Political Science Research and Methods Vol 6, No. 3, 549–569 July 2018
 doi:10.1017/psrm.2016.50

Don't Know What You Got: A Bayesian Hierarchical Model of Neuroticism and Nonresponse*

JONATHAN D. KLINGLER, GARY E. HOLLIBAUGH JR. AND ADAM J. RAMEY

Individuals who are more sensitive to negative outcomes from error are more likely to provide nonresponses in surveys. We argue Neurotics' sensitivity to negative outcomes leads them to avoid gathering costly information and forming/reporting opinions about stimuli. Using data from the 2014 Cooperative Congressional Election Study, we show Neuroticism is strongly and positively associated with NA/DK responses when placing politicians on a seven-point ideological scale. We then introduce to political science a Bayesian hierarchical model that allows nonresponse to be generated by both a lack of information as well as disincentives for response. Using this model, we show that the NA/DK responses in these data are due to inhibited information collection and indecision from error avoidance by Neurotics.

Individuals may fail to respond to political questions for many reasons, possibly including the influence of one's personality traits.[1] Personality traits are persistent individual differences, and psychologists have developed models to capture underlying structure in those differences. The five-factor model of personality has gained support from psychologists, and political scientists have been incorporating the Big Five personality traits identified by this model into the study of political behavior and institutions (Mondak and Halperin 2008; Gerber et al. 2010; Mondak et al. 2010; Dietrich et al. 2012). Neuroticism is associated with instability, suggesting it is of particular importance for information processing (Robinson and Tamir 2005; Flehmig et al. 2007; Mondak et al. 2010), and it has been associated with decreased political knowledge (Gerber et al. 2011). However, there are several mechanisms by which it may express itself through nonresponse on items requiring political knowledge—namely, inhibited data collection and indecisiveness from lower expected utilities of response. Neurotics may inhibit their exposure to contentious information, or they may be less likely to form opinions due to "mental noise" or, as we argue, error sensitivity (Robinson and Tamir 2005; Mondak and Halperin 2008; Gerber et al. 2011). However, little work has been done to understand how Neuroticism leads individuals to be less likely to respond to political questions, a question we investigate here.[2]

* Jonathan D. Klingler is a Postdoctoral Scholar in the Department of Political Science, Vanderbilt University, PMB 505, 230 Appleton Place, Nashville, TN 37203 (jonathan.d.klingler@vanderbilt.edu). Gary E. Hollibaugh, Jr. is an Assistant Professor in the Department of Political Science, University of Notre Dame, Notre Dame, IN 46556 (gholliba@nd.edu). Adam J. Ramey is an Assistant Professor in the Department of Politics, New York University Abu Dhabi, PO Box 129188, Abu Dhabi (adam.ramey@nyu.edu). Author order was determined by a spirited game of Diplomacy. All authors contributed equally to the paper. The authors thank Megan Remmel for feedback and comments. Support through ANR–Labex IAST and the University of Notre Dame Center for Research Computing is gratefully acknowledged. To view supplementary material for this article, please visit https://doi.org/10.1017/psrm.2016.50

[1] Evidence suggests personality traits are stable and causally prior to attitudes and behaviors, which draws us to them even though we acknowledge that there are many other psychological variables of importance (McCrae and Costa 1996; Roberts and DelVecchio 2000; Fraley and Roberts 2005).

[2] One notable exception is Jessee (2015), who argues personality traits are unrelated to why individuals choose "don't know" responses on surveys. Conversely, Ramey, Klingler and Hollibaugh (2016) showed that

However, doing so requires modeling the underlying nonresponse decision, as both inhibition and indecision result in responses that will be coded as NA/DK, rendering them observationally equivalent to the consumer of the resulting data, even though these responses may arise from different underlying psychological processes. Treating such responses identically has the potential to affect inferences due to underlying group-level heterogeneity. This is endemic to survey research, as many surveys give individuals the opportunity to skip questions (NA) or elicit "don't know" (DK) responses. Though deleting these observations is common, doing so can lead to biased inferences. Fortunately, Bradlow and Zaslavsky (1999) provide an approach that models missingness using a hierarchical, multiple latent variable approach. It considers individuals' responses as a product of three variables: *saliency*, *opinion*, and *decisiveness*.

Remarkably, this approach mirrors our model of decisionmaking as a function of Neuroticism. We expand on this approach and merge it with the insights of Aldrich and McKelvey's (1977) approach to modeling ideological placements of elites by survey respondents. This model enables the use of surveys with missing data to estimate the ideological placements of elites in a common space, and to examine the underlying psychological processes that result in NA/DK responses, which would not be possible with conventional methods of dealing with missingness.[3]

This paper proceeds as follows. We first discuss the literature on missing data, then the five-factor model of personality, focusing on Neuroticism, and its connection with political information. We then expand upon the core cognitive constraint framework proposed by Ramey, Klingler and Hollibaugh (2017) and articulate two mechanisms—inhibition and indecision—this framework implies for Neuroticism.[4] We then discuss our Bayesian hierarchical model. Subsequently, we investigate the relationship between Neuroticism and NA/DK responses, modeling NA/DK responses as a function of both the ability to collect adequate information to form an opinion on an item (salience) as well as sensitivity to potential error disutility from reporting clear opinions (decisiveness). The results suggest NA/DK responses are a feature of Neuroticism's inhibited information gathering as well as indecisiveness due to error sensitivity. We then discuss the implications for studying political information as well as characterizing personality traits for modeling purposes.

MISSING DATA

Missing data have long drawn the ire of social science researchers. In the case of surveys, item nonresponse can cause serious problems for multivariate analysis. Moreover, traditional remedies like listwise deletion, pairwise deletion, mean-insertion, or dummy variable adjustment have been shown to cause serious bias in estimates and/or inferences (King et al. 2001). As a result, a large literature (e.g., Heckman 1976; Rubin 1987; Schafer 1997; Gelman, King and Liu 1998; Berinsky 1999; King et al. 2001) has emerged that models missingness in ways that minimize the bias traditional remedies might induce.

This literature can be divided into two loosely defined classes. The first is that of multiple imputation models (e.g., Rubin 1987; Little and Rubin 1989; Schafer 1997; Gelman, King and Liu 1998). This paradigm seeks to "impute" the missing values by using other observed

(F'note continued)
personality traits are related to the initial decision to respond to surveys, though they do not investigate the relationship between personality traits and the eventual response.

[3] Even modifications of the Aldrich-McKelvey (AM) method that allow for missing data (e.g., Hare et al. 2015) assume the missingness is random, which we argue is an untenable assumption.

[4] For an alternative perspective, see Hall (2015).

information in the data. The method is easy to implement and available in most standard statistical packages. Though categorical variables present difficulty, imputation is about as close as one can get to a one-size-fits-all methodology.

However, this approach has limitations, two of which are of particular interest. First, the missingness must obey the so-called missing at random (MAR) assumption. Following King et al. (2001), let D_{obs} denote observed data, D_{mis} denote missing data, D denote the total data, and M represent missingness. Data are MAR if $\Pr(M|D_{obs}, D_{mis}) = \Pr(M|D_{obs})$. That is, data satisfy MAR if missingness can be modeled as a function of observed data. A canonical example of this sort of missingness is the case of high-wage earners who fail to report their income in surveys. While their income may be unobserved, several known correlates (e.g., education) are not missing. By conditioning on observed covariates, we may model missingness using existing algorithms. However, if the missing observations cannot be predicted from observed covariates, MAR is not satisfied and multiple imputation is not usable (Weisberg 2009). A second issue is the nature of the missingness itself. Indeed, when considering NA/DK responses, particularly those on opinion-oriented questions, is it the case that missing values are simply censorings or "accidents"? Perhaps it is the case that individuals who do not choose a response or elicit "don't know" may be actually making a choice in the same sense as the other categories provided. If this is true, King et al. concede that cases "… when 'no opinion' means that the respondent really has no opinion rather than prefers not to share information with the interviewer should be treated seriously and modeled directly …" (2001, 59).

Another class is deemed by King et al. (2001) as "application specific" (e.g., Heckman 1976; Bartels 1999; Berinsky 1999; Jessee 2015), and generally requires modeling of the missingness mechanism. Two such approaches (Heckman 1976; Berinsky 1999) consider data in terms of the selection model, where those choosing NA/DK select themselves out of the sample. While useful, this approach has limitations; these models require exclusion restrictions to ensure identification, and they restrict missingness to result from one choice. However, it is also plausible to think missingness results from several different factors, including insufficient information to form and report opinions, disincentives for response, or indifference due to uncertainty.[5]

Other literature on uncertainty and candidate evaluation suggests uncertainty manifests itself as response variance rather than nonresponse (e.g., Zaller and Feldman 1992; Alvarez and

[5] Some discussion of the role of Bayesian methods in missing data analysis is also warranted. The imputation posterior (IP) multiple imputation algorithm (Schafer 1997), itself a version of Tanner and Wong's (1987) data augmentation algorithm, provides a Bayesian way of calculating the exact posterior distribution of the missing values. Like all Bayesian methods, the IP algorithm is computationally intensive, and faster algorithms (e.g., King et al. 2001; Blackwell, Honaker and King 2015) have been developed. However, within the social sciences, Bayesian methods have largely been implemented within the second class of methods, the aforementioned "application-specific" ones. Indeed, nearly the entire class of Bayesian ideal point estimation models (e.g., Clinton, Jackman and Rivers 2004; Clinton et al. 2012; Tausanovitch and Warshaw 2013; Ramey 2015) are exemplars of this type. Treier and Hillygus note, "[w]ith a Bayesian approach, an individual's latent ideology scores are estimated with the data available for that individual, and those estimates are simply less precise for those with less data … In contrast, classical factor analysis would require a correction to the 'Swiss cheese' data structure, by either collapsing the different question formats, using listwise deletion, or in some way imputing data to fill in the holes" (2009, 685). This distinction can be seen in the work of Hare et al. (2015), who implement a Bayesian version of the Aldrich and McKelvey (1977) method, thus allowing it to use incomplete and missing observations. Though Hare et al. (2015) do not model the missingness mechanism, and instead assume MAR, other scholars have begun to push on this front. For example, Rosas, Shomer and Haptonstahl (2015) provide a framework for estimating ideal points using roll-call data where the missingness is modeled (as opposed to simply assuming MAR).

