Revisiting the Hedonic Travel Cost Model

by

Nancy E. Bockstael and Kenneth E. McConnell

WP 98-09

Department of Agricultural and Resource Economics
The University of Maryland, College Park

Submitted to JEEM, April 29, 1998

Revisiting the Hedonic Travel Cost Model

Nancy E. Bockstael and Kenneth E. McConnell

Department of Agricultural and Resource Economics
University of Maryland
College Park MD 20742

April 27, 1998

Possible running head: Hedonic Travel Cost Model Revisited

ABSTRACT

The hedonic travel cost model has been used to measure welfare effects associated with non-market goods. Because nature and not markets dictates the arrangement of recreational sites, a cost function with well-behaved properties is difficult to define. We have almost no *priors* on its shape and there are no forces at work to regularize the costs and characteristics relationship. Yet, the slope of this function plays the important role of price in the second stage demand estimation. At best, this variable will be measured with considerable inaccuracy and, at worst, will be non-sensical. Even if marginal values could be accurately derived, the approach does not provide useful information for many important policy questions.

Keywords: hedonic travel cost, environmental valuation, recreation

Revisiting the Hedonic Travel Cost Model

Introduction

Estimating the economic value of changes in environmentally related public goods has become an important task for environmental economists. A frequently encountered valuation problem involves deducing the value of environmental quality changes to recreationists, using data on visits to a set of geographically dispersed recreation sites. The sites embody different characteristics and different levels of some environmental amenity. Government policies can influence environmental quality at the sites, but other exogenous events such as releases of hazardous materials can affect the quality, as well. Measuring the benefits of improvements or the losses due to degradation is the central purpose of this literature. Examples of this research include Kaoru, Smith and Liu [5], Bockstael, Hanemann, and Kling [1], and Morey [6].

Economists have devised a variety of strategies for addressing this valuation problem, although a theoretically consistent and empirically tractable model appears unlikely.[1] The principal approaches are the random utility model and the hedonic travel cost model. Both models make simplifying assumptions, assumptions that would give pause to students of demand theory not familiar with recreational demand models, and both have as their goal the measurement of welfare effects resulting from changes in environmental quality characteristics at recreation sites.

Of the two models, the random utility model is the more frequently employed, but applications of the hedonic travel cost model continue to appear in the literature. The basic hedonic travel cost model was first constructed in Brown and Mendelsohn [2], although there have been several modifications stemming from problems revealed in the practical application of the approach (see Smith and Kaoru [9] for a list of these). One important adjustment has been the adoption of the 'per trip' perspective, embodied in recent applications. The purpose of this paper is to consider whether the new version of this modeling strategy has successfully resolved the problems inherent in the original. In considering this question, we compare the features of the hedonic travel cost model with the random utility version of the travel cost model and with its other namesake, the traditional hedonic (property) model of Rosen [8].

The choice of a model for a particular economic task depends on several criteria: logical consistency, tractability, simplicity, credibility and a host of lesser characteristics. But ultimately, an appropriate model is one that tells a convincing story about individual behavior. No model is exactly right but some work better than others. In this paper, we examine the hedonic travel cost model from a theoretical and practical point of view, and suggest that in many cases it fails to tell a convincing story about individual behavior. Further, because it focuses on marginal valuation it does not lend itself to the sort of valuation tasks that policy analysis or damage assessment normally entails.

The Theoretical Structure of the Model

The random utility and hedonic travel cost models, as well as the more conventional hedonic model of property values, are all members of a particular class of modeling strategies in the revealed preference literature. Unlike models that treat environmental quality as a scalar and seek to extract welfare measures from demand decisions for related goods[2], this class of models views environmental quality as a vector, where each element of the vector is a realization of environmental quality at a different location. In this case, it is the vector of environmental qualities at different locations that is exogenous to the individual and affected by policy or accident, but the actual "ambient" environmental quality that the individual exposes himself to is endogenously determined. The individual is viewed as choosing location, and this choice of location - rather than the quantity demanded of some related good - is the behavior that reveals how the individual values environmental quality changes.

The hedonic travel cost model blends aspects of the random utility model and the hedonic property value model. It takes the setting and type of data used for the former and the modeling strategy of the latter. But all three models use information on how people choose among alternative packages that cost different amounts and embody different levels of the environmental good.

