Measuring Political Attitudes and Attributes with Drift Diffusion Models

Andrew Trexler
University of Wisconsin–Madison

Christopher D. Johnston
Duke University

August 31, 2026

Abstract

Individual latent political attitudes and attributes, such as issue preferences, values, or political knowledge, are commonly measured with batteries of closed-ended survey questions. The latent variable is then modeled as a linear function of the battery items, most commonly as a simple average. We employ a novel approach to modeling such latent constructs by incorporating question response time, easily captured in web-based surveys, into the estimation of individual positions on the latent dimension of interest. Specifically, we adapt the drift diffusion model of decision-making, which has become a prominent model for describing perceptual discrimination tasks in psychology, and apply it to the respondent's task of selecting a response option on standard closed-ended questions. Examining a range of latent political constructs with data from the American National Election Studies and two original surveys, we demonstrate that this modeling approach offers gains to precision when analyzing a small number of survey items, potentially allowing survey researchers to do more with less.

Word Count: 3,732

Prepared for the 2026 Annual Meeting
of the American Political Science Association

Contact: Andrew Trexler, atrexler@wisc.edu.

Introduction

Much of the study of public opinion is concerned with what goes on in people's heads (Lippman 1922). Surveys are an attractive means of measuring many things that happen in people's heads, including such quantities as beliefs, attitudes, opinions, and knowledge. These quantities nevertheless remain latent and are not directly observable (Zaller and Feldman 1992), requiring researchers to evaluate answers to one or more survey questions (whether open- or closed-ended) to model individual differences on the latent dimension of interest. With multi-item batteries of survey questions, researchers typically (and most simply) assign numeric values to the response options and average across the items or use an additive index; other researchers engage in more complex approaches such as factor analysis, latent class analysis, or structural equation modeling.

However, with modern computer-assisted self-interviewing (CASI), the selected response option itself is not the only signal of the latent attribute available to the researcher. Survey software routinely allows easy access to item-level response times—that is, how long the respondent takes to decide on a response option and submit their answer—which may provide additional valuable information about the latent quantity of interest. For example, when asked to express a preference between two political candidates, a longer response time is likely a signal of a more difficult decision than a short response time, and this information may be helpful for estimating individual-level preferences on various dimensions of the candidates' qualities or policies.

Here, we adapt the drift diffusion model of decision-making (Ratcliff 1978), which has become a prominent model for describing perceptual discrimination tasks in psychology, and demonstrate its application to the measurement of latent individual attributes in political science. We focus on the measurement of political knowledge as an exemplar, using response times measured in an original survey ($N = 1,531$) and for online respondents to the 2024 American National Election Study (ANES, $N = 4,234$). We show that, in some settings, survey researchers can take advantage of the unobtrusive measurement of response times to improve estimation of latent quantities, potentially allowing the conservation of survey space and limiting respondent fatigue.

Applying Drift Diffusion Models to Political Knowledge

Drift diffusion models, first introduced by Ratcliff (1978), are a family of models that are well-suited to describing behavior in choice tasks that involve two alternatives by incorporating information from both the choices themselves and the time required to make the choice. Conceptually, drift diffusion models attempt to estimate the cognitive process of accumulating evidence in working memory in favor of each alternative, whether derived from perceptual stimuli (e.g., Eriksen and Eriksen 1974) or from considerations sampled from long-term memory (Zaller and Feldman 1992), until sufficient evidence is accumulated favoring one alternative over the other to effect a decision. The most basic drift diffusion model involves four key parameters: the boundary separation a that describes the distance between decision boundaries for the two alternatives (i.e., how much evidence must be accumulated in favor of one alternative over the other to effect a decision), the relative starting point w that describes the point between the boundaries the chooser begins by default (often assumed to be halfway between the boundaries), the non-decision time t_0 that describes the perceptual or reaction time required before evidence begins accumulating and moves the chooser away from w, and finally the drift rate v that describes the chooser's tendency towards one alternative over the other as evidence accumulates.

Figure 1, reproduced from (Henrich et al. 2024, Figure 1), depicts a theoretical four-parameter drift diffusion model with a separating the decision boundaries for two alternatives (A and B), the starting point w halfway between the boundaries, some amount of non-decision time t_0, and a positive drift rate v. The accumulation of evidence is assumed to be semi-random, "drifting" towards a particular alternative as a random walk with a trend (v), as depicted with the jagged black line extending from w at t_0 to the decision boundary (in favor of A). This semi-random accumulation means that, in infinite trials, a chooser with

