towards multiple different groups). I.e., in this DGP, issue questions are not nested within single policy domains but rather affected by multiple different latent attitudes.

We simulate this DGP as follows:

```
data.non.1d.constraint <- data.frame(true.type = rep("Non-1D Constraint", n.respondents))

latent.issue.question.views <- matrix(nrow = n.respondents, ncol = n.issue.questions)
for (j in 1:n.issue.questions) latent.issue.question.views[, j] <- 0

for (i in 1:n.social.groups) {
  respondents.in.group <- sample(1:n.respondents, respondents.per.group)
  relevant.issues <- sample(1:n.issue.questions, n.relevant.issues.per.group)
  respondent.level.views <- rnorm(respondents.per.group)
  for (j in relevant.issues) {
    latent.issue.question.views[respondents.in.group, j] <-
↪   latent.issue.question.views[respondents.in.group, j] + respondent.level.views
  }
}

for (j in 1:n.issue.questions) data.non.1d.constraint[, paste0("latentissueopinion_", j)]
↪  <- latent.issue.question.views[, j] / sd(latent.issue.question.views[, j]) #
↪  standardize to SD 1
```

We also again simulate a DGP of 5000 pure ideologues.

```
data.ideologues <- data.frame(true.type = rep("Pure Ideologues", n.respondents))
ideology <- rnorm(n.respondents)
for (j in 1:n.issue.questions) {
  data.ideologues[, paste0("latentissueopinion_", j)] <- ideology
}
```

We perform the same mapping of latent attitudes to binary survey questions as employed in the previous simulations.

```
# draw question difficulty parameters
bs <- rnorm(n.issue.questions)

# function to implement Rasch model
convert.column.to.binary <- function(theta.vec, b) {
  as.numeric(
    exp(theta.vec - b) / (1 + exp(theta.vec - b)) > runif(length(theta.vec))
```

```
  )
}
```

```
reverse.coded.issue.cols <- sample(1:n.issue.questions, n.issue.questions / 2) + 1 +
 ↪  n.issue.questions

implement.rasch.model <- function(df) {
  df.names <- names(df)

  # implement rasch model
  for (i in 2:ncol(df)) {
    df[, paste0("issueopinion_binary_", df.names[i])] <-
      convert.column.to.binary(df[, i], bs[i - 1]) # first column of data is the
       ↪  true.type label
  }

  # reverse code some items for realism
  for (i in reverse.coded.issue.cols) df[, i] <- 1 - df[, i]

  return(df)
}

df.list <- list(
  data.non.1d.constraint,
  data.ideologues
)
result.list <- lapply(df.list, implement.rasch.model)
```

B.4.2 Simulation Results

The results are similar to those in the previous simulations. Individuals in the DGP which have constraint in their attitudes for reasons unrelated to 1D ideology (i.e., multidimensional constraint driven by attitudes towards groups, as Converse envisioned) are again overwhelmingly miscategorized as one-dimensional Downsians despite that they are not one-dimensional ideologues. This is driven by the fact that the “Moderates” model does not allow individuals’ views within the Conversian category to be correlated across issues.

Ground Truth in DGP	Moderates Paper: Share Categorized As...		
True Type	Downsian	Conversian	Inattentive
Non-1D Constraint	90%	10%	0%
Pure Ideologues	98%	2%	0%

Table OA2: Simulation 3 Results

C Regeneration

To further demonstrate that Fowler et al.'s (2023) model does not capture what it intends to, in data that more closely approximates real data, we next "regenerate" the data that Fowler et al. (2023) themselves analyze, holding constant the number of respondents, issues, and pattern of missing data exactly, and the average level of support for each issue and the pairwise correlations between issue positions probabilistically. We do this by estimating a latent multivariate normal model for their binary issue position data, which then enables us to generate new datasets with the same pairwise issue correlation structure as the observed data. Crucially, all respondents are now of a single type: their latent responses are drawn from a common multivariate normal model rather than having Downsian/Conversian types, and this distribution is not defined by a one-dimensional ideology. There are no "Conversians" at all in this data generating process. To the extent that a one dimensional model can approximate the ways that individuals positions on different issues tend to go together, it can do so equally well for all respondents.

