

Distinguishing Occasional Abstention from Routine Indifference in Models of Vote Choice*

Benjamin E. Bagozzi[†] and Kathleen Marchetti[‡]

Abstract

Researchers commonly employ multinomial logit (MNL) models to explain individual-level vote choice while treating “abstention” as the baseline category. Though many view abstainers as a homogeneous group, we argue that these respondents emerge from two distinct sources. Some nonvoters are likely to be “occasional voters” who abstained from a given election due to temporary factors, such as a distaste for all candidates running in a particular election, poor weather conditions, or other temporary circumstances. On the other hand, many nonvoters are unlikely to vote regardless of the current political climate. This latter population of “routine nonvoters” is consistently disengaged from the political process in a way that is distinct from that of occasional voters. Including both sets of nonvoters within an MNL model can lead to faulty inferences. As a solution, we propose a baseline-inflated MNL estimator that models heterogeneous populations of nonvoters probabilistically, thus accounting for the presence of routine nonvoters within models of vote choice. We demonstrate the utility of this model using replications of existing political behavior research.

*Earlier drafts of this paper presented at the 7th Annual St. Louis Area Methods Meeting (SLAMM!) and the APSA 2013 Annual Meeting. The authors wish to acknowledge the valuable suggestions and feedback that they received from the editors, reviewers, and replication analysts at Political Science Research and Methods, as well as from Justin Esarey, John Freeman, Jeff Gill, Galin Jones, Luke Keele, Ines Levin, James McCann, Jamie Monogan, Burt Monroe, Jacob Montgomery, Bumba Mukherjee, and Chris Zorn. The authors would also like to thank Kevin Arceneaux, Robin Kolodny, David Campbell, and J. Quin Monson for sharing their replication data.

[†]Department of Political Science and International Relations, University of Delaware. Email: bagozzib@gmail.com

[‡]Department of Political Science, Dickinson College. Email: marchetk@dickinson.edu

Introduction

Outcomes of discrete, polytomous choice are central to the study of politics. For instance, scholars are frequently interested in whether individuals favor (or vote for) candidate A, candidate B, or neither candidate (i.e., abstention) in a given election. In a similar vein, many social scientists empirically examine whether citizens or governments prefer (or enact) policy A, policy B, or maintain the status quo within policy areas ranging from health care (Propper, 2000) to environmental conservation (Lehtonen et al., 2003) to minority language recognition (Liu, 2011). Each of these examples envisions a dependent variable with a small set of discrete, unordered outcomes or choices. Accordingly, quantitative studies of such questions have favored the use of polytomous choice maximum likelihood models—most frequently the multinomial logit (MNL) model¹—to estimate the effects of covariates on a respondent’s probabilities of choosing each choice option over the others.

A second commonality shared by all of these dependent variables, however, is the presence of a “status quo,” “neither,” or “abstention” category representing instances wherein a voter, a government, or another political actor favored doing nothing, or abstaining, rather than choosing either option A or option B. While the inclusion of these abstention responses is often necessary to ensure an unbiased sample for one’s dependent variable, they are usually of less interest to the researcher than are the other “active” political choices, and accordingly, the vast majority of polytomous choice models in political science treat these abstentions as the baseline (i.e., reference) category in estimation and interpretation. Although this framework may help to avoid selection on one’s outcomes of interest, observational sampling schemes of this sort risk polluting the baseline choice-category with an excess number of unrealistic observations that correspond to “inflated” individuals that will virtually *never* select a choice outcome other than “abstain”. With a baseline category inflated in this manner, MNL estimates are likely to be biased downward, as a significant portion of abstainers will be effectively impervious to the

¹As discussed below, we surveyed all recent (2009-2013) articles appearing in the *American Political Science Review*, *American Journal of Political Science* and *Journal of Politics* and found 42 studies that used a multinomial estimator in either a primary or robustness test. Of these articles, 93% favored the MNL model over other multinomial choice estimators.

effects of covariates on transition probabilities from abstention to choice. Furthermore, when a subset of one's covariates have dual effects on both the probability of routine abstention and the probability of multinomial choice, then unit homogeneity assumptions² can be violated and MNL coefficients can become biased in indeterminate directions.

More succinctly, in instances where only a subset of a population actively encounters the choice scenario of interest, naïve samples of the entire population will contain an excess number of inactive choice-maker responses. Ignoring the heterogeneity that accordingly arises within one's "abstain" response set can yield faulty inferences. To illustrate these concerns, consider the example of individual surveys of vote choice. Most major political surveys of American citizens (e.g., the American National Election Studies) include questions designed to measure citizens' voting behaviors and opinions. Often, such questions ask respondents to indicate which candidate they chose (or plan to choose) in a given election from a list of options (e.g., Barack Obama, Mitt Romney, or did not vote/neither). Researchers then typically analyze these responses using multinomial logit (MNL) models of vote choice (e.g., Arceneaux and Kolodny, 2009; Kalmoe and Piston, 2013). "Vote-abstention" responses typically arise from two distinct sources. Some nonvoters are best seen as "routine nonvoters," in that they have an abstention history that—due to slow moving or structural factors—is unlikely to change as a function of the unique characteristics of a given election (e.g., candidate personality or get out the vote [GOTV] efforts). A second subset of nonvoters are often instead characterized as "occasional voters," who have voted in recent elections but who may not make the point to vote in each and every possible election (Gerber and Green, 2000; Niven, 2004; Hillygus, 2005; Parry et al., 2008). Researchers have frequently sought to distinguish between these two different groups of nonvoters when assessing the determinants of vote-abstention, often under the contention that the salience of each determinant is likely to vary as a function of nonvoter type (Zipp, 1985; Lacy and Burden, 1999; Sanders, 1999; Plane and Gershtenson, 2004).

Hence, treating nonvoters as a single homogeneous reference category within an MNL anal-

²That is, the assumption that one's dependent variable will take on the same expected values given a particular value on an independent variable (Braumoeller, 2013; King et al., 1994, 91-93).

ysis ignores the heterogeneous effects of temporary and structural covariates on individual vote choice (relative to abstention). As such, the estimated effects of short term shocks and candidate characteristics on turnout—which likely have no effect on routine nonvoters—will be biased downwards relative to their actual effect(s) within the subsample of voters that researchers are often most interested in: occasional and routine voters. By contrast, the direct effect(s) of structural factors on vote choice or turnout may be overstated—or misattributed to active voting behavior—when routine nonvoters and occasional voters are pooled, leading to faulty conclusions and misleading GOTV-type policy prescriptions.

These dynamics have not gone unnoticed by past political science research (Beger et al., 2011; Braumoeller, 2013). Yet, MNL models continue to be the default choice for political scientists seeking to analyze survey respondents’ polytomous vote choices, as well as for researchers studying the related political-economic outcomes mentioned earlier. As the above examples demonstrate, these practices may be especially problematic when one is interested in parsing out the effects of individual and societal factors on the determinants of active political choices. To address these concerns, this paper presents a novel empirical model that accounts for baseline inflation in polytomous choice outcomes—and the heterogeneity that it causes—*probabilistically* to ensure that one’s primary choice estimates are unbiased. Drawing on a number of recently developed zero inflated estimators, the model that we propose does so by combining a logit stage for the estimation of an observation’s probability of being inflated (or not) with a second, MNL outcome stage. By estimating these two stages as a single system of equations, we are able to account for the possibility of baseline category inflation through the use of theoretically-informed covariates, and to then derive unbiased estimates of the direct effects of one’s covariates on the polytomous choice set of interest—now conditioned on the probability of each estimate being a “non-inflated” observation.

Our proposed model, the Baseline Inflated Multinomial Logit (BIMNL), addresses instances where one’s theory suggests that a *single* choice-category in a polytomous dependent variable is inflated along the lines mentioned above. As such, the BIMNL model is one that

brings a commonly used discrete choice model (the MNL) closer to our substantive knowledge of voter behavior, and political phenomena more generally. For ease of exposition, we present this new model for the case where one’s inflated (i.e., “abstention” or “status quo”) category is treated as the (omitted) baseline category in the multinomial analysis as this is the most common practice in the literature. However, it is important to note that the baseline category in multinomial models is arbitrarily determined, typically for convenience of interpretation, and our applications of the BIMNL model could be easily modified to account for single category inflation within a non-baseline response category instead.

