

Necessary Conditions for Valid Benefit Transfers

Kevin J. Boyle, Nicolai V. Kuminoff, Christopher F. Parmeter, and Jaren C. Pope

Benefit transfer is an approach to estimating costs and benefits of policies in the absence of original data collection. Many different methods have been proposed (Navrud and Ready 2007). They all use estimates from one or more previous studies to predict the benefits from a new policy at a different point in space, time, or both. Under what conditions will these predictions be valid? Previous work on benefit transfer methodology has addressed this question by suggesting informal guidelines for the selection of previous studies (*Water Resources Research* 28(3), 1992 and *Ecological Economics* 60(2), 2006). There is consensus that the study sites should be “similar” to the policy site in terms of their consumer populations and in their provision of the good being valued. The need for similarity is also emphasized in Environmental Protection Agency (EPA) guidelines for benefit-cost analysis. However, there is no formal test for judging whether two sites are sufficiently “similar” to conduct a benefit transfer.

Kevin J. Boyle, professor and Department Head, Christopher F. Parmeter, assistant professor, and Jaren C. Pope, assistant professor are in the Department of Agricultural and Applied Economics at Virginia Tech. Nicolai V. Kuminoff, assistant professor, is in the Department of Economics at Arizona State University. This work was sponsored by EPA Grant #RD-83345901-0. This article was presented in an invited paper session at the 2009 AAEA annual meeting in Milwaukee, WI. The articles in these sessions are not subjected to the journal’s standard refereeing process.

This article uses a theoretical model of consumer behavior to define necessary conditions for benefit transfer to provide consistent value estimates. In summary: (i) utility must be separable in unobserved site characteristics; (ii) the study site and policy site models must be correctly specified; (iii) people must not be sorted between the sites according to unobserved features of their preferences; and (iv) one must have adequate data on the characteristics of consumers and their choices. Our analysis makes three contributions to the literature on benefit transfer methodology. First, we identify the implicit assumptions that underlie the validity of benefit transfer. Second, we discuss the testability of these assumptions. Finally, we demonstrate how structural methods, such as *preference calibration* and *preference function transfer*, can relax the need for site similarity, providing more flexibility in how benefit transfers are conducted.

Benefit Transfer in Theory and Practice

Key issues in establishing the credibility of any benefit transfer include the definition of value, the quality of the original studies, and the need to address differences in environmental quality and consumer characteristics between the original study and new policy applications. EPA's (2000) "Guidelines for Preparing Economic Analyses" include five steps that relate to these issues:

1. Describe the policy case.
2. Identify existing, relevant studies.
3. Review available studies for quality and applicability.
 - a. the basic commodities must be essentially the equivalent;
 - b. the baseline and extent of change should be similar; and

- c. the affected populations should be similar.
4. Transfer the benefit estimates.
5. Address uncertainty.

Describing the “policy case” requires a theoretical framework that reflects the shape of the demand curve or utility function as well as the constraints faced by consumers.

Two benefit transfer methods are capable of ensuring the analysis is theoretically consistent: *preference calibration* (PC) and *preference function transfer* (PFT). Both approaches seek to estimate structural preference parameters for consumers and then transfer the corresponding utility function to the new policy application. The difference between the two methods is that PFT uses the results from a previous study to transfer a utility function (Zanderson, Termansen, and Jensen 2007) whereas PC uses previous value estimates to calibrate the parameters of a utility function posited by the benefit-transfer practitioner (Smith, Pattanayak, and Van Houtven 2006). In either case the utility function can be combined with information for the new policy (e.g. environmental quality change and consumer income and characteristics) to calculate partial equilibrium values at the policy site.

PC and PFT have both been used for benefit-cost analysis. For example benefit-cost analyses of the Environmental Quality Incentives Program and the Conservation Security Program used Feather, Hellerstein, and Hansen’s (1999) random utility model of recreation demand to estimate benefits from water quality and wildlife habitat improvements from “environmental friendly” farming practices (e.g. USDA 2005). Because PFT specifies the preference functions for consumers at the study and policy

sites, it provides a convenient framework for formalizing the conditions required for benefit transfers to provide consistent value estimates for a new policy application.

