

Vote Compass Methodology

1 Introduction

Vote Compass[®] is a voter engagement application developed by the team of social and data scientists from Vox Pop Labs. Its objective is to promote electoral literacy and public participation during election campaigns. Based on a user's responses to a series of propositions that reflect salient aspects of the campaign discourse, Vote Compass calculates the alignment between the user's personal views and the positions of the political parties. Party positions are determined by way of a comprehensive review of the public statements made by party officials on the topics included in Vote Compass. Each of the parties included in Vote Compass is directly consulted throughout this process and invited on multiple occasions to review the findings and provide feedback.

Vote Compass results are not intended and should not be interpreted as voting advice, nor as a prediction as to which party or candidate a given user intends to vote for. It is meant to serve as an entry point into an examination of parties differ across a suite of issues relevant to a given election campaign.

2 Party or Candidate Positions per Question

Development of the Vote Compass questionnaire follows a two-part research process. First, a content analysis is performed on the policy issues that figure most prominently in the platforms as well as public statements of the major political parties or candidates, and in the general media discourse. From an initial list of questions, we select those to be included in the final questionnaire based on the questions' ability to differentiate between parties or candidates and amongst voters; their breadth of coverage across multiple policy fields; and their salience in the upcoming election.

Second, party or candidate positions in the Vote Compass questionnaire are derived from the parties' or candidates' publicly-available statements. The Vote Compass research team undertakes a comprehensive review of party or candidate documents, including manifestos, election platforms, websites, speeches, press releases, legislative debates, and statements to media, in order to impute an accurate representation of parties' or candidates' stances on the policy issues explored in Vote Compass.

Preference is accorded to public statements that are recent; come from either the parties or candidates; and are directly relevant to the policy issue in question. Specifically, public statements are prioritized by date in the following order:

1. Election platforms
2. Official policy documents
3. Statements from the candidates/press releases from the party or candidate
4. Statements from Cabinet Ministers or party critics for the policy domain in question (Parliament transcripts, speeches, media etc.)
5. Statements from other elected party or candidate representatives
6. Party constitution; member-passed resolutions

Within these guidelines, allowances may be made for statements that most closely represent a party's or candidate's position on the exact phrasing of a particular Vote Compass proposition. This calibration process is followed by a consultation with the parties and/or candidates themselves. These two steps are described more in detail below.

3 The Calibration Process

Based on the collected public statements, researchers from the Vote Compass team are assigned to code or calibrate a given party's or candidate's positions on each of the final questions included. To ensure inter-coder reliability, the researchers initially undertake this task separately and subsequently compare results for consistency. All response categories are presented as Likert-type (or rating) scales. Each level of policy position is based on explicit guidelines:

- **Strongly dis-/agree, much less/more, many fewer/more, much harder/easier**
The party clearly emphasizes the issue in question and does not place any conditions, qualifications, or restrictions on its position.
- **Somewhat dis-/agree, somewhat less/more, somewhat fewer/more, somewhat harder/easier**
The party does place conditions, qualifications, or restrictions on its position; or emphasizes only part of the proposition.
- **Neutral, about the same as now**
The party addresses the issue without consistent argumentation in favour or opposition; defers taking a position; and/or mentions the issue indirectly.

Calibrations on questions pertaining to taxes and spending are based on support for nominal change. In the event that a party or candidate supports an increase/decrease in taxes or spending that was passed in a prior sitting of the legislature but has yet to come into effect, this is still considered support for a nominal change. To ensure that the results of this process are transparent for users, all party and candidate positions and supporting public statements (with URLs) are made available in the Vote Compass tool under “You vs. Party” and “Party vs. Party” on the results page. This information enables users to compare their own responses to those of the parties or candidates, and to delve deeply into party or candidate platforms and public documents.

4 Consultation with the Parties and Candidates

In a next step, Vote Compass consults with the political parties or candidates themselves. Parties or candidates are first sent a copy of the Vote Compass questionnaire, and invited to position themselves on the initial list of questions. Upon receipt of a completed questionnaire, Vote Compass reconciles the party’s or candidate’s self-placements with the calibrations determined by the research team coders. In the majority of cases, the calibrations from the party or candidate and the Vote Compass research team are in agreement. If discrepancies exist, Vote Compass sends the party or candidate a reconciliation report outlining the confirmed calibrations and the disputed ones across the final Vote Compass questionnaire. All discrepancies are flagged and justified with the party’s or candidate’s public statement collected by the research team which supports the calibration proposed by Vote Compass.

The party or candidate is able to respond to each disputed calibration by clarifying its position and providing alternate public statements which support its self-placement on the issue in question. In cases where the party or candidate provides relevant policy statements which conclusively accord with its self-placement, Vote Compass will reposition the party or candidate on this issue. Where discrepancies are not resolved by this process, the disputed placements are sent for deliberation and final ruling to the Vote Compass Advisory Board, comprised of well-established scholars in electoral politics. Parties or candidates are then sent final calibrations for review. They are able to dispute these calibrations and supporting public statements throughout the entire run of Vote Compass. If a party’s or candidate’s stance on an issue changes or if a party or candidate wishes to provide additional official documentation not considered during the reconciliation process, we will revisit the appropriate calibration to determine if a change is warranted. Whatever the reason, we encourage parties and candidates to consult with us over the course of the election campaign if necessary. Every effort is made throughout the electoral campaign to ensure the accuracy of party and candidate calibrations based on their publicly available statements.