This study proceeds as follows. In the next section we formally derive the BIMNL model and briefly mention several test statistics that can be used to evaluate the “fit” of this model relative to the MNL model. This is followed by a discussion of the BIMNL model in relation to the multinomial-choice and zero-inflation methodological literatures. We then present the results of two replication exercises in which the BIMNL model is applied to existing political science studies of vote choice. We conclude by suggesting that the BIMNL model can be extended in a number of useful directions, including additional applications to survey responses of candidate preference, as well as the creation of a baseline inflated multinomial probit model with and without correlated errors.

The BIMNL Model

The BIMNL estimator combines two latent equations: a logit equation which we denote as the “first stage” estimator and a multinomial logit (MNL) equation hereafter referred to as the “outcome stage” estimator. To motivate this model, consider a dependent variable Y_i with $i \in \{1, 2, \dots, N\}$ respondents (e.g., individuals). Suppose further that Y_i is observable and assumes the discrete *unordered* values of $0, 1, \dots, J$, with value $Y_i = 0$ representing a baseline “abstain” category. Next, let s_i denote a binary variable that indicates a *split* between regime 0 ($s_i = 0$) and regime 1 ($s_i = 1$), wherein $s_i = 0$ denotes “inflated abstainers” that will always fall within the residual baseline category of Y_i and $s_i = 1$ denotes “non-inflated abstainers”

who may choose to abstain in a given instance, but who may also potentially select an option in $1 \dots J$. In the context of a vote choice survey data set, the abstain responses in regime 0 ($s_i = 0$) would include routine nonvoters who never choose to vote, while responses in regime 1 ($s_i = 1$) include occasional nonvoters whose probability of transitioning to a candidate vote choice outcome is *not* zero. Note that s_i is related to the latent dependent variable s_i^* such that $s_i = 1$ for $s_i^* > 0$ and $s_i = 0$ for $s_i^* \leq 0$. The latent variable s_i^* represents the propensity for entering regime 1 and is given by the following linear additive specification, which we refer to as the latent inflation equation:

$$s_i^* = \mathbf{z}_i' \boldsymbol{\gamma} + u_i \quad (1)$$

The inflation equation in 1, once re-stated via a binary (logistic) response model for our dichotomous regime indicator s_i , constitutes the first stage of the BIMNL model. In equation 1, \mathbf{z}_i is the vector of covariates, $\boldsymbol{\gamma}$ is the vector of coefficients and u_i is a standard-logistic distributed error term. Hence the probability of i being in regime 1 is $\Pr(s_i = 1 | \mathbf{z}_i) = \Pr(s_i^* > 0 | \mathbf{z}_i) = \Lambda(\mathbf{z}_i' \boldsymbol{\gamma})$, and the probability that i is in regime 0 is $\Pr(s_i = 0 | \mathbf{z}_i) = \Pr(s_i^* \leq 0 | \mathbf{z}_i) = 1 - \Lambda(\mathbf{z}_i' \boldsymbol{\gamma})$ where $\Lambda(\cdot)$ is the logistic cumulative distribution function (c.d.f.).

If $s_i = 1$, then the observations in regime 1 are given by the discrete unordered variable \tilde{Y}_i which can take on any of J *unordered* values, and where $\Pr(\tilde{Y}_i = j) = P_{ij}$. In noting that by definition, $\sum_{j=0}^J P_{ij} = 1$, we then allow the probability of $\tilde{Y}_i = j \in J$ to vary as a function of some k independent variable(s) \mathbf{x}_i , indexed by a $K \times 1$ vector of parameters specific to outcome β_j by restricting the probabilities to be positive and sum to one as so:

$$\Pr(\tilde{Y}_i = j) \equiv P_{ij} = \frac{e^{\mathbf{x}_i' \beta_j}}{\sum_{j=0}^J e^{\mathbf{x}_i' \beta_j}} \quad (2)$$

Hence, observation i 's probability associated with category j is expressed as a fraction of the sum of *all* of observation i 's probabilities across the various categories J . This ensures that $\Pr(Y_i = j) \in (0, 1)$ and that $\sum_{j=0}^J \Pr(\tilde{Y}_i = j) = 1$. Under the assumption that the corresponding error term (ε_i) for equation 2 is independently and identically distributed according to a Type

I Extreme Value distribution, equation 2 denotes the primary statement of probability for the MNL model. However, this model is also *unidentified*: knowing the $(J - 1) \times k$ values of $\beta_0, \beta_1, \dots, \beta_{J-1}$, also provides one with the probability of choosing the remaining alternative. To identify the above MNL probabilities for estimation, we follow common practice and set the parameters for the first of the J alternatives to zero, i.e. β_0 , which we refer to hereafter as the *baseline* category. Doing so allows us to restate the probabilities for the baseline category ($Pr(\tilde{Y}_i = 0)$) and the other $J - 1$ categories separately as:

$$\Pr(\tilde{Y}_i) = \left\{ \begin{array}{l} \Pr(\tilde{Y}_i = 0 | \mathbf{x}_i, s_i = 1) = \frac{1}{1 + \sum_{j=1}^J e^{\mathbf{x}_i' \beta_j}} \\ \Pr(\tilde{Y}_i = j | \mathbf{x}_i, s_i = 1) = \frac{e^{\mathbf{x}_i' \beta_j}}{1 + \sum_{j=1}^J e^{\mathbf{x}_i' \beta_j}} \quad (j = 1, \dots, J) \end{array} \right\} \quad (3)$$

where $\beta_j = \beta_j - \beta_0$ are now “rescaled” parameters in that they express the influence of the various \mathbf{x} ’s on $Pr(\tilde{Y}_i = j)$ relative to $Pr(\tilde{Y}_i = 0)$.

Note that neither \tilde{Y}_i nor s_i are observable in terms of the observed baseline outcomes. However, they are observed by the criterion $Y_i = \tilde{Y}_i \times s_i$. The aforementioned expression thus implies that the (baseline) outcome $Y_i = 0$ can occur when $s_i = 0$ or when $s_i = 1$ and $\tilde{Y}_i = 0$. It also indicates that we can observe $Y_i = 1 \dots J$ only when $s_i = 1$ and $\tilde{Y}_i = 1 \dots J$. Accordingly, the baseline inflated MNL distribution arises as a mixture of a degenerate distribution in the baseline category and the assumed distribution of the variable \tilde{Y}_i as follows:

$$\Pr(Y_i) = \left\{ \begin{array}{l} \Pr(s_i = 0 | \mathbf{z}_i) + \Pr(s_i = 1 | \mathbf{z}_i) \Pr(\tilde{Y}_i = 0 | \mathbf{x}_i, s_i = 1) \text{ for } j = 0 \\ \Pr(s_i = 1 | \mathbf{z}_i) \Pr(\tilde{Y}_i = j | \mathbf{x}_i, s_i = 1) \text{ for } j = 1, 2 \dots J \end{array} \right\} \quad (4)$$

Under the assumption that u_i and ε_i identically and independently follow standard Type I Extreme Value distributions, the BIMNL model can thus be defined as:

$$\Pr(Y_i) = \left\{ \begin{array}{l} \Pr(Y_i = 0 | \mathbf{x}_i, \mathbf{z}_i) = [1 - \Lambda(\mathbf{z}_i' \gamma)] + \left(\frac{\Lambda(\mathbf{z}_i' \gamma)}{1 + \sum_{j=1}^J e^{\mathbf{x}_i' \beta_j}} \right) \\ \Pr(Y_i = j | \mathbf{x}_i, \mathbf{z}_i) = \left(\frac{\Lambda(\mathbf{z}_i' \gamma) e^{\mathbf{x}_i' \beta_j}}{1 + \sum_{j=1}^J e^{\mathbf{x}_i' \beta_j}} \right) \quad (j = 1, \dots, J) \end{array} \right\} \quad (5)$$

where, as above, $\Lambda(\cdot)$ is the logistic c.d.f. The expression in 5 thus provides the full probabilities of the BIMNL Model. Herein, the probability of observing a baseline-choice observation within the baseline equation of the BIMNL model is modeled conditional upon the probability of an observation being assigned a baseline value in the multinomial d.g.p. plus the probability of it being in regime 0 from the inflation equation. As a result, when the unordered dependent variable is baseline-inflated and thus has two types of baseline observations, the BIMNL model allows researchers to obtain more accurate estimates relative to a standard MNL model, which is to say that the BIMNL estimates are both less biased and have greater coverage probabilities. We demonstrate these points in the Monte Carlo simulations reported in our supplemental appendix. The remaining probabilities in 5 then correspond to the conditional probabilities of an observation choosing a polytomous choice value other than the baseline value, conditional on an observation being in the MNL state of the world.

