UNR Joint Economics Working Paper Series
Working Paper No. 07-012

Meta-Functional Benefit Transfer for Wetland Valuation: Making the Most of Small Samples

Klaus Moeltner and Richard Woodward

Department of Resource Economics / MS 204
University of Nevada, Reno
Reno, NV 89557
(775) 784-6701 | Fax (775) 784-1342
email: moeltner@cabnr.unr.edu

December, 2007

Abstract

This study applies functional Benefit Transfer via Meta-Regression Modeling to derive valuation estimates for wetlands in an actual policy setting of proposed groundwater transfers in Eastern Nevada. We illustrate how Bayesian estimation techniques can be used to overcome small sample problems notoriously present in Meta-functional Benefit Transfer. The highlights of our methodology are (i) The hierarchical modeling of heteroskedasticity, (ii) The ability to incorporate additional information via refined priors, and (ii) The derivation of measures of model performance with the corresponding option of model-averaged Benefit Transfer predictions. Our results indicate that economic losses associated with the disappearance of these wetlands can be substantial and that primary valuation studies are warranted.

JEL Classification: C11, C15, Q51

Keywords: Bayesian Model Averaging; t-Error Regression Model; Meta-Analysis; Benefit Transfer; Wetland Valuation

Meta-Functional Benefit Transfer for Wetland Valuation: Making the Most of Small Samples

Klaus Moeltner*
Department of Resource Economics,
University of Nevada, Reno

Richard Woodward
Department of Agricultural Economics
Texas A&M University

Please cite as
"Paper presented at the W1133 Annual Meetings in Richmond, VA, March 28-30, 2007"

Comments Welcome!

* Contact Information:

Klaus Moeltner
Department of Resource Economics / MS 204
University of Nevada, Reno
Reno, NV 89557-0105
phone: (775) 784-4803
fax: (775) 784-1342
e-mail: moeltner@cabnr.unr.edu

Meta-Functional Benefit Transfer for Wetland Valuation: Making the Most of Small Samples

Abstract:

This study applies functional Benefit Transfer via Meta-Regression Modeling to derive valuation estimates for wetlands in an actual policy setting of proposed groundwater transfers in Eastern Nevada. We illustrate how Bayesian estimation techniques can be used to overcome small sample problems notoriously present in Meta-functional Benefit Transfer. The highlights of our methodology are (i) The hierarchical modeling of heteroskedasticity, (ii) The ability to incorporate additional information via refined priors, and (ii) The derivation of measures of model performance with the corresponding option of model-averaged Benefit Transfer predictions. Our results indicate that economic losses associated with the disappearance of these wetlands can be substantial and that primary valuation studies are warranted.

Key words: Bayesian Model Averaging, t-Error Regression Model, Meta-Analysis, Benefit Transfer, Wetland Valuation

JEL codes: C11, C15, Q51

I) Introduction

There is an ongoing trend amongst resource managing agencies to use existing information on past policy outcomes to predict economic benefits for planned policy implementations. For some institutions this concept of “Benefit Transfer” (BT) has essentially become the primary tool of policy assessment. For example, in a recent insiders’ assessment of the role of BT at the U.S. Environmental Protection Agency (EPA) Iovanna and Griffiths (2006) predict that due to the triple constraints of expediency, financial strains, and administrative hurdles “original assessment studies will remain a rare exception” in future EPA valuation efforts.

It is not surprising, therefore, that the concept and methodology of BT has become of central interest to resource economists in the U.S. and abroad. In recent years much effort has been allocated to examine the theoretical underpinnings of BT (e.g. Smith and Pattanayak 2002; Smith et al. 2006) and to facilitate its econometric and computational implementation (e.g. Rosenberger and Loomis 2000a; Léon et al. 2002; Moeltner et al. 2007; Moeltner and Rosenberger 2007). Most of these studies have illustrated their respective methodologies using simulated data or hypothetical policy scenarios. However, to date there are very few contributions to the published literature that show-case the implementation of BT within the context of an actual policy implication.[1]

The first objective of this study is to provide a “real-world” example of BT within the context of wetland valuation. Existing meta-analyses of wetland valuation studies (Brouwer et al. 1999; Woodward and Wui 2001; Borisova-Kidder 2006; Brander et al. 2006), have focused primarily on the marginal effects of wetland functions and attributes on economic benefits. In contrast, this analysis illustrates how BT *was actually* used to inform decision makers on projected economic benefits related to a planned resource intervention affecting wetlands.

The second aim of this study is to demonstrate how Bayesian econometric techniques can overcome many classical estimation challenges when even the most thorough literature review produces only a small number of existing sources suitable for a meta-functional BT model. Specifically, we illustrate how (i) study-specific heteroskedasticity – ubiquitous in meta-regressions – can be addressed with only a single additional model parameter, (ii) ancillary information from source studies and meta-analyses can be used to derive more informed Bayesian priors, and (iii) Bayesian Model Averaging (BMA) can be employed to address model uncertainty in the derivation of BT predictions. We believe that our approach will widen the applicability of BT in many resource valuation settings, and provide resource planners with a tool to proceed with meta-functional BT in many cases where its implementation has to date been hampered by small sample problems.

We introduce the general policy context that triggered our BT implementation in the next Section, Section III describes the construction of the meta-dataset underlying our analysis, Section IV discussed the econometric framework and Section V presents estimation results. We summarize our findings and offer concluding remarks in Section VI.

II) Policy Context and Study Area

The Southern Nevada Water Authority (SNWA) has recently proposed a Groundwater Development Project to transfer approximately 200,000 acre-feet of groundwater per year from seven hydrographic basins in rural Eastern Nevada to the wider Las Vegas Valley to assure a reliable water supply for this fast growing urban area in future years. The construction of groundwater conveyance facilities and supporting infrastructure is proposed to start as early as 2009 (SNWA 2006).

The granting of associated water rights rests with the Nevada State Engineer ‘s Office. Over the last two years, this agency has collected scientific and economic evidence on the potential implications of this water transfer for the targeted provider areas to aid in this decision-making process. The due date for all evidence to be submitted was November 1, 2006. Approximately three weeks prior to this deadline the authors were contacted to examine if there might be any “non-market” – type economic values that could be at stake should the Project be approved.

Given the tight time frame and available scientific evidence we decided to focus on two distinct and unique wetland areas, the Swamp Cedar Natural Area and the Shoshone Ponds Natural Area. Both wetlands are located in Spring Valley in east-central NV. They are predicted to be significantly and almost immediately impacted by the planned groundwater withdrawals (Lanner 2006).

Spring Valley is approximately 9.5 miles wide (east – west) and 95 miles long (north-south). It distinguishes itself from other valleys in the Great Basin by its high elevation (5500 – 6000 feet), and its relatively abundant water resources, provided by over 100 natural springs (Charlet 2006). These springs together with snowmelt retained by a hardpan soil layer support numerous wetlands throughout the Valley (Lanner 2006).

The Swamp Cedar Natural Area (SCNA), a marshy ecosystem with natural ponds and meadows contains 3200 acres of public land administered by the Bureau of Land Management (BLM). It supports a large stand of Rocky Mountain Junipers (Juniperus scopulorum), commonly referred to as "Swamp Cedars". These Spring Valley Cedars have been described as "globally unique" as they have adapted to a distinctly different environment than is characteristic for the main population of their species (Charlet 2006; Lanner 2006). The SCNA offers recreational opportunities for hiking, primitive camping, and wildlife viewing, although it does not feature a designated access road, parking area, developed trail system or established campgrounds (BLM 1980a).

The Shoshone Pond Natural Area (SPNA) is located approximately 13 miles south of the SCNA in Southern Spring Valley. It contains 1240 acres of public land managed by the BLM. It features two important natural resources: (i) A second stand of "Swamp Cedars" of the same ecotypical variety as those found in the SCNA, and (ii) Three manmade, spring-fed pools that harbor two rare species of fish, the relict dace (Relictus solitarius) and the Pahrump poolfish (Empetrichthys latos). The relict dace is listed by the Nevada Natural Heritage database as "imperiled and vulnerable in Nevada and globally", while the Pahrump poolfish, for which the Shoshone ponds constitute one of only three remaining habitats, has been federally listed as an endangered species since 1969. While lacking maintained hiking trails or established campsites, the SPNA offers recreational opportunities for hiking, primitive camping, and wildlife viewing (BLM 1980b).

