

What are the Consequences of Consequentiality?¹

Joseph Herriges²

Catherine Kling

Chih-Chen Liu

and

Justin Tobias

Department of Economics

Iowa State University

August 8, 2007

¹The authors are listed alphabetically, without assignment of seniority. We thank without implicating Laura Taylor, Kerry Smith, and several workshop participants for valuable discussions and comments on aspects of this project. Funding for the Lakes Project data used in this analysis was provided by the Iowa Department of Natural Resources and the U.S. Environmental Protection Agency's Science to Achieve Results (STAR) program. Although the research described in the article has been funded in part by the U.S. Environmental Protection Agency's STAR program through grant R830818, it has not been subjected to any EPA review and therefore does not necessarily reflect the views of the Agency, and no official endorsement should be inferred.

²Contact author contact information: 260 Heady Hall, Department of Economics, Iowa State University, Ames, IA 50011. e-mail: jaherrig@iastate.edu. phone: (515)294-4964.

What are the Consequences of Consequentiality?

Abstract

We offer an empirical test of a theoretical result in the contingent valuation literature. Specifically, it has been argued from a theoretical point of view that survey participants who perceive a survey to be “consequential” will respond to questions truthfully regardless of the degree of perceived consequentiality. Using survey data from the Iowa Lakes Project, we test this supposition. Specifically, we employ a Bayesian treatment effect model in which the degree of perceived consequentiality, measured as an ordinal response, is permitted to have a structural impact on willingness to pay (WTP) for a hypothetical environmental improvement. We test our theory by determining if the WTP distributions are the same for each value of the ordinal response.

In our survey data, a subsample of individuals were randomly assigned supporting information suggesting that their responses to the questionnaires were important and will have an impact on policy decisions. In conjunction with a Bayesian posterior simulator, we use this source of exogenous variation to identify the structural impacts of consequentiality perceptions on willingness to pay, while controlling for the potential of confounding on unobservables. We find evidence consistent with the “knife-edge” theoretical results, namely that the willingness to pay distributions are equal among those believing the survey to be at least minimally consequential, and different for those believing that the survey is irrelevant for policy purposes.

1 Introduction

The question regarding the accuracy with which stated preference methods can be used to ascertain individual valuations for environmental goods and services continues to be hotly debated in the environmental literature. A wide variety of empirical tests have been conducted to date, including studies checking for agreement with other elicitation methods (e.g., revealed preference techniques such as recreation demand or hedonic pricing) or contrasting revealed and stated preference responses in an experimental setting. While such tests are quite informative, they are often context-specific and reduced form in nature, not addressing the fundamental issues as to why and under what conditions stated preference methods will succeed or fail.

In contrast, Carson, Groves, and Machina (CGM, 2000), by detailing the incentive and informational properties of stated preference questions, provide the basis for discerning key elements of a successful contingent valuation (CV) exercise. Specifically, CGM argue that respondents can be predicted to answer a dichotomous-choice referendum question in a manner that is consistent with expected utility maximization if they perceive the survey to be “consequential.” That is, if respondents believe the result of the survey might potentially influence an outcome they care about, they will answer CV questions truthfully since it is their dominant strategy to do so. If accurate, this result has clear and important implications regarding how a researcher should handle survey responses: all data arising from respondents who believe the survey is at least minimally consequential can be assumed to provide truthful answers to survey questions.

There have been several studies to date testing this theoretical consequentiality condition and its impact on respondents’ preference revelation. These, to our knowledge, have been almost exclusively carried out through field experiments. Though useful, these experimental studies have largely taken place outside of the context of an actual CV exercise and require the analyst to inform participants of the degree of consequentiality associated with the exercise. These field experiments have the important advantage that the researcher can directly control the degree of consequentiality, but their disadvantage is that this direct control is not typical of CV surveys. Thus, findings from such field experiments may be difficult to transfer to the survey arena.

In this paper we take a different but complementary approach, eliciting respondents' perception of consequentiality directly in a CV survey. Specifically, we use respondents' perceptions of consequentiality elicited in two CV surveys (the 2003 and 2005 Iowa Lakes Surveys) to determine whether individuals have different perceptions concerning the degree of consequentiality of the valuation exercise and whether these perceptions affects respondents' willingness to pay (WTP) in the pattern predicted by the CGM theoretical work. In both surveys, respondents were asked whether they would vote in favor of a referendum to improve water quality at a lake. Respondents were also asked to answer, on a scale from one to five, how likely it was that the survey results would influence decisions in the state concerning water quality programs. Thus, a measure of the degree to which respondents perceived the survey as consequential was directly elicited. Based on the CGM arguments, respondents who do not believe that the survey is consequential should be omitted from the sample for estimation purposes, since they do not have an incentive to respond to the referendum question truthfully. Additionally, the distribution of WTP from respondents with differing views concerning the degree of consequentiality could be tested for equality, thereby testing the fundamental CGM argument.

Since respondents who indicate a high degree of consequentiality may do so because they also place a high value in the proposed water quality improvement project, there is a potential endogeneity, or unobserved confounding problem. To address this concern, a split sample treatment was administered in the 2005 survey. Specifically, half of the sample was provided with a highlighted article from the *Iowa Conservationist* — the magazine of the Iowa Department of Natural Resources (IDNR), the state agency with primary responsibility for water quality control — indicating that IDNR was already using results from the survey in their policy decisions and planned to continue to do so. Our assumption, which is borne out empirically, is that the presentation of this information will positively affect the respondents' perceived degree of consequentiality. This exogenous treatment, which should only impact willingness to pay through its indirect effect on perceived consequentiality, will aid us in estimating the “causal” impacts of consequentiality perceptions on WTP, as we will describe below.

Making use of the information treatment in the 2005 survey, we can then explore the impact of perceived consequentiality on willingness to pay within the framework of a standard triangular treatment-response model. We proceed using a Bayesian approach and derive

and employ a new algorithm that improves upon standard estimation methods, and can be applied by other practitioners seeking to fit models with a similar structure. Specifically, since our consequentiality responses are ordinal, our model must contend with the estimation of cutpoint values. It is well-documented in the literature that standard Gibbs sampling schemes in such models can suffer from very poor mixing, particularly in moderately large data sets, thus producing imprecise and potentially inaccurate posterior inference [e.g., Cowles (1996), Nandram and Chen (1996)].³ Our proposed posterior simulator offers significant improvements by sampling the cutpoints, latent willingness to pay and latent consequentiality variables in a single step rather than sampling each component from its corresponding complete posterior conditional distribution. The application of our algorithm to this data reveals results that are broadly consistent with CGM’s theoretical predictions.

The outline of the paper is as follows. Section 2 provides a brief overview of related studies in the literature, while Section 3 describes the Iowa Lakes Project and the data used in our empirical analysis. Both the model and the simulator used to characterize the posterior distribution of the parameters of interest are described in Section 4. Details of the posterior simulator are deferred to an appendix. Section 5 provides the empirical results and the paper concludes with a summary in section 6.

2 Related Literature

Our investigation is certainly not the first to test this consequentiality condition on respondents’ preference revelation. Cummings *et al.* (1997), for example, conducted an experiment involving real and hypothetical referenda to compare how respondents behave in these two settings. They find that respondents are more likely to vote “yes” in the hypothetical setting than in the real setting and conclude that hypothetical referenda yield biased estimates of WTP. Cummings and Taylor (1998) explore this issue further by investigating how the hypothetical bias of the dichotomous choice referendum varies with the degree of consequentiality. They find that participants in different treatments, including hypothetical, probabilistic, and

³“Poor mixing” in simulation based approaches, such as Gibbs sampling, refers to the problem of the algorithm yielding high serial correlation among the draws from the distribution of interest. This high correlation, in turn, implies that a large number of simulation draws are required to accurately estimate characteristics of the distribution, such as its mean, mode or variance.

real referendum, behave significantly different from one another. Specifically, the probability of a respondent voting “yes” falls if the probability of the referendum being binding rises. They also find that respondents’ voting behavior is significantly different from an actual referendum unless there is high probability that the referendum will be binding (e.g., greater than fifty percent).

Carson *et al.* (2004) conduct a similar study with a field experiment. In the experiment, participants were informed the probability of the dichotomous choice referendum being binding. In contrast to Cummings and Taylor (1998), they find the knife-edge theoretical results suggested in CGM: as long as the probability of consequentiality exceeds zero even by a small amount, participants respond in the same manner as in an actual referendum. They suggest that results from inconsequential (hypothetical) referendum should not be used to make inference about how CV works in consequential referendum.

Bulte *et al.* (2005) explore the impact of consequentiality in a survey setting by telling respondents that the results of the survey will be available to policy makers. They do not directly elicit respondents’ reactions to this information. Their WTP estimates obtained from the survey with “cheap talk” and the survey with a consequentiality device are significantly smaller than those obtained in a purely hypothetical survey. In addition, the WTP estimates obtained from the survey with cheap talk are not significantly different from the WTPs obtained from the survey with a consequential device. Their results are also consistent with the CGM predictions.

3 The Iowa Lakes Data

This study employs data from two years of the “Iowa Lakes Project,” a four-year study and panel data collection effort aimed at understanding recreational use and the value of water quality in the primary recreational lakes of Iowa. The project began in 2002 with mail surveys sent to a random sample of 8,000 Iowa residents, obtaining detailed information regarding their visitation patterns to approximately 130 lakes, as well as standard socio-demographic data (e.g., age, education, income and gender). In subsequent years, surveys

were sent to those households completing a survey in the prior year.⁴ Standard follow-up procedures were followed in each year of the survey, including a postcard reminder mailed two weeks after the initial mailing and a second copy of the survey mailed one month later. Households were provided a \$10 incentive for completing the survey each year, helping to provide for response rates ranging 62 to 72 percent across the four years of the project.

The contingent valuation (and consequentiality) responses modeled in the current paper were included only in the 2003 and 2005 versions of the questionnaire. In 2003, a referendum style CV question was posed to estimate WTP for a water quality improvement project at one of eight focus lakes targeted in the study. These lakes were selected in consultation with the Iowa Department of Natural Resources. In addition to being geographically dispersed, the target lakes are each of policy interest since various restoration projects are being considered for them. Each respondent was asked a CV question for only one of the focus lakes. The survey described the current water quality information of the lake, including water clarity, water color, water odor, health concerns from algae blooms and bacteria level, and variety and quantity of fish. A photograph, illustrating the water clarity and water color, was provided to help respondents picture the current water quality of the focus lake vividly. Prior to the main valuation question, cheap talk text was used to remind respondents about the incentive and information properties of stated preference question in a further attempt to elicit truthful responses to the CV questions.

A water quality improvement project regarding the focus lake was then proposed that outlined the methods to achieve the water quality improvement at each lake, such as dredging and building protection strips around the perimeter. Respondents were asked whether they would vote in favor of the referendum to improve the lake where bid values (ranging from \$100 to \$600 payable over a five year period) were randomly assigned to households in the sample receiving the survey. Finally, the perceived consequentiality of the survey was elicited from respondents by asking the following question: “How likely do you think it is that the results of surveys such as this one will affect decisions about water quality in Iowa lakes?” Possible responses to this question ranged from 1 to 5 where a 1 denoted “no effect at all” (i.e., completely inconsequential) and 5 denoted “definite” effects.

⁴A second random sample was added in the panel in 2003 to fill in for non-deliverable surveys in 2003 and return the sample to a total of 8000 households. No additional households were added after 2003.

