

The Theory of Hedonic Markets:

**Obtaining welfare measures for changes in
environmental quality using hedonic market data**

March 12th 2001

Brett Day

Economics for the Environment Consultancy

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 pay 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,

$$\mathbf{z} = (z_1, z_2, \dots, z_K), \quad (\text{E1})$$

where z_i ($i = 1$ 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;

$$P = P(z) \quad (E2)$$

Which can be read as; the price of a property, P , is a function of the vector of values, 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 purchaser².

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, 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, 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.

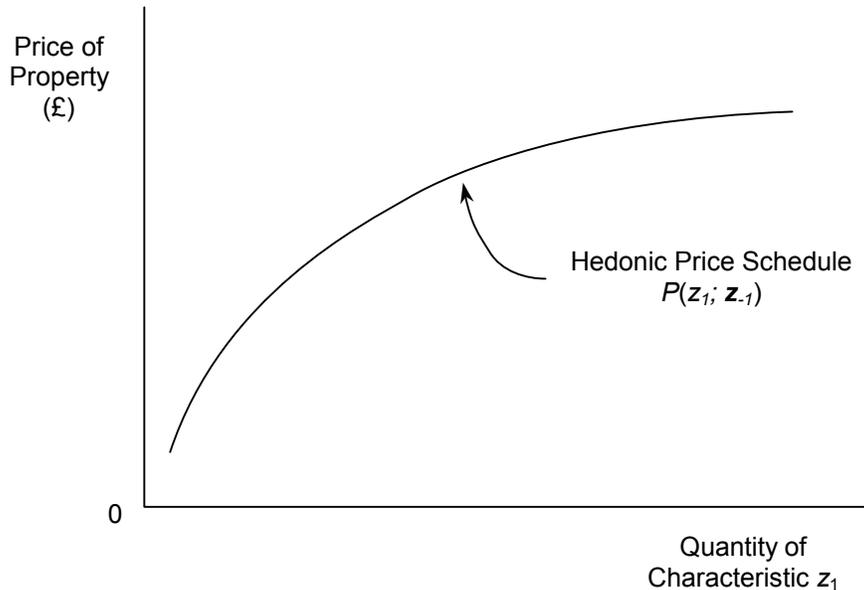
In this hypothetical case, the hedonic price schedule³ rises from left to right implying that the bigger a property’s garden the higher the price that property commands in the market.

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.

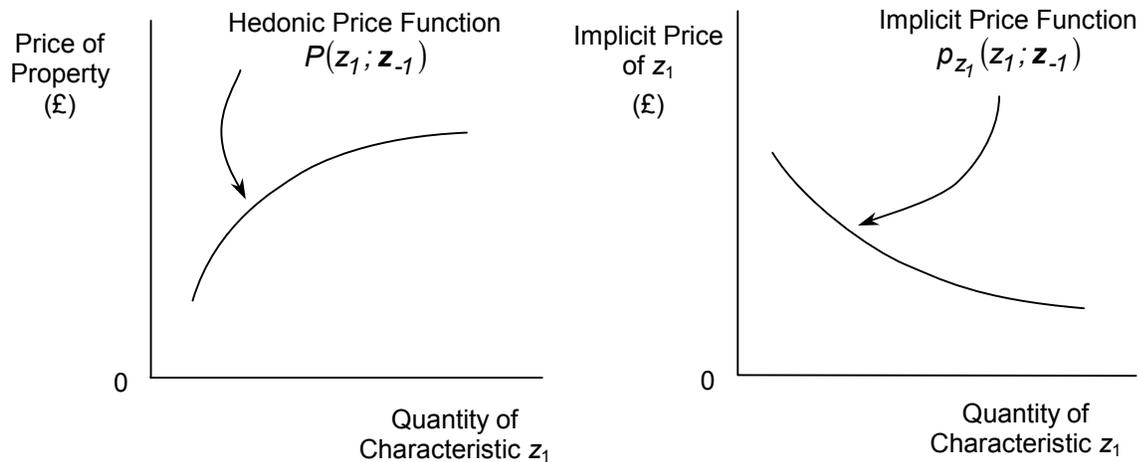
This new function is known as the *implicit price function*; *implicit* because the marginal price of a characteristic is revealed to us indirectly through the amounts households are prepared to pay for the whole property of which the particular characteristic is only a part. From Figure E2, we can see 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:

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

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 prices⁴. 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 exogenously⁵, 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.

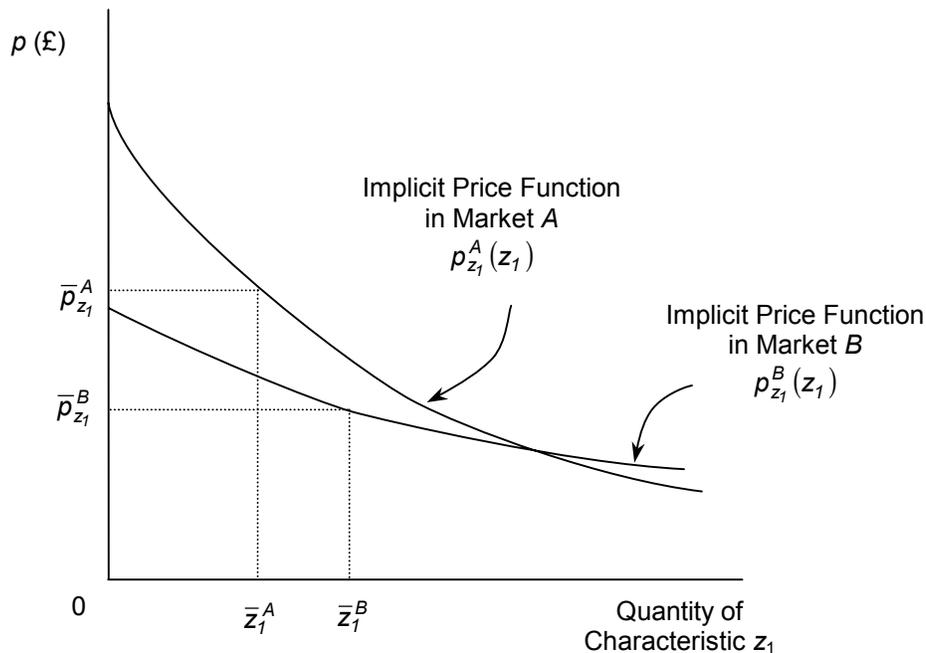
This is illustrated in Figure E3 where the implicit price function for housing attribute z_1 is shown for two separate markets, Market *A* and Market *B*. As we would expect, the two functions are quite different. Unsurprisingly, if a researcher were to summarise the implicit price functions for these 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, $\bar{p}_{z_1}^A$ and $\bar{p}_{z_1}^B$.

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.

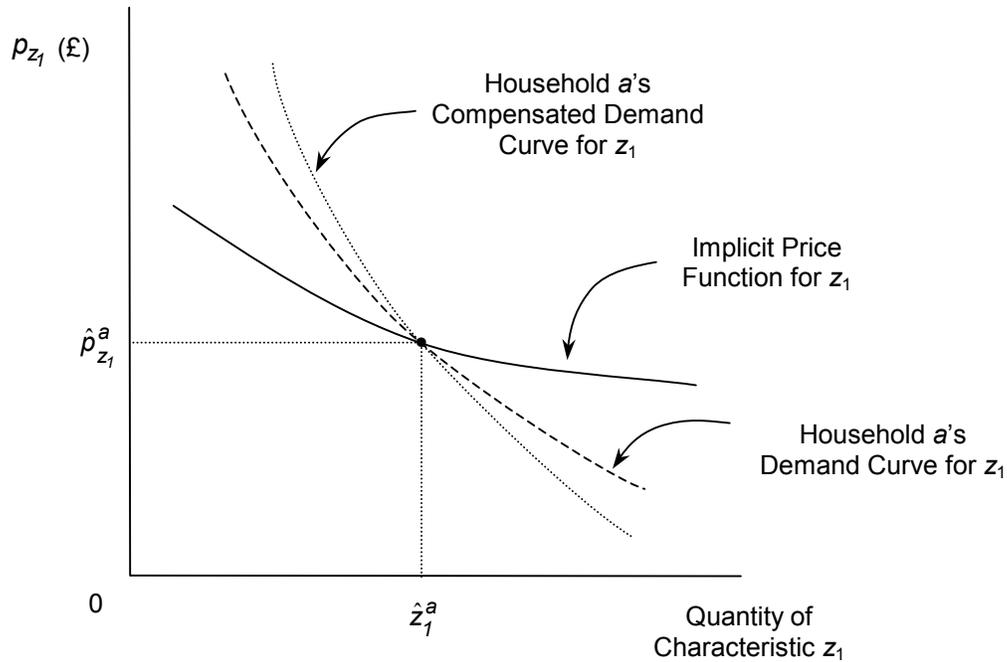
4. Household Choice in the Property Market

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 curve⁶ 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



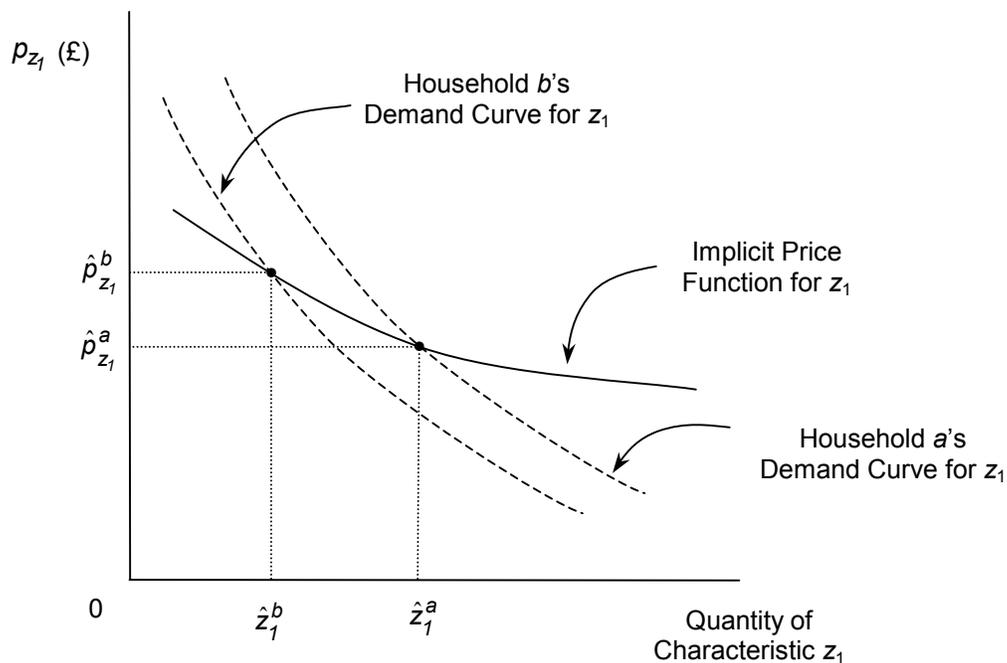
As should be familiar to those who have studied economics, the household faced by the implicit price function in this property market will choose a property with a quantity of z_1 that corresponds to the point where the market price intersects their demand curve. In the diagram this equates to choosing 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.

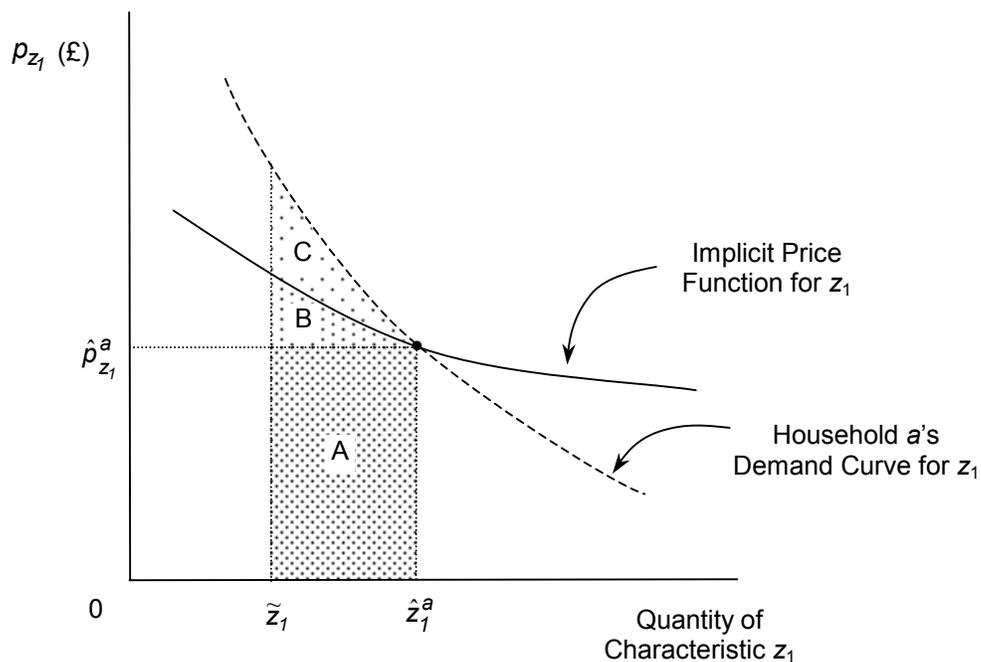
Of course, since the implicit price function will be different for each property market such welfare estimates are market specific.

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 $\hat{p}_{z_1}^a$. 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

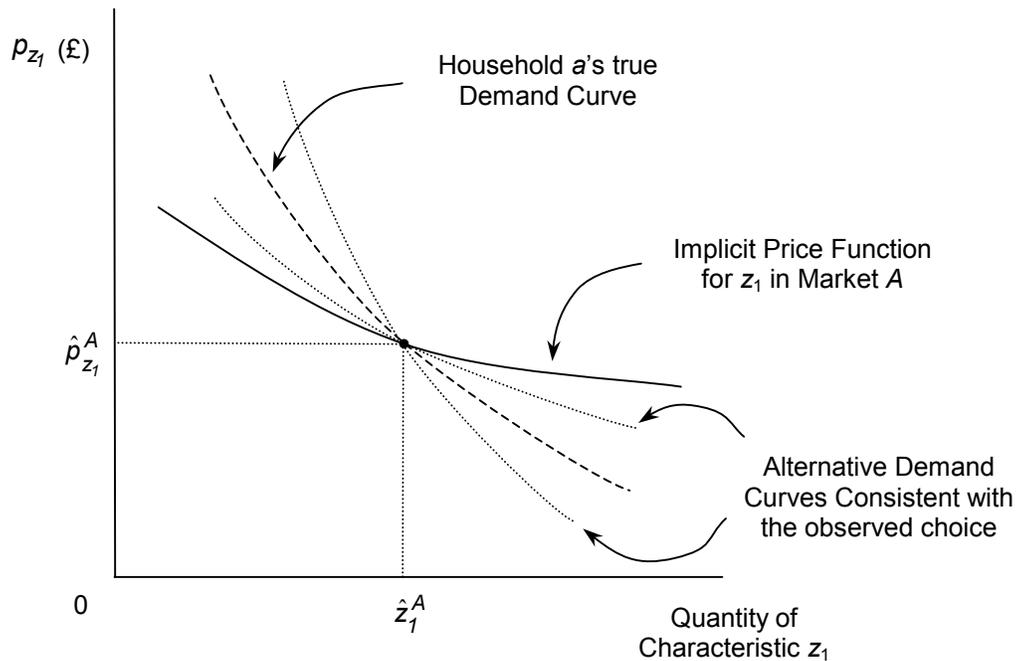
Bartik's analysis goes some way towards explaining why much of the hedonic literature has focused on the issue of estimating bid curves from empirical data. As shall become evident, however, this is not a straightforward procedure.

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

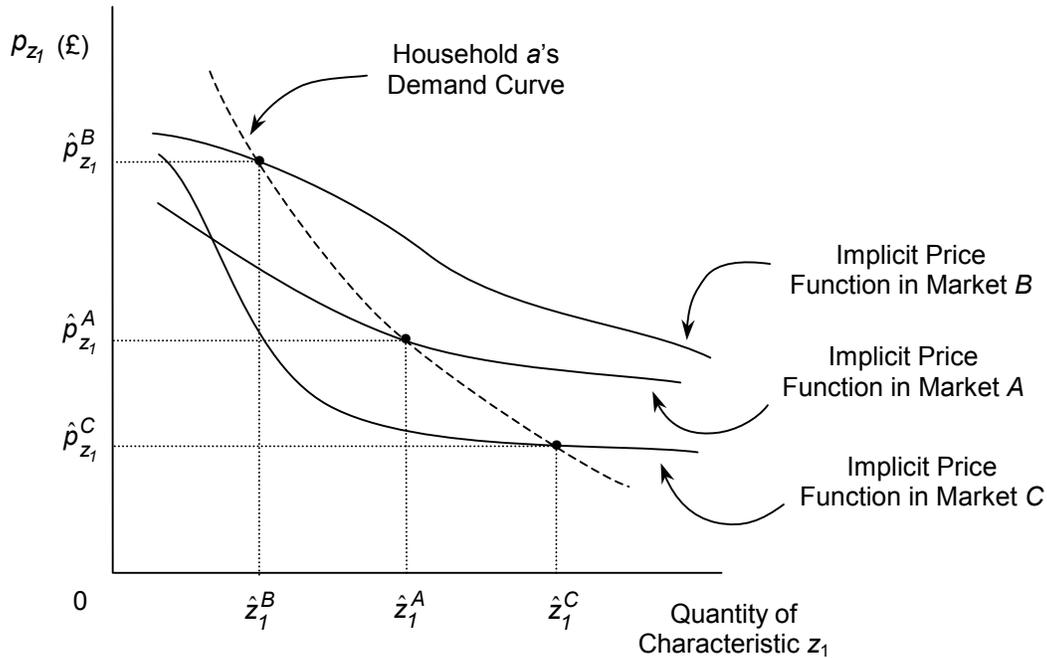


One possibility is that such information could be provided by observing the behaviour of other households in separate markets. If these households happen to have the same income and socioeconomic characteristics as the household choosing in market *A*, then it is assumed that they will have the same 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 changes⁹. 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.

CHAPTER 1. THE THEORY OF HEDONIC MARKETS

a. Introduction

One of the most familiar models in economics is that of price determination in the market. The market for a particular good consists of a large number of consumers whose demand for the good is met by the production of a large number of firms. The market mechanism works to reconcile the needs of consumers and firms by establishing the price at which aggregate demand is equal to aggregate supply and the market clears. At this price the market is said to be in equilibrium since there is no excess demand for the good and firms cannot increase their profits by changing their production of the good.

For many goods, however, this simple model is inadequate. For example, the simple model predicts that once in equilibrium the market will determine one price for the good. However, in a market such as that for housing we observe different properties commanding different prices. Indeed, housing is an example of what is called 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 and beach holidays also fit this description.

Though the simple model does not adequately explain the workings of markets in differentiated goods, it would appear that a similar market mechanism is in operation. 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 would appear to 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.

In a seminal paper, Rosen (1974) proposed a model of market behaviour that described the workings of markets for differentiated goods. The model that Rosen presented provides the theoretical underpinnings for hedonic valuation and will provide the subject matter of this first section.

b. The Property Market: The Differentiated Good

As a consequence of the fundamentally spatial nature of property, property markets are themselves defined spatially. We shall assume that at any point in time, all of the properties in one urban area represent the products in the property market. The *households* wishing to live in these properties represent the consumers in this market and the *landlords* that own the properties represent the producers in this market¹.

¹ Within this basic hedonic price model we do not consider the possibility of migration between towns.

Clearly the set of properties in the market represent a differentiated good. 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,

$$\mathbf{z} = (z_1, z_2, \dots, z_K), \quad (1)$$

where z_i ($i = 1$ 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 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”.

Further, let us assume that the set of properties in the market is fixed. That is we assume that in the short-run no new properties are built. That is not to say that the characteristics of properties do not change. A landlord maintains the quality of the property by constant renovation and maintenance. Alternatively the landlord can improve the quality of the property through investment. Building an extension, converting a loft or basement, installing double-glazing or central heating, improving the quality of the décor, indeed carrying out any number of alterations and improvements can increase the values of certain of the characteristics of the property. On the other hand, disinvesting, that is failing to maintain and renovate the property, will lead to the quality of certain of its characteristics declining. Of course, certain characteristics of the property cannot be influenced by the actions of the landlord. Most notably the landlord has little influence over the characteristics of the property that are location specific such as its proximity to places of work or to local amenities, or the property’s exposure to noise and air pollution.

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.”

c. The Property Market: 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;

$$P = P(\mathbf{z}) \quad (2)$$

Which can be read as; the price of a property, P , is a function of the vector of values, \mathbf{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.

In the property market 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. To clarify $P(\mathbf{z})$ is the per period payment made by a household to a landlord for the use of a property over that period.

Of course, many households own their own homes. In this case we treat homeowners as landlords that rent from themselves. If markets are operating perfectly, and generally we assume that they are, then the price at which the household purchases the property will be the discounted sum of all the future per period rents from that property according to;

$$\text{Purchase Price} = \sum_{t=1}^T \frac{P(\mathbf{z})}{(1+d)^t} \quad (3)$$

where t indexes each time period, T is the expected life of the property and d is the discount rate. Naturally, using Equation (3) it’s a relatively easy task to translate purchase prices into per period rentals.

Whilst the analogy of the ‘shopping cart’ is instructive in clarifying the nature of a differentiated good it is somewhat misleading in the determination of the price of that good. In particular, consider two shopping carts containing identical bundles of goods. Naturally having identical characteristics these two shopping carts would command the same price. Now imagine that a loaf of bread was added to one of the carts. The price paid for this cart would increase by the price of a loaf of bread. If another loaf were added to the same cart its price would again increase by the same amount. In general, loaves of bread don’t become cheaper or more expensive the more you purchase. Indeed, we can state that when adding a particular good (i.e. characteristic) to a shopping cart the increase in its price will be the same no matter what combination or quantity of goods are already contained in that cart. In economic terms we would say that the marginal price of each characteristic is constant.

The market forces that ensure that marginal prices are constant are known in economic terms as “arbitrage”. Returning to our example, since buying two carts both containing one loaf of bread confers the same benefit on a consumer as purchasing one cart containing two loaves of bread, arbitrage activity ensures that the marginal price of bread is constant. To expand, if the cart with two loaves were more expensive than twice the price of the carts containing one loaf each, then a rational shopper would always chose to push the two single-loafed carts through the check-out. The result would be a lack of demand for two-loafed carts and excess demand for one-loafed carts. Market forces would work so as to bring the price of two-loafed carts down and increase the price of one-loafed carts. Only when the price of the former were twice the price of the latter would the market reach an equilibrium. In this equilibrium, arbitrage activity has worked to ensure that the marginal price of extra loaves in a shopping cart is constant.

However, the same sort of activity is frequently impossible in the purchase of a truly differentiated good such as a property. This feature of hedonic markets results from the fact that households are unable to “repackage” the differentiated goods. In other words, households cannot break up the differentiated good into its constituent parts and enjoy the benefits of each characteristic separate from the whole. For example, talking in terms of just one characteristic, two houses with one bedroom are not equivalent to one house with two bedrooms since a household cannot live in both properties simultaneously. Similarly, renting a property with four bedrooms for half a year and a property with two bedrooms for the other half is not the same as renting a three-bedroom house all year round. Since these types of arbitrage activity are precluded in the housing market, market forces do not work to ensure that the marginal price of bedrooms is constant.

This observation leads to two interesting insights.

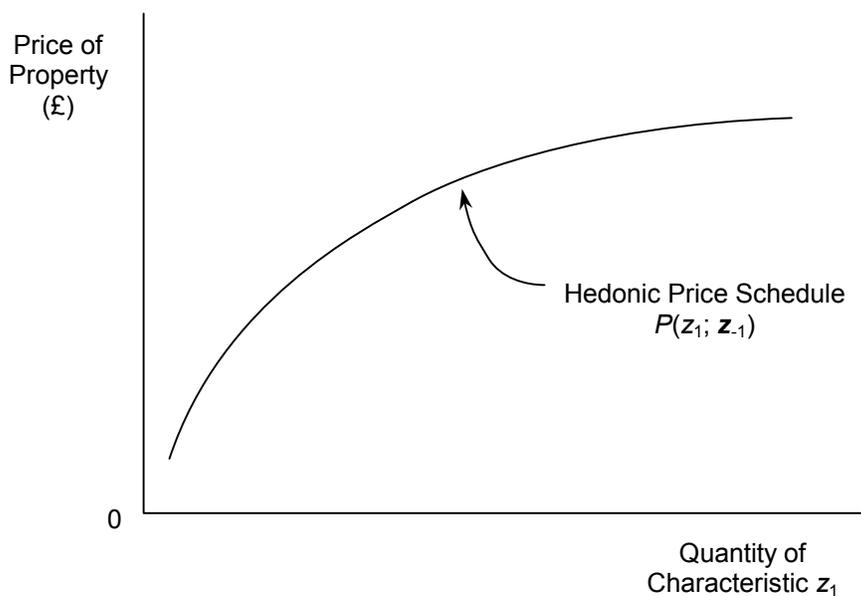
- *Marginal prices may not be constant.* To illustrate, imagine a number of properties that are identical in all characteristics except one. Within the housing market we may find that the extra paid for properties with additional units of this particular characteristic is not constant. Indeed, more typically, the additional amount paid for properties enjoying increasingly higher quantities of the characteristic (in effect the price of that characteristic) declines as the total level of that characteristic increases.
- *The price of one characteristic may depend on the quantity of another.* As an example, a house with a garden is more desirable than a house without. Further, if the aspect of the house is north-south, having a garden may be even more desirable since it will enjoy longer exposure to the sun. Now consider the extra paid for a north-south aspect, effectively the ‘price’ of north-south aspect. Without a garden, north-south aspect may be somewhat desirable, but households are unlikely to pay a great deal more for a property with this characteristic compared to an identical property with east-west aspect. For properties with a garden, on the other hand, aspect may be a much more important consideration. It would not be surprising that the price of north-south aspect will depend on whether a property has a garden or not.

These two observations will be important in empirical applications that attempt to estimate the hedonic price function from market data.

To illustrate the hedonic price function, consider the illustration in Figure 1. Plotted on the vertical axis is the price (rental per unit time) 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, 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 1, z_{-1} comes after a semicolon. This indicates that these other characteristics are held constant at some given level whilst the focus characteristic, size of garden (i.e. z_1), changes. Consequently, in this example we are not considering the interaction of different characteristics of the property.

Figure 1: The Hedonic Price Schedule for characteristic z_1



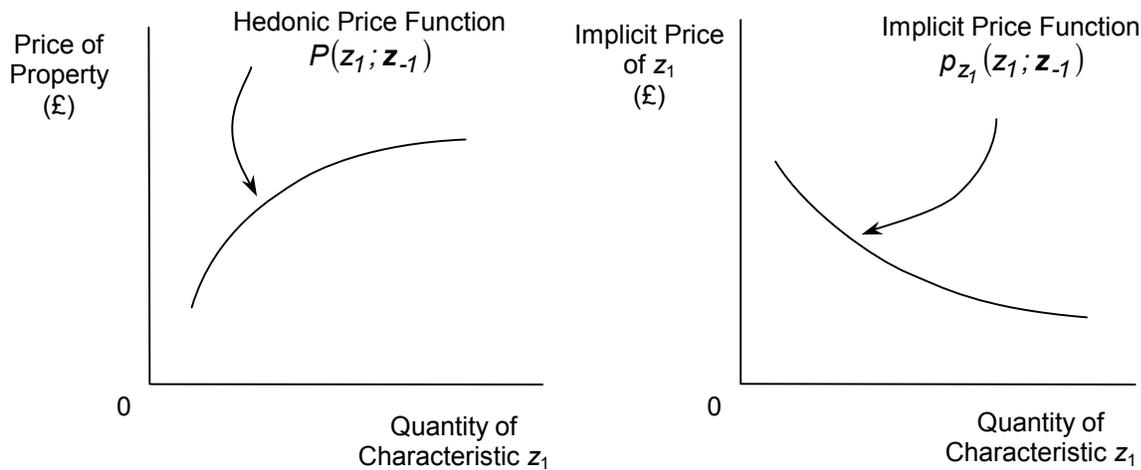
In this hypothetical case, the hedonic price function rises from left to right implying that the bigger a property's garden the higher the price that property commands in the market. Notice also that the marginal price of extra garden space is not constant. The slope of the curve becomes progressively 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.

Of course the relationship won't be identical to that graphed for every type of characteristic. For example, we might find that for another characteristic, say floor space, plotting house price against the quantity of the characteristic (again holding all other characteristics constant) results in a straight line rising from left to right. A straight line relationship suggests that there is no satiation and that the price commanded by a property increases uniformly in relation to the quantity of the characteristic that it possesses. In other words the characteristic would possess a constant marginal price.

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 2.

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 2, 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.