Both time frame and available funds were insufficient to launch a primary valuation study, which left BT as only viable alternative to produce at least approximate estimates of potential economic losses. Furthermore, given that some information on basic attributes was available for these wetlands we aimed for BT via function transfer (e.g. Kirchhoff et al. 1997; Brouwer and Spaninks 1999). In addition, since there does not exist a single valuation study that corresponds sufficiently well to the current context we decided to implement this functional BT via a Meta-Regression Model (MRM) that draws information from several underlying source studies (e.g. Rosenberger and Loomis 2000b; Shrestha and Loomis 2001; Johnston et al. 2005).

III) Data Set Construction

Suitable primary studies for the MRM were identified using the following sources: Four existing meta-analyses focusing on the economic value of wetlands (Brouwer et al. 1999; Woodward and Wui 2001; Brander et al. 2006; Borisova-Kidder 2006), the Environmental Valuation Reference Inventory (EVRI), a searchable database focusing on non-market valuation, and ECONLIT, a general searchable database for economic literature. The initial criteria for study selection were: (i) Geographic area = USA or Canada, (ii) Exclusion of coastal or marine types of wetlands, (iii) Estimated economic values must include values related to habitat, biodiversity, or species preservation. The latter two criteria flow from the nature of the current policy context: Spring Valley wetlands are distinctly different ecosystems than coastal or marine wetlands, and their economic value is primarily related to habitat and biodiversity services. Thus, we excluded studies that focused on wetlands with the *sole functions* of flood control or water quality improvements, as well as studies that *only* examined the value of specific wetlands with respect to extractive use (hunting, fishing).

This "first cut" approach produced a set of 24 initial candidate studies. Given the nature of their primary valuation objectives (habitat and biodiversity services, recreational opportunities) all of these sources use survey-based approaches to elicit households' willingness-to-pay (WTP) to preserve or expand a specific existing wetland area. A second round of screening eliminated studies that are based on identical wetlands and target populations, and studies based on surveys that produced response rates below 30 percent. In the case of duplicate studies we retained the study with the most reliable research methodology. The low-response rate criterion was applied to guard against "selection bias", i.e. the possibility that the small segment of those who participated in the survey is not representative of the underlying target population. Only one study fell into that latter category.

These selection refinements resulted in a final set of nine studies deemed suitable for the research context at hand, yielding 12 observations available for our meta-dataset (One study, Blomquist and Whitehead 1998, reports WTP estimates for four different wetlands). While this sample is not as large as would be ideal it has several desirable properties. As shown in the Table 1 the selected studies provide good coverage of the geographic target area, with applications from various parts of the United States, and one Canadian contribution (Tkac 2002). All studies were conducted within the last 15 years and thus use modern survey and estimation methodologies. The underlying target populations are of a general nature with at least regional scope. Specifically, three studies (Loomis et al. 1991; Roberts and Leitch 1997; Tkac 2002) focus on a regional population of stakeholders, while five of the studies are associated with a State-wide target population (Hanemann et al. 1991; Whitehead and Blomquist 1991; Mullarkey 1997; Blomquist and Whitehead 1998; Poor 1999) and one source (Klocek 2004) has nation-wide coverage.

The sample also exhibits a desirable mix of journal publications, book chapters, government reports, and theses or dissertations. The relatively strong representation by contributions from the "gray" literature eases the traditional concern of "publication bias" in meta-analytical research, i.e. the notion that only valuation results that are surprising or otherwise noteworthy are ever considered by journal editors.

Table 2 provides more detailed information for each observation included in our meta-dataset. Most policy scenarios presented to respondents for a given study stipulated that wetland areas would be lost (due to agricultural activities, mining, or urban sprawl) if no action was taken. With two exceptions (Roberts and Leitch 1997, who use a payment brackets-approach, and Tkac 2002 who uses a payment table with uncertainty scales á la Welsh and Poe 1998), all of the studies employed a variant of the dichotomous choice elicitation format. In most cases respondents were then asked if they would be willing to pay a specific dollar amount ("bid") into a nature conservation fund or in additional taxes to preserve these lands. The only exceptions to this “loss avoidance” approach are the studies by Poor (1999) and Klocek (2004) who asked respondents if they would be willing to contribute to a special fund to *create* additional acres of wetland (for example by converting drained agricultural areas to their original marshy conditions). The relative homogeneity of elicitation approaches is fortuitous for our meta-application since the small sample size would preempt the inclusion of study-methodological regressors, which have been found to carry considerable importance in MRMs (e.g. Johnston et al. 2005, Moeltner et al. 2007) . Given that most of our source studies use a similar elicitation approach we can argue that the effect of methodological aspects should be minor for our MRM.

All of the included studies asked respondents to value the entire bundle of wetland services, including habitat and biodiversity provision, flood control, water filtration, and opportunities for non-consumptive (wildlife viewing, hiking, photography) and consumptive (hunting, fishing) recreational activities. Some studies (Blomquist and Whitehead 1998, Tkac 2002) also stress the presence of threatened or endangered species on the wetlands under consideration. Since the surveys targeted the general population of underlying households (as opposed to a specific group of active users), only a relatively small segment of respondents indicated that they had visited the wetland under consideration in the past, as depicted in the "percentage of active users" column in Table 2. Thus, the lion's share of estimated economic benefits (i.e. reported WTP) is likely associated with non-use or existence values. This is another important and desirable feature of our data set given the current research context, since it can be expected that only a small proportion of the wider population of stakeholders will have actually visited the Spring Valley wetlands considered in this study.

As evident from the table the types of wetland, the policy scenarios in terms of wetland acres preserved or created, and the percentage of active users in the underlying sample vary widely over studies. Not surprisingly, so does aggregate WTP per household. The smallest welfare estimate (in 2006 currency) is less than a dollar per U.S.-wide household to preserve parts of the Canaan Valley National

Wildlife Refuge in West Virginia, while households in the San Joaquin Valley in California are willing to pay close to $300 per year to prevent the loss of a large share of the Valley's wetland habitats. In consequence, using any single study for a point transfer of benefits would be risky. Furthermore, using the sample mean of over $60 would likely lead to grossly inflated BT predictions given the relatively obscure and isolated nature of our policy site. In combination, these facts support the approach of functional BT via meat-regression.

IV) Econometric Framework

Classical Challenges

Virtually all existing MRMs on resource valuation have been estimated via classical least squares methods. However, in our case the small number of observations makes it difficult to take a classical estimation approach. For example, any specification test relying on asymptotic theory will be unreliable in this case. This includes all tests on heteroskedasticity, a likely occurrence in meta-regressions given that each source observation flows from a different original regression model. On the other hand, simply ignoring heteroskedasticity would cast doubt on the reliability of standard errors for estimated coefficients and associated confidence intervals for BT predictions. Traditionally, researchers have addressed heteroskedasticty problems with robust, or White-corrected, standard errors (e.g. Woodward and Wui 2001; Brander et al. 2006). However, the White correction itself rests on asymptotic theory and is thus of limited value in our application. Similarly, unless the analyst firmly believes that the basic Classical Linear Regression Model applies, there are no specification tests available for guidance on the composition or functional form of explanatory variables.

Related problems arise when the MRM is estimated with logged WTP as dependent variable, as has been the case in virtually all existing meta-analyses related to resource valuation to assure non-negativity of welfare measures, but absolute values in dollars are required for BT predictions. The 'Delta method' (e.g. Greene 2003, Ch.5) or equivalent asymptotic techniques such as bootstrap or the popular

Krinsky and Robb (1986) routine must be applied to obtain confidence intervals for the converted estimates. Again, such approaches are unreliable in our case.

We thus propose a Bayesian estimation framework for our MRM as it poses several advantages over a classical approach in this context. First, error heteroskedasticity can be modeled hierarchically with only a single additional parameter. Second, relevant information form source studies that is not captured in the actual meta-data can enter the model via the specification of prior distributions, which leads to more representative and, possibly, more efficient BT estimates. Third, a Bayesian framework allows for the incorporation of *model uncertainty* by estimating multiple candidate models and their corresponding probability weights, and then deriving a weighted model-averaged distribution for the benefit construct of interest. This circumvents the need to choose, with little guidance, a single preferred specification to generate the transfer function, as would be required in a classical framework.