We run this regeneration procedure five times on each of the datasets Fowler et al. (2023) analyze, and each time apply their model to the results. Figure OA4 plots the share of respondents who are categorized as "Downsian" in each of the regenerated datasets (y-axis), and compares the results to the the share categorized as "Downsians" in the original data (x-axis). Fowler et al.'s (2023) estimator categorizes nearly the same proportion of respondents as ideologues ("Downsians") in the regenerated datasets as when applied to each original dataset.

Figure OA5: Results When Regenerating Fowler et al. (2023)'s Data

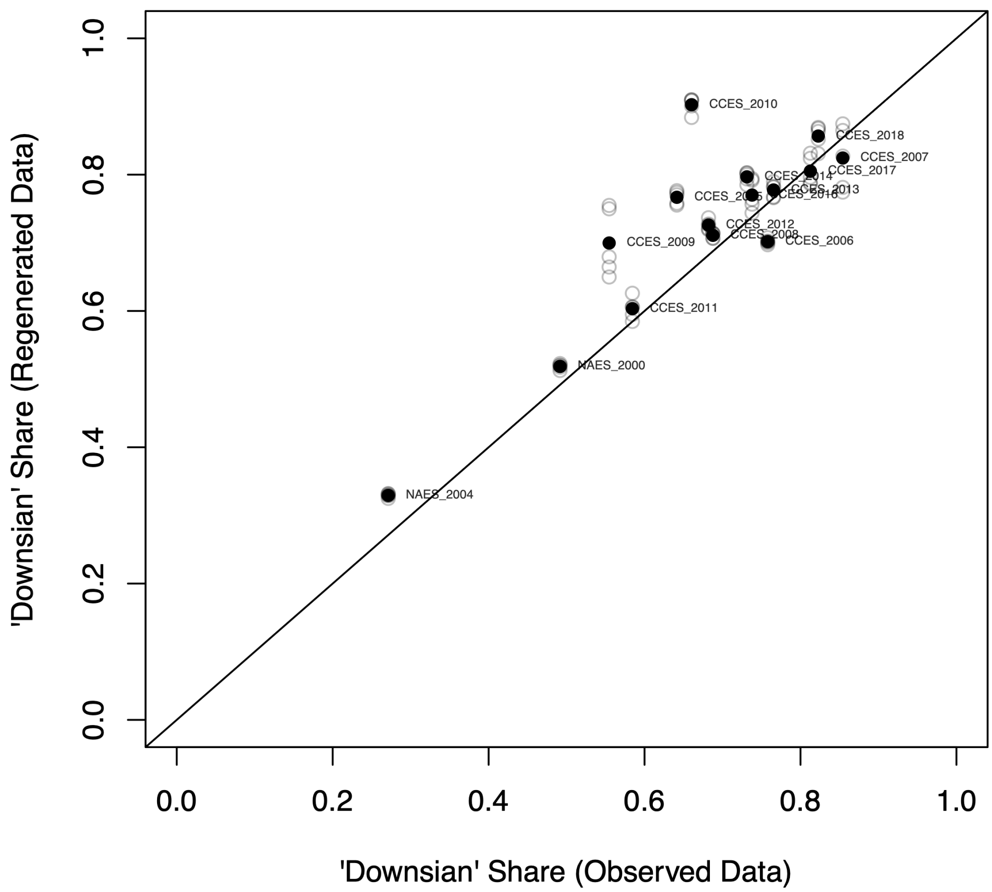

That we replicate the estimated proportions of ideologues ("Downsians") from datasets generated without different categories of respondents implies that the proportion of ideologues ("Downsians") the authors estimated in their original datasets cannot reliably indicate how many respondents truly belong to this category. Instead of depending on the proportion of ideologues and non-ideologues in the population, the proportion of each type their model estimates instead appears to depend largely on the pairwise correlation structure of issue positions in the dataset, which was preserved in these simulations.

