BRIEF SUMMARY

Introduction

The terms of reference for this project concern the application of hedonic pricing techniques to the valuation of noise pollution. The tasks described in that document can be summarised as follows;

1. To define the theoretical meaning of the willingness to pay values quoted in existing hedonic pricing studies (an appendix to the terms of reference reports a large sample of such values);
2. To make clear why theory indicates that these values will be different for different studies;
3. To define the major factors contributing to these differences;
4. To assess whether these values represent comprehensive measures of the economic welfare changes associated with changes in exposure to noise pollution;
5. To describe how such measures might be derived, and
6. To advise on whether there is a theoretical basis for a single willingness to pay value for avoidance of noise pollution which can be applied across the EU.

The details of the desk-based research addressing these issues are provided in the main report. The main report represents a comprehensive review of current thinking on the theoretical valuation of environmental goods in hedonic markets.[1] It consists of three chapters;

- Chapter 1 describes the theory of hedonic property markets;
- Chapter 2 describes how measures of welfare change resulting from changes in a housing attribute (e.g. exposure to noise pollution) might theoretically be determined in a hedonic market, and
- Chapter 3 describes the process whereby data from hedonic markets can be used to derive empirical estimates of these welfare measures.

Necessarily, the main report presents a large amount of theoretical economic material. However, in an attempt to aid understanding and accessibility, where possible, arguments have been presented diagrammatically rather than mathematically.

All the same, those with little economic training or those with little interest in the theoretical niceties may wish to focus their attention on the summary document. That document provides a shorter and more digestible version of the main report referencing the longer document where necessary and concluding on the issues described above.

This brief summary draws together the main conclusions of the research project in one place.

---

[1] In accordance with the terms of reference, purely empirical issues, such as how to measure noise exposure, are not dealt with here.

Property Markets and the Hedonic Price Function

Property is an example of a differentiated good. Such goods consist of a diversity of products that, while differing in a variety of characteristics, are so closely related in consumers' minds that they are considered as being one commodity. The price paid for a property in the property market will be determined by the particular qualities or characteristics of its structure, environs and location. Amongst these characteristics we would include the environmental quality at that particular residential location. Thus we would expect that properties in areas suffering high exposure to noise pollution will command lower prices than similar properties in peaceful locations.

As with any market, the prices that are paid in a particular property market are determined by the interacting forces of supply and demand. The market will settle on a set of prices for the numerous varieties of the differentiated good that reconcile supply with demand and clear the market. The schedule of prices determined by market forces in a particular market can be summarised by a ***hedonic price function***. This function **describes how the quantity and quality of a property's characteristics determine its price in that particular market**.

The hedonic price function for a particular property market will reflect many factors including the characteristics of the households and the availability of property characteristics. For example, we would expect properties in peaceful locations to command relatively higher prices in a generally noisy urban area than equivalent properties in a generally peaceful urban area.

As a result, **the hedonic price function for any particular property market will be unique to that market reflecting the specific conditions of supply and demand that exist at that locality**.

Implicit Prices for Property Characteristics

The hedonic price function can be used to determine how much more must be paid for a property with an each extra unit of a particular housing characteristic. This is known as the *implicit price* of a property characteristic; *implicit* because the marginal price of a characteristic is revealed to us indirectly through the amounts households are prepared to pay for the whole property of which the particular characteristic is only a part.

In hedonic markets, the price paid for extra of a characteristic may depend in part on the level of provision of that characteristic. For example, the implicit price of extra "peace and quiet" may be high if a property is in a very noisy area and relatively low if the property is in a peaceful area. Indeed, the hedonic price function can be used to determine the ***implicit price function*** which **describes the amount paid for extra of a property characteristic as a function of the level of provision of that characteristic and the level of provision of other property characteristics.**

Frequently researchers use a very simple functional form when using empirical data to estimate the hedonic price function. Typically the natural log of house price is regressed against a linear specification of the housing attributes. In this case the implicit price function for a housing attribute such as exposure to noise pollution can be represented by one figure; the percentage change in the house price brought about by a unit change in traffic noise. This is the Noise Depreciation Sensitivity Index (NDSI) measure that dominates the hedonic price studies reported in the appendix to the terms of reference. When the functional form is more complex, researchers will frequently report a less revealing summary statistic; the implicit price evaluated at the mean level of that housing attribute in the property market under study.

The values from hedonic price studies, therefore, are summaries of the implicit price of noise in a particular property market. However, remember that the hedonic price function in any property market will depend upon the particular conditions of supply and demand existing in that market. **There is no theoretical reason, therefore, to expect the summary values of the implicit price function for noise reported in hedonic analyses of different property markets to return the same value**. Indeed, we would expect them to return different values.

Welfare Measures from Hedonic Markets; Marginal Changes

Basic economic theory suggests that households possess *demand curves* for each of the characteristics of a property. **Each demand curve traces out how much the household is willing to pay for an extra unit of a housing characteristic** enjoyed at their chosen property.

The household chooses the optimal level of housing attributes by purchasing a property at which their willingness to pay for extra of a particular characteristic is equal to the amount they must pay for it in the property market. That is, they will **choose a quantity of each housing characteristic at which their demand curve for that characteristic intersects its implicit price function**.

The household will always wish to purchase properties with up to this optimal quantity of the characteristic since their willingness to pay for each of these units is greater than the price of those units. Conversely, the household would not wish to purchase a property with more of the attribute than this optimal quantity, since the price that must be paid for each unit in excess is greater than the household's willingness to pay for those units.

The important thing to note is that at the household's optimal choice, the household's willingness to pay for extra characteristic is exactly equal to the implicit price of that characteristic in the market.

In general, we could assume that **each point on the implicit price function represents an intersection with the demand curve of a particular household**. As a result, at every level of the housing attribute **the implicit price function will also give the willingness to pay of a household in the property market for extra attribute**. Consequently, the implicit price function allows researchers to determine the welfare impact of marginal changes in a housing attribute. Of course, **since the implicit price function will be different for each property market such welfare estimates are market specific**.

Welfare Measures from Hedonic Markets; Non-Marginal Changes

Unfortunately, the changes in which policy makers are interested are unlikely to be marginal. The construction of a new road through a residential area, for example, is unlikely to cause a unit change in road traffic noise and will most likely impact on a large number of households.