Having described the conditional probabilities for the BIMNL model above, we can now define the likelihood and the log-likelihood function of the BIMNL model. Specifically, let $\theta = (\gamma', \beta', \mathbf{u}', \varepsilon')'$ for the full BIMNL model. The likelihood of the BIMNL model for an i.i.d sample of $i \in \{1, 2, \dots, N\}$ observations can thus be defined as

$$\begin{aligned} \mathcal{L}(\theta) &= \prod_{i=1}^N \prod_{j=0}^J [\Pr(y_i = j | \mathbf{x}_i, \mathbf{z}_i, \theta)]^{d_{ij}} \\ &= \prod_{i=1}^N \prod_{j=0}^0 [\Pr(s_i = 0) + \Pr(s_i = 1) \Pr(\tilde{y}_i = j)]^{d_{ij}} \times \prod_{i=1}^N \prod_{j=1}^J [\Pr(s_i = 1) \Pr(\tilde{y}_i = j)]^{d_{ij}} \end{aligned} \quad (6)$$

where $(y_i = j | \mathbf{x}_i, \mathbf{z}_i)$ was described earlier and where $d_{ij} = 1$ if outcome j is realized in i and is $d_{ij} = 0$ otherwise. The log likelihood function of the full BIMNL model where $\theta = (\gamma', \beta', \mathbf{u}', \varepsilon')'$ can therefore be written succinctly as

$$\begin{aligned} \ell(\theta) &= \sum_{i=1}^N \sum_{j=0}^J d_{ij} \ln[\Pr(y_i = j | \mathbf{x}_i, \mathbf{z}_i, \theta)] \\ &= \sum_{i=1}^N \sum_{j=0}^0 d_{ij} \ln \left[1 - \Lambda(\mathbf{z}_i' \gamma) + \left(\frac{\Lambda(\mathbf{z}_i' \gamma)}{1 + \sum_{j=1}^J e^{\mathbf{x}_i' \beta_j}} \right) \right] + \sum_{i=1}^N \sum_{j=1}^J d_{ij} \ln \left(\frac{\Lambda(\mathbf{z}_i' \gamma) e^{\mathbf{x}_i' \beta_j}}{1 + \sum_{j=1}^J e^{\mathbf{x}_i' \beta_j}} \right) \end{aligned} \quad (7)$$

The log-likelihood function in 7 can be consistently and efficiently estimated using maximum likelihood which yields asymptotically normally distributed maximum likelihood estimates. We provide a preliminary R package that permits users to estimate the BIMNL model in full, and also supply the necessary R code for fully replicating and deriving the predicted probabilities that are presented in our applications below. As the BIMNL’s estimation structure is equivalent to that of extant zero inflated models, parameter identification for this estimator is technically achievable even in cases of perfect overlap among the independent variables used in each stage of the BIMNL model (i.e., without an exclusion restriction). Indeed, it has been demonstrated in both theory and practice that zero inflated count and (zero) inflated ordered estimators (*without correlated errors*) can achieve parameter identification without exclusion restrictions—although such restrictions can help guard against misspecification and computational problems in these contexts (Harris and Zhao, 2007; Bagozzi and Mukherjee, 2012; Staub and Winkelmann, 2012; Burger et al., 2009, 176). Even so, challenges to maximum likelihood estimation may arise in these contexts when outliers are present, in cases of high multicollinearity, or when parameter effects are poorly separated (Harris and Zhao, 2007; Hall and Shen, 2010). Thus, when available, the use of exclusion restrictions in the BIMNL’s inflation or outcome stage specification will likely improve estimation precision, as would alternate estimation approaches such as robust expectation solution estimation (Hall and Shen, 2010).

Though theory should guide one’s decision of when to use a BIMNL model, several model fit statistics may enable researchers to accurately test between the MNL and BIMNL models. In line with conceptualizations of extant inflated models (Greene, 2011; Harris and Zhao, 2007, 1079), the MNL model is not directly nested in the BIMNL model via parameter restrictions, though akin to the ZiOP/OP model case, the BIMNL model converges to an MNL model as $\mathbf{z}'\gamma \rightarrow \infty$ in equation 5 (i.e., as the probability of non-inflation goes to one). Hence, in choosing between BIMNL and MNL models, non-nested model test statistics are preferable to nested hypotheses tests. One frequently used test statistic for such model comparisons³ is the Vuong test for non-nested models which, in the case of the (BI)MNL model, assigns m_i as the natural

³See, e.g., Greene (2011); Harris and Zhao (2007); Bagozzi et al. (2014); Bagozzi and Mukherjee (2012).

logarithm of the ratio of the predicted probability that $Y_i = j$ for one’s MNL model (in the numerator) and one’s BIMNL model (in the denominator) and evaluates m_i via a bidirectional test statistic of

$$v = \frac{\sqrt{N}(\frac{1}{N} \sum_i^N m_i)}{\sqrt{\frac{1}{N} \sum_i^N (m_i - \bar{m})^2}} \quad (8)$$

where $v < -1.96$ favors the more general (BIMNL) model, $-1.96 < v < 1.96$ lends no support to either model, and $v > 1.96$ favors the MNL model (Vuong, 1989).

Generalized likelihood ratio (*LR*) tests,⁴ Akaike information criterion (*AIC*), and Bayesian information criterion (*BIC*) may each similarly serve as appropriate model selection criteria for our particular non-standard model comparisons given previous research on non-nested (inflated) models (Harris and Zhao, 2007, 1079).⁵ In a similar fashion to the use of the *AIC* and *BIC* in comparisons of logit and probit models, the *AIC* and *BIC* may be particularly appropriate for comparing the BIMNL and MNL models as these two models each use the categorical distribution family, albeit with different link functions for the choice probabilities.⁶ Lastly, one could also plausibly use in-sample predicted probabilities (for outcomes $Y_i = 0 \dots J$) to compare how well the MNL and BIMNL models replicate (and hence “fit”) the observed distribution of outcomes on Y_i via a proportion reduction in error (*PRE*) statistic. A *PRE* comparison for our models would first employ the probability statements appearing in equations 3 and 5 to calculate the in-sample predicted probabilities of each $Y_i = 0 \dots J$ outcome and for each observation as a function of the values on \mathbf{x}_i and \mathbf{z}_i . Classifying each observation based on the highest $Y_i = 0 \dots J$ predicted probability it receives, one can then quantify the proportion of “errors” that the MNL and BIMNL correctly classify, relative to a null model that always predicts the most frequent outcome on Y_i .

We conduct extensive Monte Carlo experiments—presented in the supplemental appendix—to evaluate the performance of each of these model fit statistics for the BIMNL and MNL mod-

⁴With degrees of freedom corresponding to the additional parameters estimated within the BIMNL model.

⁵Although the *LR* test is not strictly appropriate in this instance, given the non-nesting of inflated and comparable non-inflated models (Greene, 2011; Harris and Zhao, 2007, 1079).