Following the development of the hedonic travel cost model in a recent application (Englin and Mendelsohn [4]), we view the individual as maximizing a utility function that is a function of "ambient" levels of characteristics, Z, and a numeraire, W. The elements of Z are defined by the set of attributes that individuals value when participating in the recreation of interest. The "ambient" levels of these characteristics are the levels that exist at the *chosen* recreational site. Individuals can change their ambient levels themselves, by choosing another site. The fact that individuals choose their own Z's, combined with the appeal of the idea of

---

[1] Phaneuf, Kling, and Herriges [7] have estimated a model which shows promise in this area.
[2] These draw on assumptions of complementary or substitution between the environmental quality and marketed or household produced goods.

trading off more money for better attributes, leads to the hedonic travel cost model. In this model, the individual is believed to

$$
\text{(1)} \qquad \begin{array}{ll} \text{Max}_{Z,W} & U(Z,W) \\ \text{subject to} & Y - C(Z) - W^*P \end{array}
$$

where P is the price of the numeraire (and set to 1) and C(Z) is the hedonic travel cost function, which presumably relates the levels of attributes at a site with the cost of access. Policy or accident can alter C(Z). Demand functions for attribute levels are derived from the first order conditions implied by this underlying structure.

The intuitive appeal of the hedonic travel cost model comes from the analogy with the hedonic model of housing characteristics. People who want more of any given characteristic make the tradeoff—higher costs for more of the characteristic. In the housing market, there is no doubt that such tradeoffs are available. A house with an additional room costs more, just as a house with a view or a house with lower exposure to air pollution, *ceteris paribus*. The market equilibrium will not provide more of a desirable attribute at lower prices. Markets sometimes provide such tradeoffs in recreation. One can pay for professional guide services and probably be more successful in hunting and fishing enterprises. The question is whether the usual circumstances in which travel cost models are applied align themselves with the basic workings of the hedonic model.

The Conceptual Basis for Hedonic Models

The conventional hedonic model begins with the presumption of a continuous function relating the price of a heterogeneous good to its quality characteristics. The heterogeneous good is generally real estate, and the quality characteristics of interest often include an environmental amenity. Rosen provides the theoretical basis of this hedonic price function. It is the locus of equilibrium points in a market in which properties are allocated across households through a bidding mechanism.[3] Households are assumed to have preferences over these attributes, but it is the fact that a given property can be purchased by only one household that causes the price of properties possessing more of a desirable attribute to be bid up. Price is serving an allocation role. The underlying story is convincing. Both economic logic and empirical evidence suggest that prices will align themselves with levels of the characteristics. In the course of estimation,

[3] In the typical version of this model, the stock of housing, and hence of housing attributes, is fixed, and households bid for the fixed stock.

when a statistically significant *negative* marginal price occasionally emerges for a desirable attribute, it generally signals a specification problem. For example, levels of a desirable characteristic included in the regression may be negatively correlated with levels of another characteristic omitted from the regression. Hedonic housing prices are, however, generally quite systematic in showing positive marginal prices for desirable attributes.

Hedonic travel cost applications begin with a presumption of a similar hedonic price function for characteristics of recreational sites, the function C(Z) in model (1). This presumption implies some underlying process by which the cost of purchasing a trip depends positively on the levels of the characteristics of the trip. In the recreational case, however, no allocation mechanism usually exists to force an alignment of prices (i.e. travel costs) and characteristic levels. An individual faces an array of alternative sites with different quality characteristics and different costs of access, where cost of access is determined by how far from a site an individual happens to live, not by a market mechanism that rations the sites among users. Consequently, if we plot travel costs against attributes for all alternative sites available to the individual, there is no reason for the resulting scatter of points to be a *function* at all, let alone a well-behaved one that relates price and environmental quality. Because nature, and not a market, organizes the spatial distribution of sites and characteristics, some people may happen to live quite close to the environmentally most desirable sites, for example. Consequently, the chances are good that if one looks at *all* configurations of site attributes and travel costs, more of at least some attributes will be available at lower cost.