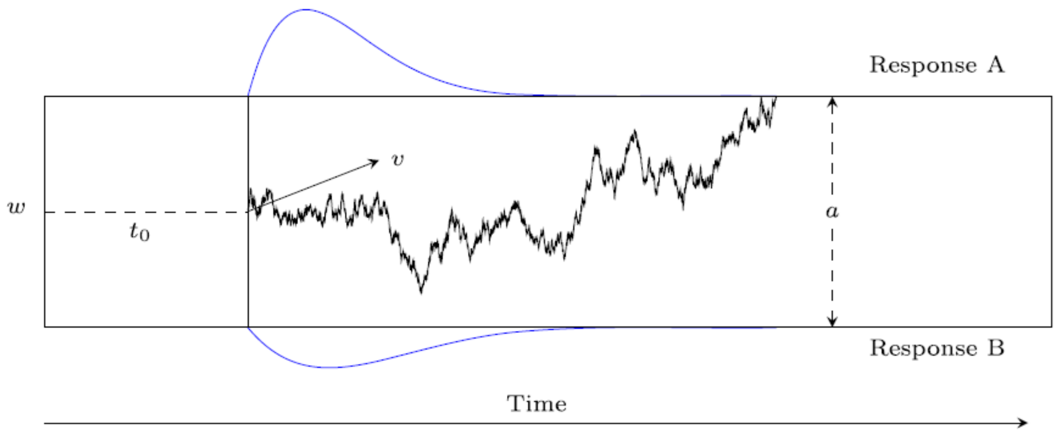

Figure 1: Reproduced from (Henrich et al. 2024, Figure 1). The figure depicts a four-parameter diffusion model of the cognitive process for choosing between two alternatives. The accumulation of evidence in favor of each alternative is depicted as the jagged black line (a random walk with a trend, the drift rate v) and the underlying distribution of expected response times for choosing each alternative are shown in blue.

a positive drift rate v will choose alternative A over B more often than not, but the semi-random accumulation of evidence may cause the chooser to select B in any given trial, if sufficient evidence is randomly accumulated to reach the lower decision boundary first. With a larger (positive) drift rate v, it is easier to accumulate evidence in favor of A to reach the decision boundary more quickly, so in infinite trials the underlying distribution of response times should be less right-skewed for selections of A than of B, as shown in blue in Figure 1.

In political science, drift diffusion models are appropriate for modeling any choice task that can be plausibly collapsed to two-alternative comparisons. Such tasks could include expressions of preferences between pairs of policies, values, traits, candidates, or composite profiles of these attributes (as commonly seen in conjoint experiments, e.g. Hainmueller, Hopkins and Yamamoto 2014); expression of beliefs through true-false decision tasks (e.g., for measuring conspiracy beliefs and the efficacy of interventions, Guay et al. 2026); and even for selection tasks on standard public opinion items such as ideological self-identification, so long as the response options are collapsed to two alternatives (e.g., liberal or conservative).

Here, we focus on an application to the measurement of political knowledge. Political knowledge is commonly measured with closed-ended survey items that ask for some factual piece of information related to politics or government, with one accurate response option and several decoy options and a "don't know" option (Delli Carpini and Keeter 1996). Collapsing across these response options is, in this case, fairly simple: the accurate response option is defined as one ("correct") alternative, and all other response options (including "don't know") are defined is the other alternative ("not correct"). For a single knowledge item, the conventional approach is to distinguish respondents into just two binary groups: those who know or guess the correct answer, and those who do not. Modeling a political knowledge item as a diffusion process, the drift rate v_j expresses the degree of knowledge for respondent j, incorporating information from the response time alongside the response selection itself to provide a continuous measure of knowledge across respondents. Correct answers with fast response times are reflected in a large and positive v_j, indicating more knowledge than correct answers with slower response times (a small but positive v_j), which suggest the respondent required more effort to work out the correct response. Likewise, faster response times with an incorrect or "don't know" answer (and thus a large and negative v_j) indicate less knowledge than a slow but inaccurate answer (a small but negative v_j), which suggests the respondent may have come closer to working out the correct answer.

To be clear, this modeling choice involves several substantive assumptions. In the case of slow but inaccurate responses, there is a very real possibility that the slow decision is a function of difficulty deciding between two equally incorrect response options, rather than between the correct option and alternative (as a small negative v_j indicates). Additionally, collapsing all decoy and "don't know" options into one category assumes that all of these options reflect the same (low) latent level of knowledge, but in reality some decoy response options are often more plausible than others and selecting a more plausible decoy generally reflects a higher latent level of knowledge than selecting a very implausible decoy or simply responding "don't know."

Despite these required assumptions, however, the granularity in estimation of latent knowledge offered by drift diffusion modeling may be a better proxy for latent political knowledge than the standard binary classification based on the response option alone. We test this possibility with a hierarchical Bayesian implementation (Henrich et al. 2024) of the four-parameter drift diffusion model in `stan`, a highly flexible probabilistic programming language for statistical modeling, using an efficient recursive algorithm (the "No-U-Turn Sampler," Hoffman and Gelman 2014) for Hamiltonian Markov Chain Monte Carlo estimation (Neal 2011). This approach enables us to fit drift diffusion models despite sparse data (that is, few knowledge items per respondent)—a common point of contrast between typical survey batteries in political science and the many-trial perceptual experiments in psychology where drift diffusion models originated (e.g., Eriksen and Eriksen 1974; Ratcliff and Smith 2004)—through partial pooling across respondents to estimate both item-level drift rates v_k for each item $k \in 1, ..., k$ and individual-level drift rates v_j for each respondent $j \in 1, ..., j$.