Focussing on the welfare impacts of such a project on the households directly effected by the change, it is simple to show that **welfare calculations based on the implicit price function are inaccurate**. In addition, since they are based on the unique implicit price function estimated for a particular market, there is **no theoretical justification for transferring them across property markets**.

Accurate welfare measures for non-marginal changes should be calculated from the demand curve. The demand curve shows the household's willingness to pay for each unit of a housing attribute. To value the welfare impact of a non-marginal change in provision of a housing attribute, we would wish to sum these willingness to pays for each unit of the attribute lost or gained. This is the ***Quantity Compensating Surplus* (*QCS*) of a welfare change. It is defined as the area under the household's demand curve between the current level of provision of the attribute and that experienced after the change.**

Further, under the assumption that preferences are stable across geographical regions, demand functions can be transferred across markets. For example, imagine that we had estimated the household demand function for environmental quality (e.g. peace and quiet). Using information on the present levels of environmental quality, the expected changes in this quality and the characteristics of the households impacted by this change, the demand function could be used to derive *QCS* measures of welfare change in any geographical area.

However, even **the *QCS* measure of welfare change is not comprehensive**. **A more comprehensive measure is that of *Total Social Benefits* (*TSB*)**. *TSB* includes benefits accruing to both households and landlords. It also accounts for changes in the hedonic price function brought about by a change in environmental quality and the responses of households and landlords to these changes. Even the *TSB* measure does not measure the benefits of an environmental improvement enjoyed by those that visit or work in the improved area.

Unfortunately the *TSB* measure requires information that is hardly ever available to researchers. In general, such complete welfare measures will only be possible *ex-post*, when researchers have information on the hedonic price function before and after the change.

Nevertheless, it can be shown that **the *QCS*, when summed over all households directly affected by the change in environmental quality, will give a lower bound to the *TSB*** of that change.

Estimating Demand Functions from Hedonic Market Data

Demand functions for environmental quality cannot be estimated from data collected in a single hedonic property market without the imposition of untestable assumptions concerning the nature of household preferences. Rather **estimation of demand functions requires data from several hedonic markets**.

Demand estimation is further complicated by the fact that marginal prices in hedonic markets are not necessarily constant; that is, the implicit price of a characteristic may vary across the range of provision of that characteristic. Whilst this complicates the procedures, it does not make the estimation of demand functions impossible. Indeed, Table 4 of the main report describes the steps that must be undertaken to overcome the problems caused by non-constant marginal prices in order to estimate demand functions for environmental quality.

Importantly, **estimated demand functions can be used as a means of transferring benefits across geographical regions**. Such transfers necessarily involve simplifications and approximations. In addition, the validity of such exercises depends on the assumption that preferences for environmental quality are stable across geographical regions. **Future research should focus on testing the accuracy of welfare measures estimated by benefits transfer**.

SUMMARY

1. The Property Market

Housing is an example of what in economics is termed a *differentiated good*. Such goods consist of a diversity of products that, while differing in a variety of characteristics, are so closely related in consumers' minds that they are considered as being one commodity. Many other goods, including breakfast cereals, cars, computers and beach holidays also fit this description.

Housing is traded in the property market. Market forces determine that different varieties of the product command different prices and that these prices depend on the individual products' exact characteristics. For example, properties that have more bedrooms will tend to command a higher price in the market than properties that have fewer bedrooms. Furthermore, the set of prices in the market define a competitive equilibrium. That is, in general, the market will settle on a set of prices for the numerous varieties of the differentiated good that reconcile supply with demand and clear the market.

As a consequence of the fundamentally spatial nature of property, property markets are themselves defined spatially. Thus at any point in time, all of the properties in one urban area represent the products in a property market. The *households* wishing to live in these properties represent the consumers in this market and they determine the level of demand in the market. The *landlords* that own the properties represent the producers in this market and consequently determine the level of supply.

We could describe any particular property by the qualities or characteristics of its structure, environs and location. A succinct means of denoting this is as a vector of values; effectively a list of the different quantities of each characteristic of the property. In general, therefore, any house could be described by the vector,

$$\mathbf{z} = (z_1, z_2, \ldots, z_K), \qquad \text{(E1)}$$

where z_i ($i = 1 \text{ to } K$) is the level or amount of any one of the many characteristics describing a property. Indeed, the vector $\mathbf{z}$ completely describes the services provided by the property to a household.

For the sake of simplicity let us assume that each of the z_i are measured in such a way that we can consider them as "goods" as opposed to "bads". For example, one of the characteristics of a property will be its exposure to road noise. Rather than measuring this as the level of "noise", we can simply invert the scale and measure it as the level of "peace and quiet".

When households select a particular property in a particular location they are selecting a particular set of values for each of the z_i. We can imagine this market for properties as being one in which the consumers consider a variety of somewhat dissimilar products which differ from each other in a number of characteristics including, amongst many characteristics, number of rooms, size of garden, distance to shops and environmental characteristics such as levels of pollution or noise. Using an analogy of Freeman (1993 p 371), "it is as if the urban area were one huge supermarket offering a wide selection of varieties. Of course, the individuals cannot move their shopping carts through this supermarket. Rather, their selections of residential locations fix for them the whole bundle of housing services. It is much as if shoppers were forced to make their choices from an array of already filled shopping carts. Individuals can increase the quantity of any characteristic by finding an alternative location alike in all other aspects but offering more of the desired characteristic."

2. The Hedonic Price Function

The price of any one of these 'shopping carts' will be determined by the particular combination of characteristics it displays. Naturally we would expect properties possessing larger quantities of good qualities to command higher prices and those with larger quantities of bad qualities to command lower prices. Again we can use a succinct piece of notation to illustrate this point;

$$P = P(\boldsymbol{z}) \qquad \text{(E2)}$$

Which can be read as; the price of a property, P, is a function of the vector of values, $\boldsymbol{z}$, describing its characteristics. This function is known as the *hedonic price function*; 'hedonic' because it is determined by the different qualities of the differentiated good and the 'pleasure' (in economic terms utility) these would bring to the purchaser[2].

To illustrate the hedonic price function, consider the illustration in Figure E1. Plotted on the vertical axis is the price of property. Along the horizontal axis is quantity of a particular housing characteristic labelled z_1. For illustrative purposes let us assume that this characteristic is the size of the property's garden. Further, let us introduce some new notation, $\boldsymbol{z}_{-1}$, which is the vector containing the levels of all property characteristics barring z_1. Notice that in the hedonic price function in Figure E1, $\boldsymbol{z}_{-1}$ comes after a semicolon. This indicates that these other characteristics are held constant at some given level whilst the focus characteristic, z_1 (size of garden), changes. Consequently, in this example we are not considering the interaction of different characteristics of the property.