Table 1: 2003 Summary Statistics By Level of Consequentiality

Variable	$c = 1$	$c = 2$	$c = 3$	$c = 4$	$c = 5$
Num. Observations (n)	307	832	1995	1294	356
Percentage YES	0.13	0.27	0.31	0.35	0.38
Percent Not Visited	0.48	0.32	0.29	0.24	0.22
Percent Male	0.70	0.65	0.68	0.69	0.65
Income/ \$1,000	54.1	58.8	55.9	58.2	49.4
Std. Income	(40.7)	(37.8)	(34.7)	(38.1)	(37.3)
Age	58.31	53.42	53.22	51.29	55.51
Std. Age	(16.41)	(15.85)	(15.61)	(15.27)	(16.73)
College	0.55	0.68	0.68	0.74	0.64

Table 1 lists a summary of the data obtained in 2003. In our sample, the more consequential the respondents view the survey, the higher the observed “yes” rate to the proposed referendum. The “yes” rate rises monotonically from 13% when respondents believe the survey is inconsequential to 38% when respondents believe the survey is definitely consequential. Also of note, respondents who do not visit the lakes tend to view the survey inconsequential. As Figure 1 illustrates, the voting pattern of respondents who report any degree of consequentiality (i.e., anyone who reports a 2, 3, 4, or 5 to the consequentiality question) is consistent with economic theory: the “yes” rate falls as the bid value rises. In this sample, it is also the case that the “yes” rate of the inconsequential group is clearly lower than that of consequential group at each bid value.

Recall that the CGM knife-edge result states that respondents’ vote on the CV question will be the same as an actual referendum as long as the respondents perceive that there is even a small chance that the survey is consequential. To test this conjecture, we could estimate a simple, single-equation WTP function for each level of reported consequentiality and test whether the distributions are statistically different from one another. However, as noted in the introduction, there is the potential endogeneity of respondents’ answers to the consequentiality question. In particular, respondents who have a high WTP may believe the survey is consequential due to their perceptions of the importance of the project. Likewise, respondents who indicated a low degree of consequentiality may do so because they place a low value in the proposed water quality improvement project.

To control for this potential endogeneity, a split sample treatment was randomly administered




Figure 1: 2003 Percentage “Yes” by Bid Value for the Five Consequentiality Groups

in the 2005 Iowa Lakes survey.⁵ In doing so, we took advantage of a magazine article and letter from the director of the state’s Department of Natural Resources that had recently been published concerning how results from the Iowa Lakes Project were being used by the department to prioritize water quality projects. Specifically, half of the sample was provided with a copy of the director’s letter and an article from the *Iowa Conservationist* indicating that the IDNR was already using results from the surveys in their policy decisions and planned to continue to do so. This information was also highlighted in the cover letter to the respondents. Thus, direct evidence of consequentiality was applied to this (randomly selected) subsample. The exogenous treatment to respondents’ perception of consequentiality allows us to separate the impact of consequentiality in the CV responses. The following section outlines a triangular treatment-response model of the 2005 data.

⁵The consequentiality questions were included in only a portion of the 2005 survey, so that the final sample available in 2005 is 1996.

4 Brief Description of Model and Associated Posterior Simulator

To test for the potential existence of differential impacts of perceived consequentiality on willingness to pay, we consider the following two equation system:⁶

$$c_i^* = \mathbf{x}_{ci}\boldsymbol{\beta}_c + \epsilon_{ci} \quad (1)$$

$$w_i^* = \mathbf{x}_{wi}\boldsymbol{\beta}_w + \bar{\mathbf{c}}_i\boldsymbol{\delta}_c + \epsilon_{wi} \quad (2)$$

where

$$\begin{bmatrix} \epsilon_{ci} \\ \epsilon_{wi} \end{bmatrix} \Big| \mathbf{x}_{ci}, \mathbf{x}_{wi} \stackrel{iid}{\sim} N \left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \sigma_{cw} \\ \sigma_{cw} & \sigma_w^2 \end{pmatrix} \right].$$

Equation (1) corresponds to the latent “consequentiality” equation, while equation (2) corresponds to the latent “willingness to pay” equation. The latent c_i^* and w_i^* are not observed, but their values generate an observed binary willingness to pay indicator w_i and reported consequentiality variable c_i . In our data, c_i takes on 5 different ordered values, ranging from regarding the survey as being completely irrelevant to the making of policy ($c_i = 1$) to regarding the survey as having a definite impact on policy ($c_i = 5$). We thus relate the observables w_i and c_i to their latent values as follows:

$$w_i = \begin{cases} 1 & \text{if } w_i^* \geq B_i \\ 0 & \text{if } w_i^* < B_i, \end{cases} \quad c_i = \begin{cases} 1 & \text{if } -\infty < c_i^* \leq 0 \\ 2 & \text{if } 0 < c_i^* \leq \alpha_3 \\ 3 & \text{if } \alpha_3 < c_i^* \leq \alpha_4 \\ 4 & \text{if } \alpha_4 < c_i^* \leq \alpha_5 \\ 5 & \text{if } \alpha_5 < c_i^* < \infty, \end{cases}$$

where B_i is a hypothetical price (bid) proposed to individual i and $\bar{\mathbf{c}}_i$ in (2) is a 1×5 vector with a one in the c_i^{th} column and zeros elsewhere.⁷ The parameter vector $\boldsymbol{\delta}_c$ (within the treatment-response framework of (1) and (2) and with a valid exclusion restriction or instrument in \mathbf{x}_c) is commonly interpreted as capturing the “causal” impact of consequentiality categories on WTP. The j^{th} element of $\boldsymbol{\delta}_c$, δ_{cj} , denotes the parameter associated with the event that $c_i = j$. For example, the willingness to pay for an individual reporting a perceived consequentiality of, say, $c_i = 3$ is given by $w_i^* = \mathbf{x}_{wi}\boldsymbol{\beta}_w + \delta_{c3} + \epsilon_{wi}$. Thus, differences in WTP among individual respondents due to differences in their perceived consequentiality of the survey should be captured by variations among the δ_{cj} ’s.

⁶We adopt the convention of using boldface to denote vector or matrix quantities and capital letters to denote matrices.

⁷Note, with this formulation of the model, an intercept must be excluded from \mathbf{x}_{wi} .

4.1 A Bayesian Approach to Estimation

In practice, a variety of methods can be employed to fit the model in (1) and (2). We have experimented with some traditional alternatives, such as simulated maximum likelihood, and found it difficult to obtain successful convergence and good performance, even under ideal conditions with correct model specification. Instead, we choose to employ a Bayesian estimation approach with data augmentation [e.g., Tanner and Wong (1987), Koop, Poirier and Tobias (2007, chapter 14)], which is quite similar in spirit to classical likelihood-based techniques, and is indistinguishable from these alternatives in large sample sizes when suitably “diffuse” or “non-informative” priors are employed.

Simply adopting the Bayesian methodology, however, is not a guarantee of good performance and, indeed, some important complications arise concerning how best to estimate the parameters of our equation system under this approach. For example, with typical conjugate priors (i.e., priors that yield posteriors of the same form), fitting the model in (1)-(2) using the standard Gibbs sampler (an algorithm that iteratively samples from the complete posterior conditionals for each set of parameters) is reasonably straight-forward, following the work of Albert and Chib (1993).⁸ However, the standard Gibbs sampler in these types of models can suffer from slow mixing, particularly in large data sets. Thus, unless something is done to improve the standard algorithm, use of Bayesian techniques offers little or no improvement over these other alternatives. To this end, we offer a refinement of the standard Gibbs algorithm which offers substantial performance improvements.

First, we introduce a reparameterization, following the suggestion of Nandram and Chen (1996) in the context of a single-equation ordered probit model. Specifically, we let $\gamma = \alpha_5^{-1}$ and define $\tilde{x} \equiv \gamma x$. [That is, the $\tilde{\cdot}$ notation simply denotes the operation of taking the original variable and multiplying it by γ .] Multiplying (1) by γ and adjusting the rule mapping the latent c_i^* into the observed c_i produces the following equivalent model:

⁸A (conditionally) *conjugate* prior is a prior that yields a (conditional) posterior distribution of the same functional form.

$$\tilde{c}_i^* = \mathbf{x}_{ci} \tilde{\boldsymbol{\beta}}_c + \tilde{\epsilon}_{ci} \quad (3)$$

$$w_i^* = \mathbf{x}_{wi} \boldsymbol{\beta}_w + \bar{\mathbf{c}}_i \boldsymbol{\delta}_c + \epsilon_{wi} \quad (4)$$

where

$$\begin{bmatrix} \tilde{\epsilon}_{ci} \\ \epsilon_{wi} \end{bmatrix} \Big| \mathbf{x}_{ci}, \mathbf{x}_{wi} \stackrel{iid}{\sim} N \left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \gamma^2 & \gamma\sigma_{cw} \\ \gamma\sigma_{cw} & \sigma_u^2 \end{pmatrix} \right],$$

$$w_i = \begin{cases} 1 & \text{if } w_i^* \geq B_i \\ 0 & \text{if } w_i^* < B_i, \end{cases} \quad c_i = \begin{cases} 1 & \text{if } -\infty < \tilde{c}_i^* \leq 0 \\ 2 & \text{if } 0 < \tilde{c}_i^* \leq \tilde{\alpha}_3 \\ 3 & \text{if } \tilde{\alpha}_3 < \tilde{c}_i^* \leq \tilde{\alpha}_4 \\ 4 & \text{if } \tilde{\alpha}_4 < \tilde{c}_i^* \leq 1 \\ 3 & \text{if } 1 < \tilde{c}_i^* < \infty. \end{cases}$$

Let

$$\tilde{\boldsymbol{\Sigma}} = \begin{pmatrix} \tilde{\sigma}_c^2 & \tilde{\sigma}_{cw} \\ \tilde{\sigma}_{cw} & \sigma_w^2 \end{pmatrix} = \begin{pmatrix} \gamma^2 & \gamma\sigma_{cw} \\ \gamma\sigma_{cw} & \sigma_w^2 \end{pmatrix}.$$

We choose to work with the parameters:⁹

$$\tilde{\boldsymbol{\beta}} = [\tilde{\boldsymbol{\beta}}'_c \quad \boldsymbol{\beta}'_w \quad \boldsymbol{\delta}'_c]' \quad \text{and} \quad \tilde{\boldsymbol{\Sigma}}^{-1}$$

rather than the original parameterization of the model. The primary reasons for doing this are threefold. First, the reparameterization helps to improve the mixing of the posterior simulations [e.g., Nandram and Chen (1996)], as freeing up the variance parameter rather than restricting it to unity tends to reduce the lagged autocorrelations among the parameter draws. Second, the reparameterization eliminates one component of the cutpoint vector (which is important in our case since $\tilde{\boldsymbol{\alpha}}$ is two-dimensional, and routines for calculating bivariate, but not trivariate, normal probabilities are readily available).¹⁰ Finally, the reparameterization eliminates the diagonal restriction on the 2×2 covariance matrix, enabling the use of a Wishart prior on $\tilde{\boldsymbol{\Sigma}}^{-1}$ and a standard Gibbs step for sampling its elements.¹¹

Since the structural parameters $\boldsymbol{\beta}_c$, σ_{cw} , σ_w^2 and $\boldsymbol{\alpha} = [\alpha_3 \quad \alpha_4 \quad \alpha_5]'$ (rather than their transformations) are ultimately of interest, we can, of course, recover these by using the appropriate

⁹We do not, however, take up the issue of the induced priors on the original structural parameters.