Bayesian Model

Our kernel MRM takes the form of a standard linear regression model with the added feature of a hierarchical distribution for the variance of the regression error to allow for observation-specific heteroskedasticity. Specifically, we stipulate individual variance weights to be drawn from an inverse-gamma distribution with equal shape and scale. Our baseline MRM can thus be written as

$$y_j = \mathbf{x}_\mathbf{j}'\boldsymbol{\beta} + \varepsilon_j \quad \text{with} \quad \varepsilon_j \sim n\left(0, \sigma^2\omega_j\right), \quad \text{and} \quad \omega_j \sim ig\left(\tfrac{v}{2}, \tfrac{v}{2}\right), \tag{1}$$

where y_j is WTP reported in study j, $\mathbf{x}_\mathbf{j}$ is a vector of population and wetland characteristics associated with study j, $\boldsymbol{\beta}$ is a corresponding vector of regression coefficients, ε_j is a zero-mean regression error with variance $\sigma^2\omega_j$, and *ig* denotes the inverse-gamma distribution.[2]

As shown in Geweke (1993) the hierarchical specification of the variance of ε_j is exactly equivalent to drawing ε_j from a *t*-distribution with mean zero, scale σ^2 and v degrees of freedom. In addition to capturing variance-inequalities across observations, this allows for higher probabilities of large error variances than would be expected for a basic normal model, a likely occurrence in a meta-regression context. To be specific, for any given σ^2 a small value of v (say 5 to 10) implies a heavy-tailed distribution, while, as is well known, the *t*-distribution approaches normality for larger values of *v*.

Conditional on the individual-specific variance weights the likelihood function for our MRM follows a multivariate normal distribution with generalized variance-covariance matrix, i.e.

$$pr\left(\mathbf{y} \mid \mathbf{X}, \boldsymbol{\beta}, \sigma^2, \boldsymbol{\omega}\right) = \left(2\pi\right)^{-n/2} \left|\left(\sigma^2 \boldsymbol{\Omega}\right)\right|^{-1/2} \exp\left(-\tfrac{1}{2}\left(\mathbf{y}-\mathbf{X}\boldsymbol{\beta}\right)' \left(\sigma^2\boldsymbol{\Omega}\right)^{-1}\left(\mathbf{y}-\mathbf{X}\boldsymbol{\beta}\right)\right)$$
$$\text{with} \quad \mathbf{X} = \begin{bmatrix}\mathbf{x}_1' & \mathbf{x}_2' & \cdots & \mathbf{x}_n'\end{bmatrix}', \quad \boldsymbol{\omega} = \begin{bmatrix}\omega_1 & \omega_2 & \cdots & \omega_n\end{bmatrix}, \text{ and } \quad \boldsymbol{\Omega} = diag\begin{bmatrix}\omega_1 & \omega_2 & \cdots & \omega_n\end{bmatrix} \qquad (2)$$

where n denotes the total number of observations. Conditioning the likelihood on $\boldsymbol{\omega}$ instead of v facilitates posterior simulation. In essence, we treat the variance weights as additional data. This is deemed 'data augmentation' in Bayesian analysis (Tanner and Wong 1987).

The specification of the Bayesian model is completed by assigning prior distributions to all model parameters. We follow standard approaches by choosing a multivariate normal distribution with mean $\boldsymbol{\mu}_0$ and variance-covariance matrix $\mathbf{V}_0$ for the vector of regression coefficients $\boldsymbol{\beta}$, and an inverse-gamma distribution with shape η_0 and scale κ_0 for the shared variance component σ^2. In addition, we specify the heteroskedasticity parameter v to follow a gamma distribution with shape 1 and inverse scale $1/v_0$. As discussed in Koop (2004), Ch. 6, this choice of hyper-prior distribution for v is computationally convenient and assures the required condition of $v > 0$. Thus, the hierarchical prior structure for our MRM can be compactly denoted as

$$\begin{aligned} &pr\left(\boldsymbol{\beta}\right) = mvn\left(\boldsymbol{\mu}_0, \mathbf{V}_0\right) \\ &pr\left(\sigma^2\right) = ig\left(\eta_0, \kappa_0\right) \\ &pr\left(\omega_j \mid v\right) = ig\left(\tfrac{v}{2}, \tfrac{v}{2}\right), \forall j \qquad p\left(v\right) = g\left(1, 1/v_0\right) \end{aligned} \qquad (3)$$

The Bayesian framework then combines likelihood function and priors to derive marginal posterior distributions for all parameters. We use a Gibbs Sampler (GS) with a built-in Metropolis-Hastings (Hastings 1970) routine for draws of v along the lines suggested in Koop 2004, Ch. 6, to simulate these distributions. The details of this algorithm are given in Appendix A.

Model weights and BT predictions

As described in more detail in the next section we estimate a set of M different MRMs, distinguished by the specification of explanatory variables and chosen parameters for prior distributions. For each model M_m, $m = 1 \ldots M$, the GS generates $r=1 \ldots R$ draws of parameters $\boldsymbol{\beta}$, σ^2, and v. For notational convenience we combine these parameters into a joint vector $\boldsymbol{\theta}$ and denote individual draws of this vector as $\boldsymbol{\theta}_{\mathbf{r}}$. For each model, our ultimate construct of interest is the posterior predictive distribution (PPD) of WTP for the policy context, conditional *only* on policy site regressors and the general model specification, i.e. $pr\left(y_p \mid \mathbf{x_p}, M_m\right)$, where subscript p indicates "policy context". The derivation of this density proceeds in two steps: First, for each round of the original GS we obtain a draw of y_p conditional on a specific set of parameters, ie. a draw from $pr\left(y_p \mid \mathbf{x_p}, \boldsymbol{\theta}_{\mathbf{r}}, M_m\right)$, as $y_{p,r} = \mathbf{x}'_{\mathbf{p}}\boldsymbol{\beta}_{\mathbf{r}} + \varepsilon_{p,r}$ where the error term $\varepsilon_{p,r}$ is drawn from $t\left(0, \sigma_r^2, v_r\right)$. Second, we repeat this process S times to obtain multiple draws of y_p for each set of original parameter draws.[3] The resulting $\left(R \cdot S\right)$ draws can thus be interpreted as flowing from $pr\left(y_p \mid \mathbf{x_p}, M_m\right)$, the desired simulated PPD of y_p for model M_m.

For each model the posterior simulator also produces the *model-conditioned marginal likelihood* i.e. $pr\left(\mathbf{y} \mid M_m\right)$, which can be used to compute the *posterior probability* for a given model, denoted as $pr\left(M_m \mid \mathbf{y}\right)$. Loosely speaking this probability indicates how likely the observed data (i.e. the dependent observations in our MRM) were generated by model M_m. Formally, the two concepts are linked through Bayes' Rule as

$$pr\left(M_m \mid \mathbf{y}\right) = \frac{pr\left(\mathbf{y} \mid M_m\right) pr\left(M_m\right)}{p\left(\mathbf{y}\right)} \tag{4}$$

where $pr\left(M_m\right)$ indicates the *prior model probability*, and $p\left(\mathbf{y}\right)$ denotes the unconditional marginal likelihood, i.e. the probability that **y** was generated by *any* of the considered models. Assuming that the

M models considered for our application constitute an exhaustive representation of all possible models (i.e. $\sum_{m=1}^{M} pr\left(M_m \mid \mathbf{y}\right)=1$), and setting $pr\left(M_m\right)=1/M, \forall m$, we can write the posterior model probability as the ratio of a specific model's marginal likelihood to the sum of all marginal likelihoods for the considered model space, i.e. as

$$pr\left(M_m \mid \mathbf{y}\right)=\frac{pr\left(\mathbf{y} \mid M_m\right)}{\sum_{m=1}^{M} pr\left(\mathbf{y} \mid M_m\right)}, \tag{5}$$

(e.g. Koop 2004, Ch. 1). Since an analytical expression for $pr\left(\mathbf{y} \mid M_m\right)$ does not exist for our kernel specification we simulate this value for each model using the approach proposed by Chib and Jeliazkov 2001.