In this hypothetical case, the hedonic price schedule[3] rises from left to right implying that the bigger a property's garden the higher the price that property commands in the market.

---

[2] In the model of the property market presented in Chapter 1 of the main document this price is the rental that a household pays to the landlord. In effect, every household in the urban area is purchasing the flow of services derived from the characteristics of the property per period of time. Of course, many households own their own homes. In this case we treat homeowners as landlords that rent from themselves.

[3] Strictly speaking, the hedonic price function is the formula that dictates the price that a property with particular characteristics will sell for in the market. The set of prices that come out of this formula are frequently referred to as the *hedonic price schedule*. However, in this document the formal distinction between function and schedule is not adhered to and the two terms are used interchangeably.

Notice also that the marginal price of extra garden space is not constant. The slope of the curve becomes progressively flatter and the incremental increase in a property's market price resulting from its possessing a bigger garden declines as gardens get progressively larger. This sort of relationship reflects a form of satiation; having a few square metres of garden will add considerably to the price of a house when compared to a house with no garden at all, whilst a few extra square metres will make a negligible difference between the selling prices of two houses which already boast football pitch-sized gardens.

Figure E1: The Hedonic Price Schedule for characteristic z_1

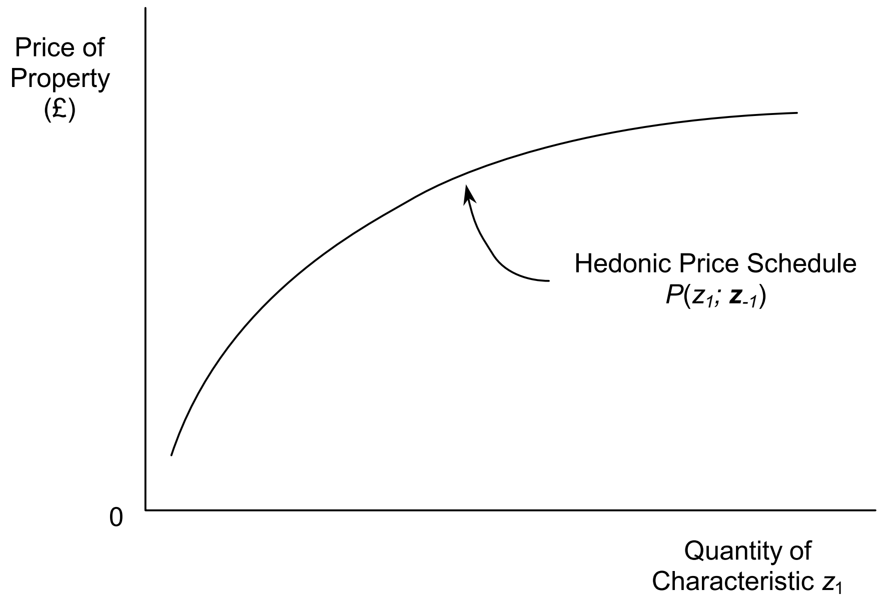

Of course the relationship won't be identical to that graphed for every type of characteristic but this declining marginal price is fairly typical of relationships observed in empirical studies.

It may be easier to illustrate the idea of non-constant marginal prices through actually plotting the additional amount that must be paid by any household to move to a bundle with a higher level of that characteristic, other things being equal. This is illustrated in the right hand panel of Figure E2.

This new function is known as the *implicit price function*; *implicit* because the marginal price of a characteristic is revealed to us indirectly through the amounts households are prepared to pay for the whole property of which the particular characteristic is only a part. From Figure E2, we can see that at first the hedonic price function rises steeply so that the implicit price of the characteristic (the extra amount paid to acquire a house with more of characteristic z_1) is also high. At higher levels of z_1 the hedonic price function is flatter so that the implicit price of the characteristic is also low.

Mathematically, the implicit price is derived as the partial derivative of the hedonic price function (Equation E2) with respect to one of its arguments, z_i, according to:

$$p_{z_i}(z_i; \mathbf{z}_{-i}) = \frac{\partial P(\mathbf{z})}{\partial z_i} \qquad \text{(E3)}$$

To re-emphasise $p_{z_i}(z_i; \mathbf{z}_{-i})$, the marginal price function of characteristic z_i, does not have to be a constant.

Figure E2: The Hedonic Price and the Implicit Price Schedules for characteristic z_1

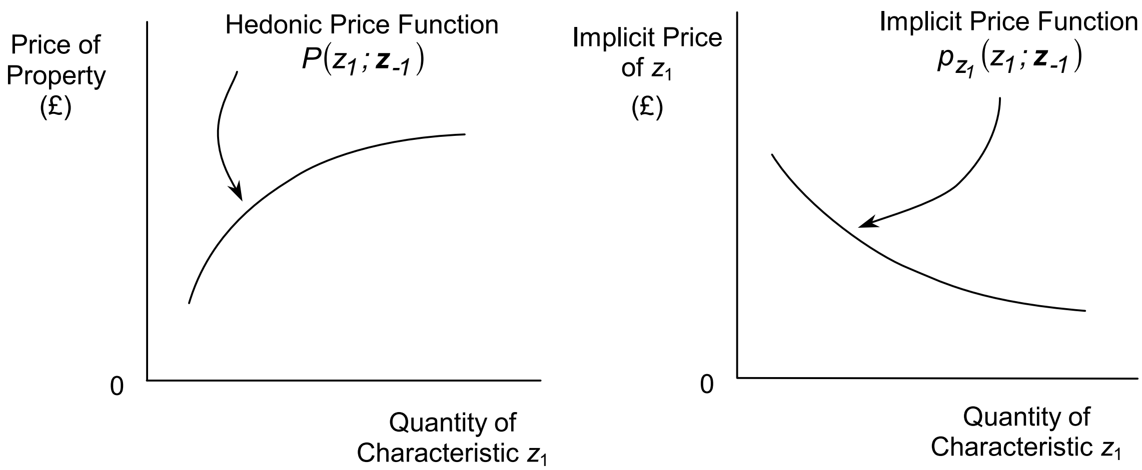

In empirical applications researchers estimate the hedonic price function of Equation (E2) by collecting data on the selling price of houses in a particular property market and regressing these on the characteristics of those properties (i.e. the z_i). To summarise the results of such a regression researchers report the implicit price of the various housing characteristics according to Equation (E3).