¹⁰The appendix provides details regarding the specifics of the required calculations for sampling the transformed cutpoints.

¹¹This latter point should probably be regarded as a minor contribution in light of the fact that one can simply write, for example, $\epsilon_{wi} = \sigma_{cw}\epsilon_{ci} + u_i$, with u_i independent of ϵ_{ci} , substitute this into equation (2), and then develop a posterior simulator. See, e.g., Li (1998) and Deb, Munkin and Trivedi (2006) for examples of this strategy.

inverse transformations at each iteration of the sampler:

$$\boldsymbol{\beta}_c = \tilde{\boldsymbol{\beta}}_c/\gamma, \quad \sigma_{eu} = \tilde{\sigma}_{eu}/\gamma, \quad \boldsymbol{\alpha} = [(\tilde{\alpha}_3/\gamma) \quad (\tilde{\alpha}_4/\gamma) \quad \gamma^{-1}]'.$$

4.2 Priors and the Joint Posterior

As fully described in the appendix, we fit this model using Gibbs sampling coupled with the Metropolis-Hastings algorithm and auxiliary variable Gibbs [e.g., Damien *et al.* (1999)]. Before talking about the specifics of this algorithm, we must first derive our *augmented posterior* distribution. This augmented posterior involves adding the latent $\tilde{\mathbf{c}}^*$ and \mathbf{w}^* to the joint posterior distribution. Under prior independence this joint distribution can be represented as:

$$\begin{aligned} p(\tilde{\mathbf{c}}^*, \mathbf{w}^*, \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}, \tilde{\boldsymbol{\alpha}} | \mathbf{c}, \mathbf{w}) &\propto p(\mathbf{c}, \mathbf{w} | \tilde{\mathbf{c}}^*, \mathbf{w}^*, \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}, \tilde{\boldsymbol{\alpha}}) p(\tilde{\mathbf{c}}^*, \mathbf{w}^* | \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}, \tilde{\boldsymbol{\alpha}}) p(\tilde{\boldsymbol{\beta}}) p(\tilde{\boldsymbol{\Sigma}}) p(\boldsymbol{\alpha}) \quad (5) \\ &= p(\boldsymbol{\alpha}) p(\tilde{\boldsymbol{\beta}}) p(\tilde{\boldsymbol{\Sigma}}) \prod_{i=1}^n p(\tilde{c}_i, w_i | \tilde{c}_i^*, w_i^*, \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}, \boldsymbol{\alpha}) p(\tilde{c}_i^*, w_i^* | \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}), \end{aligned}$$

where the product term follows from the assumed (conditional) independence across observations. For the first term on the right-hand side of this product, note

$$\begin{aligned} p(c_i, w_i | \tilde{c}_i^*, w_i^*, \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}, \tilde{\boldsymbol{\alpha}}) &= p(c_i | w_i, \tilde{c}_i^*, w_i^*, \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}, \tilde{\boldsymbol{\alpha}}) p(w_i | \tilde{c}_i^*, w_i^*, \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}) \\ &= p(c_i | \tilde{c}_i^*, \tilde{\boldsymbol{\alpha}}) p(w_i | w_i^*) \end{aligned}$$

where

$$p(c_i | \tilde{c}_i^*, \tilde{\boldsymbol{\alpha}}) = I(\tilde{\alpha}_{c_i} < \tilde{c}_i^* \leq \tilde{\alpha}_{c_i+1}) \quad (6)$$

with $I(\cdot)$ denoting the standard indicator function, and

$$p(w_i | w_i^*) = I(w_i = 0) I(w_i^* < B_i) + I(w_i = 1) I(w_i^* \geq B_i). \quad (7)$$

Apart from the priors, the remaining piece in our joint posterior in (5) follows immediately from our normality assumption:

$$\tilde{c}_i^*, w_i^* | \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}} \stackrel{ind}{\sim} N \left(\left[\begin{array}{c} \mathbf{x}_{ci} \tilde{\boldsymbol{\beta}}_c \\ \mathbf{x}_{wi} \tilde{\boldsymbol{\beta}}_w + \bar{\mathbf{c}}_i \boldsymbol{\delta}_c \end{array} \right], \tilde{\boldsymbol{\Sigma}} \right).$$

The model is completed by choosing priors of the forms:

$$\tilde{\boldsymbol{\beta}} \sim N(\boldsymbol{\mu}_\beta, \mathbf{V}_\beta) \quad (8)$$

$$\tilde{\boldsymbol{\Sigma}}^{-1} \sim W(\mathbf{R}, \nu) \quad (9)$$

$$p(\tilde{\boldsymbol{\alpha}}) = p(\tilde{\alpha}_3)p(\tilde{\alpha}_4|\tilde{\alpha}_3) = I(0 < \tilde{\alpha}_3 < 1)\frac{1}{1 - \tilde{\alpha}_3}I(\tilde{\alpha}_3 < \tilde{\alpha}_4 < 1). \quad (10)$$

For the last of these three priors, we impose the ordering restriction on the cutpoints. Unconditionally, we specify a prior for the smallest transformed cutpoint, $\tilde{\alpha}_3$ which is uniform over its support, while $\tilde{\alpha}_4$ is specified to be conditionally uniform over $(\tilde{\alpha}_3, 1)$.

The posterior simulator for fitting this model is completely described in the appendix. It is worth noting here, however, a second refinement of our algorithm relative to the standard Gibbs method. Specifically, we utilize a blocking step to sample a group of parameters jointly, rather than conditionally, thus significantly improving the mixing of our posterior simulations. That is, our algorithm employs a blocking procedure wherein the 2×1 vector of transformed cutpoints $(\tilde{\boldsymbol{\alpha}})$ and latent data $(\tilde{\mathbf{c}}^*, \mathbf{w}^*)$ are sampled together in a *single step* rather than sampled from their respective conditional posterior distributions. The clear need for such an algorithm was motivated by generated data experiments suggesting poor mixing of the simulated cutpoints when utilizing a standard Gibbs sampler for both the reparameterized and original representation of the model.¹² In cases where objects of interest are functions of the cutpoints, such as the WTP distribution described below, the improved mixing offered by our algorithm would seem to offer a substantial benefit.

4.3 Willingness to Pay

As we noted in the introduction, a primary focus of our paper involves the calculation and comparison of WTP distributions for individuals with varying degrees of perceived consequentiality. To this end, we consider a posterior predictive exercise and seek to obtain the WTP distribution for a hypothetical out-of-sample agent with observed characteristics \mathbf{x}_{cf} and \mathbf{x}_{wf} and consequentiality response $c_f = j$. The model in (3)-(4), assumed to hold for this agent's outcomes, implies that the posterior predictive willingness to pay distribution

¹²Again, we refer here to the case where no blocking steps are employed, and the cutpoints are drawn element-by-element from their (uniform) posterior conditional distributions.

for this representative agent, conditioned on the model parameters $\mathbf{\Gamma}$, is given as

$$\begin{aligned} p(w_f^* | c_f = j, \mathbf{\Gamma}) &= p(w_f^* | \tilde{\alpha}_j < \tilde{c}_f^* \leq \tilde{\alpha}_{j+1}, \mathbf{\Gamma}) \\ &= [\Pr(\tilde{\alpha}_j < \tilde{c}_f^* \leq \tilde{\alpha}_{j+1}) | \mathbf{\Gamma}]^{-1} \int_{\tilde{\alpha}_j}^{\tilde{\alpha}_{j+1}} p(w_f^*, \tilde{c}_f^* | \mathbf{\Gamma}) d\tilde{c}_f^*. \end{aligned}$$

Factoring the latter distribution into the conditional $\tilde{c}_f^* | w_f^*$ and the marginal for w_f^* enables evaluation of the integral. Specifically, we obtain¹³

$$\begin{aligned} p(w_f^* | c_f = j, \mathbf{\Gamma}) &= \left[\Phi \left(\frac{\tilde{\alpha}_{j+1} - \mathbf{x}_{cf}\boldsymbol{\beta}_c - [\tilde{\sigma}_{cw}/\sigma_w^2](w_f^* - \mathbf{x}_{wf}\boldsymbol{\beta}_w - \delta_{cj})}{\tilde{\sigma}_c \sqrt{1 - \tilde{\rho}_{cw}^2}} \right) - \right. \\ &\quad \left. \Phi \left(\frac{\tilde{\alpha}_j - \mathbf{x}_{cf}\boldsymbol{\beta}_c - [\tilde{\sigma}_{cw}/\sigma_w^2](w_f^* - \mathbf{x}_{wf}\boldsymbol{\beta}_w - \delta_{cj})}{\tilde{\sigma}_c \sqrt{1 - \tilde{\rho}_{cw}^2}} \right) \right] \frac{\phi(w_f^*; \mathbf{x}_{wf}\boldsymbol{\beta}_w + \delta_{cj}, \sigma_w^2)}{\Phi \left(\frac{\tilde{\alpha}_{j+1} - \mathbf{x}_{cf}\boldsymbol{\beta}_c}{\tilde{\sigma}_c} \right) - \Phi \left(\frac{\tilde{\alpha}_j - \mathbf{x}_{cf}\boldsymbol{\beta}_c}{\tilde{\sigma}_c} \right)}. \end{aligned} \quad (11)$$

In practice, we would like to calculate a variety of statistics associated with this willingness to pay distribution. For example, quantities like the posterior mean, posterior standard deviation, and posterior median (which is critical for determining whether or not a particular referendum is likely to pass under majority vote) are certainly of interest and are policy relevant.

Since these quantities are not easily calculated analytically, given the rather messy form of (11), it seems desirable to pursue a numerical alternative by obtaining draws from (11) and using these draws to calculate any statistic of interest. However, (11) is not of an immediately recognizable functional form, thus calling into question the feasibility of this numerical scheme. It can be shown, however, (a proof is available upon request), that a draw from (11) can be obtained by the following procedure:

First, sample

$$z_f^* \sim TN_{(z_f, \bar{z}_f)}(0, 1)$$

where $TN_{(a,b)}(\mu, \sigma^2)$ denotes a normal distribution with mean μ and variance σ^2 truncated to lie in the interval (a, b) ,

$$\underline{z}_f = \frac{\tilde{\alpha}_j - \mathbf{x}_{cf}\boldsymbol{\beta}_c}{\tilde{\sigma}_c} \quad \text{and} \quad \bar{z}_f = \frac{\tilde{\alpha}_{j+1} - \mathbf{x}_{cf}\boldsymbol{\beta}_c}{\tilde{\sigma}_c}.$$

Then, set

$$w_f^* = \pi_{0f} + \pi_{1f} z_f^* + \pi_{2f} \epsilon_f,$$

¹³Note that this specification reduces to the marginal for w_f^* when ρ_{cw} (and thus $\tilde{\sigma}_{cw}$) equals zero.

where

$$\epsilon \sim N(0, 1), \quad \pi_{0f} = \mathbf{x}_{\mathbf{w}f}\boldsymbol{\beta}_{\mathbf{c}} + \delta_{cj}, \quad \pi_1 = \sigma_w \tilde{\rho}_{cw}, \quad \text{and} \quad \pi_2 = \sigma_w \sqrt{1 - \tilde{\rho}_{cw}^2}$$

as a draw from (11).¹⁴ The above procedure provides a simple to implement series of steps (i.e., they only require drawing from standard distributions) for sampling from the WTP distribution directly, whence any feature of this distribution can be calculated. In practice, the steps above are repeated for each post-convergence simulation $\boldsymbol{\Gamma}$, and the collection of the resulting draws are then used to calculate the desired quantity or quantities. To our knowledge, a description of such a method in the presence of an endogeneity problem offers a new and hopefully valuable contribution to the literature.