Franklin 1994; Glasgow and Alvarez 2000; Alvarez and Brehm 2002). However, Bartels (1986) considers nonresponse to be a result of uncertainty, and Alvarez and Franklin note “it is natural to treat these ‘don’t know’ respondents as more uncertain than those who place the [stimulus], but then say they are not very certain of the location” (1994, 680). Thus, we find it plausible that missingness may arise from either insufficient information to form and report an opinion, from low expected utilities of response among those who have opinions, or both.[6] We believe that focusing on personality traits—in particular, Neuroticism—and modeling the decision-making process can help us understand why respondents elicit NA/DK responses.[7]

NEUROTICISM, INHIBITION, AND INDECISION

The Big Five model of personality proposes five personality traits derived from factor analysis of questionnaires as well as descriptive language (Goldberg 1981; John 1990), and these traits—Openness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism (often reverse coded as Emotional Stability)—have achieved prominence and have been used to predict life outcomes ranging from romantic fulfillment to mortality, with predictive power comparable with socioeconomic status and cognitive ability (Roberts et al. 2007). Neuroticism is associated with anxiety, depression, impulsiveness, and vulnerability to stress (Almlund et al. 2011). Related traits include external locus of control, irritability, and a sense of vulnerability (John, Robins and Pervin 2008). Neurotics tend to have low self-esteem and are unstable, withdrawn, easily angered, and difficult to motivate.

There are few clear connections between Neuroticism and political phenomena, though one that has received attention is an association with ideological extremism (Soldz and Vaillant 1999). A second line of inquiry stems from the idea that Neurotics may have more uncertainty about their attitudes (Mondak et al. 2010). Others have examined the relationship between Neuroticism and political information. Mondak and Halperin (2008, 345) hypothesized Neurotics’ instability would lead them to be more opinionated (in contrast with calm and silent emotionally stable individuals) overall, while avoiding group-based activities (such as meetings) where conflict is possible and “social distress” might be induced. These predictions were borne out by the data, with Neurotics more likely to be (and be perceived as) opinionated and to be politically knowledgeable, but less likely to engage in the political process, presumably due to the possibility of “social distress.” Conversely, Gerber et al. (2011) posited political contentiousness would prevent Neurotics from becoming interested and knowledgeable about politics in the first place (and not merely less likely to participate despite greater knowledge), a contention supported by the data. These divergent findings as to the underlying cause of underparticipation among Neurotics indicates we have not yet been able to distinguish between insufficient political knowledge and/or revealed opinionation (both of which can manifest as survey nonresponse) arising from either inhibited information collection or from

[6] We assume individuals are sensitive to error, and guessing due to uninformative beliefs might to optimal nonresponse (Holroyd and Coles 2002). For those with uncertain but informative beliefs, the expected utility of response would be a function of the penalties (rewards) for “incorrect” (“correct”) responses and the probability of providing such a response. The probability of incorrect responses would presumably increase with belief variance.

[7] While Jessee (2015) examines the relationship between personality and “don’t know” responses, he does not explicitly model the decision-making process that leads to such responses, instead preferring to rely on a multinomial probit item response model. However, this approach, while useful, is inapplicable when the decision-making process is hierarchical and NA/DK responses may arise from different processes, which is what we argue here.

refusal to reveal opinions (even if opinionation and/or knowledge is higher), through Neuroticism has been connected to both mechanisms. Importantly, both of these bodies of work agree that political contentiousness should result in Neurotic individuals being less likely to participate in the political process, even in the presence of higher levels of opinionation (as Mondak and Halperin [2008] suggest). Thus, we spend the rest of this section articulating how Neuroticism's neurological roots should lead individuals to perceive greater penalties for error (and thus more likely to refuse to form and reveal opinions) and less likely to collect information.

The association between Neuroticism and error sensitivity has been linked to serotonin (Gray and McNaughton 2003), with a broader theory suggesting a biochemically induced fixation on negative outcomes (Gray and McNaughton 2003; DeYoung and Gray 2009). In the lab, Neurotics are prone to behavioral inhibition through passive avoidance and freezing, presumably due to this fixation (DeYoung and Gray 2009). If Neurotics are fixated on error and negative outcomes, the best way to avoid negative outcomes and stress would be to withdraw and maintain the status quo. Whether through sensitivity to error, stress avoidance, or a tendency to negative self-evaluation, Neuroticism can be modeled as a *sensitivity to and fixation on prospective negative outcomes*, as proposed by Ramey, Klingler and Hollibaugh (2017).

Negative outcomes in this context refer to losses relative to a neutral reference point, similar to the approach taken by prospect theory (Kahneman and Tversky 1979; Tversky and Kahneman 1992; Derryberry and Reed 1994). Fortunately, the process by which individuals choose to form and report opinions is decision theoretic, meaning we can consider the decision by an individual to form a clear opinion on the ideological position of a political actor. In line with the relative utility structure suggested by the neuropsychology literature on Neuroticism, we assume a status quo baseline with zero utility. It is trivial to state that any losses relative to the status quo are thus negative utilities, and any gains are accordingly positive.

We assume individuals, when presented with a stimulus, must either form and report a clear opinion, or avoid doing so. In the present context, we use the term *formation* to refer to both the decision to form an opinion *and* report that opinion for evaluation. If no opinion is formed, the status quo is maintained and the individual receives neither negative outcomes nor reward. If an opinion is formed, the individual is correct with probability $p \in (0, 1)$ and receives a reward, R, or is incorrect with probability $1-p$ and receives a negative outcome L. We assume this outcome is a negative emotional state arising from experiencing error, and the reward consists of positive feelings from being correct, as well as any other gains resulting from having accurate information.

We assume a two-dimensional type space for sensitivities to reward and negative outcomes. We assume an individual's sensitivity to negative outcomes is $x \in [1, \infty)$ and the sensitivity to reward is $y \in [1, \infty)$. The negative outcome, L, is weighted by x, and the reward, R, is weighted by y. We therefore have the following utilities for opinion formation and nonformation:

$$U_N = 0,$$

$$U_F = pRy - (1-p)Lx.$$

We define m to be $\frac{R}{L}$, or the ratio of the magnitude of the reward to the magnitude of the negative outcome. If we identify the conditions under which it is optimal to not form an opinion, and substitute in m appropriately, we see nonformation is weakly optimal when

$$x \geq \left(\frac{p}{1-p}\right) my. \tag{1}$$

As the importance of negative outcomes increases, nonformation is more likely to be optimal.

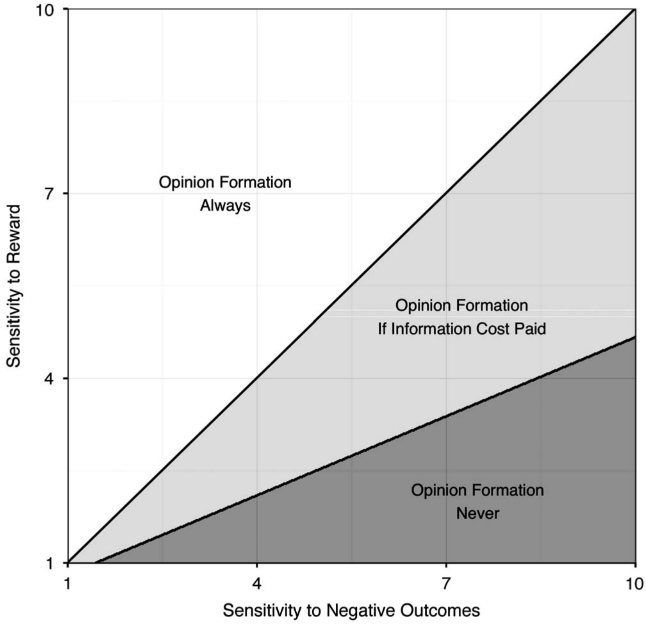

Fig. 1. Opinion formation as a function of sensitivities to reward and negative outcomes

Now, consider an extension where the player may pay a cost $c \in [0, \omega(Ry + Lx)]$ to collect additional information and increase the probability of a correct opinion from p to $p + \omega$, where $\omega \in (0, 1 - p]$.[8] The costs of information acquisition aside, the utility of nonformation is unaffected by p, while the utility of formation increases. However, we assume individuals will only pay the cost if doing so weakly increases their expected utility. Therefore, the ability to pay for information acquisition results in three cases. In the first case, where $x < \left(\frac{p}{1-p}\right)my$, opinion formation is always optimal, and there is no need to pay for more information before doing so. However, since doing so increases the probability of being correct, and therefore increases the potential rewards relative to negative outcomes, the individual will pay the cost. When $x \in \left[\left(\frac{p}{1-p}\right)my, \left(\frac{p+\omega}{1-p-\omega}\right)my - \frac{c}{L(1-p-\omega)}\right]$, opinion formation is optimal, conditional on paying for more information before doing so, as the individual is sufficiently sensitive to negative outcomes (but not so much that opinion formation is never optimal). Finally, opinion formation is never optimal when $x > \max\left\{\left(\frac{p}{1-p}\right)my, \left(\frac{p+\omega}{1-p-\omega}\right)my - \frac{c}{L(1-p-\omega)}\right\}$, because the individual is so sensitive to negative outcomes that not even the additional information that could be purchased will be sufficient (at least at the specified cost). Additionally, since there are limits to the information that may be gathered, increasing x will weakly increase the probability no opinion is formed and gathering no information becomes more optimal. Figure 1 presents an example of how the relative sensitivities to reward and negative outcomes affect the formation and acquisition decisions.[9]

We assume personality measures for the trait of Neuroticism capture its core cognitive constraint, which is a sensitivity to prospective negative outcomes; accordingly, Neuroticism is

[8] For tractability, assume both c and ω are exogenously determined. However, the substantive results are similar if the individual is allowed to pay varying costs c for varying amounts of information. Additionally, the upper bound on c ensures the cost is sufficiently reasonable relative to ω.

