Appendix 1. Estimating the Model

Given a sample of data for a dependent variable **Y**, preference variables **X**, threshold variables **Z**, and variance variables **W**, the log-likelihood function for the model of section two can be written:

$\ln L(\boldsymbol{b},\boldsymbol{g},\boldsymbol{h}) =$

$$\sum_{\text{vote a}} \ln\left(1 - \Phi\left(\frac{Z_i\boldsymbol{g} - X_i\boldsymbol{b}}{\exp(W_i\boldsymbol{h})}\right)\right) + \sum_{\text{abstain}} \ln\left(\Phi\left(\frac{Z_i\boldsymbol{g} - X_i\boldsymbol{b}}{\exp(W_i\boldsymbol{h})}\right) - \Phi\left(\frac{-Z_i\boldsymbol{g} - X_i\boldsymbol{b}}{\exp(W_i\boldsymbol{h})}\right)\right) + \sum_{\text{vote b}} \ln\left(\Phi\left(\frac{-Z_i\boldsymbol{g} - X_i\boldsymbol{b}}{\exp(W_i\boldsymbol{h})}\right)\right)$$

The parameters (**b, g, h**) can be estimated with the method of maximum likelihood. There are only two special considerations. First, if the set of variance variables **W** contains a constant then the model's parameters are not identified. This restriction is not unique to the present specification; it also applies to the heteroskedastic probit and ordered probit models (Greene 1997), and mirrors the normalization **s**=1 in an ordinary probit model (Long 1997). The second issue is that in the absence of any constraints, during the estimation procedure it is possible that the threshold values $Z_i\boldsymbol{g}$ will be negative for some observations. Negative thresholds induce P(abstain) less than zero and P(vote *a*)+P(vote *b*)>1, both of which are clearly unacceptable. To avoid these problems, during the estimation procedure I constrain the threshold values to be positive.

Inspection of this log-likelihood function reveals that the model's parameters are identified; that is, for any given set of parameters (**b, g, h**) there does not exist an alternative set of parameters (**b', g', h'**) such that L(**b, g, h**) = L(**b', g', h'**). However, because of the novelty of this model and the nonlinearity of the various components of the log-likelihood function it is worth pursuing the issue of identification as a practical matter. The question is: given a set of "true" parameters (**b, g, h**), can the method of maximum likelihood recover those parameters?

To answer this question I ran a set of simulations. I first established "true" parameter values for the utility differential, threshold, and variance components for the model of the previous section (these are presented in Appendix Table 1). Next, I generated a set of 1000 observations on four variables, { X_1, X_2, W_1, and Z_1 }, using draws from a multivariate normal distribution with mean $\{0.0 \quad 0.0 \quad 0.0 \quad 0.0\}$ and variance-covariance matrix

$$\begin{bmatrix} 1.0 & 0.2 & 0.4 & 0.6 \\ 0.2 & 1.0 & 0.2 & 0.4 \\ 0.4 & 0.2 & 1.0 & 0.4 \\ 0.6 & 0.4 & 0.4 & 1.0 \end{bmatrix}. \text{ Then, for each simulation I drew 1000 values of } \boldsymbol{e} \text{ from a normal}$$

distribution with mean zero and variance $\boldsymbol{s}_i^2$ (defined below), and calculated the various components of the model, as follows:

utility differential: $X_i\boldsymbol{b} = \boldsymbol{b}_0 + \boldsymbol{b}_1X_1 + \boldsymbol{b}_2X_2 + \boldsymbol{e}_i$

utility differential variance $\boldsymbol{s}_i^2 = \left(\exp(\boldsymbol{h}_1 X_1 + \boldsymbol{h}_2 W_1)\right)^2$

thresholds $Z_i\boldsymbol{g} = \boldsymbol{g}_0 + \boldsymbol{g}_1X_1 + \boldsymbol{g}_2Z_1$

For each simulation I generated the dependent variable in accordance with the behavioral specification described in the previous section:

vote a if $X_i\boldsymbol{b} + \boldsymbol{e}_i > Z_i\boldsymbol{g}$

vote b if $X_i\boldsymbol{b} + \boldsymbol{e}_i < -Z_i\boldsymbol{g}$

abstain if $Z_i\boldsymbol{g} \leq X_i\boldsymbol{b} + \boldsymbol{e}_i \leq Z_i\boldsymbol{g}$

I repeated the simulation procedure 1000 times, each time randomly generating another set of 1000 values for $\boldsymbol{e}$. For each of the 1000 simulations I estimated the parameters ($\boldsymbol{b}, \boldsymbol{g}, \boldsymbol{h}$) via maximum likelihood estimation, using starting values equal to a vector of zeros (except for $\boldsymbol{g}_0$, which was set to one to ensure that the initial threshold values are positive).

This simulation provides a demanding test of identification in a situation that is analogous to my application of this model to electoral choice in the 1996 American presidential election. Because some of the same variables may be used in different components of the model, I construct the simulation so that the sets of utility differential variables (X), variance variables (W), and threshold variables (Z) each contain one element in common (X_1) and one unique element (X_2, W_1, and Z_1 respectively). Further, these variables are correlated to various degrees. I chose a sample size of 1000 to test the performance of the model at a sample size that is typical for studies of mass political behavior. I also chose the parameter values so that the observed distribution of outcomes matched the distribution of outcomes in my sample, where the turnout rate among registered voters was 85.5%, and the division of the two-party vote was 43.5% for candidate *a* (Dole) and 56.5% for candidate *b* (Clinton).

[Appendix Table 1 about here]

The results from these simulations are shown in Appendix Table 1, which demonstrates that the parameter estimates, on average, match the true parameter values quite closely. In addition, the estimated standard errors of those parameter estimates, on average, are very good approximations of the actual distribution of the maximum likelihood estimates. Because the estimated standard errors are asymptotic where my sample is finite, there are some discrepancies between the actual and estimated standard errors; here the estimated standard errors are larger than the actual standard errors, though the excess is never more than about 6%. Because this difference is relatively small, and because in this particular application it tends to increase the conservatism of hypothesis tests (because estimated standard errors are a little larger than they should be, so that *z*-scores are a little smaller than they should be), I use conventionally calculated z-scores in my empirical tests.

Appendix Table 1. Simulation Results

	Parameter Value	Average Maximum Likelihood Estimate	Standard Deviation of Estimated Parameters	Average Estimated Standard Error
Preference Parameters				
b_0	0.5	0.502	0.061	0.062
b_1	2.0	2.019	0.108	0.112
b_2	-1.0	-1.008	0.064	0.068
Threshold Parameters				
g_0	0.5	0.497	0.038	0.040
g_1	1.0	1.028	0.095	0.098
g_2	-0.5	-0.519	0.054	0.057
Variance Parameters				
h_1	0.4	0.399	0.077	0.080
h_2	0.2	0.200	0.065	0.065

Note: Table entries are based on 1000 simulations as described in Appendix 1.