Beyond these specific statistics, it is also of interest to provide a plot of the full posterior predictive distribution. While (11) is conditioned on the parameters $\boldsymbol{\Gamma}$, the Bayesian approach handles parameter uncertainty by integrating the parameters out of the conditional posterior predictive. In the context of our problem, we would like to obtain

$$p(w_f^* | c_f = j, \mathbf{c}, \mathbf{w}) = \int p(w_f^* | c_f = j, \boldsymbol{\Gamma}) p(\boldsymbol{\Gamma} | \mathbf{c}, \mathbf{w}) d\boldsymbol{\Gamma}, \quad (12)$$

since w_f^* , in the case of random sampling, is independent of the past data \mathbf{c}, \mathbf{w} given $\boldsymbol{\Gamma}$. Simply plotting the left-hand side of (12) can be done without having to take draws from the (conditional) posterior predictive, as described above. Instead, we can provide a ‘‘Rao-Blackwell’’ approximation of this posterior predictive by noting

$$p(w_f^* | c_f = j, \mathbf{c}, \mathbf{w}) \approx \frac{1}{R} \sum_{r=1}^R p(w_f^* | c_f = j, \boldsymbol{\Gamma} = \boldsymbol{\Gamma}^{(r)}),$$

where $\boldsymbol{\Gamma}^{(r)}$ is the r^{th} post-convergence simulation from the posterior distribution, R denotes the total number of simulations, and $p(w_f^* | c_f = j, \boldsymbol{\Gamma} = \boldsymbol{\Gamma}^{(r)})$ is known from (11). In our empirical results provided in section 5, we use this method to plot various posterior predictives and also sample directly from (11) to calculate specific features of this distribution whose values are not known analytically.

¹⁴It is also worth noting that (11) is closely related to the *skew normal* random variable, and our procedure for sampling from it can be regarded as a slight generalization of the additive construction representation of this random [e.g., Arnold and Beaver (2003, 2004)].

5 Empirical results

Using the model described section 4 and the algorithm fully documented in the appendix, we fit our two equation triangular treatment-response model. We run our posterior simulator for 50,000 simulations and discard the first 5,000 simulations as the burn-in. Numerous generated data experiments, which are not reported here for the sake of brevity, revealed that our algorithm mixed reasonably well (i.e., the lagged autocorrelations among our parameter simulations were not severe, and for some parameters, resembled what would be obtained under iid sampling), and consistently recovered parameters of the data generating process in a variety of experimental designs (i.e., those with high and low degrees of endogeneity and those with balanced and non-balanced frequencies for the observed consequentiality responses). For our priors, we set $\boldsymbol{\mu}_\beta = \mathbf{0}$, $\mathbf{V}_\beta = (200^2)\mathbf{I}_k$, $\nu = 5$ and $R = \mathbf{I}_2$, which are reasonably non-informative choices, suggesting that information coming from the data will dominate information added through our prior.

5.1 Diagnostic Checking

Before diving into our empirical results, we first provide some information regarding the overall performance of our model. Our assumptions, significantly stronger than those required for popular alternatives like standard IV, involve a complete description of the joint distribution of observables $p(\mathbf{c}, \mathbf{w}|\boldsymbol{\Gamma})$.¹⁵ To this end, it is important to assess the appropriateness of these assumptions and to document potential deficiencies associated with the model specification.

There are numerous diagnostic checks for investigating the reasonableness of a model’s assumptions in the Bayesian paradigm, including, for example, the use of QQ plots [e.g., Lancaster (2004, Chapter 2)], posterior predictive p-values [e.g., Gelman, Carlin, Stern and Rubin (2004, section 6.3)], and other comparisons of specific features of the model to their counterparts in the observed data [e.g., Koop, Poirier and Tobias (2007, Chapter 11)]. We focus here on one such exercise, which involves simulation from the posterior predictive distribution.

¹⁵Note, given the nonlinearity of our model, standard IV would not be directly applicable.

The idea behind this exercise is to try and *replicate* the distribution of the observed data, or specific features of the data that are important to the problem at hand, based on output from the model itself. If the model is legitimate for the analysis of a given data set, then it is reasonable to require that the model will produce a distribution of predicted outcomes that mimics what is found in the raw data. For example, if a simple wage regression were run under the assumption that wages were normally distributed, the posterior predictive density of wages would be symmetric, while the observed density of hourly wages would have a pronounced right-skew. This exercise would therefore reveal that an assumption of the model - namely that of normality - is inappropriate for the observed data.

To this end, we obtain a vector $\mathbf{y}_j^{\text{rep}} = [\mathbf{c}_j^{\text{rep}'} \ \mathbf{w}_j^{\text{rep}'}]'$ from $p(\mathbf{y}^{\text{rep}}|\mathbf{y})$ where

$$p(\mathbf{y}^{\text{rep}}|\mathbf{y}) = \int_{\mathbf{R}_\Gamma} p(\mathbf{y}^{\text{rep}}|\Gamma, \mathbf{y})p(\Gamma|\mathbf{y})d\Gamma, \quad (13)$$

and $\mathbf{y}_j^{\text{rep}} \sim \mathbf{y}^{\text{rep}}|\Gamma = \Gamma_j, \mathbf{y}$, with Γ_j representing the j^{th} post-convergence draw from our posterior simulator. The density $p(\mathbf{y}^{\text{rep}}|\Gamma, \mathbf{y})$ is simply the likelihood function assumed by the given model (which does not depend on \mathbf{y} given Γ). In these simulations, we obtain a vector $\mathbf{y}_j^{\text{rep}}$ from the conditional density $p(\mathbf{y}^{\text{rep}}|\Gamma = \Gamma_j, \mathbf{y})$ by choosing exactly the same \mathbf{x}_c and \mathbf{x}_w values as those that our found in our sample of data. This motivates our use of the notation “rep” to denote *replications* of the observed data from the posterior predictive, i.e., they are “a re-run of history on the assumption that the model is what generates histories” [Lancaster (2004, pp. 90-91)].

The act of generating a series of y^{rep} variates in this way can be used to reveal how well the model fares in reproducing the actual outcome distributions in our sample. Thus, for each post-convergence simulation Γ_j , we draw the latent vector $\tilde{\mathbf{c}}^{*,\text{rep}}$ using the \mathbf{x}_c values in the sample, calculate the discrete indicator \mathbf{c}^{rep} from this latent data, and then generate a latent willingness to pay $\mathbf{w}^{*,\text{rep}}$ from the conditional distribution $p(\mathbf{w}^{*,\text{rep}}|\tilde{\mathbf{c}}^{*,\text{rep}}, \Gamma_j, \mathbf{x}_{w,\text{rep}} = \mathbf{x}_w)$. From this final draw, we calculate the replicated willingness to pay as $\mathbf{w}^{\text{rep}} = \max\{0, \mathbf{w}^{*,\text{rep}}\}$. We then compare features of these replicated simulations to those same features from the observed data.

As we should expect, the posterior predictive frequencies associated with the consequentiality responses closely match those found in the raw data. Specifically, the posterior means associated with the fraction of observations for $c^{\text{rep}} \in \{1, 2, 3, 4, 5\}$ were 4.4, 13.2, 38.9, 34.5 and

8.9, while those found in the raw data were 4.1, 13.3, 39.2, 34.5 and 8.8, respectively. At the individual level, however, the results are less encouraging. The posterior mean (and posterior standard deviation) of the fraction of \mathbf{c}^{rep} values that correctly predict their counterpart in the observed data is .31 (.011). This, of course, is an improvement over randomly guessing based on a uniform prior (which would achieve a success rate of 20 percent), but does suggest that our covariates will not play a strong role in describing variation in consequentiality perceptions. In particular, and as will be discussed in the following subsection, the instrument will only play a small role in tracing out variation in consequentiality responses. This will lead to reduced posterior precision associated with the “causal” effect parameters δ_c in (4), which will make it more difficult to get conclusive information regarding tests of parameter equality. It is important to note, however, that these results do not necessarily signal a deficiency of the model’s assumptions, but simply document that it is difficult to determine why some individuals believe the surveys are useful for policy purposes, while others believe they have little or no value for policymaking.

To assess the adequacy of the willingness to pay distribution, we consider two primary statistics: $\Pr(w^{rep} > 400|\mathbf{y})$ and $\Pr(w^{rep} < 60|\mathbf{y})$. The motivation for this focus is to determine if our model can adequately recover these “tail” probabilities, while recognizing that our model is not constructed to necessarily recover these quantities.¹⁶ In practice, it is often difficult for a model to predict such “extreme” outcomes, making statistics such as these a useful metric for assessing the performance of our model. Though the largest bid in our sample is \$600 and the smallest is \$20, we focus our calculations instead on more “interior” bid points due to small samples sizes and observed non-monotonicity in our sample. Specifically, just looking at the raw data, we find that 34.6, 28.4 and 34.5 percent of our sample say they are willing to pay bid amounts of \$400, \$500 and \$600, respectively for the proposed water quality improvement. Thus, it is impossible for *any* model to match all of these probabilities,¹⁷ leading us to average these frequencies for all bids greater than or equal to \$400 and use this as a basis for comparison. Performing the required calculations, we find a posterior mean (and standard deviation) associated with the event $w^{rep} > 400$ equal to .35 (.02), which is close to and within two standard deviations of the fraction in the raw data, .32. On the other end

¹⁶A focus on the mean, for example, is a comparison that should typically be avoided, since parameters of the model often capture the mean of the distribution of interest. The above quantities clearly rely on the accuracy of the normality assumption.

¹⁷That is, in terms of the raw data, the fraction of observations willing to pay at least \$600 will exceed the fraction of observations willing to pay at least \$500.

of the WTP distribution, we find a posterior mean (and standard deviation) associated with the event $w^{rep} < 60$ equal to .44 (.19), while the observed fraction in the raw data is quite close to this value at .39. Taken together, the results of these calculations do not seem to provide any direct evidence against our model’s assumptions or that any type of re-modeling beyond what is asserted in (3) and (4) is obviously necessary.

5.2 Parameter Estimates and Discussion of Results

Presented in Table 2 are posterior means and standard deviations associated with the parameters of our model. The continuous variables age and income are standardized to have mean zero and unit variance for interpretation purposes and to suitably scale the values of the associated slope parameters. In the willingness to pay equation, we treat a response of $c_i = 2$ as the “base category” (i.e., it takes the role of the intercept), so that all other elements of δ_c are interpreted relative to this base group.

The coefficients associated with demographic variables in the willingness to pay equation are generally consistent with our prior expectations. Females, older respondents, more educated individuals, and those with higher incomes are more willing to pay for a given water quality improvement. Specifically, females have a (latent) willingness to pay that is approximately \$73 more than males, on average, while the college educated have a latent WTP approximately \$136 higher than their non-college counterparts. Similarly, a one-standard deviation increase in age (approximately 15 years) is associated with an expected increase in latent willingness to pay of approximately \$26, while a similar standard deviation increase in income (approximately \$38,000) is associated with an expected increase in latent willingness pay of about \$92. Perhaps with the exception of the age coefficient (whose posterior probability of being positive is still reasonably high at .91), the marginal posteriors associated with these coefficients place nearly all mass over positive values.