⁶In addition to the fact that the *AIC* and *BIC* were designed in large part for the specific task of in-sample model comparison—as mentioned below.

els. Our experiments indicate that standard information-based model selection criteria, such as *BIC* and *AIC*, correctly chose between the BIMNL and MNL models (under each d.g.p.) nearly 100 percent of the time. These findings are unsurprising⁷ given that in-sample model comparisons are precisely what the *BIC* and *AIC* are designed to do, and thereby suggest that the *BIC* and *AIC* may be the best choices for applied BIMNL-MNL model comparisons among the model fit statistics discussed here. *LR* tests perform comparably, correctly choosing the BIMNL and MNL models in 95-100 percent of our simulations. Akin to our findings for the *AIC* and *BIC*, this result is consistent with the OP/ZIOP simulation results reported in Harris and Zhao (2007). Also in line with these extant findings, the generalized Vuong test statistic described above accurately chooses the BIMNL model when the d.g.p. was BIMNL, but performed poorly in selecting the MNL model when the d.g.p. was MNL. The PRE fares even worse in our simulations, exhibiting a preference for the MNL model no matter the d.g.p. and often favoring neither model over the other. This poor performance is likely attributable to (i) the PRE being a poor choice for model comparison when one's observations disproportionately fall within a single outcome category (as will typically be the case for inflated dependent variables) and (ii) the more general challenges associated with classification statistics such as the PRE in what are often irreducibly stochastic political science applications. Hence, researchers interested in using model fit statistics to supplement their (BIMNL versus MNL) model selection decisions should employ a combination of the test statistics mentioned above, and should place more weight in the *BIC*, *LR* test, and *AIC*, relative to the Vuong test and the PRE.

In its design, the BIMNL model shares a number of notable similarities with extant limited dependent variable models. As alluded to above, the BIMNL model can be most directly situated within the family of limited dependent variable finite mixture models known as “split population” or inflated models. Like the BIMNL model, these models are explicitly designed to account for inflation within a *single* category of one's dependent variable. Well known existing inflated models include zero inflated Poisson (ZIP) and negative binomial count models (ZINB; Greene, 1994), zero and middle inflated ordered probit (ZIOP and MIOP) models

⁷And are consistent with past simulation findings in this vein (Harris and Zhao, 2007).

for discrete ordered outcomes (Harris and Zhao, 2007; Bagozzi and Mukherjee, 2012; Brooks et al., 2012), and split population models for binary choice sets, such as the split population logit (Beger et al., 2011). Each of these models includes a system of two equations that allow one to estimate the effects of two potentially overlapping sets of covariates on (i) the probability of an observation arising from an inflation process specific to a single (usually $Y_i = 0$) outcome and (ii) the probability of one's discrete outcomes of interest conditional on an observation *not* being inflated. What makes the BIMNL model novel in this case is its ability to estimate these inflation processes within an unordered polytomous dependent variable, which to our knowledge is the first such estimator to explicitly do so.

These inflation processes are distinct from the choice dynamics underlying other multi-equation discrete choice models, such as selection models, as baseline inflation does not presume that one's first stage process truncates all cases for that choice outcome from one's set of outcome stage choice categories. Rather, an inflation process *augments* unwanted cases to the observed set of outcome choices. Hence, the BIMNL model conceives of vote choice differently from earlier multi-equation models of vote choice (e.g., Born, 1990; Sanders, 1999) in that it maintains "abstention" as an available multinomial outcome stage category—albeit one in which the response set has been deflated by the inflation equation estimates. This framework thereby allows researchers to correctly model vote choice in situations where candidate choice cannot be assumed to be conditional upon *all* individuals making a separate, earlier decision of whether or not to turnout for an election at all. Examples not only include measures of vote choice where distinct populations of nonvoters are believed to be present within one's abstention category, but also encompass survey questions that ask individuals to choose from a single response set that is comprised of both candidate choices and an abstention option, or instances where individuals are asked to provide their candidate preference—rather than vote choice—from a response set that includes an option of "neither" or "none of the above".

Lastly, whereas our focus here is on developing a baseline inflated MNL model, one could conceivably extend our approach to the multinomial probit (MNP) context. Doing so would

allow for a full relaxation of the MNL model's IIA assumption while also facilitating the inclusion of correlated disturbances between one's inflation and outcome equations. While this extension is intriguing, we currently choose to focus on the (BI)MNL model for the following three reasons. First, the MNL's IIA shortcomings aside, it remains the workhorse multinomial choice estimator within applied political science research. Surveying the past five years (2009-2013) of the *American Political Science Review*, *American Journal of Political Science*, and *Journal of Politics*, we found 42 articles that employed a multinomial estimator in a primary or robustness test (listed in the supplemental appendix). Of these 42 articles, only two used an MNP model as opposed to an MNL model,⁸ and in both cases the authors mentioned doing so only in a footnote as a robustness check to their primary MNL model of interest. Second, the costs of using an MNP model often outweigh its benefits. As Dow and Endersby (2004) demonstrate in this regard, the IIA assumption is not overly restrictive for most vote choice applications, while the MNP model is often fraught with estimation challenges in these contexts. Finally, although an allowance for correlated disturbances between our estimating stages is a promising extension, we do not consider this to be an overriding concern at present, given the heightened identification problems associated with this allowance (Xiang, 2010; Bagozzi and Mukherjee, 2012, 373) and the fact that an absence of this feature has far from impeded the applicability and usage of extant zero inflated models (most notably, the ZIP and ZINB).

Replication I

To further illustrate the utility of the BIMNL model, this section presents two replications of existing research. In the first, we replicate Arceneaux and Kolodny's (2009) analysis of phone and door-to-door canvassing endorsements during the 2006 Pennsylvania (PA) state House elections. One core facet of Arceneaux and Kolodny's study employs a MNL model to examine the effects of the aforementioned endorsement strategies (when undertaken by a liberal interest group) upon the following individual *vote choice* options: Democrat, Republican/Other, No response, or Abstained (their MNL baseline category). The corresponding dependent variable

⁸Or, in three instances, as opposed to the MNL model's conditional logit reformulation.

and sample include individual survey responses from approximately 2,000 randomly-selected field experiment subjects across two PA House districts. More detailed operationalizations of all variables used in the survey and analysis are reported in the supplemental appendix.

In brief, Arceneaux and Kolodny hypothesize that a liberal interest group's Democratic candidate endorsement may have crosscutting effects on potential voters depending upon these potential voters' party identification. Herein, the authors posit that a liberal group's endorsement may increase Republican-identifying individuals' support for a Democratic candidate if endorsements serve to educate potential voters on Democratic candidates' actual issue positions, but may decrease potential Republican-leaning respondents' support for a Democratic candidate if liberal interest group endorsements instead serve as a "liberal" heuristic. In testing these claims on individual vote choice, the authors find only weak support for the latter expectation. In particular, while the liberal phone and canvassing treatments each have a negative effect on Republican-identifying individuals' probabilities of Democratic candidate vote choice, each fails to achieve statistical significance at the $p < .10$ level (one-tailed). Moreover, Arceneaux and Kolodny (2009, 763) also surprisingly find that each endorsement strategy has *no effect* upon either Independent or Democratic identifying individuals' probabilities of Democratic vote choice.

Given the arguments presented above, we suspect that these muted findings may be partially attributable to a heterogeneous mixture of routine and occasional nonvoters within the baseline ("abstain") category of the authors' MNL model. As mentioned previously, we contend that nonvoters are best viewed as two distinct groups. The first group correspond to routine nonvoters who virtually never vote and pay little regard to election-specific factors in their repeated decisions to abstain from voting and from politics in general. By contrast, we characterize the second group as occasional voters, who vote intermittently but who may abstain from a given election due to extemporaneous shocks such as temporary resource or knowledge constraints, candidate specific distastes, or weather conditions.

Past research confirms that these two distinct sets of nonvoters exist, and demonstrates that

the effects of voter mobilization efforts and related variables on individual-level turnout differs according to individuals' differential levels of (non)voting propensity. For example, Parry et al. (2008) find that even in high-profile races, the most important determinant of individual-level voter turnout is vote history, and identify a positive effect for campaign communication among “seldom” voters (registered but rarely active) as opposed to intermittent (occasional) and routine voters. Hillygus (2005) draws a similar conclusion, in finding that seldom voters and self-described “unintended” voters are the most likely to be positively influenced by campaign contact. Niven (2004) also distinguishes among different (non)voter types when considering the effect that campaign tactics have on turnout, finding that while contact increases the likelihood of turnout generally, the effect was most dramatic for intermittent voters (those who cast at least one ballot in the recent past) who float in and out of the electorate. Thus, there is good reason to suspect that for any given election, voting-age abstainers will be comprised of a mixture of temporary nonvoters and routine nonvoters.

