

1
2
3 **A VARIATIONAL EM METHOD FOR MIXED MEMBERSHIP**
4 **MODELS WITH MULTIVARIATE RANK DATA: AN**
5 **ANALYSIS OF PUBLIC POLICY PREFERENCES**

6 BY Y. SAMUEL WANG, ROSS L. MATSUEDA AND ELENA A. EROSHEVA

7 *University of Washington*

8
9
10 In this article, we consider modeling ranked responses from a heteroge-
11 neous population. Specifically, we analyze data from the Eurobarometer 34.1
12 survey regarding public policy preferences toward drugs, alcohol, and AIDS.
13 Such policy preferences are likely to exhibit substantial differences within as
14 well as across European nations reflecting a wide variety of cultures, political
15 affiliations, ideological perspectives, and common practices. We use a mixed
16 membership model to account for multiple subgroups with differing prefer-
17 ences and to allow each individual to possess partial membership in more than
18 one subgroup. Previous methods for fitting mixed membership models to rank
19 data in a univariate setting have utilized an MCMC approach and do not es-
20 timate the relative frequency of each subgroup. We propose a variational EM
21 approach for fitting mixed membership models with multivariate rank data.
22 Our method allows for fast approximate inference and explicitly estimates the
23 subgroup sizes. Analyzing the Eurobarometer 34.1 data, we find interpretable
24 subgroups which generally agree with the “left versus right” classification of
25 political ideologies.

26
27 **1. Introduction.** Rank data often arise from a heterogeneous population with
28 individuals whose preferences may vary widely. In this article, we consider one
29 such example, public health policy data from the Eurobarometer 34.1. We develop
30 a computationally efficient variational EM procedure to estimate mixed member-
31 ship models with rank data. In addition to the computational aspects, we also ex-
32 tend the current literature by explicitly estimating the subgroup relative frequen-
33 cies and accommodating multivariate ranked data.

34
35 **2. Public policy preferences.** Social scientists have long held that in demo-
36 cratic societies, public opinion plays an important role in the formation of public
37 policies about important social problems [e.g., Burstein (1998), Brooks and Manza
38 (2008)]. Public opinion polls, which provide a window on the attitudes and per-
39 spectives of a nation’s citizenry, at times elicit ranked data about specific public
40 policies. An example is the Eurobarometer 34.1, a survey commissioned in 1990
41 to study European perspectives on various political and public health issues [Reif

42 Received December 2015; revised February 2017.

43 *Key words and phrases.* Mixed membership, rank data, variational inference, Eurobarometer,
public policy.

and Melich (2001)]. The Eurobarometer 34.1 collected data on a broad range of topics, using a range of question formats, including measures of health behavior, knowledge of illegal drugs, descriptions of family structure, and attitudes toward children. In particular, three of the survey questions, shown in Figure 1, asked about public policy priorities toward addressing societal/public health problems. Therefore, our analysis focuses on examining responses to these questions. We should note that, if the survey contained other pertinent questions of binary or multinomial responses, these data could have been included in the analysis using the `mixedMem` R package [Wang and Erosheva (2015), R Core Team (2016)] which allows for multivariate analysis when the variables are of different distributions.

The specific variables we consider address illicit drugs, alcoholism, and AIDS. The survey respondents were asked to rank, in order of priority, policies such as punishment for offenders, information campaigns to educate the public, rehabilitation and treatment, funding of research into causes and treatment, and fighting social causes. Such priority rankings are likely to vary across respondents within a nation, reflecting dissensus among a citizenry, as well as across nations, reflecting national differences in culture, political affiliation, ideological perspective, and common practices.

The top 5 observed responses for each question, shown in Table 1, are suggestive of significant population heterogeneity. For example, for illegal drugs, a “legal penalty for drug taking” (policy 3) is highly ranked in the first and fifth most observed permutations, but ranked last in the second, third and fourth most observed permutations. This heterogeneity is not surprising because the individuals in the survey come from a wide variety of nationalities, religious backgrounds and age groups (shown in Table 2). In analyzing such heterogeneity, an important question is whether there are underlying subgroups or policy profiles among citizens. For example, when asked how they prioritize policies about the problem of illegal drugs, do citizens form a single group that varies along a dimension of punishment to rehabilitation, which reflects a conservative-liberal continuum? Or do responses

TABLE 1

The top 5 observed rankings for each question. The place in the permutation represents the ranking level and the recorded number indicates the order in which the policy option was presented in the questionnaire. See Figure 1 for the corresponding policies

Drug	Count	Alcohol	Count	AIDS	Count
1, 2, 3, 4, 5, 6, 7	109	1, 2, 8, 9, 10	85	1, 5, 4, 2, 3	740
1, 2, 4, 5, 6, 7, 3	82	1, 5, 8, 9, 10	61	5, 4, 1, 2, 3	722
2, 1, 4, 5, 6, 7, 3	69	1, 8, 9, 10, 3	60	5, 1, 4, 2, 3	670
2, 1, 4, 6, 7, 5, 3	55	1, 2, 5, 8, 10	57	1, 5, 4, 3, 2	604
2, 3, 1, 4, 5, 6, 7	53	1, 8, 9, 10, 2	53	1, 4, 5, 2, 3	515

MIXED MEMBERSHIP MODELS FOR RANK DATA

3

1	Question 28: There are various actions that could be taken to eliminate the	1
2	drugs problem. In your opinion, what is the first priority? And the next most	2
3	urgent? (Ask respondent to rank all 7, with 1 as the most urgent.)	3
4	1. Information campaigns about the dangers of drugs.	4
5	2. Hunting down drug pushers and distributors.	5
6	3. Legal penalty for drug taking.	6
7	4. Looking after and treating drug addicts and rehabilitating them.	7
8	5. Funding research into drug substitutes, and into the treatment of drug ad-	8
9	diction.	9
10	6. Fighting the social causes of drug addiction.	10
11	7. Reinforcing the control or distribution and usage of addictive medicines.	11
12		12
13		13
14	Question 39: There are various actions that could be taken in order to ease	14
15	the problem of alcoholism and its consequences. Looking at this card, which	15
16	is the main priority in your view? and the next? (Rank up to 5.)	16
17	1. Information campaigns about the danger of alcoholism.	17
18	2. Stiffer penalties for offenses committed under the influence of alcohol.	18
19	3. Banning advertising for alcoholic drinks.	19
20	4. Increasing taxes on alcohol.	20
21	5. Restricting the sale of alcohol. Especially to young people.	21
22	6. Putting lower legal limits on alcohol content.	22
23	7. Making social outcasts of alcoholics.	23
24	8. Helping alcoholics to submit "drying out."	24
25	9. Funding medical research to develop more effective treatments.	25
26	10. Setting up more reception centers, drying out treatment centers.	26
27		27
28		28
29	Question 47: There are various actions that could be taken in order to elimi-	29
30	nate the problem of AIDS or at least to slow down its development. Looking	30
31	at this card, which is the main priority in your view? And the next? (Rank all	31
32	by giving a number from 1 to 5. with 1 as a top priority.)	32
33	1. Information campaign about the danger.	33
34	2. Punishment for behavior which increases the risk.	34
35	3. Identifying and isolating those with AIDS or those who are HIV positive.	35
36	4. Treating of those with AIDS and looking after them.	36
37	5. Funding research to find a vaccine.	37
38		38
39		39
40	FIG. 1. <i>The three questions of interest from Eurobarometer 34.1 regarding illegal drugs, alco-</i>	40
41	<i>holism, and AIDS. The survey was administered in Belgium, Denmark, East Germany, France,</i>	41
42	<i>Greece, Ireland, Italy, Luxembourg, The Netherlands, Portugal, Spain, UK, and West Germany. Fur-</i>	42
43	<i>ther demographic information is given in Table 2.</i>	43

4 Y. S. WANG, R. L. MATSUEDA AND E. A. EROSHEVA

TABLE 2
Demographic breakdown of Eurobarometer 34.1 participants

Nation	Count	Nation	Count	Nation	Count
Belgium	812	Greece	980	Northern Ireland	267
Denmark	950	Ireland	879	Portugal	978
East Germany	938	Italy	1048	Spain	864
France	967	Luxembourg	242	West Germany	975
Great Britain	956	Netherlands	1016		

Religion	Count	Religion	Count	Age group	Count
Buddhist	11	Orthodox	1057	15–24 Years	2401
Hindu	2	Other	162	25–39 Years	3489
Jewish	20	Protestant	2182	40–54 Years	2780
Muslim	23	Roman catholic	5617	55+ Years	3202
None	2692				

reflect subgroups of citizens, in which some favor punitive measures, others rehabilitation, and still others education? Do some subgroups systematically oppose some policy measures, while favoring others? Moreover, given such subgroups, are some citizens members of multiple groups, favoring, for example, both rehabilitation and information? In modern democracies characterized by a diverse citizenry, such subgroups may be likely to exist. Failure to consider such subgroups, when they in fact exist, may lead to a distorted view of a nation's public. Appropriately modeling the population heterogeneity may be particularly important when the questions cover a wide range of topics, as they do here.

3. Mixed membership models.

3.1. *Previous work.* Many approaches for modeling rank data have been proposed; for a review, see Marden (1995). In this paper, we focus on the Plackett–Luce model due to several attractive attributes (discussed further in Section 4). Assuming population homogeneity, Hunter (2004) develop a minorization-maximization method for estimating MLEs for a single Plackett–Luce distribution, and Guiver and Snelson (2009) present a Bayesian framework for estimation.

Most of the previous work that address heterogeneity in rank data consider the univariate case and specify a mixture model. Gormley and Murphy (2006) assume a mixture of Plackett–Luce distributions, while Busse, Orbanz and Buhmann (2007) assume a mixture of Mallow's distributions. In addition, Bayesian nonparametric approaches have been used to allow for an infinite number of latent subgroups [Meila and Chen (2010)] and an infinite set of alternatives [Caron, Teh and Murphy (2014)]. Gormley and Murphy (2008) propose a mixture of experts

1 model where individual level covariates specify the probability that an individual 1
2 belongs to a specific subgroup. However, each of these mixture model approaches 2
3 assume that every individual always expresses preferences consistent with only a 3
4 single subgroup. In many cases, this may be overly restrictive. 4

5 Gormley and Murphy (2009) propose a mixed membership model for univariate 5
6 rank data which allows for intra-subgroup mixing between ranking levels. Mixed 6
7 membership models extend mixture models by allowing an individual membership 7
8 to be split among multiple subgroups [e.g., Airoidi et al. (2015)]. As Gross and 8
9 Manrique-Vallier (2015), page 128, point out, the structure of a mixed member- 9
10 ship model is consistent with Zaller’s (1992) model of responses to public opinion 10
11 polls, in which “respondents randomly sample from a number of privately held 11
12 “considerations” relevant to the question at hand.” Gross and Manrique-Vallier 12
13 (2015) analyzed data on political ideology, examining core beliefs and values, find- 13
14 ing that mixed membership models reveal hidden structures in political ideologies 14
15 that previous factor analytic results missed. 15

16 Direct maximum likelihood estimation of mixed membership models is gener- 16
17 ally intractable, so MCMC or approximate inference techniques are used. In partic- 17
18 ular, Gormley and Murphy (2009) develop a Metropolis-within-Gibbs sampler. Al- 18
19 though the MCMC method allows for direct sampling from the posterior, it scales 19
20 poorly. Because of the individual level parameters, the total number of param- 20
21 eters is directly proportional to sample size and the number of variables, which can 21
22 result in slow mixing even for moderate sample sizes. MCMC methods can also 22
23 require a considerable amount of human effort since convergence diagnostics must 23
24 be checked for all parameters [Gill (2008)]. 24
25