Frequently researchers use a very simple functional form for the hedonic price function. Typically the natural log of house price is regressed against a linear specification of the housing attributes. In this case the implicit price function for a housing attribute such as exposure to traffic noise can be represented by one figure; the percentage change in the house price brought about by a unit change in traffic noise. This is the Noise Depreciation Sensitivity Index (NDSI) measure that dominates the hedonic price studies reported in the appendix to the terms of reference. When the functional form is more complex, researchers will frequently report a less revealing summary statistic; the implicit price evaluated at the mean level of that housing attribute in the property market under study.

Conclusion 1: The values from hedonic price studies contained in the appendix to the terms of reference are summaries of the implicit price of noise in a particular property market.

3. Equilibrium in the Hedonic Property Market

The property market is unusual in that it does not return a single price for each unit of attribute boasted by a property; rather it returns a continuum of prices[4]. However, we would still expect this continuum of prices to represent a *market equilibrium*. That is, at the set of prices revealed by the hedonic price schedule, demand would equal supply and the market would clear. Of course this follows basic logic, if a landlord set the rent on his/her property too high then it would remain unsold, conversely if the price were too low then he/she would risk losing out on potential profits.

In the main document the attainment of market equilibrium is explained more formerly as the interaction of households and landlords. The details of this model are beyond the scope of the present discussion, in short, however, households wish to rent the property that provides them with the greatest quality at the lowest price, whilst landlords wish to let their property at the highest price possible. The market reconciles these conflicting goals by matching households to landlords such that the households (within their limited budgets) cannot increase their utility by choosing a different property and the landlords cannot increase their profits by increasing the property's rent or changing its characteristics.

The equilibrium hedonic price schedule settled on in the market, therefore, will reflect many factors. For example, we would expect a property market in which households are generally better off to have a higher willingness to pay for property characteristics. For a property characteristic such as "peace and quiet", whose supply is determined exogenously[5], this will most likely result in generally higher implicit prices. Likewise, on the supply side, the availability of housing attributes will influence the equilibrium hedonic price schedule. Consider, for example, the price paid for waterfront properties in London and Stockholm. Whilst in both cities such properties command considerable premia, the relatively low availability of "Thames-side" properties in London means that they command highly inflated prices compared to those in Stockholm, a city built upon a series of islands.

As a result, the equilibrium hedonic price schedule for any particular housing market will be unique to that market reflecting the specific conditions of supply and demand that exist at that locality.

This is illustrated in Figure E3 where the implicit price function for housing attribute z_1 is shown for two separate markets, Market *A* and Market *B*. As we would expect, the two functions are quite different. Unsurprisingly, if a researcher were to summarise the implicit price functions for these two markets using the NDSI or by evaluating the implicit price function at the mean value of z_1, he would return very different values. Observe Figure E3 where the mean value of z_1 in the two housing markets are given by

[4] In the main document the existence of non-constant marginal prices is explained as the result of an inability to "repackage" the attributes of a property. In other words, households are unable to break up the attributes of any particular property and enjoy each independently of the whole.

[5] That is, landlords can do little if nothing to change the level of traffic noise to which their property is exposed.

$\bar{z}_1^A$ and $\bar{z}_1^B$. Summarising the implicit price function at this point would return two very different values, $\bar{p}_1^A$ and $\bar{p}_1^B$.

Figure E3: Identifying the Marginal Bid Curve

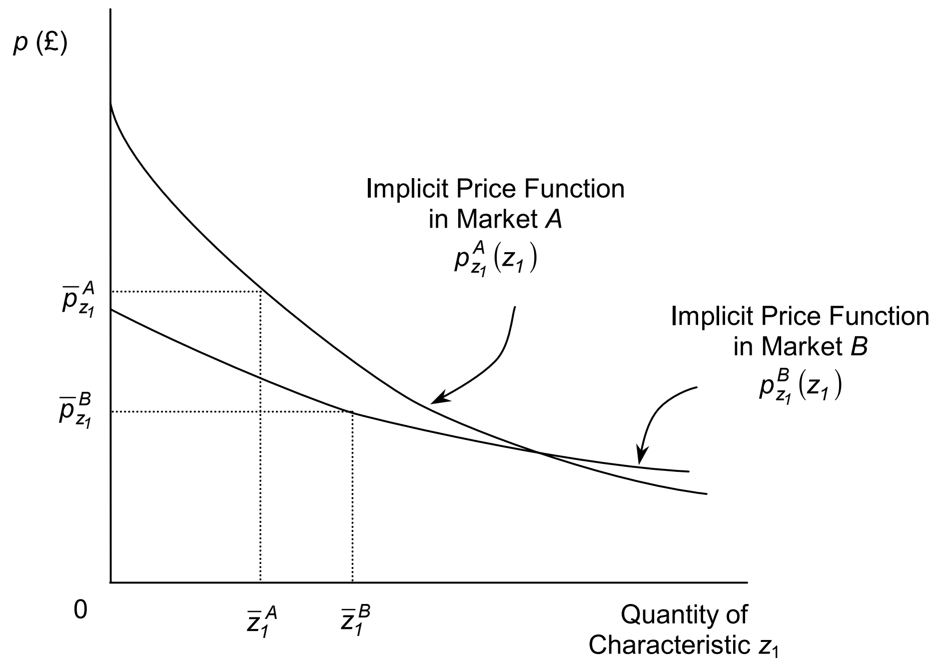

Conclusion 2: The equilibrium hedonic price function in any property market will depend upon the particular conditions of supply and demand existing in that market.

There is no theoretical reason to expect the summary values of the implicit price function for noise reported in hedonic analyses of different property markets to return the same value. Indeed, we would expect them to return different values.

There is no theoretical basis for transferring such values between different hedonic markets.

4. Household Choice in the Property Market

The hedonic price function, $P(z)$, therefore, emerges from the interaction of households (demanders) and landlords (suppliers) and represents a market clearing equilibrium that will be specific to each individual property market. Now, let us focus on how households facing such a hedonic price schedule determine their optimal residential location.