The below code shows how we created Figure OA4.

```
###
### Function to generate alternative roll call matrix
###

# matching on
#     + pattern of missingness
#     + pairwise item correlations
#     + item response frequencies
# but using latent multivariate normal response model

max_det_cor_impute <- function(Sigma) {
  n_vars <- ncol(Sigma)

  # identify the number of missing elements of Sigma
  n_missing <- sum(is.na(Sigma)) / 2

  # calculate negative determinant, based on filled in imputes
  detf <- function(imputes) {
    Sigma_Imputed <- Sigma
    i <- 1
    for (j in 1:(n_vars - 1)) {
      for (jj in (j + 1):n_vars) {
        if (is.na(Sigma[j, jj])) {
          Sigma_Imputed[j, jj] <- Sigma_Imputed[jj, j] <- imputes[i]
          i <- i + 1
        }
      }
    }
    return(-det(Sigma_Imputed))
  }

  # use optim to find maximum determinant solution
  start_tmp <- mean(as.vector(Sigma)[as.vector(Sigma) != 1], na.rm = TRUE)
  start_values <- rep(start_tmp, n_missing)
  optim_out <- optim(start_values, detf,
    method = "L-BFGS-B",
    lower = -1, upper = 1
  )

  # populate Sigma_MD with Max Det solution

  Sigma_MD <- Sigma
  i <- 1
  for (j in 1:(n_vars - 1)) {
    for (jj in (j + 1):n_vars) {
      if (is.na(Sigma[j, jj])) {
        Sigma_MD[j, jj] <- Sigma_MD[jj, j] <- optim_out$par[i]
        i <- i + 1
      }
    }
```

```r
  }

  # if not positive definite, truncate negative eigenvalues and rescale to valid correlation matrix

  n <- dim(var(Sigma_MD))[1L]
  E <- eigen(Sigma_MD)
  Sigma_CM1 <- E$vectors %*% tcrossprod(diag(pmax(E$values, 0), n), E$vectors)
  Balance <- diag(1 / sqrt(diag(Sigma_CM1)))
  Sigma_CM2 <- Balance %*% Sigma_CM1 %*% Balance

  # return valid correlation matrix

  return(Sigma_CM2)
}

polychor_pairwise <- function(mat) {
  # calculate observed polychoric correlations for items using observed data
  n_vars <- ncol(mat)
  Sigma <- matrix(NA, n_vars, n_vars)
  # populate diagonal
  for (j in 1:n_vars) Sigma[j, j] <- 1
  # calculate pairwise polychoric correlations
  for (j in 1:(n_vars - 1)) {
    for (jj in (j + 1):n_vars) {
      if (j != jj) {
        # if pair of items appears, set correlation to observed
        tmp <- suppressWarnings(polychor(mat[, j], mat[, jj]))
        if (!is.na(tmp)) Sigma[j, jj] <- Sigma[jj, j] <- tmp
      }
    }
  }

  return(Sigma)
}

mvn_rc_generator <- function(mat) {
  n_vars <- ncol(mat)

  # calculate observed polychoric correlations for observed pairs
  Sigma <- polychor_pairwise(mat)

  # impute unobserved pairwise correlations to maximise correlation matrix determinant
  # see https://royalsocietypublishing.org/doi/full/10.1098/rsos.172348#FN2R for justification
  Sigma_imputed <- max_det_cor_impute(Sigma)

  # generate multivariate normal random latent votes
  y_star <- mvrnorm(nrow(mat), rep(0, n_vars), Sigma_imputed)

  # dichotomise latent votes into observed votes
  y <- matrix(NA, nrow(mat), ncol(mat))
  for (j in 1:n_vars) y[, j] <- y_star[, j] > qnorm(prop.table(table(mat[, j]))[1])
```

```r
  # match missingness of original data set
  mat_alt <- y * (mat >= 0)

  # return new roll-call matrix
  return(mat_alt)
}



###
### load all large survey data
###

# load("bigsurveys_recoded4.RData")
data <- read.dta("bigsurveys_recoded4.dta")

policy_ind <- 3:329
years <- unique(data$source)

table(years)

###
### Remove modules from datasets
###

for (year in years) {
  tmp <- data[, policy_ind]
  tmp <- tmp[data$source == year, ]
  allna <- is.na(colMeans(tmp, na.rm = T))
  tmp <- tmp[, !allna]

  # we need fewer responses than this
  # to determine that an item was asked
  # on a module
  threshold <- 5000

  module_items <- c()
  for (i in 1:dim(tmp)[2]) {
    vec <- ifelse(tmp[, i] == 0, 1, tmp[, i])
    if (sum(vec, na.rm = T) < threshold) {
      module_items <- c(module_items, i)