As suggested previously in our section on diagnostic checking, few variables emerge as influential predictors of consequentiality perceptions. Fortunately, our instrument, the indicator denoting the receipt of the highlighted article from the *Iowa Conservationist*, is positively associated with the perceived degree of consequentiality, and the posterior places virtually all mass over positive values. Simulations from the posterior distribution were also used

to show that receipt of the highlighted article increased the probability than an “average” college educated-female will be in the highest consequentiality groups (i.e., $c_i \in \{4, 5\}$) by about 4.5 percent. Thus our instrument, though not highly influential, does play some role in explaining variation in consequentiality perceptions. The only remaining predictor which plays some role in the consequentiality equation is the college indicator, with the college educated more likely to believe that the survey is definitely consequential, and less likely to believe that it has no impact on policy decisions.

Table 2:
Posterior Means, Standard Deviations and
Probabilities of Being Positive

<i>Consequentiality Equation</i>			
Variable	$E(\cdot y)$	Std($\cdot y$)	Pr($\cdot > 0 y$)
Constant	1.64	.059	1.00
<i>Iowa Conservationist</i> Article	.103	.046	.984
Age	-.02	.025	.185
Female	.003	.051	.521
College	.071	.053	.914
Income	-.008	.026	.397
<i>Willingness to Pay Equation</i>			
Constant (Barely Consequential)	51.34	68.56	.779
Not Consequential	-192.1	106.3	.032
Moderately Consequential	34.12	60.39	.719
Consequential	57.03	98.31	.720
Definitely Consequential	-63.63	150.44	.331
Age	26.48	19.30	.909
Female	72.91	39.77	.974
College	135.99	41.56	.999
Income	91.79	22.16	1.00
<i>Covariance Matrix Parameters and Cutpoints</i>			
σ_u^2	3.82×10^5	7.46×10^4	1.00
ρ_{eu}	.099	.076	.925
α_3	.785	.045	1.00
α_4	1.89	.051	1.00
α_5	3.07	.060	1.00

Interestingly, we also find evidence suggesting an important role for unobserved confounding in our application. Specifically, the posterior mean of the correlation between the errors in (1) and (2), denoted ρ_{cw} , was found to be .099, and $\text{Pr}(\rho_{cw} > 0|y) = .925$. This suggests,

with a reasonable degree of posterior certainty, that unobserved characteristics leading the agent to believe that a survey is likely to be consequential are positively related with unobserved factors contributing to WTP. Said differently, characteristics of people who believe that their actions (such as responses to a survey) can make a difference correlate positively with (unobserved) factors making a person willing to pay to see such improvements come to fruition. Though the magnitude of this correlation is small, the fact that the marginal posterior concentrates on positive values suggests the need to control for the potential endogeneity of perceived consequentiality.

A more formal investigation regarding the importance of unobserved confounding can be conducted by providing a test of the hypothesis $\rho_{eu} = 0$. From a Bayesian point of view, this is conducted by calculating the *posterior odds ratio*, K_{12} :

$$K_{12} \equiv \frac{p(\mathcal{M}_1|y)}{p(\mathcal{M}_2|y)} = \frac{p(y|\mathcal{M}_1) p(\mathcal{M}_1)}{p(y|\mathcal{M}_2) p(\mathcal{M}_2)}, \quad (14)$$

where \mathcal{M}_j denotes *model* j , $j = 1, 2$, $p(\mathcal{M}_j|y)$ denotes the *posterior probability* of model j , $p(y|\mathcal{M}_j) = \int_{\Theta_j} p(y|\theta_j, \mathcal{M}_j)p(\theta_j|\mathcal{M}_j)d\theta_j$ is the *marginal likelihood* associated with model j and $p(\mathcal{M}_j)$ is the prior probability associated with model j . A common convention is to assume that both models are equally probable *a priori* so that the prior odds ratio cancels in (14) and thus the posterior odds ratio K_{12} equals the ratio of marginal likelihoods (also known as the Bayes factor).

To apply this model comparison method to our particular problem, let \mathcal{M}_1 denote the restricted model with $\rho_{eu} = 0$ and let \mathcal{M}_2 denote the unrestricted model in (3) and (4). For hypotheses like this one, provided parameters common to \mathcal{M}_1 and \mathcal{M}_2 are given the same priors in both models, the ratio of posterior odds in (14) reduces to: [see, Koop, Poirier and Tobias (2007, page 69)]

$$K_{12} = \frac{p(\mathcal{M}_1|y)}{p(\mathcal{M}_2|y)} = \frac{p(\rho_{eu} = 0|\mathcal{M}_2, y)}{p(\rho_{eu} = 0|\mathcal{M}_2)},$$

the ratio of the marginal posterior and prior ordinate at zero under the unrestricted model. Calculating this ratio produces a posterior odds value equal to .35, suggesting that the unrestricted model with $\rho_{eu} \neq 0$ is favored over the restricted version by a factor of approximately 2.84 to 1. This, again, provides evidence supporting the need to control for unobserved confounding and the endogeneity of consequentiality perceptions in our application.

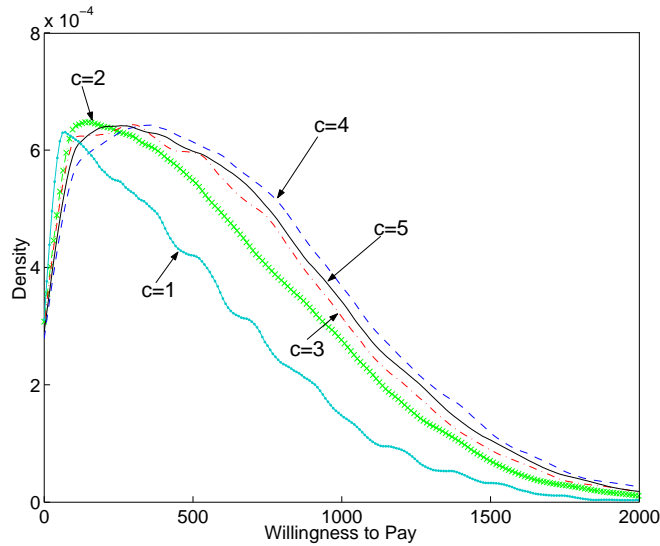


Figure 2: Willingness to Pay Posterior Predictive Distributions for 5 Consequentiality Groups

5.3 Do the data support the knife-edge theoretical predictions?

As mentioned previously in the paper, we are primarily interested in determining if the WTP distributions, and specific features of those distributions, are equal for individuals with varying, but positive, perceptions regarding the survey’s consequentiality. Specifically, previous theoretical work would seem to suggest that those individuals believing the survey to be completely inconsequential (i.e., $c_i = 1$) may report a different willingness to pay than everyone else, while those believing the survey to be at least minimally consequential (i.e., $c_i > 1$) should report the same willingness to pay.

Our first step in the investigation of this question is to plot the posterior predictive willingness to pay distributions for the various consequentiality groups, using the methodology described in section 4.3 of the paper. These plots are presented in Figure 2, and are calculated, for simplicity, for a college-educated female of average age and income who received the article touting the importance of the survey. What we take away from Figure 2 is some evidence that is broadly consistent with the knife-edge theoretical predictions. The implied WTP distributions for the group believing the survey to be completely inconsequential ($c = 1$) seems rather different from those of the other groups, and the remaining densities seem mostly similar.

Of course, the predictive WTP distributions in the figure unite two conceptually different sources of information. First, they account for the direct “structural” impacts of consequentiality perceptions on willingness to pay through the δ_c coefficients. Second, they also account for the role of unobservables and the fact that individuals believing the survey to be very consequential have, on average, unobserved characteristics that also make them more likely to pay for a given environmental improvement. As a test of the theory, what seems to be of primary interest is a comparison of the structural coefficients. If differences in the WTP distributions for those with $c \geq 2$ arise only because of differences in unobservables across groups, then the findings of our analysis are still consistent with the theoretical, knife-edge predictions. That is, any differences in willingness to pay can be interpreted as arising only from the observational nature of the survey itself, since, in the survey area, we are not able to randomly assign individuals to groups where all factors will be, on average, balanced across the groups. Conversely, if the structural parameters are significantly different, this suggests persistent WTP differences that can not be explained through differences in unobservables, and could potentially be argued to be at odds with previous theoretical and experimental predictions.

Table 3: Posterior Probabilities and Posterior Predictive WTP Statistics Across Different Levels of Consequentiality

Variable	Consequentiality Response				
	1	2	3	4	5
$\Pr(\delta_{c1} > \cdot y)$	—	—	—	—	—
$\Pr(\delta_{c2} > \cdot y)$.98	—	—	—	—
$\Pr(\delta_{c3} > \cdot y)$.97	.44	—	—	—
$\Pr(\delta_{c4} > \cdot y)$.94	.51	.64	—	—
$\Pr(\delta_{c5} > \cdot y)$.73	.29	.20	.08	—
Median WTP	0	\$171.06	\$264.9	\$355.0	\$306.2
Mean WTP	\$ 218.44	\$349.42	\$ 404.67	\$ 468.12	\$434.22

To compare these structural coefficients more formally, we report a variety of quantities of interest in Table 3. For the first set of entries of the table, we calculate posterior probabilities of the form $\Pr(\delta_{cj} > \delta_{ck} | \mathbf{y}) \forall j, k \in \{1, \dots, 5\}$.¹⁸ These parameters, and the corresponding probabilities mentioned above, are what researchers might typically look to first as a “test”

¹⁸For example, the table can be used to show that $\Pr(\delta_{c2} > \delta_{c1} | \mathbf{y}) = .98$, and similarly for the remaining entries.

of the theory. As is evident from the table, the subgroup who perceives the survey as being completely irrelevant for policy purposes clearly has lower willingness to pay than those who believe the survey to be somewhat consequential, and also a lower (though this statement is made with far less certainty) WTP than those regarding the survey as completely consequential. This result is consistent with the CGM theoretical predictions, wherein the behavior of individuals believing that the survey is irrelevant can not be characterized; our results suggest that these individuals have different WTP reports than the remaining groups, supporting the notion that these agents and their reports can be excluded in practice.

There is also reasonably strong evidence that, among those believing the survey to be at least minimally consequential, the structural impacts across groups are not obviously distinguishable. This is evident from the fact, apart from the comparison between groups 4 and 5, most of the probabilities found in the table do not seem to provide strong evidence of parameter differences.

To investigate this issue more formally, we revisit equation (14) and calculate Bayes factors associated with the individual restrictions $\delta_{cj} = 0$, for $j \in \{1, 3, 4, 5\}$. We would regard our analysis as providing support for the theoretical predictions if the Bayes factor associated with δ_{c1} favored keeping the indicator in the model, while those associated with δ_{c3} , δ_{c4} and δ_{c5} favored excluding the indicators from the model. When performing these calculations, we find Bayes factors in favor of the zero restriction equal to .36, 2.93, 1.74 and 1.20 for $j = 1, 3, 4, 5$, respectively. Thus, by a factor of approximately 2.8 to 1, we support keeping the indicator for the inconsequential group in the model, and also support (to varying degrees), dropping the remaining consequentiality indicators.¹⁹ Again, these results provide support for the knife-edge theoretical predictions; similar WTP distributions emerge for those believing the survey to be minimally consequential, while those believing the survey to be completely inconsequential have different WTP distributions, even after accounting for unobserved differences across the groups.

