An Extension and Test of Converse's "Black-and-White" Model of Response Stability

Jennifer L. Hill *
Department of Statistics
Harvard University
1 Oxford St.
Cambridge, MA 02138
phone: (617) 495-5325
fax: (617) 496-8057
email: hill@stat.harvard.edu

Prof. Hanspeter Kriesi †
Department of Political Science
University of Geneva
UNI-MAIL, 102 bd Carl-Vogt
CH-1211 Geneva 4
phone: 011 (41 22) 705 83 81
fax: 011 (41 22) 705 83 64
email: hanspeter.kriesi@politic.unige.ch

December 2, 1999

Abstract

Popular theories in political science regarding opinion-changing behavior postulate existence of one or both of two broad categories of people: those who hold their opinions over time; and those that hold no solid opinion and, when asked to make a choice, do so seemingly at random. This study explores evidence for a third category: durable changers. This group of people will change their opinion in a rational, informed manner, after being exposed to new information. Survey data collected at four time points over nearly two years tracks Swiss citizens' readiness to support pollution-reduction policies. It is analyzed using finite mixture models which allow estimation of the percentage in the population of each category for each question. Specific behaviors attributable to each of the three groups can be explicitly measured. Well-established statistical methodology (EM and Data Augmentation) is implemented to fit the models. Results provide evidence for existence of the durable-changer group.

1 Introduction

The stability of individual opinions about government policies has been a hotly debated issue among social scientists for many decades. Panel surveys have typically found a great deal of instability with respect to individual level opinions between successive interviews. In the American debate, this instability usually has not been attributed to genuine change of opinion, but has been interpreted as some kind of chance variation. The debate has focused on the type of instability and the reasons why there is so much of it. Two opposing approaches have been proposed.

*Will receive correspondence.

†The authors gratefully acknowledge funding for this project partially provided by the Swiss National Science Foundation (project 5001-035302). Thanks are due Don Rubin for his help in the initial formulation of the model, as well as John Barnard and David van Dyk for helpful comments along the way. Klaus Scherer should be acknowledged for his role in organizing the first Swiss summer school for the Social Sciences where this collaborative effort was born.

On the one hand, Philippe Converse's approach maintains that this instability is indicative of widespread "*non-attitudes*." In his seminal paper on the "Nature of belief systems in mass publics", Converse (1964) claims that many or even most people do not have any real opinions at all about what government should do in various policy areas; when survey interviewers come to the door, these people just give meaningless responses. Accordingly, large parts of the citizens are not apt for democracy: they may be susceptible to political manipulations of all kinds and their political choices may be singularly ill-considered.

Converse's point of view has been strongly challenged by the *"measurement error"* approach (see Achen 1975), which attributes the chance variations in the individuals' responses not to the lack of individual opinions, but to the inherent difficulty of survey questions in measuring them correctly. Compared to the unsettling consequences of Converse's interpretation, the implications of this approach are quite attractive. According to the measurement error approach, we can put our trust in the citizens. They have real, stable opinions. Rather than the citizens, it is our measurement instruments which are unreliable. But, this approach maintains, their deficiencies are easily remedied by standard psychometric techniques.

In Converse's *"black and white"* model (Converse 1964), there are just two groups of individuals – a perfectly stable group, and a random group: "There is first a 'hard core' of opinion on a given issue, which is well crystallized and perfectly stable over time. For the remainder of the population, response sequences over time are statistically random." Converse was led to this simple model by an inspection of his panel data from the American elections in the late fifties: surprisingly, the correlations between the opinions given at the initial interview in 1956 and the terminal interview in 1960 were of the same order of magnitude as the correlations between opinions given at each one of these moments, and opinions measured at the interview in 1958. In other words, one could just as easily predict 1960 opinions on most of the issues on the basis of individual opinions in 1956 alone as one could with a knowledge of the more proximal 1958 responses. It was this same surprising stability of correlations which Achen (1975) interpreted as a result of measurement error. Achen argued that correlations will be equal across time periods, when respondents are stable in their views and all observed variability is measurement error. This interpretation of observed instability suggests that there is just one group of citizens – everybody has basically stable opinions.

Zaller's claim (Zaller 1992) that most people's opinions on most issues are characterized by a tendency toward some degree of *ambivalence* marks an intermediary position in this debate. From his point of view, individual response instability occurs because the particular considerations that happen to be at the top of the individual's head at the time of the interview tend to vary from one occasion to the other. As he points out, this interpretation is not inconsistent with the "measurement error" model. One reason that individuals may bring different considerations to the evaluation of the same question may be that, as the "measurement error" tradition suggests, most questions are open to multiple interpretations. Another reason may be that the immediate context in the survey in which the question is posed – the order or the framing of the questions – may have an influence on the type of considerations that will be most salient in the mind of the respondents when asked to give their opinions. However, there is a fundamental difference between the two models. In contrast to the measurement error tradition, where response error is simply so much noise, in Zaller's model, "response variation is rooted in an important substantive phenomenon, namely the common existence of ambivalence in people's reactions to issues." The implication is that (p. 94) "even when people are temporally unstable, expressing completely opposing positions at different times, they may still ... be expressing real feelings, in the sense that they are responding to the issue as they see it at the moment of response. Although the perceptions of the issue may change over time, the responses they generate are not, for that reason, lacking in authenticity."

In his classical discussion of the "black and white" model, Converse (1964: 244) noted that some of the issues departed from this model "in modest degrees." He attributed these departures to the "presence of a 'third force' of people, who are undergoing meaningful conversion from one genuine opinion at t1 to an opposing but equally genuine opinion at t2." He considered this "third force” to be small, but sufficient to disrupt the fit between the data and the “black and white” model. In the subsequent discussions of the subject at stake, this “third force” disappeared, though there are reasons to believe that it may be more or less important, depending on the political events taking place in a given period.

For instance, the recent study of Sears and Valentino (1997) indicates that presidential campaigns during pre-adult socialization may have lasting effects on specific individual opinions. Of course, pre-adult socialization has long been known for its lasting effects on opinions. What is remarkable, however, is the observation that longstanding predispositions tend to be socialized episodically rather than incrementally and that specific political events play a critical role in forming new lasting individual opinions. Next, the experiences associated with great national crises, such as the Civil War and the Great Depression, have proved to be decisive forces in the shaping of lasting party identifications (Campbell *et al.* 1964, 89-92). If the impact of such events appears to have been most strongly felt by the youth, the economically underprivileged and the minority groups reacted to it quite strongly, too. Third, major catastrophes, such as the accident at Three Mile Island, also lead to major and lasting opinion changes (Baumgartner and Jones 1993, 79f.). Finally, it is at least conceivable that an intense public debate on a given policy issue may induce people to adopt a lasting new opinion on a given issue.

In this paper, we propose to go back to Converse’s distinction between “opinion holders” and “random changers”. We believe that this distinction has been too quickly abandoned. Instead of discarding the “random changers” as people with “non-attitudes”, we wish to insist with Zaller that their responses to the interview questions may be entirely meaningful to them at the time of the interview, even if, at first sight, they appear to be random to an outside observer. Their responses may, however, not only be a result of the more or less unconscious ambivalence described by Zaller. They may also result from an underlying inner conflict that cannot be resolved in the interview situation – what Alvarez and Brehm (1995, 1997) call ambivalence; or they may be the result of uncertainty due to a lack of information; or they may result from a combination of ambivalence and uncertainty (van Steenbergen *et al.* 1997). We would like to call this second group *"vacillating changers"*. In addition, we propose to add to this distinction the "third force" – the group of "durable changers". The members of this last group change their opinion just once over a more or less extended period of time.

The problem is, of course, that membership in any one of these three groups is unknown, that is, the people belonging to a given group are not in any way "tagged". *Group membership has to be inferred from the individual response behavior in a panel study.* In order to make these inferences, some straightforward assumptions about the response behavior of each group have to be introduced. Thus, Converse was unable to "tag" the members of his two groups: in his case, he attributed all the individuals who had changed sides to the "random group". However, they did not exhaust the total set of respondents following random response paths. By chance alone, some members of the "random group" successively may have given the same answers to an opinion question, which made them appear to belong to the group of "opinion holders", although in fact, they were answering in a purely random way. On the basis of the assumption that all opinion change is a purely random phenomenon, it was possible for Converse to calculate the aggregate size of the "random group", although not individual membership in this group, and to test the goodness of fit of the model against the real data.

Compared to Converse, our problem is considerably more serious: we do not only have to distinguish between two, but between three groups. In order to solve the problem of group identification, we propose to use a class of statistical models which have become known as *"finite mixture models"*; these will be discussed in greater detail in Section 3.2. For each individual and for the sample as a whole, this type of model allows to estimate the probabilities of belonging to each one of the three groups on the basis of a set of assumptions about the response behavior of the putative members of the three groups. For each one of the three groups, a separate probability model is specified describing with as few parameters as possible the response behavior of its members. The overall model (described in Section 3) can be conceived as a mixture of the three separate models. The statistical procedures employed to fit the model will be described in Sections 4.2 and 4.3. Section 5 presents results of the model fitting followed by model checks in Section 6. Concluding remarks will be given in Section 7.

2 The Data

The data come from a Swiss study on the readiness of the population to support policies aimed at curbing air pollution caused by private cars. This issue domain was chosen because it represents a crucial source of environmental pollution (traffic), it immediately concerns the individual citizens' private lives (regulation of the use of their cars), and because it was publicly debated in Switzerland at the time of the study. For six policy measures within this domain, the data comprise responses at four different time points. Each one of these measures was discussed more or less intensively at the time of the first interview, or had been the object of intense discussions before that date (December 1993). The six measures include:

- speed limits
- tax on CO_2 (implying a price increase for gas of about 10 cents/l)
- large price increase for gas (up to 2fr./l)
- promotion of electrical vehicles
- car-free zones
- parking restrictions

While all six policy measures refer to the same policy domain, they differ from each other with respect to two criteria in particular: their *familiarity* and *the degree of constraints they impose on individual behavior*. At the time of the interviews, speed limits, car free zones and the promotion of electrical vehicles were very familiar to the Swiss public, while parking restrictions and especially the proposed fiscal incentives (CO_2 tax on gas and further price increases for gas) were much less well known. As far as the second criterion is concerned, at that time fiscal incentives and speed limits were clearly considered to be more constraining by the Swiss public than the other three measures. Table 1 gives the following typology of measures.

Degree of constraints	Familiarity	
	High	Low
Low	Car-free zones Electrical vehicles	Parking restrictions
High	Speed Limits	Tax on CO_2 Gas price at 2fr/l

Table 1: Typology of policy measures

The interviews took place over a period of roughly two years: the second and third interviews were held in spring and early summer 1994, the last one in fall 1995. Between the first two interviews, a national popular vote was organized on several transportation issues related to the measures under discussion. Moreover, a large public debate on the CO_2 tax was launched by the government just before the second interview. The third interview intervened after the respondents had received a "choice questionnaire." This is a sophisticated survey tool invented by Willem Saris and his colleagues during the Dutch General Social Debate about the future of nuclear power in the Netherlands (Saris, Neijens, and de Ridder 1983; Neijens 1987; Neijens, Minkman, de Ridder, Saris, and Slot 1996). It provides respondents with both descriptive facts as well as information agreed upon by several independent experts about the issues in question; in addition, it provides them with a procedure based on decision theory to use this information efficiently in order to arrive at a "considered opinion." Between the last two interviews, no major issue-related events occurred, but half of the sample was provided with a series of articles containing additional information about the issue of the CO_2 tax.

The first two waves have complete responses from 1062 respondents. There exist missing data in the third and fourth time periods. Overall, complete data exist for 669 respondents. The number of complete observations varies slightly by question: the minimum is 669 for the car question; the maximum is 688 for the parking restrictions question. In the present study, complete case analysis is used for model fitting. Optimally, we would use a more principled approach to missing data (see, for instance, Little and Rubin 1987). Analyses comparing responses in the first and second time periods between the group who responded at all four time points and the group who failed to respond at one or both of the last two time points show relatively similar marginal distributions of responses. This is encouraging, though by no means sufficient, for the assumption of data missing completely at random needed for valid inferences in complete case analysis.

The present analysis focuses on one question at a time. The coding used for the responses Y for the i^{th} individual at the t^{th} time point is:

$$Y_{t,i} = \begin{cases} 1 & \text{for "strongly disagree"}, \\ 2 & \text{for "mildly disagree"}, \\ 3 & \text{for "no opinion"}, \\ 4 & \text{for "mildly agree"}, \\ 5 & \text{for "strongly agree"}, \end{cases}$$

The first three measures – speed limits and fiscal incentives – are quite constraining for the Swiss public and, therefore, get less support than the last three measures at any one point in time. This is illustrated by the marginal distributions presented in Table 2.

3 The Probability Model

This section describes the specification of the behavioral model we hypothesize to be true by a probability model. The description will be informal and intuitive. For a more formal, mathematical treatment of the model we refer the reader to the *Technical Appendix*. In presenting the model, we shall also specify a number of hypotheses about the values we expect to see for some of the parameters.

