

A SIMPLE METHOD OF INCORPORATING INCOME EFFECTS INTO LOGIT AND NESTED-LOGIT MODELS: THEORY AND APPLICATION

EDWARD. R. MOREY, VIJAYA R. SHARMA, AND ANDERS KARLSTROM

Substantive income effects are incorporated in a logit or nested-logit model by assuming that utility is a piece-wise linear spline function of residual income. Specific income data are not required, only income by category. Expected compensating variation is easily and accurately approximated by the difference between expected maximum utility in the proposed and initial state, multiplied by the inverse of the individual's initial marginal utility of money. This approximation is *almost* exact because although any policy can, in theory, cause an individual to jump income categories, for most policies this probability will be very small.

Key words: compensating variation, income effects, log-sum formula, nested-logit.

Substantive income effects are incorporated into a logit or nested model by assuming that utility is a piece-wise linear spline function of income; that is, marginal utility of income is assumed to be a step function of income. For estimation, specific income data are not required, only income by category; an important feature because many, if not most, data sets report income by category only. One can assume as many or as few pieces (steps) as the data and application warrant.

Given this method of incorporating income effects, an individual's expected compensating variation, $E[cv]$, is easily and accurately approximated using the standard log-sum formula for $E[cv]$, which is exact if there are no income effects. Specifically, the difference between expected maximum utility in the proposed and initial state is multiplied by the inverse of the individual's marginal utility of money in the initial state. This approximation is *almost* exact because although any policy can, in theory, cause an individual to jump income categories, for most policies this probability will be very small. Whether a policy causes an individual to jump from one income category to another is a function of the individual's specific epsilon draw, and an increas-

ing function of the magnitude of the policy change and the number of categories. Even if the policy being evaluated is trivial in an absolute sense, there is always some epsilon vector that would lead to a change in income categories.

The application is a logit model of choice of health provider by malaria patients in rural Nepal.¹ A significant proportion of the affected population is poor and one would expect that this will impact on these households' willingness to pay for improved care. Information was available to determine whether the household is poor but exact income is not known. The marginal utility from expenditures on the numeraire is assumed to be constant up to the poverty line, then a different constant for all additional expenditures on the numeraire. That is, the marginal utility takes only one step and it is at the poverty line. Estimated $E[cv]$ for providing more sites with blood-testing capabilities is approximated for each individual in the sample. These estimates vary significantly and substantively as a function of whether the household is poor.

For comparison, each household's exact estimated $E[cv]$ is also calculated using the McFadden simulation technique after randomly assigning exact income levels by

Edward Morey and Vijaya Sharma are at the Department of Economics, University of Colorado, Boulder, Colorado; Anders Karlstrom is at Systems Analysis and Economics, Royal Institute of Technology, Stockholm, Sweden.

The comments of the editor and two referees significantly improved the article, redirecting it along a more desirable path.

¹ Morey, Sharma, and Mills consider the application in detail.

category.² In terms of sample averages, there is no approximation error associated with using the log-sum formula, and, in terms of individual $E[cv]$'s, the approximation error is less than 1% for over 95% of the sample and never more than 3.5%

An important question is whether one even needs to resort to an approximation when utility is assumed as a piece-wise linear spline function of income. The answer is yes if one has income data only by category because the exact $E[cv]$ is a function of the individual's exact income. In the absence of this information, the best one can achieve is an approximation.

In cases where one has exact income data—unusual—and one assumes the spline specification, an exact closed-form solution for $E[cv]$ is preferred to the approximation, if it exists. Hanemann and others have shown with some simple examples that there can be closed-form solutions for the $E[cv]$ even when there are income effects. This led Hanemann, one of the referees, to conjecture that there might be a general closed-form solution for the $E[cv]$ that is a variant on the log-sum formula when utility is assumed as a piece-wise linear spline function of income. However, after much effort, we have not been able to find or derive a closed-form solution for the general case. The complicating factor is that an individual's marginal utility of money can vary as a function of the alternative he or she chooses; so, in most cases, an individual's marginal utility of money after the change depends on his or her epsilon draw.³

Closed-form solutions can be derived in those cases where all individuals of the same type (income, etc.) who switch alternatives, switch to the same alternative. For example, assuming only one step, a closed-form solution can be derived for the case of J alternatives when the policy involves an improvement (price decrease or quality increase) in

one and only one alternative. In this case, if one switches alternatives, one switches to the alternative that has been improved. A closed-form solution also exists in the one-step case for any policy change if the number of alternatives is limited to two. This is because only one of the alternatives has a positive transition probability.