Given the non-experimental nature of our data, it is important to note that our results are

¹⁹It is seemingly important here to mention Bartlett's paradox, where the restricted model receives larger and larger support as the prior becomes increasingly non-informative. More non-informative priors (those with prior standard deviations for the δ parameters greater than 200) will only lead to stronger evidence in favor of the restrictions, i.e., more support for the equality of WTP distributions for those believing the survey to be at least somewhat consequential.

subject to an alternative interpretation. The issue here arises from a subtle, yet conceptually important issue regarding truthful revelation of preferences versus equality of WTP distributions. With our non-experimental data, it could be the case that individuals in different consequentiality groups simply have different willingnesses to pay. This problem is likely to be substantially mitigated given our efforts to control for observed and unobserved (through the selection component of the model) differences across groups, but its presence could still remain. In other words, if we fit the model and find that the δ_c parameters are different, we could interpret this result in one of two ways: individuals are truthful (consistent with the CGM predictions) yet have different WTP distributions, or the parameter differences are suggestive that individuals are not truthful (at odds with CGM). If this were to happen, our assumption would be that we are adequately controlling for unobserved differences across groups (through controlling for the endogeneity of consequentiality) so that differences in the parameters are most likely attributable to non-truthful revelation of preferences. On the other hand, if we find evidence of parameter equality for those believing the survey is at least somewhat consequential, and inequality for those believing the survey is irrelevant, then again we have two interpretations: parameter equality implies the same behavior for the consequential groups, thus supporting the idea of truthful responses, or parameter equality persists yet the reports are not truthful. This latter explanation of the results is more difficult to rationalize, as it would require a type of misreporting on the part of agents believing the survey to be minimally consequential which cancels, on average, leading to parameter equality, and a different kind of behavior for the inconsequential agents which leads to differences in parameters. While this event is certainly possible, it seems to be reasonably unlikely. Thus, our results providing evidence of parameter equality among those believing the survey has some degree of consequentiality, and parameter inequality for those believing the survey is completely irrelevant for policy, are, we believe, consistent with and supportive of the CGM predictions. However, given the above discussion, we acknowledge the need to interpret this finding with some caution.²⁰

²⁰This interpretation concern potentially applies to field experiments as well in which respondents are informed of the consequentiality of the survey. Specifically, by informing respondents thus, individuals may alter their perception of the value of the good being considered.

6 Conclusion

In this paper we have offered an empirical evidence in support of a theoretical result in the contingent valuation literature. Specifically, we have investigated the hypothesis of whether willingness to pay distributions are equal for those individuals who believe the survey has at least some potential for shaping policy decisions.

Using a treatment-response model that controls for unobserved confounding and exploits a survey design in which a subsample of individuals are randomly provided supporting material documenting the importance of the study, we find evidence that is broadly consistent with most previous theoretical and experimental predictions. That is, we find support for the equality of WTP distributions among those believing the survey is at least minimally consequential, while those believing the survey will have no effect on policy have statistically different distributions associated with WTP. Our methodology also make use of a new Bayesian posterior simulator for fitting the nonlinear treatment-response model, which we hope will appeal to other researchers seeking to estimate models with a similar structure.

References

- [1] Albert, James H., Siddhartha Chib (1993). “Bayesian Analysis of Binary and Polychotomous Response Data,” *Journal of American Statistical Association*, 88, 669-679.
- [2] Albert, J. and S. Chib (1998). “Sequential Ordinal Modeling with Applications to Survival Data” Technical Report, Olin School of Business, Washington University.
- [3] Arnold, B.C. and Beaver, R.J. (2003). “An alternative construction of skewed multivariate distributions” Technical report 270, Department of Statistics, University of California-Riverside.
- [4] Arnold, B.C. and Beaver, R.J. (2004). “Elliptical Models Subject to Hidden Truncation or Selective Sampling” in M.C. Genton, ed., *Skew-Elliptical Distributions and Their Applications: A Journey Beyond Normality*, Chapter 6, 101-112.

- [5] Bulte, Erwin, Shelby Gerking, John A. List, and Aart de Zeeuw (2005). "The Effect of Varying the Causes of Environmental Problems on Stated WTP Values: Evidence from a Field Study," *Journal of Environmental Economics and Management*, 49, 330-342.
- [6] Carson, Richard T., Theodore Groves, and Mark J. Machina (2000). "Incentive and Informational Properties of Preference Questions," Paper presented at the Kobe Conference on Theory and Application of Environmental Valuation. Kobe: Kobe University, January.
- [7] Carson, Richard T., Theodore Groves, John A. List, and Mark J. Machina (2004). "Probabilistic Influence and Supplemental Benefits: A Field Test of the Two Key Assumptions Underlying Stated Preferences," Working Paper, University of California, San Diego.
- [8] Cowles, M. (1996). "Accelerating Monte Carlo Markov Chain Convergence for Cumulative -Link Generalized Linear Models " *Statistics and Computing* 6, 101-111.
- [9] Cummings, R. G., S. Elliott, G. Harrison, and J. Murphy (1997). "Are hypothetical referenda incentive compatible?" *Journal of Political Economy*, 105, 609-621.
- [10] Cummings, Ronald G., and Laura Osborne Taylor (1998). "Does Realism Matter in Contingent Valuation Surveys?" *Land Economics*, 74 (2), 203-215.
- [11] Damien, P., J. Wakefield and S. Walker (1999). "Gibbs sampling for Bayesian nonconjugate and hierarchical models using auxiliary variables" *Journal of the Royal Statistical Society Series B* 61, 331-344.
- [12] Deb, P., M. Munkin and P. Trivedi (2006). "Bayesian analysis of the two-part model with endogeneity: Application to health care expenditure" 21 (7), 1081-1099.
- [13] Gelman, A., J.B. Carlin, H. Stern and D. Rubin (2004). *Bayesian Data Analysis*, second ed., Chapman & Hall, London.
- [14] Haab, Timothy C., and Kenneth E. McConnell (2002). *Valuing Environmental and Natural Resources: The Econometrics of Non-Market Valuation*. Northampton, MA:Edward Elgar.
- [15] Koop, G, D.J. Poirier and J.L. Tobias (2007). *Bayesian Econometric Methods* Cambridge: Cambridge University Press.

- [16] Lancaster, T. (2004). *An Introduction to Modern Bayesian Econometrics* Blackwell, Oxford.
- [17] Li, K. (1998). “Bayesian Inference in a simultaneous equation model with limited dependent variables.” *Journal of Econometrics* 85, 387-400.
- [18] Nandram, B. and M-H Chen (1996). Reparameterizing the Generalized Linear Model to Accelerate Gibbs Sampler Convergence. *Journal of Statistical Computation and Simulation* 54, 129-144.
- [19] Tanner, M. and M. Wong (1987). The calculation of posterior distributions by data augmentation. *Journal of the American Statistical Association* 62, 810-818.

7 Appendix: The Gibbs Sampling Algorithm

As is well-known in the literature, the standard Gibbs sampler in ordered outcome analyses can suffer from slow mixing. We propose to mitigate this slow mixing problem by sampling the cutpoints $\tilde{\alpha}$ and latent variables $\tilde{\mathbf{c}}$ and $\tilde{\mathbf{w}}$ in a single step. We will do this via the method of composition, by drawing, in order, from the posterior conditionals $p(\tilde{\alpha}|\tilde{\beta}, \tilde{\Sigma}, \mathbf{c}, \mathbf{w})$, $p(\tilde{\mathbf{c}}^*|\tilde{\alpha}, \tilde{\beta}, \tilde{\Sigma}, \mathbf{c}, \mathbf{w})$ and $p(\mathbf{w}^*|\tilde{\mathbf{c}}^*, \tilde{\alpha}, \tilde{\beta}, \tilde{\Sigma}, \mathbf{c}, \mathbf{w})$.

The parameter vector is partitioned into three blocks, and a strategy for drawing from each of these blocks is described below. The first block consists of three separate pieces, crafted to draw directly from $p(\mathbf{w}^*, \tilde{\mathbf{c}}^*, \tilde{\alpha}|\tilde{\beta}, \tilde{\Sigma}, \mathbf{c}, \mathbf{w})$. Finally, in what follows, we let Γ denote all parameters in the model and Γ_{-x} represent all parameters other than x .

Step 1A: Drawing from $p(\tilde{\alpha}|\tilde{\beta}, \tilde{\Sigma}, \mathbf{c}, \mathbf{w})$.

First, note that

$$p(\tilde{\alpha}, \tilde{\mathbf{c}}^*, \tilde{\mathbf{w}}^*|\tilde{\beta}, \tilde{\Sigma}, \mathbf{c}, \mathbf{w}) \propto p(\tilde{\alpha}) \prod_{i=1}^n p(\tilde{c}_i^*, w_i^*|\tilde{\beta}, \tilde{\Sigma}) p(c_i|\tilde{c}_i^*, \tilde{\alpha}) p(w_i|w_i^*). \quad (15)$$

Thus, we obtain

$$p(\tilde{\alpha}|\tilde{\beta}, \tilde{\Sigma}, \mathbf{c}, \mathbf{w}) \propto p(\tilde{\alpha}) \prod_{i=1}^n \int_{\underline{w}_i}^{\tilde{\alpha}_{c_i+1}} \int_{\underline{w}_i}^{\bar{w}_i} p(\tilde{c}_i^*, w_i^*|\tilde{\beta}, \tilde{\Sigma}) dw_i^* d\tilde{c}_i^*.$$

In the above, we have defined

$$\underline{w}_i = \begin{cases} B_i & \text{if } w_i = 1 \\ -\infty & \text{if } w_i = 0 \end{cases} \quad \text{and} \quad \bar{w}_i = \begin{cases} \infty & \text{if } w_i = 1 \\ B_i & \text{if } w_i = 0 \end{cases}.$$

Since the cutpoints $\tilde{\alpha}_3$ and $\tilde{\alpha}_4$ are only involved in the above expression for i such that $c_i \in \{2, 3, 4\}$, we can write:

$$\begin{aligned} p(\tilde{\alpha}|\tilde{\beta}, \tilde{\Sigma}, \mathbf{c}, \mathbf{w}) &\propto p(\tilde{\alpha}) \prod_{i:c_i \in \{2,3,4\}} \Pr(\tilde{\alpha}_{c_i} < \tilde{c}_i^* \leq \tilde{\alpha}_{c_i+1}, \underline{w}_i < w_i^* \leq \bar{w}_i|\tilde{\beta}, \tilde{\Sigma}) \\ &= p(\tilde{\alpha}) \prod_{i:c_i \in \{2,3,4\}} \Pr\left(\frac{\tilde{\alpha}_{c_i} - \mathbf{x}_{ci}\tilde{\beta}_{\mathbf{c}}}{\gamma} < \frac{\tilde{c}_i^* - \mathbf{x}_{ci}\tilde{\beta}_{\mathbf{c}}}{\gamma} \leq \frac{\tilde{\alpha}_{c_i+1} - \mathbf{x}_{ci}\tilde{\beta}_{\mathbf{c}}}{\gamma}, \right. \\ &\quad \left. \frac{\underline{w}_i - \mathbf{x}_{wi}\beta_{\mathbf{w}} - \bar{\mathbf{c}}_i\delta}{\sigma_u} < \frac{w_i^* - \mathbf{x}_{wi}\beta_{\mathbf{w}} - \bar{\mathbf{c}}_i\delta}{\sigma_u} \leq \frac{\bar{w}_i - \mathbf{x}_{wi}\beta_{\mathbf{w}} - \bar{\mathbf{c}}_i\delta}{\sigma_u}\right). \end{aligned}$$

Though routines for directly calculating joint probabilities like those above are often not available in standard software packages, files for calculating the (standardized) bivariate normal cdf are often available. To make use of such routines to calculate the above, we first let $\rho_{cw} = \tilde{\Sigma}(1, 2) / \sqrt{\tilde{\Sigma}(1, 1)} \sqrt{\tilde{\Sigma}(2, 2)} = \sigma_{cw} / \sigma_w$ and define the (standardized) bivariate cdf notation:

$$\Phi(a, b; \rho) = \Pr(z_1 \leq a, z_2 \leq b; \rho)$$

where z_1 and z_2 are univariate normal random variables with zero mean, unit variances and correlation ρ .