Model-specific PPDs of the BT construct and probability weights can then be combined to produce a model-averaged predictive distribution of y_p. Analytically, this density can be expressed as

$$pr\left(y_p \mid \mathbf{x_p}\right)=\sum_{m=1}^{M}\left\{\int_{\boldsymbol{\theta}} pr\left(y_p \mid \mathbf{x_p}, M_m, \boldsymbol{\theta}\right) pr\left(\boldsymbol{\theta} \mid \mathbf{y}, \mathbf{X}, M_m\right) d\left(\boldsymbol{\theta}\right)\right\} pr\left(M_m \mid \mathbf{y}\right). \tag{6}$$

Equation (6) indicates that the posterior predictive distribution of y_p conditional only on policy descriptors $\mathbf{x_p}$ is derived by marginalizing conditional draws of y_p over (i) model parameters, and (ii) all considered models. The first marginalization is accomplished via the two-step approach outlined above. Marginalizing over model space simply involves weight-averaging model-specific draws from the PPD of y_p across all models.[4]

V) Empirical Implementation

Model space and prior refinements

Table 3 captures the model space, i.e. the exhaustive list of all MRMs considered in our analysis. We decided on this set based on plausible and available explanatory variables, available extraneous information for prior distributions, and preliminary estimation runs.[5]

The models in Table 3 differ in three dimensions: (i) The distribution of the error term, (ii) The specification of prior distributions for $\boldsymbol{\beta}$, and (iii) The functional specification of the explanatory variable "wetland acreage". As indicated in the table, Models 1- 4 have a t-distributed error as given in equation (1), while the error term in Models 5-8 follows a generic homoskedastic normal distribution. All models indexed by "a" or "b" specify a prior distribution for the constant term with a mean of zero and a variance of 10. Based on preliminary OLS results, we also allow for specifications with a large negative mean (-50) and a highly diffuse variance for the constant term (all models indexed by "c" in the table). Prior distributions for the slope coefficients contained in $\boldsymbol{\beta}$ are non-informative for "a"-type models, and refined for "b" and "c"-type MRMs. Wetland acreage enters the models in one of four possible forms: logged and logged-squared (Models 1 and 5), logged only (Models 2 and 6), linear in units of 1000 plus linear-squared (Models 3 and 7), and linear only (Models 4 and 8).

The remaining regressors include annual household income, in log-form, and the percentage of active wetland users corresponding to a given source study. Information for the refinement of priors was available for income and wetland acreage. This added information flows primarily from coefficient estimates reported in source studies or other meta-analyses. The detailed process of prior refinement is described in Appendix B. As shown in the appendix, our refined priors have considerably smaller variances than their diffuse, or non-informative, counterparts. Thus, they carry substantial weight in the posterior simulation routine and allow the added information to have a measurable effect on posterior results.

Estimation results

All models are estimated using 15000 burn-in draws and 10000 retained draws in the Gibbs Sampler. The decision on the appropriate amount of burn-ins was guided by Geweke's (1992) convergence diagnostic (CD). For each MRM, the standard deviation of the proposal density for *v* in the Metropolis Hastings algorithm contained in the GS (denoted as s_v in Appendix A) is set to achieve an optimal acceptance rate of 44-50% (see e.g Gelman et al. 2004 Ch. 11).

Table 4 presents estimation results for each model in terms of its logged marginal likelihood (denoted as *log pr(y|M)*), mean absolute percent error (MAPE), and model weight (indicated as *pr(M|y)* in the table)[6]. As discussed in Kass and Raftery (1995) marginal likelihood values are primarily useful to compare two competing models (the ratio of marginal likelihoods is often referred to as "*Bayes Factor*") or to generate model weights for each specification in a set of ex ante chosen candidates, which is the main focus in our application. The MAPE, in turn, is a universal measure of model fit that focuses primarily on the in-sample predictive ability of a given specification. For a general discussion of this and other measures of model fit in a Bayesian framework see Gelman et al. (2004), Ch. 6. Moeltner et al. (2007) compare different measures of model fit in an empirical setting.

According to Kass and Raftery (1995), a difference in marginal likelihoods of 5 or more indicates a decisive superiority of the model with the smaller value (in absolute terms). Accordingly, as can be seen from Table 4, our models with homoskedastic-normal errors score considerably better using this criterion than their t-error counterparts. This likely indicates that our data set is too small to provide substantial evidence of error heteroskedasticity. Applying equation (5), higher marginal likelihood values also translate into higher posterior model weights for the normal specifications as indicate in the last column of the table. In fact, none of the t-error models receives any appreciable posterior weight in our application.

Within each group of error distributions models with refined priors and zero-mean prior for the constant term (i.e. "b"-type models) receive slightly more favorable marginal likelihood scores than their respective non-informative versions ("a"-type models) or versions with a large negative intercept prior ("c"-type models). This effect is especially pronounced for Models 4b vs. 4a and Models 8b vs. 8a, with likelihood differences in the 4-5 point range. Not surprisingly, the refined priors generally lead to a deterioration in predictive ability with respect to the actual data as indicated by the higher MAPE scores for "b" and "c" type models compared to their "a" type counterparts. As mentioned above, our refined priors with their smaller variances absorb considerable posterior weight, while the diffuse priors for "a" – type models allocate most posterior weight to the underlying data. In general, this leads to better in-sample predictive ability, i.e. lower (= better) MAPE scores. The exceptions to this pattern can be again observed for Models 4b and 8b and, to a lesser extent, Models 4c and 8c, which receive a better MAPE score than their corresponding non-informative specifications. It thus appears that the linear formulation of wetland acreage *without* its squared form produces the best model fit for our refined MRMs. Overall, Model 8b receives by far the highest posterior model probability (0.902) and would thus be an ideal candidate to generate BT predictions if a single model had to be chosen for this task.

However, we pursue our original strategy of allowing every model to contribute to the posterior distribution of BT estimates by computing the weighted average of model-specific results. We generate predicted benefits for three possible groups of stakeholders: (i) All households in the four counties surrounding Spring Valley (11,118 units by the 2000 Census), (ii) All Nevada households (751,165 units), and (iii) All households in Nevada and Utah (1,452,446 units). We set income figures to the most recently reported Census medians, and the percentage of active users arbitrarily to 5% for the Four County-Region and 1% for wider Nevada and Utah. Naturally, a small poll of households from the target population could produce more accurate estimates of active use. Our chosen figures are at the lower end of those found in our source studies and thus appear reasonably conservative for the task at hand.

We follow the steps outlined in Section IV to generate predictive distributions. For each of the $R = 10{,}000$ parameter draws from the original GS, we draw a set of $S = 100$ predicted values for policy outcome y_p. We then keep every 20^{th} of these draws to reduce autocorrelation in our sequence. Thus, we retain 50,000 posterior predictive draws for our analysis.[7]

Table 5 captures posterior predictive distributions of benefits for the Four Counties ("Region"), Nevada ("NV") and Nevada plus Utah ("NV/UT") produced by each single MRM. Model-averaged results are given toward the bottom of the table. For each model the table reports the posterior mean and its numerical standard error (nse), a measure of simulation noise.[8] As illustrated e.g. in Koop (2004), Ch. 3, a numerical 95% confidence interval for the posterior mean can be obtained as [posterior mean $\pm 1.96 \cdot nse$]. As can be seen from the Table posterior means range from under $3/ year (Models 3a, 7a) to over $30 (Model 1c). Thus, in absence of any additional information on relative model performance it would be risky to base BT predictions on a single MRM. This clearly illustrates the benefits of guidance through posterior model weights or the convenience of model-averaging.

Using our model-averaged results we thus predict annual losses associated with the disappearance of the SCNA and SPNA wetlands of $4.8 to $5.6 per household. The numerical confidence intervals for these posterior means are in the $0.1 - $0.2 range, indicating that simulation error is a minor consideration for our application. These per-household figures translate into predictions of $62,000 for the Region, $3.6 million for Nevada, and close to $7 million for Nevada and Utah as indicated in the bottom row of the table.[9]

VI) Conclusion

This study describes an actual application of meta-functional BT to value wetland areas in Eastern Nevada. We illustrate how Bayesian estimation techniques can be used to produce reasonable BT results even with a very small underlying sample of source studies. The main advantages of our methodology compared to classical regression models are (i) The ability to capture heteroskedasticity in straightforward fashion through hierarchical modeling of the error variance, (ii) The ability to incorporate additional information not captured in the data via refined priors, and (ii) The availability of measures of model performance with the corresponding option of generating model-averaged BT predictions.