There are many ways to model longitudinal survey data. Time series models which postulate multinomial logit or probit models (reflecting the discretization of the true continuous "latent"

Speed limits	t1	t2	t3	t4
Strongly disagree	24.9	29.5	24.5	19.0
Disagree	23.4	26.2	26.0	27.5
No opinion	3.5	1.2	1.0	0.4
Agree	24.6	25.7	24.0	28.2
Strongly agree	23.7	17.4	24.5	24.9
Total	100.0%	100.0%	100.0%	100.0%

Tax on CO_2	t1	t2	t3	t4
Strongly disagree	20.5	25.4	18.3	18.5
Disagree	18.3	20.2	24.7	23.3
No opinion	6.4	3.1	3.0	2.8
Agree	34.3	32.0	30.7	32.1
Strongly agree	20.5	19.3	23.4	23.0
Total	100.0%	100.0%	100.0%	100.0%

Gas price increase	t1	t2	t3	t4
Strongly disagree	27.0	32.1	30.3	22.4
Disagree	23.5	27.6	25.4	29.4
No opinion	6.3	2.6	3.3	2.8
Agree	26.0	22.8	21.6	27.3
Strongly agree	17.2	14.8	19.4	18.2
Total	100.0%	100.0%	100.0%	100.0%

Electric vehicles	t1	t2	t3	t4
Strongly disagree	7.8	4.2	3.1	6.0
Disagree	14.8	13.5	14.7	12.4
No opinion	4.1	2.5	6.4	1.6
Agree	48.7	51.5	48.7	44.0
Strongly agree	24.6	28.3	27.2	36.0
Total	100.0%	100.0%	100.0%	100.0%

Car free zones	t1	t2	t3	t4
Strongly disagree	5.4	6.1	4.9	4.4
Disagree	9.6	13.8	13.0	12.5
No opinion	1.5	0.7	0.7	0.9
Agree	34.0	38.7	38.0	35.3
Strongly agree	49.6	40.7	43.3	46.9
Total	100.0%	100.0%	100.0%	100.0%

Parking restriction	t1	t2	t3	t4
Strongly disagree	13.5	13.8	8.2	8.6
Disagree	17.7	22.2	21.4	23.3
No opinion	3.3	1.5	3.1	1.6
Agree	34.9	39.0	45.5	33.7
Strongly agree	30.5	23.5	21.7	32.8
Total	100.0%	100.0%	100.0%	100.0%

Table 2: Marginal distribution of the six opinions at the four time points (n=688)

response) at each time point would be natural descriptive models for this data if our goal was to predict future responses. However, our goal is to *classify* different types of behavior. We therefore focus on the similarities in the different types of behaviors within an opinion-changing group rather than the individualities. While these behavioral similarities potentially could be captured through complex constraints, this approach is less straightforward given our goals and these constraints could be complicated to implement. Perhaps most importantly though, the parameters of this type of model would, for the most part, be less transparently relevant to the theory at hand. The model we propose is a direct probability representation of the theory we are trying to test.

Alternatively, we could focus on the product multinomial structure of the data – that is, we could envision a separate multinomial distribution (with constraints on the cell probabilities), where each cell corresponds to a different pattern of responses for an individual over the four time points, for each of the three behavior changing groups. The model we present is actually mathematically equivalent to the product multinomial model with a particular set of complicated constraints. We believe that our parameterization and its associated conceptual representation (see, for instance, the tree structure displayed in Figure 4.1) constitute a clearer mapping from the theoretical model to the probabilistic model.

3.1 Parameterization: The Full Model

Study participants can be characterized as belonging to one of the three groups described above with regard to each policy measure for the duration of the study period. This is an important qualification: the labels used to describe people are policy issue and time period dependent. For instance, an individual could be a durable changer regarding the CO_2 tax issue in the context of this study, but if a similar study were performed ten years later he might well be an opinion holder regarding the same issue.

Let π_j, $j = 1, 2, 3$, represent the marginal probability of belonging to each group, where opinion holders, vacillating changers and durable changers correspond to groups 1, 2, and 3, respectively.

These are the probabilities we eventually wish to estimate on the basis of the response patterns. The rest of the parameters are used primarily to characterize the response behavior of the three groups.

Hypothesis 1: We expect the size of the three groups to vary with the issue in question.

1.1 More specifically, the share of opinion holders is expected to be larger for familiar issues than for unfamiliar ones. This hypothesis is based on the idea that the crystallization of opinions depends on a durable stream of supportive information. Accordingly, policies most likely to give rise to crystallized opinions are those which have been the subject of extensive public debates in the past, those which have given rise to a policy monopoly and a corresponding policy consensus among the elites[1], and those which have been implemented by specific measures, the implications of which are well known to the respondents from personal experience. As Converse (1970, 177) has observed, “there is a very real sense in which attitudes take practice.”

1.2 However, the intensity and salience of the public campaign about the fiscal incentives during the first period of our study, as well as the intensity and salience of the experimental campaigns we introduced later in the study, were relatively limited compared to major political campaigns. Moreover, while the issue domain in question was clearly in a period of transition during the period of our study, it was not subject to a major crisis. The debate about the fiscal incentives was part of a longer process of policy change. The old policy monopoly was put into question, but the new policy consensus was slow to emerge. Thus, the debate about the fiscal incentives had not yet come to a conclusion in the course of our study. This

[1] Following the model presented by Baumgartner and Jones (1993), the policy process in a given issue domain may be conceptualized as a “punctuated equilibrium”: long periods of stability are interrupted by shorter intervals of major policy change. During extended periods, a limited number of actors develop policy in a given issue domain in an incremental fashion within established routines. These actors have a policy monopoly in their subsystem. Where such monopolies have been established, there tends to be a single understanding of the underlying policy question. Or, in other words, there tends to exist a “policy consensus.”

suggests that the number of durable changers will be rather limited: undecided public controversies foster ambivalence and uncertainty. As long as no new consensus emerges on the level of the policy elites, citizens will suspend judgment.

3.1.1 Opinion Holders

Opinion holders are defined as those who maintain an opinion either for or against an issue. Anyone who responded with a "no opinion" at any time point cannot be in this group, nor can anyone who crossed an "opinion boundary" across time points (i.e., an opinion holder cannot switch from an agree response to a disagree response, or vice-versa). Figure 1, which has boxes for each potential response over the four time periods, portrays several response patterns that could be made by an opinion holder. Boxes that an opinion holder can never occupy contain X's.

Two parameters are sufficient to describe the behavior of a member of this group:

$$
\begin{aligned}
\alpha_1 &= \Pr(\text{mildly disagree or strongly disagree}) \\
\delta_1 &= \Pr(\text{strongly disagree} \mid \text{mildly disagree or strongly disagree}) \\
&= \Pr(\text{strongly agree} \mid \text{mildly agree or strongly agree})
\end{aligned}
$$

Both of these probabilities implicitly condition on the fact that the individual whose behavior they describe is an opinion holder.

These parameters allow for differing probabilities of being for or against an issue, and conditional on being for or against an issue, differing probabilities of feeling strongly or mildly about it. Note that the parameter for the extremity or strength[2] of the reaction (δ_1) does not vary across time periods or across opinions (agree or disagree). We shall make the same assumption for the other two groups, although an inspection of Table 2 immediately shows that such an assumption approximately holds in only two out of the six cases. Thus, opposition to fiscal incentives is stronger than support, while the support of car-free zones is stronger than the respective opposition. Moreover,

[2] Extremity is one of many indicators of opinion or attitude strength (see Krosnick and Fabrigar 1995).

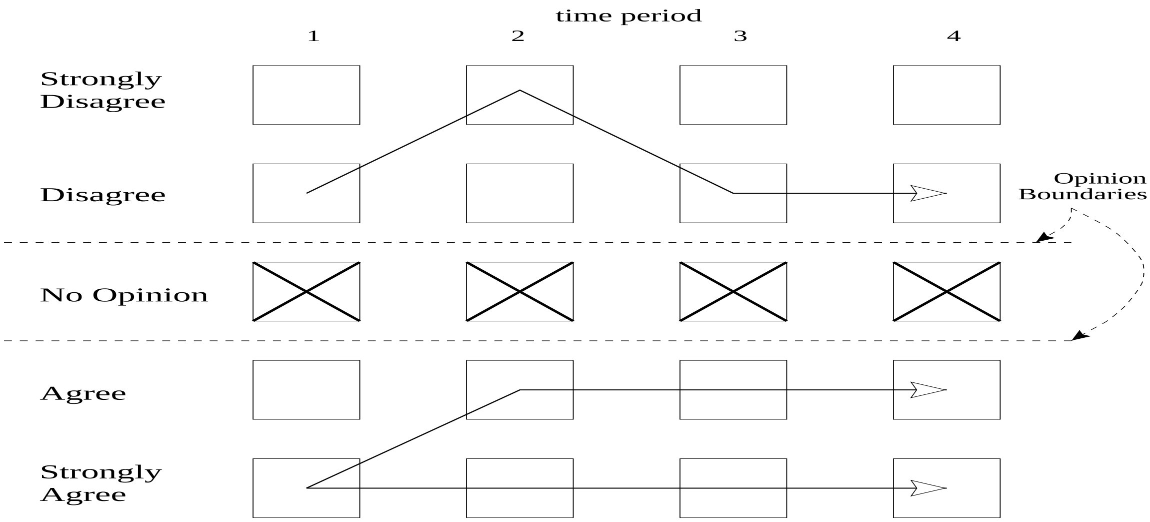

Figure 1: Potential Series of Responses for Opinion Holders

opposition to parking restrictions has become stronger in the course of the study, while the respective strength of support has fluctuated. Only in the cases of speed limits and electrical vehicles does the share of extreme opinions appear to be similar and rather stable on both sides of the opinion divide. Given that relatively little is known about the processes by which opinion strength develops (Eagly and Chaiken 1993, 681), we nevertheless prefer to make this simplifying assumption to keep the number of parameters to be estimated as small as possible.

Hypothesis 2: Among opinion holders, the share of opponents, α_1, is expected to be larger for the more constraining policy measures than for the less constraining ones. The more constraining measures impose higher costs on the individuals, which means that they will give rise to greater resistance.

3.1.2 Vacillating Changers

For the purposes of this model, we make again a simplifying assumption: the members of the group we label vacillating changers do not change their opinions in any particularly systematic way. Moreover, the responses of the members of this group are considered to be independent across time points. Therefore *any* pattern of responses could characterize a vacillating changer, for a total of 625 possible response patterns. Figure 2 illustrates just a few. These assumptions do not take into account the possibility that ambivalent individuals may change their opinion in non-random ways as a function of the contextual stimuli, such as a public debate on the issue in question.

On the basis of this assumption, the behavior of a vacillating changer at any time point can be characterized by only two parameters:

$$\begin{aligned}\varphi_2 &= \Pr(\text{no opinion}) \\ \delta_2 &= \Pr(\text{strongly disagree} \mid \text{mildly disagree or strongly disagree}) \\ &= \Pr(\text{strongly agree} \mid \text{mildly agree or strongly agree})\end{aligned}$$

We also impose the constraint

$$\frac{1-\varphi_2}{2} = \Pr(\text{mildly agree or strongly agree}) = \Pr(\text{mildly disagree or strongly disagree}).$$

Our model allows vacillating changers to have some minimal structure in their responses. They are allowed different probabilities for having no opinion, and the model also allows them to have different probabilities (constant over time) for extreme versus mild responses given that they express an opinion. However, in line with Converses black-and-white model, our model postulates that their probability to agree is identical to their probability to disagree. In other words, for those

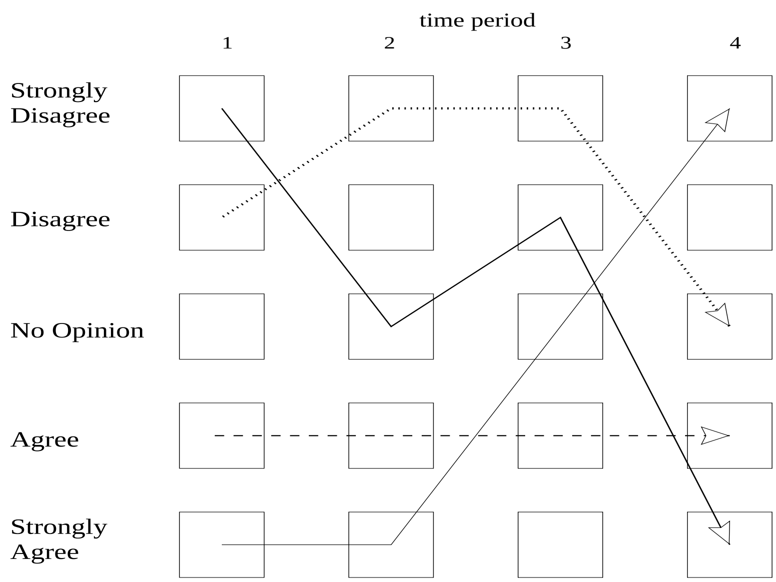

Figure 2: Potential Series of Responses for Random Changers

among the vacillating changers who voice an opinion, the probability to be for or against the issue in question is supposed to be completely random. This parameterization not only implies that we ignore contextual effects, it also implies that we do not account for acquiescence – the tendency of people with "non-attitudes" to agree with the interviewer. As far as the "no opinion" response is concerned, there are a variety of reasons for the members of this group to give such a response. Krosnick and Fabrigar (1995) mention genuine lack of opinion, ambivalence, question ambiguity, satisficing, intimidation, and self-protection. The point is that these reasons may be rather different from the considerations which lead the members of this group to choose one of the four substantive opinion categories. Note that, although these probabilities are constant over time, this does not mean that an individual will give the same response over time. In fact, responses across time periods are considered to be independent of each other for this group.