5 Vote Compass Algorithms and Visualizations

The Vote Compass results section is comprised of two main elements. Each of these elements is designed to help users understand their positions on the issues relative to the parties or candidates and to see more generally how they fit within the political landscape. The first of these is a multidimensional graph that presents the position of the user and the parties or candidates within an abstract political landscape. The second is a bar graph that displays a user’s level of agreement with each party or candidate across multiple issues.

The consequence of using multiple measures is that there will occasionally be disagreement between the party or candidate that appears “closest” in the political landscape and that which appears “closest” on the bar graph. One reason for this is that these graphs are representations of different concepts. It is also because there is no perfect measure of political distance, either ideologically or on individual issues. In a public tool of this nature, it is necessary to recognize the compromise between increased methodological sophistication and the ease with which a method can be understood by the public. The use of multiple measures admits as much. It is an acceptance of the reality that the political world — both among politicians and the public — is complex. But it is this complexity that makes politics so lively and contentious, and why successful policies and politics often require great imagination from the public and their political representatives.

The purpose of Vote Compass is therefore to encourage users to think through this complexity; to learn where parties and candidates stand on the issues and the reasons why they do so; and to raise the level and quality of political information among the public more generally. For these reasons we encourage all users to read through the accompanying documentation on the Results page that show how and why the parties and candidates were “coded” as they were on issues.

6 Representations of the Political Landscape

Users’ and parties’ placement on the graph representing the political landscape are determined by responses to Vote Compass’s attitudinal and policy-related questions. Each question is assigned to the ideological dimension(s) to which it fits best, both by theory and through analysis of survey data that we collect during the design process of the application.

6.1 Dimensionality Reduction with Factor Analysis

To represent the large number of attitudinal and policy-related questions in lower-dimensional space requires that we use a statistical dimension-reduction technique. The goal is to show users what their answers, in general, say about how they fit within an easy-to-understand political landscape. To do this, we use a statistical technique called factor analysis, which allows us to capture users’ and parties’ underlying positions on a small number of abstract political dimensions (called “factors”).

One can think of this as a position on a scale (for example, on social issues) which we cannot observe directly using a single question, but which we can measure by asking a number of questions that are connected to it. This works because people’s attitudes on one issue frequently say something about their positions on other issues. Latent-variable techniques permit analysts to uncover these positions by using many questions in concert and capitalizing on the relationships between them. Survey responses that are highly associated with the latent dimension therefore receive greater weight in determining a user’s position on that dimension than those that are only weakly associated with it. To determine these dimensions, we analyze data that are collected prior to the launch of the Vote Compass application.

In this paragraph, we lay out the model assumptions and main steps followed to derive the abstract political dimension. Denote p as the number of propositions. Let X be a vector of a user’s responses to p propositions where $X \in \mathbb{R}^p$. Assume that we can find Z is a vector of k latent variables (i.e. “factors”), which influence the users’ responses, such that $Z \in \mathbb{R}^k$. Then the relationship between X and Z can be expressed as follows:

$$X = \mu + \Lambda Z + \epsilon,$$

where $Z \sim N(0, 1)$, $\epsilon \sim N(0, \Psi)$ and $\Lambda \in \mathbb{R}^{p \times k}$ is a matrix of the factor loadings. Moreover, we also assume that $cov(X_i, Z_j) = 0$ for $i = 1, \dots, p$ and $j = 1, \dots, k$. From this follows for the distribution of $X|Z$

$$X|Z \sim N(\mu + \Lambda Z, \Psi)$$

with the parameter Ψ as the matrix $\Psi \in \mathbb{R}^{p \times p}$.

The loadings are estimated from this model using the maximum likelihood method, to which we apply the varimax rotation. We define each theoretical dimension on the basis of how well it loads on every proposition.

Once the pool of propositions for a specific dimension has been determined, we inductively project this subset of propositions into a single dimension. Each dimension can be obtained using a Bayesian ideal point model.

6.2 Bayesian Ordered Probit Model

The assumptions under the Bayesian approach focus on an individual’s position θ_j on the latent dimension in question. Implicitly, it represents respondent j ’s propensity to pick one specific answer from the ordinal choices. In the case of obtaining ideological estimates from a set of questions, θ_j can be expressed as

$$\theta_j = \beta + \epsilon_i, \text{ where } j = 1, \dots, J$$

with the respondent-specific error term $\epsilon_i \sim N(0,1)$.