Data

We explore the utility of drift diffusion models for measuring political knowledge with two datasets. The first dataset is an original survey of $N = 1638$ U.S. adults recruited via Lucid Marketplace and implemented via the Qualtrics platform, a non-probability sample provider. On this survey fielded in June 2024, we included a battery of 15 common closed-ended political knowledge questions displayed in a random order, collecting the response time (precise to the nearest microsecond) for each question. The 15 knowledge items and the percentage of the analysis sample that answered each item correctly are summarized in Table 1; these offer a valuable range of difficulty that includes both easy questions (e.g., which party controls the respondent's state governorship, 75.9% correct) and hard questions (e.g., who is the current chairperson of the U.S. Federal Reserve, 18.6% correct). The exact question wording and response options for all items is provided in Appendix B. In line with best practices in the measurement of political knowledge, we asked respondents to commit to answering the knowledge items without outside assistance (i.e., looking up the answer,

Knowledge Item	Correct (%)	Knowledge Item	Correct (%)
Party of R's State Governor	75.9	Identify U.S. Senate President	38.0
Name of First 10 Amendments	70.6	Chamber Requiring Cloture	37.2
Party Control of R's State Leg	61.2	Term of U.S. Senator	34.4
Term of Supreme Court Justice	60.7	Electoral College Tie-breaker	25.7
Margin for Veto Override	52.0	Balance of Supreme Court	20.8
Identify U.S. House Speaker	47.5	Identify UK Prime Minister	19.1
Identify 1st Amendment Right	44.5	Identify Federal Reserve Chair	18.6
Identify U.S. Chief Justice	40.1	Catch Question	6.0*

Table 1: A summary of the knowledge items fielded on the 2024 Lucid survey and the percent of the analysis sample that correctly answered each question. The * indicates the percentage of the full sample instead.

Clifford and Jerit 2016) and included a "catch" question about an obscure U.S. Supreme Court case (Motta, Callaghan and Smith 2017) at the very end of the knowledge battery (Graham 2025). We assume that respondents who fail to commit to answering the knowledge items without outside assistance or correctly answer the catch question (or both) do not provide genuine responses, and accordingly drop them from our analyses; these restrictions, alongside several other general data quality checks (described in Appendix A), reduce our analysis sample to $N = 1531$ respondents.

The second dataset we examine is the 2024 American National Election Study (ANES), a national probability sample of U.S. adults interviewed both before and after the 2024 national election. Specifically, we examine the $N = 4,234$ respondents who completed the pre-election survey via the internet mode, for which question timing data is available (precise to the nearest tenth of a second). The 2024 ANES pre-election survey includes four closed-ended political knowledge items, summarized in Table 2. Like our original survey of Lucid respondents, includes a request to commit to answering without outside assistance as well as a catch question. We exclude respondents who fail to commit to answering without outside assistance or correctly answered the catch question; these exclusions reduce our ANES analysis sample to $N = 3731$ respondents. The exact question wordings and response options for the relevant items are reproduced in Appendix B, and additional information

ANES Knowledge Items	Correct (%)
Party Control of the U.S. House of Representatives	65.2
Party Control of the U.S. Senate	62.0
Term Length for a U.S. Senator	42.8
Foreign Aid Spending < Social Security, Medicare, Defense	28.7
Catch Question	1.9*

Table 2: A summary of the knowledge items fielded on the 2024 ANES pre-election survey and the percent of the analysis sample that correctly answered each question. The * indicates the percentage of the full sample instead.

about the 2024 ANES sampling and methodology is available at electionstudies.org.

Results

We first discuss the analysis and results of our original 2024 Lucid survey. For each of the first $k = 1, ..., 15$ knowledge items presented to each respondent, we fit a unique drift diffusion model with a consistent set of (fairly permissive) priors. Specifically, we set priors for the boundary separation a to be positive and drawn from a normal distribution with a mean of 5 seconds and a standard deviation of 10 seconds, the non-decision time t_0 to be positive and drawn from a normal distribution with a mean of 0.25 seconds and a standard deviation of 1 second, the mean drift μ_{v_k} for any item k to be normally distributed with a mean of 0.16 and a standard deviation of 1, the standard deviation of drift σ_v across respondents $1, ..., j$ to be positive and normally distributed with a mean of 0 and a standard deviation of 10 (i.e., half-normal), and finally set the starting point w halfway between the decision boundaries (i.e., $a/2$). To improve estimation, we also remove observations (at the respondent-item level) with response times greater than 30 seconds (8.0 percent of observations), retaining all other observations from each respondent in the model estimation. For each drift model with k items, we conduct 1,000 warm-up iterations and 1,000 sampling iterations in `stan` to fit the model.