Figure 2: The Hedonic Price and the Implicit Price Schedules for characteristic z_1



Mathematically, the implicit price is derived as the partial derivative of the hedonic price function (Equation 2) with respect to one of its arguments, z_i , according to:

$$p_{z_i}(z_i, z_{-i}) = \frac{\partial P(z)}{\partial z_i} \quad (4)$$

Again $p_{z_i}(z_i, z_{-i})$, the marginal price function of characteristic z_i , does not have to be a constant. We have dwelt on the subject of non-constant marginal prices for characteristics since as we, shall see in a later section they are the source of much of the complications that confound the empirical estimation of welfare measures using hedonic analysis.

d. The Property Market: Household Choice

Let us take as a fact that the hedonic price function, $P(\mathbf{z})$, emerges from the interaction of households (demanders) and landlords (suppliers) and represents a market clearing equilibrium. We shall return to the mechanism by which this equilibrium is derived, but for now we shall focus on how households facing such a hedonic price schedule determine their optimal residential location.

The model that Rosen developed to explain these decisions is based on a number of assumptions. Amongst these, some of the most important are that;

- Each individual household in the market is a price taker; they make their choice of location based on the hedonic price schedule they observe in the market and cannot influence this schedule through their actions. This point has been made by several authors (McConnell and Phipps, 1987; Palmquist, 1991) and allows us to ignore the supply side of the market in modelling households' residential decisions. Given the size of the urban property markets in which hedonic pricing techniques are usually applied, such an assumption would appear reasonable.
- Each household only purchases or rents one property². If households purchase a second home, say a holiday home, then this should be considered as a separate good being purchased in a separate hedonic market. Again, this assumption is, in general, readily defensible.

Given these assumptions, Rosen sets out a model in which households choose their residential location so as to get the maximum flow of benefits or, in economic terms, utility from the property. To do this, it is assumed that households in the market have well-defined preferences over all goods and that these preferences can be represented by the utility function³;

$$U(\mathbf{z}, x; \mathbf{s}) \tag{5}$$

As should be familiar to most readers, the utility function gives the utility that a household enjoys per period of time, given the levels of the arguments contained in the brackets. In this case, there are three arguments;

- \mathbf{z} which represents the levels of the different characteristics of property that a household could purchase or rent. Naturally the utility that a household enjoys each period of time will depend on the qualities of the property in which they choose to live.

² Further, if households also act as landlords to other households, then their decisions concerning their own choice of residential location are assumed to be independent of their decisions concerning these other properties.

³ The utility function is assumed to be identical for all households in the market. However, the actual utility derived from a certain property with characteristics, \mathbf{z} , and a certain quantity of other goods, x , will depend upon the characteristics of the household, \mathbf{s} .

- x which represents all other goods outside the property market. As a matter of convenience we standardise x to have a unit price, such that, we are effectively representing all other goods by a quantity of money⁴. Again the greater this quantity of money to spend on other goods, the more utility a household will enjoy per period time.
- s which represents the characteristics of the household themselves. Clearly, the quantity of utility a household enjoys from any of the other arguments will depend on their own characteristics. For example, having a swimming pool in the back garden will confer little benefit on a household of non-swimmers. Again, notice that this argument is placed after a semicolon in the utility function. This indicates that the level of utility that a household gets from any of the arguments over which it has a choice (i.e. z and x) will be dependent upon (or conditional upon) their own characteristics.

For the purposes of developing a model of behaviour we don't specify the exact form of the utility function⁵. In other words we don't need to state that the quantity of utility that a household experiences will be, for example, 2 times the number of bedrooms plus 4 times the size of the garden, plus -.02 times the number of decibels of road traffic, plus ... etc. For our purposes, we can continue just assuming there is such a function and that it is the same for each household conditional upon their characteristics⁶.

Households choose levels of z and x to maximise $U(z, x; s)$ subject to the constraints imposed upon them by their budget. Since the price of x is taken as unity and the price of a property with characteristics z is given by the hedonic price function $P(z)$, we can represent the budget constraint as;

$$y = x + P(z) \tag{6}$$

where y is household income per period.

⁴ In the economics literature x is frequently referred to as a numeraire, a composite good or a Hicksian bundle.

⁵ We do however assume that the utility function is strictly increasing in the arguments x and z . That is we assume that the household prefers more attributes to less attributes (remember we already assumed that all attributes were measured as goods rather than bads) and more money to spend on other goods than less. Further, for mathematical simplicity, we assume that the utility function is strictly quasiconcave and twice continuously differentiable.

⁶ Some authors (e.g. Epple, 1984) do in fact impose some specific functional form on the utility function and use this as a means of investigating prices and choices in hedonic markets. We do not consider such models here.

As with the standard consumer choice problem, we can use Equations (5) and (6) to set up the Lagrangian Function;

$$L = U(\mathbf{z}, x; \mathbf{s}) + \lambda(y - x - P(\mathbf{z})) \quad (7)$$

Maximising this with respect to x , \mathbf{z} and the Lagrange multiplier λ gives rise to the first order conditions;

$$\frac{\partial L}{\partial z_i} = U_{z_i} - \lambda P_{z_i} = 0 \quad (i = 1 \text{ to } K) \quad (8)$$

$$\frac{\partial L}{\partial x} = U_x - \lambda = 0 \quad (9)$$

$$\frac{\partial L}{\partial \lambda} = y - x - P(\mathbf{z}) = 0 \quad (10)$$

Where U_{z_i} is the partial derivative of the utility function with respect to property characteristic z_i . This represents the extra utility that comes from choosing a property with one extra unit of characteristic z_i , all else being equal.

U_x is the partial derivative of the utility function with respect to the composite good. Since x is constructed so as to represent money to be spent on other goods, U_x can be interpreted as the extra utility that comes from an extra unit of money, all else being equal.

P_{z_i} is the partial derivative of the hedonic price function with respect to property characteristic z_i . Of course this is simply the implicit price function for characteristic z_i as presented in Equation (4). Indeed, $P_{z_i} = p_{z_i}(z_i, \mathbf{z}_{-i})$.

Equations (8), (9) and (10) represent the conditions that define the household's optimal choice of residential location. That is, given the constraint of their budget, the flow of utility that the household enjoys will be maximised by choosing a property whose characteristics simultaneously satisfy the conditions laid out in Equations (8), (9) and (10).

In their present form these conditions provide us with little insight into the household's choice behaviour. However, if we rearrange Equations (8) and (9) and divide one by the other (thereby eliminating the Lagrange multiplier) we reveal that one of the conditions for optimal choice is given by the expression;

$$\frac{U_{z_i}}{U_x} = p_{z_i}(z_i, z_{-i}) \quad (11)$$

To illustrate the condition laid out in Equation (11) Rosen defined a function that he termed the *bid function*, whose slope is given by the ratio of marginal utilities, U_{z_i}/U_x . Most students of economics will be familiar with terms involving ratios of marginal utilities. More usually we would expect to see the ratio of two marginal utilities preceded by a negative sign. In such a case, the expression would represent a *marginal rate of substitution*, e.g. $-U_{z_i}/U_x$; the quantity of one good that a household is willing to give up in order to obtain one more unit of another good such that their overall well-being does not change. In the same way that U_{z_i}/U_x defines the slope of Rosen's bid function, the marginal rate of substitution defines the slope of an indifference curve. In hedonic analysis there is a simple correspondence between the indifference curve and the bid function that goes some way in clarifying the nature of the latter.

Let us spend a little time considering indifference curves. In mathematical terms, the indifference curve is implicitly defined as;

$$U(z, x; s) = u \quad (12)$$

Where u is any specified level of utility. Thus, the indifference curve depicts combinations of x and z that confer the same level of well-being or utility on the household. Indeed, solving Equation (12) for x would give us a general expression for an indifference curve that we can denote;

$$x(z; s, u) \quad (13)$$

Written in this form, the indifference curve tells us what quantity of money to spend on other goods, x , would allow a household with characteristics, s , to enjoy the level of utility u , given they lived in a property with characteristics z .

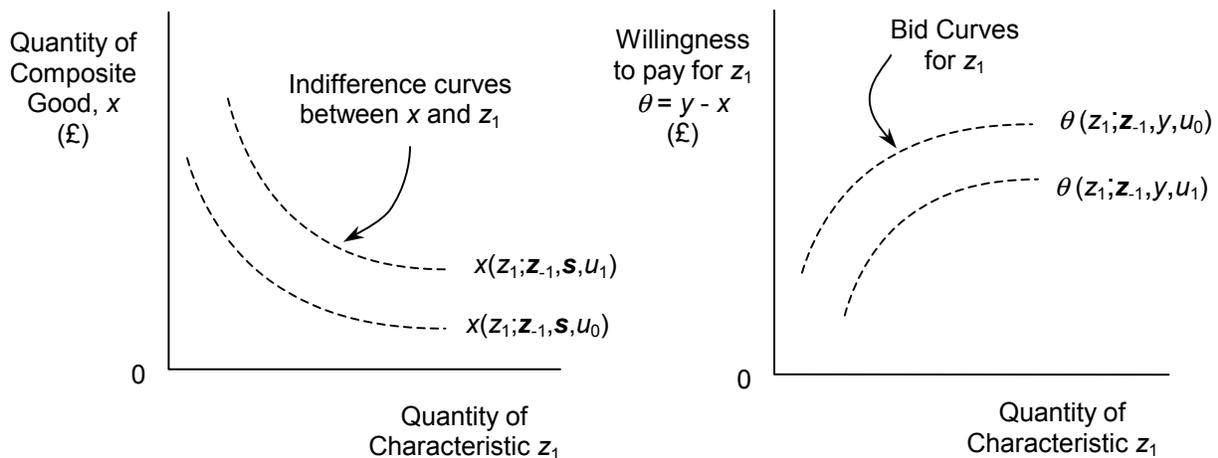
The left hand panel of Figure 3 shows a set of indifference curves between x (the quantity of money to spend on other goods) and z_1 (one of the attributes of a property)⁷. Most readers will be familiar with this diagram. Each indifference curve depicts combinations of x and z_1 that confer the same level of well-being or utility on the household.

⁷ For diagrammatic exposition, it is necessary to present indifference curves in terms of only one property attribute. The assumption in Figure 3 is that all other attributes are held at some constant level. In reality, the indifference 'curve' would be a multidimensional indifference surface plotting combinations of x and quantities of each of the attributes in z between which the household is indifferent.

Notice that when levels of attribute z_1 are low, the household is willing to give up quite a lot of their money to spend on other goods in order to acquire a property with greater levels of the attribute. Conversely, when levels of attribute z_1 are high the household is only prepared to sacrifice a small amount of money to spend on other goods in order to acquire more of the attribute⁸. The slope of the indifference curve gives the rate at which households are prepared to give up money for other goods in order to acquire more of the housing attribute whilst not changing their overall well-being. As discussed previously, the slope of the indifference curve is the marginal rate of substitution between x and z_1 , $-U_{z_1}/U_x$.

Any combination of x and z_1 that lies above and to the right of an indifference curve provides the household with more money to spend on other goods and/or more of the housing attribute. By definition they must gain more utility from such a bundle than from any bundle lying on the indifference curve. Consequently the indifference curve for u_1 must represent bundles of x and z_1 between which the household is indifferent but which all confer more utility than bundles lying along the indifference curve for u_0 .

Figure 3: Indifference Curves and the Bid Function



So far we have not considered the fact that the household is constrained in their choices of x and z by their limited budget, y . Money spent on other things is money that cannot be spent on housing attributes. Let us define, therefore, an amount that we shall call a *bid* as;

$$\theta = y - x \tag{14}$$

⁸ This “classic” shape for the indifference curve stems from our assumptions concerning the utility function described in footnote 14.

That is the bid, θ , represents the total amount a household could pay for a property given that they spent x on other goods⁹. Clearly the relationship between the bid, θ , and the amount spent on other goods, x , is very simple; as one goes up by a certain amount the other falls by the same quantity¹⁰. Indeed, using Equation (14) we could redefine the indifference relationships of Equation (13) in terms of bids rather than money spent on other goods. Replacing Equation (13) in Equation (14) gives;

$$\begin{aligned}\theta &= y - x(\mathbf{z}; \mathbf{s}, u) \\ &= \theta(\mathbf{z}; y, \mathbf{s}, u)\end{aligned}\tag{15}$$

The bids defined by Equation (15) are a special type of bid. They are bids for a property with characteristics \mathbf{z} that result in the level of utility u . Indeed, Equation (15) defines Rosen's *bid function*.

In words, the bid function depicts the maximum amount that a household would pay for a property with attributes \mathbf{z} such that they could achieve the given level of utility, u , with their income, y . Notice that increases in income translate directly (i.e. pound for pound) into increases in the bid function.

The bid function can be illustrated as *bid curves* as depicted in the right hand panel of Figure 3. In constructing these bid curves, all that has been done, in effect, is to flip the vertical axis. Bid curves still define indifference relationships. They depict combinations of property attributes, \mathbf{z} , and payments for those attributes, θ , between which the household is indifferent. Accordingly, all bid/attribute-quantity combinations on a particular bid curve provide the household with the same level of overall utility. Combinations of housing attributes and bids lying below and to the right of a particular bid curve represent bundles providing more attributes and/or lower payments. Clearly the household would gain more utility from such a bundle. Consequently the lower bid curve in Figure 3 provides the household with greater overall utility, u_1 , than the higher bid curve, u_0 .

Since the bid curve is simply an inverted indifference curve, the slope of the bid curve will be the same as the slope of the indifference curve but with the opposite sign i.e. U_{z_i}/U_x . And, of course, this ratio represents the left hand side of the condition for optimal residential location given in Equation (11).

So far our analysis has defined two closely related functions the indifference curve (Equation 13) and the bid curve (Equation 15). As yet, however, we have not determined

⁹ Note that Equation (14) does not enforce the budget constraint of Equation (10). In Equation (10) θ , the amount the household could pay for a property, is replaced by $P(\mathbf{z})$, the amount that the household would have to pay in the market for a property.

¹⁰ Though, clearly, the bid is constrained in that it cannot be greater than income, y .

how these functions are important in defining households' optimal residential choice. To do this, we must make use of the last of the first order conditions, that defining the budget constraint (Equation 10).

To plot the budget constraint we must rearrange Equation (10) to give;

$$x = y - P(z) \quad (10a)$$

In the left hand panel of Figure 4, the budget constraint has been added to the indifference diagram. The budget constraint describes all combinations of x and z_i that the household is able to buy given their income, y .¹¹ Bundles that are on the budget constraint or bundles that are to the left and below the budget constraint are affordable to the household. Those that are above and to the right of the budget constraint are too expensive for the household to purchase. In order to maximise their utility the household must choose amongst the affordable bundles of x and z_1 . This amounts to choosing the bundle amongst affordable combinations of x and z_1 that lies on the highest indifference curve. The optimal bundle (\hat{x}, \hat{z}_1) will be defined as the point of tangency between this highest indifference curve and the budget constraint. Notice that throughout this document we use a hat to represent a chosen bundle. Any other bundle in the affordable set will lie on an indifference curve that provides a lower level of utility.

The left hand panel of Figure 4 will be familiar to students of economics. However, the diagram differs from the usual consumer choice problem in that the budget constraint is not linear. For most goods marginal prices are constant, such that purchasing one more unit of a good will require a constant sacrifice in terms of ability to purchase other goods. Thus in the classic consumer choice problem, the budget constraint can be represented by a straight line. As we have already established, marginal prices may not be constant in hedonic analysis giving rise to the unfamiliar shape of the budget constraint.

The choice of an optimal bundle of housing attributes can just as easily be presented in terms of bid functions. Of course we have to transform the constraint to be expressed in the same terms as bids. Remember that the vertical axis of the bid function graph is measured in terms $y - x$; that is money available to spend on housing attributes. Rearranging the income constraint, Equation (10), gives;

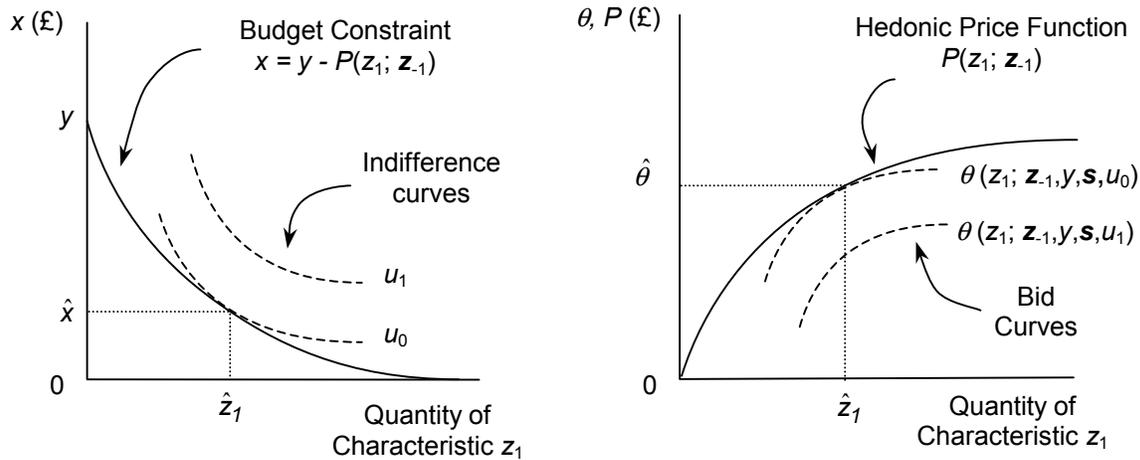
$$y - x = P(z) \quad (10b)$$

In other words, the relevant constraint is simply the hedonic price function. This is a very intuitive result. Bid functions reveal the amount that a household is willing to pay for

¹¹ Again in order to illustrate the optimal conditions graphically we are forced to present a two-dimensional analysis. The figures presented here assume that all other characteristics of the property are unchanging in the analysis.

different levels of attribute z_1 ; the hedonic price function gives the minimum price that they must pay in the market to purchase different levels of z_1 .

Figure 4: Choice of Optimal Residential Location using Indifference Curves and the Bid Function



Intersections between bid curves and the hedonic price function indicate bundles of housing attributes at which the household's willingness to pay for a property with that bundle of attributes is equal to its market price. In maximising their utility, the household will choose the bundle of housing attributes that positions them on the bid curve providing the highest level of utility whilst still being compatible with reigning market prices. To be explicit, the household will maximise their utility by moving to the lowest bid curve that is just tangent with the hedonic price function.

The point of tangency between this lowest bid curve and the hedonic price function defines the bundle of housing attributes that fulfil the first order conditions for an optimal choice (Equations 7, 8 and 9).

Combining Equations (8) and (9) into the single expression in Equation (11), provides the first condition for an optimal choice. In particular this condition states that at an optimum, the slope of the bid function and the slope of the hedonic price function must be the same. That is, the household's willingness to pay to attain a property with one more unit of z_i , U_{z_i}/U_x , must equal the market price of that extra unit, p_{z_i} . Thus a household's optimal choice of residential location will be one at which the value the household derives from the last unit of each housing attribute is exactly equal to the implicit price it had to pay for that unit. If this were not so then the household could increase their flow of utility by choosing an alternative property with different levels of attributes.

The second condition defining optimal choice is provided by the budget constraint, Equation (10). This states that the chosen bundle of attributes must be purchased at the

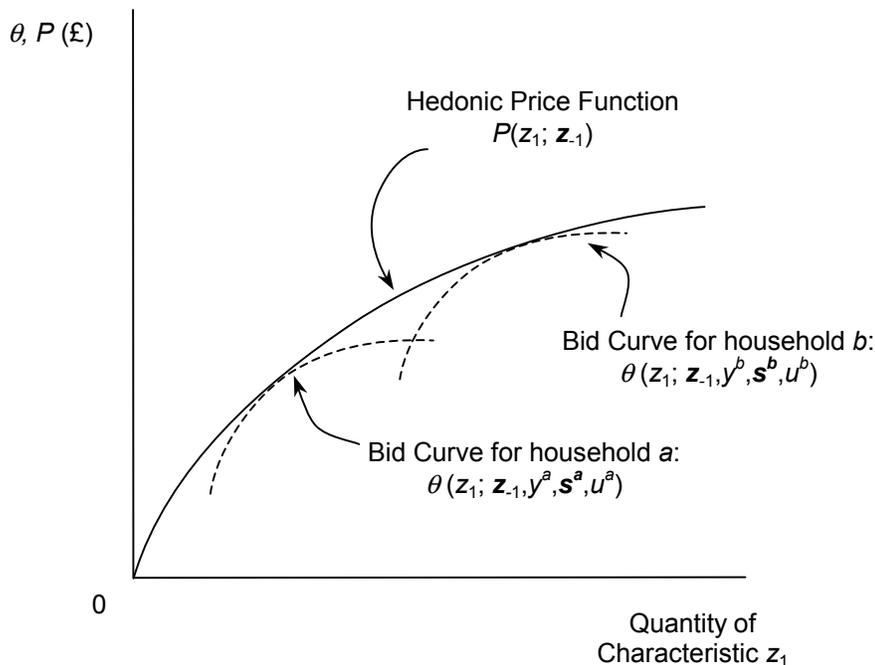
market price as defined by the hedonic price schedule. Combining this constraint with Equation (11) identifies the tangency point presented in Figure 4.

Of course, households do not have the same income nor do they have the same socioeconomic characteristics. Since both these arguments enter the bid function, we would expect bid curves to differ across households with different characteristics. Bid curves for two different households, denoted a and b , are illustrated in Figure 5. Notice that the conditioning arguments, y and s , are superscripted with this household indicator, showing that their values are specific to the particular household.

Again the optimal choice of property for each household will be defined by the tangency of one of their bid curves with the hedonic price function. Since the bid curves for the two households are different, the attributes of the property defined by this point of tangency will also differ. Notice that the utility level, u , that defines the optimising bid curve is also superscripted by a household indicator. Clearly, there is no reason to expect that the level of utility achieved by the two households would be the same.

If we were to add to Figure 5 the optimising bid curves for all the households in the market we would find that they were all tangential to the hedonic price function. Variation in household characteristics would mean that these points of tangency defined properties with different levels of the various housing attributes. In the terminology of economics, the hedonic price function forms an upper envelope to these optimising bid functions¹².

Figure 5: Choice of Property Attributes for Different Households



¹² We can think of the hedonic price function almost like an envelope into which each of the optimising curves is fitted.

e. The Property Market: Landlord Choice

So far we have examined the property market solely from the demand side; that is, in terms of consumers choosing between differentiated products. Though this decision is of greater interest for our present research objectives, it is worth taking some time to examine the supply side of the market; that is, to describe how landlords make their decisions concerning the type of properties to supply.

To simplify the analysis let us assume that each landlord rents out only one property.¹³ In each period of time a household pays the landlord rent in order to live in this property. Of course this rent does not represent pure profit to the landlord. The landlord incurs costs in supplying this property for rental;

- First and foremost, through the initial purchase of the property.¹⁴
- Second, through maintaining the quality of the property by constant renovation and maintenance.
- Finally, through investments or disinvestments designed to change the attributes of the property subsequent to its purchase.

To incorporate these costs into our per period model, all discrete investments must be converted to equivalent per period costs. For example, the purchase price of the property can be expressed as an equivalent series of per period payments using Equation (2a).¹⁵ In the same way, it is possible to express a discrete investment in the property, say the installation of double-glazing, as the discounted sum of a series of smaller equal-sized costs made over the expected lifetime of the investment.

The per period cost to the landlord of supplying a property with characteristics z , is given by the cost function¹⁶;

$$c(z; \hat{P}(\hat{z}), \bar{z}, r) \tag{16}$$

The cost of producing a property with characteristics z will differ across landlords for a number of reasons. The factors determining these differences are captured in the three conditioning arguments entering the cost function. These are;

¹³ The analysis remains relatively simple if we assume that the landlord rents out more than one property but that each property has the same characteristics. However, if the landlord rents out several properties with differing characteristics the model becomes considerably more complex whilst adding little to our understanding of the workings of the hedonic market.

¹⁴ Or even the initial construction of the property though we ignore this possibility here.

¹⁵ Indeed, in the UK, this does not represent a major abstraction from reality. It is typical for landlords to take out a mortgage in order to purchase a property. The purchase price of the property (plus interest on money borrowed) is repaid, therefore, in a series of monthly instalments.

¹⁶ This cost function is the result of a minimisation problem in which the landlord attempts to find the cheapest cost means by which to produce a property with characteristics z .

- $\hat{P}(\hat{z})$, which defines the *price paid for the property when first purchased*. Consider two landlords owning two originally identical properties. In terms of the notation, the original vector of housing attributes, \hat{z} , is identical for both landlords. Now, imagine that these two landlords purchased their properties at different times; the first during a recession that depressed the housing market and resulted in relatively low implicit prices for housing attributes, the second during a property market boom when implicit attribute prices were relatively high. Differences in the market price paid for the two properties are captured in differences in $\hat{P}(\cdot)$, the hedonic price function faced by the landlord at the time of purchase. Clearly, the cost of supplying the property is lower for the first landlord who bought when house prices were depressed than for the second landlord who bought when house prices were high.

Since this cost (expressed in equivalent per period terms) is constant for each landlord and independent of changes in the property market or in the characteristics of the property, we suppress this argument in the cost function henceforth.

- the vector \bar{z} , which defines the levels of attributes that, following the initial property purchase, are *provided costlessly to the landlord*.
 - For *structural attributes* this vector is likely to be identical to \hat{z} , the vector of housing attributes purchased by the landlord. For example, having purchased a two-bedroom house, the landlord does not have to pay further in order to maintain this number of bedrooms in the property.
 - For *locational, neighbourhood and environmental attributes*, the levels of \bar{z} will tend to be determined by factors that are beyond the control of the landlord. In economics we would describe these as *exogenous* factors. For instance, \bar{z} would include a measure of the baseline level of crime in the area. This baseline level of crime is provided costlessly to the landlord in so much as it is exogenously determined by government spending on crime prevention.¹⁷ Indeed, baseline levels of many non-structural attributes of the property will be provided costlessly to the landlord since they tend to exhibit the characteristics of public goods. Other examples include the property's proximity to recreational facilities, its access to public transport, levels of air pollution and levels of noise pollution.

For many non-structural attributes, therefore, the values in \bar{z} will not be determined solely by \hat{z} (the vector of housing attributes initially purchased by the landlord). Indeed, to a large extent, the values of \bar{z} for non-structural attributes of the property are determined by public policy. Policies that reduce crime, redirect traffic, combat air pollution or increase the quality of public transport will determine the values included in \bar{z} for certain attributes. As we shall discuss in the next chapter, the primary aim of hedonic analysis is to determine the benefits to households and landlords of public projects that change the levels of these exogenously determined housing attributes.

¹⁷ Of course, the landlord indirectly pays for public goods such as these through taxation but payments are not directly linked to the level of the attribute enjoyed at the property and payments cannot be unilaterally altered so as to influence the level of attribute enjoyed at the property.

It is relatively easy to see why the vector \bar{z} is important in determining how much it costs landlords to achieve a certain level of attribute provision at their properties.

- Consider a *structural attribute*, such as the cost to two different landlords of providing a property with two bathrooms. Imagine that the properties owned by these two landlords are identical, except that the second landlord's property was purchased with a ground floor extension fitted with a second bathroom.¹⁸ Clearly, it costs this landlord nothing to provide a two-bathroom property. On the other hand, the first landlord must invest in the building and fitting of the second bathroom. Clearly, the landlord's costs for providing a property with a certain level of structural attributes will be determined in part by the property's original level of structural attributes.
- Alternatively let us examine an *attribute with a public good element*. Take the peace and quiet attribute of a property as an example. Imagine two landlords, one whose property is in a quiet cul-de-sac and another whose property abuts a busy main road. For the sake of argument, we shall assume that their properties are identical in every other way. Now consider how much it would cost these two different landlords to achieve a certain level of peace and quiet. The landlord whose property is in a quiet road will have to spend little to attain a relatively high level of peace and quiet in the property. Possibly the most that would be needed would be to plant a few trees in the front garden to act as a barrier to traffic noise. Conversely, a landlord whose property is on a busy main road would have to invest relatively heavily in order to attain the same level of the peace and quiet attribute. Perhaps this landlord would need to install sound-proofing double-glazing. Clearly, the cost of attaining a certain level of peace and quiet will differ for the landlords of these otherwise identical properties depending on the exogenously determined level of traffic noise.
- the vector r , captures *other parameters* important in determining the landlord's costs. For example, r will include the characteristics of the landlord and the market price of investments. To illustrate, a landlord that is a capable plumber may be able to improve the quality of a property by installing an electric shower unit. The costs may be lower to this landlord than to another who has to seek professional help to achieve the same improvement.

The cost function, therefore, determines the per period cost of supply of a property with characteristics z , given the purchase price of the property $\hat{P}(\hat{z})$, the levels of exogenously determined property attributes, \bar{z} , and a number of other parameters including the characteristics of the landlord, r .

Importantly, landlords have the ability to change the characteristics of their property. For example, a landlord may choose to increase the peace and quiet at the property. For the

¹⁸ Obviously, since the two properties are not identical, we would expect their purchase prices to differ. However, this cost difference will be determined by the (possibly different) hedonic price functions faced by the two landlords at the time of purchase. These costs are already captured in the cost function by the argument $\hat{P}(\hat{z})$.

sake of argument let us assume that the least cost method of achieving this increase is to install double-glazing. In this case the difference in the cost function evaluated at the original level of peace and quiet and the cost function evaluated at the increased level of peace and quiet would be the cost of the double-glazing.¹⁹

Given the per period cost defined by the cost function, the profit that a landlord derives from renting a property with characteristics \mathbf{z} , will be determined by the rental price the landlord can charge for such a property in the market. Hence;

$$\pi(\mathbf{z}; \bar{\mathbf{z}}, \mathbf{r}) = P(\mathbf{z}) - c(\mathbf{z}; \bar{\mathbf{z}}, \mathbf{r}) \quad (17)$$

Where $\pi(\cdot)$ is the profit function defining the landlord's profit per period.