As an explanation, expected utility is given by the log-sum formula. Intuitively and simply, if there is a unique and constant marginal utility of money, it can be used to convert expected utility into money and $E[cv]$ has a closed-form solution. This is the case if there are only two alternatives or if one and only one alternative is improved. There is no unique and constant marginal utility of money if more than one alternative is improved. There is also not one alternative if an alternative deteriorates and there are more than two alternatives: one might switch to any one of the other $J - 1$ alternatives.

Our search for closed-form solutions started with the method of calculating the $E[cv]$ developed by Karlstrom (see also Karlstrom and Morey). This method builds on earlier work by Hanemann. In the presence of income effects in GEV models, $E[cv]$ can always be expressed as a sum of terms where some of the terms have obvious closed-form expressions and some are definite one-dimensional integrals. If the sum of these integrals has a closed-form solution, $E[cv]$ also has a closed-form solution. This search method leads to a closed-form solution for the case of J alternatives and one step when the policy involves an improvement (price decrease or quality increase) in one and only one alternative. Both *Mathematica* and *Maple* found the closed-form solutions to the definite integrals for this case. The solution takes the form of a "weighted log-sum formula." However, pursuing this method, we did not find a closed-form solution when there are J alternatives and one of the alternatives deteriorates. Neither *Mathematica* nor *Maple* was able to solve the definite integrals.

The rest of the article is organized as follows. The first section specifies the general form of the conditional indirect utility function assuming that utility is a piece-wise linear spline of expenditures on the numeraire. The second section presents our empirical example.

Utility assumed a piece-wise linear spline of expenditures on the numeraire.

² The exact $E[cv]$ is a function of the household's income, not just its income category (poor or not). Because income is unknown in the data set, the approximation error can only be calculated after a specific income level is assumed. For each individual in the sample, five different income levels were assumed.

³ The same outcome occurs if one introduces income effects, not by assuming a spline, but by assuming that the marginal utility of money is a function of the alternative chosen; that is, assuming varying complementarity between each alternative and the numeraire such that the individual feels differently about the numeraire depending on which alternative is chosen. See Hanemann and Dow for details on this other method of incorporating income effects. Although the two methods are similar, it is important to distinguish between them.




Figure 1. Piece-wise linear spline (example 1)

Start with the case of only one step, that is, the utility individual i gets from choosing alternative j is

(1)

$$u_{ji} = v_{ji} + \epsilon_{ji}$$

$$= \begin{cases} \alpha_0(y_i - p_{ji}) + \beta \mathbf{x}_j + \epsilon_{ji} & \text{if } (y_i - p_{ji}) \leq m_0 \\ \alpha_0 m_0 + \alpha_1(y_i - m_0 - p_{ji}) + \beta \mathbf{x}_j + \epsilon_{ji} & \text{if } (y_i - p_{ji}) > m_0 \end{cases}$$

where y_i is the income of individual i , p_{ji} is the cost of alternative j to individual i , \mathbf{x}_j is the relevant vector of characteristics of alternative j , and m_0 is the level of expenditures on the numeraire where the marginal utility from those expenditures switches from α_0 to α_1 . For example, m_0 might be the poverty line or the point where one becomes rich. In terms of the stochastic specification, assume either a logit model or a nested-logit model. See figures 1 and 2 which correspond to equa-




Figure 2. Step function (example 1)

tion (1) where μ in figure 2 is α . The figures are drawn such that $\alpha_0 > \alpha_1$, but that need not be the case. Note that equation (1) implies that the budget is exhausted and u_{ji} is a continuous function of y_i .

Because the data set does not include each household's specific income, for estimation we assume either $(y_i - p_{ki}) \leq m_0 \forall k$ or $(y_i - p_{ki}) > m_0 \forall k$; that is, whether the individual is initially poor is not a function of the alternative they choose. In applications this assumption will typically hold, and one will be able to identify those situations where it might not hold.⁴ Given the assumption, the probability that individual i chooses alternative j in the logit case is

$$(2) \quad \Pr_{ji} = \frac{\exp(-\alpha_h p_{ji} + \beta \mathbf{x}_j)}{\sum_{k=1}^J \exp(-\alpha_h p_{ki} + \beta \mathbf{x}_k)}$$

where $\alpha_h = \alpha_0$ if poor
 $\alpha_h = \alpha_1$ if not poor.