This hierarchical estimation procedure produces continuous, approximately normal estimates of latent political knowledge for each individual respondent, particularly at higher

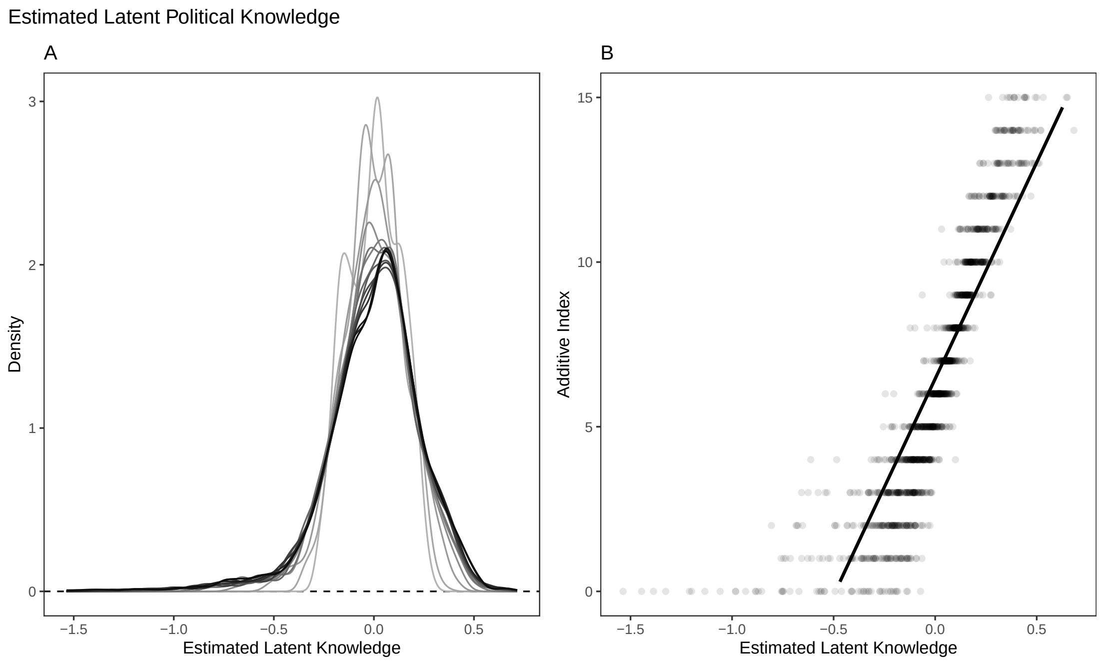

Figure 2: Data from the 2024 Lucid survey. Panel A shows the distribution of estimated latent knowledge (individual-level drift rates v_j) for drift diffusion models with the first k items for $k \in [2, 15]$. Models with higher (lower) values of k are shown in a darker (lighter) shade. Panel B shows the estimated individual-level drift rates v_j for the $k = 15$ model against the additive index of correct answers. The line of best fit between these measures is shown by the solid black line.

values of k. Panel A of Figure 2 shows the distribution of estimated latent political knowledge (that is, the estimated individual-level deviations v_j from the mean drift rate μ_v) for each model that uses the first k items for $k \in [2, 15]$.[1] Panel B shows the relationship between the estimated v_j values and the additive index values (i.e. the sum of correct answers) for the $k = 15$ model, with the line of best fit between the two measures shown by the solid black line. The estimates of latent political knowledge produced by the two measures are tightly correlated, but with much more granular variation in the estimates from the drift model.

Does this more granular variation capture information of value? We take advantage of the relatively large number of political knowledge items fielded on the survey, as well as the random order in which they were presented to the respondent, to analyze how the relative

[1] We exclude the $k = 1$ model for visual clarity, as this distribution is bimodal and tightly concentrated near zero; with only one trial to evaluate, the model weights the response choice much more heavily than the response time.

performance of a drift diffusion model versus a simple additive index changes as more items are included in the estimation. To provide an apples-to-apples comparison across models, for each increment of $k \in [1, 10]$ we regress performance (the number of correct answers) on the *subsequent* five knowledge items in the survey on v_j and, separately, on the additive index for the first k items. For example, for $k = 5$, we regress performance on items 6-10 on the v_j estimates from the $k = 5$ drift model, and separately on the additive index of performance on items 1-5. Figure 3 shows the relative performance of these two models (drift diffusion versus additive index) for each increment of $k \in [1, 10]$.