If nature arranges sites such that the site with the most desirable attributes is the closest one for an individual, then his choice of site conveys no information about how he values site attributes. This is true whether one uses the random utility model, the hedonic travel cost model, or any other model that attempts to extract information from site choice. In this case, there is simply no useful information about the individual's preferences conveyed in his choice, because no tradeoff between better attributes and higher cost is required. We are interested here in the more usual case, one in which sites are arranged rather randomly such that the recreationist need make at least some trade-offs between money and attributes, and therefore some information is conveyed by his choice.

Whatever the spatial configuration of sites of different qualities, however, there is no basis to assume that a functional relationship exists between access cost and levels of all attributes, for all sites. The absence of any convincing story to ensure such a relationship is, perhaps, the hedonic travel cost model's chief failing. An empirical manifestation of this failure is the negative marginal price problem that has arisen in many applications of the approach.

Making Sense of the Story

The users of the hedonic travel cost approach are aware of the difficulties that the distribution of attributes creates for the estimation of the cost function. However, with the exception of Smith, Palmquist and Jakus [10], no careful attempt has been made to explain the story that would make the approach more compelling. Strictly speaking, the array of sites with given characteristics and given costs of access for any one individual constitutes a constraint set to the individual. Some of the sites will pose non-binding constraints to the terms of trade available to the individual, because they will be dominated by other sites that either offer at least as high a level of attributes at a lower price or a higher level of attributes at no greater price. Thus, to make the mechanics of the problem work, we must confine ourselves to some subset of sites or a frontier. For the simple one attribute case, Figure 1 illustrates the problem. Each of the randomly scattered points in the diagram represents a site with a given travel cost to the individual and a given level of the attribute.

Figure 1

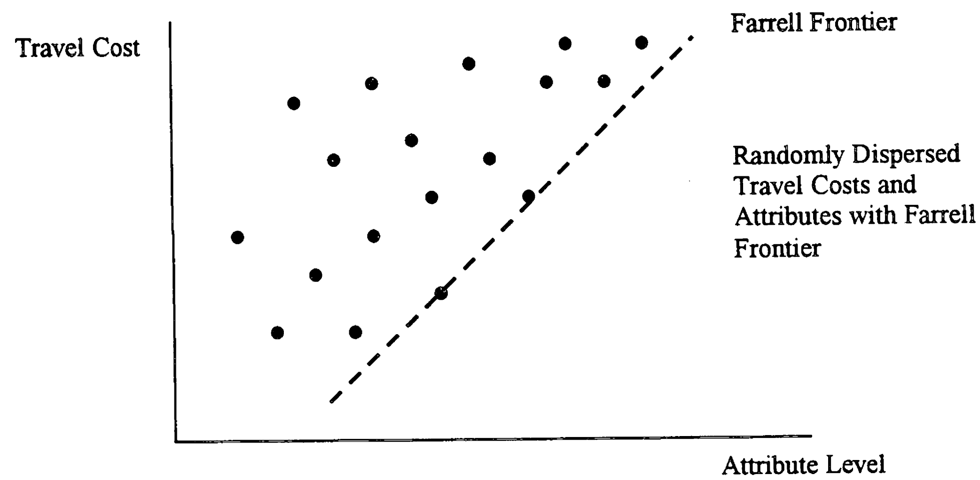

In an attempt to make sense of the hedonic travel cost model, Smith, Palmquist and Jakus make explicit the need for a further constraint on the choice set and use information on technology (cost and attribute information) to identify the set. In taking a technical approach to the formulation of a frontier, these authors avoid confounding preferences and costs. They set up a linear programming problem to define a Farrell frontier. Let Z_{ij} be the quantity of the j^{th} ($j=1,\ldots,n$) characteristic at the i^{th} site and let t_i be the travel cost to the i^{th} site. Then the frontier problem as posed by Smith, Palmquist and Jakus is to choose the costs c_0, c_1, $\ldots c_n$, in order to

$$\text{Minimize } \Sigma_i(t_i - c_0 - \Sigma_j c_j Z_{ij})$$

subject to

$$t_i - c_0 - \Sigma_j c_j Z_{ij} \geq 0$$

$$c_j \geq 0.$$

This problem is solved for each recreationist, and the solution gives the c_j's, which are treated as parametric marginal hedonic prices. In effect, the approach finds an n-dimensional plane that lies on or below the cost-attribute points such that the sum of the distances between the plane and the points is minimized, given that the slope of the plane in any z direction must be positive. (The dotted line in Figure 1 illustrates where this frontier might lie for the one attribute case). Although there are other drawbacks to this formulation, positive marginal prices are assured.