To do this we need to assume that the household has a demand curve for each housing attribute. As we shall see later this is not strictly true but this will not impede our analysis for the time being. An example of such a demand curve[6] is shown in Figure E4. This curve traces out how much a particular household (household *a*) is willing to pay for each extra unit of housing attribute z_1. As we would expect, the demand curve falls from left to right. At low levels of housing attribute z_1 the household has a higher willingness to pay to acquire a property with more of this attribute whilst at high levels the household's willingness to pay for extra attribute is relatively small.

Figure E4: Household Choice of Housing Attributes

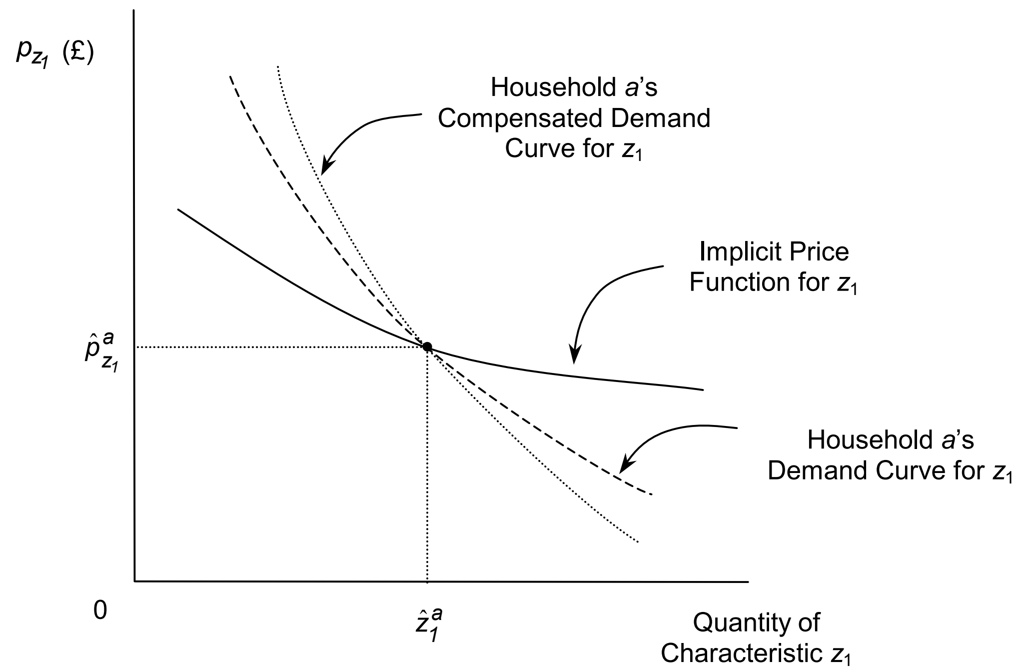

As should be familiar to those who have studied economics, the household faced by the implicit price function in this property market will choose a property with a quantity of z_1 that corresponds to the point where the market price intersects their demand curve. In the diagram this equates to choosing a property with $\hat{z}_1^a$ of the attribute at a marginal implicit price of $\hat{p}_{z_1}^a$. This is very intuitive. The household will always wish to purchase properties with up to $\hat{z}_1^a$ units of the attribute since their willingness to pay for each of these units is greater than the price of those units. Conversely, the household would not wish to purchase a property with more of attribute z_1 than $\hat{z}_1^a$, since the price that must be paid for each unit of z_1 in excess of $\hat{z}_1^a$ is greater than the household's willingness to pay for those units. The optimal level of z_1, therefore, will be found at the intersection of the demand curve and the implicit price function.

[6] Technically speaking an inverse ordinary demand curve

Notice the second curve in Figure E4. This is the household's compensated demand curve (also known as the marginal bid curve). In terms of welfare analysis, it is this function that we would seek to estimate. Whilst the definition of these two curves is covered in some detail in the main report, for the purposes of this document we shall ignore the difference between the compensated and uncompensated demand curves. Rather we shall assume that the uncompensated demand curve is a reasonable approximation to the compensated demand curve.

The important thing to note about this diagram is that at the household's optimal choice, the household's marginal willingness to pay for extra z_1 is exactly equal to the implicit price of z_1 in the market.

Figure E5: Household Choice of Housing Attributes

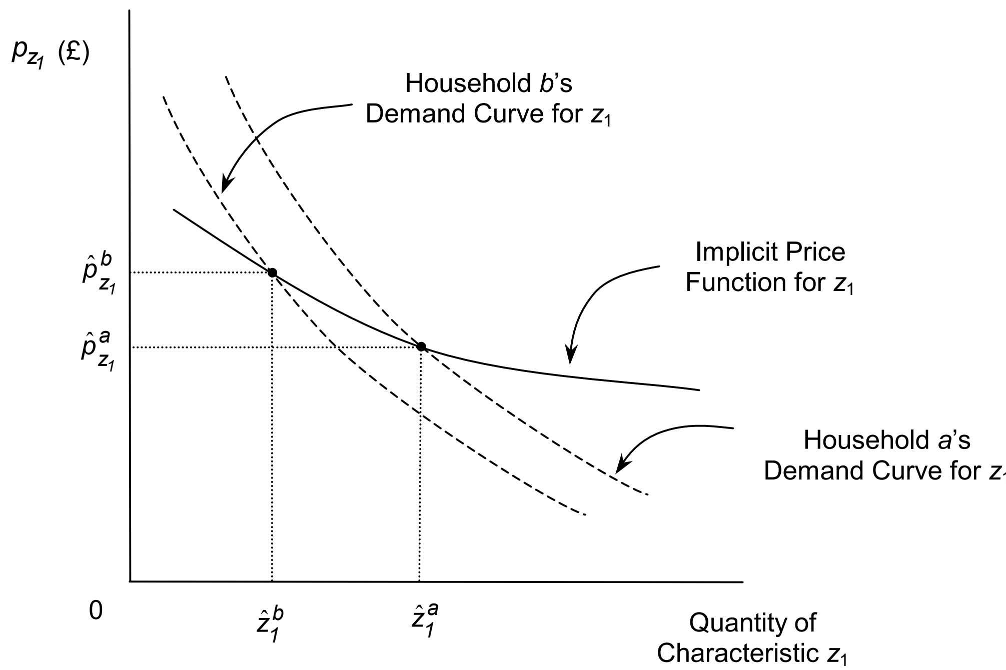

Consider Figure E5. Here the demand curve for a second household, household *b*, has been traced on to the diagram. Notice that they too choose an optimal bundle defined by the point where their demand curve intersects the implicit price function. Consequently, at a level of the housing attribute $\hat{z}_1^b$ the implicit price function will also give the willingness to pay of household *b* for extra attribute.