26 *3.2. Contribution of this work.* In this article, we propose a variational EM ap- 26
27 proach which scales well with the number of observed individuals and is capable 27
28 of tractably estimating mixed membership models for rank data with a much larger 28
29 sample size than can be handled by existing MCMC methods. Variational inference 29
30 has been used as an alternative estimation procedure in mixed membership models 30
31 where MCMC methods would be computationally infeasible [Blei, Ng and Jordan 31
32 (2003), Erosheva, Fienberg and Lafferty (2004), Airoidi et al. (2008)] and has been 32
33 shown empirically to provide results similar to MCMC in some cases [Erosheva, 33
34 Fienberg and Joutard (2007)]. A direct computational comparison between the pro- 34
35 posed variational method and the MCMC method detailed by Gormley and Mur- 35
36 phy (2009) is provided in the supplement [Wang and Erosheva (2017)]. 36

37 Motivated by Eurobarometer data on public policy priorities, we extend the 37
38 method of Gormley and Murphy (2009) to allow for multivariate data and directly 38
39 estimate the relative frequencies of each subgroup. The direct estimation of the 39
40 subgroup relative frequencies can be viewed as an empirical Bayes-type proce- 40
41 dure or robust modeling practice and has been shown to improve predictive perfor- 41
42 mance in many cases [Wang and Blei (2015)]. Indeed, we show in the supplement 42
43 that direct estimation of the subgroup relative frequencies drastically improves the 43

1 goodness-of-fit for exit poll data gathered during the 1997 Irish Presidential Elec- 1
 2 tion [Wang and Erosheva (2017)]. Finally, informed by exploratory data analyses, 2
 3 we extend the model specification to include two pre-specified subgroups. 3

4 The remainder of the article is structured as follows. We introduce the Plackett– 4
 5 Luce distribution in Section 4 and the mixed membership rank data model in Sec- 5
 6 tion 5. In Section 6, we review the variational approximation framework and detail 6
 7 the variational EM algorithm. Finally, we use the proposed method to analyze pub- 7
 8 lic policy priorities from the Eurobarometer 34.1 survey in Section 7 and conclude 8
 9 with discussion in Section 8. 9

10
 11 **4. Modeling rank data.** First, we consider univariate rank data. Suppose 11
 12 there are V alternatives in the choice set (items to be ranked). For each individual 12
 13 i , a single observation X_i is a permutation of $N_i \leq V$ of the alternatives in the 13
 14 choice set. Each alternative $v \in V$ (in this context, alternatives correspond to poli- 14
 15 cies) is assigned a nonnegative support parameter θ_v which governs how strongly 15
 16 it is preferred to other alternatives [Plackett (1975)]. The support parameters sum 16
 17 to 1 for identifiability. At each ranking level, one of the remaining alternatives is 17
 18 selected with probability proportional to its support parameter. The mass function 18
 19 for the Plackett–Luce model is defined as 19

$$20 \quad (1) \quad P(X_i) = \prod_{n=1}^{N_i} \frac{\theta_{a(n)_i}}{1 - \sum_{c=0}^{n-1} \theta_{a(c)_i}}, \quad 20$$

21 where $a(n)_i$ indicates the alternative selected at the n th ranking level, and $\theta_{a(0)} =$ 23
 24 0. The Plackett–Luce selection process can be thought of as a multinomial without 24
 25 replacement, and the support parameters represent the probability of a policy being 25
 26 selected as the top priority. 26

27 Although there are various distributions for modeling rank data, the Plackett– 27
 28 Luce model has several attractive properties. First, it can accommodate incom- 28
 29 plete rankings when not all alternatives in the choice set are selected. Assuming 29
 30 unranked policies are less preferred than ranked policies, the mass function in 30
 31 equation (1) has marginalized out any unranked policies. Second, the model satis- 31
 32 fies Luce’s Choice Axiom [Luce (1977)] which states that an individual’s relative 32
 33 preference between alternatives v and v' should not change when a third alternative 33
 34 v'' is introduced [Sen (2014)]. Finally, the parameter space of the Plackett–Luce 34
 35 model is a continuous set in the $V - 1$ simplex, while the parameter space of other 35
 36 popular models (notably Mallow’s model) may be discrete sets which can greatly 36
 37 complicate inference. 37

38
 39 **5. Generative model: Multivariate rank data.** We propose the following 39
 40 generative model for multivariate rank data. Assuming K latent subgroups, Dirich- 40
 41 let membership parameter $\alpha \in \mathbb{R}_{>0}^K$, J variables each with V_j alternatives, and a set 41
 42 of support parameters θ_{jk} for each variable j and subgroup k where $\sum_v \theta_{jkv} = 1$, 42
 43 the generative mixed membership model is: 43




FIG. 2. Plate notation of the Multivariate Mixed Membership Rank Data Model. Quantities λ , Z , and X are specific to individuals while quantities α and θ are global.

1. For each individual $i = 1, 2, \dots, T$.

1. Draw a membership vector $\lambda_i \sim \text{Dirichlet}(\alpha)$.

2. For each variable $j = 1, 2, \dots, J$ and for each ranking level $n = 1, 2, \dots, N_{ij}$.

(a) Draw a context indicator $Z_{ijn} \sim \text{multinomial}(1, \lambda_i)$.

(b) Draw a policy priority $X_{ijn} \sim \text{Plackett-Luce}(\theta, Z_{ijn}|\{X_{ijc}\}_{c < n})$,

where $\text{Plackett-Luce}(\theta, Z_{ijn}|\{X_{ijc}\}_{c < n})$ denotes that X_{ijn} is dependent on X_{ijc} for $c < n$ since the alternatives selected at a prior ranking levels cannot be selected again. The corresponding complete data likelihood and the graphical representation are shown in equation (2) and in Figure 2:

$$(2) \quad P(X, \Lambda, Z|\alpha, \theta) = \prod_i^T \left\{ \text{Dir}(\lambda_i|\alpha) \prod_j^J \prod_n^{N_{ij}} \left(\text{mult}(Z_{ijn}|\lambda_i) \times \prod_k^K \left[\frac{\theta_{jka(n)ij}}{1 - \sum_{c=0}^{n-1} \theta_{jka(c)ij}} \right]^{Z_{ijnk}} \right) \right\}.$$

In the model, λ_{ik} denotes the degree of membership of individual i within subgroup k , and in the Eurobarometer context, λ_{ik} indicates an individual's level of adherence to a policy ideology. Z_{ijn} is the subgroup governing individual's i selection for variable j at ranking level n . Note that the model assumes mixing of subgroup preferences occur both between different variables and within a single observed ranking. Thus, an individual may select their top choice according to the preferences of one subgroup, but select their second choice according to the preferences of another subgroup. We note that, conditional on the membership λ_i , there is no further dependence enforced on each Z_{ijn} across the ranking levels of a single variable. Although the Z_{ijn} are exchangeable in the generative model, in the posterior conditioned on the observations, the context indicators are no longer exchangeable because of the assumed sequential nature of the ranking procedure.

6. Variational EM approach.

6.1. *Variational approximation.* Calculating the marginal probability $P(X|\alpha, \theta) = \int_{Z, \Lambda} P(X, \Lambda, Z|\theta, \alpha)$ requires marginalizing over the simplicial membership vectors Λ and context indicators Z . Because this calculation is intractable, we

1 use a mean-field variational method which approximates the true posterior for latent variables Λ and Z . A detailed tutorial of variational inference is provided by Wainwright and Jordan (2008).

2 The variational distribution is

$$(3) \quad Q(Z, \Lambda | \phi, \delta) = \prod_i^T \left(\text{Dir}(\lambda_i | \phi_i) \prod_j^J \prod_n^{N_{ij}} \text{mult}(Z_{ijn} | \delta_{ijn}) \right)$$

3 with variational parameters ϕ and δ ($\phi_i \in \mathbb{R}_{>0}^K$ and δ_{ijn} lies in the $K - 1$ dimension simplex). Because it factors easily into functions of the variational parameters, this approximation facilitates tractable computation via the variational EM algorithm. We first derive a lower bound on the marginal distribution of the observed rankings using Jensen's inequality:

$$(4) \quad \begin{aligned} \log[p(X | \alpha, \theta)] &= \log \left[\int_{\Lambda, Z} \frac{Q(Z, \Lambda | \phi, \delta)}{Q(Z, \Lambda | \phi, \delta)} P(X, Z, \Lambda | \alpha, \theta) \right] \\ &= \log \mathbb{E}_Q \left[\frac{P(X, Z, \Lambda | \alpha, \theta)}{Q(Z, \Lambda | \phi, \delta)} \right] \\ &\geq \mathbb{E}_Q \{ \log[P(X, Z, \Lambda | \alpha, \theta)] \} - \mathbb{E}_Q \{ \log[Q(Z, \Lambda | \phi, \delta)] \}. \end{aligned}$$

4 The last line in equation (4) is often called the **Evidence Lower Bound (ELBO)** and is a function of the data, the variational parameters, ϕ and δ , as well as the global parameters α and θ . It can be shown that maximizing the ELBO with respect to the variational parameters ϕ and δ minimizes the KL-divergence between the true posterior and the variational distribution. In addition, fixing the variational parameters and maximizing the lower bound with respect to α and θ can be used as a surrogate procedure for selecting maximum likelihood estimates for α and θ [Beal (2003)]. Ultimately, by maximizing the lower bound (the ELBO), we simultaneously find pseudo-MLE estimates for α and θ and an approximate posterior distribution for the latent membership and context variables Λ and Z . This is accomplished by iterating between E-steps and M-steps as shown in Algorithm 1. The lower bound is given in equation (5), but the derivation is left for the Appendix:

$$(5) \quad \begin{aligned} \text{ELBO} &= \sum_i^T \log \left[\Gamma \left(\sum_k^K \alpha_k \right) \right] - \sum_i^T \sum_k^K \log[\Gamma(\alpha_k)] \\ &\quad + \sum_i^T \sum_k^K \left\{ (\alpha_k - 1) \left[\Psi(\phi_{ik}) - \Psi \left(\sum_{k'}^K \phi_{ik'} \right) \right] \right\} \\ &\quad + \sum_i^T \sum_j^J \sum_n^{N_{ij}} \left\{ \sum_k^K \delta_{ijnk} \left[\Psi(\phi_{ik}) - \Psi \left(\sum_{k'}^K \phi_{ik'} \right) \right] \right\} \end{aligned}$$

$$\begin{aligned}
& + \sum_i^T \sum_j^J \sum_n^{N_{ij}} \sum_k^K \delta_{ijkn} \left\{ [\log[\theta_{jka(n)ij}]] - \log \left[1 - \sum_{c=0}^{n-1} \theta_{jka(c)ij} \right] \right\} \\
& - \sum_i^T \log \left[\Gamma \left(\sum_k^K \phi_{ik} \right) \right] + \sum_i^T \sum_k^K \log[\Gamma(\phi_{ik})] \\
& - \sum_i^T \sum_k^K \left[(\phi_{ik} - 1) \left[\Psi(\phi_{ik}) - \Psi \left(\sum_{k'}^K \phi_{ik'} \right) \right] \right] \\
& - \sum_i^T \sum_j^J \sum_n^{N_{ij}} \sum_k^K \delta_{ijkn} \log[\delta_{ijkn}],
\end{aligned}$$

$\Gamma(\cdot)$ denotes the gamma function; $\Psi(\cdot)$ denotes the digamma function, the derivative of $\log \Gamma(\cdot)$.