Moreover, the broad-based sampling scheme used in Arceneaux and Kolodny's field experiment—which randomly selected registered Democrats, unaffiliated voters, and registered Republicans while placing a preference on *infrequently voting Republicans* (2009, 758)—likely reinforced these dynamics for the election at hand. If these contentions are true—and the baseline category of the aforementioned vote choice variable does indeed contain sizable proportions of both routine and occasional nonvoters—then Arceneaux and Kolodny's MNL coefficient estimates may be predominantly a reflection of the effects of their covariates on individuals' likelihoods of total disengagement, rather than on individuals' situational decisions to vote or not vote in the election at hand.⁹ Note that this will be the case even while the authors' turnout treatments are randomly assigned, given the nonrandom assignment of the authors' binary measures of survey respondents' party identifications (which are interacted with both treatments).

To evaluate these claims, we replicate the “pooled” MNL model reported in Arceneaux and Kolodny (2009, 763: Table 2) using both the MNL and BIMNL estimators. We keep the out-

⁹In this manner, the authors' MNL estimates reflect the average treatment effect of covariates on all potential voters, which may be of theoretical or practical interest in some contexts, such as in assessments of an indiscriminate turnout drive's effects—though such estimates could similarly be recovered from the BIMNL model.

come (i.e., MNL) stage specifications for both models identical to those reported in Arceneaux and Kolodny (2009, 763: Table 2) and include a set of theoretically-informed covariates in the inflation stage of our BIMNL model. While our choices for inflation stage covariates are limited to the control variables collected in Arceneaux and Kolodny’s original study, several such variables likely predict individuals’ propensities for being occasional nonvoters rather than routine nonvoters. The authors’ measure of whether or not individuals voted in the past election (*Vote 2004*) is expected to be a strong predictor of this dichotomy according to past findings (Niven, 2004; Parry et al., 2008). Likewise, the open seat race in *District 156* may constitute a relatively higher proportion of occasional nonvoters, given that the opposing district in the authors’ study is characterized as one dominated by a Republican incumbent whose district had proven safe since 1978 (Arceneaux and Kolodny, 2009, 759-760).¹⁰ We also add a number of demographic controls—such as *age* and *female*—to the BIMNL inflation stage in order to capture additional structural factors that past research suggests may be related to habitual or temporary nonvoting (Verba and Nie, 1972; Wolfinger and Rosenstone, 1980; Shields and Godel, 1997).

We report the complete table of outcome and inflation stage coefficient estimates for our MNL and BIMNL models in the supplemental appendix, and focus our attention here on discussing substantive quantities of interest and model fit statistics. To begin, the left side of Table 1 reports the estimated effects of changes in each of our BIMNL model’s inflation stage covariates upon the predicted probability of a nonvoter being an occasional, rather than routine, nonvoter.¹¹ In line with expectations, the BIMNL’s inflation stage results suggest that having *voted in 2004* is associated with a statistically significant 39 percent increase in a nonvoter’s likelihood of being an occasional nonvoter rather than a routine one. Similarly, and relative to District 161, nonvoters in *District 156* are significantly more likely to be occasional nonvoters, perhaps owing to the higher contemporary levels of competitiveness in this district. The only significant demographic predictor of nonvoter type is *age*, which here implies that older non-

¹⁰Since an absence of close or competitive elections may compel lower turnout (Nevitte et al., 2000; Blais, 2000, 60), thereby increasing the population of routine nonvoters.

¹¹Standard errors are calculated via parametric bootstraps ($m = 1000$), while holding all other variables at their means or modes.

voters are associated with higher propensities for routine nonvoting—a relationship which may be attributable to the (likely) nonlinear effects of *age* in this context. The right side of Table 1 indicates that all five of our model fit statistics consistently favor the BIMNL model over the MNL model. Taken together, Table 1 provides an assortment of evidence to suggest that the BIMNL model may offer an improvement over the MNL model in modeling Arceneaux and Kolodny’s vote choice dependent variable.

[Insert Table 1 about here]

We next consider the outcome stage vote choice estimates obtained from our MNL and BIMNL models which, as mentioned above, treat abstention as the baseline category. In deriving quantities of interest for these estimates, we hold our BIMNL inflation stage fixed at $pr(non-inflation = 1)$ which enables us to directly compare our MNL model’s estimates for all individuals to the BIMNL’s hypothetical estimates for only non-inflated individuals.¹² This matches the claims made above, as we argued that by better accounting for routine nonvoters, the BIMNL model provides more accurate estimates of covariate effects among what is typically researchers’ primary sample of interest: occasional and routine voters. In other applications, the simultaneous effects of covariates across *both* stages of one’s BIMNL model may be of most interest,¹³ and these quantities (i) can also be easily derived from the probability statements in equation 5 and (ii) have similarly been shown to exhibit less bias than comparable non-inflated model quantities within extant studies of related (zero-)inflated models (Harris and Zhao, 2007; Bagozzi et al., 2014). In line with Arceneaux and Kolodny’s original study, we focus substantively on the estimated effects of $0 \rightarrow 1$ changes in both the phone and canvassing treatment on the predicted probabilities of Republican, Democrat, and Independent individuals each voting for the *Democratic candidate*. Given the above claims, we estimate these first differences in predicted probabilities relative to (non-inflated) abstention, and plot the resultant mean estimates and confidence intervals in Figure 1.¹⁴

¹²These quantities are comparable to those reported for the MIOP(C) models in Bagozzi and Mukherjee (2012).

¹³See for example, Bagozzi et al. (2014), for a derivation of such quantities in a related zero inflated model.

¹⁴Parametric bootstraps ($m = 1,000$) are used to calculate the standard errors to these effects, while holding all

[Insert Figure 1 about here]

We begin with the Republican identifier subsample (Figure 1a). Consistent with the results reported in Arceneaux and Kolodny (2009), our MNL results suggest that phone and canvassing treatments each have a significant negative effect on Republican voters' probability of voting for the Democratic candidate. Yet, in replicating these results with the BIMNL model, we find that once routine nonvoters have been partitioned from our sample, the estimated effects of either treatment on Republican voters, while larger, are no longer statistically distinguishable from zero at traditional levels. Turning next to Independents (Figure 1b), we find that whereas our MNL model's estimated effects (like those of the authors') yield no significant findings, our BIMNL estimates for these two treatments present mean estimates that are each positive and over double in size to those of the MNL model, and which are statistically significant in the case of the phone treatment. This intuitively suggests that (i) phone treatments increase an Independent's probability of voting Democratic and (ii) the effects of liberal endorsements on this (generally uninformed, nonpartisan) group are underestimated when routine nonvoters are unaccounted for in one's analysis. In line with Arceneaux and Kolodny (2009), we find no significant effect of either treatment among Democratic identifiers (Figure 1c), a null finding that we believe, like Arceneaux and Kolodny, may be attributable to the ceiling imposed by high levels of initial Democratic candidate support among this group. Taken together, it is not evident that liberal group endorsements are a net negative for Democratic candidate voter support (as the authors' original findings imply), as our estimates indicate that, if anything, the net statistically distinguishable effect of these treatments is positive. Finally, one can also observe consistently larger confidence intervals for the BIMNL estimates in Figure 1, which are to be expected given that for the original BIMNL model, unlike the MNL model, one's choice estimates are estimated conditional on estimated probabilities of inflation.¹⁵

In sum, this replication provides evidence—in the form of model fit statistics and inflation

other variables to their means or modes. To ensure comparability to the authors' reported results (Table 4, page 765) we report 90% one-tailed confidence intervals in these simulations, though the significant BIMNL findings discussed below are also significant at the equivalent two-tailed threshold.