In contrast, the more common approach taken in the hedonic travel cost literature does not concern itself with defining a frontier. Instead the basis of the theoretical justification is that individuals will adjust their choice of attributes until their marginal costs equal their marginal values, and so the individual's actual choice must exhibit this property. In the words of Englin and Mendelsohn:

> *Users, in their choice of site, continue purchasing higher and higher quality until the value of another unit of quality is just equal to its cost. Although the cost function C(Z) is an arbitrary function, in the case of recreation depending solely on residential location, the marginal cost actually chosen by the consumer $C_z(Z)$ reflects her marginal value per trip (Englin and Mendelsohn, p 277).*

The authors argue that while we may not know anything about the configuration of the choice set, it must be true that wherever an individual ends up, at that site his marginal value for each attribute will be equal to his marginal cost of obtaining it.

In practice, Englin and Mendelsohn group individuals by origin zone and estimate separate hedonic cost functions for each zone, using only those sites that are actually chosen by recreationists from that zone. Because only the actual choices are included, one would expect that no combinations of sites that would imply negative prices ought to arise. A negative price would imply that more of an attribute could be obtained at a lower price, and no one who values these attributes would choose an inferior bundle at a higher price.

This method of circumscribing the alternative set does not, however, preclude negative prices from emerging in the estimation. When negative marginal prices arise in their empirical application, Englin and Mendelsohn attribute them to "over-satiation". Non-satiation is not an essential property of preference functions, and satiation in some characteristics is certainly plausible. To use an example from their paper, an individual might value *some* Douglas fir along a hiking trail, but might actually value a mix of tree species rather than *all* Douglas fir – implying that preferences are satiated in the amount of Douglas fir along a hiking trail. If, in addition, the cost frontier is continuous and declining in this attribute, then it is possible for the individual to optimize at a point where the value for more Douglas fir is declining.

The appeal to satiation as an explanation of negative prices is an unsatisfying story, however. It is true that satiation is a condition of preferences that is consistent with negative prices, although negative prices require not only satiation but the right configuration of sites and attributes. However, many econometric issues that could easily lead to negative prices as well, and these can not be ruled out. Perhaps of greater concern, though, is that if we admit of the possibility of satiation, the underlying story of the hedonic travel cost model further unravels, as we shall see.

Re-Examining the Underlying Theoretical Story

For the underlying story of the hedonic travel cost model to make sense, the relationship between travel costs and attributes of the subset of sites in the individual's constraint set must be reasonably approximated by a continuous, differentiable function, C(Z), with proper curvature. These sorts of conditions are common in economics, and in most circumstances they can be believed at least approximately. Usually some mechanism exists, be it in the nature of the production process, the rationality of preferences, or the mechanism of the market, that causes functions of interest to be more or less well-behaved. But in the hedonic travel cost setting there is no mechanism to regularize C(Z), so we have no guarantee that even an approximate function with useful necessary properties exists. Two issues are relevant here: Is it necessary that at the chosen site, marginal value equals marginal cost? And, assuming that marginal value *does* equal marginal cost, can we reliably estimate marginal costs and use this information to define demand curves for attributes?

For convenience, we return to the case of a single attribute, recognizing that the necessary conditions become even more tenuous with increasing numbers of characteristics. We view the problem from the perspective of a group of individuals residing at the same origin who face an array of recreational sites, where nature has endowed the geographically scattered sites with random amounts of the attribute. The cost-characteristic level pairs for these hypothetical sites are depicted in Figure 2.

First consider the case where we do not allow satiation. Preferences for characteristics are monotonically increasing over the entire range. In defining the choice set for the individual in Figure 2, the only sites that we can definitely eliminate are those that lie northwest of another available choice. The combinations lying northwest cost at least as much but provide no more Z and would never be chosen when individuals have positive marginal values for the attribute. If we connect the points in this constraint set (the non-dominated choices) and call it C(Z), then all we know for certain about C(Z) is that it must be upward sloping.