Hypothesis 3.1: Based on the arguments of Zaller which we discussed in the introduction, we would expect the assumption of equal probabilities for agreeing and disagreeing among vacillating changers to be too constraining. Instead, at any given time the vacillating changers may be expected to have some relevant issue-specific considerations, too. Accordingly, just as for the opinion holders, the share of opponents among vacillating changers, α_2 (see below), is expected to deviate significantly from the share of supporters and to be larger for the more constraining policy measures than for the less constraining ones. But the differences are expected to be less pronounced than for the opinion holders, given that the vacillating changers tend to be more ambivalent and uncertain about the issues in question than the opinion holders.

Hypothesis 3.2: We expect the share of strong opinions among vacillating changers (δ_2) to be lower than among opinion holders (δ_1). Opinion holders are expected to base their opinion on a stable set of issue-specific arguments, or "considerations", in Zaller's terms (Zaller 1992, 40), which quite explicitly favor one of the respective camps over the other. This means that they are likely to have strong opinions, whether they are in favor of or against the issue. Converse (1964) had already suggested that strong opinions are more resistant to change, and Eagly and Chaiken (1993, 580ff.) and Krosnick and Fabrigar (1995, 51f.) review a formidable array of evidence in support of this proposition. Vacillating changers, by contrast, are expected to have mixed considerations about a given issue, i.e. some considerations in favor of the issue, and some against it. This means that they will be more ambivalent about it, their opinions will generally be weaker and, as a consequence, less stable.

3.1.3 Durable Changers

Durable changers are defined as those who change their opinion or who form an opinion based on some rational decision-making process perhaps prompted by additional information or further consideration of an issue. Durable changers are allowed to change their opinion (i.e. to cross an opinion boundary as displayed in Figure 3) exactly once across the four time periods. This characteristic distinguishes them from the vacillating changers that are allowed to move back and forth, precisely as a result of their ambivalence or uncertainty. In contrast to vacillating changers, the group of durable changers is composed of individuals who adopt a new, stable opinion, either by changing sides or by forming an opinion for the first time. They are not allowed to change to the "no opinion" position, but they can move out of this position. This implies that those switching from a for position must switch to an against position and vice-versa. Possible response patterns are exhibited in Figure 3.

Durable changers have a bit more structure to their behavior than the other groups which is reflected in the increased number of parameters. First, we distinguish between a change of sides and a move out of the "no opinion" position. This is the distinction between durably changing an opinion and forming a durable opinion for the first time. Second, we distinguish between the probability that a switch of sides or a movement out of the "no opinion" position occurs after the first period, and the probability that it occurs after the second or third period. It is well known that respondents in a panel survey are clarifying their ideas in the course of the study and, as they reflect and learn more about the subject matter, get closer to their "true opinion" (McGuire 1960; Jagodzinski *et al.* 1987; Saris and van den Putte 1987). This panel or "Socratic effect" typically occurs between the first two waves of a panel study. Accordingly, we allow the probability of durable changes to be different for the first period, while we assume the respective probabilities for the two following periods to be identical. Third, we allow the direction of change to differ between the first period, on the one hand, and the second and third period, on the other hand. Since durable

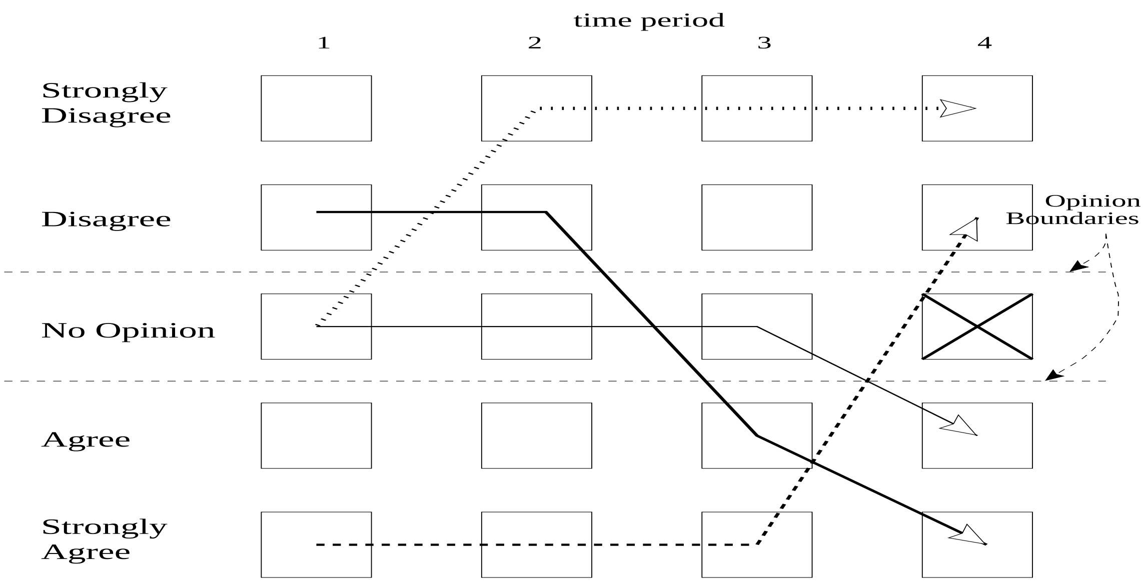

Figure 3: Potential Series of Responses for Durable Changers

changers are assumed to be strongly influenced by the additional information which they receive, the direction of their change is a function of the "tone" of the public debate. With respect to the fiscal incentives in particular, but not necessarily with regard to the other measures, this tone was rather negative during the first period, while it became, on balance, much more balanced for all issues during the following two periods. Fourth, as for the other two groups, we again allow for a stable share of strong opinions. These assumptions translate into following set of parameters:

$$\varphi_3^{(\text{pre1})} = \Pr(\text{pre-switch response is no opinion} \mid \text{switch occurs after 1st time period})$$
$$\varphi_3^{(\text{pre2})} = \Pr(\text{pre-switch response is no opinion} \mid \text{switch occurs after 2nd or 3rd time period})$$
$$\alpha_3^{(\text{pre1})} = \Pr(\text{pre-switch response is mildly disagree or strongly disagree} \mid \text{switch occurs}$$

after 1st time period)

$$\alpha_3^{\text{(pre2)}} = \Pr(\text{pre-switch response is mildly disagree or strongly disagree} \mid \text{switch occurs after 2nd or 3rd time period})$$

$$\alpha_3^{\text{(post)}} = \Pr(\text{post-switch response is mildly disagree or strongly disagree} \mid \text{switched from no opinion})$$

$$\delta_3 = \Pr(\text{strongly disagree} \mid \text{mildly disagree or strongly disagree})$$
$$= \Pr(\text{strongly agree} \mid \text{mildly agree or strongly agree})$$

$$\tau_3 = \Pr(\text{switch occurs after first time period})$$

We also impose the following constraint:

$$\frac{1-\tau_3}{2} = \Pr(\text{switch occurs after second time period})$$
$$= \Pr(\text{switch occurs after third time period})$$

Since these probabilities condition on membership in the durable-changers group, and this group comprises only individuals who switch exactly once across the four time periods, the parametric descriptions of the behavior revolve primarily around descriptions of this opinion switch. Note that only one parameter need be specified for post-switch behavior because we don't differentiate this behavior by switching time and because switches to no opinion are not allowed for this group.

3.2 The Model as a Tree

It is helpful to think of the model for the behavior of these three groups as being represented by a tree structure such as the one illustrated in Figure 4.1. This tree slightly oversimplifies the representation of our model. For instance the vacillating changer branch of the tree represents the response for just one given time period. However, it does reflect the types of behavior that are pertinent for defining each group.

[Figure 1 about here]

The type of statistical model that we use in this analysis is called a "finite mixture model" (for a description of these models and their properties, see any of Everitt and Hand 1981; Titterington, Smith, and Makov 1985; Lindsay 1995) because it can be conceived as a mixture of three separate models where the "mixing proportions" are unknown. In this case the full model is a mixture of the model for each opinion-changing behavior group. It is a mixture because, in general, we cannot deterministically separate one group or model from the next since the members of the different groups cannot be identified as such. Our model will be better behaved than some mixture models however because of the structure placed on the behavior of each group. In particular there are certain individuals whose behavior can only be accommodated by the vacillating changers group[3] Recent examples of fully-Bayesian analyses of mixture model applications include Gelman and King (1990), Turner and West (1993), and Belin and Rubin (1995).

4 Fitting the model

4.1 Re-Expressing the Data

This tree structure illustrates the types of behavior that we need to measure. Rather than defining a survey respondent by her response pattern, for example "1255", we need to characterize her in terms of a limited number of more general variables, which allow us to reproduce her trajectory, e.g. as someone who started out opposed to the issue (first extremely, "1", then not, "2"), and then crossed an opinion boundary and expressed strong agreement with the issue, "55". Therefore all of the data have been re-expressed in terms of the variables described below[4]; these variables, along with group indicators, define the elements which are used in the parameter estimates (note that t^* denotes the time period directly prior to a change in opinion). That is, while the model parameters represent the probabilities of certain types of behavior, the transformed data measure incidence of

[3]This identification creates a non-symmetric parameter space which prevents label-switching in the algorithm used to fit this model.

[4]It is computationally impossible to fit the model as we've described it without re-expressing the data. The only other option would be to fit a model (the product multinomial model described earlier) to the un-transformed data which has a different parameter for each pattern of responses (e.g. 3345) for each opinion-changing behavior group (for a total of 1875 different parameters) and then to reduce the effective number of parameters by constraining certain cells to be equal (e.g. for opinion holders 1222 is equally likely as 5444).

these same types of behavior.

$$A_i = \begin{cases} 1 & \text{if the } i^{th} \text{ person's initial response is a 4 or 5} \\ 0 & \text{otherwise} \end{cases}$$

$$B_i = \text{the number of the } i^{th} \text{ individual's responses that are either 1 or 5 across all } t$$

$$C_i = \text{number of the } i^{th} \text{ individual's responses that are 3}$$

$$D_i = \text{number of times the } i^{th} \text{ individual crosses an opinion boundary}$$

$$E_i = \begin{cases} 0 & \text{if } D_i \neq 1 \\ t^* & \text{otherwise} \end{cases}$$

$$F_i = \begin{cases} 0 & \text{if the } i^{th} \text{ individual's pre-switch response is a 1, 2, or 3, or } D_i \neq 1 \\ 1 & \text{if the } i^{th} \text{ individual's pre-switch response is a 4 or 5} \end{cases}$$

$$H_i = \begin{cases} 0 & \text{if the } i^{th} \text{ individual's pre-switch response is a 1, 2, 4, or 5, or } D_i \neq 1 \\ 1 & \text{if the } i^{th} \text{ individual's pre-switch response is a 3} \end{cases}$$

$$M_i = \begin{cases} 0 & \text{if the } i^{th} \text{ individual's post-switch response is a 1, 2, or 3, or } D_i \neq 1 \\ 1 & \text{if the } i^{th} \text{ individual's post-switch response is a 4 or 5} \end{cases}$$

$$Q_i = \begin{cases} 0 & \text{if the } i^{th} \text{ individual's post-switch response is a 1, 2, 4, or 5, or } D_i \neq 1 \\ 1 & \text{if the } i^{th} \text{ individual's post-switch response is a 3} \end{cases}$$

The vector of all of these random variables will be denoted $\mathbf{Z}_i$.

4.2 A Maximum Likelihood Algorithm – EM

The problem with using the above model to make inferences is that it relies on knowledge of group membership, which, in practice, we don't have. The EM algorithm is a method which can be used to compute maximum likelihood estimates in the presence of missing data. It is able to side-step the fact that we don't have group membership indicators, by focusing on the fact that if we had observed this "missing" data, the problem would be simple. It is an iterative algorithm with two steps: one that "fills in" the missing data, the *E-step* (expectation step); and one that estimates parameters using both the observed and filled-in data, the *M-step* (maximization step). We iterate between these two steps until an accepted definition of convergence (see Technical Appendix) is reached[5]. Starting values of parameters for the first iteration were chosen at random. Checks were performed to ensure that the same maximum likelihood estimates for each model were reached given a wide variety of (100) different randomly chosen starting values.

4.2.1 The E-step

The E-step for this model replaces the missing data (i.e. group membership indicators) with their "**e**xpected values". This means that for each person in the study, the relative probability of being in each group is calculated using conditional models which rely on the parameters described above and the behavioral data measured for that person (see Technical Appendix). Parameter values are taken from the previous iteration's M-step (starting values for the first iteration can be chosen randomly). These probabilities, denoted as $g_{k,i}$ for individual i and group k, represent our "expectation" regarding whether person i belongs to group k. They are used as weights in the M-step.