6.2. *E-step: Update ϕ and δ .* The E-step maximizes the lower bound with respect to the individual level parameters ϕ and δ . Taking the derivative of the lower bound yields the following updates for ϕ and δ :

$$\begin{aligned}
(6) \quad \delta_{ijkn}^{(s+1)} & \propto \exp \left(\Psi(\phi_{ik}^{(s)}) - \Psi \left(\sum_{k'}^K \phi_{ik'}^{(s)} \right) \right. \\
& \left. + \log[\theta_{jka(n)ij}] - \log \left[1 - \sum_{c=0}^{n-1} \theta_{jka(c)ij} \right] \right) \\
\phi^{(s+1)} & = \alpha_k + \sum_j^J \sum_n^{N_{ij}} \delta_{ijkn}^{(s+1)},
\end{aligned}$$

Algorithm 1 Variational EM for Mixed Membership Rank Data

- 1: Initialize $\theta^{(0)}$, $\alpha^{(0)}$, K , $\phi^{(0)} = \frac{1}{K}$, $\delta^{(0)} = \frac{1}{K}$
 - 2: **while** (Convergence criterion not yet satisfied) **do**
 - 3: **while** (Convergence criterion not yet satisfied) **do** ▷ E-Step
 - 4: Update δ_{ijkn}
 - 5: Update ϕ_{ik}
 - 6: **end while**
 - 7: **while** (Convergence criterion not yet satisfied) **do** ▷ M-Step
 - 8: Update α by Newton–Raphson
 - 9: Update θ by interior point method
 - 10: **end while**
 - 11: **end while**
-

1 where δ (a multinomial parameter) is normalized to sum to 1. We continue updat- 1
 2 ing each parameter in a coordinate ascent procedure until the relative increase in 2
 3 the ELBO is below a specified tolerance. 3
 4

5 **6.3. M-step: Update α and θ .** The M-Step, described in Algorithm 2, fixes 5
 6 the variational parameters ϕ and δ , and selects α , the Dirichlet parameter for the 6
 7 membership vectors, and θ , the Plackett–Luce support parameters, to maximize 7
 8 the lower bound on the marginal log-likelihood. For both parameters, there are no 8
 9 closed form solutions so we use iterative updates. 9

10 For α , we use a Newton–Raphson method to maximize the lower bound: 10
 11

$$(7) \quad \begin{aligned} \frac{\partial \text{ELBO}}{\partial \alpha_k} &= T \left(\Psi \left(\sum_{k'} \alpha_{k'} \right) - \Psi(\alpha_k) \right) + \sum_i^T \left(\Psi(\phi_{ik}) - \Psi \left(\sum_{k'} \phi_{ik'} \right) \right), \\ \frac{\partial \text{ELBO}}{\partial \alpha_{k_i} \partial \alpha_{k_j}} &= -T \left(1_{\{i=j\}} \Psi'(\alpha_{k_i}) - \Psi' \left(\sum_{k'} \alpha_{k'} \right) \right). \end{aligned}$$

12 Since θ is subject to the following constraints $\sum_v^{V_j} \theta_{jkv} = 1$ and $\theta_{jkv} \geq 0$ for 12
 13 $v = 1, \dots, V_j$, we use an interior point method to select an optimal θ [Nocedal and 13
 14 Wright (1999)]. Because the constraints are only enforced on each individual set 14
 15 $\{\theta_{jkv}\}_{v=1:V_j}$ and the objective function separates into additive terms (with respect 15
 16 to the θ_{jk}), we can select θ_{jk} for each $\{j, k\}$ separately by solving the minimization 16
 17 problem: 17

$$(8) \quad \min_{\theta_{jk}} -\text{ELBO}(\theta_{jk}) + B(\theta_{jk}) \quad \text{subject to} \quad \sum_v^{V_j} \theta_{jkv} = 1,$$

28 **Algorithm 2** M-step update for θ : interior point method 29

30 **Require:** $b_0 > 0, M \in \mathbb{Z}^+, \theta^{(0)}$ 31

```

32 1: for  $j \in [J]$  do 32
33 2:   for  $k \in [K]$  do 33
34 3:     for  $m \in [M]$  do 34
35 4:       while (Convergence criterion not yet satisfied) do 35
36 5:         Calculate step direction  $\Delta\theta_{jk}$  using penalty term  $B_{b_0^m}(\theta_{jk})$  36
37 6:         Use backtracking line search to determine step length  $\tau$  37
38 7:         Set  $\theta_{jk}^{(t+1)} = \theta_{jk}^{(t)} + \tau \Delta\theta_{jk}$  38
39 8:       end while 39
40 9:     end for 40
41 10:   end for 41
42 11: end for 42

```

1 where $B(\theta) = \infty \sum_v \mathbf{1}_{[\theta_{jkv} < 0]}$ (i.e., the nonnegativity constraint on θ has been
 2 converted into a penalty term which assigns infinite loss to infeasible points). Be-
 3 cause B is not a smooth function of θ , we approximate it with the smooth function
 4 $B_b = \frac{-1}{b} \sum_v \log[\theta_{jkv}]$ and solve the relaxed minimization problem instead.

5 For notational ease, we use $\Phi_{jk}(\theta_{jk})$ to denote the objective function; H
 6 denotes Hessian of Φ_{jk} ; g denotes the gradient of Φ_{jk} and $\mathbf{1}$ denotes a row vector of
 7 1's with length V_j ;
 8

$$\begin{aligned} \frac{\partial \Phi_{jk}(\theta)}{\partial \theta_{jkv_1}} &= - \sum_i^T \sum_n^{N_{ij}} \delta_{ijnk} \left[\frac{\mathbf{1}_{\{X_{ijn}=v_1\}}}{\theta_{jkv_1}} + \frac{\sum_{c=1}^{n-1} \mathbf{1}_{\{X_{ijc}=v_1\}}}{1 - \sum_{c=0}^{n-1} \theta_{jva(c)ij}} \right] - \frac{1}{b\theta_{jkv_1}}, \\ \frac{\partial^2 \Phi_{jk}(\theta)}{\partial \theta_{jkv_2} \partial \theta_{jkv_1}} &= - \sum_i^T \sum_n^{N_{ij}} \delta_{ijnk} \left[\frac{\sum_{c=1}^{n-1} \mathbf{1}_{\{X_{ijc}=v_1\}} \sum_{c=1}^{n-1} \mathbf{1}_{\{X_{ijc}=v_2\}}}{(1 - \sum_{c=0}^{n-1} \theta_{jka(c)ij})^2} \right], \\ \frac{\partial^2 \Phi_{jk}(\theta)}{\partial \theta_{jkv_1}^2} &= - \sum_i^T \sum_n^{N_{ij}} \delta_{ijnk} \left[\frac{-\mathbf{1}_{\{X_{ijn}=v_1\}}}{\theta_{jkv_1}^2} + \left(\frac{\sum_{c=1}^{n-1} \mathbf{1}_{\{X_{ijc}=v_1\}}}{1 - \sum_{c=0}^{n-1} \theta_{jka(c)ij}} \right)^2 \right] \\ &\quad + \frac{1}{b\theta_{jkv_1}^2}. \end{aligned}$$

9 Satisfying the Karush–Kuhn–Tucker conditions with the remaining equality con-
 10 straint yields the update direction $\Delta\theta$ for $\theta^{(s+1)} = \theta^{(s)} + \Delta\theta$, where
 11

$$(9) \quad \Delta\theta = -H^{-1} \left(g - \mathbf{1}^T \frac{\mathbf{1}H^{-1}g}{\mathbf{1}H^{-1}\mathbf{1}^T} \right).$$

12 Because the Newton step in equation (10) uses a quadratic approximation of
 13 the objective function, the proposed $\Delta\theta$ increment may be ill-sized. If the step
 14 size is too large, the increment may actually lead to a larger value of the objective
 15 function or infeasible updates where $\theta_{jkv} < 0$. Thus, we use a backtracking line
 16 search, detailed in Algorithm 3, to ensure that each update will always increase the
 17 lower bound and remain in the feasible set.
 18

19 **6.4. Algorithm discussion.** For numerical stability, we first solve the mini-
 20 mization for a small value of b , and then use that solution to initialize subsequent
 21

22 **Algorithm 3** M-step update for θ : backtracking line search

23 **Require:** $\theta_{jk}, \Delta\theta_{jk}, \tau_0 \in (0, 1)$

24 1: $s = 0$

25 2: **while** Any($\theta_{jkv} + \tau_0^s \Delta\theta_{j,k} < 0$) or [$\text{ELBO}(\theta_{jkv} + \tau_0^s \Delta\theta_{jk}) < \text{ELBO}(\theta_{jkv})$] **do**

26 3: $s++$

27 4: **end while**

28 5: **return** $\tau = \tau_0^s$

1 minimizations with increasingly larger values of b [Nocedal and Wright (1999)]. 1
 2 In addition, if a particular variable has a large number of options, inverting the 2
 3 $V_j \times V_j$ Hessian can become computationally expensive. A quasi-Newton or gra- 3
 4 dient ascent procedure may require less overall computation by avoiding large 4
 5 matrix inversions. We also note that the sub-routines for α and each θ_{jk} can be 5
 6 computed completely in parallel which may be fruitful if the number of subgroups 6
 7 or variables is large. 7

8 As with almost all variational methods, the objective function is multi-modal, so 8
 9 only convergence to a local maximum is guaranteed. Using available prior knowl- 9
 10 edge can be very helpful in determining reasonable initializations for a specific 10
 11 problem; however, multiple initialization points are recommended to increase the 11
 12 probability of finding the global maximum. 12

13 We found empirically that initializing α and θ with the following heuristically 13
 14 driven two step procedure generally resulted in stationary points with a larger 14
 15 ELBO. Using a random initialization of θ and α and setting all values of ϕ and 15
 16 δ to $1/K$, iterate the variational EM procedure until reaching a stationary point $\tilde{\alpha}$ 16
 17 and $\tilde{\theta}$. Then, use the resulting global parameters $\tilde{\alpha}$ and $\tilde{\theta}$ to initialize a second run 17
 18 (with ϕ and δ reset to $1/K$). The result of the first run is only a stationary point 18
 19 with respect to all the parameters (both global and individual), so resetting the ϕ 19
 20 and δ parameters will generally result in a new stationary point where the global 20
 21 parameters $\hat{\theta}$ and $\hat{\alpha}$ are different than the intermediate initialization points. 21

22 Because the dimension of even just the α and θ parameters can be quite large, a 22
 23 huge number of random restarts may be needed to explore the space well when se- 23
 24 lecting uniformly. We posit that using this two-step procedure to find initialization 24
 25 points concentrates the search in areas where the ELBO is likely to be larger. 25

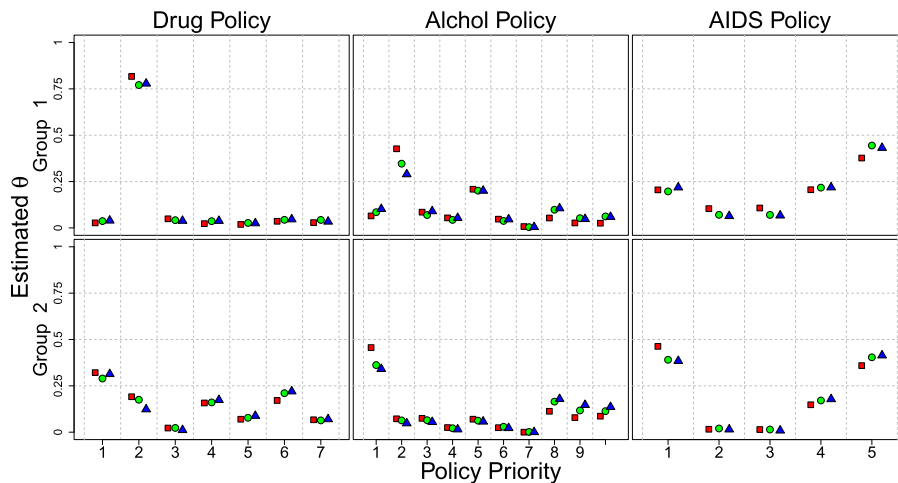
26
 27 **7. Eurobarometer analysis.** We now analyze rank data from the Eurobarom- 27
 28 eter 34.1 survey [Reif and Melich (2001)]. We removed individuals with missing 28
 29 data (i.e., anyone who did not respond to all 3 questions of interest) and individuals 29
 30 who had reported ties in any of their rankings, leaving 11,872 individuals of the 30
 31 original 12,733. 31