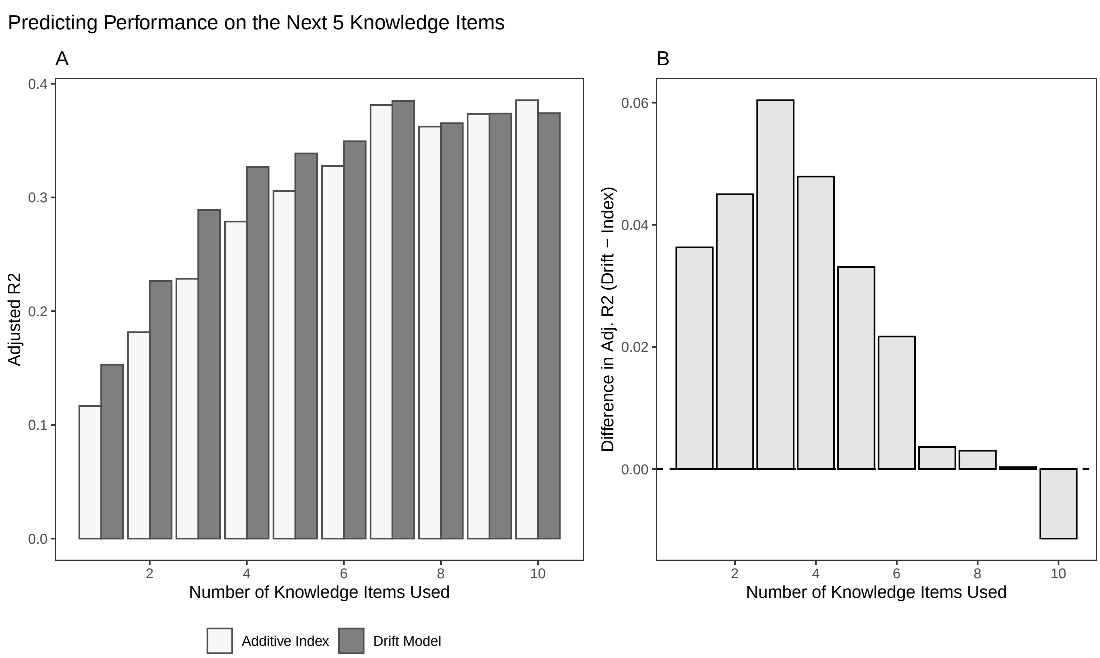

Figure 3: Data from the 2024 Lucid survey. Panel A shows the adjusted R-squared value for the drift diffusion estimates and additive index estimates predicting performance on 5 subsequent knowledge items at each increment of $k \in [1, 10]$. Panel B shows the difference in adjusted R-squared values (drift model minus additive index model) for these same increments of k.

Panel A shows the comparative adjusted R-squared value of these two models at each increment of $k \in [1, 10]$, while Panel B shows the difference in adjusted R-squared values (drift model minus additive index model) for these same increments of k. For low values of k, the drift diffusion model consistently out-performs the additive index, with substantial gains in the adjusted R-squared. Indeed, when $k \leq 5$, the adjusted R-squared value for the drift diffusion model is roughly the same as the adjusted R-squared of an additive index model that uses one to two additional knowledge items—suggesting that drift diffusion modeling may allow survey researchers to obtain similar performance to standard approaches with one or two fewer survey items, conserving survey space and limiting respondent fatigue. These gains appear to quickly peter out and even reverse at higher values of k; a comparison of the Bayesian Information Criterion (BIC) of the two models strongly favors the drift diffusion model for $k \in [1, 8]$, weakly favors the drift diffusional model for $k = 9$, and strongly favors the additive index when $k = 10$. This pattern perhaps suggests that, for batteries with many knowledge items, the amount of noise from the response times begins to overwhelm the signal in the drift diffusion models, and thus researchers with large batteries would be better off simply using the information from the response choices themselves.

2024 ANES

Given the potential promise of drift diffusion modeling shown in our original 2024 Lucid survey, we next turn to the 2024 ANES data. The 2024 ANES has far fewer political knowledge items (4), providing a suitable test of potential gains from drift diffusion modeling at low levels of k, but this also requires us to adopt a slightly different approach to model comparison. For each of the four knowledge items, we fit a unique drift diffusion model using the other three items (only) and calculate an equivalent additive index of correct answers. We then regress the drift model estimates of v_j on binary performance (correct, not correct) on the hold-out item, and the do same using the calculated additive index. We fit the drift diffusion models with the same priors and procedures as for the 2024 Lucid survey data (see above), and again restrict our data to observations with response times of 30 seconds or less (at the respondent-item level, dropping 7.5% of observations) to improve estimation. A comparison of model performance for each of the four held-out items (one per panel) is shown in Figure 4.