To make our analysis compatible with that for the demand side of the market, let us define a function that joins all combinations of \mathbf{z} and $P(\mathbf{z})$ that return the same profit for the landlord. To do this, set $\pi(\mathbf{z}; \bar{\mathbf{z}}, \mathbf{r})$ equal to the constant π in Equation (17). Then solve for the market prices that would be required in order to realise the profit π for different levels of \mathbf{z} . Mathematically, this amounts to;

$$\phi(\mathbf{z}; \bar{\mathbf{z}}, \mathbf{r}, \pi) = \pi + c(\mathbf{z}; \bar{\mathbf{z}}, \mathbf{r}) \quad (18)$$

This function is what Rosen terms the *offer function*. Simply put, the offer function describes the rent the landlord would need to receive in order to achieve a profit of π if he were to provide his property with a level of characteristics given by the vector \mathbf{z} .

As with their counterpart, bid curves, the offer function can be illustrated as *offer curves*. Each offer curve combines rental prices and levels of attribute provision that result in the same level of profit. The left hand panel of Figure 6 plots one landlord's offer curves for attribute z_1 . The upper offer curve represents combinations of attributes and prices that would return a profit of π_2 , the middle curve combinations giving a profit of π_1 and the lower curve a profit of π_0 .

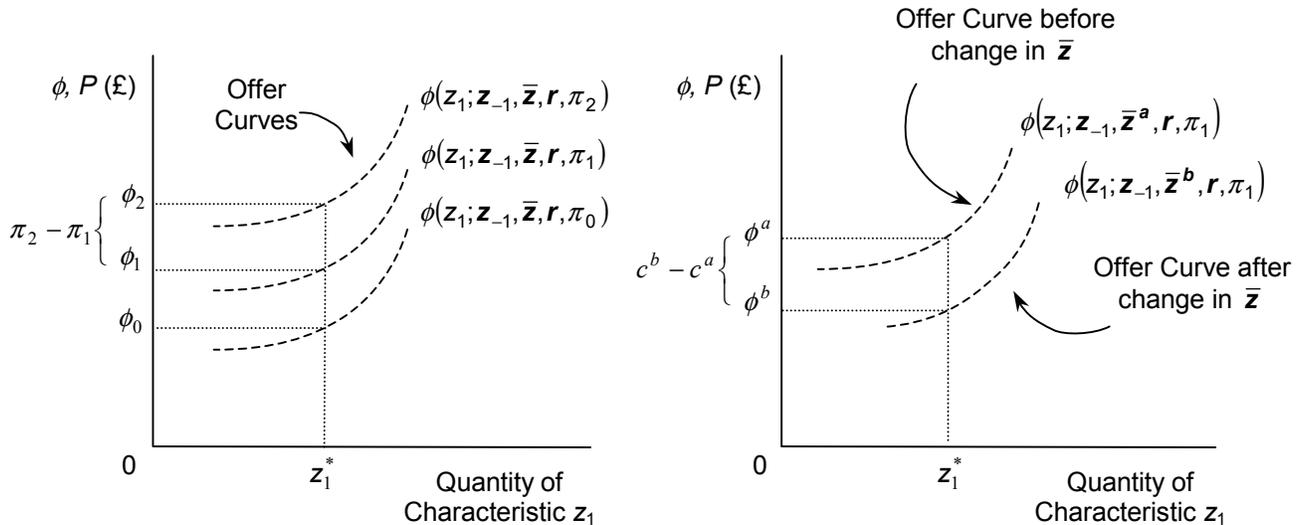
As we would expect, higher offer curves define higher levels of profit for the landlord. Take for example one particular level of provision of z_1 , let us say z_1^* . At this level of provision, the offer curves illustrated in Figure 6 evaluate to the three different offers ϕ_0 , ϕ_1 and ϕ_2 . Since \mathbf{z} , $\bar{\mathbf{z}}$, and \mathbf{r} are identical for each of these evaluations, the costs of provision must also be identical. All that changes between these offers is the level of

¹⁹ An investment such as the installation of double-glazing would require a one-off payment. Of course, it is possible to express this one-off payment as the discounted sum of a series of smaller equal-sized costs made over the expected lifetime of the investment. The increase defined by the cost function would be equal to this extra per period cost.

profit accruing to the landlord. Thus the vertical distance between two offer curves (in which the arguments entering the cost function remain unchanged) measures the difference in profit associated with the two curves. Of course, this should be evident from Equation (18).

In the left hand panel of Figure 6, therefore, the vertical distance between the middle and upper offer curves is the difference in profits associated with the two curves i.e. $\pi_2 - \pi_1$.

Figure 6: The landlord's offer curves



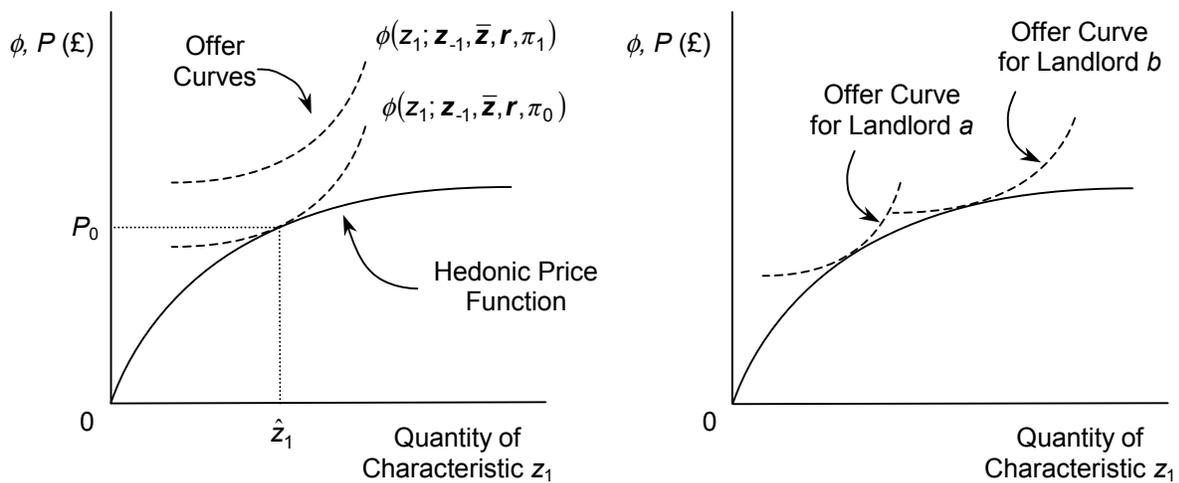
The right hand panel of Figure 6 depicts a comparable analysis, but this time we compare offer curves that differ in the level of the exogenously determined levels of \bar{z} . For example, we might imagine that the only difference between the upper and lower offer curves is the baseline level of crime in the area around the property. The upper curve is an offer curve for the landlord when faced by a low level of crime the, lower offer curve would represent the landlord's situation faced by a higher level of crime. Notice that the two offer curves are associated with the same level of profit to the landlord, π_1 . At a particular level of provision of provision of z_1 , again let us focus on z_1^* , these two curves evaluate to two different offers, ϕ^a and ϕ^b . Clearly, the change in crime has influenced the costs of the landlord in providing a property with attribute levels z_1^*, z_{-1} . For example, the lower the level of crime, the lower the level of vandalism, the less the landlord would have to spend in maintaining and repairing the property. Since the two offer curves represent the same profit to the landlord, the vertical distance between them must indicate this cost saving, i.e. $c^b - c^a$. Notice also that this cost saving does not necessarily have to result from costs incurred in the provision of attribute z_1 .

In maximising profit, the landlord seeks to provide the bundle of housing attributes that positions them on the offer curve providing the highest level of profit whilst still being compatible with reigning market prices. Similar to the choices made by households, this entails a tangency condition. In the left hand panel of Figure 7, the highest offer curve

compatible with the hedonic price function is that returning a profit of π_0 . The best this landlord can achieve, therefore, is to alter the level of attribute z_1 to \hat{z}_1 , and charge a rent of P_0 .

Again offer curves differ across landlords due to differences in purchase prices (though, for simplicity, this argument is suppressed in the analysis), the vector of parameters r and the exogenously determined levels of attributes \bar{z} . As a consequence, different landlords will choose to supply properties with different bundles of attributes. This is illustrated in the right hand panel of Figure 7. Indeed, the hedonic price function forms the lower envelope to the set of all landlords' optimising offer curves.

Figure 7: Landlord's Optimising Choices of Housing Attributes to Supply



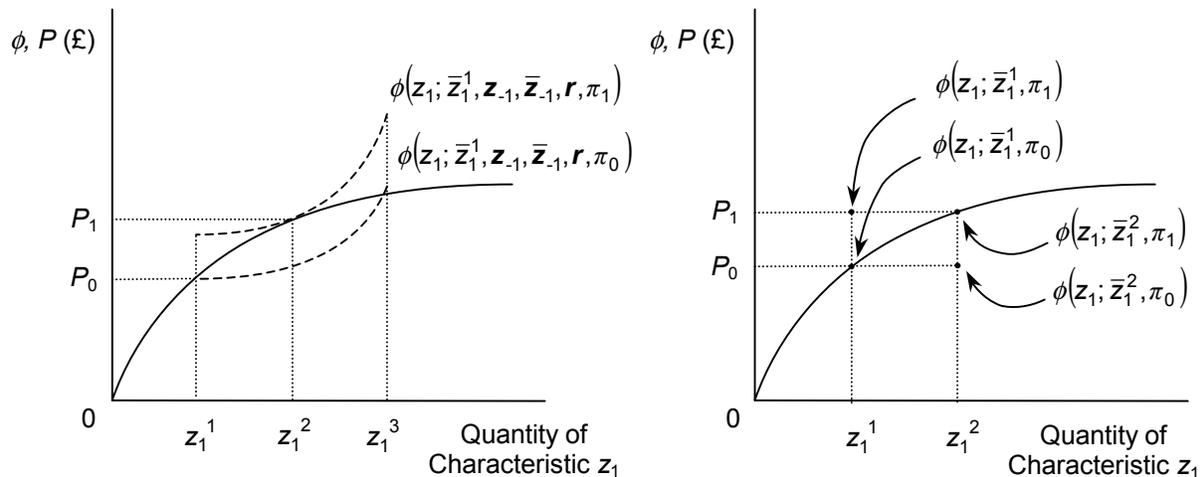
Offer curves differ slightly from their counterpart, bid curves, in so much as they will frequently be defined over quite a small range of attribute space. That is to say, for any one landlord the ability to change the attributes of the property may be relatively limited.

Let us return to our example of the peace and quiet attribute of a property. To a large extent, this is fixed for the property by exogenous factors, most notably proximity to noisy roads. The left hand panel of Figure 8 shows offer curves for a landlord whose property is exposed to an exogenously determined level of peace and quiet, z_1^1 . As such, this quantity enters the offer function where it is labelled \bar{z}_1^1 . At this level of the attribute the landlord's property would command a rent of P_0 that would reward the landlord with a profit of π_0 .

However, this is not the most profit that the landlord could earn on this property. As mentioned previously, the landlord could increase the level of peace and quiet at the property by, for example, planting trees in the front garden to act as a barrier to traffic noise and/or fitting sound-proofing double-glazing. The best the landlord could do is to increase the peace and quiet attribute of the property to z_1^2 . Here he would be able to charge a higher rent, P_1 , and would enjoy a profit of π_1 .

In the example shown in the left hand panel of Figure 8, no amount of investment will increase the peace and quiet at this property beyond z_1^3 , such that the offer function is not defined for values of z_1 in excess of this quantity. As it happens, these restrictions will not concern this landlord as he maximises his profit by providing a property with z_1^2 of the attribute.

Figure 8: Landlords' Optimising Choice when the level of provision of an attribute is constrained



However, the possibility arises that landlords may face what are termed corner solutions. That is, the profit maximising level of provision of a housing attribute may be at one of the constraints of the offer curve. To illustrate with our current example, the offer curve is constrained by the exogenously determined level of ambient noise pollution, which places a lower limit of z_1^1 on the peace and quiet of the property, and by the limitations of noise avoidance technology, which places an upper limit of z_1^3 on the peace and quiet of the property. If either of these limits were the profit maximising solution for the landlord then the simple tangential condition shown in Figure 7 would not hold.

In the extreme, landlords may have no control over the level of an attribute. In the terminology of the last paragraph the lower and upper limits for an attribute are one in the same. Take for example the proximity of the property to local amenities such as a shopping centre or school. Since the property is fixed in space there is no way in which the landlord can influence the time it would take a household living in that property to access these facilities.²⁰ In this case, the bid curves will shrink to a point above the exogenously determined level of the attribute.

²⁰ Once again, this is an example of an attribute whose value is exogenously determined and would enter the cost function and hence offer function in the vector \bar{z} .

Such a situation is illustrated in the right hand panel of Figure 8. Again to illustrate in one dimension we assume that the levels of all other property attributes do not change and for ease of exposition we suppress the vectors z_{-1} , \bar{z}_{-1} and r , from the offer function. Consider first the situation where the level of attribute z_1 is exogenously constrained to a level z_1^1 . The best this landlord can do is to charge a rent of P_0 , which returns a profit of π_0 . Charging anything lower would necessitate missing out on possible profits, charging anything greater would make it impossible to rent out the property.

The only way in which this landlord could increase profits would be if there were an exogenous change in the level of z_1 enjoyed at the property. For example, if attribute z_1 represents accessibility to the town centre, then the building of an urban tram link that passed near the property would increase the accessibility of the property and consequently increase the level of attribute z_1 . In Figure 8, this is represented by an increase from z_1^1 to z_1^2 .

If we assume that this change doesn't influence the cost of providing other property attributes then the landlord could continue to charge P_0 and earn a profit of π_0 . Of course, as illustrated in the figure, the landlord could make more profit than this by increasing the rent on the property to P^1 . Charging this rent the landlord's profits increase to π_1 .

We have dwelt on corner solutions such as those illustrated in the right hand panel of Figure 8 because such solutions typify environmental attributes and it is these attributes that are our central concern.²¹ However, it is fair to assume that in the short run the levels of all attributes are constrained in a similar manner. For example, given the hedonic price function, it may increase the profitability of a particular property if it were to possess an extra bedroom. Of course, such changes do not happen overnight. The landlord might have to employ an architect to design an extension to the property, apply for planning permission, employ builders and finally have the proposed extension constructed, decorated and furnished. We shall return to such considerations in the next chapter.

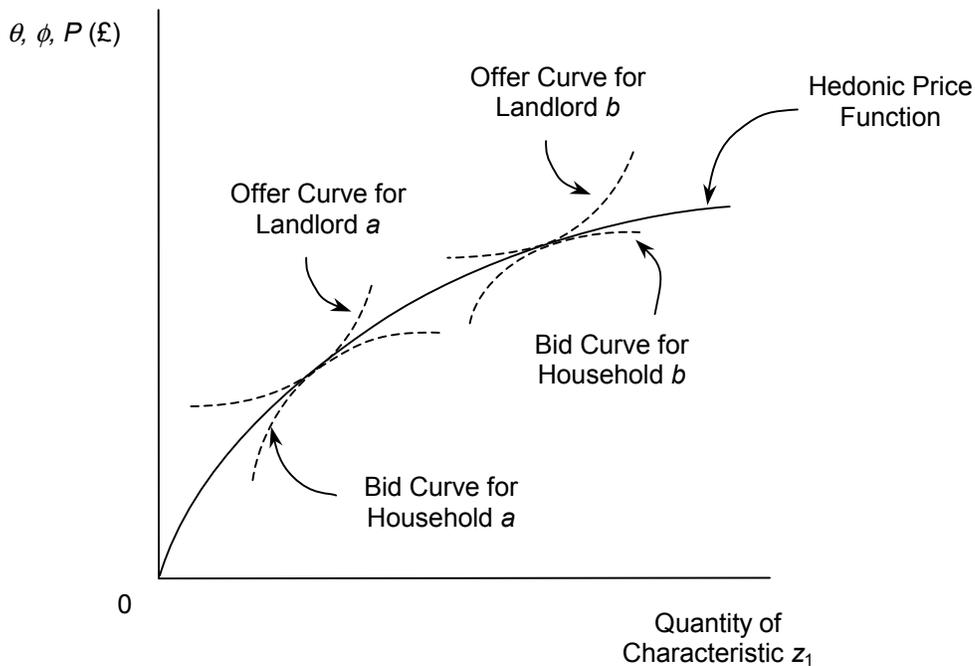
²¹ Since many environmental qualities have the properties of public goods, their level of provision tends, to a greater extent, to be determined exogenously. Further, landlords frequently have limited scope for adjusting the levels of these attributes through investment in private goods for the property.

f. The Property Market: Equilibrium

So far we have examined the choices of consumers (households) and suppliers (landlords) in the property market independently. Figure 9 presents both sets of decisions combined in the same diagram. Households define their optimal residential location by choosing a property that boasts the set of attributes that coincide with the tangent of their lowest bid curve with the hedonic price function. The household cannot increase their utility by bidding for a property with different characteristics. Simultaneously, landlords maximise their profits by choosing to supply a property with the set of attributes that allows them to move to their highest offer curve that is still compatible with market prices. Supplying an alternative set of attributes would result in the landlord receiving offers for the property that resulted in lower levels of profits.

As depicted in Figure 9, the bid curves of households and the offer curves of landlords will “kiss” along the hedonic price function. At each coincidence of bid and offer curves, a landlord and household are paired; the landlord can do no better than to accept the household’s offer who in turn can do no better than to rent the property from that landlord.

Figure 9: Choice of Property Attributes for Different Households



The situation we have described is one of market equilibrium. At the reigning market prices revealed by the hedonic price schedule, demand for properties is equal to the supply of properties and the market clears. Hitherto our analysis has been at the level of individual households and landlords. At this disaggregate level we have assumed that each individual economic agent, being only a small player in the entire market, takes the equilibrium hedonic price schedule as given. To understand how the equilibrium is reached in the first place, we need to look at aggregate demand and supply.

Let us denote by $Q^d(z)$ the aggregate market demand for properties with characteristics z . Similarly let $Q^s(z)$ be the aggregate market supply for properties with characteristics z .²² The market will be in equilibrium when the hedonic price function, $P(z)$, is such that $Q^d(z) = Q^s(z)$ for all z . At this schedule of prices, individual landlords and households will act as has been described above.

Suppose that the quantities demanded and supplied for a particular combination of housing attributes, say z^* , do not match at the prevailing market prices. In such a case, we would expect market forces to act so as to change the hedonic price schedule and bring the market back into equilibrium. In the hedonic property market, however, this will not be so simple as to only involve change in the price of properties with attributes z^* . Instead, a price change for properties with this particular bundle of attributes will induce changes everywhere in the hedonic price schedule.

For example, imagine that excess demand for attributes z^* caused the price of properties with this bundle to increase. Households, who at the old hedonic price schedule would have maximised their utility by choosing a property with these attributes, will substitute away from that particular bundle thereby increasing demand for properties with different characteristics. Similarly, landlords noting the extra profits to be made by supplying properties with attributes z^* , will modify their properties to take advantage of the higher market prices. This will induce reductions in the supply of properties with alternative bundles of attributes. The process of substitution and relocation ripples through the entire market until a new hedonic price schedule is established that reconciles aggregate demand with aggregate supply for all z .

A number of researchers (e.g. Rosen, 1974; Epple, 1987) have attempted to analytically model equilibrium in hedonic markets. To do this it is necessary to make specific assumptions concerning the various behavioural functions that determine household and landlord behaviour. Specifically, researchers must assert a particular functional form for the utility function of households, $u(\cdot)$, and the cost function of landlords, $c(\cdot)$. Further they have to make assumptions concerning the distribution in the population of household characteristics, s , and landlord characteristics, r . Given specific forms for each of these different functions it should be possible to solve for an expression that gives the equilibrium hedonic price function. This expression will be a function of the arguments in the underlying functions. Hence, using such models it is possible to investigate how changes in the underlying arguments influence the hedonic price schedule.

Unfortunately, the complexity of the hedonic market is such that one must make very restrictive assumptions concerning the various functions in order to end up with an expression for the hedonic price schedule that is reasonably tractable. In general, therefore, research has concentrated on empirical analyses of hedonic markets.

²² If the attributes of properties cannot be influenced by landlords (as is probably true for many environmental attributes and for most property attributes in the short run) then the supply of properties with certain characteristics will be determined by the state of the current housing stock in the urban area. In this case aggregate supply will be perfectly inelastic and the equilibrium hedonic price schedule will be determined solely by the level of aggregate demand for properties with different characteristics.

In many ways, the problem for empirical analysis is the reverse of analytical research. Rather than assuming the functional forms of the underlying behavioural functions and working through the problem to solve for the hedonic price schedule, researchers estimate the hedonic price schedule from real world market data and work back through the problem to discover the form of the underlying behavioural functions. We shall return to how this might be achieved in Chapter 3.

CHAPTER 2. WELFARE MEASUREMENT IN HEDONIC MARKETS

a. Introduction

Our interest in hedonic markets stems from the fact that environmental quality can be counted amongst the attributes of a property. Whilst the various attributes which make up environmental quality (e.g. peace and quiet, clean air, access to recreational areas etc.) are frequently not directly traded in markets, hedonic property markets provide an indirect means by which households can express preferences for such goods.

For example, though a household may wish to increase the quality of the air they breathe, there is no independent market in which they could express this preference. Households couldn't, for instance, call up a firm and purchase a month's supply of clean air. On the other hand, the property market provides one channel through which households can express preferences for environmental quality. If a household wishes to improve the quality of the air they breathe they can do so by purchasing a property located in a less polluted area.

Like the other attributes of a property, differences in environmental quality will be reflected in differences in the price paid for a property. Indeed, with information on the implicit price of environmental quality and the residential locations chosen by different households, analysts have access to information from which they can deduce household preferences for environmental goods.

The search for these underlying preferences is the key goal of empirical analysis of hedonic market data. Specifically, establishing the structure of preferences makes it possible to estimate the impact on economic welfare of a change in environmental quality. We shall return to the issue of estimating household preferences from empirical data in the final chapter.

In this chapter we show how the theory of hedonic markets outlined in Chapter 1 allows us to describe the welfare effects of changes in environmental quality.

b. The Hedonic Market and Changes in Environmental Quality

Before we embark on an analysis of welfare measures it is worth developing some intuitive understanding of the impact a change in environmental quality might have in the property market. Let us consider the impacts of an environmental improvement such as a reduction in noise pollution, a fall in levels of crime or an increase in air quality. Of course this change may be a relatively minor or alternatively may represent a significant environmental improvement. Also, the improvement may take place uniformly across the urban area or be restricted to specific neighbourhoods.

Marginal, localised changes may have little impact on the property market as a whole. Of course landlords will be able to increase the rent they charge on properties in the improved area since those properties now enjoy a higher level of environmental quality. As a result, the household's living in those properties will no longer be at their optimal location. Indeed, they could well elect to move to a new house possessing the original

bundle of characteristics enjoyed at their previous property. In the real world, however, there are considerable transaction costs associated with moving house. For relatively small changes in rent, therefore, households may elect to remain where they are. Nevertheless, in the long run, perhaps prompted by other changes in the property market, we would expect households to move to a property with an optimal bundle of characteristics.

If the environmental change is neither marginal nor localised then the pattern of changes in the property market may be far more complex. In the simple case discussed previously, the environmental change is not substantial enough to significantly effect the market clearing implicit prices; the hedonic price function for properties in that market is unaffected by the change.

Of course, if the environmental improvement is sufficiently large (in degree and/or geographical area) then this is unlikely to be true. As in markets for any good, changing the conditions of supply and demand will change the market-clearing prices. Naturally for goods traded in large, possibly world markets localised changes in the conditions of supply and demand are unlikely to effect prices. In property markets, on the other hand, the reverse is true. Property markets are inherently constrained to specific geographical regions. As such even relatively small changes in the conditions of supply in one part of that region may well effect the market clearing implicit prices across the whole market. Indeed, we might expect property markets to be particularly responsive to even relatively minor changes.

Bartik (1988) provides a detailed description of how an environmental improvement might impact on property rents, location choices and housing supply. He envisages an environmental improvement in one part of a hypothetical property market. Obviously, the improvement increases the environmental quality of properties in the affected area. Now, if the hedonic price function were unaffected then, as described before, landlords would be able to increase the rents they charge for their properties. However, imagine now that the improvement were sufficiently large to precipitate a shift in the hedonic price function. That is, the added supply of environmental quality in the market would, in general, necessitate a reduction in price per unit (implicit price) of environmental quality across the entire market in order to ensure the market cleared. For any one property, therefore, the change in rent will be determined by the opposing forces of a location-specific environmental improvement that would tend to push rents up and a market-wide increase in supply of environmental quality that would tend to push rents down. Thus even though some properties may not be directly effected by the environmental improvement, market adjustments may well result in changes in their rental value.

Of course the overall impact on the hedonic price function will not be limited to adjustments in the price of environmental quality. It seems likely that a number of concomitant effects will cause shifts in the supply and demand for housing characteristics. For a start, demand for property characteristics that are substitutes for the environmental attribute will decline. For instance, demand for double-glazed properties will decline in an area in which noise pollution has been reduced. Similarly, demand for complementary attributes will increase. For example, a reduction in air pollution might increase demand for houses with gardens. The implicit prices for these substitutes and

complements will themselves have to adjust in order to ensure that the demand for these attributes is balanced by the supply.

Further, in response to the shifts in the hedonic price function, households will no longer be at their optimal residential location and will choose to move to a new property. Indeed, we would expect that landlords at certain locations would find that the characteristics of the households wanting to rent their property would change. For example, reductions in the implicit price of environmental quality will encourage lower income households to demand properties in areas that they previously could not afford. Such that at any given level of environmental quality, there will be an increase in demand from lower income households. Bartik (1988) hypothesises, that lower income households will have lower demands for other housing characteristics and landlords will change their levels of investment in properties to maximise their profits. For areas that experience large increases in environmental quality the reverse may be true. High income households will be attracted to the area and their higher demands for other property characteristics will encourage landlords to invest in property improvements that will increase their rental value.

It is evident that the overall change in the hedonic price function and the resulting change in rents and locational choice are extremely complex. For any one property, the eventual rental value will not be determined solely by the change in environmental quality experienced at that location. Instead it will be determined by the complex interaction of supply and demand across the entire market¹. For our purposes it is sufficient to note that a change in environmental quality will lead to a shift in the hedonic price function.

c. Measuring Changes in Economic Welfare in Hedonic Markets

Our goal in analysing data from hedonic markets is to establish how changes in environmental quality impact upon economic welfare. Of course, before we embark on showing how this might be achieved, it is essential that we establish exactly what is meant by the term ‘a change in economic welfare’.

Essentially what we are seeking to measure is how greatly changes in environmental quality change the well-being of economic agents in society. In terms of the property market we have defined these economic agents as *households* and *landlords*. Further, we have defined household well-being as the *utility* they derive from their choice of residential location and expenditure on other goods, whilst landlord well-being is defined as the *profits* they realise from rental of their property.

For landlords then, the effect on economic welfare resulting from a change in environmental quality can be measured as the change in their profits (Δp) from renting out a property.

For households, the measure of change in economic welfare is not so obvious. Ideally, we would want to measure the change in utility (ΔU) that the household experienced as a

¹ It is of little surprise that economists have had difficulty developing analytical models that adequately reflect the complexity of these adjustments to the hedonic price function.

result of the change in environmental quality. Of course, describing changes in welfare as changes in utility is merely a theoretical construct. It is not possible to independently measure a household's level of utility before and after a change in environmental quality, nor is it possible to ask a household to report their change in welfare in units of utility.

Instead, economists have posited an alternative measure, a household's own monetary valuation of the change in welfare they experience. A monetary measure of welfare change has a number of advantages chief amongst which is that it can be summed across households to form an aggregate measure that can be used in cost/benefit analysis.

In the following discussion we focus on one such monetary measure known as a *compensating measure* of welfare change. Compensating measures take the current level of household utility as a baseline.

- For an environmental improvement, the compensating measure would be the maximum quantity of money that the household would willingly give up in order to ensure that they enjoyed the environmental improvement. This is often referred to as the household's maximum willingness to pay (WTP) to achieve an improvement.
- For a fall in environmental quality, the compensating measure would be the minimum amount of money that the household would accept in order to endure the deterioration in environmental quality. This is often referred to as the household's minimum willingness to accept (WTA) compensation for a deterioration

With these measures in mind, let us consider a property market and examine the welfare impacts of a change in environmental quality. As shall become apparent in the following two sections, this is not as straightforward as might be hoped. It turns out that there are a number of ways in which the change in economic welfare might be evaluated; each evaluation differing according to the assumptions that are made concerning the response of households and landlords to the change in environmental quality. As might be expected, the fewer assumptions we make, the more comprehensive the measure of the welfare change. At the same time, however, the fewer assumptions made the greater the informational requirements involved in calculating the welfare measure.

d. Changes in Economic Welfare for Households

Let us assume that the property market we are considering is in equilibrium. In this market both landlords and households are assumed to have made optimal decisions; landlords can't improve profits by altering the characteristics of their property and households can't increase their utility by choosing to rent a different property.