32
 33 *7.1. Model selection.* Table 1 shows the top 5 observed rankings for each 33
 34 question. In particular, we observe that the response which ranks the policy pri- 34
 35 orities in the exact order of presentation is the most common pattern for drug 35
 36 priorities and is also the 6th and 28th most common pattern for alcohol and AIDS. 36
 37 If some individuals used this ordering out of convenience, the resulting responses 37
 38 may not be informative of true policy preferences. To capture this tendency, we 38
 39 include a subgroup whose preferences correspond with the presentation-ordered 39
 40 permutation. Following Gormley and Murphy (2006), we also include a “noise” 40
 41 subgroup whose preferences are uniform across all policy priorities. We fix the 41
 42 support parameters, θ , for these groups, but estimate their relative frequencies (the 42
 43 corresponding elements of α). This approach is similar to the extended grade of 43

1 membership model [Erosheva, Fienberg and Joutard (2007)] which also models 1
 2 specific response patterns with unusually high counts; however, here we allow 2
 3 for partial membership in the fixed subgroups, whereas Erosheva, Fienberg and 3
 4 Joutard (2007) assume that some individuals are full members of the fixed sub- 4
 5 groups. 5

6 We use a held out ELBO procedure to select an appropriate number of sub- 6
 7 groups K . We randomly split the sample in half to create a training set and test 7
 8 set. For each $K = 3, 4, \dots, 10$ (we do not include $K = 2$ because that would only 8
 9 be the 2 fixed groups), we first fit a model to the training set and select global pa- 9
 10 rameters $\hat{\alpha}$ and $\hat{\theta}$. Then, to compute the held out ELBO on the test set, we use a 10
 11 single E-step which fits the individual variational parameters ϕ and δ for the test 11
 12 set given the $\hat{\alpha}$ and $\hat{\theta}$ from the training set. We use 40 different initialization points 12
 13 $\theta_0 \sim \text{Dir}(\{a, a, \dots, a\})$ at each $a = 0.6, 1.1, \text{ and } 1.5$ and select the stationary point 13
 14 across all K with the highest resulting held out ELBO. We then used the stationary 14
 15 point selected by the procedure to initialize a final run with the results presented 15
 16 below. 16

17 To ensure that the model interpretation is not dependent on the selected train- 17
 18 ing/test set, we repeated this procedure with 3 different training/test splits. For 18
 19 each batch, the held out ELBO values do vary widely within a fixed K due to the 19
 20 multi-modality of the ELBO. However, we see the same trend in all 3 cases; the 20
 21 largest held out ELBO values for each $K = 3, 4, 5, 6$ are somewhat close and peak 21
 22 at either 4 or 5 and the maximum held out ELBO values decrease rapidly as K 22
 23 increases beyond 6. Of the three batches, the first batch selects a 5 subgroup 23
 24 (including the two fixed groups) and the other two batches select a 4 subgroup 24
 25



26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41 FIG. 3. The estimated support parameters of two largest subgroups of the selected 5 subgroup 41
 42 model from batch 1 (squares) is plotted against the estimated support parameters from batch 2 (cir- 42
 43 cles) and batch 3 (triangles). 43




FIG. 4. 1000 simulated top priority rankings using the fitted $\hat{\alpha}$ and $\hat{\theta}$ values are shown with the boxplots. The observed counts are indicated by the x's.

model. Figure 3 shows that the two largest subgroups of the first batch (5 subgroup model) are very similar in structure to the two nonfixed subgroups of the second and third batches (4 subgroup models). Since the 5 subgroup model has the largest ELBO, we describe that model in the remainder of this article.

7.2. *Goodness of fit.* To check goodness-of-fit, we generate 1000 simulated data sets using the fitted values $\hat{\alpha}$ and $\hat{\theta}$. Figure 4 shows that the model captures the general trend of observed counts for the first place rankings of each variable. These plots are not quite posterior predictive checks because α and θ are fixed so the variability of the simulated outcomes is smaller than if α and θ were also considered random quantities.

7.3. *Model interpretation.* Table 3 presents the ratio of the estimates $\hat{\theta}$ and uniform support parameters (i.e., $\hat{\theta}_{jkv}/\frac{1}{V_j}$). Thus, the reported values represent how many times more likely a full member of a subgroup would be to select a specific policy as their top priority compared to an individual selecting policies randomly. A value larger than 1 indicates that the policy is more popular than average for the variable and subgroup, and a value less than 1 indicates that the policy is less popular than average. The \log_{10} of these values are also represented in Figure 5 where priorities favored more than average have a positive bar height and priorities favored less than average have a negative bar height. Furthermore, within each subgroup, priorities are sorted by estimated support allowing readers to more easily characterize subgroup preferences.

Subgroup 1 generally favors punitive policies. For illegal drugs, the top two priorities are “Punish dealers” and “Penalize users,” and subgroup 1 is 44 times more likely to select “Punish dealers” than the least favored option of “Funding research.” Similarly, for alcohol, the most popular policies are “Stricter penalties for offenses” and “Restricting sale.” Although “Ostracizing alcoholics” is the least popular policy for all subgroups, in subgroup 1, it is only 55 times less likely than the top priority while it is 260 times less popular than the top option for subgroup 3

1 and numerically zero for subgroup 2.¹ For AIDS, although the two punitive options 1
 2 (“Punish behavior” and “Isolate patients”) are the least favored policies, subgroup 2
 3 1 is only roughly 3.5 times less likely to select these two options relative to the 3
 4 most popular option of “Funding research.” 4

5 Subgroup 2 generally prioritizes “Information campaigns.” “Information cam- 5
 6 paigns” are roughly 1.7 times more likely to be selected as the top priority than the 6
 7 second most popular option of “Penalizing dealers” and 14 times more likely than 7
 8 the least popular option “Penalize users.” For alcohol, “Information campaigns” 8
 9 are 4 times more likely than the second most popular alternative “Rehabilitate al- 9
 10coholics.” The least popular option is “Penalize users,” with an estimated support 10
 11 parameter that is numerically zero; the two other least popular policies include 11
 12 “Increasing taxes” and “Lowering limits.” The dislike for these options is consis- 12
 13 tent with the idea of limited government social intervention. For AIDS, subgroup 13
 14 2 is the only subgroup for which “Funding research” is not the most popular op- 14
 15 tion, with “Information campaigns” roughly 1.3 times more likely than “Funding 15
 16 research.” 16

17 Finally, subgroup 3 typically supports rehabilitation and treatment, as well as 17
 18 funding research. For illegal drugs, although the most popular policy is “Treat- 18
 19ing addicts,” “Punishing dealers,” and “Addressing social causes” are also popular 19
 20 policies. For alcoholism, “Rehabilitation” is by far the most popular policy with 20
 21 “Funding research” and “Increasing resources for rehabilitation” as the only other 21
 22 options with substantial support. For AIDS, this subgroup expresses strong support 22
 23 for “Treating AIDS” and “Funding research.” 23

24 Broadly speaking, the identified groups are consistent with the typical Left (lib- 24
 25eral) versus Right (conservative) political ideology archetypes. The focus on puni- 25
 26 tive measures is consistent with a right leaning approach toward governance while 26
 27 the focus on information and rehabilitation typifies a more left leaning approach 27
 28 [Cavadino and Dignan (2006)]. 28

29 Shown in Table 4, the small magnitude of $\hat{\alpha}$, the Dirichlet membership param- 29
 30eter, suggests relatively low levels of intra-individual mixing. However, the modal 30
 31 grade of membership in Figure 6 shows that a quarter of all individuals still ex- 31
 32hibit significant intra-individual mixing. We also see that the nontrivial relative 32
 33 frequency estimates of the noninformative fixed groups justify their inclusion in 33
 34 the analysis. 34

35 The Dirichlet distribution for λ_i implicitly enforces negative dependence be- 35
 36tween subgroup memberships, although positive dependencies could be modeled 36
 37 using distributions considered by Blei and Lafferty (2005). However, the magni- 37
 38tude of correlations between subgroup memberships is still informative. In Table 5, 38
 39 the estimated membership in the informative subgroups is more strongly correlated 39
 40 40

41
 42 ¹For alcohol, there are 10 options and at most 5 ranking levels. Thus, it may be possible for an 42
 43 option to appear extremely infrequently or not at all. 43

16 Y. S. WANG, R. L. MATSUEDA AND E. A. EROSHEVA

TABLE 3

The estimated support parameters divided by the support parameter under uniform selection. The bootstrapped 95% confidence intervals are shown in parentheses

	Subgroup 1 <i>Punitive subgroup</i>	Subgroup 2 <i>Info subgroup</i>	Subgroup 3 <i>Rehab subgroup</i>
Drug policy			
Inform public	0.191 (0.175, 0.209)	2.249 (2.17, 2.324)	0.745 (0.694, 0.8)
Punish dealers	5.719 (5.623, 5.808)	1.338 (1.27, 1.404)	1.411 (1.304, 1.527)
Penalize users	0.344 (0.316, 0.376)	0.157 (0.147, 0.167)	0.225 (0.204, 0.246)
Treat addicts	0.162 (0.149, 0.176)	1.103 (1.064, 1.144)	1.456 (1.373, 1.542)
Fund research	0.132 (0.122, 0.143)	0.489 (0.469, 0.509)	0.966 (0.913, 1.017)
Social causes	0.251 (0.231, 0.272)	1.197 (1.149, 1.247)	1.386 (1.313, 1.461)
Control medicine	0.202 (0.186, 0.219)	0.468 (0.449, 0.489)	0.81 (0.767, 0.853)
Alcohol policy			
Inform public	0.648 (0.609, 0.689)	4.565 (4.421, 4.709)	0.324 (0.296, 0.355)
Penalize off	4.269 (4.104, 4.432)	0.722 (0.681, 0.767)	0.192 (0.168, 0.218)
Ban ads	0.85 (0.798, 0.904)	0.744 (0.702, 0.787)	0.104 (0.089, 0.121)
Inc taxes	0.541 (0.509, 0.576)	0.25 (0.23, 0.271)	0.045 (0.036, 0.055)
Rest sale	2.086 (1.994, 2.181)	0.7 (0.662, 0.741)	0.391 (0.347, 0.438)
Low limits	0.472 (0.445, 0.5)	0.243 (0.224, 0.262)	0.134 (0.117, 0.154)
Ostze alcs	0.074 (0.065, 0.084)	0.00 (0.00, 0.00)	0.022 (0.017, 0.028)
Rehab alcs	0.535 (0.505, 0.563)	1.128 (1.085, 1.171)	5.194 (5.011, 5.38)
Fund research	0.27 (0.25, 0.288)	0.787 (0.751, 0.822)	1.788 (1.69, 1.887)
Inc resources	0.254 (0.234, 0.274)	0.862 (0.825, 0.898)	1.805 (1.718, 1.893)
AIDS policy			
Inform public	1.026 (0.984, 1.067)	2.314 (2.234, 2.393)	0.646 (0.598, 0.7)
Punish behavior	0.521 (0.495, 0.548)	0.079 (0.072, 0.086)	0.251 (0.228, 0.276)
Isolate patients	0.537 (0.508, 0.57)	0.072 (0.065, 0.079)	0.199 (0.176, 0.225)
Treat AIDS	1.031 (0.993, 1.068)	0.74 (0.705, 0.775)	1.607 (1.524, 1.685)
Fund research	1.885 (1.819, 1.951)	1.796 (1.725, 1.869)	2.296 (2.195, 2.4)

TABLE 4

The estimated α Dirichlet membership parameter and the 95% bootstrapped confidence intervals. The estimated relative frequency of each subgroup is given by $\frac{\alpha_k}{\sum_{k'} \alpha_k}$

	$\hat{\alpha}$	Relative frequency
Subgroup 1	0.05 (0.048, 0.053)	0.322 (0.314, 0.329)
Subgroup 2	0.048 (0.047, 0.051)	0.31 (0.302, 0.318)
Subgroup 3	0.024 (0.023, 0.026)	0.154 (0.149, 0.16)
Subgroup 4	0.014 (0.013, 0.015)	0.088 (0.084, 0.092)
Subgroup 5	0.02 (0.019, 0.021)	0.126 (0.122, 0.131)