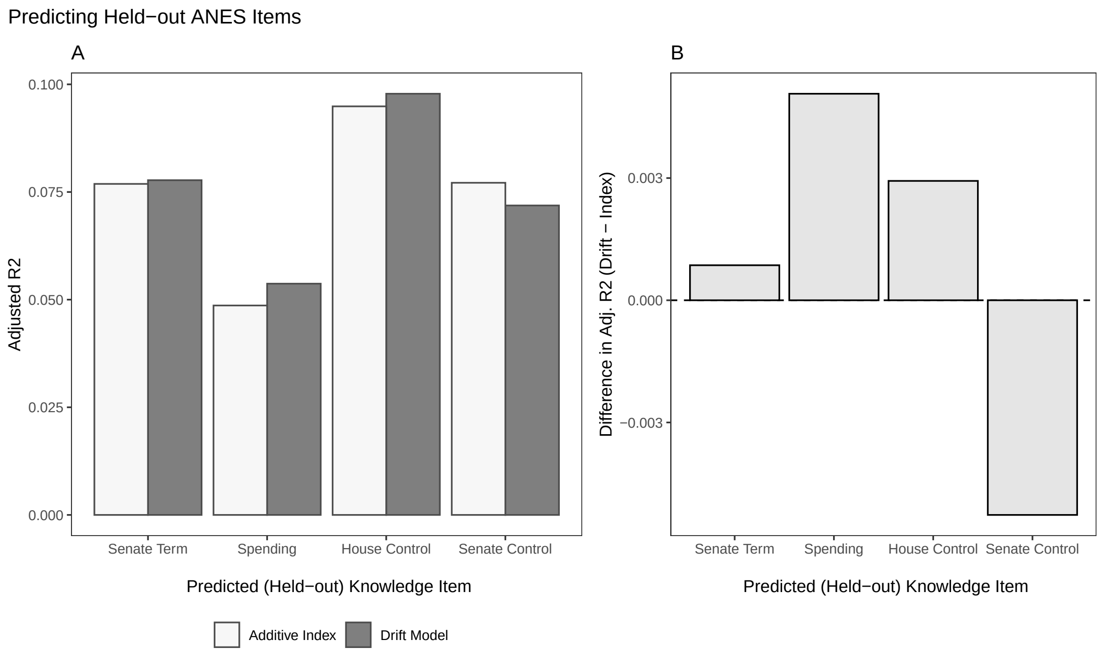

Figure 4: Data from the 2024 ANES pre-election survey. Panel A shows the comparative adjusted R-squared for each 3-item model predicting a held-out fourth knowledge item. Panel B shows the difference in adjusted R-squared values (drift model minus adaptive index model) for each of the same held-out items.

The results of this analysis shows that the drift diffusion models generally offer some limited gains to predictive power, but the evidence is somewhat mixed. In three of four cases, the drift diffusion model outperforms the additive index on predicting the held-out knowledge item; in the fourth, the additive index provides a higher adjusted R-squared value, although in model comparisons the BIC strongly favors the drift diffusion model in all four cases. In all cases, the differences are proportionally much smaller than those observed in the 2024 Lucid sample.

Discussion

Across two large survey datasets, we have conducted an exploratory foray into the promise and limitations of drift diffusion modeling for measuring latent quantities in political science, such as political knowledge. In our original 2024 survey conducted with a non-probability Lucid sample, we observed substantial gains over a conventional additive index in predicting performance on future political knowledge items for short batteries, but this benefit vanished and even reversed for larger batteries. With the 2024 ANES political knowledge items, we observed smaller and less consistent gains from drift diffusion modeling.

The differences between the two samples may help explain some of the disparity in our results across the two datasets. Opt-in online panelists recruited via platforms like Lucid are typically highly professionalized (Hillygus, Jackson and Young 2014; Stagnaro et al. 2026), meaning they are accustomed to the CASI surveying experience and incentivized to complete surveys with a balance of speed and effortful attention to successfully earn a monetary payment. In comparison, respondents to the ANES, a very lengthy survey with a gold-standard probability-based sampling design that recruits primarily by mail, are much less likely to be acclimatized to the survey-taking experience and much more likely to exhibit variable fatigue when completing the political knowledge items. As such, response times in the Lucid survey might provide a slightly less noisy signal of aptitude on the knowledge items than on the ANES survey. Indeed, we observe that response times are slightly shorter in the Lucid sample (media of 9.47 versus 10.20 seconds), although the spans of the inter-quartile ranges are quite similar (10.02 seconds versus 9.90 seconds).

For survey researchers constrained to small batteries by limited survey space, and particularly those working with non-probability online panels, drift diffusion modeling may offer improved measurement of latent quantities of interest. Though we focus here on the measurement of political knowledge, drift diffusion models can be a potentially fruitful tool for many other latent constructs in the study of individual political behavior, from policy preferences and candidate choices to political values, group identification, and more.

References

Clifford, Scott and Jennifer Jerit. 2016. "Cheating on Political Knowledge Questions in Online Surveys: An Assessment of the Problem and Solutions." *Public Opinion Quarterly* 80(4):858–887.

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

Eriksen, Barbara A. and Charles W. Eriksen. 1974. "Effects of Noise Letters Upon the Identification of a Target Letter in a Nonsearch Task." *Perception Psychophysics* 16(1):143–149.

Graham, Matthew H. 2025. "Catch Questions Can Increase Lookups on Political Knowledge Questions." *Public Opinion Quarterly* 89(2):415–423.

Guay, Brian, Adam J. Berinsky, Gordon Pennycook and David Rand. 2026. "Examining Partisan Asymmetries in Fake News Sharing and the Efficacy of Accuracy Prompt Interventions." *The Journal of Politics* p. forthcoming.

Hainmueller, Jens, Daniel J. Hopkins and Teppei Yamamoto. 2014. "Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments." *Political Analysis* 22(1):1–30.