¹⁵And the additional stochastic component, and parameters, that this estimated inflation process entails.

stage estimates—to suggest that the BIMNL model may outperform the MNL model in modeling the heterogeneous population of abstainers within Arceneaux and Kolodny’s vote choice dependent variable. In contrast to Arceneaux and Kolodny’s findings, our BIMNL replication suggests that liberal interest group-directed phone and door-to-door canvassing endorsements of PA Democratic House candidates now have a larger, but less precise, negative effect on Republican-identifiers’ choice of voting for a Democratic candidate, and a positive (and in the case of phone treatment, significant) effect on the likelihood of Democratic candidate vote choice among Independents. In both cases, we believe that these differences are attributable to the fact that, once completely disengaged voters have been excluded, the BIMNL model is better able to estimate the effects of Arceneaux and Kolodny’s treatments upon those voters who are susceptible to such interventions. Given the substantive magnitude of our BIMNL findings for Independents, this replication has important theoretical and policy implications for interest groups’ candidate endorsement tactics, a point to which we return in the conclusion.

Replication II

Our second application replicates a recent analysis of individuals’ 2004 U.S. presidential vote choice (Campbell and Monson, 2008). In this study, Campbell and Monson evaluate the interactive effects of (i) gay marriage ban (GMB) ballot initiatives and (ii) potential voters’ religious affiliations on an unordered polytomous dependent variable measuring survey respondents’ stated 2004 vote choice of: Bush, Kerry, or abstention (the baseline category). Drawing on a representative sample of approximately 1,500 individuals from a 2004 Election Panel Study (EPS), Campbell and Monson ultimately find that the presence of GMBs increased mobilization for Bush among evangelical Protestants but increased abstention levels (relative to voting for either candidate) among secular voters. We provide detailed operationalizations for these, and all other, covariates in our supplemental appendix.

While Campbell and Monson’s findings are insightful, we suspect that the same heterogeneous mixture of routine nonvoters and occasional nonvoters discussed above may be present

in the baseline abstention category of the authors' polytomous dependent variable. The EPS survey likely includes substantial proportions of both types of nonvoters, given its nationally-representative sample of potential voters, and given the extant literature on distinct nonvoter types (summarized earlier). If this is the case, then the authors' estimates may be misattributing routine nonvoting effects as having direct effects on vote choice due to heterogeneity in their nonvoting subsample of respondents. We examine whether this is the case by replicating Campbell and Monson's primary specification using MNL and BIMNL models below.¹⁶

As above, our choice of BIMNL inflation stage covariates is limited to the variables available in Campbell and Monson's study. We accordingly identify a number of plausible inflation stage predictors from Campbell and Monson's available covariates. First, we include *education*, as it has proven to be a consistent predictor of voting behavior and political participation in past research (Smets and van Ham, 2013). We posit that educated nonvoters will have higher likelihoods of being temporary nonvoters. For similar reasons, we include the authors' *mobilization index* measure which is also anticipated to be a positive predictor of occasional, rather than routine, nonvoting. We next add Campbell and Monson's dichotomous indicator of individuals reporting their religion as *secular* with the expectation that nonreligious individuals may also be less active in other areas of associational life, such as politics or voting.¹⁷ Finally, we include *age*, but omit *female*, given the inflation stage findings obtained in our earlier application.

The left portion of Table 2 reports estimated changes in the predicted probabilities of non-inflation (i.e., of being an occasional nonvoter rather than routine nonvoter), given reasonable changes in each of our inflation stage covariates.¹⁸ First, we find that a realistic increase in *education* is associated with a small, but significant, positive increase in a nonvoter's likelihood of being an occasional, rather than routine, nonvoter, which is in line with our expectations. Similarly, increases in an individual's *mobilization index* also positively affect the likelihood

¹⁶Our (BI)MNL replications differ from Campbell and Monson in several respects. Most notably, Campbell and Monson use an MNP model to estimate the effects of their covariates on *2004 vote choice*, not a MNL model. The authors also use robust standard errors, whereas we do not. Given that (i) this is a methods exercise and (ii) the authors state that their main substantive results remain unchanged even when employing a MNL model (Campbell and Monson, 2008, Footnote 13: 408), we do not believe our departures to be particularly problematic.

¹⁷For arguments to this effect, see, e.g., Verba et al. (1995); Jones-Correa and Leal (2001).

¹⁸Standard errors are calculated via parametric bootstraps with all other variables at their means or modes.

of non-inflation. In contrast to our earlier application, *age* (now measured on a limited ordinal scale) is also a positive predictor of occasional nonvoting—a departure we find somewhat unsurprising given the different operationalizations and nonmonotonicities that likely underlie this particular relationship. Finally, *secularism* is associated with a large (nine percent) and significant decrease in the probability that a nonvoter is an occasional—rather than routine—nonvoter, which is consistent with our expectation that nonreligious individuals are less likely to be politically active on the whole.

[Insert Table 2 about here]

As the model fit statistics in the right portion of Table 2 reveal, we find only mixed support for the BIMNL model’s “goodness of fit” in this application. In particular, two of our three primary fit statistics favor the BIMNL model, while the third favors the MNL model. The less accurate PRE and Vuong statistics are similarly split between the two models. These inconclusive findings are relatively unsurprising given the low N for the study, but nevertheless underscore the need to be cautious in drawing broader theoretical conclusions from the outcome stage results presented below.

We next turn to the outcome stages of our MNL and BIMNL models. As before, the coefficient estimates for our models’ respective estimation stages are reported in the supplemental appendix. The present section compares our MNL and BIMNL derived first differences in predicted probabilities of votes for Bush and votes for Kerry,¹⁹ given reasonable changes in two key outcome stage covariates: *secularism* and *education*. For our BIMNL estimated first differences in predicted probabilities, we again hold the BIMNL’s inflation stage equation fixed at $pr(non-inflation = 1)$ to facilitate comparisons of our “all respondents” MNL estimates to the BIMNL model’s more specific estimated subset of “voters and occasional voters”. We report these quantities of interest, along with 90% confidence intervals, in Figure 2.

[Insert Figure 2 about here]

¹⁹As above, each first difference is estimated relative to a (non-inflated) voter choosing to abstain, with all other variables held to means or modes and the effects of our covariates on the other choice outcomes restricted to zero.

In addition to identifying the abstention inducing effects of GMB among secular individuals, Campbell and Monson's reported estimates also suggest a puzzling finding, in that the individual coefficient estimate for *secularism* is negative and significant within the "Vote Kerry versus Abstain" equation. This implies that in states without GMB's, secular individuals are less likely (than mainline Protestant individuals) to vote for Kerry than to abstain. As Figure 2a indicates, we find an equivalent effect in our MNL replication. However, once routine abstainers are partitioned from the sample, our comparable BIMNL estimates reveal that this effect disappears, as the estimates are now statistically indistinguishable from zero. Our estimated quantities of interest for a key control variable also highlights interesting differences between the MNL and BIMNL estimates. Specifically, the quantities of interest in Figure 2b demonstrate that a well known positive determinate of turnout, *education*, has no direct effect on vote choice (relative to abstention), once its significant effects on the probability of abstention inflation are accounted for—a finding that potentially alters how we interpret the moderating effects of education within GOTV type mobilization efforts.

In sum, our second application further highlights the benefits of the BIMNL model over the MNL model, but also underscores some potential challenges. While we obtained a number of theoretically consistent findings for the inflation stage of our BIMNL model—such as those related to *education*, *mobilization index*, and *secularism*—our model fit statistics were ambiguous as to whether the BIMNL did, in fact, provide a better model fit than a comparable MNL model. Because such inconclusive findings are apt to occur in practice—especially when applied to data sets with fewer than 2,000 observations (as was the case here)—we encourage researchers to be cautious and to rely on theory when selecting between the BIMNL and MNL model in such instances. These weak foundations notwithstanding, our outcome stage estimates do suggest that the BIMNL model offers novel improvements in theory testing. Specifically, the puzzling estimated effect for *secularism* among individuals' propensities to vote for Kerry dissipates once baseline inflation is accounted for, while the effect of *education* becomes more nuanced.