Indeed, we could continue tracing demand curves for each of the households in the property market onto the figure. Eventually, each point on the implicit price function would be intersected by the demand curve of a particular household. As a result, at every level of the housing attribute the implicit price function will also give the willingness to pay of a household in the property market for extra attribute.

Conclusion 3: The implicit price function for a particular market will also trace out the willingness to pay of households in that market for extra z_1.

The implicit price function, therefore, allows researchers to determine the welfare impact of marginal changes in a housing attribute.

Of course, since the implicit price function will be different for each property market such welfare estimates are market specific.

5. Welfare Measures for Non-Marginal Changes

Unfortunately, the changes in which policy makers are interested are unlikely to be marginal. The construction of a new road through a residential area, for example, is unlikely to cause a unit change in road traffic noise and will most likely impact on a large number of households.

For now, let us focus on the welfare impacts that such a change would have on one household. Figure E6 illustrates the demand function and optimal choice of residential location as chosen by household *a* for attribute z_1. To focus ideas, let us assume that z_1 represents the environmental quality (e.g. peace and quiet) enjoyed at a property.

Figure E6: Household Choice of Housing Attributes

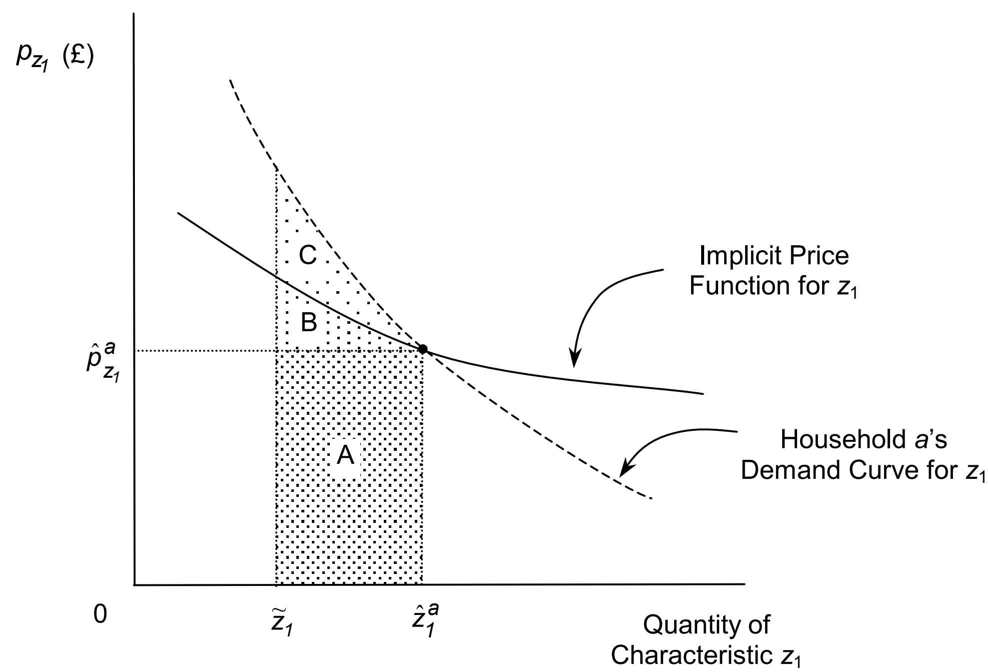

Facing the implicit price function in this market, the household chooses a property with a level of environmental quality $\hat{z}_1^a$ where the implicit price is $\hat{p}^a_{z_1}$. Imagine that an exogenous change in environmental quality, say the opening of a new road, resulted in environmental quality at this location falling to $\tilde{z}_1$. Three possible measures of the welfare change experienced by the household are illustrated in the figure.

- The first amounts to valuing each unit of environmental quality lost at the household's original marginal willingness to pay for environmental quality. This amounts to area A.
- The second involves measuring willingness to pay as the area under the implicit price function between the two levels of environmental quality and amounts to area A + B.
- The third measures the welfare change as the area under the demand curve between the two levels of environmental quality and amounts to areas A + B + C. Palmquist (1988) has labelled this measure the *Quantity Compensating Surplus* (*QCS*)[7].

Of the three measures, the *QCS* is the most correct measure of the welfare change experienced by the household (we shall qualify this statement shortly). Each unit of change in environmental quality is valued at the household's willingness to pay for that unit as traced out by the demand curve.

Notice that the other two measures, based on implicit prices, will underestimate the welfare impacts of a deterioration in environmental quality. Similarly these measures would overestimate the welfare impacts of an improvement in environmental quality.

Further, since these two measures are based solely on the implicit price function estimated for a particular market, they are not transferable across markets. The *QCS* measure, on the other hand, is based on the underlying preferences of households. If we were to make the assumption that households have the same preferences for environmental quality then it would follow that a household in another property market with identical characteristics to household *a* would possess an identical demand curve. The *QCS* measure of a deterioration in environmental quality from $\hat{z}_1^a$ to $\tilde{z}_1$ would be identical for all such households, no matter where they lived.

Indeed, if it were possible to estimate the demand curve for environmental quality as a function of household characteristics, then it would be possible to transfer this function across households and markets. Using information on the present levels of environmental quality, the expected changes in this quality and the characteristics of the households impacted by this change, the transferred function could be used to derive *QCS* measures of welfare change.

Conclusion 4: For non-marginal changes in environmental quality, welfare calculations based on the implicit price function are inaccurate. In addition, since they are based on the unique implicit price function estimated for a particular market, there is no theoretical justification for transferring them across property markets.

[7] More correctly, this measure is the area under the compensated demand curve between the two level of environmental quality.

Welfare calculations based on the demand curve, the quantity compensating surplus (*QCS*), give theoretically justifiable (to be qualified shortly) estimates of the impacts of changes in environmental quality. Moreover, since these estimates are based on underlying preferences they are not specific to a particular market.

Under the assumption that preferences are stable across geographical regions, demand functions can be transferred across markets.

Using information on the present levels of environmental quality, the expected changes in this quality and the characteristics of the households impacted by this change, the transferred function could be used to derive *QCS* measures of welfare change in any geographical area.