[9] Figure 1 was created using $R = 3$, $L = 3$, $p = 0.5$, $\omega = 0.2$, $x \in [1, 10]$, $y \in [1, 10]$, and $c = 1$.

https://doi.org/10.1017/psrm.2016.50 Published online by Cambridge University Press

parameterized as x in the above model. In the empirical model that follows, the relative utility of response is described as decisiveness.[10] As increases in x are associated with an decreased likelihood of choosing to form an opinion, we obtain the following hypothesis:

HYPOTHESIS 1: More Neurotic individuals should be less be decisive.

Next, collecting information on a particular stimulus can be described as finding that stimulus salient. In the decision described above, as x increases, it is more likely to be suboptimal for an individual to collect information on a stimulus. This generates a second hypothesis:

HYPOTHESIS 2: More Neurotic individuals should find the response stimuli less salient.

As Neurotic individuals collect less information and are less likely to choose to form and report opinions, we expect them to be more likely to not have clear evaluations of stimuli and therefore present NA/DK responses, suggesting the following hypothesis:[11]

HYPOTHESIS 3: More Neurotic individuals should provide more NA/DK responses.

Finally, as Neurotics will collect less information, they will be less informed about broad sets of stimuli, including the ideological scale itself, thus generating our final hypothesis:

HYPOTHESIS 4: More Neurotic individuals should incorrectly perceive the ideological scale.

A STATISTICAL MODEL OF THE DECISION-MAKING PROCESS

The most common approach to modeling ideological placements of elites by voters was pioneered by Aldrich and McKelvey (1977). The AM algorithm assumes an arbitrary individual i's placement of an elite stimulus j on an ordinal ideological scale is given by

$$y_{ij} = a_i + b_i x_j + \varepsilon_{ij}, \tag{2}$$

[10] This terminology used by Bradlow and Zaslavsky (1999) in the discussion of the hierarchical model is used here. This terminology is limited to this model and not the broader concept of decisiveness in the public opinion literature.

[11] Upon initial reading, Hypotheses 2 and 3 might seem incompatible with Zaller and Feldman's (1992) model of survey response. However, it is consistent with Zaller (1990), which discusses nonresponse to a far greater degree than in Zaller and Feldman (1992), which is explicitly concerned about response stability. From the 1990 paper:

> [Deduction] 1 claims that people will not, on average, be very critical in deciding which of the messages encountered they will accept. Some persons, however, may have been exposed to few or no persuasive arguments on some issues; or, because they rarely think about the arguments they have accepted, they may be unable (via A3) to call any considerations to mind in the short time they give themselves for answering survey questions. Such persons must presumably answer questions with "no opinion." Because people who are low on political awareness would tend to think less about politics than other persons would, they should be less likely to offer opinion statements or, conversely, more likely to offer no opinion responses (D2). Krosnick and Milburn [1990] review the considerable evidence supporting this deduction (Zaller 1990, 129).

Consider a hypothetical person who is so Neurotic that she has failed to gather significant information. In this case, when confronted with a survey item, the number of considerations available for recall would be very low and thus she would be more likely to offer a NA/DK (or in Zaller's terms, "no opinion") response.

where a_i and b_i are individual distortion parameters and x_j the latent ideological locations of the elite stimuli.[12] The distortion parameters capture the idea that individuals may perceive the underlying ideological space differently. Aldrich and McKelvey (1977) estimate this model using a singular value decomposition and demonstrate it accurately recovers the locations of stimuli as well as the information possessed by the survey respondents.

Unfortunately, this procedure cannot handle missing values and removes individuals who fail to place even just one stimulus. Addressing this shortcoming, Hare et al. (2015) develop a Bayesian version that can incorporate missing values.[13] Their approach assumes missing placements are MAR and are drawn from the assumed distribution of the placements.

However, this assumption is generally problematic and particularly so with respect to Neuroticism. Specifically, we do not believe missing placements are MAR. Instead, we believe missing placements are influenced by individuals' personality traits—namely their varying degrees of Neuroticism. If this is the case, then we should actively model the decision-making process, whereby individuals decide to answer (or not to answer) elite placement questions.

Our approach does just this. We expand upon a model developed by Bradlow and Zaslavsky (1999), which assumes NA/DK responses are driven by latent psychological processes. We then merge this approach with the insights of the AM algorithm. While our interests are in the effects of Neuroticism, the framework we develop can be applied to any similar decision-making setup.

To begin, let $i = 1, 2, \ldots, N$ denote the set of respondents to a survey and let $j = 1, 2, \ldots, J$ denote the a set of stimuli that they are asked to rate on an ordinal scale. For each individual i, y_{ij} is his ordinal response to item j. Typically, these sorts of items involve five- or seven-point scales. If i skipped the item, his response is coded as either NA or DK. In most analyses of these data, ordered probit or logit are employed, with the probabilities of the various y_{ij}s modeled as functions of covariates X and cutpoints c_q. Estimation is either achieved by maximizing a likelihood function or, with assignment of priors, a sampling from posterior distributions.

This can be interpreted as a generalization of the ordered probit, with the main departure being the multiple latent variables. In the ordered probit, the y_{ij}s are viewed as realizations of an underlying y^*_{ij}, where the ordinal values are determined by cutpoints on the latent scale. A hierarchical approach views an individual's response as a product of three latent processes: *saliency*, *opinion*, and *decisiveness*. *Saliency*, given by ψ_{ij}, is the first latent factor in the decision-maker's process. If the item is not salient, $\psi_{ij} < 0$ and the respondent will elicit a NA/DK response.

If the item is salient, the next stage involves computing *i's opinion* about the location of j, ϑ_{ij}. Since placing stimuli at extremes of the scale might be systematically different from placing them at the center, respondents whose latent opinion is more extreme are assumed to have definitive opinions. The extremity is defined in terms of cutpoints, c_q, where $q = 1, \ldots, Q-1$ represents the ordinal response category; additionally, $c_0 = -\infty$ and $c_Q = \infty$. An opinion ϑ_{ij} is considered extreme if

$$\vartheta_{ij} \notin [c_L, c_H], \tag{3}$$

where the cutpoints c_L and c_H depend on the number of possible ordinal responses given on the particular item. In particular, we assume $c_L = c_{q_L - 1}$ and $c_H = c_{q_H}$. The indifference zone boundaries q_H and q_L are typically chosen so that they straddle the cutpoint that corresponds to the middle category. For example, if the observed data are from a seven-point scale, category 4

[12] This formulation is simplified from the original paper, but is equivalent with appropriate substitutions.

[13] However, it still requires a minimum number of stimuli are scaled by a given respondent.

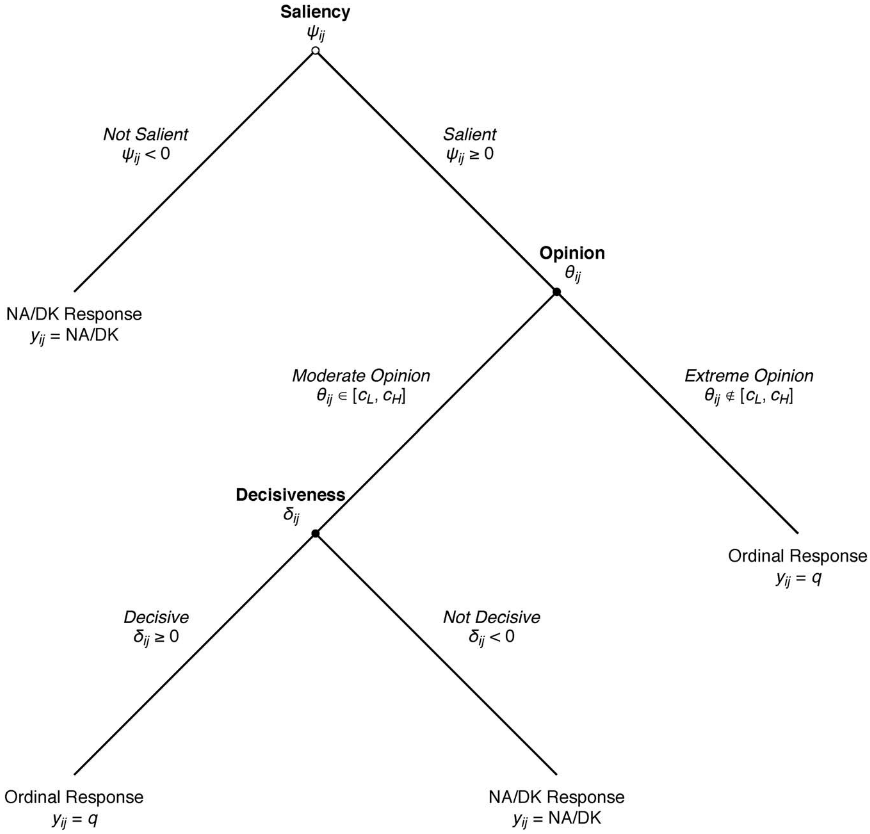

Fig. 2. The hierarchical model

is at the center. This makes $q_L = 3$ and $q_H = 5$ ideal candidates for $c_L = c_2$ and $c_H = c_5$, respectively.[14]

Should i have a ϑ_{ij} that satisfies the above condition, we assume he will elicit an ordinal response. However, if $\vartheta_{ij} \in [c_L, c_H]$, we say that i is in the *indifference zone*. This, in turn, leads to the last stage in the decision tree. Latent opinions in this range lead to one of two observed behaviors. If the individual is decisive, then he would be more inclined to elicit an ordinal response than if he were indecisive. This notion is formalized in the third latent variable δ_{ij}, where $\delta_{ij} \geq 0$ implies i's decisiveness on the item and hence, an ordinal response will be given. If he is not decisive, $\delta_{ij} < 0$ and the NA/DK response is given. The entire process is depicted in Figure 2.