Conclusion

Questions of individual vote choice, when applied via representative surveys of voting age citizens, are apt to include heterogeneous mixtures of nonvoters. Extant research suggests that one subset of these self-reported nonvoters will correspond to occasional nonvoters, who abstain from voting in a specific election due to temporary factors such as short term economic shocks, weather conditions, or personal factors. Yet, many other nonvoters are likely to be routine nonvoters who, due to structural or personal reasons, have chosen to abstain from the political process entirely. These nonvoters are thus largely immune to the short-term factors affecting the occasional abstainers mentioned above. Ignoring this distinction can lead one to conflate the estimated direct effect of both sets of determinants on nonvoting and, by implication, the factors affecting actual voters' candidate choices in MNL models. Hence, it is imperative that public opinion scholars account for these disparate populations of nonvoters when "did not vote/abstain" serves as the reference category in empirical models of vote choice.

This paper presents a new discrete choice estimator—the BIMNL model—to allow researchers to more accurately account for this variation in populations of nonvoters. After deriving this model, we demonstrate that ignoring baseline-choice category inflation within empirical studies of vote choice can potentially affect one's parameter estimates and theoretical conclusions. Here, our replication findings offer new theoretical insights into the determinants of individual vote choice in American elections. For instance, our first application demonstrates that the effects of a liberal interest group's Democratic candidate endorsement may in fact produce a net positive, rather than a net negative, outcome for Democratic candidates across all potential voters, once routine nonvoters are accounted for in the BIMNL model. Similarly, our second application suggests that a number of commonly understood direct predictors of turnout and vote choice, such as education and secularism, may instead primarily affect these outcomes through their effects on routine versus occasional nonvoting, rather than through any direct effect on individuals' active decisions to vote or abstain in a given election.

This study can be extended in several directions. As mentioned above, the BIMNL model

could be generalized to the MNP setting. Notwithstanding the MNP concerns raised earlier, such an extension would afford researchers the opportunity to properly account for baseline category inflation in instances where they believe their polytomous dependent variable to be both unordered and in broader violation of the IIA assumption. Furthermore, a BIMNP model could also provide a feasible framework for the provision of correlated error terms between the two estimation stages of an inflated polytomous choice model (under assumptions of bivariate normality). Allowances for correlated disturbances between inflation and outcome equations, although challenging for estimation and identification, have been shown to be extremely useful to both theory and estimation within inflated *ordered* probit (IOP) estimators (Harris and Zhao, 2007; Bagozzi et al., 2014; Bagozzi and Mukherjee, 2012), as well as in applications of inflated probit estimators (e.g., Xiang, 2010), and thus would likely be of great interest to survey researchers within the polytomous choice setting as well.

The insights presented above may also facilitate future research within the area of political behavior by allowing researchers to better test and develop theories of routine and occasional nonvoting. As such, our paper has important implications for political behavior research both within academia as well as in the more applied settings of political campaigns and activism. By accounting for routine nonvoters, the BIMNL model provides more nuanced estimates of how various factors affect the voting behavior of the very individuals that campaigns and scholars (arguably) want to understand and motivate the most: those who have voted in past elections but not necessarily routinely. It is this population of disproportionately nonpartisan, uninformed, and heuristic-dependent (non)voters who, if effectively mobilized, can have a substantively large effect on an electoral outcome. As the 2008 election demonstrated, the ability to encourage voter turnout among populations of Americans with inconsistent voting records (e.g., young voters between the ages of 18-29) can significantly bolster support for one party, or candidate, over another. Thus, accurately understanding how GOTV activities differentially affect occasional versus routine nonvoters could be crucial for many political campaigners, interest organizations, and scholars of political behavior.

REFERENCES

- Arceneaux, K. and R. Kolodny (2009). Educating the least informed: Group endorsements in a grassroots campaign. *American Journal of Political Science* 53(4), 755–770.
- Bagozzi, B. E., D. W. J. Hill, W. H. Moore, and B. Mukherjee (2014). Modeling two types of peace: The zero-inflated ordered probit (ziop) model in conflict research. *Journal of Conflict Resolution*.
- Bagozzi, B. E. and B. Mukherjee (2012). A mixture model for middle category inflation in ordered survey responses. *Political Analysis* 20(3), 369–386.
- Beger, A., J. H. DeMeritt, W. Hwang, and W. H. Moore (2011). The split population logit (spoplogit): Modeling measurement bias in binary data. working paper.
- Blias, A. (2000). *To Vote or Not to Vote: The Merits and Limits of Rational Choice Theory*. Pittsburgh: The University of Pittsburgh Press.
- Born, R. (1990). Surge and decline, negative voting, and the midterm loss phenomenon: A simultaneous choice analysis. *American Journal of Political Science* 36, 615–645.
- Braumoeller, B. F. (2013). The anna karenina principle in international relations. Presented at the annual convention of the International Studies Association.
- Brooks, R., M. N. Harris, and C. Spencer (2012). Inflated ordered outcomes. *Economic Letters* 117, 683–686.
- Burger, M., F. van Oort, and G.-J. Linders (2009). On the specification of the gravity model of trade: Zeros, excess zeros and zero-inflated estimation. *Spatial Economic Analysis* 4(2), 167–190.
- Campbell, D. E. and J. Q. Monson (2008). The religion card: Gay marriage and the 2004 presidential election. *Public Opinion Quarterly* 72(3), 399–419.

- Dow, J. K. and J. W. Endersby (2004). Multinomial probit and multinomial logit: A comparison of choice models for voting research. *Electoral Studies* 23(1), 107–222.
- Gerber, A. S. and D. P. Green (2000). The effect of a nonpartisan get-out-the-vote drive: An experimental study of leafletting. *Journal of Politics* 62(3), 846–857.
- Greene, W. H. (1994). Some accounting for excess zeros and sample selection in poisson and negative binomial regression models. Working Paper EC-94-10: Department of Economics, New York University.
- Greene, W. H. (2011). *Econometric Analysis* (Seventh ed.). Englewood Cliffs, NJ: Prentice-Hall.
- Hall, D. B. and J. Shen (2010). Robust estimation for zero-inflated poisson regression. *Scandinavian Journal of Statistics* 37, 237–252.
- Harris, M. N. and X. Zhao (2007). A zero-inflated ordered probit model, with an application to modelling tobacco consumption. *Journal of Econometrics* 141(2), 1073–1099.
- Hillygus, D. S. (2005). Campaign effects and the dynamics of turnout intention in election 2000. *Journal of Politics* 67(1), 50–68.
- Jones-Correa, M. A. and D. L. Leal (2001). Political participation: Does religion matter? *Political Research Quarterly* 54(4), 751–770.
- Kalmoe, N. P. and S. Piston (2013). Is implicit prejudice against blacks politically consequential? evidence from the amp. *Public Opinion Quarterly* 77(1), 305–222.
- King, G., R. O. Keohane, and S. Verba (1994). *Designing Social Science Inquiry: Scientific Inference in Qualitative Research*. Princeton, NJ: Princeton University Press.
- Lacy, D. and B. C. Burden (1999). The vote-stealing and turnout effects of Ross Perot in the 1992 US presidential election. *American Journal of Political Science* 42, 233–255.

- Lehtonen, E., J. Kuuluvainen, E. Pouta, M. Rekola, and C.-Z. Li (2003). Non-market benefits of forest conservation in southern finland. *Environmental Science & Policy* 6, 195–204.
- Liu, A. H. (2011). Linguistic effects of political institutions. *Journal of Politics* 73(1), 125–139.
- Nevitte, N., A. Blais, E. Gidengil, and R. Nadeau (2000). *Unsteady State: The 1997 Canadian Federal Election*. Oxford University Press.
- Niven, D. (2004). The mobilization solution? face-to-face contact and voter turnout in a municipal election. *Journal of Politics* 66(3), 868–884.
- Parry, J., J. Barth, M. Kropf, and E. T. Jones (2008). Mobilizing the seldom voter: Campaign contact and effects in high profile elections. *Political Behavior* 30, 97–113.
- Plane, D. L. and J. Gershtenson (2004). Candidates' ideological locations, abstention and turnout in u.s. midterm elections. *Political Behavior* 26(1), 69–93.
- Propper, C. (2000). The demand for private health care in the uk. *Journal of Health Economics* 19(6), 855–876.
- Sanders, M. S. (1999). Unified models of turnout and vote choice for two-candidate and three-candidate elections. *Political Analysis* 7, 89–115.
- Shields, T. G. and R. K. Godel (1997). Participation rates, socioeconomic class biases, and congressional elections: A crossvalidation. *American Journal of Political Science* 41(2), 683–691.
- Smets, K. and C. van Ham (2013). The embarrassment of riches? a meta-analysis of individual-level research on voter turnout. *Electoral Studies* 32, 344–359.
- Staub, K. E. and R. Winkelmann (2012). Consistent estimation of zero-inflated count models. *Health Economics* 22(6), 673–686.
- Verba, S. and N. H. Nie (1972). *Participation in America: Political Democracy and Social Equality*. New York: Harper & Row.