In this market we shall denote the original equilibrium hedonic price function by $P^b(z)$, where the superscript b indicates that this is *before* any changes in conditions in the hedonic market. Following a change in environmental quality, the market settles at a new hedonic price function that we shall denote $P^a(z)$. Once again the superscript a indicates that this is *after* the change in conditions in the hedonic market.

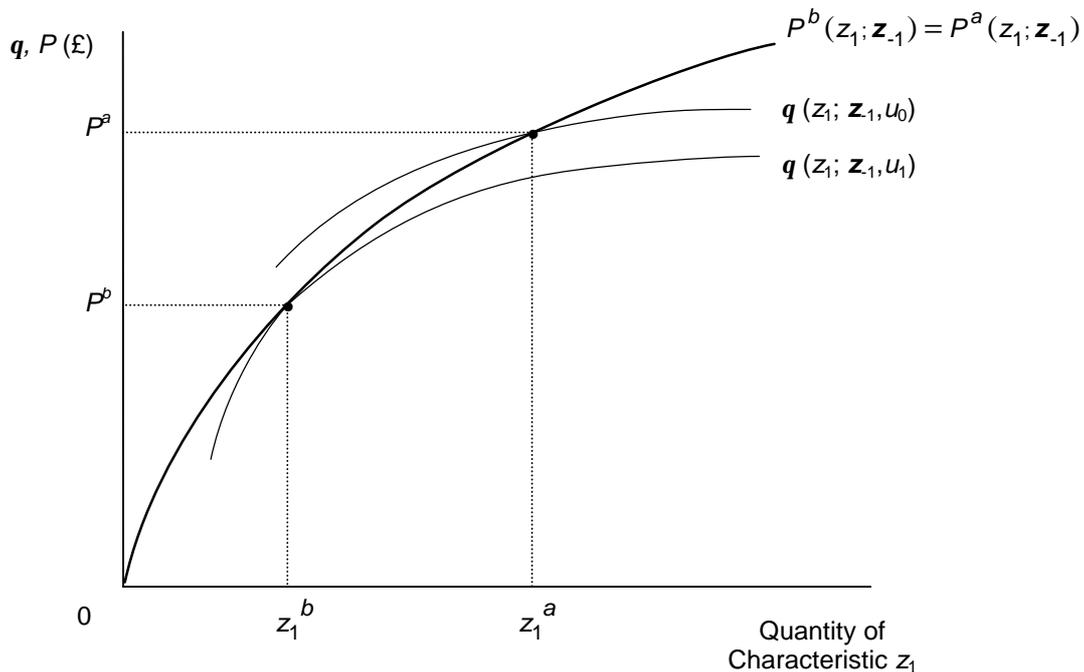
i. Household welfare changes from a localised environmental improvement

To begin with let us consider the welfare impact of a localised environmental improvement. As discussed in the previous section such a change should not impact on the property market as a whole and the hedonic price function will not need to adjust in order to clear the market. Thus, in this case, $P^a(z) = P^b(z)$.

Let us focus on just one property located in the area experiencing the environmental improvement. If we designate attribute z_1 to be the level of environmental quality, then the initial level of environmental quality at this property can be represented by z_1^b . As illustrated in Figure 10, at this level of environmental quality the property commands a rental price of P^b .

The household choosing to reside in this property will have a bid curve tangential to the hedonic price function at this level of environmental quality. In Figure 10, this utility maximising choice places the household on their lowest bid curve compatible with the hedonic price function and results in a level of utility u_1 .

Figure 10: Change in household welfare from a localised change in environmental quality and costless moving



Now, the exogenous improvement increases the environmental quality of the property from z_1^b to z_1^a (where once again b superscript stands for *before* and a superscript stands for *after*). Of course, improving the attributes of the property enables the landlord to charge a higher rental price. Indeed, the rent on the property following the environmental

improvement would increase from P^b to P^a . Clearly, this represents a benefit to the landlord but we shall postpone a discussion of this welfare gain until the next section.

What then are the welfare impacts on the household residing at this location? Clearly, the household enjoys an improvement in environmental quality, however this is accompanied by an increase in rental price. As illustrated in Figure 10, the household will find itself at a less than optimal residential location. Indeed, continuing to live at the property would mean their level of utility would fall from u_1 to u_0 .

Since the hedonic price function has not changed, we know that the household's optimal choice of property would be one boasting the original level of environmental quality at that property. Indeed, if we assume that moving house is costless then the household would be best off simply moving to a property with identical characteristics to their present property, except with the level of environmental quality enjoyed previous to the improvement. Moving to such a property would return them to their original level of utility, u_1 . Overall then, under the assumption of costless moving, the environmental improvement will have no impact on the welfare of households.

In the real world, however, there are considerable transaction costs associated with moving house. Incorporating such transaction costs complicates the analysis. For a start, we should note that households only envisage living in any one property for a limited period of time. At the end of such a period, the household foresees that changes in their characteristics (e.g. marriage, the birth of children, retirement etc.) will have changed the nature of their preferences for properties. We can assume, therefore, that the household foresees a series of property relocations each incurring a transaction cost. Consequently, we can express the sum of these payments as an equivalent per period cost, TC , such that the per period price of living in a particular residential location is in fact the market rental price plus this added cost (i.e. $P(z) + TC$). In effect, foreseeable changes in preferences allow the household to write-off moving costs over the duration of their expected tenureship of a series of properties.

Moving house in response to an exogenous change would mean incurring unexpected transaction costs at an earlier date causing the value of TC to increase. Let us call this added per period transaction cost tc . Rather than follow through the complex arguments that including transaction costs entails, let us simply note that two possibilities present themselves;

- If the benefits of moving outweigh the moving costs then the household will relocate to a new property with the attributes of their original choice.² In welfare terms, the household ends up enjoying the same level of utility as prior to the environmental improvement but are worse off by an amount equal to the costs of moving. Thus the quantity tc , measures the per period welfare loss of the environmental improvement.

² More correctly, moving house in response to an exogenous change causes the value of TC to increase by an amount labelled tc . It would appear to a household considering a change of location as if the per period price of property rental had shifted upwards. Since the per period price of property rental is different from that faced in making their original decision, it is unlikely that the household's optimal response will be to relocate to a property with identical attributes to those enjoyed at their previous optimal choice. We do not follow these considerations further here. The interested reader is referred to Freeman (1993, p398-400).

- If, on the other hand, the benefits of moving do not exceed the transaction costs then the household will decide to remain in their original, though now sub-optimal, residential location.³ Clearly, the loss in welfare associated with remaining in this improved property paying a higher rent is not as great as the transaction costs. Consequently, tc must represent an upper bound on the welfare loss to the household.

To summarise, the environmental improvement will result in households in the improved region being at less than optimal residential locations. If we ignore transaction costs then households will relocate to properties with identical attributes as those enjoyed at their original residential locations. The environmental improvement will have no impact on their welfare. If we include transaction costs then we can assume that the environmental improvement may cause households to move property earlier than they would otherwise have anticipated. Such premature relocation would increase the equivalent per period costs of moving house by a quantity tc . This quantity must represent an upper bound on the household's welfare loss resulting from the environmental improvement since they could always pay this amount so as to relocate to a property offering the level of welfare enjoyed prior to the change.

If the total number of properties in the market is labelled H then the small subset of properties affected by the environmental improvement can be labelled H_1 . Further, if we index all the households in the market by $h = 1$ to H , then the welfare change experienced by households from a localised environmental improvement can be expressed;

$$\sum_{h \in H_1} -tc_h \leq W_H \leq 0 \quad (19)$$

Where W_H is the total welfare change experienced by households in the market and the expression $h \in H_1$ tells us to include only households living in the H_1 properties affected by the environmental improvement.

³ Of course remaining at a sub-optimal residential location is only a short-term solution. Indeed, at some point in the future we would expect the household to move to an optimal residential location. Two possible stimuli may precipitate such a move;

- First, it can be assumed that a household foresees that changes in their characteristics (e.g. marriage, the birth of children, retirement etc.) will change the nature of their preferences for properties. At some point in the future, therefore, the household would expect to move and will have anticipated the transaction costs of such relocation. As such, at some point in the future transaction costs will no longer represent a barrier to relocation.
- Second, unforeseeable changes such as further exogenous changes in the attributes of the property or unexpected changes in the household characteristics (e.g. becoming unemployed) may tip the balance in favour of relocation. In other words, cumulative unforeseen changes may mean that the benefits of relocation outweigh the costs of moving.

In either case, the household will have incurred a welfare loss from the change in environmental quality that must be less than the full transaction costs of moving. Hence we can always take the value tc as the upper bound of this welfare loss.

ii. Household welfare changes from a non-localised environmental improvement

Imagine now, that we are dealing with an environmental improvement that has more than a purely localised impact. If the change we are considering represents a major improvement and/or is widely spread across the urban area then the consequences for the property market may extend beyond a simple increase in the price of affected properties.

One possibility is to assume away these wider implications and use Equation (19) to measure the welfare change of households. In economics such a measure would be described as a partial equilibrium solution since it does not allow for the complex pattern of changes in the hedonic price function and choices of residential location that would allow the market to come back into a state of general equilibrium.

In this section we discuss welfare measures that account for these general equilibrium effects. Perhaps surprisingly, therefore, we begin this section by introducing another partial equilibrium measure of household welfare. This simple measure will prove to be of major importance since it can be shown to represent a lower bound to the entire general equilibrium welfare impact experienced by both households and landlords. But we shall return to demonstrate this anon.

Figure 11 presents the situation facing a household living in the area witnessing an environmental improvement. At the original level of environmental quality the household faces the old hedonic price schedule, $P^b(z)$, and maximises its utility by choosing a property with a level of environmental quality indicated by z_1^{bo} . Here we have expanded the notation such that the superscript *bo* indicates that this is the quantity chosen *before* the change in environmental quality in the households *old* choice of property. At this point, the household reaches its lowest bid curve that is still compatible with the prices it faces in the market, $q(z; u_1)$.⁴ The household's WTP or bid, indicated by q^{bo} , is equal to the market price, P^{bo} , and the household enjoys a level of utility labelled u_1 .

An exogenous change increases the environmental quality enjoyed at this location to z_1^{ao} , where the superscript *ao* indicates that this is the environmental quality *after* the change in the household's *old* choice of property.

Since we are now considering a non-localised change we would imagine that the hedonic price function would shift in response to this environmental improvement. However, for the moment, we shall ignore this general equilibrium response. Further we shall consider the situation in which landlords continue to charge the rental price associated with old level of provision of z_1 . In this case, the household has effectively been given the benefits that come from living in a location with an improved environment. Indeed, at this location paying the original level of rent for that property they would find themselves on the bid curve $q(z; u_2)$ realising a higher level of utility labelled u_2 .

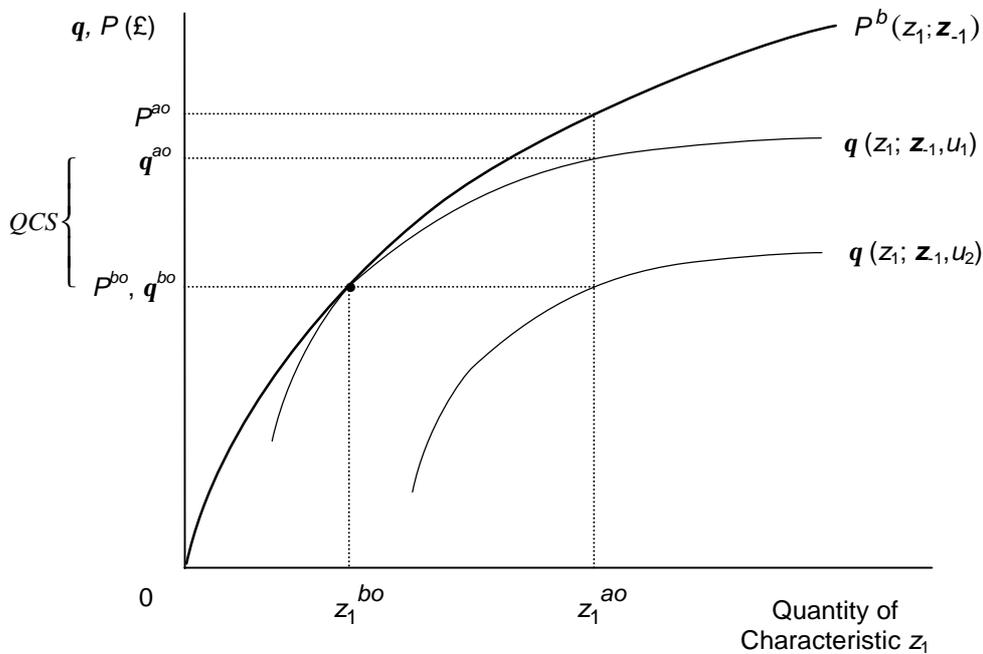
One possible compensating monetary measure of the welfare that the household gains from this improvement is the amount of money that if taken away from the household whilst living in the property in the improved location would make them as well off as

⁴ Income and socioeconomic characteristics have been suppressed to simplify notation.

they had been previous to the change. In other words, the household's WTP to achieve the improvement in environmental quality.

This measure can be shown simply in Figure 11. The bid curve on the diagram traces out all combinations of WTP and levels of environmental quality that result in the household enjoying a level of utility labelled u_1 . Of course this is also the level of utility that the household realised prior to the environmental change by locating at their optimal residential location. To achieve this level of utility the household was willing to pay q^{bo} . Following the environmental improvement the household would be willing to pay q^{ao} to achieve the same level of utility. A measure of the household's WTP for the change in environmental quality is the difference between these two amounts.

Figure 11: The Quantity Compensating Surplus measure of the welfare change resulting from an improvement in environmental quality



This amount has been termed the *quantity compensating variation*, by Palmquist (1988). However, following Freeman's definitions (see Freeman, 1993; p 48-9) this is probably best thought of as a *compensating surplus* measure since it allows for no adjustment in household residential location following the change in environmental quality. Hence here we label this amount as the *quantity compensating surplus (QCS)*. This amount is shown graphically in Figure 11 and can be stated mathematically as;

$$QCS = ?(z_1^{ao}, z_{-1}^{bo}; y, s, u_1) - ?(z_1^{bo}, z_{-1}^{bo}; y, s, u_1) \quad (20)$$

Since, the *QCS* measure assumes there are no adjustments in the hedonic property market the welfare change is assumed to impact only households in the affected area. Indeed, using this measure, the total welfare impact of the environmental improvement is given by;

$$W_H = \sum_{h \in H_1} QCS = \sum_{h \in H_1} ?(z_{1h}^{ao}, z_{-1h}^{bo}; y, s, u_{1h}) - ?(z_{1h}^{bo}, z_{-1h}^{bo}; y, s, u_{1h}) \quad (21)$$

Notice that the informational requirements of the *QCS* measure are relatively undemanding. To evaluate W_H using this measure, a researcher would simply have to know the environmental quality at all affected properties before and after the improvement and be able to evaluate the bid function for each household at these two values of environmental quality.

However, the *QCS* measure of welfare change is relatively restrictive in the assumptions it makes concerning how the market and the economic agents in the market react to a change in environmental quality. All the benefits of the change accrue to households occupying properties in improved locations since landlords are assumed not to change property rents. Of course this is most unrealistic; a landlord is hardly likely to remain charging the same rent (P^{bo}) when the market price for a property with that level of environmental quality is actually (P^{ao}). Further, 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, our analysis in the previous section indicates that an increase in environmental quality in the urban area may well precipitate a shift in the hedonic price function.

Figure 12 shows just such a shift. The environmental improvement has led to an adjustment in the hedonic market that has reduced the price of property at any given level of environmental quality. As described earlier, the hedonic function *after* this adjustment is labelled $P^a(z)$.

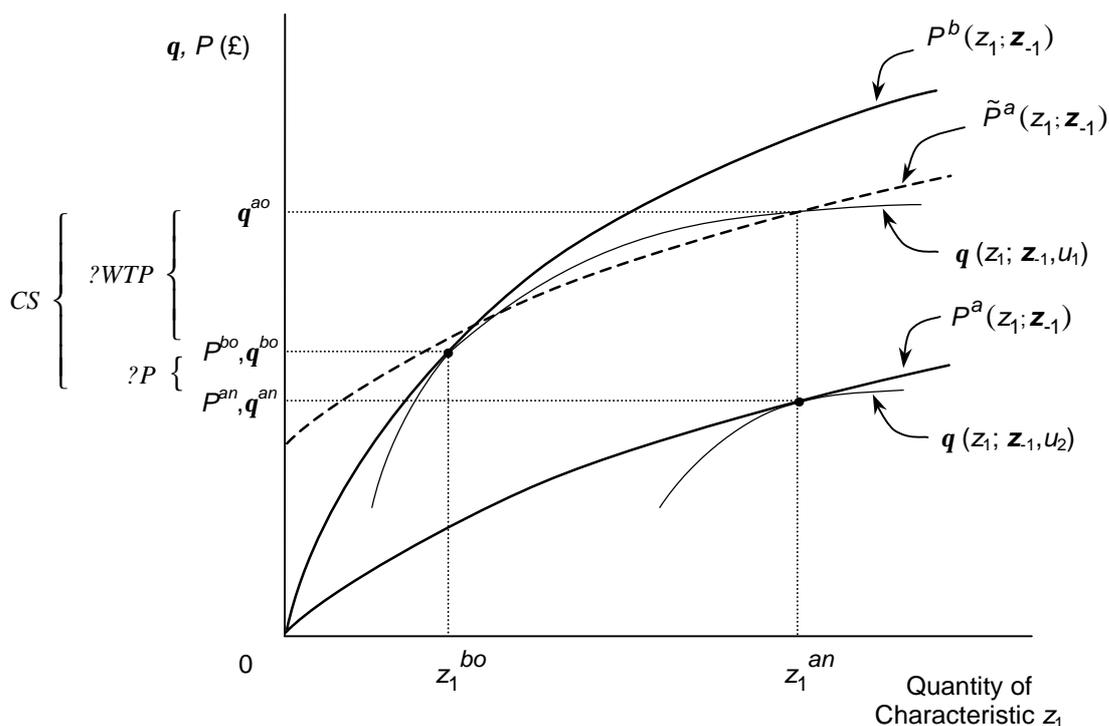
It is important to note that since the hedonic function has changed the environmental improvement has an impact on all households in the property market. Indeed, as a consequence we would expect each household in the property market to adjust to the new hedonic price function by choosing a new residential location. As before we assume they move to the property amongst those that they can afford which provides them with the highest level of utility.

In the figure this is illustrated for one household as the tangent of a bid curve and the new hedonic price function. This particular household will move to a new property with a level of environmental quality given by z_1^{an} . Where the superscript *an* indicates that this is the level of environmental quality enjoyed *after* the change at their *new* choice of residential location. By moving property, the household moves on to a lower bid curve and manages to achieve a higher level of utility, u_2 .

Since, households are allowed to respond optimally to the changes in the hedonic market by moving residential location, our previous measure of welfare change, the QCS , is no longer an adequate measure of the benefits of an environmental improvement. Figure 12 can be used to illustrate a second compensating measure of welfare change that accounts for household relocation.

As the household has been made better off we assume that they would be willing to pay out some money to ensure that they continued to enjoy their new level of well-being, u_2 , rather than returning to that enjoyed prior to the environmental improvement in their original location, u_1 . Let us constrain the household to remain at their new choice of property. Thus the compensating monetary measure we seek is the amount of money that once taken away from the household in their new residential location would return them to their original level of well-being.

Figure 12: The Compensating Surplus measure of the welfare change resulting from an improvement in environmental quality



To illustrate this measure, examine Figure 12. Here the change in the household's income that would result from paying out a compensating monetary measure, is shown as a vertical shift in the hedonic price function. In effect, paying out money is equivalent to making all properties more expensive⁵. The maximum amount the household would be

⁵ Readers familiar with the illustration of welfare measures in diagrams with indifference curves and income constraints will recognise this procedure. Indeed, this parallel is made explicit by remembering that

willing to pay to ensure the change in environmental quality whilst constrained to remain at their new residential location, will be the amount that shifts the hedonic price function to the point where it intersects the original bid curve.

As illustrated in Figure 12, the vertical distance between the hedonic function $P^a(\mathbf{z})$ and the hedonic function as it would appear to the household once it had paid out its maximum WTP, $\tilde{P}^a(\mathbf{z})$, gives a second measure of welfare change. This distance is the *compensating surplus* (CS) measure of the household's welfare change described by Bartik (1988).

It can be shown that this CS measure can be decomposed, in an intuitively appealing manner, into two separate values. The first value is the household's WTP for the change in housing attributes. That is, the difference between the household's WTP to achieve a level of well-being u_0 , at the old and new residential locations ($DWTP$)⁶. The second value is simply the difference in rental payments at the old and new residential locations (DP). In mathematical terms, therefore, CS can be written as;

$$CS = \Delta WTP - \Delta P = \left(\mathbf{q}(z_1^{an}, \mathbf{z}_{-1}^{an}; u_1) - \mathbf{q}(z_1^{bo}, \mathbf{z}_{-1}^{bo}; u_1) \right) - \left[P^a(z_1^{an}, \mathbf{z}_{-1}^{an}) - P^b(z_1^{bo}, \mathbf{z}_{-1}^{bo}) \right] \quad (22)$$

Since all households are assumed to relocate in response to the shift in the hedonic price function the total welfare benefits of the environmental improvement will include a measure for each of the H households in the urban area;

$$W_H = \sum_{h=1}^H CS_H = \sum_{h=1}^H \left(\mathbf{q}(z_{1h}^{an}, \mathbf{z}_{-1h}^{an}; u_{1h}) - \mathbf{q}(z_{1h}^{bo}, \mathbf{z}_{-1h}^{bo}; u_{1h}) \right) - \left[P^a(z_{1h}^{an}, \mathbf{z}_{-1h}^{an}) - P^b(z_{1h}^{bo}, \mathbf{z}_{-1h}^{bo}) \right] \quad (23)$$

Notice that in comparison with the QCS measure, evaluating the CS measure of welfare change imposes far greater informational requirements on the researcher. Not only must the researcher be able to evaluate the bid function, but also predict how the hedonic price function will adjust in response to the environmental improvement. Further, the researcher must anticipate the characteristics of the property that each household will choose to rent in response to the new hedonic price function. If the welfare evaluation is

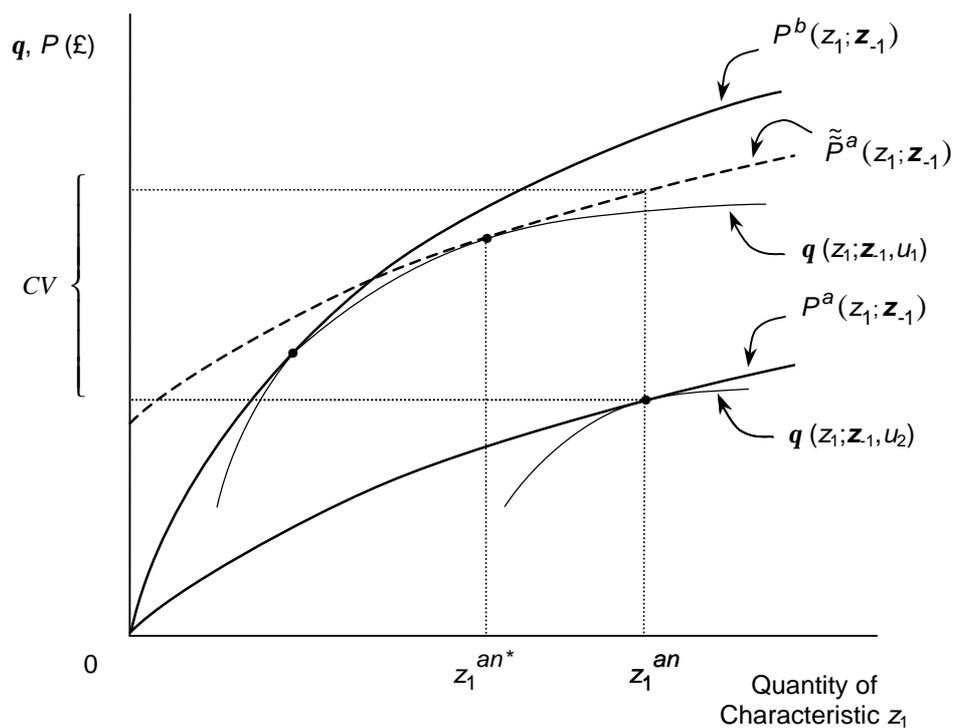
the bid curve and hedonic price function are simply inversions of corresponding indifference curves and income constraints (see Chapter 1).

⁶ This is similar to though not the same as the QCS measure described above, but here the household is no longer constrained to the level of environmental quality at their original location. Rather the household selects a new level of environmental quality by selecting a new property which maximises their well-being in response to the new hedonic price function.

to be carried out prior to the environmental improvement, as would be the case in a cost-benefit analysis, these requirements are so onerous as to make the measure practically impossible to evaluate in the real world.

For the sake of completeness we present one further measure of household welfare change. It transpires that even Bartik's *CS* measure of household welfare change is not the most comprehensive measure. In paying out the amount *CS* the household is experiencing a change in income. As their income changes, their optimal choice of residential location will also change. However, in measuring *CS* we have constrained the household to remain in the same residential location. If we relax this constraint then the household can respond optimally by changing their location in response to a change in income. Indeed, allowing the household to respond optimally means that they would be able to pay out a greater amount to achieve the improvement in environmental quality⁷.

Figure 13: The Compensating Variation measure of the welfare change resulting from an improvement in environmental quality



In Figure 13 we have again illustrated the change in income that would result from paying out a compensating measure as a vertical shift in the hedonic price function.. The

⁷ As Palmquist (1986) points out, whenever, we release a constraint on household behaviour we increase their ability to react optimally, thus increasing the quantity of money they would be willing to pay to secure an improvement in environmental quality.

maximum amount the household would be willing to pay to ensure the change in environmental quality will be the amount that shifts the hedonic price function to the point where it is just tangent with the original bid curve. The point of this tangency would determine the characteristics of the property that the household would decide to rent if it were forced to pay out its maximum willingness to pay to achieve the improvements in environmental standards. We denote the characteristics of this property z^{an*} .

As illustrated in Figure 13, the vertical distance between the hedonic function $P^a(z)$ and the hedonic function as it would appear to the household once it had paid out its maximum WTP, $\tilde{P}^a(z)$, gives a third measure of welfare change that we shall identify as the *compensating variation (CV)*. This is the measure presented in Palmquist (1986).

CV is the most comprehensive measure of welfare change since it allows the household to react optimally in adjusting to changes in the prices it faces in the market and in adjusting to changes in its own income. The *CV* measure of a welfare change resulting from an improvement in environmental quality will always be greater than the *CS*. However, the informational requirements of the *CV* measure are even greater than those of the *CS* measure. As a consequence we do not consider this measure further.

The various measures of household welfare discussed in this section are summarised in Table 1.

Table 1: Measures of household welfare change

Welfare Measure	Description	Computation of Total Welfare Change for Households	Informational Requirements
Localised:			
No Moving Costs	<ul style="list-style-type: none"> No shift in hedonic Households incur no transaction costs in moving property 	$W_H = 0$	<ul style="list-style-type: none"> None
Moving Costs	<ul style="list-style-type: none"> No shift in hedonic Households incur transaction costs in moving property 	$\sum_{h \in H_1} -tc_h \leq W_H \leq 0$	<ul style="list-style-type: none"> Only affected households Increase in equivalent per period transaction costs
Non-Localised:			
Quantity Compensating Surplus	<ul style="list-style-type: none"> Hedonic shifts Landlords do not change rental on properties Households remain in their original properties 	$W_H = \sum_{h \in H_1} \left[q(z_{1h}^{ao}, z_{-1h}^{bo}; u_{1h}) - q(z_{1h}^{bo}, z_{-1h}^{bo}; u_{1h}) \right]$	<ul style="list-style-type: none"> Only affected households Environmental quality at each affected property before and after improvement Household bid function
Compensating Surplus	<ul style="list-style-type: none"> Hedonic shifts Landlords change property rents in accordance with the new hedonic Households relocate to optimal residential locations 	$W_H = \sum_{h=1}^H \left(q(z_{1h}^{an}, z_{-1h}^{an}; u_{1h}) - q(z_{1h}^{bo}, z_{-1h}^{bo}; u_{1h}) - \left[P^a(z_{1h}^{an}, z_{-1h}^{an}) - P^b(z_{1h}^{bo}, z_{-1h}^{bo}) \right] \right)$	<ul style="list-style-type: none"> All households Hedonic before and after change Environmental quality at each affected property before and after improvement Households choice of residential location in response to new hedonic Household bid function

e. Changes in Economic Welfare for Landlords

So far we have considered only the demand side of the market. A comprehensive measure of the welfare change resulting from an exogenous environmental improvement should also take account of changes in the profits realised by landlords.

As Bartik (1988) points out, there are four reasons why we would expect a landlord's profits to change after a change in environmental quality;

- If environmental quality at the property's location changes, the property's rental value will change even if the overall hedonic price schedule does not shift
- Environmental quality changes may affect a landlord's costs (e.g. an increase in air pollution may necessitate more frequent cleaning of the property).
- Any shift in the hedonic function resulting from the environmental improvement affects rents received by all landlords, even those whose property did not directly experience a change in environmental quality
- Landlords may respond to all these changes by altering the levels of attributes associated with their property. In so doing they will alter the rental price of the property and also the cost of supplying this property to the market.