While the hierarchical error variance approach turned out to be of limited importance in our application, the refinement of prior distributions led to clearly superior posterior results for several of our specifications. In addition, guidance through marginal likelihood values and associated posterior model weights proved critical in identifying promising MRMs and in properly weighting individual models in the generation of model-averaged predictive distributions.

Naturally, we also ought to stress the limitations of our approach compared to a primary valuation study for the policy area. Our small sample size and the lack of detailed information on specific attributes of wetland areas considered in original studies preempts a more thorough examination of the effect of various wetland features (other than acreage) on WTP. Each of the wetlands in our meta-data is unique in some sense, and wetland size in acres alone is not necessarily a reliable proxy for wetland quality attributes. For example, it is quite possible that the Spring Valley wetlands are valued more highly than predicted in our analysis given their function as habitats for a globally unique stand of trees, and two threatened / endangered fish species. On the other hand, many of the included wetlands in our meta-regression offer richer recreational opportunities than the Spring Valley areas. This, in turn, could inflate our BT estimates.

Perhaps the most meaningful way to interpret our secondary-data results is to use them as a strong indication that the economic losses associated with a potential disappearance of Spring Valley wetlands could be of substantial magnitude, and that therefore primary economic research is both warranted and justified. Given the large geographic scale of the proposed groundwater extraction project, and the potentially irreversible nature of its environmental implications, it is imperative that decision makers be informed of all economic benefits and costs involved. These considerations should also include non-market type values associated with affected natural areas. We hope that our preliminary results via BT will aid in creating awareness that such values exist and that they can be of important magnitude.

Appendix A: Posterior Simulation

This Appendix outlines the detailed steps of the Gibbs Sampler (GS) for the regression model with t-distributed errors. It is convenient to apply Tanner and Wong 1987's concept of data augmentation and treat draws of $\boldsymbol{\omega} = \begin{bmatrix} \omega_1 & \omega_2 & \cdots & \omega_n \end{bmatrix}$ as additional data in the likelihood function. This leads to the augmented joint posterior $pr\left(\boldsymbol{\beta}, \sigma^2, v, \boldsymbol{\omega} \mid \mathbf{y}, \mathbf{X}\right)$, which the GS breaks down into consecutive draws of conditional components.

Step 1: Draw $\boldsymbol{\beta}$

Given our multivariate-normal choice of prior for $\boldsymbol{\beta}$ the conditional posterior for this vector can be derived in straightforward fashion (e.g. Lindley and Smith 1972) as:

$$pr\left(\boldsymbol{\beta} \mid \mathbf{y}, \mathbf{X}, \sigma^2, \boldsymbol{\omega}\right) = mvn\left(\boldsymbol{\mu}_1, \mathbf{V}_1\right) \qquad \text{where}$$

$$\mathbf{V}_1 = \left(\mathbf{V}_0^{-1} + \mathbf{X}'\left(\sigma^2\boldsymbol{\Omega}\right)^{-1}\mathbf{X}\right)^{-1} \text{ and } \quad \boldsymbol{\mu}_1 = \mathbf{V}_1\left(\mathbf{V}_0^{-1}\boldsymbol{\mu}_0 + \mathbf{X}'\left(\sigma^2\boldsymbol{\Omega}\right)^{-1}\mathbf{y}\right).$$

Step 2: Draw σ^2

Applying again standard results for generalized regression models, we obtain

$$pr\left(\sigma^2 \mid \mathbf{y}, \mathbf{X}, \boldsymbol{\beta}, \boldsymbol{\omega}\right) = ig\left(\eta_1, \kappa_1\right) \quad \text{with} \quad \eta_1 = \left(n + 2\eta_0\right)/2 \quad \text{and} \quad \kappa_1 = \frac{1}{2}\left(\left(\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\right)'\boldsymbol{\Omega}^{-1}\left(\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\right) + 2\kappa_0\right).$$

Step 3: Draw v

The relevant kernel for draws of v is its prior times the density of $\boldsymbol{\omega}$, i.e. $pr\left(v \mid \boldsymbol{\omega}\right) = \frac{1}{v_0}\exp\left(-\frac{v}{v_0}\right) \cdot \prod_{j=1}^{n} \frac{\left(\frac{v}{2}\right)^{\frac{v}{2}}}{\Gamma\left(\frac{v}{2}\right)} \omega_j^{-\left(\frac{v}{2}+1\right)} \exp\left(-\frac{v}{2\omega_j}\right)$. This is a non-standard density, and we use a random walk Metropolis-Hastings algorithm (MH, Hastings 1970 Chib and Greenberg 1995) to take draws from this kernel. Specifically, we draw a candidate value of v_c in the r[th] round of the GS from a truncated-at-zero normal proposal density with mean v_{r-1}, i.e. the current value of *v*, and standard deviation s_v, and accept the draw as the new current value with probability $\alpha_v = \min\left(\frac{pr\left(v_c \mid \boldsymbol{\omega}\right)}{pr\left(v_{r-1} \mid \boldsymbol{\omega}\right)}, 1\right)$.

The standard deviation of s_v is chosen (after some trial and error in preliminary runs) to yield an acceptance probability in the 45-50% range, as suggested by Gelman et al. 2004. Ch. 11.

<u>Step 4:</u> Draw $\boldsymbol{\omega}$

For this step we note that $\frac{\varepsilon_j}{\sigma} \sim n\left(0, \omega_j\right)$. We can then use again standard results for the Gaussian regression model to obtain $pr\left(\omega_j \mid y_j, \mathbf{x_j}, \boldsymbol{\beta}, \sigma^2, v\right) = ig\left(\psi, \zeta\right)$ with $\psi = \left(v+1\right)/2$ and $\zeta = \frac{1}{2}\left(\left(y_j - \mathbf{x'_j}\boldsymbol{\beta}\right)^2 / \sigma^2 + v\right)$.

Appendix B: Specification of Refined Priors:

This appendix describes the detailed derivation of refined priors for slope coefficients in models indexed with "*b*" or "*c*" in Table 3. For ease of exposition we label the relevant components of $\boldsymbol{\beta}$ as follows:

β_{inc} slope coefficient for log(income)
β_{lna} slope coefficient for log(acres), where applicable
β_{lna2} slope coefficient for $[\log(\text{acres})]^2$, where applicable
β_{a} slope coefficient for (acres, in 000), where applicable
β_{a2} slope coefficient for $[(\text{acres, in } 000)]^2$,where applicable

Refinements for β_{inc}:

Given the log-form for the dependent variable in our MRMs β_{inc} can be interpreted as income elasticity with respect to WTP. Three of the source studies underlying our meta-dataset include a statistically significant income variable in their model. Blomquist and Whitehead 1998 regress log(WTP) against income (in \$000) and estimate a coefficient of 0.03. However, for our specification we need a prior distribution for $\frac{\partial \log(wtp)}{\partial \log(inc)}$. This requires a conversion of Blomquist and Whitehead 1998's result. Noting that in their model $wtp = \exp\left(\gamma_{inc} \cdot (inc/1000)\right)\exp(\text{"rest"})$ we can write $\frac{\partial wtp}{\partial inc} = wtp \cdot (\gamma_{inc}/1000)$ and

$\frac{\partial wtp}{\partial inc}\frac{inc}{wtp} \approx \frac{\partial \log(wtp)}{\partial \log(inc)} = (\gamma_{inc}/1000) \cdot inc$. Using their sample mean of income, converted to 2006 dollars, we derive an approximated point estimate for income elasticity of 1.146. The second source study that relates WTP to income is Poor 1999. This study also estimates a model of log(wtp) on log(income), so we can directly adopt the reported coefficient of 0.12 as a second prior point estimate for β_{inc}.