4.2.2 The M-step

The M-step finds **m**aximum likelihood estimates for all of our parameters. All of these parameter estimates are quite intuitive. For instance, the estimate for φ_2 is,

$$\hat{\varphi}_2 = \frac{\sum_{i=1}^{N} g_{2,i} C_i}{4T_2} .$$

where T_2 is the sum of the individual-specific weights, $g_{2,i}$ (calculated in the E-step), corresponding to the vacillating-changers group. This estimate takes the weighted sum of "no opinion" responses (C_i) across the four time periods (where each weight reflects the probability that the person is a vacillating changer) and divides it by the number of people we expect to belong to this group (the sum of the weights, T_2) multiplied by four (for the four time periods; four possible responses for

[5]One definition of convergence might be when a particular parameter (perhaps the one that changed the most from the previous iteration to the current) changes by less than a sufficiently small value (see, for instance, van Dyk, Meng, and Rubin 1995).

each person). This is the logical estimate for the parameters representing the probability that a vacillating changer will respond with "no opinion" at any given time point.

4.3 The Data Augmentation Algorithm

While point estimates of the parameters are helpful, they are insufficient to answer all the questions we might have about the parameters. For instance, it is useful to understand how much uncertainty there is about the parameter estimate so that we can perform hypothesis tests or create confidence intervals. One way to do this is to estimate the entire distribution of each parameter given the data we've observed; this distribution is called the *posterior* distribution. Posterior distributions formally combine the distribution of the data given unknown parameter values with a *prior* distribution on the parameters[6]. This prior distribution quantifies our beliefs about the parameter values before we see any data. The priors used in this analysis reflect our lack of *a priori* information about the parameter values and are thus "non-informative."

The posterior distribution can often be difficult to calculate or manipulate in closed form. Luckily algorithms exist that allow us to draw samples from these distributions. It is then trivial to use these samples to create accurate Monte Carlo estimates of confidence intervals or other statistics needed to make inferences. A useful algorithm for this problem is called the Data Augmentation (DA) algorithm[7] (Tanner and Wong 1987) and it has a particularly intuitive form; the intuition is parallel to that of the EM algorithm[8].

The DA algorithm has two basic steps in this problem:

[6] The formal mathematics to calculate these posterior distributions is straightforward and described in the Technical Appendix.

[7] For a general resource which provides a description of EM and the DA algorithm, as well as some interesting applications, refer to a text on Bayesian statistics such as Gelman, Carlin, Stern, and Rubin (1995). This text is also a useful reference for general Bayesian calculations and definitions.

[8] While EM provides only point estimates, it is still an integral part of the model-fitting process: it is generally easier to implement and more stable than DA thus it is a useful tool for generating starting values, exploring whether or not there are multiple modes, and even for checking the accuracy of DA code.

1. Sample the "missing data", group membership indicators, given the parameters.

2. Sample the parameters, $\theta = (\pi_1, \pi_2, \pi_3, \alpha_1, \delta_1, \dots)$, given the group membership indicators, from their posterior distribution.

Iterations continue until we have converged to a stationary distribution, the posterior distribution, which summarizes our knowledge about the parameters given the observed data. The parameter draws before convergence are treated as a "burn-in" period and discarded[9]. Then as many draws as are desired to sufficiently estimate the empirical distribution are taken.

The DA algorithm was used in this problem because non-Bayesian techniques have generally been found to be flawed when applied to mixture models, particularly when calculating standard errors. In addition the approximations which have been derived to accommodate testing of certain hypotheses are rather limited (see, for example, Titterington, Smith, and Makov 1985) and cannot approach the flexibility in the types of inferences that can be performed trivially once we can sample from the posterior distribution (for a discussion, see van Dyk and Protasso 1999). Assuming that the correct model is used, the DA algorithm will converge to the correct posterior distribution. The properties exhibited in our simulations lead us to believe that our DA algorithms had converged by the time we started saving draws from the posterior distribution.

4.3.1 Drawing Group Indicators Given Parameters

We can use the observed data for a given person along with parameter values to determine the probability that he falls in each group simply by plugging these values into the models we've specified for each group[10]. Then we can use these probabilities to temporarily (i.e., for one iteration) classify people into groups. For each person we sample from a Trinomial distribution of sample

[9] Convergence can be assessed using a variety of diagnostics. We have chosen to use the $\hat{R}$ statistic proposed by Gelman and Rubin (1992) and its multivariate extension discussed in Brooks and Gelman (to appear). These diagnostics monitor the mixing behavior of several chains, each originating from a different starting point.

[10] We obtain parameter values from the second step in each iteration, so for the first iteration we just start at as random place in the parameter space.

size one with probabilities equal to (draws of) the relative probabilities of belonging to each group (given individual characteristics). This draw specifies a group membership label.

4.3.2 Drawing Parameters Given Group Indicators

We sample parameters from their distribution conditioning on the data (i.e., using the information we have about our survey participants through their response behavior as measured by Z) and the group indicators we drew in the previous step. This is akin to fitting a separate model for each group using only those people classified in the previous step to that group for each analysis.

5 Results

In this section we report the results of the model fit via the DA algorithm for each of the different policy measures. Each table presents the point estimate of the mean for each parameter as well as a 95% confidence interval from the empirical posterior distribution (each with 10,000 draws). Table 3 displays results from the standard model. Unfortunately this model did not fit the data from all questions. In particular, problems arose in estimating parameters of the highly structured durable-changers group which for each question represents only a small fraction of the population. Additional models, involving constraints to the initial specification, were developed and used in three cases. For the three non-constraining issues we had to in each case set certain parameters equal to either zero or one[11]. The standard model fit the three constraining issues. Note that the confidence intervals often are very large, indicating very low precision of some of the parameter estimates. This reflects our uncertainty regarding some of these parameters caused by a lack of a sufficient number of people who engaged in the types of behaviors to which these parameters correspond.

[11] We were alerted to these boundary problems when we ran the initial EM algorithms.

Parameters	Speedlimits		C02 tax		Gas Price Increase	
	Mean	95% C.I.	Mean	95% C.I.	Mean	95% C.I.
π_1	0.53	(0.48, 0.57)	0.44	(0.40, 0.49)	0.44	(0.40, 0.49)
π_2	0.40	(0.35, 0.45)	0.49	(0.44, 0.54)	0.48	(0.43, 0.53)
π_3	0.07	(0.04, 0.11)	0.07	(0.04, 0.10)	0.07	(0.04, 0.11)
α_1	0.52	(0.47, 0.57)	0.41	(0.36, 0.47)	0.62	(0.56, 0.68)
δ_1	0.67	(0.64, 0.70)	0.61	(0.57, 0.65)	0.68	(0.65, 0.71)
φ_2	0.03	(0.02, 0.04)	0.06	(0.05, 0.07)	0.07	(0.05, 0.08)
δ_2	0.21	(0.18, 0.24)	0.27	(0.24, 0.30)	0.25	(0.22, 0.28)
$\varphi_3^{(pre1)}$	0.37	(0.18, 0.60)	0.45	(0.27, 0.66)	0.45	(0.26, 0.66)
$\alpha_3^{(pre1)}$	0.13	(6e-4, 0.33)	0.28	(0.11, 0.47)	0.11	(9e-4, 0.27)
$\varphi_3^{(pre2)}$	0.09	(4e-4, 0.29)	0.46	(3e-3, 0.99)	0.07	(1e-5, 0.34)
$\alpha_3^{(pre2)}$	0.80	(0.52, 0.98)	0.28	(2e-5, 0.92)	0.84	(0.43,.998)
$\alpha_3^{(post)}$	0.33	(0.09, 0.65)	0.61	(0.36,0.86)	.80	(0.52,0.999993)
δ_3	0.54	(0.39, 0.69)	0.50	(0.39, 0.61)	0.58	(0.45, 0.72)
τ_3	0.60	(0.42, 0.78)	0.90	(0.75, 0.99)	0.76	(0.57, 0.93)

Parameters	Electric		Car		Parking	
	Mean	95% C.I.	Mean	95% C.I.	Mean	95% C.I.
π_1	0.53	(0.49, 0.55)	0.65	(0.61, 0.69)	0.41	(0.37, 0.45)
π_2	0.41	(0.37, 0.43)	0.33	(0.29, 0.37)	0.53	(0.48, 0.58)
π_3	0.06	(0.04, 0.07)	0.02	(0.01, 0.05)	0.06	(0.03, 0.10)
α_1	0.05	(0.03, 0.07)	0.05	(0.03, 0.07)	0.18	(0.13, 0.23)
δ_1	0.46	(0.43, 0.47)	0.63	(0.61, 0.66)	0.57	(0.54, 0.61)
φ_2	0.08	(0.06, 0.08)	0.02	(0.015, 0.04)	0.04	(0.03, 0.05)
δ_2	0.23	(0.20, 0.24)	0.26	(0.23, 0.30)	0.25	(0.22, 0.28)
$\varphi_3^{(pre1)}$	0.25	(0.10, 0.33)	0.24	(3e-4, 0.85)	0.40	(0.15, 0.76)
$\alpha_3^{(pre1)}$	0.71	(0.51, 0.80)	0.57	(0.01, 0.98)	0.42	(0.07, 0.72)
$\varphi_3^{(pre2)}$	1*	NA	0.17	(9e-4, 0.71)	0.00*	NA
$\alpha_3^{(pre2)}$	0*	NA	0*	NA	0.63	(0.10 0.98)
$\alpha_3^{(post)}$	0*	NA	0*	NA	0.24	(3e-4,0.63)
δ_3	0.19	(0.10, 0.23)	0.57	(0.27, 0.82)	0.37	(0.18, 0.58)
τ_3	0.93	(0.82, 0.97)	0.47	(0.12, 0.83)	0.63	(0.36, 0.89)
* This parameter fixed by model constraints, not estimated.						

Notational Convention

Subscripts	1 = opinion holders	2 = vacillating changers	3 = durable changers
Greek letters	π = group membership δ = extreme	α = disagree τ = switched after 1st time period	φ = no opinion

Table 3: Estimates of Parameters and their Uncertainty for Standard Model

5.1 Model simplifications and extensions

Model checks which will be discussed in Section 6 led us to consider some variations on the basic three-group model, described in Section 3, that we originally hypothesized. The alternatives to the standard model were the following:

1. Two-group model. This model includes only the opinion holders and the vacillating changers and uses the same parameterization for these groups as described in Section 3.
2. Unconstrained Model. This model loosens the standard model constraint for vacillating changers that

   $$\frac{1-\varphi_2}{2} = \mathrm{Pr}(\text{mildly agree or strongly agree}) = \mathrm{Pr}(\text{mildly disagree or strongly disagree}).$$

   This model therefore adds an additional parameter,

   $$\alpha_2 = \mathrm{Pr}(\text{mildly disagree or strongly disagree}).$$

Comparisons between the standard model and the two-group model will help us to examine the evidence for the existence of the durable changer category (at least for such a group as the one defined in Section 3) given existence of Converse's hypothesized categories, which we've labeled the opinion holders and vacillating changers. Comparisons between the unconstrained and the standard model can be used to examine the evidence for the strict definition of the vacillating-changers group. If this constraint does not appear to fit the data adequately then there is support for the theory that the behavior of this group is not truly random in choosing between agree (mildly and strongly) and disagree (mildly and strongly) responses.

It is probable that none of these models is detailed enough to capture the subtleties in opinion-changing behavior that exist in this time period. However, there was not enough data to support the more complicated and highly-parameterized models that we attempted to fit.

Table 4 displays the results from the unconstrained model (the two-group model didn't fit well enough to warrant reporting the results). Notice that the point estimates (means) of vacillating changers are larger for the non-constraining issues with this model. The results for the electric vehicles question are the most different across models, though the differences are not sizeable. Overall, the parameter estimates are quite similar to those in the standard model. The parameters for the durable changers for all questions most often display differences in point estimates across models; however, the confidence intervals are generally wide enough to overlap each other, which means that true differences may often be negligible. These differences in results across models do not change the substantive views regarding the implications of the model with the obvious exception that to accept the unconstrained model is to accept a view of vacillating changers that does not allow for purely random choice between agreeing and disagreeing with an issue.

6 Diagnostics

We examine the adequacy of our model and estimation algorithm in three ways: statistical checks of the model, statistical checks of the algorithm used to fit the model, and substantive checks of the model.

6.1 Statistical Diagnostics

We used two standard statistical diagnostics to assess model adequacy: posterior predictive checks test the adequacy of specific aspects of the model; Bayes factors test which of the postulated models fits the data better.

6.1.1 Posterior Predictive Checks

To assess statistically how well specific aspects of each of our models fit the data, we have performed *posterior predictive checks* (Rubin 1984; Gelman *et al.* 1996) which generally take the following form:

1. For each draw of model parameters from the posterior distribution, generate a *new* dataset.
2. For each dataset calculate a statistic which measures a relevant characteristic of the model.