This model provides a rich description of behavior. For example, the NA/DK response can be observed if the respondent has a high expected value of nonresponse (indecisiveness) or if the respondent lacks enough information to form an opinion (salience). These are different kinds of NAs and are modeled as such. Additionally, if there are no NAs, this model reduces to a simple ordered probit; this model is therefore always preferred, as it picks up effects that the ordered probit would miss, but still produces the same results when NAs are absent.

[14] Depending on the data, researchers may modify the indifference zone bounds accordingly.

To ensure identification, it is assumed the three latent variables, ψ_{ij}, ϑ_{ij}, and δ_{ij} are distributed normally with variance 1.[15] Saliency, ψ_{ij}, is assumed have a mean μ_{ij}^{ψ} such that

$$\mu_{ij}^{\psi} = \eta_i + X_{ij}^{\psi}\beta^{\psi}, \tag{4}$$

where η_i is a random intercept allowing individuals to vary in terms of saliency and X_{ij}^{ψ} a vector of person-item covariates thought to influence saliency. This captures the idea that when certain properties of the stimuli match certain properties of the respondent, the item might then be more (or less) salient. For the latent opinion ϑ_{ij}, the mean is given by

$$\mu_{ij}^{\vartheta} = \alpha_i + \gamma_i \xi_j, \tag{5}$$

where ξ_j is the true latent position of stimulus j and α_i and γ_i are individual-specific distortion parameters. This approach deviates from the original Bradlow and Zaslavsky (1999) model, but is in keeping with the political science literature using the Aldrich and McKelvey (1977) technique (e.g., Hollibaugh, Rothenberg and Rulison 2013; Hare et al. 2015; Ramey 2016).

The final latent variable, decisiveness δ_{ij} has a mean

$$\mu_{ij}^{\delta} = Z_i^{\delta}\beta^{\delta}, \tag{6}$$

and Z_i^{δ} is a vector of covariates affecting decisiveness. These covariates are not indexed by j, as decisiveness is assumed to be a property of the individual and not the items.

These latent draws may be summarized as follows:

$$\psi_{ij} \sim \mathcal{N}\left(\mu_{ij}^{\psi}, 1\right), \tag{7}$$

$$\vartheta_{ij} \sim \mathcal{N}\left(\mu_{ij}^{\vartheta}, 1\right), \tag{8}$$

$$\delta_{ij} \sim \mathcal{N}\left(\mu_{ij}^{\delta}, 1\right). \tag{9}$$

The complete data likelihood, based on the above definitions, is given by

$$\mathcal{L}(\phi_1, \phi_2 \,|\, y_{ij}, X, Z) \propto \prod_{i=1}^{N}\prod_{j=1}^{J} p_{ij}(\phi_1, \phi_2 \,|\, y_{ij}, X), \tag{10}$$

where the p_{ij} are probabilities associated with the ordinal outcomes. For simplicity, Figure 3 presents a simplified decision tree broken down into three different probabilities: r, s, and t. To evaluate p_{ij} in these terms, we need to look at the various responses that could be provided on the ordinal items. First, we consider the case where i elicits NA/DK. This could have resulted

15 Also see Rivers (2003). Our decision to employ normally distributed latent variables is driven largely by convention. Given the binary nature of the decision to elicit a response and the ordinal nature of a particular response on the seven-point scale, it seems natural to formulate this problem in a manner similar to the latent variable formulations of (ordered) probit, which allows a data-augmented Gibbs sampler to be constructed based on the latent induced normality. In sum, we believe the gains from normality, coupled with its longstanding use in this class of models, outweigh potential costs associated with functional form dependence (e.g., Clinton, Jackman and Rivers 2004).

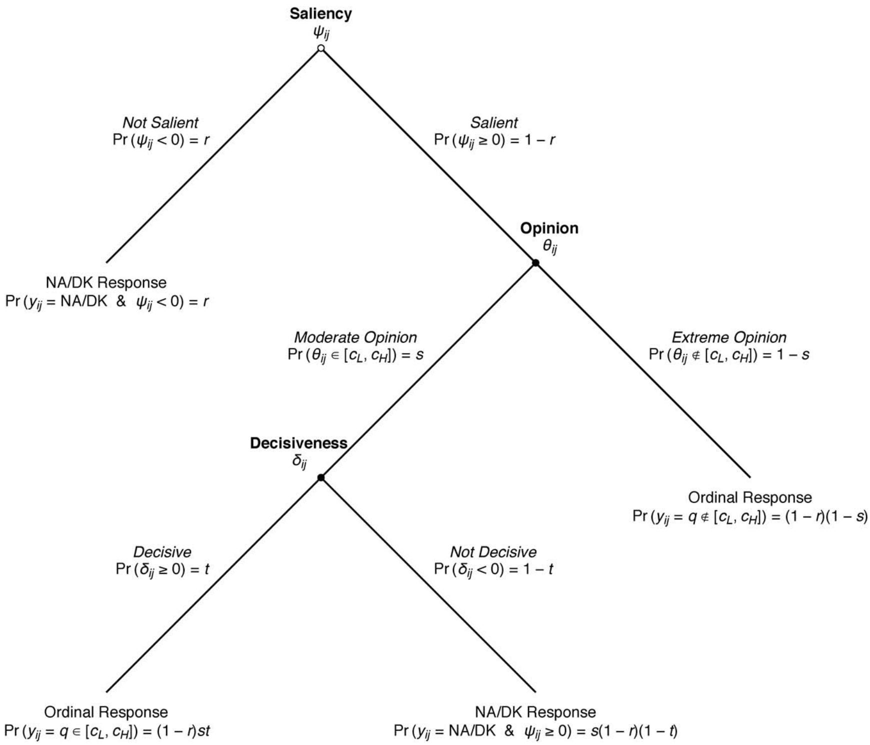

Fig. 3. The hierarchical model with probabilities

in two ways, as seen in Figure 3. Thus, the probability of observing a NA/DK response is

$$Pr\left(y_{ij} = \text{NA} / \text{DK}\right) = r + s(1-r)(1-t)$$
$$= \Phi\left(-\mu_{ij}^{\psi}\right) + \left(\Phi\left(c_{q_H} - \mu_{ij}^{\vartheta}\right) - \Phi\left(c_{q_L} - \mu_{ij}^{\vartheta}\right)\right)\left(1 - \Phi\left(-\mu_{ij}^{\psi}\right)\right)\Phi\left(-\mu_{ij}^{\delta}\right). \tag{11}$$

Second, we look at non-NA/DK responses that fall outside of the indifference zone. The probability of observing a response outside of this zone is given by

$$Pr\left(y_{ij} = q \notin [q_L, q_H]\right) = (1-r)(1-s)$$
$$= \left(1 - \Phi\left(-\mu_{ij}^{\psi}\right)\right)\left(\Phi\left(c_q - \mu_{ij}^{\vartheta}\right) - \Phi\left(c_{q-1} - \mu_{ij}^{\vartheta}\right)\right). \tag{12}$$

Finally, there is the probability of observing a non-NA/DK response that is within the indifference zone. Examining Figure 3, this is given by

$$\Pr\left(y_{ij} = q \in [q_L, q_H]\right) = (1-r)st$$
$$= \left(1 - \Phi\left(-\mu_{ij}^{\psi}\right)\right)\left(\Phi\left(c_q - \mu_{ij}^{\vartheta}\right) - \Phi\left(c_{q-1} - \mu_{ij}^{\vartheta}\right)\right)\left(1 - \Phi\left(-\mu_{ij}^{\delta}\right)\right). \tag{13}$$

We can assemble the pieces in Equations 11 through 13 into a single statement as follows:[16]

$$p_{ij}(y_{ij}) = \begin{cases} \Phi\left(-\mu_{ij}^{\psi}\right) + \left(\Phi\left(c_{q_H} - \mu_{ij}^{\vartheta}\right) - \Phi\left(c_{q_L} - \mu_{ij}^{\vartheta}\right)\right)\left(1 - \Phi\left(-\mu_{ij}^{\psi}\right)\right)\Phi\left(-\mu_{ij}^{\delta}\right), & \text{if } y_{ij} = \text{NA / DK} \\ \left(1 - \Phi\left(-\mu_{ij}^{\psi}\right)\right)\left(\Phi\left(c_q - \mu_{ij}^{\vartheta}\right) - \Phi\left(c_{q-1} - \mu_{ij}^{\vartheta}\right)\right), & \text{if } y_{ij} = q \notin [q_L, q_H] \\ \left(1 - \Phi\left(-\mu_{ij}^{\psi}\right)\right)\left(\Phi\left(c_q - \mu_{ij}^{\vartheta}\right) - \Phi\left(c_{q-1} - \mu_{ij}^{\vartheta}\right)\right)\left(1 - \Phi\left(-\mu_{ij}^{\delta}\right)\right), & \text{if } y_{ij} = q \in [q_L, q_H]. \end{cases}$$