Figure 2

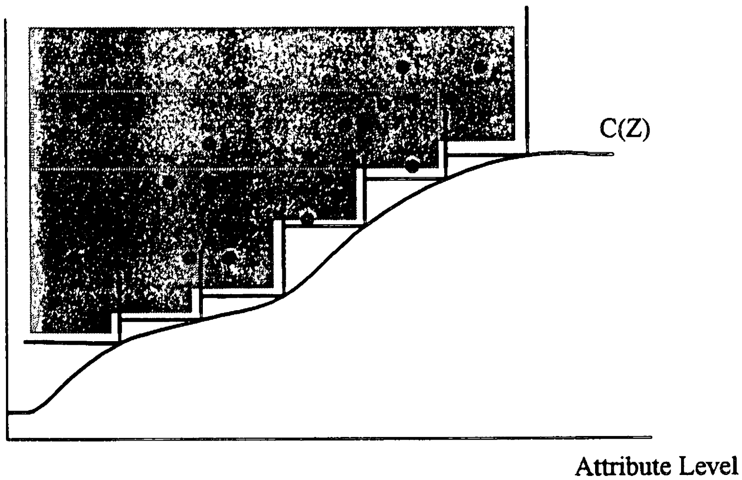

The assumption that the slope of this fitted function at the chosen site actually equals the individual's marginal value for Z becomes more tenuous the sparser and/or more irregular the frontier or constraint set. Unless a reasonably complete and well-shaped constraint set exists, the individual's marginal value may deviate substantially from the slope of the fitted function. This is most obviously true when sites are relatively few or, equivalently, when groups of sites are quite similar in characteristics. If the individual can not marginally adjust his chosen level of Z but must decide among a finite set of discrete choices - then his optimum will have the nature of a corner solution; it will be characterized by a set of inequalities, not the equality of marginal costs and values. Expanding to the case of more than one characteristic, if sets of characteristics tend to be bundled such that one can not be chosen independently of another, again marginal conditions will not likely hold.

The approach can fail even if the individual does equate his marginal value with marginal cost, however. If the constraint set is much more irregular than the fitted functional form, large errors will arise from imputing marginal values from the slope of the estimated hedonic travel cost function. In a hypothetical example given in Table 1, all sites are in the dominant set in that no site costs more than another that embodies at least as much of every attribute. Any one of these sites could be chosen by an individual whose utility was increasing in the Z's over their entire range. Given a diverse set of preferences among those in the origin group – each site *could* be chosen by someone. If so, then all twelve sites would appear in the hedonic travel cost regression using the Englin and Mendelsohn approach.[4] However, this set describes a highly non-linear surface; and fitting a linear regression, for example, would yield estimated negative marginal prices as shown in the results reported at the bottom of the table.

Table 1: A negative price with no dominated sites.

Site number	Travel Costs	z_1	z_2	z_3
1	10	15	15	15
2	15	16	14	15
3	20	17	13	15
4	25	18	12	15
5	30	20	20	20
6	35	21	17	20
7	36	22	15	20
8	38	23	14	20
9	10	14	14	20
10	15	14	11	21
11	20	14	9	22
12	25	14	7	23

$$\text{Travel costs} = -36.44 + 3.11z_1 - .847z_2 + .907z_3$$
(7.3) (10.1) (2.68) (3.12)

If we admit the possibility of satiation, the hedonic travel cost method becomes even more problematic. Now the only requirement on the constraint set is that no site be included that would involve paying more than necessary for any level of Z. But if the individual is allowed to become "over-satiated", then we can not further constrain the technology set, *a priori*, since we can not rule out ranges in which less of the attribute costs more. Any of the sites marked with an X in Figure 3 could be chosen by someone, depending on their preferences for Z relative to the numeraire and depending on the level, if any, at which they become satiated in Z. No meaningful cost function can be fit to these points.