A Conceptual Framework for Preference Function Transfer

This section introduces a unified conceptual framework for preference function estimation and transfer based on five assumptions that are jointly sufficient for PFTs to provide consistent estimates of partial equilibrium values. The benefit transfer objective is to compute a compensating variation (CV) value estimate for a change in environmental quality at the policy site. An individual consumer's utility can be written as a function of environmental quality (q), income (y), preferences (α), and a vector of all other characteristics describing the consumer and the site (Z). Using this notation, CV is implicitly defined by equation (1):

$$(1) \quad V(q_{PS}^0, Z_{PS}, y_{PS}; \alpha) = V(q_{PS}^1, Z_{PS}, y_{PS} - CV_{PS}; \alpha).$$

The 0 and 1 superscripts denote environmental quality before and after the change, and the PS subscripts indicate that variables are measured at the policy site.

A PFT predicts CV at the policy site by combining a parametric estimate of $V(\cdot)$ from the study site with data on characteristics describing the policy site and its affected consumer population $(q_{PS}^0, q_{PS}^1, Z_{PS}, y_{PS})$. We discuss the estimation and transfer processes in turn and define the assumptions that underlie each step.

Estimating a Preference Function at the Study Site

Consumers at the study site reveal their preferences by choosing among $j \in 1, \dots, J$ alternatives. The object of choice may be an actual choice such as housing, job, or recreation. Or it may be a choice revealed in a stated preference survey. In general we

can express consumer i 's utility from choice j as $V(q_j, Z_{ij}, y_i; \alpha_i)$. Notice that the vector of characteristics (Z_{ij}) may vary across consumers and choices (e.g. travel costs, number of children, etc.), and there is heterogeneity in preferences and income (y_i, α_i).

To recognize that the researcher is unable to observe all characteristics of sites and people, we can partition the full vector of characteristics $Z = [X, d, \xi]$, where X denotes observed site characteristics, d denotes observed demographic characteristics, and ξ denotes all unobserved characteristics of sites and people. This distinction motivates the first assumption that applies to the study site in a PFT framework. Utility is assumed to be separable in unobserved characteristics:

$$\text{ASSUMPTION SS.1: } V_{SS} = G_{SS}(q_j, X_{ij}, d_i, y_i; \beta_i) + H_{SS}(\xi_j; \gamma_i).$$

Under the separability assumption, the full vector of preferences α_i can be partitioned into preferences for observed site characteristics β_i and preferences for unobserved site characteristics γ_i . The role of separability is to guarantee that the marginal utility from environmental quality does not depend on variables which the researcher is unable to measure for the initial study.

The goal of estimation is to use the available information on the characteristics of people and their choices over the J alternatives to recover a vector of parameters θ that characterizes the distribution of preferences for observed site characteristics $F_{SS}(\beta; \theta)$. This can be done using one of several econometric methods such as a random-utility model, a corner solution model, a structural hedonic model, or an equilibrium sorting

model. In all of these frameworks, one must impose more structure to identify the distribution of preferences. This is typically done by placing parametric restrictions on F_{SS} , G_{SS} , and H_{SS} . For example random utility models often assume: (i) G_{SS} is linear and additively separable in X_{ij} ; (ii) H_{SS} follows an *iid* extreme value distribution, $H_{SS}(\xi_j; \gamma_i) = \varepsilon_{ij} \sim \text{Type I EV}$; and (iii) F_{SS} is defined such that β_i is a linear additively separable function of d_i , y_i , and an *iid* random shock from a normal distribution.

The second assumption is that all of the structural components of the study site preference function are correctly specified and estimated using appropriate methods.

ASSUMPTION SS.2: $F_{SS}(\cdot)$, $G_{SS}(\cdot)$, $H_{SS}(\cdot)$ are specified and estimated correctly.

Assuming the specification is correct and data are measured without error, it is well known that one can recover a consistent estimate for θ_{SS} if an appropriate estimator is used. In order to transfer the resulting preference function to the policy site, three additional assumptions are needed.

Transferring Preferences to the Policy Site

The first assumption is a stronger version of separability. Utility at the policy site is assumed to be separable in all unobserved characteristics at the policy *and* study sites.