Parameters	Speedlimits Mean	Speedlimits 95% C.I.	C02 tax Mean	C02 tax 95% C.I.	Gas Price Increase Mean	Gas Price Increase 95% C.I.
π_1	0.53	(0.48, 0.57)	0.40	(0.36, 0.45)	0.44	(0.40, 0.49)
π_2	0.39	(0.33, 0.45)	0.53	(0.48, 0.58)	0.48	(0.43, 0.53)
π_3	0.08	(0.05, 0.12)	0.07	(0.04, 0.10)	0.08	(0.04, 0.11)
α_1	0.53	(0.48, 0.59)	0.48	(0.41, 0.54)	0.63	(0.57, 0.69)
δ_1	0.67	(0.64, 0.70)	0.65	(0.61, 0.69)	0.68	(0.65, 0.71)
φ_2	0.03	(0.02, 0.04)	0.06	(0.04, 0.07)	0.07	(0.05, 0.08)
α_2	0.45	(0.39, 0.50)	0.37	(0.33, 0.41)	0.45	(0.42, 0.49)
δ_2	0.22	(0.19, 0.25)	0.27	(0.24, 0.29)	0.25	(0.22, 0.28)
$\varphi_3^{(pre1)}$	0.34	(0.16, 0.57)	0.45	(0.27, 0.66)	0.44	(0.26, 0.65)
$\alpha_3^{(pre1)}$	0.10	(1e-4, 0.30)	0.17	(1e-3, 0.37)	0.10	(2e-4, 0.26)
$\varphi_3^{(pre2)}$	0.08	(5e-4, 0.26)	0.36	(3e-4, 0.98)	0.06	(3e-6, 0.32)
$\alpha_3^{(pre2)}$	0.84	(0.56, 0.99)	0.48	(2e-4,0.98)	0.14	(0.47,0.999)
$\alpha_3^{(post)}$	0.35	(0.11, 0.64)	0.71	(0.44, 0.997)	0.81	(0.53,.999999)
δ_3	0.49	(0.35, 0.67)	0.50	(0.39, 0.60)	0.56	(0.43,0.71)
τ_3	0.59	(0.42, 0.76)	0.88	(.71, .99)	0.74	(0.56,0.92)

Parameters	Electric Mean	Electric 95% C.I.	Car Mean	Car 95% C.I.	Parking Mean	Parking 95% C.I.
π_1	0.40	(0.34, 0.46)	0.59	(0.53, 0.64)	0.36	(0.32, 0.41)
π_2	0.57	(0.51, 0.64)	0.40	(0.34, 0.45)	0.58	(0.52, 0.63)
π_3	0.03	(0.01, 0.05)	0.02	(5e-3, 0.04)	0.06	(0.03, 0.09)
α_1	0.10	(0.06, 0.14)	0.07	(0.05, 0.10)	0.23	(0.17, 0.29)
δ_1	0.57	(0.52, 0.62)	0.68	(0.64, 0.71)	0.62	(0.57, 0.66)
φ_2	0.06	(0.05, 0.07)	0.02	(0.01, 0.03)	0.04	(0.03, 0.05)
α_2	0.25	(0.22, 0.28)	0.32	(0.27 ,0.37)	0.36	(0.32, 0.40)
δ_2	0.20	(0.18, 0.23)	0.26	(0.23, 0.29)	0.24	(0.22, 0.27)
$\varphi_3^{(pre1)}$	0.14	(1e-5, 0.56)	0.14	(1e-5, 0.68)	0.41	(0.11, 0.83)
$\alpha_3^{(pre1)}$	0.33	(3e-4, 0.56)	0.19	(1e-5, 0.81)	0.12	(1e-5, 0.48)
$\varphi_3^{(pre2)}$	1*	NA	.36	(2e-3, 0.99)	0.00*	NA
$\alpha_3^{(pre2)}$	0*	NA	0*	NA	1.00*	NA
$\alpha_3^{(post)}$	0*	NA	0*	NA	0.43	(0.01,0.997)
δ_3	0.25	(0.07, 0.54)	0.44	(0.06, 0.80)	0.40	(0.25, 0.59)
τ_3	0.83	(0.54, 0.97)	0.63	(0.14, 0.97)	0.54	(0.30, 0.78)

* This parameter fixed by model constraints, not estimated.

Notational Convention

Subscripts	1 = opinion holders	2 = vacillating changers	3 = durable changers
Greek letters	π = group membership δ = extreme	α = disagree τ = switched after 1st time period	ϕ = no opinion

Table 4: Estimates of Parameters and their Uncertainty for Unconstrained Model

3. Plot the sampling distribution (histogram) of these statistics and see where the observed value of the statistic (i.e. the statistic calculated from the data that was actually observed) lies in relation to this distribution.
4. If this observed value appears consistent enough with the statistics calculated from the generated data (e.g. it falls reasonably well within the bounds of the histogram) then we won't reject the model. Lack of consistency with, or extremity compared to, the generated statistics can be characterized by the percentage of the generated statistics that are more extreme than the observed statistic. We will use the convention of referring to this percentage as the posterior predictive p-value.

Of course, as usual, failure to reject the model does not imply full acceptance of the model, but it heightens our confidence in the model. Posterior predictive checks are easy to implement and they allow for use of a flexible class of statistics without having to analytically calculate sampling distributions for each.

One statistic used to check model adequacy is the percentage of people who get classified as durable changers given that they switch opinions exactly once ($\hat{\pi}_3 / \sum_i I(D_i = 1)$). This statistic reflects the classifying behavior of the model. The p-values are presented in Table 5. Each column displays the p-values corresponding to statistics calculated on datasets generated under one of our models: standard, two-group, and unconstrained.

Issue $\setminus$ Model	Standard	Two-group	Unconstrained
Speed limits	0.42	0.01	0.46
Tax on CO_2	0.43	0.00	0.33
Gas price at 2fr/l	0.29	0.00	0.28
Electrical vehicles	0.41	0.07	0.15
Car free zones	0.45	0.05	0.41
Parking restrictions	0.21	0.00	0.30

Table 5: Posterior predictive p-values for durable-changer-classification statistic

This statistic provides evidence for lack of fit of the two-group model. If the two-group model were an adequate representation of the data, then generating data under this smaller model would yield statistics from the same distribution as our observed statistic. These p-values contradict this hypothesis of the adequacy of the two-group model.

Not surprisingly, these p-values provide no evidence that either three-group model (standard or unconstrained) fails to fit. However, it is important to remember that this statistic represents only one measure of goodness-of-fit, albeit one closely related to the issue with which we are concerned: evidence for existence of a third group, the durable changers. These results do not necessarily imply that either three-group model *in general* fits better than the two-group model. Posterior predictive checks are not generally intended for choosing one model over another unless it is possible to check every aspect of the model about which we might be concerned. They are generally intended to help investigate the evidence for lack of fit of a particular characteristic of a given model. Of course if we are satisfied with the overall fit of two models, would like to choose between them, and there is a limited number of differences between them, we could use posterior predictive checks to test the implications of these differences.

A statistic which targets the difference between the standard and unconstrained three-group models is

$$\hat{\alpha}_2 - (1 - \hat{\varphi}_2 - \hat{\alpha}_2).$$

This represents the discrepancy between the estimated probability of disagreeing or agreeing with an issue for a vacillating changer (which should be 0 on average for the standard model). Data were generated under the standard model and the statistics calculated from these datasets form the null distribution. The observed data statistic should fall well within the bounds of the null distribution (i.e. we should see high p-values) if the constraint seems reasonable for our data.

The only dataset for which the constraint appears reasonable is the gas prices dataset. In all the other cases the imposed constraint does not appear to be consistent with the data, which means that we cannot presuppose purely random (agree/disagree) behavior on the part of the vacillating changers. Even in the particular case of the price of gas, the negative result does not necessarily imply such random behavior on the part of the vacillating changers. We have estimates of φ_2

Issue \ Model	Unconstrained
Speed limits	0.02
Tax on CO_2	0.00
Gas price at 2fr/l	0.49
Electrical vehicles	0.00
Car free zones	0.00
Parking restrictions	0.00

Table 6: Results for the $\hat{\alpha}_2$ statistic

and α_2 in the unconstrained model of .07 and .45 respectively. This implies that the estimated probability of agreeing (mildly or strongly) is .48 which is nearly equal to .45. It could just be that these proportions reflect the "true" opinion distribution of the moment among the vacillating changers with respect to a substantial increase in the price of gas.

Several other posterior predictive checks were performed (including a more global check using the log-likelihood statistic) which indicate good fit for all of the models. Altogether, the posterior predictive checks provide evidence regarding the superior fit of the three-group model versus the two-group model. This in turn provides indirect evidence for the existence of durable changers. In addition, these checks provide little support for the constrained definition of the vacillating changers consistent with completely random agree/disagree responses.

6.1.2 Bayes Factors

A Bayes factor is a measure of the weight of evidence for one model versus another.[12] The higher the value of the Bayes factor, B, (or some appropriate function thereof), the more we believe that our alternative model provides a better fit than the null model, H_0. One suggestion for the interpretation of these values comes from Kass and Raftery (1995) and is displayed in Table 7.

[12] We used the harmonic mean of likelihood values evaluated at parameter draws from our posterior distribution to estimate $\Pr(\mathbf{Z}) = \int \Pr(\mathbf{Z} \mid \theta)\pi(\theta)d\theta$ for each model (Newton and Raftery 1994). For a general description of Bayes factors and methodology used to calculate them see Kass and Raftery (1995).

Bayes Factor ($2\log_e(B)$)	Evidence against H_0
0 to 2	not worth more than a bare mention
2 to 6	positive
6 to 10	strong
$>$10	very strong

Table 7: Bayes Factors Interpretation

We calculate the Bayes factors for each survey question to test the weight of evidence for the standard model versus the two-group model and to test the unconstrained model versus the standard model.

Issue	standard versus two-group	unconstrained versus standard
Speed limits	67	2
Tax on CO_2	58	23
Gas price at 2fr/l	77	0
Electrical vehicles	51	100
Car free zones	6	42
Parking restrictions	16	36

Table 8: Bayes Factors ($2\log_e(B)$)

The Bayes factors which test the weight of evidence for the standard model against the null hypothesis of the two-group model appear to provide strong or very strong evidence in favor of the former. The second column displays similar results for the test of the unconstrained model versus the null model as the standard model. Only the gas price question appears to provide no greater evidence for the unconstrained model. The speed limits question provides only “positive” evidence in favor of the unconstrained over the standard model. These results are consistent with the results we obtained using posterior predictive checks.

6.2 Assessing Frequency Properties of Posterior Confidence Intervals

It is advisable when using any statistical technique to be aware of its frequency properties. For instance, if we formed 95% confidence intervals over repeated samples from the true distribution we would like to know that these intervals would cover the true value at least 95% of the time. Since we never really know the true distribution we can only approximate this scenario. However, such an exercise should still be quite informative.

A simulation was performed which generated 100 datasets using the full model with the maximum likelihood estimates from the speedlimits data as parameters. A DA algorithm (1000 steps) was run on each dataset and 95% confidence intervals were calculated for each parameter in all datasets. Then whether or not the interval covered the "true" value of the parameter from our constructed model was recorded. On average, both across parameters and across datasets, the intervals covered the "true" parameter values slightly more than 95% of the time. This is reassuring evidence about the DA algorithm used in this problem.

6.3 Substantive Checks of the Model

In addition to the statistical checks, we would like to discuss the plausibility of the results by discussing the empirical distributions of some aspects of the response behavior in light of the hypotheses formulated in Section 3. These checks are mostly illustrative but should serve to increase our confidence in the model. The results used in the section are those from the unconstrained model (though there are few substantive differences from those reached if we used the results from the standard model).

6.3.1 The estimated sizes of the three groups

The most important result concerns the estimates of the sizes of the three groups, π_1, π_2, and π_3. Table 9 presents a summary of these estimates for the six issues together with the corresponding shares of respondents who did not change sides throughout our study, those who changed sides once and those who changed them more than once. Changing sides corresponds to a change in opinion between the three major categories: no opinion, agree (mild or strong), or disagree (mild or strong). As can be seen from this table, the model's estimates of opinion holders vary between roughly three fifths (car-free zones) and over one-third (parking restrictions) of the respondents. The corresponding empirical shares of respondents having never changed their opinion throughout our study are consistently higher, but not by very much. This is quite plausible, given that, just by chance, some vacillating changers may never have changed their opinion during the period of our study, although they do not really hold a stable opinion. Car-free zones, the issue with the largest share of opinion holders is also the one most familiar to the Swiss public. All larger Swiss cities have had car-free zones for years, so people have experience with this measure. The other two familiar issues also have comparatively high shares of opinion holders, which confirms Hypothesis 1.1.