MIXED MEMBERSHIP MODELS FOR RANK DATA

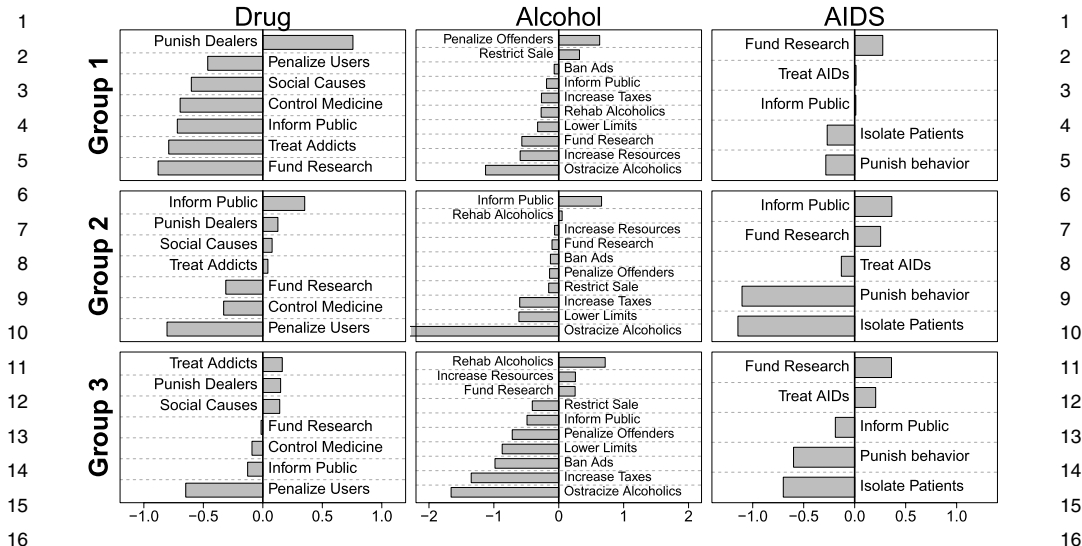


FIG. 5. Each barplot in the figure shows $\log_{10}(\hat{\theta}_{jkv}/(1/V_j))$ so that priorities which are more likely to be selected than average have a positive bar height and priorities which are less likely to be selected than average have a negative bar height. The priorities for each subgroup are reordered vertically by largest estimated support to smallest estimated support. Note that the bar for “Ostracize alcoholics” for subgroup 2 has been truncated.

with membership in the other informative subgroups (the least negative correlation between subgroups 1–3 is -0.29) than membership in the fixed groups (the most negative correlation between subgroups 1–3 versus subgroups 4–5 is -0.25). This is not surprising because subgroups 1–3 indicate a particular ideology on public policy, while subgroups 4 and 5 essentially represent the lack of preferences which align with the dominant subgroups.




FIG. 6. The estimated degree of membership in each individual’s modal subgroup.

TABLE 5
Correlation between the estimated subgroup memberships for each subgroup

	Subgroup 1	Subgroup 2	Subgroup 3	Subgroup 4	Subgroup 5
Subgroup 1	1.00	-0.55	-0.29	-0.25	-0.07
Subgroup 2	-0.55	1.00	-0.36	-0.25	-0.18
Subgroup 3	-0.29	-0.36	1.00	-0.16	-0.03
Subgroup 4	-0.25	-0.25	-0.16	1.00	-0.04
Subgroup 5	-0.07	-0.18	-0.03	-0.04	1.00

7.4. *Uncertainty estimates.* Because $\hat{\alpha}$ and $\hat{\theta}$ were selected through a pseudo-MLE procedure, there are no readily available model-based standard errors; however, we estimate standard errors via an empirical bootstrap procedure. For each bootstrap sample $b = 1, \dots, 5000$, we select 11,872 individuals with replacement from the observed sample and use the variational EM procedure to select pseudo-MLE estimates $\hat{\alpha}^{(b)}$ and $\hat{\theta}^{(b)}$. Each bootstrap sample run is initialized at the same starting points used for the full model. This initialization avoids overestimating variability in stationary points due to multi-modality of the objective function and seeks to only capture sampling variability. To form 95% confidence intervals, we take the 0.025 and 0.975 quantiles of the bootstrapped estimates.

7.5. *Multivariate versus univariate model.* We also acknowledge the implicit decision to use a multivariate model instead of fitting a univariate model for each question. Under a univariate models, subgroup membership λ for each individual is estimated independently of the responses to other questions. By contrast, in a multivariate specification, individuals can still exhibit a different mix of subgroups across each question and ranking level, but the posterior estimates are shrunk toward the individual's overall membership λ_i .

As a sensitivity analysis, we fit univariate models for each question. Fewer subgroups may be necessary when considering a univariate model when compared a multivariate model, but since the univariate models are only used to validate the structure of subgroups identified with multivariate data, we fit models with 5 subgroups and do not repeat the model selection procedure. The estimated support parameters θ do not differ substantially from the multivariate model, but the estimated membership parameters α differ across univariate models in an informative way. Table 6 shows a much higher proportion of membership in the information subgroup for the AIDS univariate model than we see in the drugs univariate, alcohol univariate, or full multivariate models. This suggests that, on average, individuals have a stronger preference for "Information campaigns" to address AIDS, which is not seen as strongly in addressing alcohol or drugs. As expected, the relative frequencies when averaged across all three univariate models are similar to the relative frequencies of the full model.

MIXED MEMBERSHIP MODELS FOR RANK DATA

TABLE 6
The estimated $\hat{\alpha}$ and corresponding relative frequency for each of the univariate and full multivariate models

	Subgroup 1	Subgroup 2	Subgroup 3	Subgroup 4	Subgroup 5
Drug univariate est.	0.031	0.020	0.015	0.010	0.011
Alcohol univariate est.	0.062	0.042	0.042	0.021	0.027
AIDS univariate est.	0.006	0.060	0.018	0.014	0.012
Full model est.	0.050	0.048	0.024	0.014	0.020
Drug univariate rel. freq.	0.362	0.229	0.169	0.113	0.128
Alcohol univariate rel. freq.	0.321	0.217	0.218	0.107	0.138
AIDS univariate rel. freq.	0.052	0.549	0.160	0.126	0.113
Full model rel. freq.	0.322	0.310	0.154	0.088	0.126
Average univariate rel. freq.	0.259	0.319	0.181	0.115	0.127

7.6. *Individual membership estimates.* We examine two specific individuals that illustrate the richness of description afforded by using a mixed membership model.

Table 7 shows the observed responses from a 68-year-old British male. For addressing drugs, his responses follow the presentation ranking, but for alcohol and AIDS, he indicates a preference for research and rehabilitation. We estimate the membership for this man to be 66% subgroup 3 (Rehab and Research) and 34% subgroup 5 (Presentation Ordering). Because the mixed membership framework allows for intra-individual mixing, the rank ordered response for drug policy is attributed to the noninformative subgroup. This contrasts with a finite mixture model approach, which would otherwise include this noisy response for drug policy in the estimates for subgroup 3.

Table 8 shows the responses of a 40-year-old Spanish female. Her perspectives on alcohol policy differ drastically from her preferences for drugs and AIDS. We see that her top policies for drugs and AIDS are highly punitive, but information

TABLE 7
Observed Responses from a 68-year-old British male

	Drug policy	Alcohol policy	AIDS policy
Priority 1	Inform public	Fund research	Isolate patients
Priority 2	Punish dealers	Increase resources	Treat AIDS
Priority 3	Penalize users	Rehabilitate alcoholics	Fund research
Priority 4	Treat addicts	Ban advertisements	Inform public
Priority 5	Fund research	Inform public	Punish behavior
Priority 6	Fight social causes		
Priority 7	Control medicine		

TABLE 8
Observed responses from a 40-year-old Spanish female

	Drug policy	Alcohol policy	AIDS policy
Priority 1	Punish dealers	Inform public	Punish behavior
Priority 2	Penalize users	Rehabilitate alcoholics	Isolate patients
Priority 3	Treat addicts	Restrict sale	Treat AIDS
Priority 4	Inform public	Ban advertisements	Fund research
Priority 5	Fight social causes	Increase resources	Inform public
Priority 6	Control medicine		
Priority 7	Fund research		

campaigns and rehabilitation are preferred for alcoholism. These different perspectives are captured in the model with an estimated membership of 63% in subgroup 1 and 37% in subgroup 2.

7.7. *Membership by demographic subgroup.* The broad interpretation of our results agree with previous studies which have identified demographic characteristics associated with general dispositions toward penal ideology. To examine these demographic trends clearly, we filter out individuals whose membership in subgroups 4 and 5 (the noninformative subgroups) is over 50%. This leaves 10,448 of the 11,872 original individuals. We then examine the conditional membership of the remaining individuals in subgroups 1, 2, and 3 (i.e., $\tilde{\lambda}_i = \frac{\lambda_i}{\lambda_1 + \lambda_2 + \lambda_3}$).

We first examine a self reported measure of Left versus Right political ideology. Individuals were asked: “In political matters, people talk of the left or the right. How would you place your views on this scale?” In the recorded scale, 1 indicates far left and 10 indicates far right. This is not a perfect analog since each individual likely responded in reference to their national definition of “center” whereas the subgroup membership estimated from the rank data is a global measure. Nonetheless, we see that there is a very significant Spearman’s rank-order correlation of 0.15 (p-value $< 2e-16$) between self-reported Left versus Right score from Eurobarometer 34.1 and membership in the “Punitive subgroup.”

We also examine the average membership across religious affiliation. Some speculate that Anglo-Saxon cultures are particularly punitive because of Protestant religions with strong Calvinistic overtones [Tonry (2007)] or fundamentalist beliefs [Grasmick et al. (1992)]. As shown in Table 2, almost all individuals in the survey report their religion as either Roman catholic, None, Protestant, or Orthodox. In addition to denomination, the Eurobarometer also recorded how often an individual attends religious services. We collapse the original categories of “Several times a week” and “Once a week” to a single “Regular attendance” category and collapse “Few times a year,” “Once a year,” and “Never” into an “Irregular attendance” category. Table 9 shows that those who attend religious services

MIXED MEMBERSHIP MODELS FOR RANK DATA

21

TABLE 9

Average conditional membership by denomination and regularity of attending religious services. Regular attendance indicates attending a religious service at least once a week; irregular attendance indicates attending a religious service less than once a week

Religion	Attendance	Subgroup 1	Subgroup 2	Subgroup 3
None	NA	0.36	0.45	0.19
Orthodox	Irregular	0.21	0.58	0.21
Orthodox	Regular	0.30	0.47	0.23
Protestant	Irregular	0.40	0.40	0.20
Protestant	Regular	0.53	0.31	0.16
Roman catholic	Irregular	0.40	0.39	0.21
Roman catholic	Regular	0.47	0.34	0.18

regularly have a much higher average membership in subgroup 1. Also, Roman Catholics and Protestants are much more likely to belong to subgroup 1 than individuals who report no religion or Orthodox Christians. We note that roughly 90% of the individuals who responded as Orthodox Christians were Greek and roughly 97% of Greek respondents reported their religion as Orthodox Christianity. This confounding may be the cause of the particularly low subgroup 1 membership for Orthodox Christians.

In addition, we examine the average estimated membership across levels of education. The Eurobarometer asks “How old were you when you finished full-time education?” The average membership in subgroup 1 (punitive) decreases steadily as education increases, a finding that is consistent with previous research on industrialized countries, including Western Europe [e.g., Kitschelt and Rehm (2014), Mayhew and Van Kesteren (2002)].