[4] Diversity of preferences appears to contribute to the problem described here, but similarity of preferences causes its own problems. If individuals from the same origin have similar preferences, only a few sites will

Figure 3

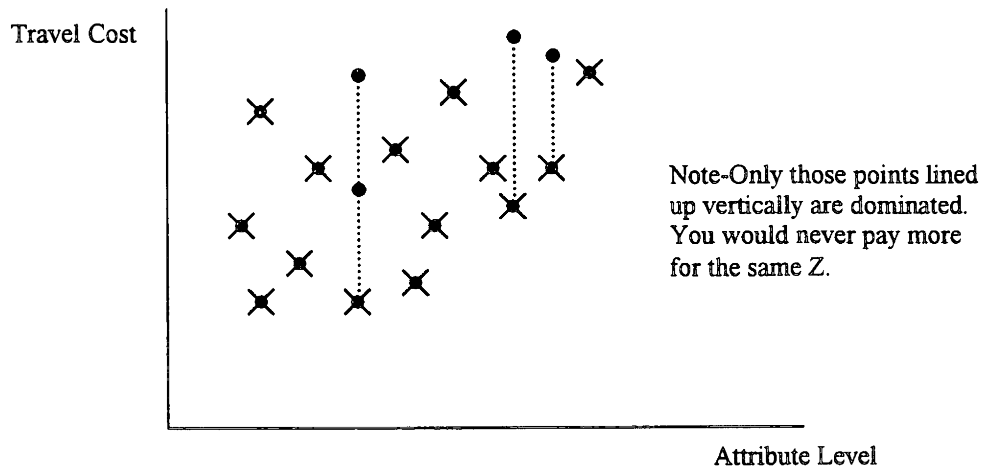

Even if nature arranges sites of different quality in some regular pattern, we no longer have any assurance of a well-behaved problem. Consider the two cost functions in Figure 4.

Figure 4

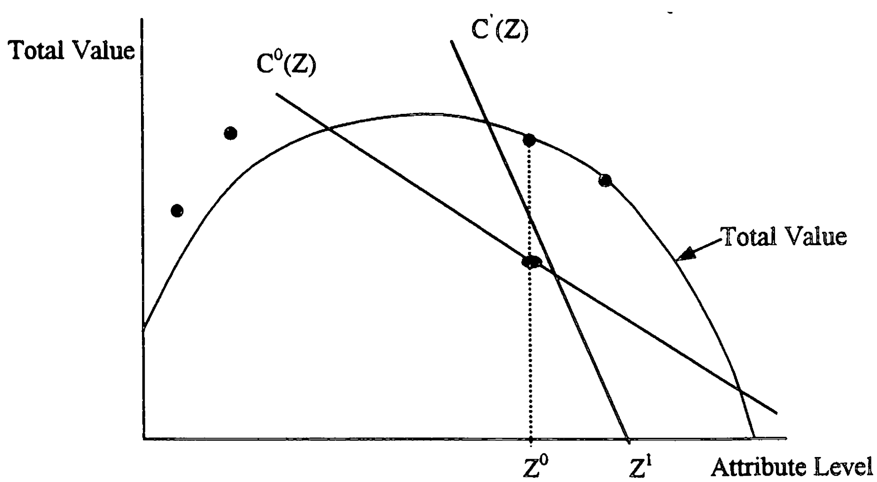

For $C^0(Z)$, the attribute level Z^0 is optimal. But if $C^1(Z)$ is the cost function, then attribute level Z^1 is optimal, and at that point marginal cost and marginal value are not equal. Second order conditions do not hold where marginal conditions are equated. Since the technology set boundary can now have any shape and total value functions can be non-

---

be chosen and, then, even fewer travel cost/attribute points will contribute to the estimation of the hedonic monotonic, it is easy to produce patterns of value functions and boundaries (or "cost functions") such that the optimum choice is at a site where the marginal conditions are not satisfied.

Discontinuities and non-convexities wreak havoc in economics wherever they arise. There are at least two reasons why they need to be taken particularly seriously here. First, there is no basis on which to argue away their existence. Unlike most examples in economics, we have almost no *priors* on the shape of this frontier and there are no inherent forces at work that would regularize the relationship between costs and characteristics. Second, the slope of this elusive cost function ends up playing the momentous role of price in the second stage demand estimation. If the problem is such that individuals do not really equate marginal value and marginal cost then the second stage of the analysis does not make sense. If they do equate these marginal terms, but our estimate of the irregular relationship between costs and characteristics is a poor approximation of the true one, then price in the second stage will be measured with considerable inaccuracy.