The third study is Klocek 2004. The author estimates a linear bid function model relating WTP to income (in dollars). Dividing by the coefficient on "bid" we derive a scale-corrected income coefficient of 0.00022, i.e. $\frac{\partial wtp}{\partial inc} = 0.00022$. Using again sample means for income and WTP we convert this figure to an approximate point estimate for income elasticity of $\frac{\partial wtp}{\partial inc}\frac{inc}{wtp} = 1.23$.

We then treat these three estimates as draws form a normal prior distribution with mean μ_0 and variance V_0. Further, we impose the constraint of very small probability mass (say <0.01) for negative values given the vast evidence from the empirical valuation literature of non-decreasing WTP over income. We thus employ a constrained maximum likelihood routine to estimate the most likely values for μ_0 and variance V_0 given the three data points. We then use the resulting estimates of $\hat{\mu}_0 = 0.9413$ and $\hat{V}_0 = 0.4046^2$ as prior mean and variance for β_{inc} in our refined models.[10]

Refinements for $\beta_{\ln a}$

Given their focus on single wetland applications none of our source studies explicitly include a variable corresponding to wetland size in their regression model. However, most of the existing meta-analyses on wetland valuation include this regressor. For example, Woodward and Wui 2001 relate log(wtp/acre) to log(acres) with an estimated coefficient of -0.286 (model C). We can write their model as $\log\left(\frac{wtp}{acres}\right) = \gamma_{ac}\log(acres) + \text{"rest"} \xrightarrow{\text{ergo}} \log(\text{wtp}) = (1+\gamma_{ac})\log(acres) + \text{"rest"}$. It follows that $\frac{\partial \log(wtp)}{\partial \log(acres)} = (1+\gamma_{inc}) = 0.714$, which serves as our first point estimate for $\beta_{\ln a}$. Borisova-Kidder 2006 regresses log(wtp) against acres, with an estimated coefficient of 0.000000965. Following the arguments for the Blomquist and Whitehead 1998 study above we derive an expression for "acreage elasticity" as $\frac{\partial wtp}{\partial acres}\frac{acres}{wtp} = \gamma_{ac} \cdot acres$. Using the sample mean of 270,758 acres this yields a second point estimate for $\beta_{\ln a}$ *of* 0.26. Brander et al. 2006 regress log(wtp/hectares) against log(hectares) and estimate a coefficient of -0.11. Using the logic applied to Woodward and Wui 2001 this implies a point estimate of 0.89 for log(wtp) against log(hectares). This marginal effect is the same for "acres" since the conversion factor would be absorbed in the intercept. Feeding these three point estimates into our constrained ML routine described above yields estimates for the mean and variance of the refined prior distribution for $\beta_{\ln a}$ of 0.6213 and 0.2654^2, respectively.[11]

Refinements for $\beta_{\ln a2}$

Using a change-in-variable approach it is straightforward to show that if $\frac{\partial \log(y)}{\partial \log(x)} = a$,

$\frac{\partial \log(y)}{\partial (\log(x))^2} = \frac{1}{2} a \cdot (\log(x))^{-1}$. Thus, using a mean acreage of 963,466 for Woodward and Wui 2001 we obtain an approximate estimate of $\frac{\partial \log(wtp)}{\partial (\log(acres))^2} = \frac{1}{2} 0.714 \cdot (\log(963{,}466))^{-1} = 0.026$. The analogous values for Borisova-Kidder 2006 and Brander et al. 2006 are 0.01 (study mean = 270,758 acres) and 0.054 (study mean = 4049 acres), respectively. Unconstrained ML applied to these three point estimates produces estimated prior moments of $\hat{\mu}_0 = 0.03$ and $\hat{V}_0 = 0.0065^2$.

Refinements for β_a

We use the same three meta-studies to derive prior distributions for this coefficient. For Woodward and Wui 2001 we need to convert $\frac{\partial \log(wtp)}{\partial \log(acres)}$ into $\frac{\partial \log(wtp)}{\partial (acres/1000)}$. Using again a change-in-variable approach it can be shown that if $\frac{\partial \log(wtp)}{\partial \log(acres)} = a$, $\frac{\partial \log(wtp)}{\partial (acres/1000)} = 1000a/acres$.

Using *a* = 0.714 and the sample mean of 963,466 acres we obtain a first point estimate for β_a of 0.00074. For the Borisova-Kidder 2006 study we simply need to multiply the reported coefficient by 1000 to obtain our second point estimate of 0.000965. For Brander et al. 2006 we proceed as for Woodward and Wui 2001 to derive a third estimate of 0.022. Imposing again the constraint of a positive marginal effect of acreage on WTP our ML routine produces estimates of 0.0151 and 0.0065^2 for the prior mean and variance of β_a.

Coefficient for β_{a2}

Given our results for β_a mean estimates for β_{a2} would be arbitrarily close to zero. We thus retain a non-informative prior mean of 0 for this coefficient but reduce the prior variance to 0.1 in our refined specifications.

Notes

[1] We are only aware of one such study, where BT is employed to decide on critical habitat designations (Loomis 2006).

[2] Since only one study furnishes multiple observations on WTP in this application we abstract from the modeling of panel-structures and treat each observation as flowing from a different source.

[3] While this *S*-fold replication is optional it is computationally inexpensive and improves the efficiency of the predictive distribution.

[4] For a comprehensive discussion of Bayesian Model Averaging (BMA) see e.g. Raftery (1995), Hoeting et al. (1999), and Chipman et al. (2001).

[5] Naturally, one could widen this model space with additional variants of our chosen specifications. However any additional feasible model is unlikely to carry considerable posterior weight.

[6] Given our focus on model weights and posterior *predictive* distributions detailed results for posterior distributions of individual parameters for each model have been omitted from this text. They are available from the authors upon request.

[7] To guard against dramatic outliers we further truncate the *exponentiated* distribution of our logged predictions at the 99th percentile, i.e. we discard the 500 largest observations. This final adjustment is implemented in identical fashion for all models. Intuitively, this correction could be interpreted as "imposing income constraints" on the predicted WTP values.

[8] The *nse* is computed as $std/\sqrt{(R_p)}$ where *std* is the standard deviation of the predicted distribution and R_p is the number of simulated draws for the predicted series.

[9] One could argue that the SCNA and SPNA ought to be valued separately. We decided to pool the two areas for valuation purposes since this strategy best corresponds to the bulk of scenarios underlying our meta-data. In most of these studies, respondents were asked to value groups, bundles, or large areas of

non-contiguous wetlands. A separate valuation of the Spring Valley areas using our MRM would likely lead to an over-estimation of combined economic benefits. Naturally, it would be straightforward to design a primary valuation study that elicits separate benefit figures for the two areas.

[10] One might alternatively consider using reported standard errors in source studies to derive prior variances. However, standard errors (i.e. empirical variability of estimated coefficients) in a classical framework simply indicate noise resulting from sampling error, i.e. the notion of extrapolating from a finite sample to an underlying population. This concept is completely absent in Bayesian methodology, where stipulated prior distributions already correspond to the underlying population of interest.

[11] Intuitively, WTP should be non-decreasing in wetland size. However, in our case the non-negativity constraint emerged as non-binding in the ML routine.

<u>References</u>

Bureau of Land Management (1980a). Swamp Cedar Instant Study Area, Ely Field Office Report.

Bureau of Land Management (1980b). Shoshone Pond Instant Study Area, Ely Field Office Report.

Blomquist, G. C. and J. C. Whitehead (1998), 'Resource Quality Information and Validity of Willingness to Pay in Contingent Valuation', *Resource and Energy Economics* **20** 179-196.

Borisova-Kidder, A. (2006). Meta-Analytical Estimates of Values of Environmental Services Enhanced by Government Agricultural Conservation Programs. Dissertation, Department of Agricultural, Environmental, and Development Economics, Ohio State University.

Brander, L. M., R. J. G. Florax and J. E. Vermaat (2006), 'The Empirics of Wetland Valuation: A Comprehensive Summary and a Meta-Analysis of the Literature', *Environmental and Resource Economics* **33** 223-250.

Brouwer, R., I. H. Langford, I. J. Bateman and R. K. Turner (1999), 'A Meta-Analysis of Wetland Contingent Valuation Studies', *Regional Environmental Change* **1**(1), 47-57.

Brouwer, R. and F. Spaninks (1999), 'The Validity of Environmental Benefits Transfer: Further Empirical Testing', *Environmental and Resource Economics* **14** 95-117.