As with the discussion for households, we shall work from less comprehensive measures of landlords' welfare change through to a fully comprehensive measure.

i. Landlord welfare changes from a localised environmental improvement

To begin with let us consider the welfare impact of a localised environmental improvement. As before, such a change is insufficient to provoke a change in the hedonic price function. This then represents our first assumption.

- **Assumption 1:** The *environmental improvement is localised* and hence does not change the market clearing hedonic price function.

Further, let us assume that the level of this environmental attribute at any property is entirely determined by exogenous factors.

- **Assumption 2:** The *landlord cannot independently influence the property's environmental quality*. It is entirely determined by exogenous factors.

Assumption 2 results in the corner solution discussed in relation to the right hand panel of Figure 8. A similar diagram is reproduced here as Figure 14 where z_1 represents levels of environmental quality. Since the landlord is unable to alter the level of environmental quality through his own actions, the offer curves in Figure 14 reduce to points above the exogenously determined level.

Let us focus on the property of one landlord in the area experiencing the environmental improvement. Initially, the landlord's property enjoys a level of environmental quality z_1^b , where, once again the b superscript indicates that this is *before* the environmental improvement. Since this is supplied without cost to the landlord, the quantity $\bar{z}_1^b = z_1^b$ is

the baseline level of environmental quality. This quantity enters the cost and thence offer functions as an element in the vector \bar{z} .

Given the hedonic price function $P^b(z)$, the best the landlord can do is move to the point labelled X , coinciding with the offer curve $f(z; \bar{z}_1^b, \bar{z}_{-1}, p^b)$. Here the landlord supplies his property with z_1^b of the environmental attribute and levels of the other property attributes given by the vector, z_{-1}^b . As a result, the landlord can charge a rent of P^b and earns a profit of p^b .

Now, imagine a public programme that increases the level of environmental quality enjoyed at the landlord's property to z_1^a , where the a superscript indicates that this is *after* the environmental improvement. Let us make a further assumption;

- **Assumption 3:** *The level of environmental quality does not affect the optimal level of provision of other property characteristics.* Technically, this amounts to assuming that the attribute z_1 does not interact with other arguments in the hedonic price function.

Thus after the environmental improvement, the landlord will maintain the levels of other environmental attributes at z_{-1}^b .

The first welfare measure we consider requires one further assumption;

- **Assumption 4:** *The level of environmental quality does not affect the costs of supplying other property attributes.* Technically this amounts to assuming that the attribute z_1 does not interact with other arguments in the cost function.

Given our four assumptions, measuring the benefits to landlords of the environmental improvement is a relatively straightforward task.

To illustrate the welfare change experienced by a landlord owning a property in the improved area, observe Figure 14.

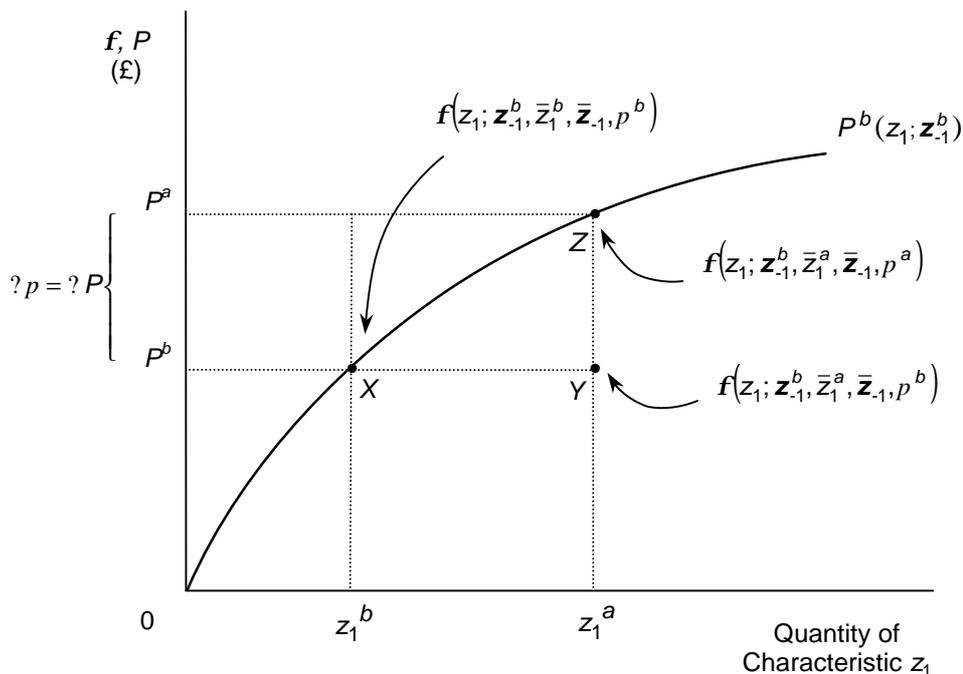
Following the environmental improvement, the landlord could continue to charge a rental price of P^b . This would correspond to the point marked Y in Figure 14. There are a number of things to note about this point.

- First, since the improvement is determined by exogenous factors (Assumption 2), the landlord incurs no added cost in supplying the extra environmental quality.
- Second, we have assumed that the environmental improvement would not encourage the landlord to change levels of supply of other attributes (Assumption 3). Thus following the environmental improvement, the landlord continues to supply the other housing attributes at levels given by the vector z_{-1}^b .
- Finally, we have assumed that changes in environmental quality do not change the costs of supplying the other property characteristics (Assumption 4). Since these are still supplied at z_{-1}^b , the landlords costs in supplying other property attributes will also remain unchanged.

We can conclude that the landlord incurs the same costs after the improvement as before. As a result, the profit associated with point Y is identical to that associated with point X, namely p^b .

Of course, the property now boasts a higher level of environmental quality. Indeed, the landlord is in a position whereby he can increase profits by increasing the rental price of the property. Indeed, given the hedonic price function, the landlord could increase the rental price up to the point marked Z. Notice that this increase in rental price adds directly to the landlord's profits. At Z, the landlord charges a rental price of P^a and realises a profit p^a .

Figure 14: Landlord welfare change for a localised change in an exogenously determined environmental attribute when costs do not change



The welfare measure we seek, therefore, is the difference between profits before the improvement, p^b , and profits after the improvement, p^a . We know from the previous chapter that, provided all else stays the same, the vertical distance between two offer curves equates to the difference in profits associated with the two curves (see Figure 6). Accordingly, the vertical distance YZ measures the increase in profits enjoyed by the landlord. Conveniently, this vertical distance is also the difference between the hedonic price function evaluated at the original and improved levels of attribute z_1 .

Given our four original assumptions, therefore, the change in profits for the landlord can be written;

$$\Delta \mathbf{p} = \mathbf{p}^a - \mathbf{p}^b = P^b(z_1^a, \mathbf{z}_{-1}^b) - P^b(z_1^b, \mathbf{z}_{-1}^b) \quad (24)$$

Of course, we could also derive this result analytically. We know from Equation (15) that the profit realised by the landlord for a property with characteristics \mathbf{z} will equal the rental price of such a property minus the cost of providing the property. Thus we could just as easily write;

$$\Delta \mathbf{p} = \mathbf{p}^a - \mathbf{p}^b = \left(\begin{array}{l} P^b(z_1^a, \mathbf{z}_{-1}^b) - c(z_1^a, \mathbf{z}_{-1}^b; \bar{z}_1^a, \bar{\mathbf{z}}_{-1}^a) \\ - [P^b(z_1^b, \mathbf{z}_{-1}^b) - c(z_1^b, \mathbf{z}_{-1}^b; \bar{z}_1^b, \bar{\mathbf{z}}_{-1}^b)] \end{array} \right) \quad (25)$$

Now, we have already assumed that attribute z_1 is provided without cost to the landlord (Assumption 2) and that the level of this attribute has no effect on the costs of providing other property attributes (Assumption 4). As a result, we can conclude that $c(z_1^a, \mathbf{z}_{-1}^b; \bar{z}_1^a, \bar{\mathbf{z}}_{-1}^a)$ and $c(z_1^b, \mathbf{z}_{-1}^b; \bar{z}_1^b, \bar{\mathbf{z}}_{-1}^b)$ take on the same value and fall out of Equation (25) leaving the desired result, Equation (24).

This is, of course, very intuitive. If the improvement allows the landlord to increase the rental price from P^b to P^a but leaves all costs unchanged, the increase in profits for the landlord will simply be the increase in rental price charged on the property.

Given our assumptions, the total welfare gain to landlords will be given by summing Equation (24) across all landlords. Of course, one of those assumptions is that there are no adjustments in the hedonic property market (Assumption 1). Consequently the welfare change will only be experienced by landlords owning properties in the affected area. In the previous section, we denoted this set of properties H_1 . Thus, indexing landlords in the market by $l = 1$ to L , the welfare change experienced by landlords can be expressed;

$$W_L = \sum_{l \in H_1} P^b(z_{1l}^a, \mathbf{z}_{-1l}^b) - P^b(z_{1l}^b, \mathbf{z}_{-1l}^b) \quad (26)$$

where W_L is the total welfare change experienced by landlords in the market.

One of the advantages of this welfare measure is that it requires relatively little information. To use this measure, a researcher would simply need an estimate of the hedonic price function and details of the level of the environmental attribute at affected properties before and after the improvement.

Of course, the assumptions made in deriving Equation (26) were very restrictive. In what follows, we shall present three more measures of landlord welfare change that successively relax these assumptions.

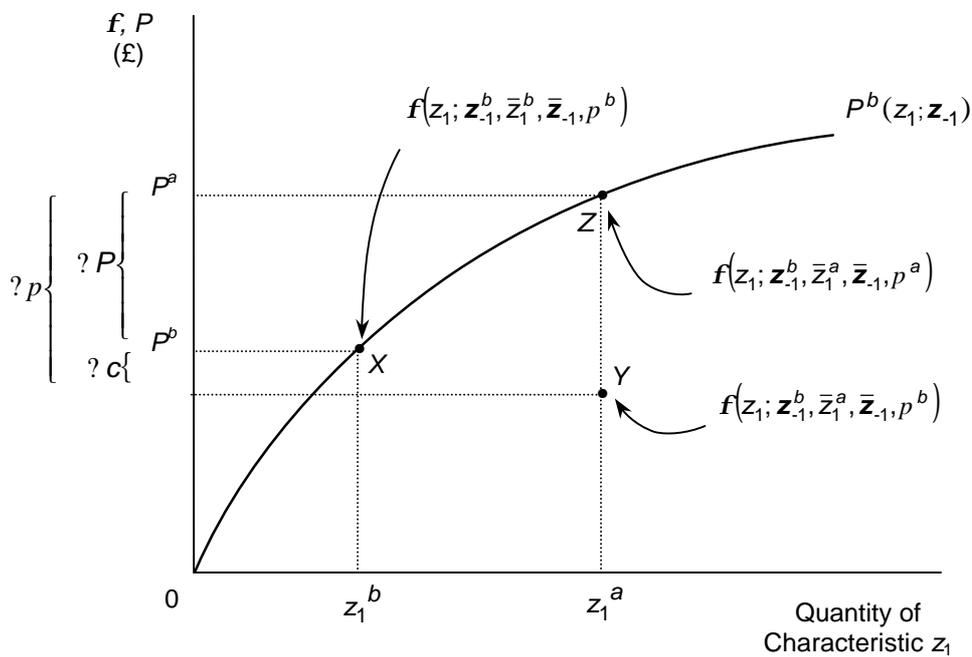
First, consider the situation where the level of the environmental attribute z_1 influences the landlords' costs. In other words, let us relax Assumption 4. Examples of environmental improvements that might result in concomitant cost savings include;

- a reduction in air pollution which reduces the necessity to clean or repaint the property
- a reduction in crime which reduces the need for repairs caused by vandalism

This case is depicted in Figure 15. Again the environmental improvement has only a local impact (Assumption 1), the level of the attribute is entirely determined by exogenous factors (Assumption 2) and the landlord persists in supplying other property attributes at the same level after the improvement (Assumption 3).

Before the improvement, the landlord chooses to locate at point X. Here the landlord supplies a property with the exogenously determined level of environmental quality z_1^b and chosen levels of other property attributes given by the vector z_{-1}^b . At this combination of attributes the landlord maximises profits by charging a rent P^b of which p^b is profit.

Figure 15: Landlord welfare change for a localised change in an exogenously determined environmental attribute when costs change



Following an environmental improvement, the level of z_1 is increased to z_1^a at no cost to the landlord. Further according to Assumption 3, the landlord continues to provide other property attributes at the same levels, that is, z_{-1}^b . However, by relaxing Assumption 4,

we allow for the possibility that the environmental improvement may reduce the cost of providing the other housing attributes at these levels.

Indeed, following the environmental improvement the landlord could locate at point Y . Here, the landlord could charge a lower price yet, as a result of cost savings, achieve the same level of profits as previous to the environmental improvement. The vertical distance between X and Y measures the cost savings brought about by the environmental improvement.

Of course the landlord will not locate at Y . Instead, he will maximise his profits by locating at point Z . Here the landlord charges a rent P^a of which p^a is profit.

The environmental improvement increases the landlord's profits from p^b to p^a . Again, this increase can be measured as the vertical distance between the offer curves, YZ . Notice that allowing for cost changes expands our measure of the welfare gains for landlords. Not only does the landlord enjoy an increase in rent, ΔP , but also experiences a reduction in costs Δc .

Accordingly, this broader welfare measure can be calculated as;

$$\Delta p = \Delta P + \Delta c = \left(\begin{array}{l} [P^b(z_1^a, z_{-1}^b) - P^b(z_1^b, z_{-1}^b)] \\ + [c(z_1^b, z_{-1}^b; \bar{z}_1^b, \bar{z}_{-1}^b) - c(z_1^a, z_{-1}^b; \bar{z}_1^a, \bar{z}_{-1}^b)] \end{array} \right) \quad (27)$$

Since, this measure continues to assume that there are no adjustments in the hedonic property market the welfare change is only experienced by landlords owning properties in the affected area. Using this measure, the total welfare impact of the environmental improvement is given by;

$$W_L = \sum_{i \in H_1} \left(\begin{array}{l} [P^b(z_{1i}^a, z_{-1i}^b) - P^b(z_{1i}^b, z_{-1i}^b)] \\ + [c(z_{1i}^b, z_{-1i}^b; \bar{z}_{1i}^b, \bar{z}_{-1i}^b) - c(z_{1i}^a, z_{-1i}^b; \bar{z}_{1i}^a, \bar{z}_{-1i}^b)] \end{array} \right) \quad (28)$$

Notice that this measure of welfare change is informationally more exacting since it demands that the researcher has knowledge of the landlords cost function.

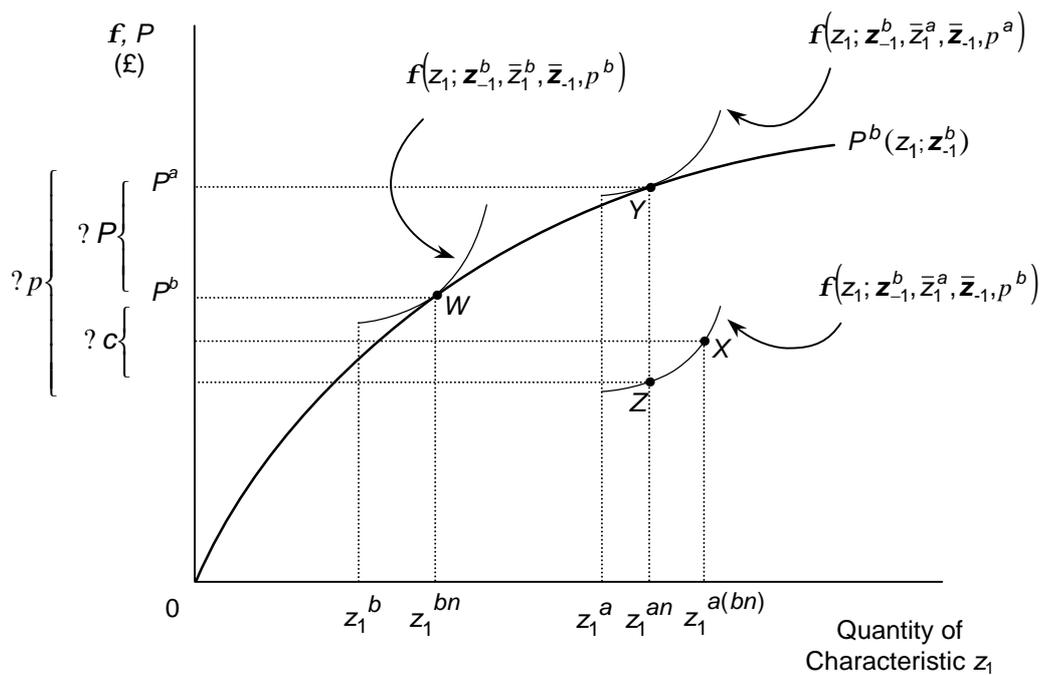
The two welfare measures that we have developed so far, have both assumed that landlords are not able to influence the level of environmental quality of their properties. Whilst this may be true in the short-term, we have already cited counter examples. For instance, a landlord can change a property's exposure to noise pollution by installing double-glazing.

Our next task, therefore, is to relax Assumption 2 and consider the situation where the level of environmental quality is not entirely determined by exogenous factors. For now, however, we maintain Assumption 3. That is, following an environmental improvement, we allow landlords to alter the level of environmental quality of their properties but not alter the levels of other property attributes. Compared with the last two scenarios, this is

more indicative of a landlord's medium to long term response to changes in property market conditions.

The pattern of responses is fairly complex and is laid out in Figure 16. In the first instance the landlord is faced by the hedonic price function $P^b(\cdot)$ and the exogenously determined level of the environmental attribute z_1^b . To illustrate let us assume that z_1 is the level of crime in the area. Faced with these two restrictions, the landlord maximises profits by investing in private goods that expand the level of attribute z_1 to z_1^{bn} . Here the superscript n indicates the *new* level of the property attribute once the investments have been undertaken. For instance the landlord could further reduce the risk of crime by installing a burglar alarm monitored by a private security company. Following these investments, the landlord achieves point W where the rental value of the property is P^b and the landlord earns a profit of p^b .

Figure 16: Landlord welfare change for a localised change in environmental attribute



Now let us consider a public programme that leads to an increase in the exogenously supplied level of z_1 from z_1^b to z_1^a . In our example, the level of criminal activity in the area falls. For the sake of argument, imagine that the landlord did not adjust to this change. In our example, the landlord might continue to employ the private security firm despite the fact that crime risks in the area have fallen. Following the change the landlord's property would boast a level of environmental quality given by $z_1^{a(bn)}$ where

the superscript $a(bn)$ indicates that this is the level of provision *after* the change but whilst maintaining the *new* level of property investments undertaken *before* the change.

Thus if the landlord wished to maintain the same level of profit as previous to the change, he would end up at point X which lies on the new offer curve providing the original level of profit, \mathbf{p}^b .

Notice that, as in the previous scenario, the increased environmental quality has resulted in immediate reductions in the costs of providing other housing attributes. Indeed, the vertical distance between W and X measures the cost savings brought about by the environmental improvement.

Of course X is by no means the landlord's optimal location. Indeed, given $P^b(\cdot)$ and the exogenously determined level of the environmental attribute z_1^a , the landlord would be best advised to increase the rent on the property and consider the potential benefits of changing the property's level of environmental quality.

In Figure 16 the best the landlord could do would be to relocate to point Y . Here, the level of the environmental attribute z_1 has been altered to z_1^{an} and the landlord maximises profits at \mathbf{p}^a by charging a rental of P^a . Continuing our example, in response to the fall in crime in the area, the landlord may decide to increase the rent on the property whilst terminating his employment of the private security company.

Once again, the increase in the landlord's profits will be the vertical distance between Y and the point on the equivalent offer curve delivering the original level of profits, point Z . In Figure 16, therefore, the increase in the landlord's profits is the distance ZY .

Again this increase in profits can be decomposed into a change of price and a change in costs according to;

$$\Delta \mathbf{p} = \Delta P + \Delta c = \left(\begin{array}{l} [P^b(z_1^{an}, z_{-1}^b) - P^b(z_1^{bn}, z_{-1}^b)] \\ + [c(z_1^{bn}, z_{-1}^b; \bar{z}_1^b, \bar{z}_{-1}^b) - c(z_1^{an}, z_{-1}^b; \bar{z}_1^a, \bar{z}_{-1}^a)] \end{array} \right) \quad (29)$$

This measure is broader than those that were discussed previously, because it allows for landlords to adjust the levels of environmental quality after the exogenous change. Since we are still dealing with a localised environmental improvement, this broader measure will still only be defined for properties in the affected area. The total welfare change is given by

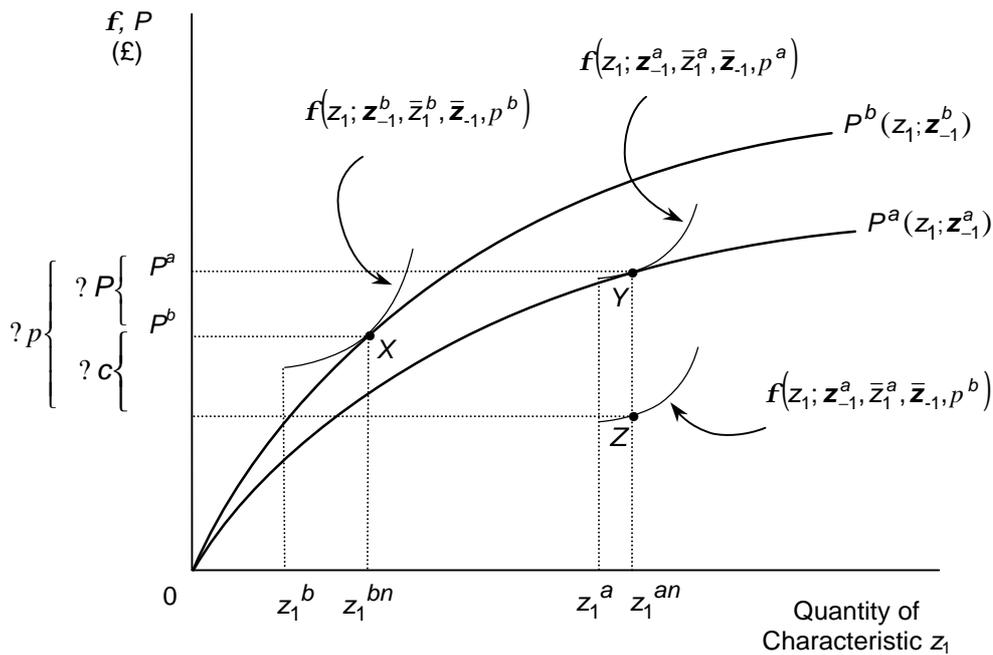
$$W_L = \sum_{l \in H_1} \left(\begin{array}{l} [P^b(z_{1l}^{an}, z_{-1l}^b) - P^b(z_{1l}^{bn}, z_{-1l}^b)] \\ + [c(z_{1l}^{bn}, z_{-1l}^b; \bar{z}_{1l}^b, \bar{z}_{-1l}^b) - c(z_{1l}^{an}, z_{-1l}^b; \bar{z}_{1l}^a, \bar{z}_{-1l}^a)] \end{array} \right) \quad (30)$$

ii. Landlord welfare changes from a non-localised environmental improvement

The final welfare measure we discuss relaxes all four assumptions simultaneously. This scenario, therefore, includes situations where the environmental improvement is substantial enough to result in a shift in the hedonic price function. Further, unlike the measure described by Equation (30), we allow for the fact that the landlord may decide to change the levels of provision of all the housing attributes as a result of the environmental improvement and subsequent shift in the hedonic price function. This case is depicted in Figure 17.

The landlord starts off with an exogenously determined level of environmental quality \bar{z}_1^b and baseline levels of other property attributes given by the vector \bar{z}_{-1} . In the first instance the landlord is faced by the hedonic price function $P^b(\cdot)$. In order to maximize profits the landlord wishes to move to point X by altering the environmental quality of the property to z_1^{bn} and the levels of other property attributes z_{-1}^b . Here the landlord can charge a price of P^b and earns profits from the property of p^b .

Figure 17: Landlord welfare change for a non-localised change in an environmental attribute



Now a public programme results in an environmental improvement in the urban area. At the landlord's property this manifests itself as an increase in the exogenously determined level of environmental quality from \bar{z}_1^b to \bar{z}_1^a . However, this is not a merely localised change. Indeed, the set of prices given by the old hedonic price function would no longer clear the market. Thus in response to the environmental improvement, the market adjusts, establishing equilibrium at the new hedonic price function given by $P^a(\cdot)$.

The landlord is faced by a number of simultaneous changes;

- environmental quality at their property increases
- as result of the environmental improvement the costs of providing different combinations of property attributes reduce
- the hedonic price function changes

In response the landlord will maximise profits by moving to point Y by altering the provision of environmental quality to z_1^{an} and the levels of other property attributes to z_{-1}^a . Notice that we have allowed for the fact that it may be optimal to adjust the level of all housing attributes in response to the environmental improvement.

Following the same argument as that used previously, the relevant welfare measure is the vertical distance between the points marked Z and Y .

This measure is the landlords' equivalent to the Compensating Surplus measure defined for households. As with that measure, the landlord is allowed to respond optimally to the change in environmental quality and the shift in the hedonic price function. For this reason we label this comprehensive welfare measure the Compensating Profit (CP). In mathematical terms it is defined as;

$$CP = \Delta p = \Delta P + \Delta c = \left(\begin{array}{l} [P^a(z_1^{an}, z_{-1}^a) - P^b(z_1^{bn}, z_{-1}^b)] \\ + [c(z_1^{bn}, z_{-1}^b; \bar{z}_1^b, \bar{z}_{-1}^b) - c(z_1^{an}, z_{-1}^a; \bar{z}_1^a, \bar{z}_{-1}^a)] \end{array} \right) \quad (31)$$

If Δp is negative then the change in environmental quality reduces the welfare of the landlord. If Δp is positive then the change in environmental quality increases the welfare of the landlord.

Since all landlords are assumed to respond to the shift in the hedonic price function the total welfare benefits of the environmental improvement will include a measure for each of the H landlords in the urban area;

$$W_L = \sum_{l=1}^H CP_l = \sum_{l=1}^H \left(\begin{array}{l} [P^a(z_{1l}^{an}, z_{-1l}^a) - P^b(z_{1l}^{bn}, z_{-1l}^b)] \\ + [c(z_{1l}^{bn}, z_{-1l}^b; \bar{z}_{1l}^b, \bar{z}_{-1l}^b) - c(z_{1l}^{an}, z_{-1l}^a; \bar{z}_{1l}^a, \bar{z}_{-1l}^a)] \end{array} \right) \quad (32)$$

Notice that the informational requirements of the CP measure are extremely onerous. Not only must the researcher be able to predict how the hedonic price function will change in response to a non-localised change in environmental quality, but must also be able to predict the optimal response of each landlord to the change in market conditions.

Table 2 summarises the various measures of landlord welfare change described in this section.

Table 2: Measures of landlord welfare change

Welfare Measure	Description	Computation of Total Welfare Change for Landlords	Informational Requirements
Localised:			
Exogenous Attribute, no Cost Changes	<ul style="list-style-type: none"> No shift in hedonic Rent increase for improved properties 	$W_L = \sum_{l \in H_1} P^b(z_{1l}^a, z_{-1l}^b) - P^b(z_{1l}^b, z_{-1l}^b)$	<ul style="list-style-type: none"> Only affected landlords Environmental quality before and after change Original hedonic
Exogenous Attribute, with Cost Changes	<ul style="list-style-type: none"> No shift in hedonic Landlords in improved areas experience cost changes Rent increase for improved properties 	$W_L = \sum_{l \in H_1} \left(P^b(z_{1l}^a, z_{-1l}^b) - P^b(z_{1l}^b, z_{-1l}^b) + [c(z_{1l}^b, z_{-1l}^b; \bar{z}_{1l}^b, \bar{z}_{-1l}^b) - c(z_{1l}^a, z_{-1l}^b; \bar{z}_{1l}^a, \bar{z}_{-1l}^a)] \right)$	<p>As previous, plus:</p> <ul style="list-style-type: none"> Changes in exogenous levels of other attributes Landlord cost function
Any attribute	<ul style="list-style-type: none"> No shift in hedonic Landlords in improved areas experience cost changes Landlords optimise level of provision of environmental quality attribute Rent change for improved properties 	$W_L = \sum_{l \in H_1} \left([P^b(z_{1l}^{an}, z_{-1l}^b) - P^b(z_{1l}^{bn}, z_{-1l}^b)] + [c(z_{1l}^{bn}, z_{-1l}^b; \bar{z}_{1l}^b, \bar{z}_{-1l}^b) - c(z_{1l}^{an}, z_{-1l}^b; \bar{z}_{1l}^a, \bar{z}_{-1l}^a)] \right)$	<p>As previous, plus:</p> <ul style="list-style-type: none"> Landlords' choices of environmental quality attribute after improvement
Non-Localised:			
Compensating Profit	<ul style="list-style-type: none"> Hedonic shifts Landlords in improved areas experience cost changes Landlords optimise property attributes Rent change for all properties 	$W_L = \sum_{l=1}^H \left(P^a(z_{1l}^{an}, z_{-1l}^a) - P^b(z_{1l}^{bn}, z_{-1l}^b) + [c(z_{1l}^{bn}, z_{-1l}^b; \bar{z}_{1l}^b, \bar{z}_{-1l}^b) - c(z_{1l}^{an}, z_{-1l}^a; \bar{z}_{1l}^a, \bar{z}_{-1l}^a)] \right)$	<p>As previous, plus:</p> <ul style="list-style-type: none"> All landlords Landlords' choices of all attributes after improvement Hedonic before and after change

f. Combining Household and Landlord Welfare Measures

The total benefits to households and landlords resulting from an environmental improvement are found simply by adding W_H to W_L . Of course, this total welfare measure will depend on which assumptions are made and hence which of the formulas in Tables 1a and 1b are chosen to represent the households' and landlords' welfare changes.