Issue	π_1	No change	π_3	1 change	π_2	2+ changes
Speed limits	.53	.62	.08	.21	.39	.17
Tax on CO_2	.40	.52	.07	.21	.53	.27
Gas price at 2fr/l	.44	.52	.08	.26	.48	.23
Electrical vehicles	.40	.57	.03	.17	.57	.26
Car free zones	.59	.70	.02	.13	.40	.17
Parking restrictions	.36	.48	.06	.23	.58	.28

Table 9: The estimated size of the three groups, and shares of corresponding number of respondents who changed sides

As we have observed, the public campaign which took place against the fiscal incentives during the first period of our study, and the experiments mimicking to a certain extent the impact of a political campaign during the second and third periods of the study were of somewhat limited intensity and salience, compared to major electoral campaigns, scandals or catastrophes. Accordingly, as expected on the basis of Hypothesis 1.2, only very few people definitely changed their opinion during the period of our study: the respective estimates vary between 2% (car-free zones) and 8% (speed limits,gas price at 2fr/l). Comparing these estimates for durable changers with the corresponding empirical shares of respondents who have changed their opinion just once during the period of our study, we note that our estimates are all much smaller than the corresponding percentages of people who switched just once. This may be explained by the fact that, by pure chance, many vacillating changers, too, may have changed sides just once during the period in question. Alternately, it may be that our model is incorrectly specified or insufficiently complex to effectively distinguish between vacillating and durable changers. Our results confirm the idea that during a period of controversial, but inconclusive public debate about a policy domain, durable changes of individual opinions occur only rarely and public opinion progresses incrementally, at a rather slow pace.

We are left with very sizeable groups of vacillating changers, ranging from two fifths (.39) of the citizens in the case of speed limits to more than one half (.57) in the case of electrical vehicles. Even if we take into account that the parameterization has been very liberal with respect to this group, and that we, therefore, probably tend to overestimate its size, these estimates are very impressive. By and large they correspond to the empirical percentage of the respondents who at one point or another changed their mind about a given issue in the course of our study.

6.3.2 The implied response behavior of the three groups

As it turns out, among opinion holders, the estimated share of opponents (α_1) is, indeed, larger for the more constraining policy measures than for the less constraining ones, which confirms Hypothesis 2. This is shown by Table 10. As expected by Hypothesis 3.1, the estimated share of opponents among vacillating changers (α_2) is also larger for the more constraining measures than for the less constraining ones. But, as expected, the respective differences are much more pronounced among opinion holders than among vacillating changers. For reasons of ambivalence, uncertainty or both, the vacillating changers' opinions are less discriminating than those of the opinion holders. The shares of opponents among the population at large turn out to be mixtures of the corresponding shares of the two groups.

Issue	α_1	Population at large*	α_2
Gas price at 2fr/l	.63	.54	.45
Speed limits	.53	.50	.45
Tax on CO_2	.48	.42	.37
Parking restrictions	.23	.32	.36
Electrical vehicles	.10	.19	.25
Car free zones	.07	.17	.32

*Average number of opponents across the four waves.

Table 10: The estimated proportion of opponents among opinion holders and the population at large

As Table 11 shows, Hypothesis 3.2 is also confirmed, since in each and every case, opinion holders are much more likely to hold strong opinions (δ_1) than vacillating changers (δ_2). We have not formulated any hypotheses about the durable changers in this respect. In their case, the respective shares of strong opinions (δ_3) turn out to be more variable than in the other groups. Nor have we formulated any hypotheses about issue-specific variations. For all of the issues, strong opinions prevail among opinion holders. In fact in most instances, approximately two-thirds of the opinion holders have a strong opinion, compared to only roughly one fourth of the vacillating changers.

Group Issue	Opinion Holders δ_1	Durable Changers δ_3	Vacillating Changers δ_2
Speed limits	.67	.49	.22
Tax on CO_2	.65	.50	.27
Gas price at 2fr/l	.68	.56	.25
Electrical vehicles	.57	.25	.20
Car free zones	.68	.44	.26
Parking restrictions	.62	.40	.24

Table 11: The estimated share of strong opinions among the three groups

6.3.3 Measures of extremity of reaction

As we discussed in Section 3.1.1 the model simplification that is least tenable is the constraint placed on the extremity reaction parameters, δ_1, δ_2, and, δ_3, to be constant within group. It is possible to fit a more general model that allows δ to vary between opinion boundaries. We have not pursued this due to sample size constraints.

7 Conclusion

We conclude from the statistical and substantive checks that our unconstrained parameterization of the model (i.e. adding the α_2 parameter) as well as the estimates it provides have a high degree of plausibility. We believe we have shown that finite mixture models provide a framework which allows us to reasonably estimate proportions of stable opinion holders, vacillating changers, and durable changers. According to our estimates, on average between 36% and 59% of the Swiss sample have stable opinions about the policy measures dealing with environmental protection in the domain of the use of private cars. These figures compare favorably to those of Converse (1970, 176), who calculated that, in the case of the attitude item fitting the black-and-white model, "80% of the responses represented confessions of 'no opinion' or statistically random responses". Our estimation procedure confirms the impression of a comparatively high level of stability, which is already provided by a simple inspection of the correlation matrices. The bivariate correlations between our Swiss items display the same temporal pattern as that which led Converse to his "black-and-white" model in the first place, but they suggest a rather high level of stability: they are located in the range (.45 to .50) of the correlations for American social welfare items reported by Converse and Markus (1979) as far as the less constraining issues are concerned, and in the range of the American moral issues (.62 to .64) as for the more constraining ones[13].

[13]Note, however, that the issue-specific stability observed for our Swiss policy measures falls still far short of the stability (.81 to .83) that has been measured in the U.S. for a basic political orientation such as one's party identification.

But even in this case of comparatively high issue-specific stability, the limited share of opinion holders leaves much room for political campaigns and major events to influence the public opinion. Our results suggests that, short of major events, durable changes in individual opinions occur only rarely. Most individual opinion change is likely to consist of short-term reactions to external stimuli. Our model estimations show, however, that contrary to Converse's original "black-and-white" model, the vacillating changers who are responsible for these short-term reactions do not behave completely at random. Thus, their tendency to agree or disagree with a given issue depends, to some extent at least, on the degree to which the issue imposes individual constraints. This result not only turns against the notion of completely arbitrary responses on the part of the vacillating changers, it also undermines the notion that the apparent chance variation in the individual responses is essentially due to measurement error, as maintained by Achen. By contrast, this important result is in line with Zaller's notion that people are expressing real feelings and respond to the issue as they see it at the moment of response, even if they are temporally unstable, expressing completely opposing positions at different times. While we have not done so in the present paper, on the basis of an extended analysis of the data used here it is possible to show that the short-term reactions of the vacillating changers do not behave randomly, but respond to the information they get during public debates.

References

Achen, C. H. (1975), "Mass political attitudes and the survey response," *American Political Science Review* 69, 1218–1231.

Alvarez, R. M. and Brehm, J. (1995), "American ambivalence towards abortion policy: Development of a heteroskedastic probit model of competing values," *American Journal of Political Science* 39, 1055–82.

Alvarez, R. M. and Brehm, J. (1997), "Are Americans ambivalent towards racial policies?" *American Journal of Political Science* 41, 345–374.

Baumgartner, F. R. and Jones, B. D. (1993), *Agendas and Instability in American Politics*, Chicago: The University of Chicago Press.

Belin, T. and Rubin, D. (1995), "The analysis of repeated-measures data on schizophrenic rection times using mixture models," *Statistics in Medicine* 90, 694–707.

Brooks, S. and Gelman, A. (to appear), "General methods for monitoring convergence of iterative simulations," *Journal of Computational and Graphical Statistics* .

Campbell, A., Converse, P. E., Miller, W. E., and Stokes, D. E. (1964), *The American Voter*, Chicago: University of Chicago Press.

Converse, P. E. (1964), "The Nature of Belief Systems in Mass Publics," in *Ideology and Discontent*, ed. D. Apter, pp. 206–61, New York: The Free Press.

Converse, P. E. (1970), "Attitudes and non-attitudes: continuation of a dialogue," in *The quantitative analysis of social problems,*, ed. E. R. Tufte, pp. 168–189, Reading, MA: Addison-Wesley.

Converse, P. E. and Markus, G. B. (1979), "Plus ca change: The new CPS election study panel," *American Political Science Review* 73, 1, 32–49.

Eagly, A. H. and Chaiken, S. (1993), *The Psychology of Attitudes*, New York: Harcourt, Brace, Jovanovich.

Everitt, B. S. and Hand, D. J. (1981), *Finite Mixture Distributions*, Chapman & Hall.

Gelman, A., Carlin, J., Stern, H., and Rubin, D. (1995), *Bayesian Data Analysis*, London: Chapman & Hall.

Gelman, A. and King, G. (1990), "Estimating the electoral consequences of legislative redistricting," *Journal of the American Statistical Association* 85.

Gelman, A., Meng, X.-L., and Stern, H. (1996), "Reply to Comments on "Posterior Predictive Assessment of Model Fitness Via Realized Discrepancies"," *Statistica Sinica* 6, 796–807.

Gelman, A. and Rubin, D. B. (1992), "Inference From Iterative Simulation Using Multiple Sequences (Disc: P483-501, 503-511)," *Statistical Science* 7, 457–472.

Jagodzinski, W., Khnel, S. M., and Schmidt, P. (1987), "Is there a 'Socratic effect' in nonexperimental panel studies? Consistency of an attitude toward guestworkers," *Sociological Methods and Research* 15, 3, 259–302.

Kass, R. E. and Raftery, A. E. (1995), "Bayes Factors," *Journal of the American Statistical Association* 90, 773–795.

Krosnick, J. A. and Fabrigar, L. R. (1995), "No Opinion Filters and Attitude Strength," Tech. rep., Department of Psychology, The Ohio State University.

Lindsay, B. G. (1995), *Mixture Models: Theory, Geometry and Applications*, Institute of Mathematical Statistics.

Little, R. J. A. and Rubin, D. B. (1987), *Statistical Analysis With Missing Data*, New York: John Wiley & Sons.

McGuire, W. J. (1960), "A syllogistic analysis of cognitive relationships," in *Attitude Organization and Change*, eds. M. J. Rosenberg, C. Hovland, W. McGuire, R. Abelson, and J. Brehm, pp. 65–111, Westport, CT: Greenwood Press.

Neijens, P. (1987), *The Choice Questionnaire. Design and Evaluation of an Instrument for Collecting Informed Opinions of a Population*, Amsterdam: Free University Press.

Neijens, P., Minkman, M., de Ridder, J., Saris, W., and Slot, J. (1996), "A Decision Aid in a Referendum," *International Journal of Public Opinion Research* 8, 1, 83–90.

Newton, M. A. and Raftery, A. E. (1994), "Approximate Bayesian Inference With the Weighted Likelihood Bootstrap (Disc: P26-48)," *Journal of the Royal Statistical Society, Series B, Methodological* 56, 3–26.

Rubin, D. B. (1984), "Bayesianly Justifiable and Relevant Frequency Calculations for the Applied Statistician," *The Annals of Statistics* 12, 1151–1172.

Saris, W. E., Neijens, P., and de Ridder, J. A. (1983), *Kernenergie: ja of nee?*, Amsterdam: SSO.

Saris, W. E. and van den Putte, B. (1987), "True score or factor models. A secondary analysis of the ALLBUS-test-retest data," *Sociological Methods and Research* 17, 2, 123–157.

Sears, D. O. and Valentino, N. A. (1997), "Politics Matters: Political Events as Catalysts for Preadult Socialization," *American Political Science Review* 91, 1, 45–65.

Tanner, M. A. and Wong, W. H. (1987), "The Calculation of Posterior Distributions By Data Augmentation (C/R: P541-550)," *Journal of the American Statistical Association* 82, 528–540.

Titterington, D., Smith, A., and Makov, U. (1985), *Statistical Analysis of Finite Mixture Distributions*, New York: .

Turner, D. and West, M. (1993), "Bayesian analysis of mixtures applied to postsynaptic potential fluctuations," *Journal of Neuroscience Methods* 47, 1–23.

van Dyk, D. A., Meng, X.-L., and Rubin, D. B. (1995), "Maximum Likelihood Estimation Via the ECM Algorithm: Computing the Asymptotic Variance," *Statistica Sinica* 5, 55–75.

van Dyk, D. A. and Protasso, R. (1999), "Statistics: Handle with Care," Tech. rep., Harvard University.

van Steenbergen, M. R., Brewer, P., and Vercellotti, T. (1997), "Generalized Value Conflict and Uncertainty in Public Opinion," Paper prepared for the Conference "No opinion, Instability and Change in Public Opinion Research", Amsterdam.

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

Technical Appendix

A Full Model Parameterization

If we denote the vector random variable for group membership for the i^{th} person as $\mathbf{G}_i$, the probability of falling into each group can be described by the following parameters:

$$\begin{aligned}
\pi_1 &= \text{Pr(individual } i \text{ belongs to Opinion-Holder Group)} = \Pr(\mathbf{G}_i = v_1), \\
\pi_2 &= \text{Pr(individual } i \text{ belongs to Vacillating-Changer Group)} = \Pr(\mathbf{G}_i = v_2), \\
\pi_3 &= \text{Pr(individual } i \text{ belongs to Durable-Changer Group)} = \Pr(\mathbf{G}_i = v_3),
\end{aligned}$$

where, $\pi_1 + \pi_2 + \pi_3 = 1$, and the v_j are simply vectors of length 3 with a one in the j^{th} position and zeros elsewhere.