Charlet, D. (2006). Effect of Interbasin Water Transport on Ecosystems of Spring Valley, White Pine County, NV, Report submitted to the Office of the State Engineer of the State of Nevada, June 2006 (Exhibit # 3030)

Chib, S. and E. Greenberg (1995), 'Understanding the Metropolis-Hastings Algorithm', *American Statistician* **49** 327-335.

Chib, S. and I. Jeliazkov (2001), 'Marginal Likelihood Form the Metropolis-Hastings Output', *Journal of the American Statistical Association* **95**(453), 270-281.

Chipman, H., E. I. George and R. E. McCulloch (2001). The Practical Implementation of Bayesian Model Selection. IMS Lecture Notes - Monograph Series**:** 67-116.

Gelman, A., J. B. Carlin, H. S. Stern and D. B. Rubin (2004), *Bayesian Data Analysis*. Boca Raton / London / New York / Washington, D.C.: Chapman&Hall/CRC.

Geweke, J. (1992). Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments, in Bayesian Statistics 4. J. M. Bernardo, J. O. Berger, A. P. Dawid and A. F. M. Smith (Eds.). Oxford, UK, Oxford University Press.

Geweke, J. (1993), 'Treatment of the Independent Student-T Linear Model', *Journal of Applied Econometrics* **8** S19-S40.

Greene, W. H. (2003), *Econometric Analysis*: Prentice Hall, 5th Edition.

Hanemann, M., J. Loomis and B. Kanninen (1991), 'Statistical Efficiency of Double-Bounded Dichotomous Choice Contingent Valuation', *American Journal of Agricultural Economics* **73**(4), 1255-1263.

Hastings, W. K. (1970), 'Monte Carlo Sampling Methods Using Markov Chains and Their Applications', *Biometrika* **57** 97-109.

Hoeting, J. A., D. Madigan, A. E. Raftery and C. T. Volinsky (1999), 'Bayesian Model Averaging: A Tutorial', *Statistical Science* **14**(4), 382-417.

Iovanna, R. and C. Griffiths (2006), 'Clean Water, Ecological Benefits, and Benefits Transfer: A Work in Progress at the U.S. EPA', *Ecological Economics* **60**(2), 473-482.

Johnston, R. J., E. Y. Besedin, R. Iovanna, C. J. Miller, R. F. Wardwell and M. H. Ranson (2005), 'Systematic Variation in Willingness to Pay for Aquatic Resource Improvements and Implications for Benefit Transfer: A Meta-Analysis', *Canadian Journal of Agricultural Economics* **53**(2-3), 221-248.

Kass, R. E. and A. E. Raftery (1995), 'Bayes Factors', *Journal of the American Statistical Association* **90**(430), 773-795.

Kirchhoff, S., B. G. Colby and J. T. LaFrance (1997), 'Evaluating the Performance of Benefit Transfer: An Empirical Inquiry', *Journal of Environmental Economics and Management* **33**(1), 75-93.

Klocek, C. A. (2004). Estimating the Economic Value of Canaan Valley National Wildlife Refuge: A Contingent Valuation Approach. Dissertation. Davis College of Agricultural Forestry and Consumer Sciences. Morgantown, WV, West Virginia University.

Koop, G. (2004), *Bayesian Econometrics*: John Wiley & Sons, Ltd.

Krinsky, I. and A. L. Robb (1986), 'On Approximating the Statistical Properties of Elasticities', *Review of Economics and Statistics* **68**(4), 715-719.

Lanner, R. M. (2006). The Effect of Groundwater Pumping Proposed by the Southern Nevada Water Authority on The "Swamp Cedars" (*Juniperus Scopulorum*) of Spring Valley, Nevada.

Léon, C. J., F. J. Vazquez-Polo, N. Guerra and P. Riera (2002), 'A Bayesian Model for Benefits Transfer: Application to National Parks in Spain', *Applied Economics* **34**(6), 749-757.

Lindley, D. V. and A. F. M. Smith (1972), 'Bayes Estimates for the Linear Model', *Journal of the Royal Statistical Society, Series B (Methodological)* **34**(1), 1-41.

Loomis, J. (2006), 'Estimating Recreation and Existence Values of Sea Otter Expansion in California Using Benefit Transfer', *Coastal Management* **34** 387-404.

Loomis, J., M. Hanemann, B. Kanninen and T. Wegge (1991). WTP to Protect Wetlands and Reduce Wildlife Contamination from Agricultural Drainage, in The Economics and Management of Water and Drainage in Agriculture. A. Dinar and D. Zilberman (Eds). Norwell, Massachusetts, Kluwer Academic Publishers**:** 411-429.

Moeltner, K., K. Boyle and R. Paterson (2007), 'Meta-Analysis and Benefit Transfer for Resource Valuation - Addressing Classical Challenges with Bayesian Modeling', *Journal of Environmental Economics and Management* **53** 250 - 269.

Moeltner, K. and R. S. Rosenberger (2007). Meta-Regression and Benefit Transfer: Data Space, Model Space, and the Quest for 'Optimal Scope', Paper presented at the Triangle Resource & Environmental Economics Seminar Series, Raleigh, NC, April 27, 2007.

Mullarkey, D. (1997). Contingent Valuation of Wetlands: Testing Sensitivity to Scope. Dissertation. Department of Agricultural and Applied Economics. Madison, University of Wisconsin.

Poor, P. J. (1999), 'The Value of Additional Central Flyway Wetlands: The Case of Nebraska's Rainwater Basin Wetlands', *Journal of Agricultural and Resource Economics* **24**(1), 253-265.

Raftery, A. E. (1995), 'Bayesian Model Selection in Social Research', *Sociological Methodology* **25** 111-163.

Roberts, L. A. and J. A. Leitch (1997). Economic Valuation of Some Wetland Outputs of Mud Lake, Minnesota-South Dakota, Department of Agricultural Economics, North Dakota Agricultural Experiment Station, North Dakota State University, Fargo, ND 58105.

Rosenberger, R. S. and J. B. Loomis (2000a), 'Panel Stratification in Meta-Analysis of Economic Studies: An Investigation of Its Effects in the Recreation Valuation Literature', *Journal of Agricultural and Applied Economics* **32**(2), 459-470.

Rosenberger, R. S. and J. B. Loomis (2000b), 'Using Meta-Analysis for Benefit Transfer: In-Sample Convergent Validity Tests of an Outdoor Recreation Database', *Water Resources Research* **36**(4), 1097-1107.

Shrestha, R. K. and J. B. Loomis (2001), 'Testing a Meta-Analysis Model for Benefit Transfer in International Outdoor Recreation', *Ecological Economics* **39** 67-83.

Smith, V. K. and S. K. Pattanayak (2002), 'Is Meta-Analysis a Noah's Ark for Non-Market Valuation?' *Environmental and Resource Economics* **22**(1-2), 271-96.

Smith, V. K., S. K. Pattanayak and G. V. Houtven (2006), 'Structural Benefits Transfer: An Example Using VSL Estimation', *Ecological Economics* **60**(2), 361-371.

Southern Nevada Water Authority (SNWA) (2006). Clark, Lincoln and White Pine Counties Groundwater Development Project. http://www.nvgroundwaterproject.com/html/index.html.

Tanner, M. A. and W. H. Wong (1987), 'The Calculation of Posterior Distributions by Data Augmentation (with Discussion)', *Journal of the American Statistical Association* **82** 528-550.

Tkac, J. M. (2002). Estimating Willingness to Pay for the Preservation of the Alfred Bog Wetland in Ontario: A Multiple Bounded Discrete Choice Approach. Master's Thesis, Department of Agricultural Economics. Montreal, Quebec, Macdonald Campus, Gill University.

Welsh, M. P. and G. L. Poe (1998), 'Elicitation Effects in Contingent Valuation: Comparisons to a Multiple Bounded Discrete Choice Approach', *Journal of Environmental Economics and Management* **36**(2), 170-185.

Whitehead, J. C. and G. C. Blomquist (1991), 'Measuring Contingent Values for Wetlands: Effects of Information About Related Environmental Goods', *Water Resources Research* **27**(10), 2523-2531.