Before discussing these measures further, we should note that such welfare estimates;

- measure the welfare benefits to both households and landlords for changes in environmental quality in their residential location
- ignore the benefits to visitors that travel by the improved area.
- ignore the benefits to those who work in the improved area⁸.
- ignore the costs of causing the environmental improvement. For example, no account is taken of the costs to industry of reducing emissions or the cost to the tax payer of traffic calming schemes designed to reduce traffic noise.

In the simplest case, the environmental improvement is a localised phenomena that causes no change in the hedonic price function. If we assume that households incur no moving costs then they will be relocated to a property offering the attributes of their original location prior to the improvement and experience no welfare change. Further, if we assume that landlords cannot affect the level of environmental quality at their properties, that the level of environmental quality does not influence the optimal level of provision of other attributes and that their costs of providing other property attributes are unaffected by the improvement, then the welfare gain for the landlords is simply the change in the rental price of their properties. The total welfare change is given by the sum of Equation (26) and the upper bound of Equation (19);

$$W_L + W_H = \sum_{l \in H_1} P^b(z_{1l}^a, z_{-1l}^b) - P^b(z_{1l}^b, z_{-1l}^b) + 0 \quad (33)$$

In other words, under certain restrictive assumptions, the total welfare change can be measured as the change in price of affected properties. What is more, to calculate this measure requires only two pieces of information;

- the current hedonic price function.
- the level of environmental quality at each affected property following the environmental improvement.

For any one market, welfare changes as measured by Equation (33) should be relatively simple to estimate. Unfortunately, it is not possible to transfer such estimates to different property markets. Remember from Chapter 1 that the hedonic price function is

⁸ Of course, there is no reason why we shouldn't be able to measure the benefits to these individuals reflected in other hedonic markets such as the hedonic market for office space or the hedonic wage market.

determined by the unique conditions of supply and demand existing in a particular market. As a result, hedonic price functions will differ across property markets. A welfare measure calculated using the hedonic price function in one particular market would only be relevant to that market. It would make no sense to transfer such evaluations across different markets.

Of course, Equation (33) is by no means a comprehensive measure of the welfare change associated with a localised change in environmental quality. Indeed, by relaxing some of the assumptions underlying Equation (33) we could expand our measure of the welfare gain. For example, we might wish to allow for the fact that households face transaction costs when moving properties, that landlords might wish to optimally adapt the level of environmental quality at their properties and that changes in environmental quality might affect the costs of providing other property attributes. In this case our welfare measure would be the sum Equation (30) and the lower bound of Equation (19);

$$W_L + W_H = \sum_{l \in H_1} \left(\begin{aligned} & \left[P^b(z_{1l}^{an}, z_{-1l}^b) - P^b(z_{1l}^{bn}, z_{-1l}^b) \right] \\ & + \left[c(z_{1l}^{bn}, z_{-1l}^b; \bar{z}_{1l}^b, \bar{z}_{-1l}^b) - c(z_{1l}^{an}, z_{-1l}^b; \bar{z}_{1l}^a, \bar{z}_{-1l}^a) \right] \end{aligned} \right) - \sum_{h \in H_1} tc_h \quad (34)$$

Of course, this may be a more comprehensive measure of the welfare change brought about by the environmental improvement, but it is also considerably harder for a researcher to estimate. Compared to Equation (33) the researcher would now need to estimate the moving costs for each household affected by the environmental change, the landlords' cost function and the adaptations made by landlords to the environmental quality attribute following the improvement. Indeed, attempting to estimate Equation (33) prior to a change in environmental quality is almost an impossible task.

In the extreme, we could relax all assumptions and allow for changes in environmental quality that are non-localised and precipitate alterations in the hedonic price function. Ignoring transaction costs, this measure would be derived by adding Equation (32) to Equation (23);

$$\begin{aligned} TSB &= \sum_l CP_l + \sum_h CS_h \\ &= \sum_{l=1}^H \left(\begin{aligned} & \left[P^a(z_{1l}^{an}, z_{-1l}^a) - P^b(z_{1l}^{bn}, z_{-1l}^b) \right] \\ & + \left[c(z_{1l}^{bn}, z_{-1l}^b; \bar{z}_{1l}^b, \bar{z}_{-1l}^b) - c(z_{1l}^{an}, z_{-1l}^a; \bar{z}_{1l}^a, \bar{z}_{-1l}^a) \right] \end{aligned} \right) \\ &+ \sum_{h=1}^H \left(\begin{aligned} & \left(q(z_{1h}^{an}, z_{-1h}^{an}; u_{1h}) - q(z_{1h}^{bo}, z_{-1h}^{bo}; u_{1h}) \right) \\ & - \left[P^a(z_{1h}^{an}, z_{-1h}^{an}) - P^b(z_{1h}^{bo}, z_{-1h}^{bo}) \right] \end{aligned} \right) \end{aligned} \quad (35)$$

This would give us our most comprehensive measure of the welfare change⁹ and hence is labelled the *Total Social Benefits (TSB)* of the of the change in environmental quality. Notice that this measure is summed over all households and landlords in the urban area even those not originally located in the improved area. This is important since the latter group may be affected by subsequent changes in the hedonic market.

Further, and most importantly with regards to the present discussion, the measure is almost impossible to calculate. To assess Equation (35) 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. As discussed previously, the complexity of the hedonic market equilibrium precludes analytical solutions to this problem. As such, Equation (35) is of little use to practitioners attempting to measure the benefits derived from a program designed to change environmental quality in an urban area.

g. A Quantifiable Lower Bound

Since the informational requirements for measuring *TSB* are prohibitive, economists have looked to define a simpler measure that might lend itself to estimation in the real world. It turns out that one such measure is the sum of *QCS* measures presented in the previous discussion (Equation 21). All that is required to calculate this measure is knowledge of the bid function of households in the affected area, details of their current residential choices and information on the level of environmental change experienced by each household. The great advantage of using this measure is that estimates of welfare changes can be made without knowledge of how households, landlords and the hedonic price function react to a change in environmental quality.

Encouragingly, Bartik (1988) has given a theoretical justification for choosing to measure the welfare changes resulting from a change in environmental quality as the sum of households' *QCS*. He shows that the sum of *QCS* across all affected households provides a lower bound estimate of the *TSB*. That is, if we calculate the sum of *QCS* resulting from a change in environmental quality, we will be calculating a figure that whilst not reflecting the full benefits of the change will provide a lower estimate of these benefits.

Bartik's intuitive proof involves partitioning the welfare changes affecting households and landlords into a series of three stages. Whilst these stages help in the analysis of welfare changes they are not meant to represent a realistic sequence of events. The three stage decomposition is presented in Table 3.

In the *first stage*, some or all of the residential locations in the urban area experience an improvement in environmental quality. At this stage, we assume that neither landlords, nor households nor the hedonic market adjust in response to this change. Thus the household stays in the same property, the landlord does not increase the rent nor adjust

⁹ Though remember this formula is based on the less comprehensive *CS* measure that does not allow for the adjustments in residential location in response to changes in income.

the property's attributes and the hedonic price function does not change from its previous form.

- Since households cannot move property, the benefit to households will be simply their WTP for the environmental improvement at their original location. This is the *QCS* measure presented in Figure 11 and Equation (20).
- Since landlords cannot change rents or adjust the attributes of their properties, they will only be affected by the change in environmental quality if it affects their costs. Since we assume they make no changes to their properties at this stage, the measure of cost savings is that given by the vertical distance between *W* and *X* in Figure 16.

In the *second stage*, the hedonic price function shifts to its final form but we still constrain households and landlords to their original location and supply choices. Since households and landlords remain in the same location the change in rent associated with the shift in the hedonic price function acts to simply transfer money from one to the other. Indeed, whatever the pattern of rent changes in the second stage, there is no overall welfare effect.

Notice, however, that though in stage 2 the aggregate welfare change across the whole urban area is zero, welfare changes for each individual household and landlord may be positive or negative depending on the particular pattern of rent changes. Landlords at unimproved sites, for example, will almost certainly experience some reduction in rent and hence profits.

In the *third stage*, households are allowed to move and landlords are allowed to change the attributes of their properties in response to the new hedonic price function. Since both households and landlords are allowed to respond optimally, they must, by definition, experience an increase in welfare. Households will move to the property that offers them the highest possible utility. This must be at least as beneficial as remaining in the original property since they could always opt not to move house. A similar story can be made for landlords' supply decisions. In effect, therefore, compared to stage 2, both households and landlords must witness an increase in welfare. Again this is not to say that every household and landlord experiences an increase in welfare over all three stages. Whilst households and landlords only benefit in stages 1 and 3, they may just as well lose benefit as gain benefit in the rent changes isolated in stage 2.

As shown in Table 3, summing all three stages for households results in the total welfare gains given by the sum of household *CS*'s given in Equation (23). Similarly, summing all three stages results in the sum of landlords *CP*'s given in Equation (32). Thus the three stage decomposition, whilst not reflecting the simultaneous nature of responses to the change in environmental quality, accurately represents the overall change in welfare.

Table 3: A decomposition of the welfare effects of a change in environmental quality from Bartik (1988)

Benefits at Various Stages			
	Households	Landlords	Net Efficiency Benefits
<p>Stage 1: Amenity changes, no adjustment or rent change</p>	$\sum_h [\mathbf{q}(z_{1h}^{ao}, \mathbf{z}_{-1h}^{bo}; u_{1h}) - \mathbf{q}(z_{1h}^{bo}, \mathbf{z}_{-1h}^{bo}; u_{1h})]$ <p>Household WTP at original location: zero for unimproved sites, positive for improved sites</p>	$\sum_l - [c(z_{1l}^{a(bn)}, \mathbf{z}_{-1l}^b; \bar{z}_{1l}^a, \bar{\mathbf{z}}_{-1l}) - c(z_{1l}^{bn}, \mathbf{z}_{-1l}^b; \bar{z}_{1l}^b, \bar{\mathbf{z}}_{-1l})]$ <p>Landlord cost savings: assumed non-negative for improved sites, zero for unimproved sites</p>	<p>Sum of all households' WTP plus all landlords' cost savings</p>
<p>Stage 2: Rent Change</p>	$\sum_h - [P^a(z_{1h}^{ao}, \mathbf{z}_{-1h}^{bo}) - P^b(z_{1h}^{bo}, \mathbf{z}_{-1h}^{bo})]$ <p>Rent change at both improved and unimproved sites</p>	$\sum_l P^a(z_{1l}^{a(bn)}, \mathbf{z}_{-1l}^b) - P^b(z_{1l}^{bn}, \mathbf{z}_{-1l}^b)$ <p>Rent change at both improved and unimproved sites</p>	<p>Zero efficiency benefits; pecuniary transfer between households and landlords</p>
<p>Stage 3: Adjustment</p>	$\sum_h [\mathbf{q}(z_{1h}^{an}, \mathbf{z}_{-1h}^{an}; u_{1h}) - P^a(z_{1h}^{an}, \mathbf{z}_{-1h}^{an}) - \mathbf{q}(z_{1h}^{ao}, \mathbf{z}_{-1h}^{bo}; u_{1h}) - P^a(z_{1h}^{ao}, \mathbf{z}_{-1h}^{bo})]$ <p>Measure of household utility increase from adjustment, for households originally at both improved and unimproved sites</p>	$\sum_l [P^a(z_{1l}^{an}, \mathbf{z}_{-1l}^a) - c(z_{1l}^{an}, \mathbf{z}_{-1l}^a; \bar{z}_{1l}^a, \bar{\mathbf{z}}_{-1l}) - P^a(z_{1l}^{a(bn)}, \mathbf{z}_{-1l}^b) - c(z_{1l}^{a(bn)}, \mathbf{z}_{-1l}^b; \bar{z}_{1l}^a, \bar{\mathbf{z}}_{-1l})]$ <p>Landlord profit increase from adjustment to new hedonic: applies to landlords at all sites</p>	<p>Net gain from adjustment must be non-negative for all</p>
<p>Sum of three stages</p>	$\sum_h [\mathbf{q}(z_{1h}^{an}, \mathbf{z}_{-1h}^{an}; u_{1h}) - \mathbf{q}(z_{1h}^{bo}, \mathbf{z}_{-1h}^{bo}; u_{1h}) - P^a(z_{1h}^{an}, \mathbf{z}_{-1h}^{an}) + P^b(z_{1h}^{bo}, \mathbf{z}_{-1h}^{bo})]$ <p>Net household gain: sum over all households, Equation (23) in text</p>	$\sum_l [P^a(z_{1l}^{an}, \mathbf{z}_{-1l}^a) - c(z_{1l}^{an}, \mathbf{z}_{-1l}^a; \bar{z}_{1l}^a, \bar{\mathbf{z}}_{-1l}) - P^b(z_{1l}^{bn}, \mathbf{z}_{-1l}^b) - c(z_{1l}^{bn}, \mathbf{z}_{-1l}^b; \bar{z}_{1l}^b, \bar{\mathbf{z}}_{-1l})]$ <p>Net landlord gain, sum over all landlords, Equation (32) in text</p>	<p>Sum of 1st and 2nd columns is same as Equation (35)</p>

The insight of Bartik's decomposition is to isolate all individual welfare losses as price changes in stage 2. Since price changes simply represent pecuniary transfers between agents in the property market, these losses must be offset by equivalent gains elsewhere. In other words when we are interested in the aggregate welfare change, we can ignore the losses incurred by certain landlords and households by netting these out as a price change.

As a result *TSB*, that is the total welfare change experienced by all households and landlords in the urban area, can be regarded as the sum of the four non-negative values defined in stages one and three. In words, these are;

1. WTP of households at improved locations to enjoy the change in environmental quality whilst staying in their original property ($\sum_{h \in H_1} QCS_h$)
2. cost savings for landlords at stage 1
3. household utility gains from relocation at stage 3
4. landlord profit gains from changes in supply at stage 3

Since all four values are non-negative, $\sum_{h \in H_1} QCS_h$ must also be a lower bound to *TSB*.

This is an extremely important insight since it gives us a good theoretical reason for using $\sum_{h \in H_1} QCS_h$ to measure the welfare change resulting from an environmental improvement.

There are a number of reasons why this might be desirable.

- First, since the *QCS* measure does not require information on how the market price or agents in the market adjust to a change in market conditions, it can be calculated in advance of a public programme to improve environmental quality.
- Second, the *QCS* is a measure of household welfare change. Consequently using the sum of *QCS*s as a lower bound estimate of *TSB* removes the need to examine the supply side of the market. Researchers can ignore the considerable difficulties associated with estimating landlord cost and offer functions.
- Third, *QCS* is only defined for households in an affected area. As a result, the researcher only requires information on which households will be affected by the environmental improvement and the extent of improvement enjoyed by each.
- Finally, the *QCS* measure is based solely on underlying preferences for environmental quality as captured in the bid function. The measure is not particular to a specific property market. Indeed, if a researcher could derive the bid function from one market then this could be used to evaluate the *QCS* in another property market, provided the researcher was prepared to assume that preferences for environmental quality were stable across the two markets.

Clearly, using the sum of households' *QCS* as a lower bound approximation to the *TSB* makes it practical to carry out *ex-ante* assessments of the welfare gains from environmental improvements. Obviously, the accuracy of this approximation will depend on the size of the values taken by the other three elements of *TSB* isolated in Bartik's

analysis. Certainly, the approximation will tend to be more accurate when the environmental change is less extensive as the benefits of household relocation and landlord change in supply will tend to be smaller.

h. Conclusions

This chapter has demonstrated how the benefits of an environmental improvement can be measured in the property market. The benefits captured in this market are those accruing to households and landlords of a particular residential location. The measures described here do not capture the benefits to visitors that travel by the improved area nor do they capture the benefits to those who work in the improved area.

In the simplest case, the environmental improvement is a localised phenomena that causes no change in the hedonic price function. If households can move freely and landlords do not enjoy cost savings and are constrained not to alter the supply of property attributes, then the welfare benefits of the improvement accrues to landlords as the *change in the rental price* of their properties (Equation 33).

This measure is easy to calculate for any property market for which the hedonic price function is known. Unfortunately, the fact that the measure is based on the unique hedonic price function of a particular market means that there is no theoretical substance to transferring such values across property markets.

Clearly, estimating the welfare change of an environmental improvement by the increase in prices of affected properties is to impose severe restrictions on the reactions of the economic agents in the market to the improvement. Indeed, a completely comprehensive measure of the welfare benefits of an environmental improvement is given by the *Total Social Benefits (TSB)* measure of Equation (35).

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.

Nevertheless, in an important analysis, Bartik (1988) showed how a third measure the *QCS*, when summed over all households directly affected by the change in environmental quality, could always be taken as a lower bound to the *TSB*. There are a number of reasons why using the *QCS* measure might be desirable. In particular, the *QCS* measure is based solely on the household bid function. As a result, it is not necessary to consider the supply side of the market nor predict market conditions following environmental change. Further, the *QCS* measure is not particular to a specific property market. Indeed, if a researcher could derive the bid function from one market then, provided the researcher

was prepared to assume that preferences for environmental quality were stable across the two markets, this could be used to evaluate the *QCS* in another property market.

In the next chapter, therefore, we investigate the possibilities for deriving estimates of the bid function from which the *QCS* measure of welfare change can be derived.

CHAPTER 3. DEMAND ANALYSIS USING HEDONIC MARKET DATA

a. Introduction

Bartik's analysis presented in the last section, 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. Over the last twenty or so years, researchers have raised some major problems concerning the possibility of identifying bid functions from observations of households' behaviour in hedonic property markets. In short this research has amounted to answering three major questions;

- First, whether the bid function or its derivative the marginal bid function, could ever be identified from data on residential choices in a single hedonic market in which all households face the same hedonic price schedule. It turns out that to learn anything about household demand for property characteristics, one must observe household choices in response to a variety of different hedonic price schedules. That is, a prerequisite for identifying the bid function is that data is available from multiple hedonic property markets.
- Second, whether the marginal bid function can be directly observed through household choices in multiple markets. Again, it is relatively simple to show that the household's actual choices of attribute quantities in response to different hedonic price schedules do not trace out the marginal bid function.
- The third question then, is whether it is possible to infer the bid function from observed choices in hedonic markets. Fortunately the answer to this question is that we can use the information provided by observed behaviour to deduce the bid function, though the techniques are relatively complex.

In this section we address each of the questions raised above. Again, the focus of this discussion will be theoretical, though of course the end objective will be to produce theoretical results that allow estimation from market data.

b. The Marginal Bid Function

The bid function, $\theta(z; y, s, u)$ describes the amount of money that a household would be prepared to pay for a property with attributes z in order to enjoy the level of utility, u . Of course, the amount that a household would bid for a particular property will not depend solely on the level of utility specified in the bid function. Rather, the household's income, y , and socioeconomic characteristics, s , will also influence their bid.

As we have shown previously, the bid function can be illustrated as bid curves. Bid curves depict combinations of property attributes, z , and payments for those attributes, θ , between which the household is indifferent (i.e. combinations that confer the same utility on the household).

For our present purposes, it frequently proves more convenient to work with the marginal bid function. That is, a function that shows how much a household is willing to pay for each extra unit of housing attribute z_i , so as to maintain the same level of utility, u . Mathematically the marginal bid function is the partial derivative of the bid function. Remember from Equation (15) that the bid function is defined as;

$$\theta(z; y, s, u) = y - x(z; s, u) \quad (15)$$

Thus the marginal bid function is given by;

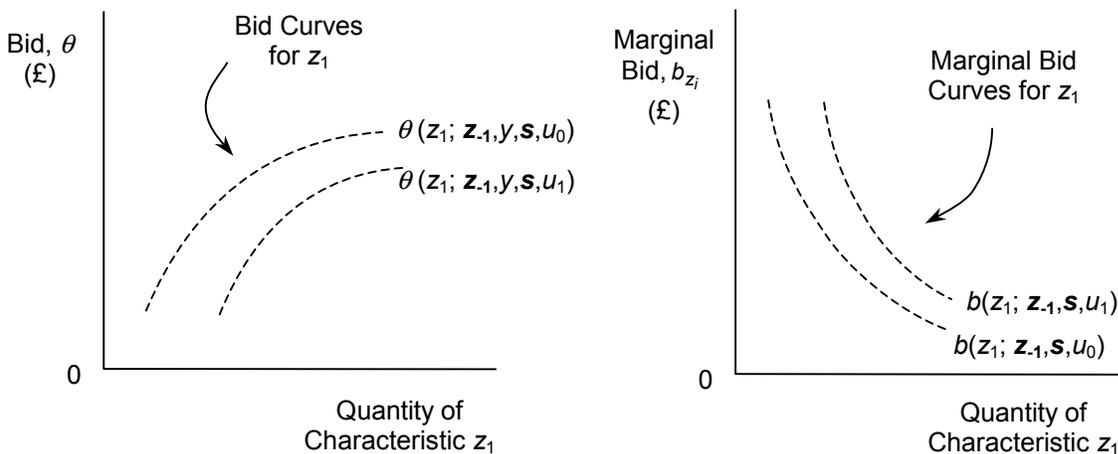
$$b_{z_i}(z_i; z_{-i}, s, u) = \frac{\partial \theta(z; y, s, u)}{\partial z_i} \quad (36)$$

Notice that the household's income y falls out of the marginal bid function. Everything else being equal, the amount that a household is prepared to pay for a property with one extra unit of an attribute in order to maintain the same level of utility is independent of their income.

The marginal bid function can itself be illustrated as a *marginal bid curve* which describes the slope of an equivalent bid curve.

Two bid curves and the equivalent marginal bid curves for a household are illustrated in Figure 18. In the left hand panel, the higher bid curve corresponds to combinations of payments and housing attribute z_1 that result in a utility level u_0 . The lower bid curve corresponds to a higher level of utility, u_1 , since each level of attribute z_1 is associated with a lower payment.

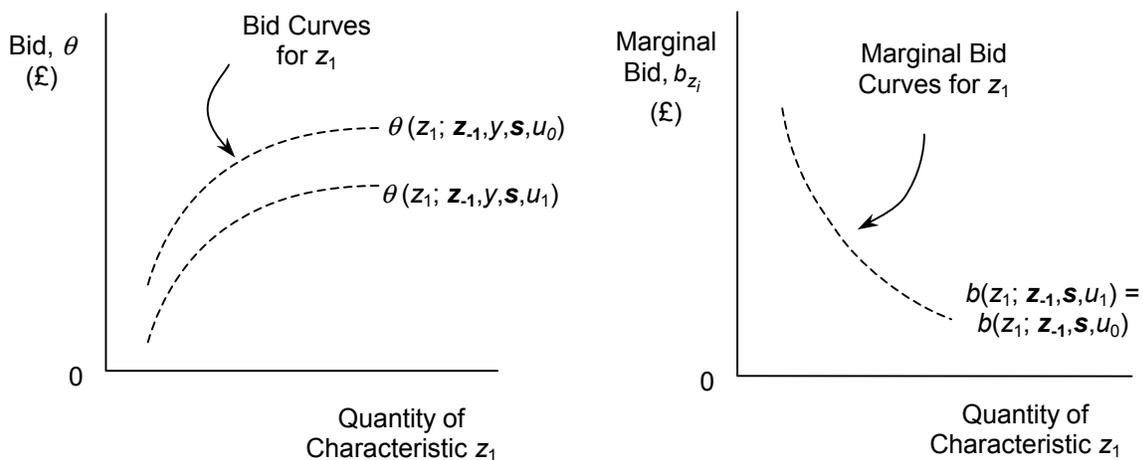
Figure 18: Bid Curves and Marginal Bid Curves



As we would expect, the marginal bid curves in the right hand panel of Figure 18 slope down from left to right. The household is prepared to pay less for each successive unit of attribute z_1 . Though not shown in the figure, at some level of z_1 the marginal bid curves will intercept the horizontal axis. This intercept would reflect the point of satiation at which paying anything for more z_1 would reduce the household's utility below that described by the particular marginal bid curve.

One special case of which we should be aware is when households have quasilinear preferences. This is the case shown in Figure 19. Quasilinear preferences describe indifference curves which are simply vertical translations of each other. Since bid curves are inverted indifference curves, quasilinear preferences can be illustrated as in the left panel of Figure 19 where the bid curves are just vertical translations of each other. Notice that in this case, the slope of the bid curve at all levels of z_1 , is identical for all bid curves no matter what level of utility they represent. With quasilinear preferences, therefore, the household's marginal bid functions lie on top of one another. The relevance of this particular form of preferences will become apparent later.

Figure 19: Bid Curves and Marginal Bid Curves with Quasilinear preferences



In Chapter 1 we showed how the household's choice of property characteristics could be illustrated using bid functions and the hedonic price function. As shown in the left hand panel of Figure 20, the household chooses the bundle of housing attributes that positions them on the bid curve providing the highest level of utility whilst still being compatible with reigning market prices. In other words, the household maximises their utility by moving to the lowest bid curve that is just tangent with the hedonic price function. In the illustration the household's optimal choice is to select a property with \hat{z}_1 of housing attribute z_1 . (Notice that we use a hat to signify optimal choices). This property provides the household with their maximum possible utility, u_1 .

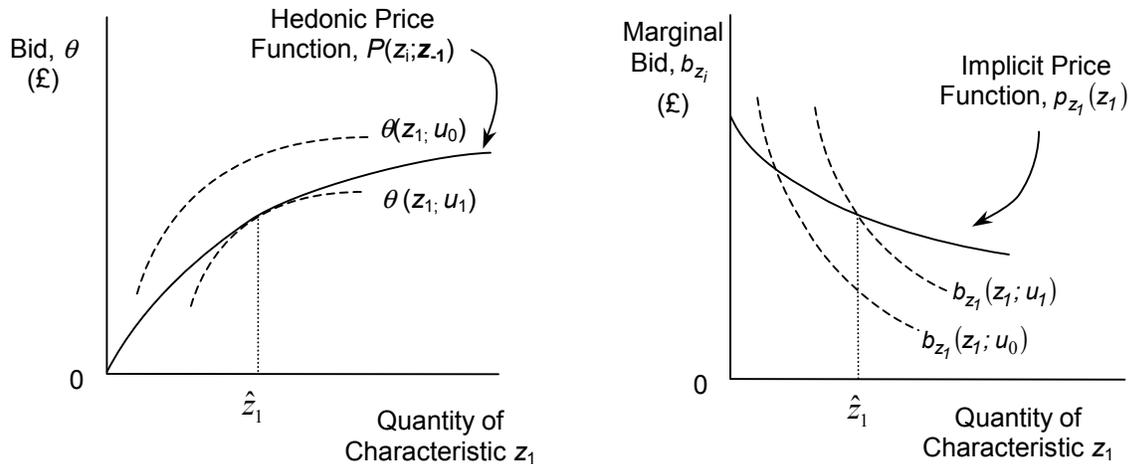
The optimal choice can also be illustrated using marginal bid curves. The right hand panel of Figure 20 plots marginal bid curves corresponding to levels of utility u_0 and u_1 .

On the same graph is drawn the implicit price function for attribute z_1 , $p_{z_1}(z_1)$. Casting our minds back to Chapter 1, remember that the implicit price function describes the additional amount that must be paid by any household in the property market to move to a property with a higher level of characteristic z_1 , other things being equal (see Figure 2). The implicit price function is defined mathematically as the derivative of the hedonic price function with respect to attribute z_i . That is;

$$p_{z_i}(z_i; z_{-i}) = \frac{\partial P(z)}{\partial z_i} \quad (4)$$

Thus $p_{z_1}(z_1)$ is the function giving the marginal price of extra z_1 . Notice that the implicit price is a function and depends on the level of z_1 . (Of course it may also depend on the levels of other housing attributes, z_{-1} , but for simplicity we have suppressed these arguments.) As emphasised in Chapter 1 and illustrated in Figure 20, the implicit price of an attribute does not have to be constant for all levels of z_1 .