At the national level, the average memberships are also consistent with qualitative characterizations of national policy. The United Kingdom and Ireland have high average memberships in subgroup 1 (punitive), while Denmark and France are among a cluster of countries with low average memberships in subgroup 1. These findings are generally consistent with previous research using the International Crime Victimization Surveys, which finds punitive attitudes in the United

TABLE 10

Average conditional membership by “Last age of formal education”

Last age of formal education	Subgroup 1	Subgroup 2	Subgroup 3
16 or less	0.47	0.33	0.19
17 to 19	0.40	0.40	0.20
20 to 21	0.33	0.48	0.20
22 or older	0.26	0.54	0.20

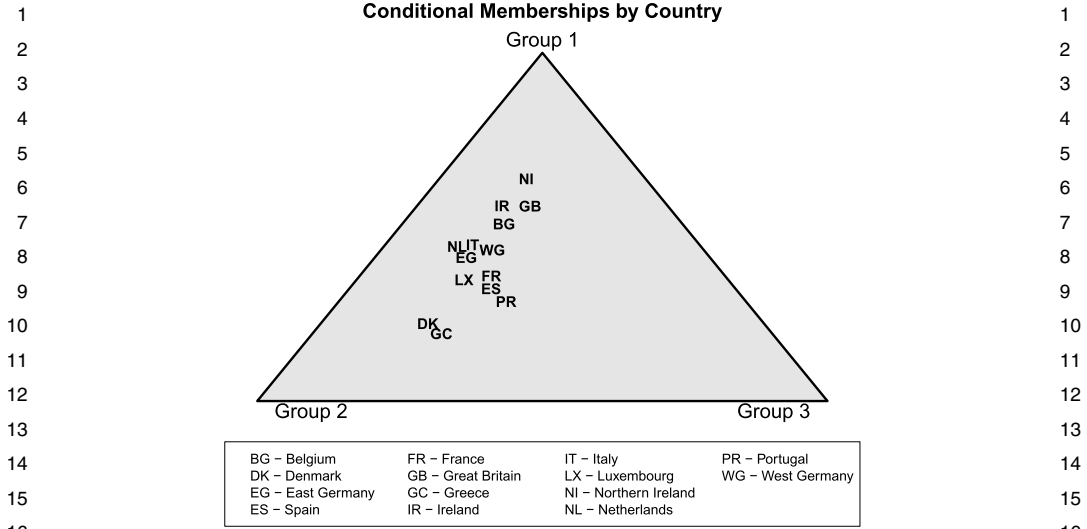


FIG. 7. The estimated average membership of each country where the noninformative fixed groups (subgroups 4 and 5) have been marginalized out.

Kingdom and Ireland, and nonpunitive attitudes in Denmark and France [e.g., Roberts (2013)]. In slight contrast to that work, we find Belgium to have a relatively high membership in subgroup 1, and the Netherlands to have a relatively low membership.

8. Discussion. In this article, we propose a mixed membership model for multivariate rank data and develop a variational EM estimation approach that is a computationally attractive alternative to fully Bayesian estimation for large scale rank data. Mixed membership models provide valuable insights into latent structure within a heterogeneous population and allow for a richer description when compared to previous mixture model approaches. When MCMC is tractable for smaller data sets as in Gormley and Murphy (2009), the results provide direct samples from the posterior. Nevertheless, the demands placed on human and computer time to conduct such an analysis can be substantial, and scalability of MCMC methods is poor. Ultimately, a mixed membership analysis of larger data sets necessitate other approaches. Of course, what actually qualifies as “large scale” or “big data” is dependent on the complexity of analysis. In rank data, the complexity quickly grows as the number of variables and alternatives increase.

In addition to the computational gains, the proposed method extends the method of Gormley and Murphy (2009) to explicitly fit the Dirichlet membership parameter α . Unless there is strong prior knowledge about subgroup sizes, this extension can result in better fitting models by directly capturing from the data differences in the subgroup structure and the level of intra-group mixing. A direct comparison of

1 both goodness of fit and computational effort is provided in the supplement [Wang 1
2 and Erosheva (2017)]. 2

3 To accommodate multivariate ranked data, our model makes the simplest as- 3
4 sumption that all context indicators Z_{ijn} are drawn from the same multinomial 4
5 distribution governed by a single membership vector λ_i . An alternative and more 5
6 complex model might include an additional layer of hierarchy between λ_i and 6
7 Z_{ijn} for each variable j . This would allow context indicators Z_{ijn} from separate 7
8 variables to be drawn from different distributions, while still respecting the multi- 8
9 variate structure. 9

10 There are drawbacks, however, to the variational approximation. Because of 10
11 the multi-modal objective function, many random restarts should be used. Prior 11
12 knowledge can be used to select good initialization points, but finding a global 12
13 maximum is not guaranteed. We propose a two-step procedure for initialization, 13
14 but addressing multi-modality through stochastic optimization methods [Bottou 14
15 (2010)] or placing strong priors on the support parameters to induce “smoothness” 15
16 in the ELBO are two natural extensions. 16

17 Also, unlike a full Bayesian specification, the variational EM method does not 17
18 provide a posterior for the global parameters. Frequentist uncertainty estimates, 18
19 however, can still be achieved through a bootstrap procedure, but each boot- 19
20 strapped model must be carefully initialized to avoid overestimating variability. To 20
21 our knowledge, bootstrapping with variational estimation has not been previously 21
22 used in the existing literature. 22

23 As with any mixture or mixed membership model, selecting the number of 23
24 subgroups is difficult. Our model selection procedure involves cross-validation of 24
25 the held-out ELBO. This procedure, however, can be complicated by the multi- 25
26 modality of the objective function and the selected model might depend on the 26
27 specific test and training sets. BIC procedures are also widely used although the 27
28 theoretical justification does not hold in mixed membership models [Airoldi et al. 28
29 (2015)]. Alternative approaches include stability-based measures [Lange et al. 29
30 (2002)], direct goodness-of-fit measures [Cohen and Mallows (1983)], and non- 30
31 parametric model extensions such as those based on Dirichlet processes [Teh et al. 31
32 (2006)]. 32

33 Analyzing the Eurobarometer 34.1 data, we find three informative policy pref- 33
34 erence subgroups as well as substantial support for a uniform ranking group and a 34
35 presentation-ordered group. The three informative subgroups primarily favor puni- 35
36 tive policies, information campaigns, and rehabilitation and research, respectively. 36
37 When comparing subgroup membership to educational, religious and national de- 37
38 mographic information, we see trends which generally agree with the existing lit- 38
39 erature. In particular, fewer years of formal education and more religious partici- 39
40 pation is generally associated with more punitive attitudes toward social issues. 40
41 In addition, at the national level, average subgroup membership roughly agrees with 41
42 previous characterizations of national punitive attitudes. 42
43 43

1 Finally, our analysis has implications for survey development. Because of a siz- 1
 2 able presentation-ordered subgroup in our analysis, we recommend randomizing 2
 3 the presentation of choices when collecting rank data to decrease bias due to non- 3
 4 informative responses where respondents rank choices by simply following the 4
 5 presentation order. We also note that the variable with the largest proportion of 5
 6 presentation-ordered responses is the question regarding illegal drugs which also 6
 7 happens to allow up to 7 rankings, while the other two questions only allow up to 7
 8 ranking levels. This observation naturally leads to speculation of whether decreas- 8
 9 ing the number of ranking levels and cognitive load may ultimately lead to more 9
 10 “informative” responses. 10

11 Although our analysis focused on issues within political science, sociology, and 11
 12 public health, multivariate rank data can elicit and capture a rich representation of 12
 13 individual preferences. We believe that the proposed methodology will be of broad 13
 14 interest. Psychologists, economists, other social scientists, and marketing profes- 14
 15 sionals who analyze large scale rank data can rely on the proposed methodology 15
 16 to represent large scale ranked preferences with realistic models which are still 16
 17 parsimonious and easily interpretable. 17

18 APPENDIX 18

19
 20 **Derivation of lower bound on marginal log-likelihood (ELBO).** The 20
 21 derivation of the lower bound from equation (4) is shown here. The lower bound 21
 22 is 22

$$23 \log[P(X|\alpha, \theta)] \geq \mathbb{E}_Q(\log(p(X, Z, \Lambda))) - \mathbb{E}_Q(\log(Q(Z, \Lambda))). 24$$

25 Note that $\Gamma(\cdot)$ denotes the gamma function, $\Psi(\cdot)$ denotes the digamma function, 25
 26 and f_j denotes the Plackett-Luce mass function of variable j . X_{ij} denotes the 26
 27 observation of N_{ij} level rankings and $a(n)_{ij}$ indicates the alternative selected by 27
 28 individual i for variable j at ranking level n . Note that for all multinomial mass 28
 29 functions shown below, the size = 1. 29

30 We consider each piece of the lower bound separately. The log likelihood for 30
 31 the complete data is 31

$$32 \log[p(X, Z, \Lambda|\alpha, \theta)] 32$$

$$33 = \log \left[\prod_i^T \left(\text{Dir}(\lambda_i|\alpha) \prod_j^J \left[\prod_n^{N_{ij}} \text{mult}(Z_{ijn}|\lambda_i) \right] f_j(X_{ij}|\theta, Z_{ij}\cdot) \right) \right] 34$$

$$35 (11) = \sum_i^T \log[\text{Dir}(\lambda_i|\alpha)] + \sum_i^T \sum_j^J \sum_n^{N_{ij}} \log[\text{mult}(Z_{ijn}|\lambda_i)] 35$$

$$36 + \sum_i^T \sum_j^J \log[f_j(X_{ij}|\theta, Z_{ij}\cdot)]. 36$$

37 37
 38 38
 39 39
 40 40
 41 41
 42 42
 43 43

1 The expectation of the first term with respect to the variational distribution Q 1
 2 becomes 2

$$\begin{aligned}
 & \sum_i^T \mathbb{E}_Q \log[\text{Dir}(\lambda_i | \alpha)] \\
 &= \sum_i^T \mathbb{E}_q \left\{ \log \left[\frac{\Gamma[\sum_k^K \alpha_k]}{\prod_k^K \Gamma(\alpha_k)} \prod_k^K \lambda_{ik}^{\alpha_k - 1} \right] \right\} \\
 &= \sum_i^T \log \left[\Gamma \left(\sum_k^K \alpha_k \right) \right] - \sum_i^T \sum_k^K \log[\Gamma(\alpha_k)] \\
 (12) \quad &+ \sum_i^T \sum_k^K (\alpha_k - 1) \mathbb{E}_q \{ \log[\lambda_{ik}] \} \\
 &= \sum_i^T \log \left[\Gamma \left(\sum_k^K \alpha_k \right) \right] - \sum_i^T \sum_k^K \log[\Gamma(\alpha_k)] \\
 &+ \sum_i^T \sum_k^K (\alpha_k - 1) \left[\Psi(\phi_{ik}) - \Psi \left(\sum_k^K \phi_{ik} \right) \right].
 \end{aligned}$$

24 The expectation of the second term with respect to the variational distribution 24
 25 Q becomes 25

$$\begin{aligned}
 & \sum_i^T \sum_j^J \sum_n^{N_{ijr}} \mathbb{E}_Q \{ \log \text{mult}(z_{ijn} | \lambda_i) \} \\
 &= \sum_i^T \sum_j^J \sum_n^{N_{ij}} \mathbb{E}_Q \left\{ \log \left[\prod_k^K \lambda_{ik}^{z_{ijnk}} \right] \right\} \\
 (13) \quad &= \sum_i^T \sum_j^J \sum_n^{N_{ij}} \sum_k^K \mathbb{E}_Q \{ z_{ijnk} \log(\lambda_{ik}) \} \\
 &= \sum_i^T \sum_j^J \sum_n^{N_{ij}} \sum_k^K \mathbb{E}_Q \{ z_{ijnk} \} \mathbb{E}_Q \{ \log[\lambda_{i,k}] \} \\
 &= \sum_i^T \sum_j^J \sum_n^{N_{ij}} \sum_k^K \delta_{ijnk} \left[\Psi(\phi_{ik}) - \Psi \left(\sum_{k'}^K \phi_{ik'} \right) \right].
 \end{aligned}$$