The second layer looks at the vector of prior parameters $\phi_1 = (\eta, \alpha, \gamma, \xi)$. Each of these is normally distributed as follows:

$$\eta_i \sim \mathcal{N}(Z_i^{\eta}\beta^{\eta}, \sigma_{\eta}^2), \tag{14}$$

$$\alpha_i \sim \mathcal{N}(Z_i^{\alpha}\beta^{\alpha}, \sigma_{\alpha}^2), \tag{15}$$

$$\gamma_i \sim \mathcal{N}(Z_i^{\gamma}\beta^{\gamma}, \sigma_{\gamma}^2), \tag{16}$$

$$\xi_j \sim \mathcal{N}(0, 1). \tag{17}$$

The terms Z_i^{η}, Z_i^{α}, and Z_i^{γ} are matrices of covariates assumed to influence saliency and scale usage. For identification, we place standard Normal priors on the stimuli. In the following application, we assume covariates are the same across parameters (i.e., $Z_i^{\eta} = Z_i^{\alpha} = Z_i^{\gamma}$).

The final layer is the vector of hyperparameters for the coefficients, variances, and cutpoints:

$$\phi_2 = (\beta^{\psi}, \beta^{\delta}, \beta^{\eta}, \beta^{\alpha}, \beta^{\gamma}, \sigma_{\eta}^2, \sigma_{\alpha}^2, \sigma_{\gamma}^2, \boldsymbol{c}).$$

Each group of coefficients for each latent parameter $m \in \{\psi, \delta, \eta, \alpha, \gamma\}$ are drawn from a multivariate normal distribution with mean 0 and covariance Σ:

$$\beta^m \sim \mathcal{MVN}(0, \Sigma). \tag{18}$$

The choice of Σ can be as large or as small as necessary.[17] For the variance of the saliency intercept and the distortion vectors, we employ an uninformative conjugate prior. Standard results show this distribution to be the inverse Gamma. Thus, for $k \in \{\eta, \alpha, \gamma\}$, the prior for σ_k^2 is

$$\sigma_k^2 \sim \mathcal{IG}\left(\frac{\rho}{2}, \frac{1}{2}\right), \tag{19}$$

and ρ is chosen to be $\frac{1}{2}$.

Last is the vector of cutpoints. All cutpoints are assumed to be drawn uniformly from the last cutpoint to the current cutpoint. More specifically

$$c_q \mid c_{q-1}, c_{q+1} \sim U\left(c_{q-1}, c_{q+1}\right), q = 1, 2, \ldots, Q-1, \tag{20}$$

$c_0 = -\infty$, $c_Q = \infty$, and, for identification, some $q' \in \{1, 2, \ldots, Q-1\}$, $c_{q'} = -1.5$.

[16] In Equation 12, $\Phi\left(c_q - \mu_{ij}^{\vartheta}\right) - \Phi\left(c_{q-1} - \mu_{ij}^{\vartheta}\right)$ is denoted as $1 - s$, whereas in Equation 13 it is defined as s. This is because s represents an arbitrary choice inside the indifference zone, whereas $\Phi\left(c_q - \mu_{ij}^{\vartheta}\right) - \Phi\left(c_{q-1} - \mu_{ij}^{\vartheta}\right)$ represents a particular choice q assumed to be either outside (Equation 12) or inside (Equation 13) the indifference zone.

[17] Here, the diagonal elements of Σ are chosen to be 25. This is substantially larger than the value of 4 chosen by Bradlow and Zaslavsky (1999) and ensures that the Markov Chain Monte Carlo [MCMC] chains are free to traverse a wide swath of the parameter space. However, the substantive effects of this decision are minimal.

We combine the expressions for the likelihood and priors to form the complete posterior:

$$\pi(\phi_1, \phi_2 \mid y, X) \propto \mathcal{L}(y \mid \phi_1, \phi_2, X) p(\phi_1 \mid \phi_2, y, X) p(\phi_2). \tag{21}$$

As is the case with all hierarchical models, there is a large number of parameters to estimate; we are required to estimate a minimum of $3N+J+Q-3$ parameters (not including the βs or σ^2s). Fortunately, Bayesian methods are well suited for these sorts of models, and we use Plummer's (2003) JAGS software to sample from the full posterior.

PREDICTING NONRESPONSE WITH PERSONALITY

To examine the relationships between Neuroticism and nonresponse, we use the 2014 Cooperative Congressional Election Study (CCES). We asked 1000 respondents to place themselves and nine political figures on a seven-point ideological scale. These included Barack Obama, Hillary Clinton, Jeb Bush, Rand Paul, Ted Cruz, the Democratic and Republican Parties, the Tea Party, and the Supreme Court. We also asked them to take the Ten-Item Personality Inventory ('TIPI') to estimate their Big Five traits on a 1–7 scale, which were normalized to a 0–1 scale (Gosling, Rentfrow and Swann 2003).[18] We dropped those who failed to place themselves on either scale.

We first estimate binomial probit models where the dependent variable is the number of NA/DK responses elicited.[19] Along with the Big Five, we include as covariates respondents' age, gender, income, and education, indicator variables for whether they identified as Black, Hispanic, or some other race, a variable (*High News Interest*) equaling 1 if the respondent indicated he or she "follow[s] what's going on in government" most of the time, and an additional variable (*Unknown News Interest*) equaling 1if the respondent did not know how often he or she followed current events, as these concepts have been shown to explain ideological uncertainty (e.g., Jackson 1993; Alvarez and Franklin 1994; Delli Carpini and Keeter 1996).[20,21]

We find Neuroticism to be positively associated with higher proportions of NA/DK responses. Figure 4 illustrates how the predicted proportion of NA/DK responses increases as Neuroticism increases.[22] As can be seen, the predicted proportion of NA/DK responses increases from about 0.338 to about 0.455 as Neuroticism increases across the plotted range.[23]

[18] While the TIPI is shorter than standard instruments, it is well suited to time-limited tasks like the CCES, and results from the TIPI tend to be highly correlated with the results one would get from longer question batteries (Gosling, Rentfrow and Swann 2003; Ehrhart et al. 2009).

[19] In total, 16 respondents were not asked about Rand Paul, and 84 were not asked about Ted Cruz. This is not a problem for the current analysis, as the binomial framework weights respondents based on the number of questions asked. However, for those that follow, we remove those who were not *asked* about all stimuli. Including them would result in incorrect estimates, as we would be including sources of missingness not due to nonresponse.

[20] Though other research in this vein includes political knowledge as an additional covariate (e.g., Alvarez and Franklin 1994; Glasgow and Alvarez 2000; Gerber et al. 2011), this variable is, unfortunately, not available in our data. However, *News Interest* should be highly correlated with the underlying trait.

[21] The results of these models are presented in the supplementary Appendix.

[22] In Figure 4, the estimates from Model 5 from Table C-1 in the supplementary Appendix were used, as it had the lowest Bayesian Information Criterion [BIC] of those models estimated. All continuous [categorical] variables were set to their means [modes].

[23] That the effects of Agreeableness are comparable in both direction and size to those of Neuroticism (see Appendix) should be noted. While we do not focus on this trait here, previous research (e.g., Mondak 2010) has suggested more agreeable individuals tend to have lower levels of political knowledge, in part because they are less likely to discuss politics or be exposed to political disagreement. As we see later, the hierarchical model indicates Agreeableness is negatively associated with the probability of finding the stimuli salient, presumably due to the lower levels of discussion.

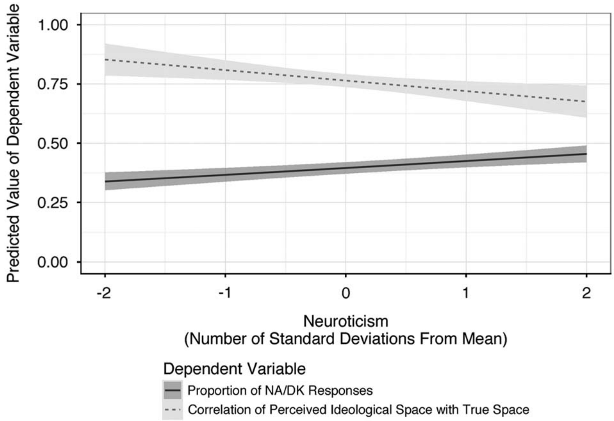

Fig. 4. Predicted results from binomial probit and tobit regression models

We next use AM estimation to recover the ideological space. Though this method does not allow one to include NA/DK responses, one benefit is that it provides estimates of respondents' political information, based on the correlation of the "true" ideological space with how they perceive it.[24] Higher values indicate respondents' perceptions correlate highly with reality. A series of tobit models were run where the dependent variable is the respondents' estimated level of political information and the independent variables are those used in the prior binomial probit regressions.[25]

We find that higher levels of Neuroticism are correlated with lower correlations between the "true" ideological space and respondents' perceptions. Figure 4 illustrates how the predicted correlation decreases as Neuroticism increases.[26] Over the plotted range, the predicted correlation of the perceived space with the true space decreases from about 0.853 to about 0.676.