Figure 20: Choice of Optimal Attribute Levels using Bid Functions and Marginal Bid Functions

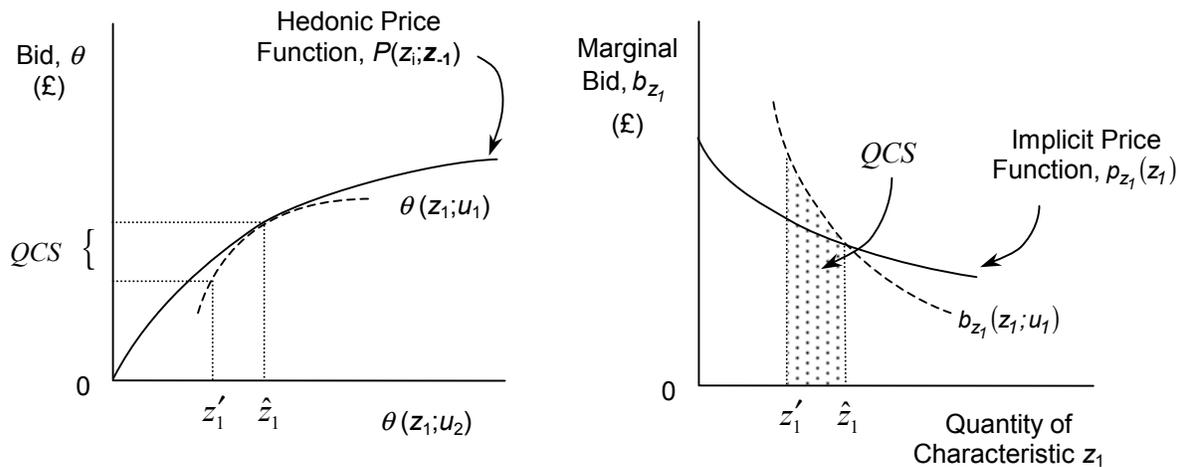


To establish the choice of attribute levels in the marginal analysis one must know in advance the maximised level of utility, u_1 . Then the optimal bundle can be found by moving down the marginal bid curve corresponding to u_1 until the household's marginal willingness to pay for extra z_1 is identical to the marginal price of z_1 in the market¹. This

¹ In some presentations of hedonic theory, it is not made clear that except for the case of quasilinear preferences, there are an infinite number of marginal bid curves each corresponding to a different level of

is very intuitive. The household will always wish to purchase properties with up to \hat{z}_1 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 , since the price that must be paid for each unit of z_1 in excess of \hat{z}_1 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 marginal bid function corresponding to maximised utility and the implicit price function.

Figure 21: Welfare Analysis using Bid Functions and Marginal Bid Functions



The quantity compensating surplus (QCS) defined in Chapter 2 can also be illustrated using marginal bid functions. Imagine a household whose optimal residential location has a level of attribute z_1 given by \hat{z}_1 . An exogenous change decreases the level of z_1 enjoyed at this location to z'_1 . The QCS measure of welfare change is defined as the amount of money that if given to the household whilst living in the same property would make them as well off as they had been previous to the change. In other words, the household's willingness to accept compensation for suffering the fall in the level of z_1 . In the left hand panel of Figure 21 this is illustrated as the difference between the optimising bid curve at \hat{z}_1 and z'_1 .

Now, since, the marginal bid curve is simply the derivative of the bid curve, this amount is exactly equivalent to the shaded area in the right hand panel of Figure 21. That is, the QCS can be measured as the area under the marginal bid curve (corresponding to maximum utility) between the two levels of attribute z_1 .

utility. Moreover, to define the household's optimal choice of housing attributes using marginal bid curves, one must know which of these marginal bid curves corresponds to the maximising level of utility.

c. Identification of the Marginal Bid Function in Multiple Markets

For a moment, let us consider the problem faced by a researcher investigating a hedonic market. To undertake the project, the researcher collects together information on the selling prices of properties in a single market and records details of the attributes of the property and the characteristics of the purchasing household. Using the data on property prices and attributes, the researcher uses multiple regression techniques to estimate the hedonic price function. This is often referred to as the *first stage* of hedonic analysis.

However, the researcher's objective is to estimate *QCS* measures of welfare changes brought about by changes in the environmental attributes of properties. To estimate such welfare measures the researcher needs to know more than the shape of the hedonic price function. As we have seen, *QCS* measures can be defined in terms of the *bid function* or the *marginal bid function*. Consequently, the researcher must undertake further analysis to estimate either of these two functions. This is often referred to as the *second stage* of hedonic analysis.

Theory tells the researcher that at the optimal choice of attributes the slope of the bid function (corresponding to maximised utility) is equal to the slope of the hedonic price function. Thus, second stage analysis proceeds through the researcher calculating the slope of the hedonic price function at each household's choice of property attributes².

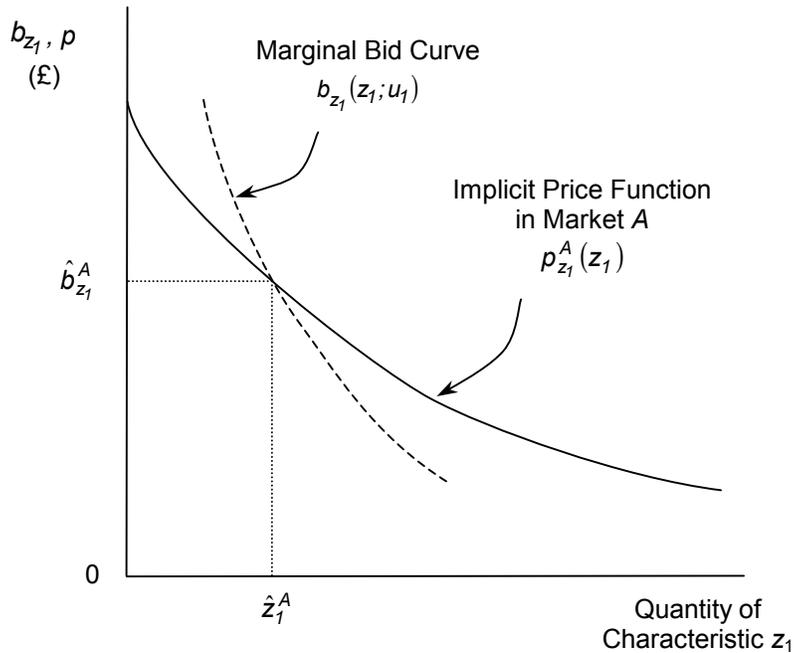
Of course, the slope of the hedonic price function is simply the implicit price of each housing attribute (see Equation 3). Further, as discussed in the previous section, the household's optimal choice of residential location will be such that they equate the implicit price of each housing attribute with the marginal bid curve corresponding to maximised utility (see Figure 20). In short, implicit prices calculated from the first stage analysis provide information on the marginal bid curve. Second stage hedonic analysis, therefore, generally seeks to use the information provided by implicit prices to estimate the marginal bid function.

Consider Figure 22. Here the household choosing a property in Market *A* is faced by the implicit price function for attribute z_1 labelled $p_{z_1}^A(z_1)$. The household chooses a residential location that maximises their utility at level u_1 which corresponds to the marginal bid function shown in the figure. Observing this behaviour in the market, the researcher records just one point on the marginal bid curve. That is, the household's behaviour reveals that for a property boasting \hat{z}_1^A of attribute z_1 the household will be willing to pay $\hat{b}_{z_1}^A$ per unit of z_1 in order to achieve a level of utility u_1 . Unfortunately, knowing one point on the marginal bid curve for u_1 is not sufficient to define the whole curve. Indeed, as various authors have pointed out (e.g. Brown and Rosen, 1982; Murray,

² Of course, the slope of the hedonic price function will be multi-dimensional, having as many dimensions as there are housing attributes. In other words, the slope of the hedonic price function, evaluated at any particular combination of property attributes, will describe the implicit price of an extra unit of each housing attribute.

1982; McConnell and Phipps, 1987) any shaped curve is compatible with this one point provided it passes through $(\hat{z}_1^A, \hat{b}_{z_1}^A)$.

Figure 22: Identifying the Marginal Bid Curve



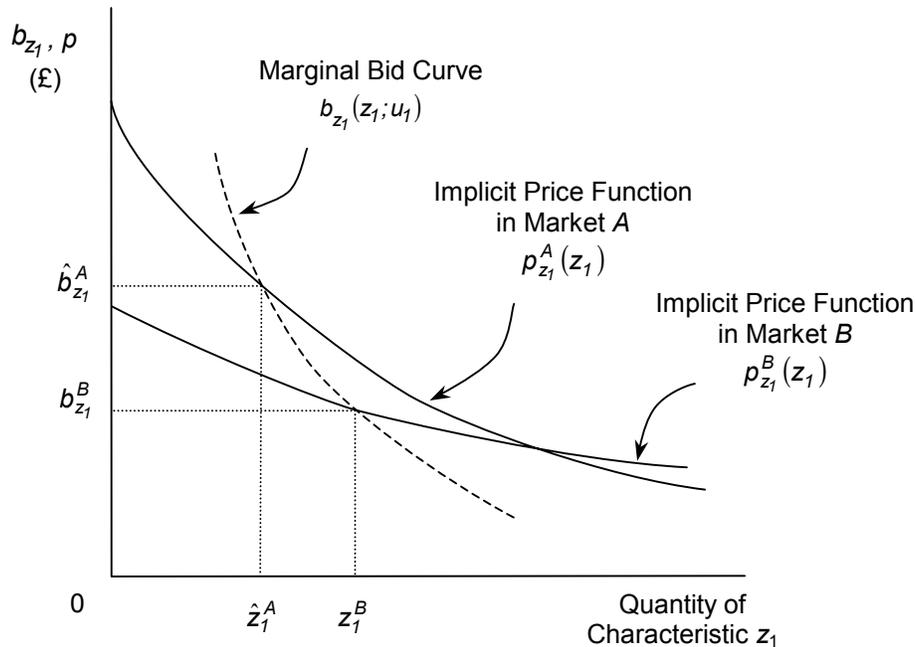
To identify the marginal bid function we would require further information. Specifically, we would need to know the household's marginal bids at alternative levels of z_1 that kept the household at a level of utility u_1 .

One possibility is that such information could be provided by observing the behaviour of another household in a separate market, market B . If this household happens to have identical income (y) and socioeconomic characteristics (s) to the household choosing in market A , then it is assumed that they will have the same preferences. Thus, if both households faced 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 two markets would almost certainly ensure that the equilibrium hedonic price function in market B was different from that in market A .

This is illustrated in Figure 23 where the non-linear implicit price function for market B , $p_{z_1}^B(z_1)$ is also shown. Notice the second implicit price function cuts the marginal bid function for $b_{z_1}(u_1)$ at $(z_1^B, b_{z_1}^B)$. If this were the bundle chosen by the household in market B , then we would have information on the shape of the marginal bid function. Indeed if we could observe the intersection of $b_{z_1}(u_1)$ with a number of different implicit

price functions then we would have the required information to trace out the shape of the marginal bid function.

Figure 23: Identifying the Marginal Bid Curve



Unfortunately, this is not the case. Since the hedonic price function is different in the second market, the second household's optimal choice of residential location may not afford the same level of utility. For example, if prices are generally lower, then the household's maximised level of utility might also be greater, say u_2 . What the researcher would observe in the second market would be the intersection of $b_{z_1}(u_2)$ with $p_{z_1}^2$, and no information would be gained on the shape of $b_{z_1}(u_1)$ ³.

We shall return to discuss this predicament in more detail shortly. For now, however, we can draw the following conclusions;

- In order to estimate the marginal bid function, researchers require information on the choices made by similar households faced by different implicit prices. Estimation of marginal bid curves, therefore, requires data from multiple markets.
- The observed behaviour of households' choices in different markets does not provide the information needed to directly estimate the marginal bid function.

³ Unless of course $b_{z_1}(u_2)$ and $b_{z_1}(u_1)$ were identical. This will only happen in the special case where households have *quasilinear* preferences.

d. Marginal Bid functions and Demand Curves with Linear Hedonic Price Functions

Chapter 1 highlighted the fact that households are unable to “repackage” the different attributes of a property. In other words, households cannot break up a property into its constituent parts and enjoy the benefits of each characteristic separate from the whole. It was shown that the one of the consequences of this feature of hedonic markets is that the hedonic price function may not be linear. That is, it is possible for the price that is paid for each extra unit of a particular housing attribute to vary according to the level of that attribute. Indeed, typically the additional amount paid for properties enjoying increasingly higher quantities of a characteristic (the implicit price of that characteristic) declines as the total level of that characteristic increases. In this section, we return to the issue of non-constant implicit prices and show why this causes problems in the second stage of hedonic analysis.

To illustrate the problem, it is easiest to begin in the counterfactual and assume, for the time being, that implicit prices are constant. Figure 24 depicts the choices made by three identical households⁴ selecting a property in three different markets (markets A , B and C). To simplify the problem further, we shall study only one dimension of the households’ choice problem; their selection of a level of housing attribute z_1 .

Let us focus for the moment, on the choice made by the household in Market A . Here the household faces the hedonic price function P^A . Notice that this is a straight line; the hedonic price function is said to be *linear*. Since the hedonic price is linear its slope is constant. Moreover, if the hedonic price function has a constant slope the implicit price of z_1 in market A , is simply the constant $p_{z_1}^A$.⁵ To emphasise this point, when the hedonic price function is linear, the implicit price function can be described by just one parameter, in this case the constant $p_{z_1}^A$.

The household in market A maximises their utility by moving to the lowest bid curve that is just tangent with the hedonic price function, $\theta(u_1)$. In the illustration the household’s optimal choice is to select a property with \hat{z}_1^A of housing attribute z_1 . This property provides the household with their maximum possible utility, u_1 . This choice point is marked with a dot (as are all other actual choices made by households in the following discussion). We can trace this choice of z_1 down into the lower panel of Figure 24 which shows a marginal analysis of the same information. As discussed in the previous section, the household’s marginal bid is given by the implicit price of z_1 at a level of \hat{z}_1^A . Since,

⁴ That is, each household has the same income, y , and socioeconomic characteristics, s . Since the households are identical, we could alternatively treat them as the same household choosing a property in three different markets. Further, since y and s are identical, these arguments are suppressed in the bid functions and marginal bid functions presented in the text and figures.

⁵ Notice that the implicit price is no longer shown as the function $p_{z_1}(z_1)$, where z_1 in brackets indicates that the implicit price depends on the level of z_1 .

the hedonic is linear the implicit price is simply the constant $p_{z_1}^A$. Hence we can plot one point on the household's marginal bid curve $b_{z_1}(z_1; u_1)$, $(\hat{z}_1^A, p_{z_1}^A)$.

Now let us turn to the household in market B . Notice that the linear hedonic price function in market B , P^B , has a shallower slope than that in market A . Consequently, the constant implicit price of z_1 , $p_{z_1}^B$, in this market is itself lower. Of course, if the price of each unit of z_1 is lower, the household will be able to reach a higher level of overall utility. Indeed, as illustrated in the top panel of Figure 24, the household maximises utility by choosing \hat{z}_1^B of housing attribute z_1 . At this choice point the household is on their highest bid curve consistent with the hedonic price function, $\theta(z_1; u_2)$, where they realise the higher level of utility u_2 . Again we can plot this choice point on the lower panel at $(\hat{z}_1^B, p_{z_1}^B)$.

Notice, however, that $(\hat{z}_1^B, p_{z_1}^B)$ is not a point on the marginal bid curve $b_{z_1}(z_1; u_1)$.⁶ As suggested in the last section, observing the household's choice of z_1 in a second market with a different implicit price does not provide the researcher with the information necessary to trace out the marginal bid curve $b_{z_1}(z_1; u_1)$.

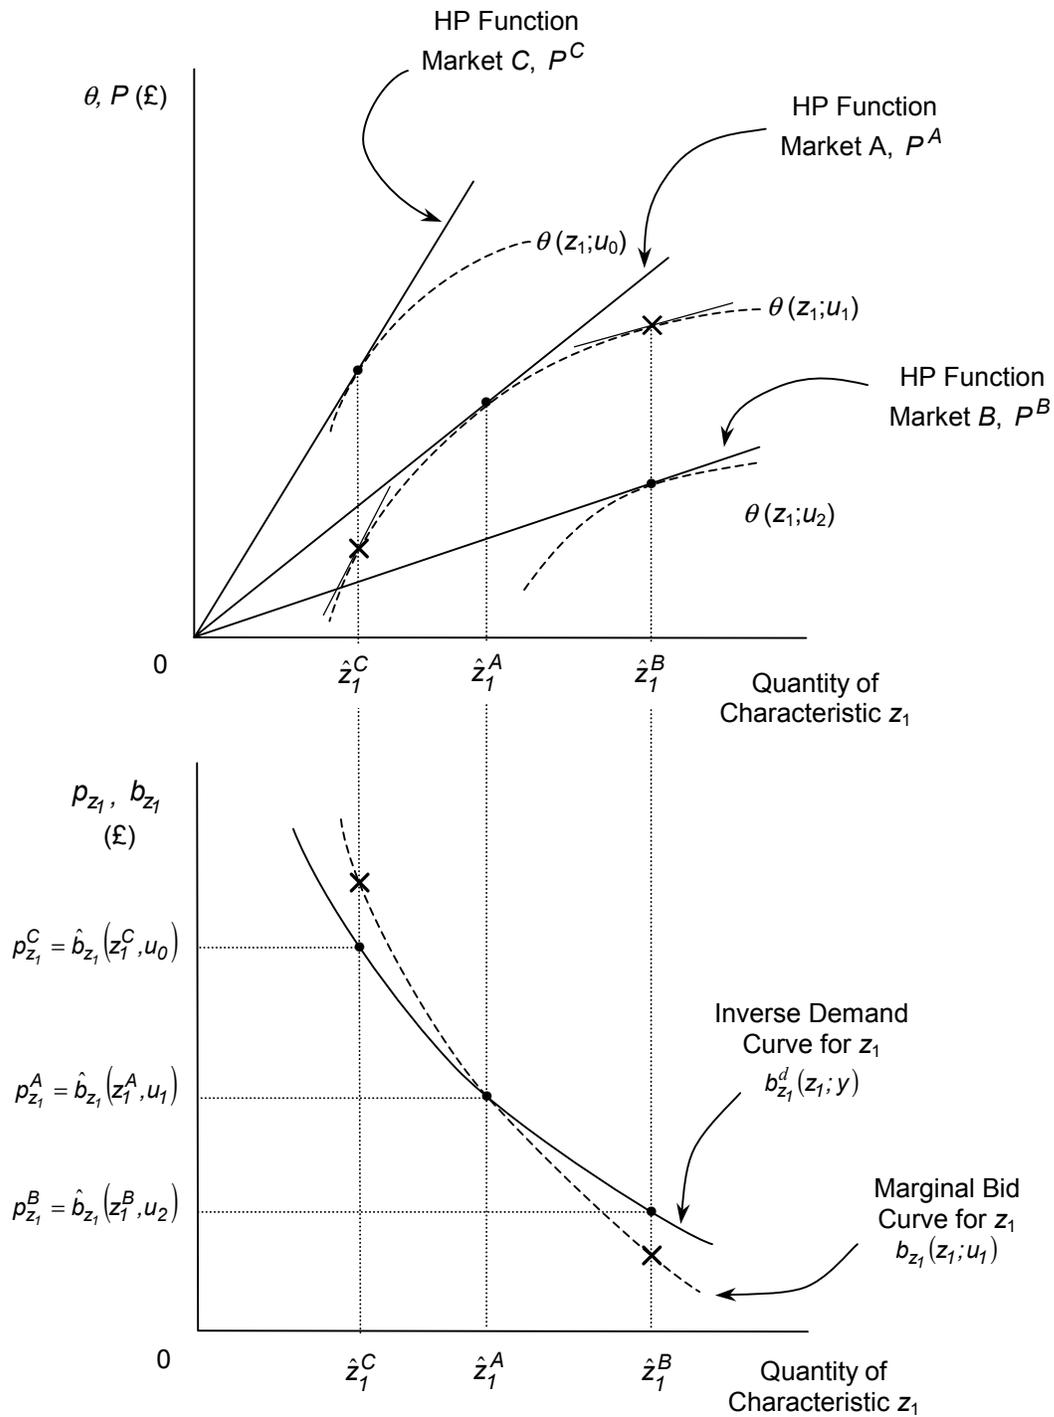
Nevertheless, in our diagrammatic presentation we can locate the point on $b_{z_1}(z_1; u_1)$ corresponding to \hat{z}_1^B . The implicit price in market B , $p_{z_1}^B$, is the household's observed willingness to pay for extra z_1 at \hat{z}_1^B . The amount we are looking for, however, is the household's marginal willingness to pay for extra z_1 at \hat{z}_1^B whilst maintaining a level of utility u_1 .

On the diagram this corresponds to the slope of the bid function $\theta(z_1; u_1)$ at \hat{z}_1^B . This point is marked by a cross on the diagram through which a line tangential to the bid function has been drawn. (In the following discussion crosses indicate behaviour not actually observed in markets). Notice that the slope at this point is slightly shallower than that of the hedonic price function in market B . Consequently, the marginal bid curve $b_{z_1}(z_1; u_1)$ at \hat{z}_1^B will itself be slightly lower than the observed marginal bid at \hat{z}_1^B (i.e. $p_{z_1}^B$). This point is marked on the lower diagram in Figure 24 with a cross.

In general, this will be the case for any attribute if it behaves like a normal good. Only if the household has quasilinear preferences will the two slopes be identical at \hat{z}_1^B . If this were the case the dot and cross in the lower diagram would coincide.

⁶ Rather it is a point on the marginal bid curve $b_{z_1}(z_1; u_2)$. Again, the marginal bid curve $b_{z_1}(z_1; u_2)$ will be different to $b_{z_1}(z_1; u_1)$ unless the household has quasilinear preferences.

Figure 24: Linear hedonic price function and inverse demand curves



Finally, observe the choice made by the household in market C . Here the implicit price of z_1 is the constant $p_{z_1}^C$. Since this is higher than that observed in either of the other markets, the household in market C must make do with a lower level of utility. Indeed, the utility maximising choice of z_1 , \hat{z}_1^C , only affords a level of utility u_0 . Again we can plot the observed behaviour in the lower panel as the point $(\hat{z}_1^C, p_{z_1}^C)$. Meanwhile, the point corresponding to \hat{z}_1^C on the marginal bid curve $b_{z_1}(z_1; u_1)$ is the slope of $\theta(z_1; u_1)$ at \hat{z}_1^C . Notice that this is slightly steeper than the hedonic price function in market C . Hence the marginal bid for z_1 that maintains the level of utility u_1 is higher than the marginal bid observed in the market $p_{z_1}^C$. This point is also plotted in the lower panel of Figure 24. Again if preferences were quasilinear then the dot and cross would coincide.

So far we have managed to plot five points in the lower panel of Figure 24. Those marked with dots represent choices actually observed in the market, those marked with crosses represent behaviour not actually observed.

In fact these five points trace out two separate curves. The first, constructed by joining the dots, is what we would actually observe if we were to collect data on household's property choices from different markets with linear hedonic functions. This curve traces out household's marginal willingness to pay for extra z_1 at different levels of z_1 . For those familiar with economics, this is simply an *inverse ordinary demand curve*. We denote this function;

$$b_{z_1}^d(z_1; y) \quad (37)$$

Where $b_{z_1}^d(\cdot)$ is the inverse ordinary demand function for housing attribute z_1

z_1 is the level of the housing attribute and

y is the household's income

With a linear hedonic price function, the inverse ordinary demand function takes a very simple form sloping down from left to right. As we might expect, at higher levels of z_1 the household is willing to pay less for each extra unit.

The second curve is that which the researcher wishes to identify, the marginal bid curve. This traces out household's marginal bids at different levels of z_1 that maintain a level of utility u_1 . For those familiar with economics, this is simply an *inverse compensated demand curve*. As already stated, we denote this function;

$$b_{z_1}(z_1; u) \quad (36)$$

Where $b_{z_1}(\cdot)$ is the marginal bid curve or inverse compensated demand function for housing attribute z_1

z_1 is the level of the housing attribute and

u is the level of utility

Unfortunately, this second curve is not observed in market behaviour. Crucially, however, the inverse ordinary demand curve and the marginal bid curve will generally be fairly similar (as shown pictorially in the figure).

Indeed, they will be identical if the household has quasilinear preferences. Quasilinear preferences represent the special case where the household has a zero income elasticity of demand for the housing attribute. Remember from Equation (15) that increases in income translate directly (i.e. pound for pound) into increases in the bid function. In effect, increases in income cause the bid curves to shift vertically upwards. Since quasilinear preferences give rise to bid curves that are themselves vertical translations of each other the net effect of an increase in income is that the household moves onto a bid curve representing a higher level of utility but does not change their demand for the good.

In the real world, however, quasilinear preferences are the exception rather than the rule. One might reasonably expect that as a household's income increases their demand for housing attributes would itself increase. Moreover, the greater the income elasticity of demand for the particular attribute the greater the difference between the ordinary inverse demand curve and the marginal bid curve.⁷ On the other hand, theoretical research suggests that within reasonable bounds for the income elasticity of demand the slopes of the two curves will be reasonably similar (Willig, 1976).

One possibility, therefore, is that researchers use market data to estimate the ordinary inverse demand curve. Approximate *QCS* welfare measures can be estimated as the area under the inverse demand curve between the two levels of attribute z_1 . Further, if this approximation is thought to result in serious error, there are techniques by which the researcher can retrieve the marginal bid curve from an estimated inverse demand curve, we shall return to this in later discussion.

e. Marginal Bid Functions and Demand Curves with Nonlinear Hedonic Price Functions

In a world with purely linear hedonic price functions, therefore, everything seems rosy. Market data can be used to estimate the inverse demand function and this should provide a reasonably good approximation to the marginal bid function. However, in the real world, hedonic price functions are not linear and there's the rub. When implicit

⁷ The difference between the slopes of the two curves will also depend on the significance of expenditure on that attribute as a part of the consumer's budget.

prices are not constant and preferences are not quasilinear, the inverse demand curve as we have illustrated it does not exist.

To illustrate observe Figure 25. Here we have done away with the assumption of linear hedonic price functions and quasilinear preferences. Now the hedonic price functions in markets A , B and C are all non-linear. The figure has been constructed such that the households in the three markets maximise their utility by choosing the exact same levels of z_1 as were illustrated in the linear case of Figure 24. Further, the diagram has been drawn such that the household in market A achieves the same level of utility, u_1 , at their optimal choice of z_1 as was chosen facing the linear hedonic price function in Figure 24.

By construction, therefore, the point in the lower panel of Figure 25 corresponding to the choice of the household in market A , is identical to that in Figure 24; $(\hat{z}_1^A, p_{z_1}^A)$. Once again, this describes one point on the marginal bid function $b_{z_1}(z_1; u_1)$.

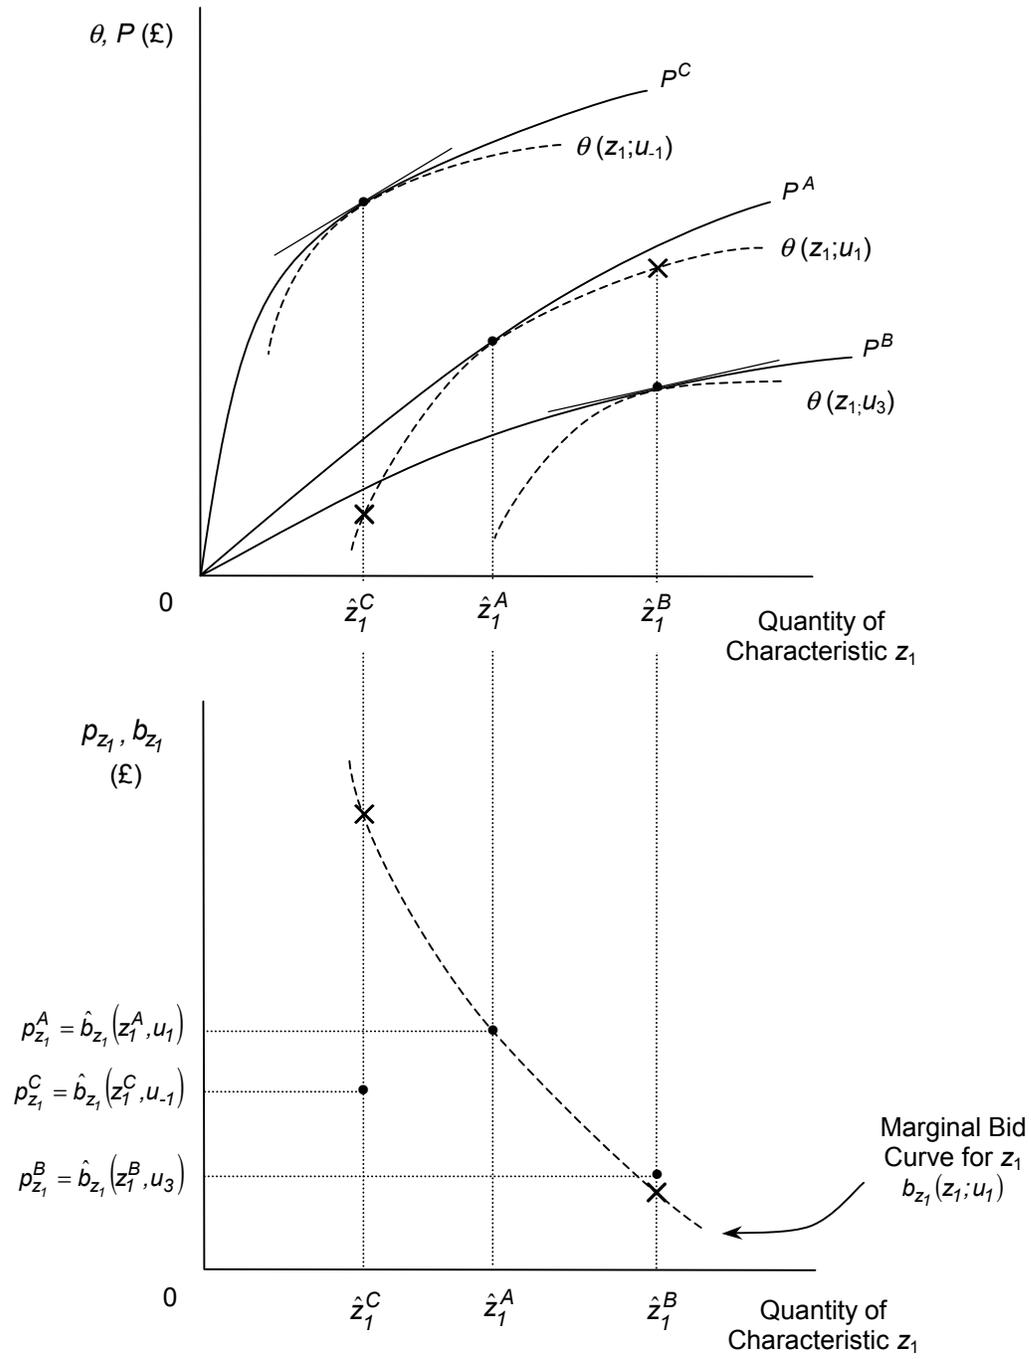
Consider now the choice of the household in market B . The non-linear hedonic price function in this market is in all places lower than that in market A . Consequently, the price paid for any level of z_1 in market B is less than that paid for the same level of z_1 in market A . Not surprisingly, therefore, the household in market B , manages to achieve a higher level of utility, u_3 , whilst choosing a higher level of z_1 , \hat{z}_1^B .