ASSUMPTION PS.1: $V_{PS} = G_{PS}(q, X, d, y; \beta) + H_{PS}(\xi; \gamma)$, where $\xi = \xi_{PS} \cup \xi_{SS}$.

Together, the two separability assumptions (SS.1 and PS.1) guarantee that unobserved characteristics cancel out of the expression for compensating variation. In other words our inability to observe ξ_{PS} and ξ_{SS} does not have the potential to contaminate our

measure for CV from equation (1).

Given separability in unobserved characteristics, the PFT takes the estimate for $G(q, X, d, y; \beta)$ from the study site and adapts it to the policy site using PS data on q , X , d , and y . This process must define the preference function for people at the policy site and their joint distribution of income and demographic characteristics. First, consider the preference function. We assume the empirical specification for preferences from the study site is equally valid for the policy site:

$$\text{ASSUMPTION PS.2: } G_{SS}(q_j, X_{ij}, d_i, y_i; \beta_i) = G_{PS}(q_j, X_{ij}, d_i, y_i; \beta_i) \quad \forall i, j, \\ F_{SS}(\beta; \theta) = F_{PS}(\beta; \theta), \quad \text{and} \quad \theta_{SS} = \theta_{PS}.$$

Equality of G_{SS} and G_{PS} means the parametric structure used to represent utility for the study site is sufficiently general that differences between the study site population and the policy site population can be fully captured by differences in preference parameters, incomes, and demographic characteristics. Equality of F_{SS} and F_{PS} means the same statistical distribution (normal, beta, gamma) can be used to represent the form of preference heterogeneity at each site. Equality of θ_{SS} and θ_{PS} means the two distributions have exactly the same shape. Overall, assumption PS.2 can be interpreted as a “sorting” restriction. It guarantees that people are not systematically sorted between the study site and the policy site according to unobserved features of their preferences that influence their willingness-to-pay for changes in environmental quality.

The final assumption is that the researcher knows the joint distribution of income and demographic characteristics describing the consumer population at the policy site.

$$\text{ASSUMPTION PS.3: } \quad d_{PS}, y_{PS} \text{ are known.}$$

At first glance this may seem like a harmless data requirement. It is certainly easy to obtain spatially delineated census data on income and demographics. This may be sufficient for applications where the quality change affects all residents living in a particular set of neighborhoods that map into Census areas. However, if the change affects a particular recreation site, such as a beach, lake, or park, and the probability of attending that site is a function of income, demographics, or distance, then knowledge of the site selection process is required to predict d_{PS}, y_{PS} from the Census data.

Consistency of Partial Equilibrium Welfare Measures

Together, the five PFT assumptions allow us to develop a consistent estimator of compensating variation for the policy site population CV_{PS} . Equation (2) provides a formal definition for expected individual CV_{PS} written in terms of the study site preference function and the available data for the policy site and its population:

$$(2) \quad E(CV_{PS} | d, y) = \int_{\beta} CV_{PS}(q^0, q^1, X, d, y; \beta) F_{SS}(\beta; \hat{\theta}_{SS}) d\beta,$$

where $CV_{PS} : G_{SS}(q_{PS}^0, X_{PS}, d_{PS}, y_{PS}; \beta) = G_{SS}(q_{PS}^1, X_{PS}, d_{PS}, y_{PS} - CV_{PS}; \beta)$.

Notice that the expectation is taken with respect to the random component of preference heterogeneity. If people have homogeneous preferences for observed characteristics, equation (2) would provide an exact measure for CV_{PS} .

We can make a simple argument for the consistency of transferred benefit estimates. Assumptions SS.1 and PS.1 jointly guarantee that unobserved characteristics drop out of the expression for CV_{PS} , so it can be defined in terms of G_{PS} . Assumption PS.2 states that the structure of G_{PS} is given by G_{SS} , and any unobserved preference

heterogeneity will be captured by $F_{SS}(\beta; \theta_{SS})$. While the true value of θ_{SS} is unknown, assumption SS.2 implies the study site estimator $\hat{\theta}_{SS}$ is consistent. Thus, equation (2) is a correctly specified function of a vector of random variables, $\hat{\theta}_{SS}$. As long as this function is real valued and continuous at θ_{SS} , theorem 3.2.6 in Amemiya (1985) implies (2) will provide a consistent estimate for individual CV. Assumption PS.3 ensures we have the data on income and demographics needed to actually calculate the expected CV for each individual. Thus, the five PFT assumptions are jointly sufficient for the transfer process to provide consistent partial equilibrium benefit measures for the policy site.