Since

$$\begin{aligned} \Pr(a_l < z_1 < a_u, b_l < z_2 < b_u; \rho) &= \Pr(z_1 < a_u, z_2 < b_u; \rho) - \Pr(z_1 < a_l, z_2 < b_u; \rho) \\ &\quad - \Pr(z_1 < a_u, z_2 < b_l; \rho) + \Pr(z_1 < a_l, z_2 < b_l; \rho) \\ &= \Phi(a_u, b_u; \rho) - \Phi(a_l, b_u; \rho) - \Phi(a_u, b_l; \rho) + \Phi(a_l, b_l; \rho), \end{aligned}$$

we can write

$$p(\tilde{\boldsymbol{\alpha}} | \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}, \mathbf{c}, \mathbf{w}) \propto p(\tilde{\boldsymbol{\alpha}}) \prod_{i: c_i \in \{2, 3, 4\}} \left[\Phi(\bar{A}_i, \bar{D}_i; \rho_{cw}) - \Phi(\bar{A}_i, \underline{D}_i; \rho_{cw}) - \Phi(\underline{A}_i, \bar{D}_i; \rho_{cw}) + \Phi(\underline{A}_i, \underline{D}_i; \rho_{cw}) \right] \quad (16)$$

where

$$\underline{A}_i = \frac{\tilde{\alpha}_{c_i} - \mathbf{x}_{ci} \tilde{\boldsymbol{\beta}}_c}{\gamma}, \quad \bar{A}_i = \frac{\tilde{\alpha}_{c_{i+1}} - \mathbf{x}_{ci} \tilde{\boldsymbol{\beta}}_c}{\gamma}, \quad \underline{D}_i = \frac{w_i - \mathbf{x}_{wi} \boldsymbol{\beta}_w - \bar{\mathbf{c}}_i \boldsymbol{\delta}}{\sigma_u} \quad \text{and} \quad \bar{D}_i = \frac{\bar{w}_i - \mathbf{x}_{wi} \boldsymbol{\beta}_w - \bar{\mathbf{c}}_i \boldsymbol{\delta}}{\sigma_u}.$$

The conditional posterior above is not of a recognizable form, and thus we employ a Metropolis-Hastings step for sampling from (15). Specifically, building upon the idea of Cowles (1996), we make use of a random-walk type chain which incorporates the ordering and truncation restrictions on the elements of $\tilde{\boldsymbol{\alpha}}$.

To this end, we let $\tilde{\boldsymbol{\alpha}}^{(t-1)}$ denote the current value of the chain. Implementation of the M-H step requires the specification of a *proposal density* or *jumping distribution*, which we denote as $q(\tilde{\boldsymbol{\alpha}} | \tilde{\boldsymbol{\alpha}}^{(t-1)})$. This proposal density governs the likelihood of movement to $\tilde{\boldsymbol{\alpha}} = \tilde{\boldsymbol{\alpha}}^{(t)}$ given

that the chain is currently at $\tilde{\boldsymbol{\alpha}}^{(t-1)}$. For this application, we choose

$$\begin{aligned}
q(\tilde{\boldsymbol{\alpha}}|\tilde{\boldsymbol{\alpha}}^{(t-1)}) &= q(\tilde{\alpha}_3|\tilde{\boldsymbol{\alpha}}^{(t-1)})q(\tilde{\alpha}_4|\tilde{\alpha}_3, \tilde{\boldsymbol{\alpha}}^{(t-1)}) \\
&= \left[\frac{\phi(\tilde{\alpha}_3; \tilde{\alpha}_3^{(t-1)}, d_1^2)}{\Phi\left(\frac{\tilde{\alpha}_4^{(t-1)} - \tilde{\alpha}_3^{(t-1)}}{d_1}\right) - \Phi\left(\frac{-\tilde{\alpha}_3^{(t-1)}}{d_1}\right)} \right] I\left[0 < \tilde{\alpha}_3 < \tilde{\alpha}_4^{(t-1)}\right] \\
&\quad \times \left[\frac{\phi(\tilde{\alpha}_4; \tilde{\alpha}_4^{(t-1)}, d_2^2)}{\Phi\left(\frac{1 - \tilde{\alpha}_4^{(t-1)}}{d_2}\right) - \Phi\left(\frac{\tilde{\alpha}_3 - \tilde{\alpha}_4^{(t-1)}}{d_2}\right)} \right] I\left[\tilde{\alpha}_3 < \tilde{\alpha}_4 < 1\right].
\end{aligned}$$

In other words, we decompose our transition kernel into a marginal for $\tilde{\alpha}_3$ and a conditional for $\tilde{\alpha}_4$ given $\tilde{\alpha}_3$ (where both are conditioned on the current value of the chain, $\tilde{\boldsymbol{\alpha}}^{(t-1)}$). For the first of these, $\tilde{\alpha}_3$ is sampled from a normal distribution with mean equal to the chain's current value (denoted $\tilde{\alpha}_3^{(t-1)}$) and variance d_1^2 which is truncated to the interval $(0, \tilde{\alpha}_4^{(t-1)})$. For the second step, we sample $\tilde{\alpha}_4$ from a normal density with mean $\tilde{\alpha}_4^{(t-1)}$ and variance d_2^2 which is truncated to the interval $(\tilde{\alpha}_3, 1)$ (where $\tilde{\alpha}_3$ was obtained from the first step).

With this choice of proposal density, we sample $\tilde{\boldsymbol{\alpha}}^{cand}$ from $q(\tilde{\boldsymbol{\alpha}}|\tilde{\boldsymbol{\alpha}}^{(t-1)})$ and accept $\tilde{\boldsymbol{\alpha}}^{cand}$ as a draw from the conditional with probability:

$$\min \left\{ \frac{p(\tilde{\boldsymbol{\alpha}}^{cand}|\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}, \mathbf{c}, \mathbf{w})}{p(\tilde{\boldsymbol{\alpha}}^{t-1}|\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}, \mathbf{c}, \mathbf{w})} \frac{q(\tilde{\boldsymbol{\alpha}}^{t-1}|\tilde{\boldsymbol{\alpha}}^{cand})}{q(\tilde{\boldsymbol{\alpha}}^{cand}|\tilde{\boldsymbol{\alpha}}^{(t-1)})}, 1 \right\},$$

where the target and proposal density ordinates are obtained from (15) and the form of the transition kernel described above. If the candidate draw is not accepted, we set the current value of the chain $\tilde{\boldsymbol{\alpha}}$ equal to its previous value, i.e., $\tilde{\boldsymbol{\alpha}} = \tilde{\boldsymbol{\alpha}}^{t-1}$. In practice, we set $d_1 = d_2 = .01$, which produced reasonable acceptance rates (near 40 percent) and also seemed to perform well in generated data experiments.

Step 1B: Drawing from $p(\tilde{\mathbf{c}}^*|\tilde{\boldsymbol{\alpha}}, \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}, \mathbf{c}, \mathbf{w})$

To implement full blocking procedure, where $\tilde{\boldsymbol{\alpha}}, \tilde{\mathbf{c}}^*$ and \mathbf{w}^* are sampled together in a single step, we would need to draw from the posterior conditional:

$$p(\tilde{c}_i^*|\tilde{\boldsymbol{\alpha}}, \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}, \mathbf{c}, \mathbf{w}), \quad i = 1, 2, \dots, n,$$

where \mathbf{w}^* has been integrated out of the model.

In this regard, note that, from (5),

$$p(\tilde{c}_i^* | \tilde{\boldsymbol{\alpha}}, \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}, \mathbf{c}, \mathbf{w}) \propto p(c_i | \tilde{c}_i^*, \tilde{\boldsymbol{\alpha}}) \int_{\underline{w}_i}^{\bar{w}_i} p(\tilde{c}_i^*, w_i^* | \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}) dw_i^*.$$

We perform this integration by breaking the joint distribution of $p(\tilde{c}_i^*, w_i^* | \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}})$ into the conditional $p(w_i^* | \tilde{c}_i^*, \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}})$ times the marginal $p(\tilde{c}_i^* | \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}})$.

The latter of these is simply:

$$\tilde{c}_i^* | \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}} \sim N(\mathbf{x}_{\mathbf{ci}} \tilde{\boldsymbol{\beta}}_{\mathbf{c}}, \tilde{\sigma}_{\mathbf{c}}^2)$$

and the former is:

$$w_i^* | \tilde{c}_i^*, \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}} \sim N\left(\mathbf{x}_{\mathbf{wi}} \boldsymbol{\beta}_{\mathbf{w}} + \bar{\mathbf{c}}_{\mathbf{i}} \boldsymbol{\delta} + [\tilde{\sigma}_{cw} / \tilde{\sigma}_{\mathbf{c}}^2](\tilde{c}_i^* - \mathbf{x}_{\mathbf{ci}} \tilde{\boldsymbol{\beta}}_{\mathbf{c}}), \sigma_w^2 (1 - \tilde{\rho}_{cw}^2)\right).$$

So, when performing the necessary integration, we obtain, when $w_i = 1$:

$$p(\tilde{c}_i^* | \tilde{\boldsymbol{\alpha}}, \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}, \mathbf{c}, \mathbf{w}) \propto \phi(\tilde{c}_i^*; \mathbf{x}_{\mathbf{ci}} \tilde{\boldsymbol{\beta}}_{\mathbf{c}}, \tilde{\sigma}_{\mathbf{c}}^2) \left[1 - \Phi\left(\frac{B_i - \mathbf{x}_{\mathbf{wi}} \boldsymbol{\beta}_{\mathbf{w}} - \bar{\mathbf{c}}_{\mathbf{i}} \boldsymbol{\delta} - [\tilde{\sigma}_{cw} / \tilde{\sigma}_{\mathbf{c}}^2](\tilde{c}_i^* - \mathbf{x}_{\mathbf{ci}} \tilde{\boldsymbol{\beta}}_{\mathbf{c}})}{\sigma_w \sqrt{1 - \tilde{\rho}_{cw}^2}}\right) \right] \times I(\tilde{\alpha}_{c_i} < \tilde{c}_i^* \leq \tilde{\alpha}_{c_i+1}). \quad (17)$$