    }
  }
  if (length(module_items) > 1) {
    module_people <- which(!is.na(rowMeans(tmp[, module_items], na.rm = T)))
  } else {
    module_people <- c()
  }


  print(year)
  print(c("Total people, items:", dim(tmp)))
```

```r
  print(c("Module people:", length(module_people)))
  print(c("Module items:", length(module_items)))

  # now that module items have been identified,
  # NA those items for those modules
  if (length(module_items) > 0) {
    ind <- which(names(data) %in% names(tmp)[module_items])
    data[data$source == year, ind] <- NA
  }
}


###
### run em_mix_irt on every included data set, and append estimates
###

# store all of the results in a list
reslist <- list()
reslist_alt <- list()
# store ideal points
data$x <- rep(NA, dim(data)[1])
# store probability of each voter type
data$w1 <- rep(NA, dim(data)[1])
data$w2 <- rep(NA, dim(data)[1])
data$w3 <- rep(NA, dim(data)[1])
# store individual log likelihoods
data$lk1 <- rep(NA, dim(data)[1])
data$lk2 <- rep(NA, dim(data)[1])


for (i in 1:length(years)) {
  year <- years[i]

  print(year)
  mat <- apply(data[data$source == year, policy_ind], 2, as.numeric)
  exclude <- !colnames(mat) %in% c("cces2012_immigration5", "cces2012_immigration6")
  notna <- which(!is.na(colMeans(mat, na.rm = T)) & exclude)
  print(length(notna))

  mat <- mat[, notna]
  res <- em_mix_irt(mat, iter = 100)
  reslist[[year]] <- res
  data$x[data$source == year] <- res$irt$x
  data$w1[data$source == year] <- res$w[, 1]
  data$w2[data$source == year] <- res$w[, 2]
  data$w3[data$source == year] <- res$w[, 3]
  data$lk1[data$source == year] <- res$irt$lk
  data$lk2[data$source == year] <- res$ivp$lk

  # analysis on re-generated roll-call matrix

  reslist_alt[[year]] <- list()
  reslist_alt[[year]]$observed_pairwise_polychor <- polychor_pairwise(mat)
```

```
  reslist_alt[[year]]$replicate_weight_estimates <- list()
  for (r in 1:5) {
    mat_alt <- mvn_rc_generator(mat)
    res_alt <- em_mix_irt(mat_alt, iter = 100)
    reslist_alt[[year]]$replicate_weight_estimates[[r]] <- colMeans(res_alt$w)
  }
}

###
### Identify the direction using pid7
###

data$pid7[data$pid7 == "__NA__"] <- NA
data$pid7 <- as.numeric(data$pid7)
sources <- unique(data$source)

for (i in sources) {
  ind <- data$source == i
  direction <- sign(cor(data$pid7[ind], data$x[ind], use = "pairwise.complete.obs"))
  data$x[ind] <- direction * data$x[ind]
}

for (i in sources) {
  ind <- data$source == i
  print(i)
  print(cor(data$pid7[ind], data$x[ind], use = "pairwise.complete.obs"))
}

###
### Save the results
###

save(reslist, reslist_alt, file = "all_results4_alt.RData")

save(data, file = "bigsurveys_recoded_plus_estimates4.RData")

write.dta(data, file = "bigsurveys_recoded4.dta")



###
### Comparison of Estimates on observed vs regenerated Data
###

if (FALSE) {
  average_weight_observed <- matrix(NA, length(reslist), 3)
  average_weight_regenerated <- matrix(NA, length(reslist), 3)

  for (i in 1:length(reslist)) {
    average_weight_observed[i, ] <- colMeans(reslist[[i]]$w)
    average_weight_regenerated[i, ] <- colMeans(reslist_alt[[i]]$w)
  }
```

```r
  plot(average_weight_observed[, 1], average_weight_regenerated[, 1],
    xlim = c(0, 1), ylim = c(0, 1), xlab = "'Downsian' Share (Observed Data)", ylab =
      ↪  "'Downsian' Share (Regenerated Data)"
  )
  text(average_weight_observed[, 1], average_weight_regenerated[, 1], names(reslist), pos
   ↪  = 4, cex = 0.4, col = rgb(0, 0, 0, 0.6))
  abline(0, 1)
}
```

References

Converse, Philip E. 1964. The nature of belief systems in mass publics. In *Ideology and its discontents*, ed. David Apter. New York: Glencoe Free Press.

Fowler, Anthony, Seth J Hill, Jeffrey B Lewis, Chris Tausanovitch, Lynn Vavreck and Christopher Warshaw. 2023. “Moderates.” *American Political Science Review* 117(2):643–660.