The third term is

$$\begin{aligned}
& \sum_i^T \sum_j^J \mathbb{E}_Q \{ \log [f(X_{ij} | \theta, Z_{ij \cdot})] \} \\
&= \sum_i^T \sum_j^J \mathbb{E}_Q \left\{ \log \left[\prod_n^{N_{ij}} \left[\prod_k^K \left[\frac{\theta_{jka(n)ij}}{(1 - \sum_{c=0}^{n-1} \theta_{jka(c)ij})} \right]^{Z_{ijnk}} \right] \right] \right\} \\
(14) \quad &= \sum_i^T \sum_j^J \sum_n^{N_{ij}} \sum_k^K \mathbb{E}_Q \left\{ Z_{ijnk} \left(\log [\theta_{jka(n)ij}] \right. \right. \\
&\quad \left. \left. - \log \left[1 - \sum_{c=0}^n \theta_{jka(c)ij} \right] \right) \right\} \\
&= \sum_i^T \sum_j^J \sum_n^{N_{ij}} \sum_k^K \delta_{ijnk} \left(\log [\theta_{jka(n)ij}] - \log \left[1 - \sum_{c=0}^n \theta_{jka(c)ij} \right] \right).
\end{aligned}$$

Now for the second term of the ELBO

$$\begin{aligned}
& \mathbb{E}_Q \{ \log [Q(Z, \Lambda)] \} \\
&= \mathbb{E}_Q \left\{ \log \left[\prod_i^T \left(\text{Dir}(\lambda_i | \phi_i) \prod_j^J \prod_n^{N_{ij}} \text{mult}(Z_{ijn} | \delta_{ijn}) \right) \right] \right\} \\
&= \mathbb{E}_Q \left\{ \sum_i^T \log [\text{Dir}(\lambda_i | \phi_i)] + \sum_i^T \sum_j^J \sum_n^{N_{ij}} \log [\text{mult}(Z_{ijn} | \delta_{ijn})] \right\} \\
&= \mathbb{E}_Q \left\{ \sum_i^T \log \left[\frac{\Gamma(\sum_k^K \phi_{ik})}{\prod_k^K \Gamma(\phi_k)} \prod_k^K \lambda_{ik}^{\phi_{ik}-1} \right] + \sum_i^T \sum_j^J \sum_n^{N_{ij}} \log \left[\prod_k^K \delta_{ijnk}^{Z_{ijnk}} \right] \right\} \\
(15) \quad &= \mathbb{E}_Q \left\{ \sum_i^T \log \left[\Gamma \left(\sum_k^K \phi_{ik} \right) \right] - \sum_i^T \sum_k^K \log [\Gamma(\phi_{ik})] \right\} \\
&\quad + \mathbb{E}_Q \left\{ \sum_i^T \sum_k^K (\phi_{ik} - 1) \log [\lambda_{ik}] + \sum_i^T \sum_j^J \sum_n^{N_{ij}} \sum_k^K Z_{ijnk} \log [\delta_{ijnk}] \right\} \\
&= \sum_i^T \log \left[\Gamma \left(\sum_k^K \phi_{ik} \right) \right] - \sum_i^T \sum_k^K \log [\Gamma(\phi_{ik})]
\end{aligned}$$

$$\begin{aligned}
 & + \sum_i^T \sum_k^K (\phi_{ik} - 1) \left[\Psi(\phi_{ik}) - \Psi\left(\sum_k^K \phi_{ik}\right) \right] \\
 & + \sum_i^T \sum_j^J \sum_n^{N_{ij}} \sum_k^K \delta_{ijnk} \log[\delta_{ijnk}].
 \end{aligned}$$

SUPPLEMENTARY MATERIAL

Supplement to “A variational EM method for mixed membership models with multivariate rank data: An analysis of public policy preferences” (DOI: 10.1214/17-AOAS1034SUPP; .pdf). By analyzing 1997 Irish presidential election data, we provide a direct computational and goodness-of-fit comparison to the MCMC method of Gormley and Murphy (2009).

aoas1034supp.pdf

REFERENCES

- AIROLDI, E. M., BLEI, D. M., FIENBERG, S. E. and XING, E. P. (2008). Mixed membership stochastic blockmodels. *J. Mach. Learn. Res.* **9** 1981–2014.
- AIROLDI, E. M., BLEI, D. M., EROSheVA, E. A. and FIENBERG, S. E. (2015). Introduction to mixed membership models and methods. In *Handbook of Mixed Membership Models and Their Applications*. Chapman & Hall/CRC Handb. Mod. Stat. Methods 3–13. CRC Press, Boca Raton, FL. MR3380022
- BEAL, M. J. (2003). Variational algorithms for approximate Bayesian inference Ph.D. thesis, Univ. College, London, UK.
- BLEI, D. M. and LAFFERTY, J. D. (2005). Correlated topic models. In *Advances in Neural Information Processing Systems 18 [Neural Information Processing Systems, NIPS 2005, December 5–8, 2005, Vancouver, British Columbia, Canada]* 147–154.
- BLEI, D. M., NG, A. Y. and JORDAN, M. I. (2003). Latent Dirichlet allocation. *J. Mach. Learn. Res.* **3** 993–1022.
- BOTTOU, L. (2010). Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010* 177–186. Physica-Verlag/Springer, Heidelberg. MR3362066
- BROOKS, C. and MANZA, J. (2008). *Why Welfare States Persist: The Importance of Public Opinion in Democracies*. Univ. Chicago Press, Chicago.
- BURSTEIN, P. (1998). Bringing the public back in: Should sociologists consider the impact of public opinion on public policy? *Social Forces* **77** 27–62.
- BUSSE, L. M., ORBANZ, P. and BUHMANN, J. M. (2007). Cluster analysis of heterogeneous rank data. In *Machine Learning, Proceedings of the Twenty-Fourth International Conference (ICML 2007), Corvallis, Oregon, USA, June 20–24, 2007* 113–120.
- CARON, F., TEH, Y. W. and MURPHY, T. B. (2014). Bayesian nonparametric Plackett–Luce models for the analysis of preferences for college degree programmes. *Ann. Appl. Stat.* **8** 1145–1181. MR3262549
- CAVADINO, M. and DIGNAN, J. (2006). Penal policy and political economy. *Criminology and Criminal Justice* **6** 435–456.
- COHEN, A. and MALLOWS, C. (1983). Assessing goodness of fit of ranking models to data. *The Statistician* 361–374.
- EROSHEVA, E. A., FIENBERG, S. E. and JOUTARD, C. (2007). Describing disability through individual-level mixture models for multivariate binary data. *Ann. Appl. Stat.* **1** 502–537. MR2415745

- 1 EROSHEVA, E., FIENBERG, S. and LAFFERTY, J. (2004). Mixed-membership models of scientific
2 publications. *Proc. Natl. Acad. Sci. USA* **101** 5220–5227. 2 <author>
- 3 GILL, J. (2008). Is partial-dimension convergence a problem for inferences from MCMC algorithms?
4 *Polit. Anal.* **16** 153–178. 3 <author>
- 5 GORMLEY, I. C. and MURPHY, T. B. (2006). Analysis of Irish third-level college applications data.
6 *J. Roy. Statist. Soc. Ser. A* **169** 361–379. MR2225548 4 <author>
5 <mr>
- 7 GORMLEY, I. C. and MURPHY, T. B. (2008). A mixture of experts model for rank data with appli-
8 cations in election studies. *Ann. Appl. Stat.* **2** 1452–1477. MR2655667 6 <mr>
7 <mr>
- 9 GORMLEY, I. C. and MURPHY, T. B. (2009). A grade of membership model for rank data. *Bayesian*
10 *Anal.* **4** 265–295. MR2507364 8 <mr>
9 <mr>
- 11 GRASMICK, H. G., DAVENPORT, E., CHAMLIN, M. B. and BURSİK, R. J. (1992). Protestant fun-
12 damentalism and the retributive doctrine of punishment. *Criminology* **30** 21–46. 10 <author>
11 <author>
- 12 GROSS, J. H. and MANRIQUE-VALLIER, D. (2015). A mixed membership approach to the assess-
13 ment of political ideology from survey responses. In *Handbook of Mixed Membership Models*
14 *and Their Applications. Chapman & Hall/CRC Handb. Mod. Stat. Methods* 119–139. CRC Press,
15 Boca Raton, FL. MR3380027 12 <mr>
13 <mr>
- 16 GUIVER, J. and SNELSON, E. (2009). Bayesian inference for Plackett–Luce ranking models. In
17 *Proceedings of the 26th Annual International Conference on Machine Learning, ICML 2009,*
18 *Montreal, Quebec, Canada, June 14–18, 2009* 377–384. 14 <author>
15 <author>
- 19 HUNTER, D. R. (2004). MM algorithms for generalized Bradley–Terry models. *Ann. Statist.* **32**
20 384–406. MR2051012 16 <author>
17 <mr>
- 21 KITSCHHELT, H. and REHM, P. (2014). Occupations as a site of political preference formation. *Com-*
22 *parative Political Studies* **47** 1670–1706. 18 <mr>
19 <author>
- 23 LANGE, T., BRAUN, M. L., ROTH, V. and BUHMANN, J. M. (2002). Stability-based model selec-
24 tion. In *Advances in Neural Information Processing Systems 15 [Neural Information Processing*
25 *Systems, NIPS 2002, December 9–14, 2002, Vancouver, British Columbia, Canada]* 617–624. 20 <author>
21 <author>
- 26 LUCE, R. D. (1977). The choice axiom after twenty years. *J. Math. Psych.* **15** 215–233. MR0462675 22 <author>
23 <mr>
- 27 MARDEN, J. I. (1995). *Analyzing and Modeling Rank Data. Monographs on Statistics and Applied*
28 *Probability* **64**. Chapman & Hall, London. MR1346107 24 <mr>
25 <mr>
- 29 MAYHEW, P. and VAN KESTEREN, J. (2002). Cross-national attitudes to punishment. *Changing*
30 *Attitudes to Punishment* 63–92. 26 <author>
27 <author>
- 31 MEILA, M. and CHEN, H. (2010). Dirichlet process mixtures of generalized mallows models. In
32 *UAI 2010, Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence,*
33 *Catalina Island, CA, USA, July 8–11, 2010* 358–367. 28 <author>
29 <author>
- 34 NOCEDAL, J. and WRIGHT, S. J. (1999). *Numerical Optimization. Springer Series in Operations*
35 *Research.* Springer, New York. MR1713114 30 <mr>
31 <mr>
- 36 PLACKETT, R. L. (1975). The analysis of permutations. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **24** 193–
37 202. MR0391338 32 <mr>
33 <mr>
- 38 R CORE TEAM (2016). R: A language and environment for statistical computing. R Foundation for
39 Statistical Computing, Vienna, Austria. 33 <author>
34 <unstr>
- 40 REIF, K. and MELICH, A. (2001). Euro-barometer 34.1: Health Problems, Fall 1990. 34 <unstr>
35 <unstr>
- 41 ROBERTS, J. V. (2013). Public opinion and the nature of community penalties: Nternational findings.
42 *Changing Attitudes to Punishment* 33. 36 <author>
37 <author>
- 43 SEN, A. K. (2014). *Collective Choice and Social Welfare* **11**. Elsevier, Amsterdam. 37 <author>
38 <author>
- 44 TEH, Y. W., JORDAN, M. I., BEAL, M. J. and BLEI, D. M. (2006). Hierarchical Dirichlet processes.
45 *J. Amer. Statist. Assoc.* **101** 1566–1581. MR2279480 38 <mr>
39 <mr>
- 46 TONRY, M. (2007). Determinants of penal policies. *Crime and Justice* **36** 1–48. 39 <author>
40 <author>
- 47 WAINWRIGHT, M. J. and JORDAN, M. I. (2008). Graphical models, exponential families, and vari-
48 ational inference. *Found. Trends Mach. Learn.* **1** 1–305. 41 <author>
42 <author>
- 49 WANG, C. and BLEI, D. M. (2015). A general method for robust Bayesian modeling. *ArXiv preprint.*
50 *Available at arXiv:1510.05078.* 42 <author>
43 <author>