However, these results do not allow us to analyze the latent decision-making process, nor do they allow us to distinguish between different kinds of nonresponse. Thus, we shift our focus to the hierarchical model. For the following analyses, we ran the model for 90,000 draws each for eight chains (with an additional tuning period of 10,000 draws/chain) and burned the first 40,000 draws/chain, leaving us with 50,000 draws/chain from the posterior distribution (for a total of 400,000 draws across all eight chains). We then applied a thinning interval of four, leaving us with 12,500 thinned draws/chain (for a total of 100,000 draws total). As we show in the Appendix, the Gelman-Rubin (1992) and Brooks-Gelman (1998) diagnostics—as well as traceplots and running mean plots—suggests convergence was achieved.

For this model, the regressors included in the Opinion Intercept, Opinion Slope, Saliency Intercept, and Decisiveness regressions are all those contained in the most fully specified

[24] Hare et al.'s (2015) Bayesian version of the AM algorithm allows for missing values, but relies on the MAR assumption.

[25] The results of these tobit models are presented in Table C-2 in the Appendix. Tobit models are used because the dependent variable lies in the $[-1, 1]$ interval by construction; substantively identical results are found if a zero-one inflated β regression model—with the dependent variable rescaled to lie within the $[0, 1]$ interval—is used. Though the (rescaled) β model is arguably more appropriate due to the tobit's censoring assumptions, the tobit model is more easily interpreted.

[26] In the tobit results in Figure 4, the estimates from Model 7 from Table C-2 in the supplementary Appendix were used, as it had the lowest BIC of all models estimated. All continuous [categorical] variables were set to their means [modes].

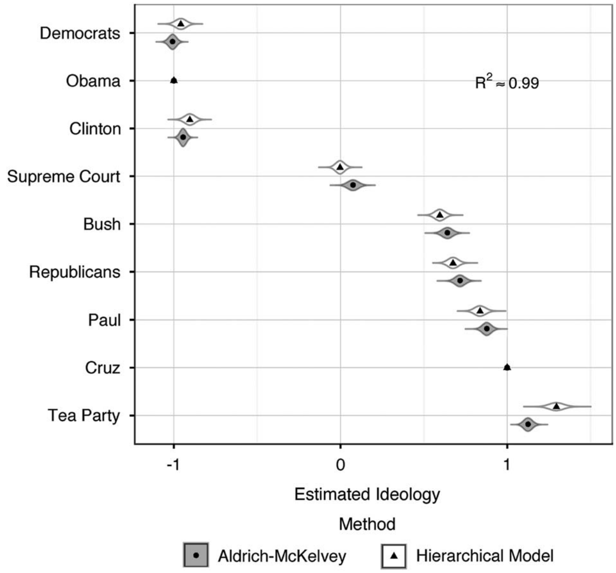

Fig. 5. Comparison of estimated ideology of stimuli using the Aldrich-McKelvey method versus the hierarchical model
Note: Points for the hierarchical model indicate median estimates, and the violin plots indicate the posterior distributions; placements of Obama and Cruz are fixed at −1 and 1, respectively, to identify the scale.

binomial probit and tobit models discussed earlier.[27] To estimate the stimuli-dependent slopes in the Saliency regression, we create an indicator variable that equals 1 if the stimulus is of the same party as the respondent, and 0 otherwise (*Copartisan*).[28] Finally, on the seven-point scale, we set 3 ("Somewhat Liberal") and 5 ("Somewhat Conservative") to be the bounds of the indifference zone.

For our purposes, the most relevant results from the hierarchical model are the estimated posterior distributions of the estimated effects of the Big Five traits and news interest on the Opinion Slopes, Saliency Intercepts, and Decisiveness, all of which are displayed in Figure 6.[29] However, we first examine model fit, through comparison with some "baseline" model.[30] Figure 5 presents the estimated posterior distributions for the ideologies of the stimuli, as well as bootstrapped distributions (based on 100,000 draws) from the AM estimation.[31] To identify the scale for the hierarchical model, the ideologies of President Obama and Senator Ted Cruz are fixed at −1 and 1, respectively.[32] All posterior distributions are tight, and correlate highly

[27] These are Models 6 and 12 from Tables C-1 and C-2 in the Appendix.
[28] The Democratic Party, Barack Obama, and Hillary Clinton were coded as Democratic. The Republican Party, the Tea Party, the Supreme Court, Ted Cruz, Jeb Bush, and Rand Paul were coded as Republican.
[29] Full hierarchical model results are presented in Table C-3 in the Appendix.
[30] We also perform a series of posterior predictive checks, with the results thereof in the supplementary Appendix.
[31] Violin plots like those presented in Figures 5 and 6 are similar to boxplots, with the additional advantage of providing the entire kernel density, and allow for a compact way of comparing multiple distributions in a small amount of space (Hintze and Nelson 1998).
[32] For presentational purposes, the AM estimates are linearly rescaled to ensure that the estimates of President Obama and Ted Cruz are −1 and 1, respectively. This rescaling has zero effect on the reported correlation.

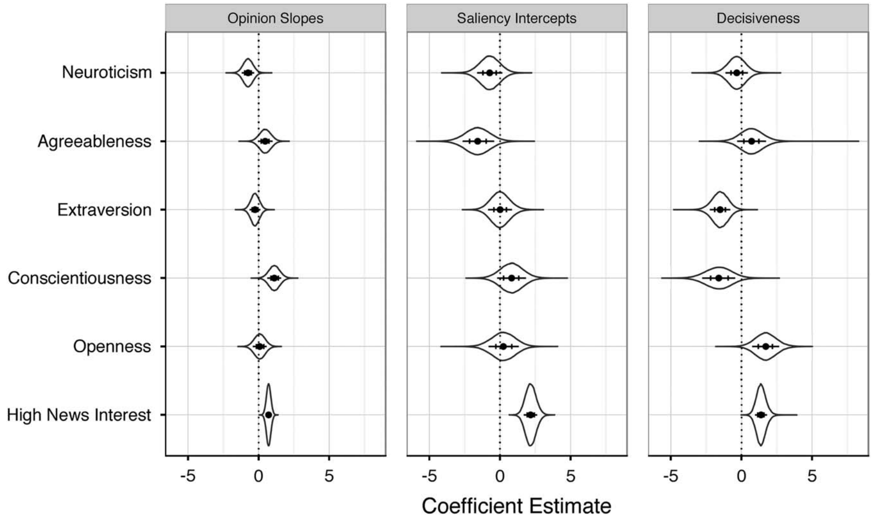

Fig. 6. Posterior distributions of estimated effects of traits and high news interest on Opinion Slopes, Saliency Intercepts, and Decisiveness
Note: Points indicate medians. Bars indicate 80 percent highest posterior density [HPD] intervals. Ticks indicate 50 percent HPD intervals.

($R^2 = 0.99$) with those recovered from the AM estimation, suggesting the hierarchical model taps into the same dimension as the AM method, with the advantage of being able to provide information about the latent decision-making process. These results provide us with a high degree of confidence in our model's performance.[33]

Turning to Figure 6, the hierarchical model's estimates indicate Neuroticism has a negative effect in the Opinion Slope equation; indeed, the vast majority of the posterior distribution lies to the left of 0. This suggests Neuroticism is associated with ideological placement being less highly correlated with the underlying ideological space, a result consistent with those from the AM models, further supporting Hypothesis 4.

One draw of the hierarchical model is its ability to model the decision-making process and pinpoint why Neurotic individuals are more likely to provide NA/DKs. As seen in Figure 6, the coefficient on Neuroticism in the Saliency Intercept equation is negative, and over 85 percent of the posterior distribution lies to the left of 0. This result provides evidence suggesting Neurotic individuals are less likely to find the stimuli salient, increasing the probability of NA/DKs, in support of Hypothesis 2. Additionally, political interest is associated with higher Saliency Intercepts, suggesting those who are more politically interested are more likely to find the stimuli salient and therefore less likely to provide NA/DKs.[34]

[33] This is especially true since our posterior predictive checks (see the Appendix) indicate high degrees of fit in terms of the predicted survey responses, and Figure 5 indicates the same in terms of the recovered ideologies.
[34] We also find a strong negative relationship between Agreeableness and the Saliency Intercept. This is consistent with our result from the binomial probit regressions (see Appendix) which found that more Agreeable individuals are more likely to provide NA/DK responses. This is also largely consistent with—and provides a psychological microfoundation for—previous research on political knowledge (e.g., Mondak 2010) that has suggested more Agreeable individuals are less politically knowledgeable, as they are less likely to discuss politics or be exposed to political disagreement. Therefore, these individuals are less likely to even *know* about political figures, therefore depressing their Saliency values, and increasing the rates at which they provide NA/DKs. However, we leave this question for further research.

Turning to the Decisiveness panel in Figure 6, we find consistent—albeit weaker—support for our relevant hypothesis (Hypothesis 1). For the Neuroticism coefficient, the majority of the posterior distribution is to the left of 0. These results provide evidence that Neurotics' increased rate of NA/DK responses (in support of Hypothesis 3) might also arise from indecision as well as inhibition.

There may be concerns that our models suffer from omitted variable bias. This concern cannot be definitively ruled out, but we have reason to believe it is unfounded. Foremost, our analyses incorporate the variables prior research suggests are associated with item salience and decisiveness. Bradlow and Zaslavsky (1999) argue that respondents find questions to be more salient when those they fall into their areas of expertise, and we include education and news interest variables, thus capturing respondents' expertise on politicians' ideological positions. They also argue that decisiveness is associated with the market share of a customer, but as all voters have one vote, this variable has no obvious equivalent for ideological placement of politicians.