Henrich, Franziska, Raphael Hartmann, Valentin Pratz, Andreas Voss and Karl Christoph Klauer. 2024. "The Evolution of Election Polling in the United States." *Behavior Research Methods* 56:3102–3116.

Hillygus, D. Sunshine, Natalie Jackson and McKenzie Young. 2014. Professional Respondents in Nonprobability Online Panels. In *Online Panel Research: A Data Quality Perspective*, ed. Mario Callegaro, Reg Baker, Jelke Bethlehem, Anja S. Göritz, Jon A. Krosnick and Paul J. Lavrakas. Wiley pp. 219–237.

Hoffman, Matthew D. and Andrew Gelman. 2014. "The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo." *The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo* 15:1593–1623.

Lippman, Walter. 1922. *Public Opinion*. New York: Harcourt and Brace.

Motta, Matthew P., Timothy H. Callaghan and Brianna Smith. 2017. "Looking for Answers: Identifying Search Behavior and Improving Knowledge-Based Data Quality in Online Surveys." *International Journal of Public Opinion Research* 29(4):575–603.

Neal, Radford M. 2011. MCMC Using Hamiltonian Dynamics. In *Handbook of Markov Chain Monte Carlo*, ed. Steve Brooks, Andrew Gelman, Galin Jones and Xiao-Li Meng. CRC Press pp. 113–162.

Ratcliff, Roger. 1978. "A Theory of Memory Retrieval." *Psychological Review* 85(2):59–108.

Ratcliff, Roger and Philip L. Smith. 2004. "A Comparison of Sequential Sampling Models for Two-Choice Reaction Time." *Psychological Review* 111(2):333–367.

Stagnaro, Michael Nicholas, James Druckman, Adam J. Berinsky, Antonio Alonso Arechar, Robb Willer and David Rand. 2026. "Representativeness versus Response Quality: Assessing Nine Opt-in Online Survey Samples." *Nature Human Behavior* 10:1441–1454.

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.

Online Appendix to

Measuring Political Attitudes and Attributes with Drift Diffusion Models

Contents

Study Information

This section provides additional information about the two data sources used in this study. For both data sources, as with all survey research, the design and collection of data has limitations, and resulting estimates may involve unmeasured error that limits representativeness to the target population.

2024 Lucid Survey

The data for the 2024 Lucid survey constitute a non-probability convenience sample of the U.S. general adult population recruited from the Lucid Marketplace respondent pool via quota sampling. Lucid is a platform which links researchers with over 250 sample providers. Providers direct their panelists to Lucid's Marketplace and Lucid directs these panelists to available surveys for which they qualify.

The 2024 Lucid survey was conducted by the authors and funded through internal University research funds. The survey was fielded from June 1-6$^{\text{th}}$, 2024, via the Qualtrics online survey platform. The study was approved by Duke University's Institutional Review Board under protocol 2024-0473. Respondents received an incentive (i.e. cash or similar equivalents) to participate in our survey, but researchers are not provided specific information about these incentives. According to Lucid's ESOMAR 28, "All incentives are handled by our supply partners. Each incentive program is unique. Incentives come in the form of cash, gift cards, loyalty rewards, and charitable donations."

A total of $1,883$ individuals entered the survey. Along with providing consent to participate in the study, participants were screened for eligibility in several ways. Respondents were required to disable any virtual private network (VPN) in active use, allowing geolocation to ensure that participants were located in the United States. Respondents were also required to successfully pass a Captcha check and an explicitly labeled attention check before proceeding to the survey. Additionally, we removed respondents who indicated that they were less than 18 years of age, resided outside the United States, critically failed a subsequent explicitly labeled attention check, exhibited extreme speeding (completion in less than one third of the median time), or failed to reach the end of the survey. We also removed respondents who failed at least two of the following quality checks: self-reported age and birth year do not correspond, within a tolerance of +/- 2 years; self-reported state of residence and zip code do not match; non-sequitur or item non-response to an open-ended question about preferred news source; speeding, defined by completion in less than half of the median time; scoring less than 0.65 on Qualtrics' internal reCaptcha measure; or partially failing the second attention check. These exclusions reduce the sample to $N = 1,638$ respondents. Finally, for the purposes of this analysis, we exclude participants who refused to commit to answering our political knowledge items without outside assistance, or who correctly answered an obscure "catch" question. These restrictions reduce our final analysis sample to $N = 1,531$ respondents.

2024 ANES

The American National Election Studies (ANES) conducts two-wave panel surveys of US adult citizens in each presidential election cycle, with one pre-election wave and one post-election wave. Although the ANES sometimes includes additional sampling frames (such as sampling from prior ANES respondents to generate a multi-year panel dataset), in each election cycle the ANES generally seeks a fresh cross-sectional sample of the target population via address-based sampling. Additional information on ANES data and methodology is available at electionstudies.org.

For this analysis, we restrict our analysis to web-based participants in the 2024 ANES pre-election survey ($N = 4,234$), for whom item-level response timing information is available. We additionally exclude participants who refused to commit to answering the political knowledge items without outside assistance, or who correctly answered an obscure "catch" question. These restrictions reduce our final analysis sample to $N = 3,731$ respondents.