- Verba, S., K. L. Schlozman, and H. Brady (1995). *Voice and Equality: Civic Voluntarism in American Politics*. Cambridge, MA: Harvard University Press.
- Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica* 57(2), 307–333.
- Wolfinger, R. E. and S. J. Rosenstone (1980). *Who Votes?* New Haven: Yale University Press.
- Xiang, J. (2010). Relevance as a latent variable in dyadic analysis of conflict. *Journal of Politics* 72(2), 484–498.
- Zipp, J. F. (1985). Perceived representativeness and voting: An assessment of the impact of “choices” vs. “echoes”. *American Political Science Review* 79(1), 50–61.

Table 1: Inflation Stage & Model Fit Results for 2006 PA Vote Choice

Inflation Covariate	First Differences in Pr(Occasional Nonvoter)	Model Fit Statistics	
		MNL	BIMNL
0 → 1 Δ in <i>Vote 2004</i>	0.39 (0.18 ↔ 0.60)	lnL -2110.59	-2069.97
53 → 71 Δ in <i>Age</i>	-0.05 (-0.08 ↔ -0.02)	Vuong	✓
0 → 1 Δ in <i>Female</i>	0.01 (-0.02 ↔ 0.04)	PRE	✓
2 → 4 Δ in <i>House Size</i>	0.01 (-0.02 ↔ 0.04)	AIC	✓
		BIC	✓
0 → 1 Δ in <i>District 156</i>	0.06 (0.02 ↔ 0.10)	LR	✓

Changes in non-binary variables \approx mean +1SD changes. Values in parentheses are 90% confidence intervals.

Table 2: Inflation Stage & Model Fit Results for 2004 Presidential Vote Choice

Inflation Covariate	First Differences in Pr(Occasional Nonvoter)	Model Fit Statistics		
			MNL	BIMNL
5 → 7 Δ in <i>Education</i>	0.01 (0.001 ↔ 0.03)	lnL	-553.50	-545.27
		Vuong		✓
1 → 2 Δ in <i>Mobilization</i>	0.02 (0.003 ↔ 0.06)	PRE	✓	
		AIC		✓
2 → 3 Δ in <i>Age</i>	0.01 (0.001 ↔ 0.04)	BIC	✓	
		LR		✓
0 → 1 Δ in <i>Secular</i>	-0.09 (-0.28 ↔ -0.001)			

Changes in non-binary variables \approx mean +1SD changes. Values in parentheses are 90% confidence intervals.

Figure 1: Effects of Phone & Canvassing on Predicted Probabilities of 2006 PA Vote Choice

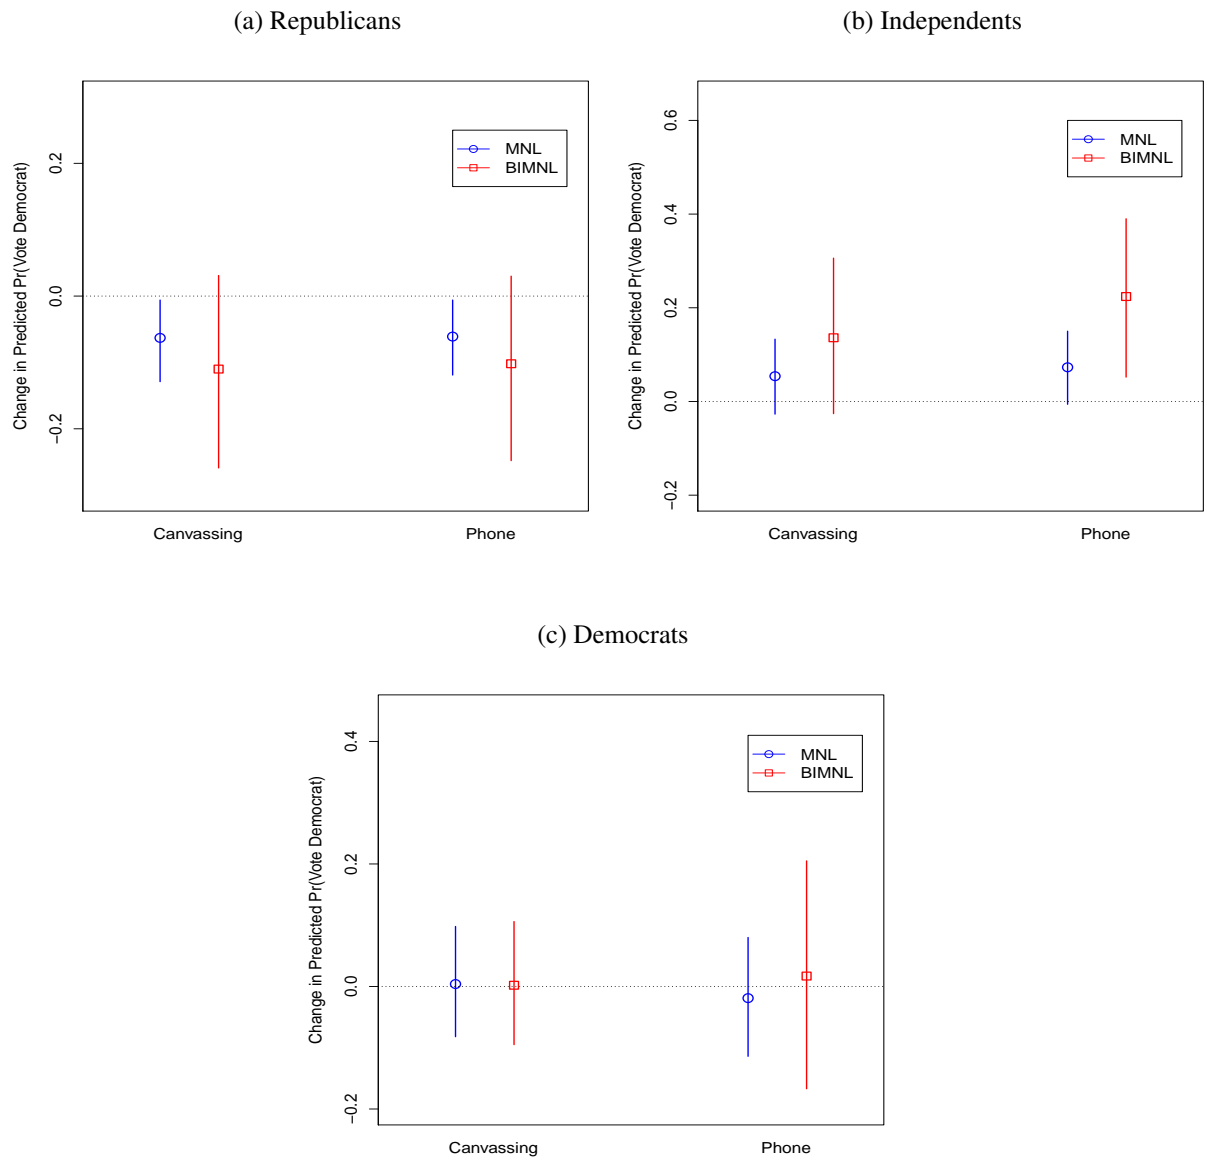


Figure 2: Effects of Selected Covariates on Predicted Probabilities of 2004 Vote Choice

(a) $0 \rightarrow 1$ change in *Secular* (When $GMB=0$)

(b) $5 \rightarrow 7$ change in *Education*