Second, our analyses also include the variables held to be important for the prediction of political information items using the Big Five. Mondak et al. (2010) predict political knowledge using the Big Five along with controls for age, education, race, and sex, while Gerber et al. (2011) use variations of them along with an array of controls based on income and employment status. Our analyses (see full tables in the Appendix) include all of the knowledge-related independent variables used by Gerber et al. (2011) alongside some others. Other work on uncertainty related to ideological placement suggests partisanship is important, as well as gender, political information, and the politician's ideology, which we either estimate or control for in our hierarchical model (Bartels 1986; Alvarez and Franklin 1994). Third, as discussed in the Appendix, examination of predicted versus actual prevalence of response, as well as comparison with AM estimates, suggests that our model exhibits good fit and we have reason to have confidence in its performance. Omitted variable bias is difficult to diagnose, but we find little reason for concern in this analysis.

Overall, our analysis suggests more Neurotic individuals are more likely to provide NA/DK responses to ideological placement questions on political surveys, are less likely to accurately perceive the underlying ideological space, and that this is in part due to these questions being less salient to these individuals. We have also uncovered evidence that this response pattern might also be due to more Neurotic individuals being less decisive, in line with our expectations; however, evidence for this latter finding is somewhat weaker. Overall, these results support our conception of Neuroticism as being a proxy for sensitivity to negative outcomes. Finally, the hierarchical model is able to go further than traditional methods of addressing NA/DK responses and look at the factors influencing saliency and decisiveness across individuals; the significant coefficients in all equations are *prima facie* evidence that more conventional methods might result in incorrect inferences, suggesting that this model will help scholars better understand the generation of NA/DK responses in surveys.

DISCUSSION AND CONCLUSION

This paper introduces to political science, a Bayesian hierarchical model for ordinal data that allows for NA/DK responses, driven by latent psychological processes. While our interests lie in the effects of Neuroticism, this framework can be applied to any similar setup, allowing scholars to better model decision-making processes and generate more reliable estimates. Importantly, the ideological estimates produced by this model are nearly identical to those of the AM model, suggesting that both tap into the same ideological dimension. However, the ability of the hierarchical model to provide information about the underlying decision-making process gives it an advantage beyond the traditional AM method (and ordered probit). That said, it should be noted that the original Bradlow and Zaslavsky (1999) model was designed for consumer satisfaction surveys, and is therefore an imperfect fit to the data used here; arguably, the middle categories in satisfaction surveys more clearly represent “indifference” than those on ideological scales. However, that we generate nearly identical estimates of perceived ideology across both methods suggests this is not a major problem. Nonetheless, in the future we plan to leverage this model and apply it to the study of political approval. Furthermore, the empirical and theoretical models may be adapted to situations in which penalties originate in social indesirability rather than inaccuracy of opinion.

Substantively, we have found evidence Neuroticism is associated with higher rates of NA/DK responses on surveys and incorrect placement of political figures on an ideological scale. Additionally, we have argued that this is in part due to reduced salience for more Neurotic individuals as well as higher rates of indecisiveness. More broadly, these results are consistent with a theory of opinion formation (and reporting) and information acquisition based on modeling Neuroticism as a core cognitive constraint of sensitivity to negative outcomes. The effects of Neuroticism on salience and decisiveness are consistent with our theory that more Neurotic individuals identify contexts in which they will be indecisive and avoid collecting costly information in anticipation of that decision.

That Neuroticism is associated with lower salience and decisiveness provides support for the model-derived hypothesis that Neuroticism is associated with a decision to avoid paying for costly information, and also provides some support for the hypothesis that Neuroticism is associated with the decision not to form a belief. The relative weakness of this latter finding would appear to be a puzzle for our theory. However, our model assumes an incorrect response carries a penalty, which may not hold for our internet survey data. In the absence of a personal interviewer, there may be no penalty for incorrect responses, which would result in no theoretical relationship between Neuroticism and decisiveness. In the course of everyday conversations where individuals dispute political facts and provide social penalties for incorrect assertions, the outlined mechanism would strengthen, promoting indecisiveness and a lack of information gathering among Neurotics. Future work incorporating clear negative outcomes for incorrect assertions through face-to-face survey administration, verification, and/or material negative outcomes for incorrect responses is necessary to clearly resolve this puzzle.

Additionally, our results speak to the larger literature on opinion uncertainty. While much previous psychological research has focused on personality and its role in opinion uncertainty (e.g., Robinson and Tamir 2005; Flehmig et al. 2007; Mondak et al. 2010), political science research has emphasized opinion uncertainty as a function of education, political interest, and other demographic variables (e.g., Bartels 1986; Jackson 1993; Alvarez and Franklin 1994; Delli Carpini and Keeter 1996).[35] One of the benefits of our model is that it estimates NA/DK responses as functions of three latent factors (salience, opinion, and decisiveness), and these factors themselves are estimated as functions of personality traits and demographic variables (while allowing for the inclusion of other variables), thus unifying the previous literature on nonresponse.

We plan on leveraging the framework of Ramey, Klingler and Hollibaugh (2017) to link personality traits to a wide variety of political behaviors via core cognitive constraints, with the intent of framing underlying psychological processes in terms suitable for formal modeling. Indeed, this paper has shown one way they may be formalized and tested, and provides a blueprint for future scholars to incorporate personality traits into their own research.

[35] Other research investigates the role of value conflict (e.g., Alvarez and Brehm 2002).

REFERENCES

Aldrich, John H., and Richard D. McKelvey. 1977. 'A Method of Scaling With Applications to the 1968 and 1972 Presidential Elections'. *American Political Science Review* 71(1):111–30.

Almlund, Mathilde, Angela Lee Duckworth, James Heckman, and Tim Kautz. 2011. 'Personality Psychology and Economics'. In Eric A. Hanushek, Stephen Machin and Ludger Woessmann (eds), *Handbook of the Economics of Education*, 1–181. New York: Elsevier.

Alvarez, R. Michael, and Charles H. Franklin. 1994. 'Uncertainty and Political Perceptions'. *Journal of Politics* 56(3):671–88.

Alvarez, R. Michael, and John Brehm. 2002. *Hard Choices, Easy Answers: Values, Information, and American Public Opinion*. Princeton, NJ: Princeton University Press.

Bartels, Larry M. 1986. 'Issue Voting Under Uncertainty: An Empirical Test'. *American Journal of Political Science* 30(4):709–28.

Bartels, Larry M. 1999. 'Panel Effects in the American National Election Studies'. *Political Analysis* 8(1):1–20.

Berinsky, Adam J. 1999. 'The Two Faces of Public Opinion'. *American Journal of Political Science* 43(4): 1209–30.

Blackwell, Matthew, James Honaker, and Gary King. 2015. 'A Unified Approach to Measurement Error and Missing Data: Details and Extensions.' *Sociological Methods and Research*. Advance Online Publication. doi: http://dx.doi.org/10.1177/0049124115589052.

Bradlow, Eric T., and Alan M. Zaslavsky. 1999. 'A Hierarchical Latent Variable Model for Ordinal Data From a Customer Satisfaction Survey With "No Answer" Responses'. *Journal of the American Statistical Association* 94(445):43–52.

Brooks, Stephen P., and Andrew Gelman. 1998. 'General Methods for Monitoring Convergence of Iterative Simulations'. *Journal of Computational and Graphical Statistics* 7(4):434–55.

Clinton, Joshua, Simon Jackman, and Douglas Rivers. 2004. 'The Statistical Analysis of Roll Call Data'. *American Political Science Review* 98(2):355–70.

Clinton, Joshua D., Anthony Bertelli, Christian R. Grose, David E. Lewis, and David C. Nixon. 2012. 'Separated Powers in the United States: The Ideology of Agencies, Presidents, and Congress'. *American Journal of Political Science* 56(2):341–54.

Delli Carpini, Michael X., and Scott Keeter. 1996. *What Americans Know About Politics and Why it Matters*. New Haven, CT: Yale University Press.

Derryberry, Douglas, and Marjorie A. Reed. 1994. 'Temperament and Attention: Orienting Toward and Away From Positive and Negative Signals'. *Journal of Personality and Social Psychology* 66(6): 1128–39.

DeYoung, Colin G., and Jeremy R. Gray. 2009. 'Personality Neuroscience: Explaining Individual Differences in Affect, Behavior, and Cognition'. In Philip J. Corr and Gerald Matthews (eds), *The Cambridge Handbook of Personality Psychology*, 323–46. New York: Cambridge University Press.

Dietrich, Bryce J., Scott Lasley, Jeffery J. Mondak, Megan L. Remmel, and Joel Turner. 2012. 'Personality and Legislative Politics: The Big Five Trait Dimensions Among U.S. State Legislators'. *Political Psychology* 33(2):195–210.

Ehrhart, Mark G., Karen Holcombe Ehrhart, Scott C. Roesch, Beth G. Chung-Herrera, Kristy Nadler, and Kelsey Bradshaw. 2009. 'Testing the Latent Factor Structure and Construct Validity of the Ten-Item Personality Inventory'. *Personality and Individual Differences* 47(8):900–905.

Flehmig, Hagen C., Michael Steinborn, Robert Langner, and Karl Westhoff. 2007. 'Neuroticism and the Mental Noise Hypothesis: Relationships to Lapses of Attention and Slips of Action in Everyday Life'. *Psychology Science* 49(4):343–60.

Fraley, R. Chris, and Brent W. Roberts. 2005. 'Patterns of Continuity: A Dynamic Model for Conceptualizing the Stability of Individual Differences in Psychological Constructs Across the Life Course'. *Psychological Review* 112(1):60–74.