Woodward, R. T. and Y.-S. Wui (2001), 'The Economic Value of Wetland Services: A Meta-Analysis', *Ecological Economics* **37** 257-270.

Table 1: Source Studies Used in Meta-Regression

study ID	Authors	Publication type	Publication year	Study year	Study Area	Underlying target population	response rate
1	Loomis et al .	journal article	1991	1989	Wetlands in the San Joaquin Valley, CA	San Joaquin Valley households	35%
2	Hanemann et al.	journal article	1991	1989	Wetlands in the San Joaquin Valley, CA	CA households outside the San Joaquin Valley	51%
3	Whitehead and Blomquist	journal article	1991	1989	Clear Creek wetland area in Western KY	Kentucky households	31%
4	D. Mullarkey	PhD dissertation	1997	1994	110 acres of undesignated wetlands in northwest Wisconsin	Wisconsin households	60%
5	Roberts and Leitch	Government Report	1997	1996	Mud Lake wetland area on SD / MN border	Households within 30 miles of study area	62%
6	Blomquist and Whitehead	journal article	1998	1990	Various wetland habitats in Western KY	Kentucky households	70%
7	J. Poor	Journal article	1999	1996	Rainwater Basin Wetlands, NB	Nebraskan households	46%
8	J. M. Tkac	Master's thesis	2002	2001	Alfred Bog, Ontario, CA	Households in the United Counties of Prescott and Russell, Ontario	57%
9	C. A. Klocek	PhD dissertation	2004	1996	Canaan Valley National Wildlife Refuge	U.S. Households	74%

Table 2: Observation-Specific Details for Meta-Regression

study ID	wetland type	original wetland area (acres)	policy scenario (acres preserved or created)	official resource designation	WTP per HH and year	HH Income	percentage of active users
1	unspecified	85,000	58,000	includes several NWRs and WMAs	284.15	66,776	46%
2	unspecified	85,000	58,000	includes several NWRs and WMAs	248.23	82,061	38%
3	bottomland hardwood forests wetlands	84,000	5,000	none	17.39	52,258	16%
4	unspecified	110	110	none	1.7 (a)	43,880 (m)	1%
5	permanently, semi-permanently, or seasonally flooded lacustrine wetlands	5,000	5,000	none	3.03	38,745 (m)	18%
6a	permanently flooded freshwater marsh	3,968	500	none	2.62	38,207	14,2%
6b	temporarily flooded bottomland hardwoods	70,080	500	none	7.27	38,207	14,2%
6c	seasonally flooded bottomland hardwoods	25,216	500	none	5.7	38,207	14,2%
6d	permanently flooded bottomland hardwood	1,408	500	none	17.37	38,207	14,2%
7	unspecified	34,000	41000 (c)	none	27.18	41,238	52%
8	domed peat bog with boreal forest	10,378	10,378	class 1 Wetland / ANSI	4.66 (a)	46,024 (m)	29%
9	high elevation moist valley	708	23,292	NWR	0.63	64,532	2%
	means:	***33,739***	***14,707***		***61.36***	***51,077***	***25%***

All monetary figures are in 2006 U.S. dollars
(a)= originally elicited as lump sum payment; annualized using a discount rate of 6%
(s) = sample mean as reported in source study, converted to 2006 dollars
(m) = census median (sample income not reported)
HH = household
NWR = National Wildlife Refuge / WMA = Wildlife Management Area / ANSI = Area of Natural and Scientific Interest

Table 3: Model Specifications

model label	error distribution	prior mean and variance for constant term	priors for slope coefficients	log (acres)	log (acres), squared	acres (in units of 1000)	acres (in units of 1000), squared
M1a	t	0, 10	diffuse	x	x		
M1b	t	0, 10	refined	x	x		
M1c	t	-50, 100	refined	x	x		
M2a	t	0, 10	diffuse	x			
M2b	t	0, 10	refined	x			
M2c	t	-50, 100	refined	x			
M3a	t	0, 10	diffuse			x	x
M3b	t	0, 10	refined			x	x
M3c	t	-50, 100	refined			x	x
M4a	t	0, 10	diffuse			x	
M4b	t	0, 10	refined			x	
M4c	t	-50, 100	refined			x	
M5a	normal	0, 10	diffuse	x	x		
M5b	normal	0, 10	refined	x	x		
M5c	normal	-50, 100	refined	x	x		
M6a	normal	0, 10	diffuse	x			
M6b	normal	0, 10	refined	x			
M6c	normal	-50, 100	refined	x			
M7a	normal	0, 10	diffuse			x	x
M7b	normal	0, 10	refined			x	x
M7c	normal	-50, 100	refined			x	x
M8a	normal	0, 10	diffuse			x	
M8b	normal	0, 10	refined			x	
M8c	normal	-50, 100	refined			x	

Table 4: Comparison of Model Fit

t - errors

model	log pr(y|M)	MAPE	pr(M|y)
M1a	-53.869	0.735	0.000
M2a	-50.686	0.679	0.000
M3a	-55.964	0.471	0.000
M4a	-52.002	0.814	0.000
M1b	-52.142	0.937	0.000
M2b	-50.145	0.858	0.000
M3b	-53.262	0.820	0.000
M4b	-47.565	0.742	0.000
M1c	-56.943	1.067	0.000
M2c	-55.041	1.008	0.000
M3c	-59.316	0.802	0.000
M4c	-53.453	0.756	0.000

normal errors

model	log pr(y|M)	MAPE	pr(M|y)
M5a	-32.679	0.728	0.002
M6a	-29.566	0.682	0.037
M7a	-35.126	0.482	0.000
M8a	-31.511	0.748	0.005
M5b	-31.045	0.866	0.009
M6b	-29.509	0.817	0.040
M7b	-32.279	0.767	0.002
M8b	-26.385	0.729	0.902
M5c	-35.882	0.965	0.000
M6c	-34.405	0.930	0.000
M7c	-38.736	0.760	0.000
M8c	-32.226	0.734	0.003

pr(y|M) = marginal likelihood
MAPE = mean absolute percent error
pr(M|y) = posterior model weight

Table 5: Comparison of BT Predictions

t - errors

model	Region mean	Region nse	NV mean	NV nse	NV/UT mean	NV/UT nse
M1a	4.150	0.046	3.372	0.039	3.477	0.040
M2a	5.643	0.066	4.094	0.049	4.039	0.047
M3a	3.357	0.018	2.753	0.015	2.810	0.016
M4a	8.177	0.085	7.586	0.083	7.623	0.083
M1b	20.505	0.293	21.284	0.332	21.726	0.339
M2b	13.704	0.172	12.349	0.160	12.721	0.166
M3b	7.911	0.070	7.802	0.069	8.091	0.075
M4b	6.635	0.060	5.655	0.056	5.604	0.052
M1c	22.554	0.353	30.670	0.502	31.503	0.543
M2c	16.522	0.230	19.170	0.269	20.414	0.296
M3c	5.917	0.055	6.694	0.066	6.815	0.066
M4c	4.972	0.041	5.332	0.047	5.261	0.045

normal errors

model	Region mean	Region nse	NV mean	NV nse	NV/UT mean	NV/UT nse
M5a	3.480	0.031	2.801	0.025	2.787	0.025
M6a	4.540	0.037	3.368	0.029	3.426	0.030
M7a	3.266	0.015	2.699	0.013	2.704	0.013
M8a	5.674	0.042	4.766	0.037	4.816	0.037
M5b	11.223	0.115	10.300	0.113	10.341	0.115
M6b	8.874	0.078	7.740	0.073	7.893	0.076
M7b	6.017	0.041	5.766	0.042	5.808	0.042
M8b	5.427	0.037	4.624	0.033	4.688	0.034
M5c	10.108	0.117	11.575	0.142	11.768	0.145
M6c	8.236	0.079	9.023	0.093	9.329	0.097
M7c	4.410	0.030	4.659	0.034	4.774	0.034
M8c	4.029	0.026	4.098	0.027	4.224	0.028

weighted average, all models

	Region mean	Region nse	NV mean	NV nse	NV/UT mean	NV/UT nse
per HH	5.577	0.040	4.749	0.035	4.817	0.036
total	62,005	442	3,567,616	26,329	6,996,222	52,042

nse = numerical standard error of the posterior mean