6. The Comprehensiveness of the *QCS* Welfare Measure

QCS is not a comprehensive measure of the welfare change resulting from a change in environmental quality. For a start, it only measures the welfare impacts experienced by households. No account is taken of the welfare impacts of the change on landlords (i.e. how landlords' profits change in response to the change in environmental quality).

In addition, the *QCS* measure takes no account of the fact that an exogenous change in the level of environmental quality enjoyed at some (or possible all) locations in the urban area will have the effect of changing supply conditions in the market. Indeed, we might expect that a change in environmental quality in the urban area would precipitate a shift in the hedonic price function. Moreover, the measure does not allow for the fact that the household may react to changes in environmental quality at their residential location and to changes in the hedonic price function by choosing to move to an alternative property. Chapter 3 of the main report describes a comprehensive measure of welfare change, the *Total Social Benefits* (*TSB*), which takes account of all these factors.

However, the *TSB* measure is little more than a theoretical construct. To estimate such a measure researchers would require detailed knowledge of how the equilibrium hedonic price function would be affected by changes in environmental quality and how households' and landlords' choices would respond to both changes in environmental quality and changes in the hedonic price schedule.

Unfortunately, hedonic market equilibria are too complex to derive satisfactory analytical solutions by which to predict such outcomes. Indeed, the *TSB* measure is almost impossible to calculate *ex-ante*, making it of little use to practitioners attempting to measure the potential benefits of a program seeking to change environmental quality in an urban area.

It is worth noting that even this *TSB* measure of welfare change ignores the benefits to visitors that travel through the improved area. Similarly it ignores the benefits to those who work in the improved area. Moreover, the measure ignores the costs (savings) of causing the environmental improvement (deterioration). For example, no account is taken of the cost to the taxpayer of traffic calming schemes designed to reduce traffic noise.[8]

Nevertheless, in an important analysis, Bartik (1988) showed that the *QCS* measure when summed over all households directly affected by the change in environmental quality could always be taken as a lower bound to the *TSB*. For this reason, much of the theoretical work on hedonic analysis has focussed on the task of using data from property markets to estimate demand curves for environmental quality.

Conclusion 5: The *QCS* measure of welfare change is not comprehensive. A more comprehensive measure is that of *Total Social Benefits* (*TSB*). *TSB* includes benefits accruing to both households and landlords. It also accounts for changes in the hedonic price function brought about by a change in environmental quality and the responses of households and landlords to these changes.

Even the *TSB* measure does not measure the benefits of an environmental improvement enjoyed by those that visit or work in the improved area.

Unfortunately the *TSB* measure requires information that is hardly ever available to researchers. In general, such complete welfare measures will only be possible *ex-post*, when researchers have information on the hedonic price function before and after the change.

Nevertheless *QCS*, when summed over all households directly affected by the change in environmental quality, can be shown to be a lower bound to the *TSB* of that change.

7. Estimating Demand Curves using Hedonic Market Data

Bartik's analysis goes some way towards explaining why much of the hedonic literature has focused on the issue of estimating bid curves from empirical data. As shall become evident, however, this is not a straightforward procedure.

Consider Figure E7a which presents our familiar diagram of household *a*'s optimal choice of residential location in Market *A*. In this market the household chooses $\hat{z}_1^A$ of the attribute at a marginal implicit price of $\hat{p}_{z_1}^A$. Observing this behaviour in the market, the researcher records just one point on the demand curve. Unfortunately, knowing one point on the demand curve is not sufficient to define the whole curve. Indeed, as illustrated in Figure E7, any shaped curve is compatible with this one point provided it passes through ($\hat{z}_1^A$, $\hat{p}_{z_1}^A$).

To identify the demand curve we would require further information. Specifically, we would need to know the household's choices of z_1 at alternative prices.

---

[8] Though these costs/savings would usually be estimated from other data as part of a general cost-benefit analysis

Figure E7a: Information on the Demand Curve from one Market

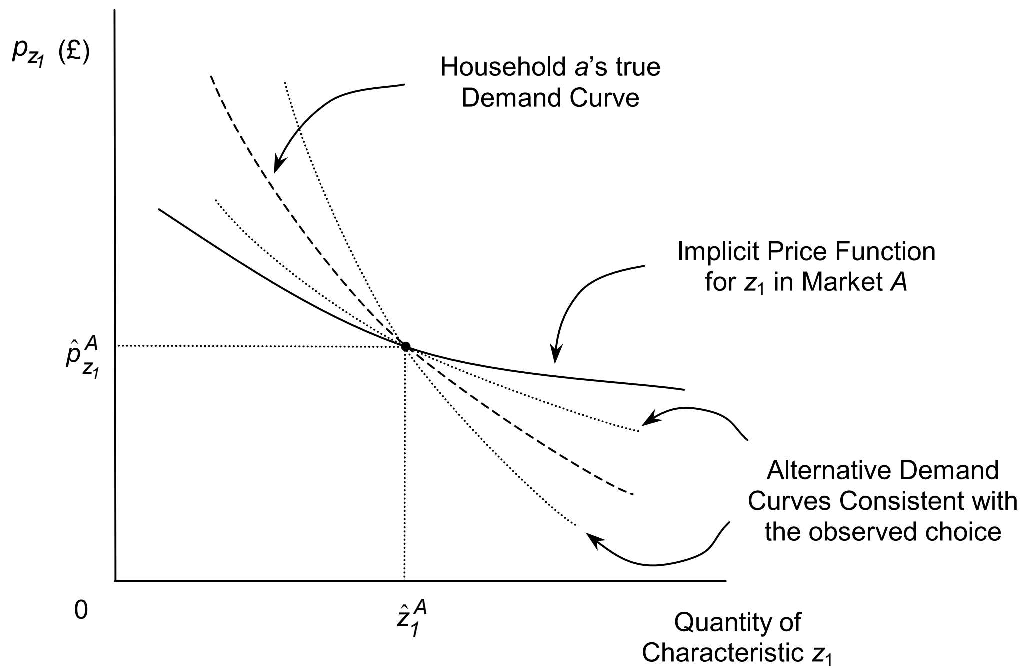

One possibility is that such information could be provided by observing the behaviour of other households in separate markets. If these households happen to have the same income and socioeconomic characteristics as the household choosing in market *A*, then it is assumed that they will have the same preferences and hence the same demand curve. Thus, if these different households were faced by the same hedonic price schedule they would choose the exact same bundle of attribute levels in their optimal residential location. However, differences in the conditions of supply and demand in the different markets in which they reside will almost certainly ensure that they are faced by different equilibrium hedonic price functions.