The same holds for the nested-logit model.

Now consider a policy that changes costs and characteristics of individual i from $\{\mathbf{p}_i^0, \mathbf{x}^0\}$ to $\{\mathbf{p}_i^1, \mathbf{x}^1\}$. The compensating variation individual i associates with this change, cv_i , is the amount of money that must be subtracted from his or her income in the new state to equate utility in the new state with utility in the initial state. It is a random variable from the researcher's perspective in that it is a function of individual i 's epsilon draw. The search is therefore for its expectation, $E[cv_i]$. Given equation (1), $E[cv_i]$ is, in most cases, closely approximated with

$$(3) \quad \tilde{E}[cv_i] = \frac{1}{\alpha_h} \left\{ \ln \left[\sum_{j=1}^J \exp(-\alpha_h p_{ji}^1 + \beta \mathbf{x}_j^1) \right] - \ln \left[\sum_{j=1}^J \exp(-\alpha_h p_{ji}^0 + \beta \mathbf{x}_j^0) \right] \right\}$$

That is, $\tilde{E}[cv_i]$ is the change in expected maximum utility, taking account of the spline, weighted by the marginal utility of money in the initial state.

⁴ If the assumption does not hold and specific income levels are known, one can estimate

$$\Pr_{ji} = \frac{\exp(v_{ji})}{\sum_{k=1}^J \exp(v_{ki})}$$

directly.

Consider why equation (3) is almost equal to $E[cv_i]$ for many policies considered by environmental economists? Equation (3) is not exact because there is always some probability that the policy will cause the individual to change income category. Whether the policy causes a specific individual to change income category depends on the magnitude and extent of the price and characteristic changes, and the individual's epsilon draw. For policies that do not have a major effect on welfare (many environmental projects), only individuals with unlikely epsilon draws will change income category.

Before generalizing equation (1) to include more categories, note that is not more or less general than assuming that utility from consumption of the numeraire is some nonlinear but a differentiable function of $(y_i - p_{ji})$. It is just different.

Equation (1) can be generalized to include as many income categories as the application warrants. For example, if there were three categories (e.g. poor, rich, and middle class)

$$\begin{aligned}
 (4) \quad v_{ji} &= \alpha_0(y_i - p_{ji}) + \beta \mathbf{x}_j && \text{if } (y_i - p_{ji}) \leq m_0 \\
 &= \alpha_0 m_0 + \alpha_1(y_i - m_0 - p_{ji}) + \beta \mathbf{x}_j && \text{if } m_0 < (y_i - p_{ji}) \leq m_1 \\
 &= \alpha_0 m_0 + \alpha_1 m_1 + \alpha_2(y_i - (m_0 + m_1) - p_{ji}) + \beta \mathbf{x}_j && \text{if } (y_i - p_{ji}) > m_1
 \end{aligned}$$

where $m_1 > m_0$. It might look at figure 3. Note that as the number of categories increases, the probability that a policy will cause an individual to jump categories increases, and the approximation becomes less exact. However, in practice, this is not a concern as long as the range of each income category is large relative to the magnitude of the $E[cv]$'s.

An empirical example: estimating malaria patients' household compensating variations for health care proposals in rural Nepal.

Only the minimum details of the empirical example are presented. The full model and results are described in Morey, Sharma, and Mills. The spline model was developed because it was our expectation that $E[cv]$ for improved treatment options in rural Nepal would vary significantly with whether the household was poor, our desire was to be utility theoretic, and the fact was that from the data we could only tell if the household was poor or not. When our data were collected, malaria patients in




Figure 3. Piece-wise linear spline (example 2)

the survey area had the choice of six types of malaria treatment providers: four types of government providers and two types of private providers. An important component of the treatment by the government providers is the time it takes to get the blood test results. This can be days, or even a couple of weeks; the blood samples have to be transported long distances over poor roads. The blood test identifies the type of malaria, and government malaria clinics do not prescribe the "cure" until they have this information. Each patient's household compensating variation, is estimated for providing more government malaria clinics with blood-testing capabilities. These additional labs will save substantial time between the visit to the care provider and final treatment; 0.9 to 3.4 hours in Dhanusha district, 5 to 8.4 hours in Nawalparasi district.