MIXED MEMBERSHIP MODELS FOR RANK DATA

29

1 WANG, Y. S. and EROSHEVA, E. A. (2015). mixedMem: Tools for discrete multivariate mixed mem- 1
2 bership models. R package version 1.1.2. Available at [https://cran.r-project.org/web/packages/](https://cran.r-project.org/web/packages/mixedMem/index.html) 2
3 mixedMem/index.html. 3 <author>
4 WANG, Y. S. and EROSHEVA, E. A. (2017). Supplement to “A variational EM method for mixed 4
5 membership models with multivariate rank data: An analysis of public policy preferences.” 5
6 DOI:10.1214/17-AOAS1034SUPP. 6 <unstr>
7 ZALLER, J. (1992). *The Nature and Origins of Mass Opinion*. Cambridge Univ. Press, Cambridge, 7 <author>
8 MA. 8

9 Y. S. WANG	R. L. MATSUEDA	9
10 E. A. EROSHEVA	UNIVERSITY OF WASHINGTON	10
11 UNIVERSITY OF WASHINGTON	DEPARTMENT OF SOCIOLOGY	11
12 DEPARTMENT OF STATISTICS	SEATTLE, WASHINGTON 98195	12
13 SEATTLE, WASHINGTON 98195	USA	13
14 USA	E-MAIL: matsueda@uw.edu	14
15 E-MAIL: ysamwang@uw.edu		15
16 erosheva@uw.edu		16
17		17
18		18
19		19
20		20
21		21
22		22
23		23
24		24
25		25
26		26
27		27
28		28
29		29
30		30
31		31
32		32
33		33
34		34
35		35
36		36
37		37
38		38
39		39
40		40
41		41
42		42
43		43

THE ORIGINAL REFERENCE LIST

The list of entries below corresponds to the original Reference section of your article. The bibliography section on previous page was retrieved from MathSciNet applying an automated procedure.

Please check both lists and indicate those entries which lead to mistaken sources in automatically generated Reference list.

- AIROLDI, E. M., BLEI, D. M., FIENBERG, S. E. and XING, E. P. (2008). Mixed membership stochastic blockmodels. *Journal of Machine Learning Research* **9** 1981–2014. 10.1145/1390681.1442798
- AIROLDI, E. M., BLEI, D. M., EROSHEVA, E. A. and FIENBERG, S. E. (2015). Introduction to mixed membership models and methods. In *Handbook of mixed membership models and their applications*. Chapman & Hall/CRC Handb. Mod. Stat. Methods 3–13. CRC Press, Boca Raton, FL. 3380022
- BEAL, M. J. (2003). Variational algorithms for approximate Bayesian inference PhD thesis, University College London, UK.
- BLEI, D. M. and LAFFERTY, J. D. (2005). Correlated topic models. In *Advances in Neural Information Processing Systems* 18 [*Neural Information Processing Systems, NIPS 2005, December 5-8, 2005, Vancouver, British Columbia, Canada*] 147–154.
- BLEI, D. M., NG, A. Y. and JORDAN, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research* **3** 993–1022.
- BOTTOU, L. (2010). Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT 2010* 177–186. Springer.
- BROOKS, C. and MANZA, J. (2008). *Why welfare states persist: the importance of public opinion in democracies*. University of Chicago Press.
- BURSTEIN, P. (1998). Bringing the public back in: should sociologists consider the impact of public opinion on public policy? *Social forces* **77** 27–62.
- BUSSE, L. M., ORBANZ, P. and BUHMANN, J. M. (2007). Cluster analysis of heterogeneous rank data. In *Machine Learning, Proceedings of the Twenty-Fourth International Conference (ICML 2007), Corvallis, Oregon, USA, June 20-24, 2007* 113–120. 10.1145/1273496.1273511
- CARON, F., TEH, Y. W. and MURPHY, T. B. (2014). Bayesian nonparametric Plackett-Luce models for the analysis of preferences for college degree programmes. *Ann. Appl. Stat.* **8** 1145–1181. 10.1214/14-AOAS717 3262549
- CAVADINO, M. and DIGNAN, J. (2006). Penal policy and political economy. *Criminology and Criminal Justice* **6** 435–456.
- COHEN, A. and MALLOWS, C. (1983). Assessing goodness of fit of ranking models to data. *The Statistician* 361–374.
- EROSHEVA, E. A., FIENBERG, S. E. and JOUTARD, C. (2007). Describing disability through individual-level mixture models for multivariate binary data. *Ann. Appl. Stat.* **1** 502–537. 10.1214/07-AOAS126 2415745
- EROSHEVA, E., FIENBERG, S. and LAFFERTY, J. (2004). Mixed-membership models of scientific publications. *Proceedings of the National Academy of Sciences of the United States of America* **101** 5220–5227.
- GILL, J. (2008). Is partial-dimension convergence a problem for inferences from MCMC algorithms? *Political Analysis* **16** 153–178.
- GORMLEY, I. C. and MURPHY, T. B. (2006). Analysis of Irish third-level college applications data. *J. Roy. Statist. Soc. Ser. A* **169** 361–379. 10.1111/j.1467-985X.2006.00412.x 2225548
- GORMLEY, I. C. and MURPHY, T. B. (2008). A mixture of experts model for rank data with applications in election studies. *Ann. Appl. Stat.* **2** 1452–1477. 10.1214/08-AOAS178 2655667

- 1 GORMLEY, I. C. and MURPHY, T. B. (2009). A grade of membership model for rank data. *Bayesian* 1
 2 *Anal.* **4** 265–295. 10.1214/09-BA410 2507364 2 <author>
- 3 GRASMICK, H. G., DAVENPORT, E., CHAMLIN, M. B. and BURSIK, R. J. (1992). Protestant fun- 3
 4 damentalism and the retributive doctrine of punishment. *Criminology* **30** 21–46. 4 <author>
- 5 GROSS, J. H. and MANRIQUE-VALLIER, D. (2015). A mixed membership approach to the assess- 5
 6 ment of political ideology from survey responses. In *Handbook of mixed membership models and* 6
 7 *their applications*. Chapman & Hall/CRC Handb. Mod. Stat. Methods 119–139. CRC Press, Boca 7
 8 Raton, FL. 3380027 8 <author>
- 9 GUIVER, J. and SNELSON, E. (2009). Bayesian inference for Plackett-Luce ranking models. In 9
 10 *Proceedings of the 26th Annual International Conference on Machine Learning, ICML 2009,* 10
 11 *Montreal, Quebec, Canada, June 14-18, 2009* 377–384. 10.1145/1553374.1553423 11 <author>
- 12 HUNTER, D. R. (2004). MM algorithms for generalized Bradley-Terry models. *Ann. Statist.* **32** 384– 12
 13 406. 10.1214/aos/1079120141 2051012 13 <author>
- 14 KITSCHOLT, H. and REHM, P. (2014). Occupations as a site of political preference formation. *Com-* 14
 15 *parative Political Studies* **47** 1670–1706. 15 <author>
- 16 LANGE, T., BRAUN, M. L., ROTH, V. and BUHMANN, J. M. (2002). Stability-based model selec- 16
 17 tion. In *Advances in Neural Information Processing Systems 15 [Neural Information Processing* 17
 18 *Systems, NIPS 2002, December 9-14, 2002, Vancouver, British Columbia, Canada]* 617–624. 18 <author>
- 19 LUCE, R. D. (1977). The choice axiom after twenty years. *J. Mathematical Psychology* **15** 215–233. 19
 20 0462675 20 <author>
- 21 MARDEN, J. I. (1995). *Analyzing and modeling rank data*. Monographs on Statistics and Applied 21
 22 *Probability* **64**. Chapman & Hall, London. 1346107 22 <author>
- 23 MAYHEW, P. and VAN KESTEREN, J. (2002). Cross-national attitudes to punishment. *Changing* 23
 24 *attitudes to punishment* 63–92. 24 <author>
- 25 MEILA, M. and CHEN, H. (2010). Dirichlet process mixtures of generalized mallows models. In 25
 26 *UAI 2010, Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence,* 26
 27 *Catalina Island, CA, USA, July 8-11, 2010* 358–367. 27 <author>
- 28 NOCEDAL, J. and WRIGHT, S. J. (1999). *Numerical optimization*. Springer Series in Operations 28
 29 *Research*. Springer-Verlag, New York. 10.1007/b98874 1713114 29 <author>
- 30 PLACKETT, R. L. (1975). The analysis of permutations. *J. Roy. Statist. Soc. Ser. C Appl. Statist.* **24** 30
 31 193–202. 0391338 31 <author>
- 32 R CORE TEAM (2016). R: a language and environment for statistical computing R Foundation for 32
 33 Statistical Computing, Vienna, Austria. 33 <author>
- 34 REIF, K. and MELICH, A. (2001). Euro-barometer 34.1: Health Problems, Fall 1990. 34
 35 10.3886/ICPSR09577.v1 35 <unstr>
- 36 ROBERTS, J. V. (2013). Public opinion and the nature of community penalties: nternational findings. 36
 37 *Changing Attitudes to Punishment* 33. 37 <author>
- 38 SEN, A. K. (2014). *Collective choice and social welfare* **11**. Elsevier. 38 <author>
- 39 TEH, Y. W., JORDAN, M. I., BEAL, M. J. and BLEI, D. M. (2006). Hierarchical Dirichlet processes. 39
 40 *J. Amer. Statist. Assoc.* **101** 1566–1581. 10.1198/016214506000000302 2279480 40 <author>
- 41 TONRY, M. (2007). Determinants of penal policies. *Crime and Justice* **36** 1–48. 41 <author>
- 42 WAINWRIGHT, M. J. and JORDAN, M. I. (2008). Graphical models, exponential families, and vari- 42
 43 ational inference. *Foundations and Trends in Machine Learning* **1** 1–305. 10.1561/2200000001 43 <author>
- 44 WANG, C. and BLEI, D. M. (2015). A general method for robust Bayesian modeling. *arXiv preprint* 44
 45 *arXiv:1510.05078*. 45 <author>
- 46 WANG, Y. S. and EROSHEVA, E. A. (2015). mixedMem: tools for discrete multivariate mixed mem- 46
 47 bership models R package version 1.1.2. 47 <author>
- 48 WANG, Y. S. and EROSHEVA, E. A. (2016). Supplement to “A Variational EM method for mixed 48
 49 membership models with multivariate rank Data: an analysis of public policy preferences”. 49 <author>
- 50 ZALLER, J. (1992). *The nature and origins of mass opinion*. Cambridge university press. 50 <author>

META DATA IN THE PDF FILE

Following information will be included as pdf file Document Properties:

Title : A variational EM method for mixed membership models with multivariate rank data: An analysis of public policy preferences
Author : Y. Samuel Wang, Ross L. Matsueda, Elena A. Erosheva
Subject : The Annals of Applied Statistics, 0, Vol. 0, No. 00, 1-29
Keywords: Mixed membership, rank data, variational inference, Euro-barometer, public policy

THE LIST OF URI ADDRESSES

Listed below are all uri addresses found in your paper. The non-active uri addresses, if any, are indicated as ERROR. Please check and update the list where necessary. The e-mail addresses are not checked – they are listed just for your information. More information can be found in the support page:
<http://www.e-publications.org/ims/support/urihelp.html>.

- 200 <http://www.imstat.org/aoas/> [2:pp.1,1] OK
- 200 <http://www.imstat.org> [2:pp.1,1] OK // <http://www.imstat.org/>
- 404 <http://dx.doi.org/10.1214/17-AOAS1034SUPP> [4:pp.26,26,27,27] Not Found
- 301 <http://arxiv.org/abs/arXiv:1510.05078> [2:pp.27,27] Moved Permanently
- <mailto:ysamwang@uw.edu> [2:pp.28,28] Check skip
- <mailto:erosheva@uw.edu> [2:pp.28,28] Check skip
- <mailto:matsueda@uw.edu> [2:pp.28,28] Check skip