Survey Questionnaires

This section provides the exact question wording and response options of the political knowledge items used on each survey.

2024 Lucid Survey

Note: for all multiple choice questions below, the order of the response options was randomized, unless there was a clear logical ordering to the options (i.e., length of the Senate term, proportion required to override a veto, number of Republican-appoint justices). For questions regarding party control, the major parties were randomly positioned first, followed by "Neither." In all cases the "I don't know" option always appeared last. All 15 analysis items were presented in a random order, book-ended by the pledge question and the catch question.

Next, we're going to ask you some questions about current affairs. It is important to us that you do NOT use outside sources like the Internet to search for the correct answer. Will you answer the following questions <u>without help</u> from outside sources?

- Yes
- No

Who is the current Speaker of the U.S. House of Representatives?

- Nancy Pelosi
- Marco Rubio
- Mitch McConnell
- Mike Johnson
- I don't know

How long is the term of office for a U.S. Senator?

- 2 years
- 4 years
- 6 years
- 8 years
- I don't know

What job or political office is now held by Rishi Sunak?

- Prime Minister of the United Kingdom
- Secretary of the U.S. Treasury
- U.S. Senate Minority Leader
- Prime Minister of India
- I don't know

Who currently holds the title "President of the U.S. Senate"?

- Kamala Harris
- Mike Pence
- Mitch McConnell

- Elizabeth Warren
- I don't know

Who is the current Chairperson of the U.S. Federal Reserve?
- Ben Bernanke
- Jerome Powell
- Janet Yellen
- Paul Krugman
- I don't know

What proportion of each chamber of the U.S. Congress is required to overturn a Presidential veto?
- One-half
- Two-thirds
- Three-quarters
- Four-fifths
- I don't know

Is the Governor [DC only: Mayor] of [R's state of residence] a member of the Democratic Party, the Republic Party, or neither?
- Democratic Party
- Republican Party
- Neither
- I don't know

Which political party currently controls the most seats in the [lower chamber of R's state legislature]?
- Democratic Party
- Republican Party
- Neither
- I don't know

In an election for U.S. president, which of the following chooses the next president if the vote in the Electoral College is tied?
- The presidential cabinet
- The U.S. House of Representatives
- The U.S. Senate
- A convention of the states
- I don't know

Wich of the following rights is guaranteed by the <u>First</u> Amendment to the U.S. Constitution?
- The right to freedom of religion
- The right to bear arms
- The right to privacy
- The right to remain silent

- I don't know

In which part of the federal government is there a rule that requires a 60% majority vote to end a filibuster?

- The U.S. House of Representatives
- The U.S. Senate
- The U.S. Supreme Court
- The presidential cabinet
- I don't know

When a justice is confirmed to the U.S. Supreme Court, how long are they appointed for?

- 18 years
- Until a president decides to replace them
- Until they are 75 years old
- It is a lifetime appointment
- I don't know

Of the nine (9) justices currently appointed to the U.S. Supreme Court, how many were appointed by a Republican president?

- 3
- 4
- 5
- 6
- I don't know

What are the first Ten Amendments to the U.S. Constitution called?

- The Bill of Rights
- The Federalist Papers
- The Declaration of Independence
- The Articles of Confederation
- I don't know

In what year did the U.S. Supreme Court decide the case *Hill v. Wallace*?

- 1851
- 1877
- 1912
- 1922
- 1938
- 1964
- 1975
- 2002
- I don't know

2024 ANES

Note: the order of the response options was randomized for the "spending" question below; all others used a fixed order when presenting closed-ended options. All items were presented in sequential order as shown below.

We are interested in the guesses people make when they do not know the answer to a question. We will ask you several questions. Some may be easy, but others are meant to be so difficult that you will have to guess. In fact, for some of these questions, if you answer correctly, we will know that you probably looked up the answer. Please do not look up the answers you do not know. Instead, please just make your best guess. Will you please promise to try your best without looking up any answers? Or do you not want to make that promise?

- I promise to try my best without looking up any answers
- I do not want to make that promise

In what year did the Supreme Court of the United States decide Halstead v. Grinnan? Type the year.

- (Numeric entry)

Note: if respondents entered "1894," they were asked "You are right! Did you look up the answer to that question, or did you already know it yourself?" and told "Please do not look up the answers yet! For the next questions, please make your best guess without any help. After you do that, we will show you all the right answers." We exclude all respondents with entered "1894" or answered "I do not want to make that promise" in the preceding question.
For how many years is a United States Senator elected – that is, how many years are there in one full term of office for a U.S. Senator? Type the number.

- (Numeric entry)

On which of the following does the U.S. federal government currently spend the least?

- Foreign aid
- Medicare
- National defense
- Social Security

Do you happen to know which party currently has the most members in the U.S. House of Representatives in Washington?

- Democrats
- Republicans

Do you happen to know which party currently has the most members in the U.S. Senate?

- Democrats
- Republicans