The reported model is simple logit; nested specifications were estimated but did not significantly improve the fit. Providers, by type, were characterized with alternative specific intercepts and required travel time from that facility to the nearest lab. Patient characteristics that influence choice include travel costs, income category, household size, gender of patient, whether the patient is a child, type and severity of malaria, district of residence, and expected wait time between the onset of symptoms and the next expected home visit by a malaria worker.⁵

⁵ These occur on a semiregular basis but the length between them varies across villages. District of residence is relevant because the part of Nawalparasi district surveyed is on the border with India and private practitioners from India cross the border and roam from village to village on bicycles soliciting and treating patients.

Table 1. Sample Mean: Approximate Versus Simulated $E[cv]$

	$\tilde{E}[cv]$	$E[cv]$ r 1	$\tilde{E}[cv]$ r 2	$\tilde{E}[cv]$ r 3	$\tilde{E}[cv]$ r 4	$\tilde{E}[cv]$ r 5
Dhanusha	0.8779	0.8778	0.8779	0.8779	0.8777	0.8777
Nawalparasi	3.072	3.072	3.073	3.072	3.072	3.073

The specification of α_h is a bit more general than that presented above. Specifically, estimated a_h is found to be a function of household size, the patient's gender, and whether the household is poor. Poverty, not surprisingly, significantly and substantially increases the marginal utility of money. Household size and a male patient decrease it.

Two things need to be demonstrated: (a) that on both an aggregate and individual household level $\tilde{E}[cv]$ is close to exact, and (b) that including the step at the poverty line significantly affects the household's estimated $E[cv]$. Begin with the first issue. Without full income data, it is not possible to determine how close each household's $\tilde{E}[cv]$ is to its $E[cv]$. To assess this, we randomly assigned incomes to households by income category and simulated the $E[cv]$ using a method developed by McFadden.⁶ See also Herriges and Kling, and Morey. In summary, cv_i is a random variable from the researcher's perspective, it depends on ϵ_i . Because there are six provider types, six epsilons are randomly drawn for each household from the simple extreme value distribution. Conditional on this vector of epsilons, one calculates $cv_i(\epsilon_i)$. This procedure is repeated 20,000 times and $E[cv_i] = 0.00005 \sum_{d=1}^{20,000} cv_i(\epsilon_i^d)$. We repeated these 20,000 simulations for each household five times, each time with a different random assignment of income. For each round of income assignment, we calculate the mean of the $E[cv]$ for all of the households in the sample (314 in Dhanusha and 175 in Nawalparasi).

Table 1 reports the sample mean of the $\tilde{E}[cv]$ for each district and the corresponding mean of the simulated $E[cv]$ in each of the five rounds. The unit of currency is the Nepalese rupee (Rs.). For comparison, the average wage rate for farm workers is Rs 1.50 per hour for an adult male. These estimated compensating variations are not large but are sufficient to cover the cost of the required additional lab

instruments and training. Note that these estimates are not for treatment but for faster treatment; everyone is eventually treated because everyone is periodically visited by a malaria worker.

Now consider the approximation error household by household. Table 2 reports for the five rounds the range on the five frequency distributions of the approximation errors (differences between $\tilde{E}[cv_i]$ and simulated $E[cv_i]$ for each of the five rounds). For example, over the five rounds, the percentage of households with approximation errors less than or equal to 0.1% varied between 51.9% and 59.3% in Dhanusha. The important thing to note is that the approximation error was never more than 3.3% and always less than 1% for at least 95.9% of the households.