One final feature of the model exacerbates the problem. In practice the hedonic travel cost function is usually estimated as a *linear* function in attributes, not because this best fits the frontier, but because two problems arise if the cost function is not linear in attributes. The marginal prices are no longer exogenous to the individual in the first stage of estimation, and "prices" in the second stage will not be parameters, making it difficult to define conventional demand functions. But fitting a linear-in-attributes function to the complex and irregular shape that the cost frontier is likely to exhibit will yield estimated attribute prices that bear no resemblance to actual marginal costs. In fact, as we saw in the example in Table 1, there is no guarantee that even the signs of these marginal prices can be trusted.

Negative marginal prices are the Achilles' heal of the hedonic travel cost model. To justify them by an appeal to satiation undermines the conceptual support of the model. In the absence of a justification of negative marginal prices, however, one must either admit that a crucial dependent variable has been seriously misestimated or that the underlying story about individual behavior is simply untenable.

Welfare Measurement Using the Hedonic Travel Cost Model

The hedonic travel cost approach attempts to map site access costs and the distribution of attributes into a cost function for attributes, such that demand for attributes, rather than demand for sites, can be estimated. Let us abstract from the difficulties in the previous section and assume that the condition that marginal value equals marginal cost *does* characterize the

---

cost function.

individual's site choice. Let us also assume that we have accurately recovered these marginal costs and, using this information, we have successfully estimated demands for attributes. This last step is not trivial, by the way, but is potentially easier than in the hedonic property value model because the availability of data for several origins can potentially provide the variation necessary to identify the parameters in the demand function.[5] Given that all of these steps can be successfully accomplished, how can the results be used for environmental valuation?

Even if the hedonic travel cost analysis allows the recovery of marginal values for the environmental amenity, it can do little more. Both policy evaluation and damage assessment, as distinct from more abstract questions of resource allocation, generally require the measurement of welfare changes associated with discrete changes in environmental quality at specific locations. To use the estimated attribute demand function to analyze such changes, we need to translate the discrete changes at given sites into initial and final positions on the environmental quality demand function.

The logic of the story is analogous to that in the hedonic property value literature. Policy changes or accidents alter the distribution of attribute levels over sites, which in turn changes the cost function for attributes. With a new cost function, individuals may or may not choose a different site, and at this optimum the individual's ambient level of the relevant characteristic may (or may not) change. Presumably if we knew the initial and subsequent ambient levels of a characteristic, then we could approximate the welfare change by the change in the area behind the demand function for the characteristic between the two equilibrium positions.

The difficulty of this approach to valuation is in finding the individual's subsequent choice. We have no *a priori* means of doing so, unless the hedonic travel cost model can provide us with this prediction. In order to use the model to predict the subsequent choice, we must first translate the policy change into a change in the cost function. With a new estimated cost function, the subsequent ambient level of the characteristic would be determined by the intersection of the demand and new marginal cost functions (assuming, of course, the availability of a continuous array of characteristic levels).

There is one case in which the apparent necessity of obtaining the new cost function is circumvented. With knowledge of the characteristic demand equation only, we could value the *complete* elimination of an attribute at *all* sites. Suppose the attribute in question was the amount of old growth forest encountered along all hiking trails in a choice set. If we had a demand function for the attribute – old growth forest, we would be able to measure the losses associated

[5] We repeat, however, that if the hedonic cost function is not linear, recovering conventional demand functions becomes problematic.

with elimination of old growth forest at all recreational sites as the area behind the demand curve above the marginal price line at the initial equilibrium quantity. This is because we implicitly know what we need to know about the new cost function. There is no price at which any of the characteristic can be purchased.

In the general case, however, the policy change must be translated into a change in the hedonic cost function; but this is very difficult to do. It involves defining a new functional relationship between travel costs and site characteristics. We can not estimate the new cost function by the conventional means because that requires first restricting the set of sites (and their associated costs and qualities) to those sites actually chosen, and we can not know the chosen sites until we estimate the cost function. This procedure precludes the possibility of using the hedonic travel cost approach to predict the effects of policies or accidents that alter the cost of acquiring a characteristic. And since it is solely through the cost function for characteristics that the individual can be affected by exogenous changes, this apparently precludes the use of the approach for anything but *ex post* welfare analysis.[6]