Similarly, when $w_i = 0$, we get

$$p(\tilde{c}_i^* | \tilde{\boldsymbol{\alpha}}, \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}, \mathbf{c}, \mathbf{w}) \propto \phi(\tilde{c}_i^*; \mathbf{x}_{\mathbf{ci}} \tilde{\boldsymbol{\beta}}_{\mathbf{c}}, \tilde{\sigma}_{\mathbf{c}}^2) \left[\Phi\left(\frac{B_i - \mathbf{x}_{\mathbf{wi}} \boldsymbol{\beta}_{\mathbf{w}} - \bar{\mathbf{c}}_{\mathbf{i}} \boldsymbol{\delta} - [\tilde{\sigma}_{cw} / \tilde{\sigma}_{\mathbf{c}}^2](\tilde{c}_i^* - \mathbf{x}_{\mathbf{ci}} \tilde{\boldsymbol{\beta}}_{\mathbf{c}})}{\sigma_w \sqrt{1 - \tilde{\rho}_{cw}^2}}\right) \right] \times I(\tilde{\alpha}_{c_i} < \tilde{c}_i^* \leq \tilde{\alpha}_{c_i+1}). \quad (18)$$

This density is, again, not of a standard form. To make this blocking step feasible and preferable relative to a standard Gibbs algorithm, we require a fast method for generating draws from (16) and (17), since this must be done for each observation in the sample. Here we make use of the auxiliary variable Gibbs sampling technique of Damien et al. (1999) as a quick and efficient way of generating these draws. To illustrate its use, we first consider the case where $w = 1$ so that our goal is to sample from (16). To this end, we introduce a set of auxiliary variables u_i , $i = 1, 2, \dots, n$ where

$$u_i | c_i^*, \tilde{\boldsymbol{\alpha}}, \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}, \mathbf{c}, \mathbf{w}) \stackrel{ind}{\sim} U\left(0, \left[1 - \Phi\left(\frac{B_i - \mathbf{x}_{\mathbf{wi}} \boldsymbol{\beta}_{\mathbf{w}} - \bar{\mathbf{c}}_{\mathbf{i}} \boldsymbol{\delta} - [\tilde{\sigma}_{cw} / \tilde{\sigma}_{\mathbf{c}}^2](\tilde{c}_i^* - \mathbf{x}_{\mathbf{ci}} \tilde{\boldsymbol{\beta}}_{\mathbf{c}})}{\sigma_w \sqrt{1 - \tilde{\rho}_{cw}^2}}\right) \right] \right), \quad (19)$$

a series of independent, uniform random variables over the given interval. It follows that the joint distribution of u_i and c_i^* is

$$p(u_i, c_i^* | \tilde{\boldsymbol{\alpha}}, \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}, \mathbf{c}, \mathbf{w}) \propto \phi(\tilde{c}_i^*; \mathbf{x}_{\mathbf{ci}} \tilde{\boldsymbol{\beta}}_{\mathbf{c}}, \tilde{\sigma}_c^2) I(\tilde{\alpha}_{c_i} < \tilde{c}_i^* \leq \tilde{\alpha}_{c_{i+1}}) \\ \times I\left(0 < u_i < \left[1 - \Phi\left(\frac{B_i - \mathbf{x}_{\mathbf{wi}} \boldsymbol{\beta}_{\mathbf{w}} - \bar{\mathbf{c}}_i \boldsymbol{\delta} - [\tilde{\sigma}_{cw}/\tilde{\sigma}_c^2](\tilde{c}_i^* - \mathbf{x}_{\mathbf{ci}} \tilde{\boldsymbol{\beta}}_{\mathbf{c}})}{\sigma_w \sqrt{1 - \tilde{\rho}_{cw}^2}}\right)\right]\right).$$

At this point, a simple Gibbs sampler can be constructed, which draws from $u_i | \tilde{c}_i^*, \cdot$ and $\tilde{c}_i^* | u_i, \cdot$. At the end of this subroutine, the last \tilde{c}_i^* simulation is retained and kept as a draw from the desired conditional distribution. The first of these conditionals is easy to draw from, and is given by (18). The second involves working out the conditional for \tilde{c}_i^* . To this end, let

$$h_{1i}^* \equiv \mathbf{x}_{\mathbf{ci}} \boldsymbol{\beta}_{\mathbf{c}} + \frac{\tilde{\sigma}_c}{\tilde{\rho}} \left[\frac{B_i - \mathbf{x}_{\mathbf{wi}} \boldsymbol{\beta}_{\mathbf{w}} - \bar{\mathbf{c}}_i \boldsymbol{\delta}}{\sigma_w} + \Phi^{-1}(u_i) \sqrt{1 - \tilde{\rho}^2} \right].$$

Then,

$$\tilde{c}_i^* | u_i, \cdot, w_i = 1 \sim \begin{cases} TN_{(\max\{\tilde{\alpha}_{c_i}, h_{1i}^*\}, \tilde{\alpha}_{c_{i+1}})}(\mathbf{x}_{\mathbf{ci}} \boldsymbol{\beta}_{\mathbf{c}}, \tilde{\sigma}_c^2) & \text{if } \rho > 0 \\ TN_{(\tilde{\alpha}_{c_i}, \min\{\tilde{\alpha}_{c_{i+1}}, h_{1i}^*\})}(\mathbf{x}_{\mathbf{ci}} \boldsymbol{\beta}_{\mathbf{c}}, \tilde{\sigma}_c^2) & \text{if } \rho < 0 \end{cases}.$$

Note, of course, that when $\rho = 0$, the above is undefined. However, unless $\rho = 0$ is set as the initial condition of the chain, a simulated correlation identically equal to zero will never occur and moreover, when $\rho = 0$ the desired conditionals in (16) and (17) reduce to truncated normal distributions, and auxiliary variable methods are not required.

In a similar manner, when $w_i = 0$, we first define

$$h_{0i}^* \equiv \mathbf{x}_{\mathbf{ci}} \boldsymbol{\beta}_{\mathbf{c}} + \frac{\tilde{\sigma}_c}{\tilde{\rho}} \left[\frac{B_i - \mathbf{x}_{\mathbf{wi}} \boldsymbol{\beta}_{\mathbf{w}} - \bar{\mathbf{c}}_i \boldsymbol{\delta}}{\sigma_w} - \Phi^{-1}(u_i) \sqrt{1 - \tilde{\rho}^2} \right].$$

It follows that

$$\tilde{c}_i^* | u_i, \cdot, w_i = 0 \sim \begin{cases} TN_{(\max\{\tilde{\alpha}_{c_i}, h_{1i}^*\}, \tilde{\alpha}_{c_{i+1}})}(\mathbf{x}_{\mathbf{ci}} \boldsymbol{\beta}_{\mathbf{c}}, \tilde{\sigma}_c^2) & \text{if } \rho < 0 \\ TN_{(\tilde{\alpha}_{c_i}, \min\{\tilde{\alpha}_{c_{i+1}}, h_{1i}^*\})}(\mathbf{x}_{\mathbf{ci}} \boldsymbol{\beta}_{\mathbf{c}}, \tilde{\sigma}_c^2) & \text{if } \rho > 0 \end{cases}.$$

In practice, when we get to Step 1B of our algorithm, we run this subroutine for 20 iterations to obtain the desired $\tilde{\mathbf{c}}^*$ draws. Our initial step in this process samples $c_i^* \stackrel{ind}{\sim} TN_{(\tilde{\alpha}_{c_i}, \tilde{\alpha}_{c_{i+1}})}(\mathbf{x}_{\mathbf{ci}} \boldsymbol{\beta}_{\mathbf{c}}, \tilde{\sigma}_c^2)$, guaranteeing that the regions of truncation defined above are well-defined. We evaluated the performance of this method in numerous generated data experiments, and found that it converged within only a few iterations, with 20 being adequate for a burn-in. Moreover, the code for this step can be completely vectorized, and the required distributions are easily sampled.

Step 1C: Drawing from $p(\mathbf{w}^* | \tilde{\mathbf{c}}^*, \tilde{\boldsymbol{\alpha}}, \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}}, \mathbf{c}, \mathbf{w})$.

With draws obtained from Steps 1A and 1B, we can then draw each w_i^* from its complete posterior conditional. That is, for $i = 1, 2, \dots, n$,

$$w_i^* | \Gamma_{-w_i^*}, \mathbf{c}, \mathbf{w} \sim \begin{cases} TN_{[B_i, \infty)} \left(\mathbf{x}_{wi} \boldsymbol{\beta}_w + \bar{\mathbf{c}}_i \boldsymbol{\delta}_c + \frac{\tilde{\sigma}_{\epsilon u}}{\tilde{\sigma}_\epsilon^2} [\tilde{c}_i^* - \mathbf{x}_{ci} \tilde{\boldsymbol{\beta}}_c], \sigma_u^2 (1 - \tilde{\rho}_{\epsilon u}^2) \right) & \text{if } w_1 = 1 \\ TN_{(-\infty, B_i)} \left(\mathbf{x}_{wi} \boldsymbol{\beta}_w + \bar{\mathbf{c}}_i \boldsymbol{\delta}_c + \frac{\tilde{\sigma}_{\epsilon u}}{\tilde{\sigma}_\epsilon^2} [\tilde{c}_i^* - \mathbf{x}_{ci} \tilde{\boldsymbol{\beta}}_c], \sigma_u^2 (1 - \tilde{\rho}_{\epsilon u}^2) \right) & \text{if } w_1 = 0. \end{cases}$$

Step 2: Drawing from $p(\tilde{\boldsymbol{\beta}} | \Gamma_{-\tilde{\boldsymbol{\beta}}}, \mathbf{c}, \mathbf{w})$.

First, define

$$\mathbf{X}_i = \begin{bmatrix} \mathbf{x}_{ci} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{x}_{wi} & \bar{\mathbf{c}}_i \end{bmatrix}.$$

Then, similar to a SUR model [e.g., Koop, Poirier and Tobias (2007, pages 138-138)], we obtain

$$\tilde{\boldsymbol{\beta}} | \Gamma_{-\tilde{\boldsymbol{\beta}}}, \mathbf{c}, \mathbf{w} \sim N(\mathbf{D}_\beta \mathbf{d}_\beta, \mathbf{D}_\beta)$$

where

$$\mathbf{D}_\beta = \left(\sum_i \mathbf{X}_i' \tilde{\boldsymbol{\Sigma}}^{-1} \mathbf{X}_i + \mathbf{V}_\beta^{-1} \right)^{-1} \quad \text{and} \quad \mathbf{d}_\beta = \sum_i \mathbf{X}_i' \tilde{\boldsymbol{\Sigma}}^{-1} \begin{bmatrix} \tilde{c}_i^* \\ w_i^* \end{bmatrix} + \mathbf{V}_\beta^{-1} \boldsymbol{\mu}_\beta.$$

Step 3: Drawing from $p(\tilde{\boldsymbol{\Sigma}}^{-1} | \Gamma_{-\tilde{\boldsymbol{\Sigma}}^{-1}}, \mathbf{c}, \mathbf{w})$.

Finally, again making use of techniques like those employed in the SUR model, we obtain

$$\tilde{\boldsymbol{\Sigma}}^{-1} | \Gamma_{-\tilde{\boldsymbol{\Sigma}}^{-1}}, \mathbf{c}, \mathbf{w} \sim W \left(\left[\sum_i \begin{bmatrix} \tilde{\epsilon}_i \\ u_i \end{bmatrix} [\tilde{\epsilon}_i \ u_i] + \mathbf{R}^{-1} \right]^{-1}, n + \nu \right).$$

A Gibbs sampler proceeds by simulating (in order) from these posterior conditional distributions.