An example may help to fix ideas. Suppose our policy site data are limited to information on price (p) and quality (q), and the study site preference function can be written as follows:

$$(3) \quad V_{ij} = \beta_1(y_i - p_{ij}) + \beta_{2,i}q_j + \beta_3x_j + \varepsilon_{ij},$$

where $\beta_1 = \theta_1$, $\beta_{2,i} = \theta_2d_i + \eta_i$, and $\eta_i \sim N(\theta_3, \theta_4)$.

The preference function clearly satisfies the separability requirement (SS.1). At first it may seem like a problem that we do not have data on x at the policy site, x_{PS} . However, because (3) is additively separable in x_{PS} , satisfying assumption PS.1, β_3x_j will drop out of the expression for CV_{PS} .¹ Assumption PS.2 implies that we can use the structure of (3) to define the following measure for individual CV_{PS} :

$$(4) \quad CV_{PS,i} = \frac{\theta_2d_i + \eta_i}{\theta_1} (q_{PS}^0 - q_{PS}^1), \text{ where } \eta_i \sim N(\theta_3, \theta_4).$$

Assumption SS.2 implies that the study site parameter estimates $(\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3, \hat{\theta}_4)$ are

consistent for the policy site population, and the remaining information on individual income and demographics needed to calculate CV_{PS} is provided by assumption PS.3.

Discussion

Testing the Four “S” Assumptions of Benefit Transfer

Perhaps the key implication of our conceptual framework is the need to test the four “S” assumptions that underlie benefit transfers: *separability*, *specification*, *sorting*, and *selection*. Assumptions SS.1 and PS.1 state that utility is separable in unobserved characteristics. Separability is routinely invoked in random utility models for computational convenience. At present there is no formal test for separability, and the extent to which this assumption drives benefit estimates is unclear. Assumptions SS.1 and PS.1 are likely to be valid for passive use values, such as the WTP to prevent a major oil spill in Alaska or to establish a marine preserve in the Atlantic Ocean. In both of these examples, a consumer’s WTP for the prospective quality change is unlikely to depend on characteristics of the study site where the survey is implemented or the policy site where values are transferred. While WTP may depend on a consumer’s income and education, these variables are typically recorded during the initial survey and their distributions at the policy site are observable as long as that site can be defined in terms of Census geography. The separability assumptions are more problematic for applications where WTP depends on the health of consumers, the health of their family members, or other variables that are rarely observed outside of specialized surveys.

In principle accomplishing assumption SS.2 is possible. We expect this assumption will be better met as the empirical literature evolves and better data

management procedures are employed. Careful documentation of data and estimation will help others to judge a preference function specification. Objective evaluation is possible if, for example, the study site data are generated by a quasi-experimental process that can be exploited to measure the relative performances of competing specifications (Provencher and Bishop 2004). Even without such a design, the researcher can assess the sensitivity of parameter estimates to subjective modeling decisions (Kuminoff 2009). While there is no bright-line rule here for when a study meets these criteria, the analyses cited above and other forms of sensitivity analysis conducted for the study site model can help to investigate the uncertainty associated with using a study site preference function to accomplish a PFT, addressing the fifth step in the EPA guidelines.

Assumption PS.2 states that people are not systematically “sorted” between the study site and the policy site according to unobserved features of their preferences for environmental quality. The hedonic property value literature, for example, maintains that market forces sort people within a single housing market, yet there is no consensus about the extent of sorting that occurs at the national level (Zabel and Kiel 2000; Chay and Greenstone 2005). The point is that any systematic variation in preferences between the study site population and the policy site population can invalidate a benefit transfer, and testing for spatial variation in environmental preferences at the national and sub-national levels is an important topic for future research on the validity of benefit transfers. A trivial case that rules out sorting is when the study and policy sites refer to the same geographic region at nearly the same point in time. Assumption PS.2 is also more likely to be valid for transfers based on unexpected events, such as the discovery of a cancer

cluster (e.g. Churchill County, Nevada, in 2000), where households do not have an opportunity to sort themselves according to their preferences for environmental risk prior to the event. The no sorting assumption is more problematic for transferring estimates between the site of an unexpected event and a site with long-term contamination, such as a superfund site, where people have had considerable time to adjust their behavior.