In this empirical example, the approximation error in percentage terms is bounded from above by the difference in percentage terms between $1/\alpha_0$ and $1/\alpha_1$. For example, if $\alpha_0 = 0.5$ and $\alpha_1 = 0.4$ (the approximate estimated values in the application), and if a household had an income of $m_0 + \textit{epsilon}$, calculation of $\tilde{E}[cv]$ would multiply the difference between expected maximum utility in the two states by 2.5 rather than 2.0, so $\tilde{E}[cv]$ will overestimate $E[cv]$ by 25%. This extreme case did not arise in our simulations and is unlikely to arise in any application where exact income is known because it is highly unlikely that there will be households with income right at m_0 . In addition, unless the policy being evaluated has a very large effect on welfare, there are few households with an income within their $E[cv_i]$ of m_0 . We conclude that the approximation is *close enough for government work*.

Table 2. Cummulative Frequency of Approximation Error in the 5 Rounds

Approx Error	Dhanusha	Nawalparasi
0.1% or less	51.9% to 59.3%	20.6% to 29.7%
0.5% or less	98.1% to 99.4%	78.8% to 81.1%
1% or less	99.7% to 100%	95.9% to 97.7%
3.3% or less	100%	100%

⁶ Relevant ranges were chosen for each income category; one of the authors is Nepalese. Over each range a uniform distribution was assumed.

Consider now the importance of including income effects in general, and specifically doing it by assuming utility is a linear spline of expenditures on the numeraire. Typically willingness to pay is an important determinant of ability to pay, so cv_i should typically be a function of income. In our application the simple spline model predicts significantly better than a no-income effects model ($\alpha_0 = \alpha_1$). With the spline model, we calculated $\hat{E}[cv_i]$ first assuming that the household is poor, then not poor. Both of these calculations are then compared to the estimated $E[cv_i]$ from the no-income-effects model. The no-income effects model understates $E[cv]$ by as much as 19% and overstates it by as much as 33%. The mean of the deviations was small in Nawalparasi and large in Dhanusha.

Summary

The article outlines a simple method of incorporating income effects into logit and nested-logit models. Specifically, utility is assumed to be a linear spline of expenditures on the numeraire; that is, the marginal utility from the numeraire is a step function. This method of incorporating income effects does not require exact income data; it only requires income by category. Given that exact income data are hardly ever available, the spline is one of the few ways of incorporating income effects, unless one is willing to convert income categories into exact incomes, which would generate significant approximation errors.

Although one can incorporate as many steps as desired, we suspect that one or two steps will typically do the trick: whether the household is or is not poor, or whether it is poor, rich, or other. These are the steps to check first.

When income effects are incorporated in this way, our expectation of a household's compensating variation is accurately approximated by the change in their expected utility weighted by the inverse of their marginal utility of money in the initial state.

[Received January 2000; final revision received May 2002.]

References

- Dow, W.H. "Flexible Discrete Choice Demand Models Consistent With Utility Maximization: An Application to Health Care Demand." *American Journal of Agricultural Economics* 81(1999):680–85.
- Hanemann, W.M. "Welfare Analysis with Discrete Choice Models." In J.A. Herriges and C.L. King, eds. *Valuing Recreation and the Environment*. Cheltenham, UK: Edward Elgar, 1999, pp. 33–63.
- Herriges, J.A. and C.L. King. "Nonlinear Income Effects in Random Utility Models." *The Review of Economics and Statistics* 81(1999):62–72.
- Karlstrom, A. "Four Essays on Spatial Modelling and Welfare Analysis." PhD thesis, Department of Infrastructure and Planning, Royal Institute of Technology, Stockholm, Sweden, 1999.
- Karlstrom A. and E. Morey. "Quickly and Accurately Calculating Welfare Measures in GEV Random Utility Models with Income Effects: A Primer." Working paper, Department of Economics, University of Colorado, Boulder, 2002.
- McFadden, D. "Computing Willingness-to-Pay in Random Utility Models." In J.C. Moore, R.G. Riezman, and J.R. Melvin, eds. *Trade, Theory and Econometrics: Essays in Honour of John S. Chipman*. London: Routledge, 1999.
- Morey, E.R. "TWO RUMS unCLOAKED: Nested-Logit Models of Site Choice and Nested-Logit Models of Participation and Site Choice." In J.A. Herriges and C.L. King, eds. *Valuing Recreation and the Environment*. Cheltenham, UK: Edward Elgar, 1999, pp. 65–120.
- Morey, E.R., V.R. Sharma and A. Mills. "Willingness to Pay and Determinants of Choice for Improved Malaria Treatment in Rural Nepal." *Social Science and Medicine*, forthcoming 2002.