This is illustrated in Figure E7b where the implicit price functions for markets *B* and *C* are also shown. Notice that these implicit price functions intersect the demand curve at different levels of z_1. The points $(\hat{z}_1^B, \hat{p}_{z_1}^B)$ and $(\hat{z}_1^C, \hat{p}_{z_1}^C)$ define two more locations on the demand curve. Observing the choices of households living in different markets provides the information required to trace out the shape of the demand curve.

Unfortunately the procedure for estimating demand curves is not as simple as collecting data from multiple markets and running pooled regressions of observed levels of quantity against observed implicit prices. Further complications arise as a result of the nonlinear form of the hedonic price function. However, none of these complications are insurmountable and solutions to the problems of demand estimation are discussed in the main report Chapter 3 sections *f*, *g* and *h*. Further, a step by step guide to demand estimation using data from hedonic property markets is provided in Table 4 of the main report.

Figure E7b: Information on the Demand Curve from Multiple Markets

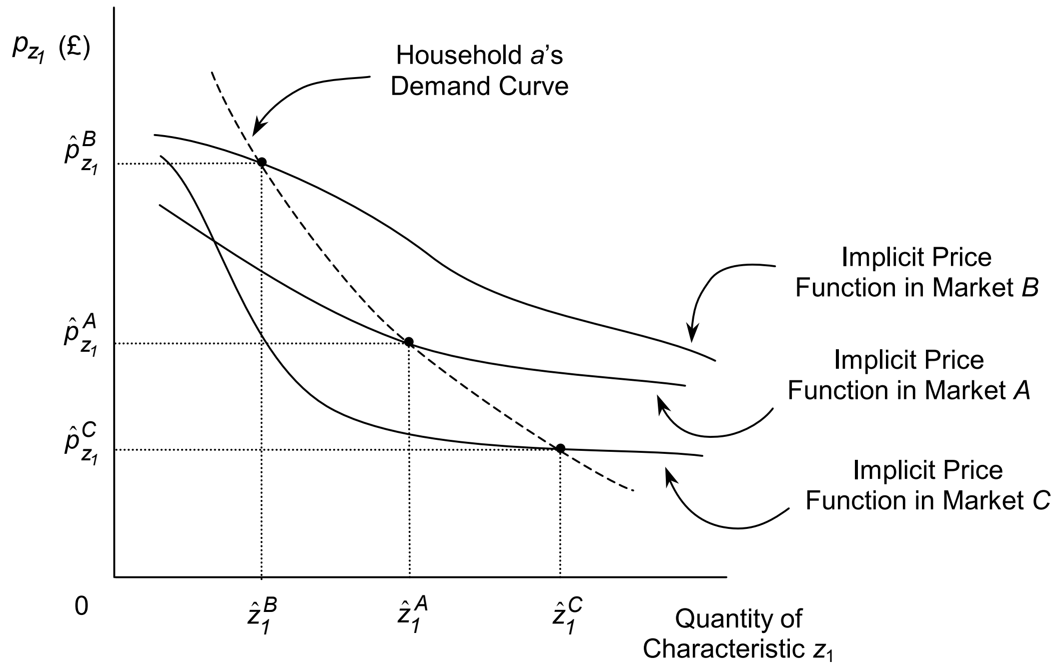

Conclusion 6: Demand functions for environmental quality cannot be estimated from data collected in a single hedonic property market without the imposition of untestable assumptions concerning the nature of household preferences.

Rather estimation of demand functions requires data from several hedonic markets.

Demand estimation is further complicated by the fact that marginal prices in hedonic markets are not linear. Table 4 of the main report describes the steps that must be undertaken to overcome these problems in order to estimate demand functions for environmental quality.

8. Conclusions on the Possibilities for Benefits Transfer

Whilst the techniques of demand estimation from hedonic analysis have been known for some years, the majority of empirical applications have stopped short of estimating demand curves. Rather researchers have gone no further than estimating the hedonic price function and reporting the implicit price of housing attributes. Whilst implicit prices can be used for measuring the welfare impacts of marginal changes in housing attributes in a particular market, they will not be accurate indicators of the welfare impacts for large changes in the housing attribute or when changes occur over a wide geographic area (see discussion in Chapter 2). Further, these implicit prices are specific to a particular housing market since they are determined by the particular circumstances of supply and demand operating in that market. Consequently, there is no theoretical basis for transferring implicit prices from one market to another. Benefits transfer using implicit prices is meaningless.

Recently, a number of research articles have reported more thorough hedonic analyses in which demand curves have been estimated (e.g. Cheshire and Sheppard, 1998; Palmquist and Isangkura, 1999; Boyle et al., 1999 and Zabel and Kiel, 2000). Demand curves, represent underlying household preferences for housing attributes. As a result they can be used to derive theoretically consistent estimates of household's welfare changes[9]. Further, under the assumption that household preferences for housing attributes are stable across different property markets, such demand functions should be transferable across property markets.

Since such transfers do not involve a single figure but an entire function, the data requirements may be intense. As described in Chapter 3 of the main report it should be possible to make some approximations that reduce these requirements. In this case, the researcher need only collect information on the income, socioeconomic characteristics and proposed change in attribute levels to be experienced in the transfer location.

As yet we are unaware of any work that has been undertaken to test the validity of such benefits transfer exercises. Indeed, a fundamental area of future research should be to investigate the accuracy of such benefit transfer measures by comparing estimated welfare values using a benefit transfer function with those derived from a separate hedonic analysis for that market. Particular attention should be paid to testing the assumption of stable preferences for environmental quality across geographical regions.

Conclusion 7: Suitably estimated demand functions could be used as a means of transferring benefits across geographical regions.

Such transfers necessarily involve simplifications and approximations. In addition, the validity of such exercises depends on the assumption that preferences for environmental quality are stable across geographical regions.

Future research should focus on testing the accuracy of welfare measures estimated by benefits transfer.

[9] As discussed previously and outlined in detail in Chapter 2 of the main report, these welfare estimates represent only those accruing to households and not those accruing to landlords. Moreover, they are only lower bounds for this value.