Testing the “no sorting” assumption is also important for evaluating the possibilities for using random parameter models for benefit transfer. Two key conditions must be satisfied for a random parameter model to provide valid transfer estimates. The random component of heterogeneity in preferences for environmental quality must be described by the same statistical distribution at the study site and the policy site, and the two populations must share the same functional relationship between preferences and demographics. For example, in order to transfer a preference function for beach restoration from Florida to southern California, we must be willing to assume that both populations share the same distribution of marginal utility for beach length after we condition on their demographic characteristics such as age, income, and race. This is implied by assumption PS.2.

Finally, assumption PS.3 requires that we know the incomes and demographic characteristics of the policy site population. This is trivial to test—we either have these data or we do not. It is less clear what can be done without the data. The seemingly obvious solution is to use Census data on people who live near the policy site. Yet, this approach raises the possibility of selection bias. We are unaware of any past work on the “site selection bias” that would arise from using proxy demographics from the Census (or

other sources) to predict the characteristics of the policy site population. It is common in the random utility literature to specify the probability of site selection as a function of demographic characteristics, and recent work has also suggested a role for social interactions (Timmins and Murdock 2007). The challenge is to find a way to use this information to predict the subset of the local population that will select the location of the policy site area over available substitutes. This is another important topic for future research.²

Implications for EPA's Benefit Transfer Guidelines

By adding rigor to the description of consumer behavior, our conceptual framework relaxes the rigidity of EPA's "Guidelines for Preparing Economic Analyses" to provide more flexibility in how benefits transfers can be conducted. The EPA guidelines require: (i) the *same* basic commodities to be consumed at the study and policy sites; (ii) *similar* baseline levels of those commodities; and (iii) *similar* consumer populations. These requirements can be relaxed if one is willing to specify the preference functions that guide consumer behavior. First, consider the basic commodities. The study site and policy site commodities must both convey the same measure of environmental quality (q), for example, lake water clarity. The lake sites may differ in the other characteristics that they convey, as long as each lake can be consistently described by a set of characteristics (Z), and utility is separable in the characteristics that cannot be measured at all lakes (ξ). Differences in commodities are taken into account during the transfer process by calibrating the utility function using the available policy site data.

Second, consider the levels of each commodity and the size of the change at the

policy site. As long as our structural specification for the preference function is accurate, we can use it to assess the partial equilibrium welfare implications from any change at the policy site. Of course this is predicated on the assumption that the range of environmental quality at the study site is sufficient to identify the structural parameters of the preference function.

Finally, our restrictions on sorting and selection relax the need for similar populations. The study site and policy site populations may differ as long as the systematic variation in their preferences for environmental quality is fully explained by observable demographic characteristics (i.e., the transfer estimate can be calibrated for such differences through d). The main implication here is that the EPA transfer conditions do not need to be treated as strict requirements as long as the benefit transfer method is capable of calibrating for the relevant differences in the characteristics of people and their choices at the study and policy sites.

Conclusions and Next Steps

Previous work on the methodology of benefit transfer has sought to establish broad, informal guidelines for the transfer process. In contrast we have defined the specific assumptions that make benefit transfers work. While we have defined these assumptions in the context of PFT, they are not unique to this method. They serve as necessary conditions for *any* theoretically consistent approach to benefit transfer. Preference function transfer simply provides a convenient framework in which to formalize the assumptions and consider their implications.

In reduced-form frameworks of PFTs, such as meta-analysis, the four “S” (4S)

assumptions—separability, specification, sorting, and selection—are implicit in the transfer process. Violations of the study site assumptions (SS.1-2) will generate measurement error in the dependent variable of a meta-regression, and violations of the policy site assumptions (PS.1-3) will confound the model’s external validity. For example a meta-analysis that uses past estimates for welfare measures from one site to predict welfare at a different site implicitly assumes that any systematic variation in preferences between the two populations is explained by the observable characteristics of the population that serve as regressors. While the 4S assumptions are sufficient for the validity of a structural PFT, they are not sufficient to ensure the validity of a reduced-form transfer. In addition to satisfying the 4S assumptions, a meta-regression must also be correctly specified and sample selection issues must be addressed (Nelson and Kennedy 2009; Rosenberger and Johnston 2009).