β represents a vector of parameters, containing separate coefficients, or discrimination parameters, for each item i as β_1, \dots, i . θ_j remains unobservable. The relationship between θ_j and a respondent's observed answer y_{ji} can be determined with the help of $\kappa = 1, \dots, 5$, which index the ordered response options to each proposition. The probability that a respondent chose a specific category can be represented by an ordered probit model as

$$\begin{aligned}
 P(y_{ji} = 1) &= \Phi(\tau_{i,1} - \theta_{ji}) \\
 &\vdots \\
 P(y_{ji} = k) &= \Phi(\tau_{ik} - \theta_{ji}) - \Phi(\tau_{i,k-1} - \theta_{ji}) \\
 &\vdots \\
 P(y_{ji} = 5) &= 1 - \Phi(\tau_{i,5-1} - \theta_{ji})
 \end{aligned}$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function. τ_i is a vector of unobserved threshold values or "cut points" for i . The number of threshold values for each item is one less of the number of answer options on the ordered scale. Propositions in Vote Compass with their five-level Likert-type scales require four cut points ($\tau_{i,5-1} = \tau_{i,4}$).

6.3 Maximum A-Posteriori Estimates

For efficiency when calculating a user's placement on the scale in the application, Vote Compass uses point estimates for the parameters in the Bayesian ordered probit model to calculate maximum a-posteriori (MAP) estimates for the respondents' values of θ . These are effectively numerically equivalent to estimates obtained from a fully Bayesian model.

The necessary posterior distributions to produce MAP estimates for new users are derived from the pilot sample. We use "Stan" and the R-Program for statistical computing to fit an ordered probit model and identify its posteriors.

MAP estimations requires a log-likelihood function to estimate user positions θ_j . This log-likelihood function is:

$$\text{LogLik} = \sum_j \ln[P(y_{ij})] = \sum_j \ln[\Phi(\tau_{ik} - \theta_{ji}) - \Phi(\tau_{i,k-1} - \theta_{ji})]$$

7 Representations of the Issues

Users scores on the bar graph are calculated using the absolute distance of the user's positions on the issues from those of each party or candidate (also known as the Manhattan distance). For example, let us say that for a given question, there are 5 possible answers as follows, each of which are assigned a number (in parentheses below):

Answer choice:	Strongly Negative	Somewhat Negative	Neutral	Somewhat Positive	Strongly Positive
Answers:	...	User (2)	...	Party A (4)	Party B (5)

Given these hypothetical user and party positions, we calculate the user to be 2 units away from Party A (i.e. $|2 - 4|$) and 3 units away from Party B (i.e. $|2 - 5|$). The sum of these distances, as calculated for each issue, thus represents how far a user is from a given party overall. Because we prefer to measure the bar graph distance in terms of how close a user is to a party rather than how far, we subtract this distance from that which would be given to a party that is as far as possible from the user and then divide by this distance. If a party takes a position that is as far from the user as possible, the bar will read 0; if a party takes a position that exactly mirrors that of the user, the bar will read 100.

We also account for the fact that users frequently care about some issues more than others. We do this by allowing the user to indicate how much more strongly they care about some issue categories compared to the rest. Once a user has input these weights, the distances on the bar graph are weighted accordingly: the distances from a user to a party or candidate on the issues about which a user cares about strongly are magnified and those about which a user cares little are shrunk.

To be more specific, we are going to formalize the steps used to calculate the bar graph. Let n be the number of propositions, and x_{ij} and x_{ip} be a user's and a party's or a candidate's position to the i -th issue respectively. Hence, the absolute distance δ_{jp} of the user to the party or the candidate p can be expressed as follows:

$$\delta_{jp} = \sum_{i=1}^n |x_{ij} - x_{ip}|$$

Next, we calculate the maximum possible distance the party or the candidate can be from the user $max(\delta_{jp})$, given the user's responses (which are coded on a scale from 1 through 5):

$$max(\delta_{jp}) = \sum_{i=1}^n (|x_{ij} - 3| + 2)$$

This equation centres the scale and takes the absolute value of a user's response to determine its distance from the centre (=3). It adds 2, which is the maximum distance the party or the candidate can be from the centre. For example, if a user answers '1', a party position of '5' is the maximum distance $|5 - 1| = 4$ from the user. If a user answers '2', a party position of '5' is the maximum distance $|5 - 2| = 3$ from the user. If a user answers '3', a party position of '1' or '5' is the maximum distance $|1 - 3| = |5 - 3| = 2$ from the user, etc. This equation finds the sum of the maximum distances a party can be from the user on each of the questions.

The final agreement score between the party or the candidate p and the user j is

$$agreement_{jp} = \left(1 - \frac{\delta_{jp}}{max(\delta_{jp})}\right) \times 100$$

When select issues are weighted by the user, with issue weight ω_{ij} as $0 \leq \omega_{ij} \leq 10$, the disagreement δ_{jp} and the maximum possible distance to the user's position are determined via these adjusted formulas:

$$\delta_{weighted}^{jp} = \sum_{i=1}^n |\omega_{ij}x_{ij} - \omega_{ij}x_{ip}|$$

and

$$max(\delta_{weighted}^{jp}) = \sum_{i=1}^n (|\omega_{ij}x_{ij} - 3 \times \omega_{ij}| + 2 \times \omega_{ij})$$

Calculation of the weighted agreement score between the party or the candidate p and the user j is analogue to the equation for $agreement_{jp}$ above.