Conclusions

On the surface, it might appear that the attractiveness of the hedonic travel cost model turns on the completeness of the recreationist's choice set. Where there are relatively few sites and large gaps in the map of available attributes, it would appear more difficult to approximate the relationship between the travel cost and attribute level with a continuous function and impossible to measure the welfare change. As the number of sites increases and the choice of attribute levels looks more continuous, the hedonic travel cost model might seem increasingly appropriate. This is the usual argument in the housing literature, where Rosen's traditional hedonic model assumes approximately continuous levels of characteristics.[7] The continuous model is most natural for economists who are used to envisioning individuals acting as though they use calculus to solve their decision problems when faced with a continuum of alternatives. When faced with a finite, discrete array, however, we characterize the choice process using discrete choice models in which the decision is viewed as a simple comparison of utility evaluated at alternative options.

---

[6] Englin and Mendelsohn recognize the need for identifying the chosen sites before and after the change (p. 278), but they do not suggest any way in which the hedonic travel cost model could be used to predict changes in site choice.

[7] Discrete analogues of these models have been explored, however. See Cropper, Deck, Kishor and McConnell [3].

Unlike the hedonic property value model, however, even a relatively complete alternative set is not a sufficient condition for marginal value and marginal cost to be equated at the optimum in the hedonic travel cost problem. And even where it is, there is no guarantee that the relationship will be sufficiently regular to be approximated by an estimated function. This is because no mechanism exists to ensure that even the boundary of the "technology" set is a well-behaved function with useful curvature. Since we have no *a priori* information on the shape of the boundary of the technology set, we have no hope of estimating reliably an approximation of marginal cost. This is especially the case when satiation of preferences is permitted, i.e. when preferences can not be assumed to be monotonically increasing in each characteristic over its entire range.

In using the hedonic travel cost model for welfare assessment, the construction of the model dictates that the policy intervention work through the cost function. This limits the model's applicability, because actual policies work by changing the attribute levels at a set of sites, and there is no obvious way in which this change in attribute levels at various sites can be transformed into a change in the hedonic travel cost function. An exception is the situation where the attribute is removed completely, because in this case the cost of the attribute becomes infinite. If one could be satisfied with the estimation of the characteristic demand functions, this would appear to be a useful application of the hedonic travel cost model. But in the general case, a policy change will cause a change in the cost function and a change in the equilibrium. Unless we can translate the policy change into a change in the cost function or produce an independent means of predicting the new site choice, *ex ante* welfare predictions are precluded.

References

1. N. Bockstael, W. M. Hanemann, and C. Kling, Modelling recreational demand in a multiple site framework, *Water Resources Research*, **23**, 951-960 (1987).

2. G. Brown and R. Mendelsohn. The hedonic travel cost method, *Review of Economics and Statistics*, **66**, 427-433 (1984).

3. M. Cropper, L. Deck, N. Kishor, and K. McConnell, Estimating the Value of Attributes with Single Market Data, *Review of Economics and Statistics*, **75**, 225-232 (1993).

4. J. Englin and R. Mendelsohn, A hedonic travel cost analysis for valuation of multiple components of site quality, *Journal of Environmental Economics and Management*, **21**, 275-290 (1991).

5. Y. Kaoru, V. K. Smith, and J.-L. Liu, Using random utility models to estimate the recreational value of estuarine resources, *American Journal of Agricultural Economics*, **77**, 141-51 (1995).

6. E. Morey, The choice of ski areas: estimation of a CES preference ordering with characteristics, *Review of Economics and Statistics*, **66**, 584-590 (1984).

7. G. Phaneuf, C. Kling, and J. Herriges, Estimation and welfare calculations in a generalized corner solution model with an application to recreation demand, manuscript, Iowa State University (1997).

8. S. Rosen, Hedonic prices and implicit markets: product differentiation in pure competition, *Journal of Political Economy*, **82**, 34-55 (1974).

9. V. K. Smith and Y. Kaoru, The hedonic travel cost: a view from the trenches, *Land Economics*, **63**, 180-192 (1987).

10. V. K. Smith, R. Palmquist, and P.Jakus, Combining Farrell Frontier and hedonic travel cost models for valuing estuarine quality, *Review of Economics and Statistics*, **64**, 694-699 (1991).