It is also important to underscore that our analysis reflects a partial equilibrium perspective. We have held prices fixed, and we have implicitly assumed that people do not adjust their behavior in response to changes in quality. While this seems extreme, it may provide a reasonable approximation for the short-run response to moderate changes. A partial equilibrium perspective is clearly less appropriate for large changes that may force people to adjust immediately or even relatively small changes where the incremental cost of adjustment is minimal. Extending our analysis to consider general equilibrium adjustment would require us to address a fifth “S”—*substitution*. We conjecture that one could make consistent predictions for general equilibrium welfare measures using data on the substitution possibilities faced by the policy site population.

Overall, our analysis suggests a new direction for future research on benefit transfer methodology. By pursuing an agenda of testing the structural assumptions that underlie benefit transfers, we can identify which (if any) assumptions are systematically violated, assess the consequences, and look for solutions. In ongoing research we are using simulation methods to assess the relative importance of violating each “S” assumption in a policy relevant application of PFT. Preliminary results suggest that welfare measures are particularly sensitive to violations of the “no sorting” condition.

References

- Amemiya, T. 1985. *Advanced Econometrics*. Cambridge: Harvard University Press.
- Chay, K.Y., and M. Greenstone. 2005. "Does Air Quality Matter? Evidence from the Housing Market." *Journal of Political Economy* 113(2):376-424.
- Feather, P., D. Hellerstein, and L. Hansen. 1999. *Economic Valuation of Environmental Benefits and the Targeting of Conservation Programs: The Case of the CRP*. Washington DC: U.S. Department of Agriculture, ERS Agr. Econ. Rep. 778, April.
- Kuminoff, N.V. 2009. "Decomposing the Structural Identification of Nonmarket Values." *Journal of Environmental Economics and Management* 57(2):123-39.
- Navrud, S., and R.C. Ready. 2007. *Environmental Value Transfer: Issues and Methods*. The Netherlands: Springer.
- Nelson, J.P., and P.E. Kennedy. 2009. "The Use (and Abuse) of Meta-Analysis in Environmental and Natural Resource Economics: An Assessment." *Environmental and Resource Economics* 42(3):345-77.
- Provencher, B., and R.C. Bishop. 2004. "Does Accounting for Preference Heterogeneity Improve the Forecasting of a Random Utility Model? A Case Study." *Journal of Environmental Economics and Management* 48(1):793-810.
- Rosenberger, R.S., and R.J. Johnston. 2009. "Selection Effects in Meta-Analysis and Benefit Transfer: Avoiding Unintended Consequences." *Land Economics*, 85(3):410-428.
- Smith, V.K., S. Pattanayak, and G. Van Houtven. 2006. "Structural Benefit Transfer: An

- Example Using VSL Estimates." *Ecological Economics* 60(2):361-71.
- Timmins, C., and J. Murdoch. 2007. "A Revealed Preference Approach to the Measurement of Congestion in Travel Cost Models." *Journal of Environmental Economics and Management* 53(2):141-290.
- U.S. Department of Agriculture (USDA). 2005. *Conservation Security Program: Amendment to the Interim Final Rule Benefit Cost Assessment*. Washington DC, March.
- U.S. Environmental Protection Agency (EPA). 2000. *Guidelines for Preparing Economic Analysis*. Washington DC, September.
- Zabel, J.E., and K.A. Kiel. 2000. "Estimating the Demand for Air Quality in Four U.S. Cities." *Land Economics* 76(2):174-94.
- Zanderson, M., M. Termansen, and F.S. Jensen. 2007. "Testing Benefits Transfer of Forest Recreation Values over a Twenty-Year Time Horizon." *Land Economics* 83(3):412-40.

¹ Conversely, assumption PS.1 implies that any variable observed for the policy study, but not the initial study, will drop out of the expression for compensating variation.

² Another challenge is that practitioners are not consistent in the demographics included from one study to another, and the demographics chosen are often not the same as those maintained in large, national databases such as the census data.