Following a now familiar procedure, we can plot this choice point in the lower panel of Figure 25 by determining the implicit price of z_1 at \hat{z}_1^B as the slope of the bid function $\theta(z_1; u_2)$ at \hat{z}_1^B . Notice that because of the non-linear hedonic price function, the implicit price at \hat{z}_1^B is not necessarily the same as the implicit price at other levels of z_1 .

In the linear case, this choice point defined a second point on the inverse ordinary demand curve. Indeed, we might expect that in this non-linear case we could trace out a similar shaped curve. Certainly this second point in the lower panel of Figure 25 would seem to be following the correct pattern. As we would expect, the household's willingness to pay for z_1 at this higher level of provision is lower than that observed at the lower level of provision chosen in market A . Further, if we plot the marginal bid function $b_{z_1}(z_1; u_1)$ at this level of provision it falls below that observed in market choices. Again the result observed in the linear hedonic price function case.

However, observe the choice made by the household in market C . Since the hedonic price function is in all places higher than that in market A , it comes as no surprise that the household's optimal choice, is at a lower level of provision and affords them a lower level of overall utility, u_1 . When we come to plot this choice point in the lower panel, however, we are struck by an anomaly. At \hat{z}_1^C facing the hedonic price function in market C , the household's marginal willingness to pay for extra z_1 is lower than that recorded in market A . This is despite the fact that the household in market C has chosen a property with lower levels of z_1 than that chosen in market A .

Figure 25: Non-linear hedonic price function and inverse demand curves (1)



Clearly, with non-linear hedonic price functions and preferences that are not quasilinear, observed choices do not plot out a nice downward sloping inverse ordinary demand curve⁸.

To emphasise this point consider Figure 26 where a fourth identical household is observed choosing a property in market *D*. Here, the household maximises their utility by choosing \hat{z}_1^D of the housing attribute. Whilst this is an identical quantity to that chosen by the household in market *A*, the slope of the hedonic price function in market *D* is shallower than that in market *A*. Plotting this on the marginal analysis diagram we see that with nonlinear hedonic price functions, the same level of demand can be associated with two different implicit prices. To summarise, when implicit prices are non-constant and preferences are not quasilinear, the inverse ordinary demand curve as normally conceived is not well defined. A household's marginal willingness to pay for extra z_1 at any level of z_1 will depend on the shape of the entire hedonic price function faced in that market.

The problem is further complicated when we move out of the unidimensional problem of choosing just one housing attribute level and consider choice across many attributes. In this case, if patterns of substitutability and complementarity exist between the attribute of interest and the other attributes, then the household's marginal willingness to pay for extra z_1 at any level of z_1 will also depend on the shape of the hedonic price function for all these attributes.

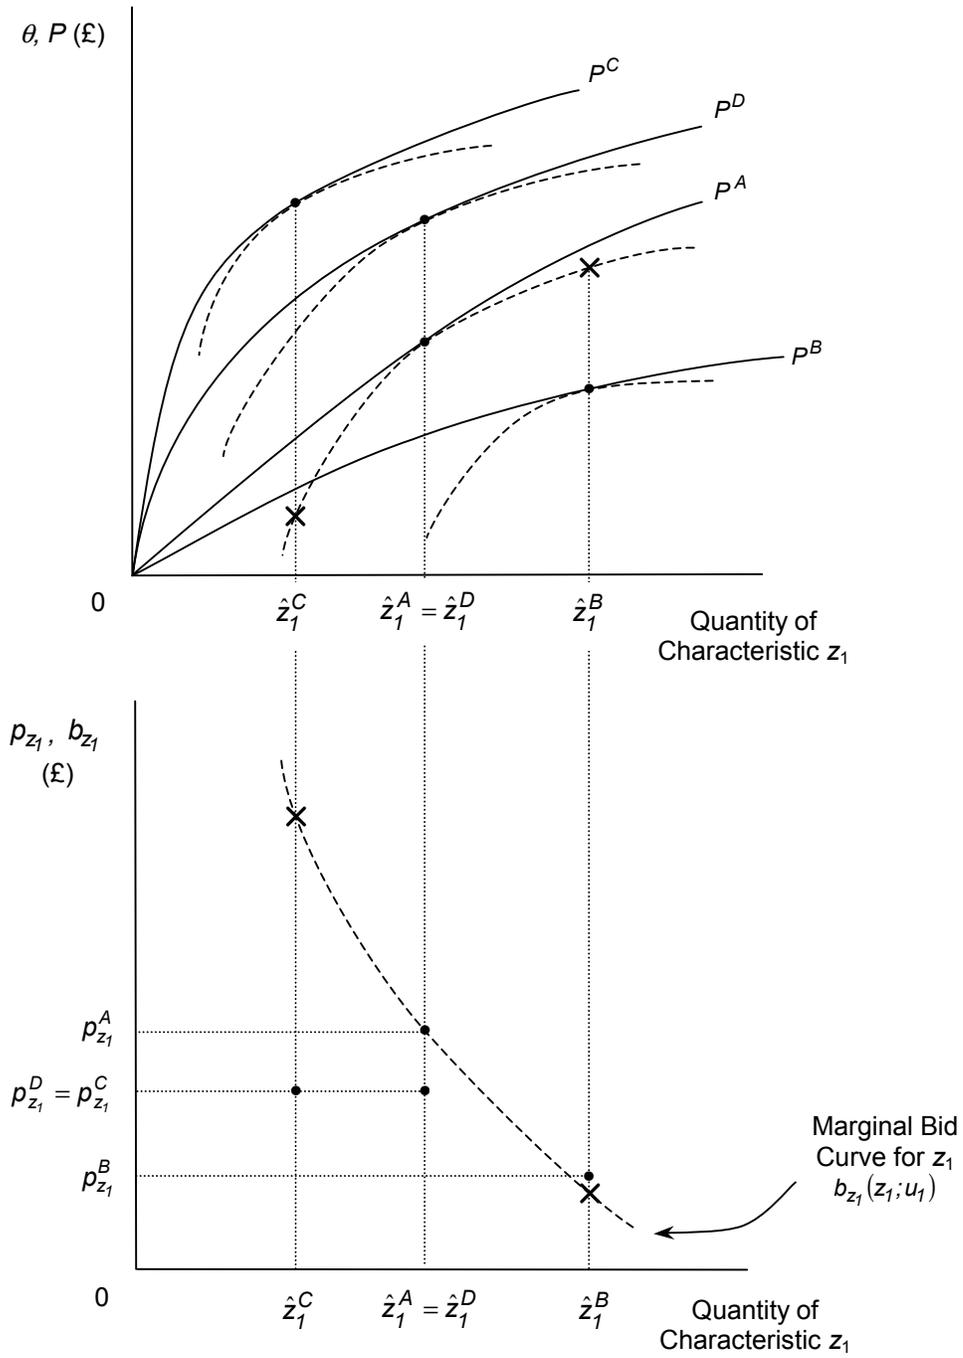
This presents a considerable problem for welfare analysis in hedonic markets. Specifically, it becomes impossible to estimate a simple inverse ordinary demand function for an attribute of interest. That is, a simple regression of the implicit prices paid for an attribute by different households against quantities of this attribute, quantities of other attributes and household income will not yield a classic downward sloping inverse demand curve⁹. Indeed, when marginal prices are non-constant there is no reason for us to expect any relationship between marginal willingness to pay for an attribute and the quantity of that attribute presently enjoyed¹⁰.

⁸ Note that if preferences were quasilinear then the slope of the bid function at any particular level of housing attribute would be the same for all bid curves. In this special case, the existence of nonlinear hedonic price functions does confound the existence of a downward sloping inverse ordinary demand curve.

⁹ Remembering that identification of such a function would require data on households in different markets facing different hedonic price functions

¹⁰ This observation suggests that simple meta-analyses of the summary results of hedonic analyses have little theoretical basis. For example, a number of authors (Smith and Huang, 1995; Schipper, 1996; Bertrand, 1997) have carried out meta-analyses using results from various hedonic property price studies reporting households' marginal willingness to pay to avoid pollution (i.e. 'average' implicit prices for pollution). Amongst other things, these meta-analyses have sort to establish the relationship between marginal willingness to pay to avoid pollution and current levels of pollution. The discussion in this section shows that in the face of non-linear hedonic price functions, no simple relationship between the two exists.

Figure 26: Non-linear hedonic price function and inverse demand curves (2)



f. Mythical Demand Curves: Linearising the Budget Constraint

Fortunately, as pointed out by Murray in 1983 and later by Palmquist (1988) the problems introduced by nonlinear hedonic price functions can be overcome. In short, the solution requires the budget constraint to be linearised around the optimal choice of housing attributes. This linearised budget constraint is defined by a set of constant implicit prices and an income level that we shall call the household's "mythical" income (Murray's terminology). It so happens that the bundle of housing attributes chosen by the household faced with the nonlinear hedonic price function will be the same as that they would have chosen if they had this mythical income and were faced by the linear hedonic price function. In effect, the technique of linearising the budget constraint allows the researcher to treat the choices made by households as if they were choices made in response to constant implicit prices. Of course, with constant implicit prices the inverse ordinary demand function is defined by Equation (37) and takes on its classic downward sloping curve. This "mythical" inverse ordinary demand function should be a reasonable approximation to the household's marginal bid curve.

The technique of linearising the budget constraint is illustrated in Figure 27. The top panel of this diagram depicts the choice of housing attribute z_1 made by two households faced by the same nonlinear hedonic price function. Let us assume that these two households have the same socioeconomic characteristics, s , but that household b has a higher income than household a . That is y_b is greater than y_a .

We can just as well illustrate these choices in the indifference diagram in the lower panel. This diagram plots indifference relationships between money to spend on other goods, the numeraire, and the level of housing attribute z_1 . Since the hedonic price function is nonlinear, the budget constraints faced by the two households are themselves nonlinear. Notice that the budget constraint for household b is simply a vertical translation of that faced by household a . The actual incomes of the two households will be given by the point where the budget constraints intercept the y-axis and these two amounts are labelled on the diagram¹¹.

Consider now the choice made by household a . This household optimises their utility by choosing a level of the housing attribute labelled \hat{z}_1^a at which the implicit price of z_1 is $\hat{p}_{z_1}^a$. At this point we wander into the realms of the "mythical" rather than real worlds.

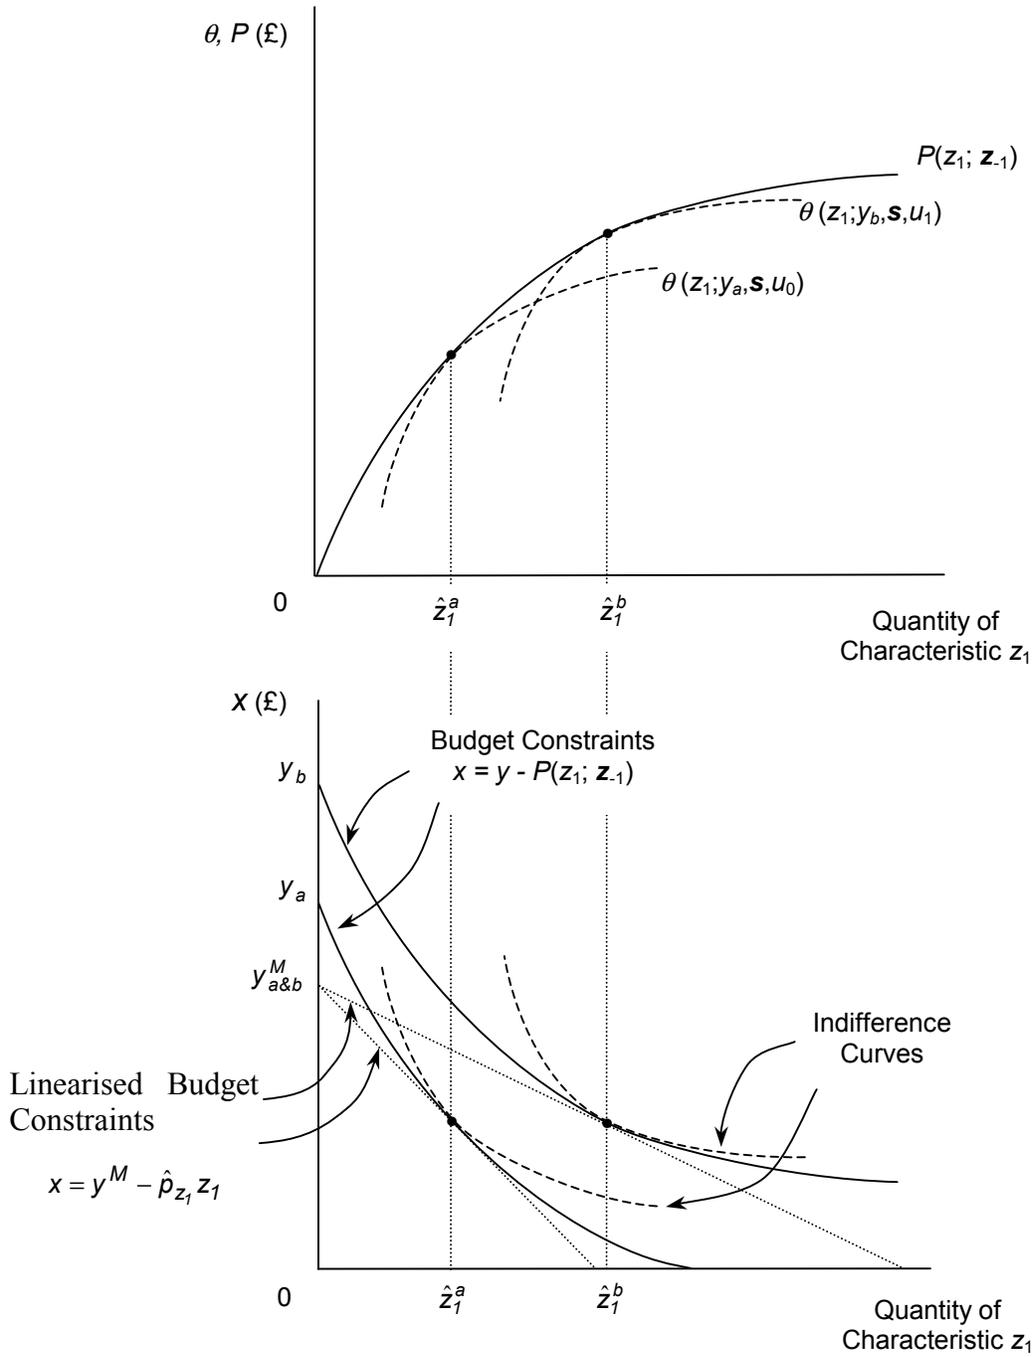
Imagine that the implicit price at this optimal choice of housing attributes was actually a constant marginal price coming from a linear hedonic function. If this were so we could construct a budget constraint running through the household's optimal choice with a slope of $\hat{p}_{z_1}^a$. The intercept of this mythical budget constraint gives household a 's

mythical income y_a^M . The important thing to note is that the choice of property attributes made by household a with income y_a facing the nonlinear hedonic price function is

¹¹ We assume that households would not be willing to pay anything for a house with no z_j . For example, if z_j represents "peace and quiet", then this assumption amounts to saying that there is a point where a household would not purchase a property because it is too noisy to live in.

identical to that which they would have made if they had an income of y_a^M and faced a linear hedonic function with constant marginal price $\hat{p}_{z_1}^a$.

Figure 27: Linearising the budget constraint



Now consider the choice made by household b . Following the same procedure, we can construct a mythical linear budget constraint whose slope is defined by the implicit price of the attribute at household b 's optimal choice, $\hat{p}_{z_1}^b$. The intercept of this budget constraint with the y-axis gives household b 's mythical income y_b^M . Again, the bundle of attribute quantities chosen by household b will be identical whether they are making choices in the real world with the nonlinear hedonic function and income y_b or in the mythical world with the linear hedonic price function and income y_b^M .

The diagram has been constructed such that both households have the same mythical income. Notice that the decisions made by these two households could just as well be treated as the those made by a single mythical household with income $y_{a\&b}^M$ choosing a property in two separate markets. In the first market this mythical household faces a linear hedonic price function in which z_1 has the constant implicit price $\hat{p}_{z_1}^a$ in the second the household faces a linear hedonic price function with the slightly lower constant implicit price $\hat{p}_{z_1}^b$. As we would expect, the household facing the lower price chooses more z_1 . Indeed, given observations from many households with the same mythical income we could trace out the entire mythical ordinary demand curve. Since in the mythical world all hedonic price functions are linear the mythical ordinary demand curve is well defined. In fact this mitigates a simple procedure for estimation;

- Estimate the hedonic price function¹²,
- Calculate the implicit price for each housing attribute
- Calculate the implied mythical income at these implicit prices according to;

$$y^M = y - P(\hat{z}) + \sum_{i=1}^K \hat{p}_i \hat{z}_i \quad (38)$$

- Estimate the mythical inverse ordinary demand curve by regressing the implicit price for an attribute on the chosen quantities of the attribute, chosen quantities of other attributes and mythical income;

$$\hat{p}_{z_1} = b_{z_1}^M(\hat{z}_1, \hat{z}_{-1}, y^M, \mathbf{s}) \quad (39)$$

More typically, researchers estimate the mythical ordinary demand function;

¹² Note that we still require data from more than one market to ensure identification of the mythical ordinary inverse demand curve.

$$\hat{z}_1 = z_{z_1}^M(\hat{p}_{z_1}, \hat{\mathbf{p}}_{z_{-1}}, y^M, \mathbf{s}) \quad (40)$$

where $\hat{\mathbf{p}}_{z_{-1}}$ is the vector of all other attribute chosen implicit prices.

Equation (40) tends to be seen as a more natural specification than Equation (39) since it is the z_i rather than the p_{z_i} which are the observed outcome of household's choices in hedonic markets. Note carefully, however, that in hedonic markets, where marginal prices are nonlinear household's actually simultaneously choose both the quantities and the marginal price of attributes.

g. Mythical demand curves: Estimation and welfare analysis

Ideally, the researcher would estimate a system of demand curves for all property attributes. In reality, however, the usual procedure is to concentrate on one or a number of attributes that form the focus of the research programme. Further, rather than including all attribute quantities in the regression and imposing the theoretical restrictions on Equations (39) and (40) required by demand theory, researchers employ fairly simple functional forms, including only a handful of other attribute quantities.

Econometric estimation of mythical ordinary demand curves is further complicated by problems of endogeneity. As we have seen, in hedonic markets, the marginal price of housing attributes will generally not be constant. In maximising their utility from the choice of residential location, the household chooses both the quantity of housing attributes and the marginal price of the attributes. In estimating, Equation (40), therefore, the implicit prices of housing attributes on the right hand side of the equation are *endogenous*. Further, since mythical income is calculated using the chosen level of marginal price (Equation 38), this too is endogenous. Unless researchers account for this endogeneity, the parameter estimates from the econometric estimation of the mythical inverse ordinary demand curve will be biased.

Typically, endogeneity is handled through the application of instrumental variable techniques. The trick here is to regress each of the endogenous variables in the demand equation on a set of exogenous variables that in this context are referred to as instruments. The results of these ancillary regressions are used to calculate predicted values for the endogenous variables. The demand equations are estimated using these predicted rather than the actual values of the endogenous variables. Avoiding the econometric details, the instrumental variables technique removes the problem of biased parameter estimates caused by the inclusion of endogenous regressors in the demand equations.

This all seems very straightforward, however, difficulties arise in choosing suitable instruments. These variables should be highly correlated with the endogenous variable they are being used to predict but at the same time should not be correlated with the error term entering the demand equation. For example, imagine that we were choosing

instruments for the household's mythical income. Suitable candidates might include the household's socioeconomic characteristics including the number of members of the household, their ages and educational status. Suitable instruments for implicit prices could once again include socioeconomic traits but authors have also suggested using the marginal price paid by similar households, where similarity is determined either in terms of these household's socioeconomic characteristics (Murray, 1983) or their spatial proximity (Cheshire and Sheppard, 1998).

With the mythical ordinary demand curve estimated, approximate *QCS* measures of welfare change can be obtained by integrating under this curve between the initial level of the attribute and that following some external change.

Some authors have taken the process one step further and attempted to derive exact *QCS* measures by estimating the household's marginal bid function. Such approaches rely upon duality results between the inverse ordinary demand curve and the inverse compensated demand curve. However, we do not discuss these issues further in this document.

Table 4 presents a step by step guide to hedonic analysis, from collecting data through to welfare estimation.

Table 4: Steps to Perform a Hedonic Analysis

<i>Step 1</i>	<p><i>Collect data</i></p> <p>This should include;</p> <p>Property sales prices and</p> <p>the socioeconomic characteristics of purchasing households</p> <p>Data should provide information on the choices made by households in two or more independent hedonic property markets.</p>
<i>Step 2</i>	<p><i>Estimate Hedonic Price Function for each market</i></p> <p>Regress property prices on property characteristics according to;</p> $P = P(z_1, z_2, \dots, z_K)$ <p>Repeat for each separately identified property market</p> <p>Test for market segmentation with each property market</p>
<i>Step 3</i>	<p><i>Calculate Implicit Prices chosen by Households</i></p> <p>For each household, calculate the implicit price of housing attributes according to;</p> $p_{z_i}(z_i; \mathbf{z}_{-i}) = \frac{\partial P(\mathbf{z})}{\partial z_i}$

<p><i>Step 4</i></p>	<p><i>Calculate each Household's Mythical Income</i></p> <p>Using the implicit prices estimated in step 3 calculate each household's mythical income according to;</p> $y^M = y - P(\hat{z}) + \sum_{i=1}^K \hat{p}_i \hat{z}_i$
<p><i>Step 5</i></p>	<p><i>Calculate instruments for Implicit Prices and Mythical Income</i></p> <p>Select instruments for implicit prices. Suitable candidates include;</p> <ul style="list-style-type: none"> • Socioeconomic characteristics • Implicit prices chosen by similar (demographic traits/spatial proximity) households <p>Select instruments for Mythical Income. Suitable candidates include;</p> <ul style="list-style-type: none"> • Socioeconomic characteristics <p>Using data from all markets estimate two ancillary equations regressing observed implicit prices and mythical income on instruments</p> <p>Use the regression results to calculate predicted values for implicit prices and mythical income for each household. Call these; \tilde{y}^M and \tilde{p}_{z_i}</p>
<p><i>Step 6</i></p>	<p><i>Estimate Mythical Ordinary Demand Function</i></p> <p>Using predicted values calculated in step 5 estimate the demand function according to;</p> $\hat{z}_1 = z_{z_1}^M(\tilde{p}_{z_1}, \tilde{p}_{z_{-1}}, \tilde{y}^M, s)$
<p><i>Step 7</i></p>	<p><i>Calculate QCS welfare measures</i></p> <p>Integrate under the mythical demand curve between the initial level of the attribute and that following some external change</p>

h. Mythical Demand Curves: 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 mythical demand curves. Rather researchers have gone no further than Step 3, 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 mythical demand curves have been estimated (e.g. Cheshire and Sheppard, 1998; Palmquist and Isangkura, 1999; Boyle et al., 1999 and Zabel and Kiel, 2000). Mythical demand curves, represent underlying household preferences for housing attributes. Consequently they can be used to derive theoretically consistent estimates of household's welfare changes¹³. 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. Specifically, to calculate implicit prices and mythical income (the arguments of the mythical demand function) in the transfer location would require knowledge of the hedonic price function in that market.

However, it may be possible to make approximations that avoid the need to estimate the hedonic price function in the transfer location. First, the mythical *inverse* demand function should be estimated as in Equation (40). The transfer equation will then contain housing attribute levels as its arguments rather than implicit prices. Further, future hedonic analyses should report relatively simple specifications of the mythical inverse demand function. For example, the function could be estimated using just the quantity of the housing attribute of interest, mythical income and socioeconomic variables that might easily be recovered in the transfer location.

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. Such a procedure would necessarily generate welfare values that are an approximation to the true change (most notably in that the transfer function is unrealistically simple and that actual rather than mythical income is used in the transfer location). Future research should 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.

¹³ As discussed in Chapter 2, these welfare estimates represent only those accruing to households and not those accruing to landlords. Moreover, they are only lower bounds for this value. Complete welfare estimates require information on the response of the hedonic price function to changes in the conditions of supply and demand brought about by a change in the provision of a housing attribute. The complexity of the market mechanism in hedonic markets means that it is rarely possible to predict such changes. In general, complete welfare measures will only be possible ex-post, when researchers have information on the hedonic price function before and after the change.

References

- Bartik, T. J. (1988). "Measuring the Benefits of Amenity Improvements", *Land Economics*, 64(2), pp 172-183.
- Bertrand, N. F. (1997). "Meta-analysis of studies of willingness to pay to reduce traffic noise", unpublished MSc dissertation, University College London.
- Boyle, K. J., P. J. Poor and L. O. Taylor (1999). "Estimating the demand for protecting freshwater lakes from eutrophication", *American Journal of Agricultural Economics*, 81(5), pp1118-22.
- Brown, J. N. and H. S. Rosen (1982). "On the estimation of structural hedonic price models", *Econometrica*, 50(3), pp 765-768.
- Cheshire, P. and S. Sheppard (1998). "Estimating the demand for housing, land and neighbourhood characteristics", *Oxford Bulletin of Economics and Statistics*, 60(3), pp357-382
- Epple D. (1987). "Hedonic prices and implicit markets: Estimating demand and supply functions for differentiated products", *Journal of Political Economy*, 95, pp 59-80.
- Freeman, A. M. (1993). *The measurement of environmental and resource values: Theory and methods*, Resources for the Future, Washington, D.C.
- Harrison, D. and D. L. Rubinfeld (1978). "Hedonic housing prices and the demand for clean air", *Journal of Environmental Economics and Management*, 5, pp 81-102.
- McConnell, K. E. and T. T. Phipps (1987). "Identification of preference parameters in hedonic models: Consumer demands with nonlinear budgets", *Journal of Urban Economics*, 22, pp 35-52.
- Murray, M. P. (1983). "Mythical demands and supplies for proper estimation of Rosen's hedonic price model", *Journal of Urban Economics*, 14, pp 327-337.
- Palmquist, R. B. (1984). "Estimating the demand for the characteristics of housing", *Review of Economics and Statistics*, 66, pp 394-404.
- Palmquist, R. B. (1988). "Welfare measurement for environmental improvements using the hedonic model: The case of nonparametric marginal prices", *Journal of Environmental Economics and Management*, 15, pp 297-312.
- Palmquist, R. B. (1991). "Hedonic methods" in *Measuring the Demand for Environmental Quality*, J. B. Braden and C. D. Kolstad (eds.), Amsterdam: North Holland.
- Palmquist, R. B. and A. Isangkura (1999). "Valuing air quality with hedonic and discrete choice models", *American Journal of Agricultural Economics*, 81(5), pp1128-33.
- Rosen, S. (1974). "Hedonic prices and implicit markets: Production differentiation in pure competition", *Journal of Political Economy*, 82, pp 34-55.

- Schipper, Y. J. J. (1996). "On the valuation of aircraft noise: A meta-analysis", unpublished paper, Titenberg Institute, Free University of Amsterdam.
- Smith V. K. and J. C. Huang (1995). "Can markets value air quality? A meta-analysis of hedonic property values", *J. Pol. Econ.*, 103, pp 209-227.
- Willig, R. D. (1976). "Consumer's surplus without apology", *American Economic Review*, 66(4), pp 589-597.
- Zabel, J. E. and K. A. Kiel (2000). "Estimating the demand for air quality in four U.S. cities", *Land Economics*, 76(2), pp 174-19