A.1 Opinion Holders

The two parameters,

$$\begin{aligned}
\alpha_1 &= \Pr(Y_{1,i} = 1 \text{ or } 2 \mid \mathbf{G}_i = v_1), \text{ and} \\
\delta_1 &= \Pr(Y_{t,i} = 1 \mid Y_{1,i} = 1 \text{ or } 2, \mathbf{G}_i = v_1) = \Pr(Y_{t,i} = 5 \mid Y_{1,i} = 4 \text{ or } 5, \mathbf{G}_i = v_1),
\end{aligned}$$

are sufficient to describe the behavior of opinion holders across the four time points, given the constraints

$$\begin{aligned}
\Pr(Y_{t,i} = 3) &= 0 \text{ , } \forall t, \\
\Pr(Y_{t,i} \subset \{3,4,5\} \mid Y_{1,i} \subset \{1,2\}, \mathbf{G}_i = v_1) &= 0 \text{ , } \forall t \neq 1 \text{ , and,} \\
\Pr(Y_{t,i} \subset \{1,2,3\} \mid Y_{1,i} \subset \{4,5\}, \mathbf{G}_i = v_1) &= 0 \text{ , } \forall t \neq 1,
\end{aligned}$$

which formally set

$$\begin{aligned}
1 - \alpha_1 &= \Pr(Y_{1,i} = 4 \text{ or } 5 \mid \mathbf{G}_i = v_1), \text{ and} \\
1 - \delta_1 &= \Pr(Y_{t,i} = 2 \mid Y_{1,i} \subset \{1,2\}, \mathbf{G}_i = v_1) = \Pr(Y_{t,i} = 4 \mid Y_{1,i} \subset \{4,5\}, \mathbf{G}_i = v_1).
\end{aligned}$$

A.2 Vacillating Changers

The behavior of the members of this group can be expressed through two parameters as well.

$$\begin{aligned}
\varphi_2 &= \Pr(Y_{t,i} = 3 \mid \mathbf{G}_i = v_2), \text{ and} \\
\delta_2 &= \Pr(Y_{t,i} = 1 \mid Y_{t,i} \subset \{1,2\}, \mathbf{G}_i = v_2) = \Pr(Y_{t,i} = 5 \mid Y_{t,i} \subset \{4,5\}, \mathbf{G}_i = v_2).
\end{aligned}$$

We assume that

$$\frac{1-\varphi_2}{2} = \Pr(Y_{t,i} \subset \{1,2\} \mid \mathbf{G}_i = v_2) = \Pr(Y_{t,i} \subset \{4,5\} \mid \mathbf{G}_i = v_2).$$

A.3 Durable Changers

Durable changers' more structured behavior is reflected in the increased number of parameters:

$$\begin{aligned}
\varphi_3^{\text{(pre1)}} &= \Pr(Y_{t^*,i} = 3 \mid t^* = 1, \mathbf{G}_i = v_3) \\
\varphi_3^{\text{(pre2)}} &= \Pr(Y_{t^*,i} = 3 \mid t^* \subset \{2,3\}, \mathbf{G}_i = v_3) \\
\alpha_3^{\text{(pre1)}} &= \Pr(Y_{t^*,i} \subset \{1,2\} \mid t^* = 1, \mathbf{G}_i = v_3) \\
\alpha_3^{\text{(pre2)}} &= \Pr(Y_{t^*,i} \subset \{1,2\} \mid t^* \subset \{2,3\}, \mathbf{G}_i = v_3) \\
\alpha_3^{\text{(post)}} &= \Pr(Y_{(t^*+1),i} \subset \{1,2\} \mid Y_{t^*} = 3, \mathbf{G}_i = v_3) \\
\delta_3 &= \Pr(Y_{t,i} = 1 \mid Y_{t,i} \subset \{1,2\}, \mathbf{G}_i = v_3) \\
&= \Pr(Y_{t,i} = 5 \mid Y_{t,i} \subset \{4,5\}, \mathbf{G}_i = v_3), \text{ and} \\
\tau_3 &= \Pr(t^* = 1 \mid \mathbf{G}_i = v_3).
\end{aligned}$$

where t^* represents the time period directly prior to a change in opinion. Here, assume that

$$\Pr(t^* = 2) = \Pr(t^* = 3) = \frac{1-\tau_3}{2}.$$

B The Likelihoods

There are 625 different response patterns possible in the data. Therefore a simple model for the observed data is a multinomial distribution where each person has a certain probability of falling in each one of 625 response-pattern bins. This model would ignore the group structure we're interested in however. Extending this idea to a model for the complete data which includes not only the response patterns, $\mathbf{X}$, but also the group membership indicators, $\mathbf{G}$, would be a product multinomial model with a separate multinomial model for each group.

Let X_i denote a vector random variable of length 625 with elements $X_{k,i}$, where $X_{k,i} = 1$ if individual i has response pattern k and zero otherwise. If the group membership of each study participant were known the likelihood function would be

$$L(\theta \mid \boldsymbol{X}, \mathbf{G}) = \prod_{i=1}^{N} \prod_{k=1}^{625} \prod_{j=1}^{3} (\pi_j p_{k.j})^{x_{k,i} g_{j,i}},$$

where θ represents the model parameters, and $p_{k.j}$ denotes the probability[14] of belonging to cell k (having response pattern k) given that one is in group j. This is called the "complete-data likelihood."

The likelihood function given the observed data, however, which does not include the unknown group membership indicators[15], is

$$L(\theta \mid \mathbf{X}) \;=\; \prod_{i=1}^{N} \Big(\sum_{j=1}^{3} p(X_i, G_i \mid \theta) \Big) = \prod_{i=1}^{N} \Big(\pi_1 \prod_{k=1}^{625} p_{1,k}^{x_{k,i}} + \pi_2 \prod_{k=1}^{625} p_{2,k}^{x_{k,i}} + \pi_3 \prod_{k=1}^{625} p_{3,k}^{x_{k,i}} \Big) \quad (1)$$

Maximum likelihood estimation, which requires us to maximize Equation 1 as a function of θ, is complicated by the summation.

Using the new variables described in Section 4.1 (and for the reasons described in Section 3), we can re-express the complete-data likelihood as:

$$L(\theta \mid \boldsymbol{Z}, \mathbf{G}) = \prod_{i=1}^{N} \prod_{j=1}^{3} \pi_j^{g_{j,i}} {p_{j,i}^{*}}^{g_{j,i}},$$

where $p_{j,i}^{*}$ is the probability that individual i belongs to group j conditional on his observed data, Z_i .

[14] Note that although many of the $p_{k.j}$ are structural zeros, the corresponding $x_{k,i} g_{j,i}$ will always be zero as well, and $0^0 = 1$.

[15] It should be noted, however, that individuals who crossed opinion boundaries two or three times, or who crossed once but into the "no opinion" category, must be vacillating changers. All others would be one two of the groups since opinion holders and durable changers have mutually exclusive behavior patterns. So, in a sense, group membership labels are "observed" for these people. This is accounted for in the algorithm.

The conditional probability of individual i being an opinion holder ($j = 1$) given her responses (Z_i) can be calculated as

$$p^*_{1,i} = (\alpha_1^{1-A_i}(1-\alpha_1)^{A_i}(1-\delta_1)^{(4-B_i)}\delta_1^{B_i}) * \mathrm{I}(C_i = 0)\mathrm{I}(D_i = 0) \; ,$$

where $\mathrm{I}(\cdot)$ is an indicator function which equals 1 if the condition in parentheses holds and equals 0 otherwise. The indicator functions constrain this probability to be zero for behavior that is disallowed for this group: responding with no opinion during at least one time period, $\mathrm{I}(C_i > 0)$; and, switching opinions, $\mathrm{I}(D_i \neq 0)$.

Similarly, the conditional probability of belonging to each of the other groups given the observed responses can be expressed by the functions

$$\begin{aligned} p^*_{2,i} &= \varphi_2^{C_i}\left(\frac{1-\varphi_2}{2}\right)^{4-C_i}\delta_2^{B_i}(1-\delta_2)^{(4-C_i-B_i)}, \qquad \text{and,} \\ p^*_{3,i} &= \left[\tau_3^{\mathrm{I}(E_i=1)}\frac{(1-\tau_3)}{2}^{\mathrm{I}(E_i\subset\{2,3\})}\left(\varphi_3^{(\text{pre1})^{H_i}}\alpha_3^{(\text{pre1})^{(1-F_i)(1-H_i)}}(1-\alpha_3^{(\text{pre1})}-\varphi_3^{(\text{pre1})})^{F_i}\right)^{\mathrm{I}(E_i=1)}\right. \\ &\quad \left(\varphi_3^{(\text{pre2})^{H_i}}\alpha_3^{(\text{pre2})^{(1-F_i)(1-H_i)}}(1-\varphi_3^{(\text{pre2})}-\alpha_3^{(\text{pre2})})^{F_i}\right)^{\mathrm{I}(E_i\subset\{2,3\})} \\ &\quad \left.\left((1-\alpha_3^{(\text{post})})^{M_i}(\alpha_3^{(\text{post})})^{(1-M_i)}\right)^{H_i}\delta_3{}^{B_i}(1-\delta_3)^{(4-B_i-C_i)}\right]\mathrm{I}(D_i=1)\mathrm{I}(Q_i=0) \; . \end{aligned}$$

Clearly this model ignores potentially relevant background information such as gender, age, political affiliation, income, etc. Later efforts might seek to incorporate this information in more complicated models.

C The EM algorithm

C.1 The E-step

The E-step for this model replaces the missing data with their expected values. Specifically, we take the expected value of the complete data log likelihood,

$$Q = E[logL(\theta \mid \boldsymbol{Z}, \mathbf{G}) \mid \boldsymbol{Z}, \theta^{(\ell)}]$$

where the expectation is taken over the distribution of the missing data,

$$p(\mathbf{G}_i \mid \mathbf{Z}_i, \theta^{(\ell)}) = \prod_{j=1}^{3} \left(\frac{\pi_j^{(\ell)} p_{j,i}^{*(\ell)}}{\sum_{j=1}^{3} \pi_j^{(\ell)} p_{j,i}^{*(\ell)}} \right)^{g_{j,i}}$$

conditional on the observed data, $\mathbf{Z}_i$, and the parameters from the previous ($(\ell-1)^{st}$) iteration of the M-step. Q is linear in the missing data (the group indicators), so the E-step reduces to finding the expectation of the missing data and plugging it into the complete-data log likelihood. The expectation of the indicator for group j and individual i is,

$$E[\mathbf{G}_i = v_j \mid \boldsymbol{Z}_i, \theta^{(\ell)}] = \frac{\pi_j^{(\ell)} p_{j,i}^{*(\ell)}}{\sum_{j=1}^{3} \pi_j^{(\ell)} p_{j,i}^{*(\ell)}},$$

which is the probability, given her response pattern, of falling into group j relative to the other groups.

C.2 The M-step

In the ℓ^{th} iteration, the M-step finds the parameter estimates, $\theta^{(\ell)}$ that maximize Q.

The following are the equations for the parameter estimates that maximize Q given the estimates of $\mathbf{G}_i = (g_{1,i}, g_{2,i}, g_{3,i})^T$ from the E-step (where $T_j = \sum_{i=1}^{N} g_{j,i}$):

$$\begin{aligned}
\hat{\pi}_j &= \frac{T_j}{\sum_{j=1}^{3} T_j}, \qquad \text{j = 1,2,3,} \\
\hat{\alpha}_1 &= \frac{\sum_{i=1}^{N} g_{1,i} \mathrm{I}(C_i = 0)\mathrm{I}(D_i = 0)(1 - A_i)}{\sum_{i=1}^{N} g_{1,i} \mathrm{I}(C_i = 0)\mathrm{I}(D_i = 0)}, \\
\hat{\delta}_1 &= \frac{\sum_{i=1}^{N} g_{1,i} \mathrm{I}(C_i = 0)\mathrm{I}(D_i = 0) B_i}{4\sum_{i=1}^{N} g_{1,i} \mathrm{I}(C_i = 0)\mathrm{I}(D_i = 0)}, \\
\hat{\varphi}_2 &= \frac{\sum_{i=1}^{N} g_{2,i} C_i}{4T_2}, \\
\hat{\delta}_2 &= \frac{\sum_{i=1}^{N} g_{2,i} B_i}{\sum_{i=1}^{N} g_{2,i}(4 - C_i)}, \\
\hat{\varphi}_3^{\text{(pre1)}} &= \frac{\sum_{i=1}^{N} g_{3,i} \mathrm{I}(E_i = 1)\mathrm{I}(D_i = 1)\mathrm{I}(Q_i = 0) H_i}{\sum_{i=1}^{N} g_{3,i} \mathrm{I}(E_i = 1)\mathrm{I}(Q_i = 0)\mathrm{I}(D_i = 1)},
\end{aligned}$$