Gelman, Andrew, and B. Rubin Donald. 1992. 'Inference From Iterative Simulation Using Multiple Sequences'. *Statistical Science* 7(4):457–511.

Gelman, Andrew, Gary King, and Chuanhai Liu. 1998. 'Not Asked and Not Answered: Multiple Imputation for Multiple Surveys'. *Journal of the American Statistical Association* 93(443):846–57.

Gerber, Alan S., Gregory A. Huber, David Doherty, and Conor M. Dowling. 2011. ‘Personality Traits and the Consumption of Political Information’. *American Politics Research* 39(1):32–84.

Gerber, Alan S., Gregory A. Huber, David Doherty, Conor M. Dowling, and Shang E. Ha. 2010. ‘Personality and Political Attitudes: Relationships Across Issue Domains and Political Contexts’. *American Political Science Review* 104(1):111–33.

Glasgow, Garrett, and R. Michael Alvarez. 2000. ‘Uncertainty and Candidate Personality Traits’. *American Politics Research* 28(1):26–49.

Goldberg, Lewis R. 1981. ‘Language and Individual Differences: The Search for Universals in Personality Lexicons’. In Ladd Wheeler (ed.), *Review of Personality and Social Psychology*, vol. 2, 141–65. Beverly Hills, CA: Sage.

Gosling, Samuel D., Peter J. Rentfrow, and William B. SwannWilliam B. Jr. 2003. ‘A Very Brief Measure of the Big-Five Personality Domains’. *Journal of Research in Personality* 37(6):504–28.

Gray, Jeffrey A., and Neil McNaughton. 2003. *The Neuropsychology of Anxiety: An Enquiry Into the Functions of the Septo-Hippocampal System*. New York: Oxford University Press.

Hall, Matthew E. K. 2015. ‘Judging With Personality: A Psychoeconomic Model of Supreme Court Behavior’. Manuscript.

Hare, Christopher, David A. Armstrong, Ryan Bakker, Royce Carroll, and Keith T. Poole. 2015. ‘Using Bayesian Aldrich-McKelvey Scaling to Study Citizens’ Ideological Preferences and Perceptions’. *American Journal of Political Science* 59(3):759–74.

Heckman, James J. 1976. ‘The Common Structure of Statistical Models of Truncation, Sample Selection and Limited Dependent Variables and a Simple Estimator for Such Models’. *Annals of Economic and Social Measurement* 5(4):475–92.

Hintze, Jerry L., and Ray D. Nelson. 1998. ‘Violin Plots: A Box Plot-Density Trace Synergism’. *The American Statistician* 52(2):181–84.

Hollibaugh, Gary E. Jr., Lawrence S. Rothenberg, and Kristin K. Rulison. 2013. ‘Does it Really Hurt to be Out of Step?’. *Political Research Quarterly* 66(4):856–67.

Holroyd, Clay B., and Michael G. H. Coles. 2002. ‘The Neural Basis of Human Error Processing: Reinforcement Learning, Dopamine, and the Error-Related Negativity’. *Psychological Review* 109(4):679–709.

Jackson, John E. 1993. ‘Attitudes, No Opinions, and Guesses’. *Political Analysis* 5(1):39–60.

Jessee, Stephen A. 2015. ‘“Don’t Know” Responses, Personality, and the Measurement of Political Knowledge’. *Political Science Research and Methods.* Advance Online Publication. doi: http://dx.doi.org/10.1017/psrm.2015.23

John, Oliver P. 1990. ‘The “Big Five” Factor Taxonomy: Dimensions of Personality in the Natural Language and in Questionnaires’. In Lawrence A. Pervin and Oliver P. John (eds), *Handbook of Personality: Theory and Research*, 66–100. New York: Guilford Press.

John, Oliver P., Richard W. Robins, and Lawrence A. Pervin. 2008. *Handbook of Personality: Theory and Research* 3rd ed. New York: Guilford Press.

Kahneman, Daniel, and Amos Tversky. 1979. ‘Prospect Theory: An Analysis of Decision Under Risk’. *Econometrica* 47(2):263–91.

King, Gary, James Honaker, Anne Joseph, and Kenneth Scheve. 2001. ‘Analyzing Incomplete Political Science Data: An Alternative Algorithm for Multiple Imputation’. *American Political Science Review* 95(1):49–69.

Krosnick, Jon A., Michael A. Milburn. ‘Psychological Determinants of Political Opinionation’. *Social Cognition* 8(1):49–72.

Little, Roderick J. A., and Donald B. Rubin. 1989. ‘The Analysis of Social Science Data With Missing Values’. *Sociological Methods & Research* 18(2–3):292–326.

McCrae, Robert R., and Paul T. Costa Jr. 1996. ‘Toward a New Generation of Personality Theories: Theoretical Contexts for the Five-Factor Model’. In Jerry S. Wiggins (ed.), *The Five-Factor Model of Personality: Theoretical Perspectives*, 51–87. New York: Guilford Press.

Mondak, Jeffery J. 2010. *Personality and the Foundations of Political Behavior*. New York: Cambridge University Press.

Mondak, Jeffery J., and Karen D. Halperin. 2008. 'A Framework for the Study of Personality and Political Behaviour'. *British Journal of Political Science* 38(2):335–62.

Mondak, Jeffery J., Matthew V. Hibbing, Damarys Canache, Mitchell A. Seligson, and Mary R. Anderson. 2010. 'Personality and Civic Engagement: An Integrative Framework for the Study of Trait Effects on Political Behavior'. *American Political Science Review* 104(1):85–110.

Plummer, Martyn. 2003. 'JAGS: A Program for Analysis of Bayesian Graphical Models Using Gibbs Sampling'. In Kurt Hornik, Friedrich Leisch and Achim Zeileis (eds), *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*. Vienna, Austria: Technische Universitat Wien. Available at https://www.r-project.org/conferences/DSC-2003/Proceedings/.

Ramey, Adam. 2015. 'Weighing the Alternatives: Preferences, Parties, and Constituency in Roll Call Voting'. *Journal of Politics* 77(2):421–32.

Ramey, Adam J. 2016. 'Vox Populi, Vox Dei? Crowd-Sourced Ideal Point Estimation'. *Journal of Politics* 78(1):281–95.

Ramey, Adam J., Jonathan D. Klingler, and Gary E. Hollibaugh Jr. 2016. 'Measuring Elite Personality Using Speech'. *Political Science Research and Methods.* Advance Online Publication. doi: http://dx.doi.org/10.1017/psrm.2016.12.

Ramey, Adam J., Jonathan D. Klingler, and Gary E. Hollibaugh Jr. 2017. *More Than a Feeling: Personality, Polarization, and the Transformation of the U.S. Congress*. Chicago, IL: University of Chicago Press.

Rivers, Douglas. 2003. 'Identification of Multidimensional Spatial Voting Models'. Manuscript.

Roberts, Brent W., Nathan R. Kuncel, Rebecca Shiner, Avshalom Caspi, and Lewis R. Goldberg. 2007. 'The Power of Personality: The Comparative Validity of Personality Traits, Socioeconomic Status, and Cognitive Ability for Predicting Important Life Outcomes'. *Perspectives on Psychological Science* 2(4):313–45.

Roberts, Brent W., and Wendy F. DelVecchio. 2000. 'The Rank-Order Consistency of Personality Traits From Childhood to Old Age: A Quantitative Review of Longitudinal Studies'. *Psychological Bulletin* 126(1):3–25.

Robinson, Michael D., and Maya Tamir. 2005. 'Neuroticism as Mental Noise: A Relation Between Neuroticism and Reaction Time Standard Deviations'. *Journal of Personality and Social Psychology* 89(1):107–14.

Rosas, Guillermo, Yael Shomer, and Stephen R. Haptonstahl. 2015. 'No News is News: Nonignorable Nonresponse in Roll-Call Data Analysis'. *American Journal of Political Science* 59(2):511–28.

Rubin, Donald B. 1987. *Multiple Imputation for Nonresponse in Surveys*. New York: Wiley.

Schafer, Joseph L. 1997. *Analysis of Incomplete Multivariate Data*. Boca Raton, FL: CRC Press.

Soldz, Stephen, and George E. Vaillant. 1999. 'The Big Five Personality Traits and the Life Course: A 45-Year Longitudinal Study'. *Journal of Research in Personality* 33(2):208–32.

Tanner, Martin A., and Wing Hung Wong. 1987. 'The Calculation of Posterior Distributions by Data Augmentation'. *Journal of the American Statistical Association* 82(398):528–40.

Tausanovitch, Chris, and Christopher Warshaw. 2013. 'Measuring Constituent Policy Preferences in Congress, State Legislatures, and Cities'. *The Journal of Politics* 75(2):330–42.

Treier, Shawn, and D. Sunshine Hillygus. 2009. 'The Nature of Political Ideology in the Contemporary Electorate'. *Public Opinion Quarterly* 73(4):679–703.

Tversky, Amos, and Daniel Kahneman. 1992. 'Advances in Prospect Theory: Cumulative Representation of Uncertainty'. *Journal of Risk and Uncertainty* 5(4):297–323.

Weisberg, Herbert F. 2009. *The Total Survey Error Approach: A Guide to the New Science of Survey Research*. Chicago, IL: University of Chicago Press.

Zaller, John. 1990. 'Political Awareness, Elite Opinion Leadership, and the Mass Survey Response'. *Social Cognition* 8(1):125–53.

Zaller, John, and Stanley Feldman. 1992. 'A Simple Theory of the Survey Response: Answering Questions Versus Revealing Preferences'. *American Journal of Political Science* 36(3):579–616.