$$
\begin{aligned}
\hat{\alpha}_3^{\text{(pre1)}} &= \frac{\sum_{i=1}^N g_{3,i}\mathrm{I}(E_i = 1)\mathrm{I}(D_i = 1)\mathrm{I}(Q_i = 0)(1 - F_i)(1 - H_i)}{\sum_{i=1}^N g_{3,i}\mathrm{I}(E_i = 1)\mathrm{I}(D_i = 1)\mathrm{I}(Q_i = 0)}\,, \\
\hat{\varphi}_3^{\text{(pre2)}} &= \frac{\sum_{i=1}^N g_{3,i}\mathrm{I}(E_i \subset \{2,3\})\mathrm{I}(D_i = 1)\mathrm{I}(Q_i = 0)H_i}{\sum_{i=1}^N g_{3,i}\mathrm{I}(E_i \subset \{2,3\})\mathrm{I}(D_i = 1)\mathrm{I}(Q_i = 0)}\,, \\
\hat{\alpha}_3^{\text{(pre2)}} &= \frac{\sum_{i=1}^N g_{3,i}\mathrm{I}(E_i \subset \{2,3\})\mathrm{I}(D_i = 1)\mathrm{I}(Q_i = 0)(1 - F_i)(1 - H_i)}{\sum_{i=1}^N g_{3,i}\mathrm{I}(E_i \subset \{2,3\})\mathrm{I}(D_i = 1)\mathrm{I}(Q_i = 0)}\,, \\
\hat{\alpha}_3^{\text{(post)}} &= \frac{\sum_{i=1}^N g_{3,i}H_i\mathrm{I}(D_i = 1)\mathrm{I}(Q_i = 0)(1 - M_i)}{\sum_{i=1}^N g_{3,i}H_i\mathrm{I}(D_i = 1)\mathrm{I}(Q_i = 0)}\,, \\
\hat{\delta}_3 &= \frac{\sum_{i=1}^N g_{3,i}\mathrm{I}(D_i = 1)\mathrm{I}(Q_i = 0)B_i}{\sum_{i=1}^N g_{3,i}\mathrm{I}(D_i = 1)\mathrm{I}(Q_i = 0)(4 - C_i)}\,, \\
\hat{\tau}_3 &= \frac{\sum_{i=1}^N g_{3,i}\mathrm{I}(D_i = 1)\mathrm{I}(Q_i = 0)\mathrm{I}(E_i = 1)}{\sum_{i=1}^N g_{3,i}\mathrm{I}(D_i = 1)\mathrm{I}(Q_i = 0)}\,.
\end{aligned}
$$

C.3 Iterations

EM has the desirable quality that the value of the log-likelihood increases at every step. Iterations continue until the log-likelihood increased by less than 1×10^{-10}.

D The Data Augmentation Algorithm

The DA algorithm, whose goal it is to get draws from $p(\theta \mid \mathbf{Z})$, has two basic steps in this problem:

1. Draw the “missing data”, group membership indicators, given the parameters.
2. Draw the parameters, $\theta = (\pi_1, \pi_2, \pi_3, \alpha_1, \delta_1, \dots)$, given the group membership indicators.

D.1 Drawing Group Indicators Given Parameters

The first step is to assign group membership labels to all observations conditional on the parameters. We draw from

$$
\text{Mult}\left(\frac{\pi_1^{(\ell)} p_{1,i}{}^{*(\ell)}}{\sum_{j=1}^3 \pi_1^{(\ell)} p_{1,i}{}^{*(\ell)}}, \frac{\pi_2^{(\ell)} p_{2,i}{}^{*(\ell)}}{\sum_{j=1}^3 \pi_2^{(\ell)} p_{2,i}{}^{*(\ell)}}, \frac{\pi_3^{(\ell)} p_{3,i}{}^{*(\ell)}}{\sum_{j=1}^3 \pi_3^{(\ell)} p_{3,i}{}^{*(\ell)}}\right).
$$

D.2 Drawing Parameters Given Group Indicators

Next we draw parameters from their distribution conditioning on the data and the group membership indicators drawn in the previous step. The posterior distribution can be expressed as $p(\theta \mid \mathbf{Z}, \mathbf{G}) = L(\theta \mid \mathbf{Z}, \mathbf{G})p(\theta)$, where $p(\theta)$ is the prior distribution on the parameters, θ. The complete data likelihood can be expressed as

$$\begin{aligned}
L(\theta \mid \mathbf{Z}, \mathbf{G}) &= \prod_{i=1}^{N}\prod_{j=1}^{3} \pi_j^{g_{j,i}} {p_{j,i}^*}^{g_{j,i}} \\
&= \prod_{i=1}^{N} {\pi_1}^{g_{1,i}} \left(\left[\alpha_1^{(1-A_i)} (1-\alpha_1)^{A_i} (1-\delta_1)^{(4-B_i)} \delta_1^{B_i} \right] (\mathrm{I}(C_i = 0)\mathrm{I}(D_i = 0)) \right)^{g_{1,i}} \\
& \quad {\pi_2}^{g_{2,i}} \left(\varphi_2^{C_i} \left(\frac{1-\varphi_2}{2} \right)^{4-C_i} \delta_2^{B_i} (1-\delta_2)^{(4-C_i-B_i)} \right)^{g_{2,i}} \\
& \quad {\pi_3}^{g_{3,i}} \Bigg\{ \Bigg[\tau_3^{\mathrm{I}(E_i=1)} \frac{(1-\tau_3)}{2}^{\mathrm{I}(E_i \subset \{2,3\})} \\
& \quad \left({\varphi_3^{\text{(pre1)}}}^{H_i} {\alpha_3^{\text{(pre1)}}}^{(1-F_i)(1-H_i)} (1-\alpha_3^{\text{(pre1)}} - \varphi_3^{\text{(pre1)}})^{F_i} \right)^{\mathrm{I}(E_i=1)} \\
& \quad \left({\varphi_3^{\text{(pre2)}}}^{H_i} {\alpha_3^{\text{(pre2)}}}^{(1-F_i)(1-H_i)} (1-\varphi_3^{\text{(pre2)}} - \alpha_3^{\text{(pre2)}})^{F_i} \right)^{\mathrm{I}(E_i \subset \{2,3\})} \\
& \quad \left((1-\alpha_3^{\text{(post)}})^{M_i} (\alpha_3^{\text{(post)}})^{(1-M_i)} \right)^{H_i} {\delta_3}^{B_i} (1-\delta_3)^{(4-B_i-C_i)} \Bigg] \mathrm{I}(D_i = 1)\mathrm{I}(Q_i = 0) \Bigg\}^{g_{3,i}}
\end{aligned}$$

If I assume *a priori* independence of appropriate parameters, and choose conjugate priors, I can factor $p(\theta)$ into seven independent Beta distributions (for α_1, δ_1, φ_2, δ_2, τ_3, $\alpha_3^{\text{(post)}}$, and δ_3) and three independent Dirichlet distributions (for (π_1, π_2, π_3), $(\varphi_3^{\text{(pre1)}}, \alpha_3^{\text{(pre1)}}, (1-\varphi_3^{\text{(pre1)}} - \alpha_3^{\text{(pre1)}})$, and $(\varphi_3^{\text{(pre2)}}, \alpha_3^{\text{(pre2)}}, (1-\varphi_3^{\text{(pre2)}} - \alpha_3^{\text{(pre2)}}))$.

Parameters can then be drawn from the appropriate posterior distributions (found by standard conditional probability calculations). For example, if $p(\alpha_1, 1-\alpha_1)$ is specified as a Beta(a, b), then we would draw α (and $(1-\alpha)$) from Beta$[(a + N - \sum_{i=1}^{N} A_i), (b + \sum_{i=1}^{N} A_i)]$.

Given the tree structure of the model and the conditional independence that it implies, prior independence of the parameters does not seem an unwarranted assumption. Priors were chosen to be as non-informative as possible. Beta and Dirichlet priors can be conceptualized as "pseudo-counts." For instance, using a Beta(1,1) prior for the distribution of α_1 can be thought of as adding one person to the group of opinion holders who were against the issue and one person to the group of opinion holders who were for the issue. The prior specifications used in these analyses gave equal weight *a priori* to both (or all three) possibilities modeled by a particular distribution and kept the hyperparameters (parameters of the priors) as small as possible. It is sometimes necessary in models with boundaries to have hyperparameters larger than zero just to ensure the algorithm doesn't fail because it attempts to draw parameters values that are so small that the computer can only see them as zero (or the same issue with values close to one). The prior used most often in this analysis added six pseudo-counts to the analysis, split equally across categories of behavior (note that fractional hyperparameter values can be used).

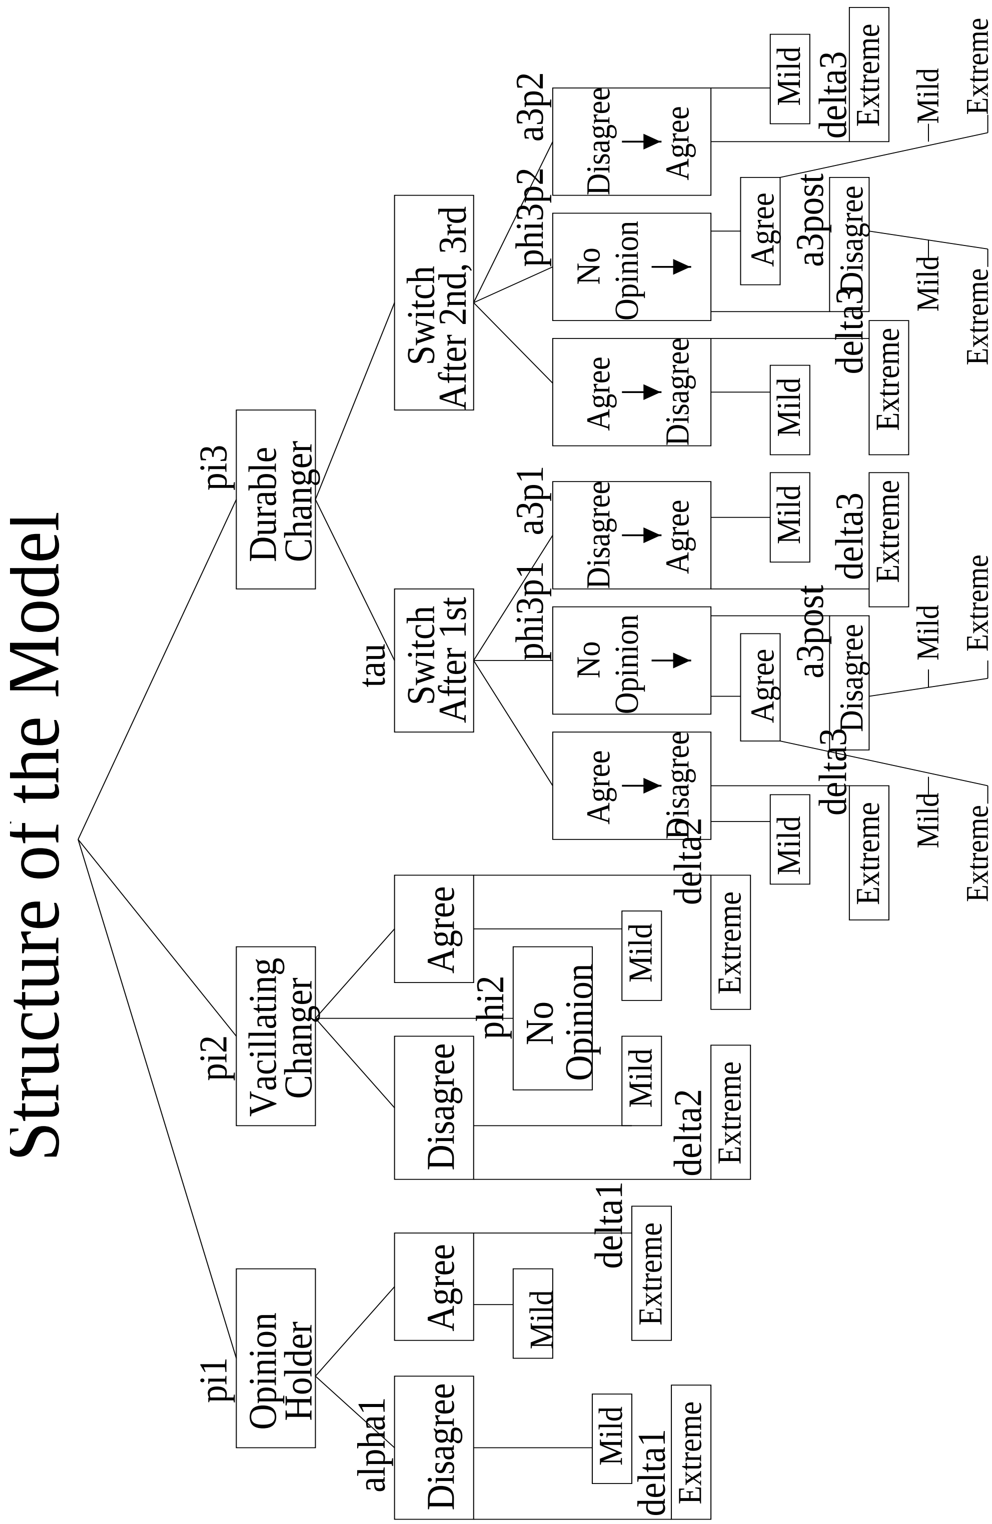

Figure 4: Tree Structure of the Model