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 $\boldsymbol{z}$ 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, $\boldsymbol{s}$, 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;

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

Thus the marginal bid function is given by;

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

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

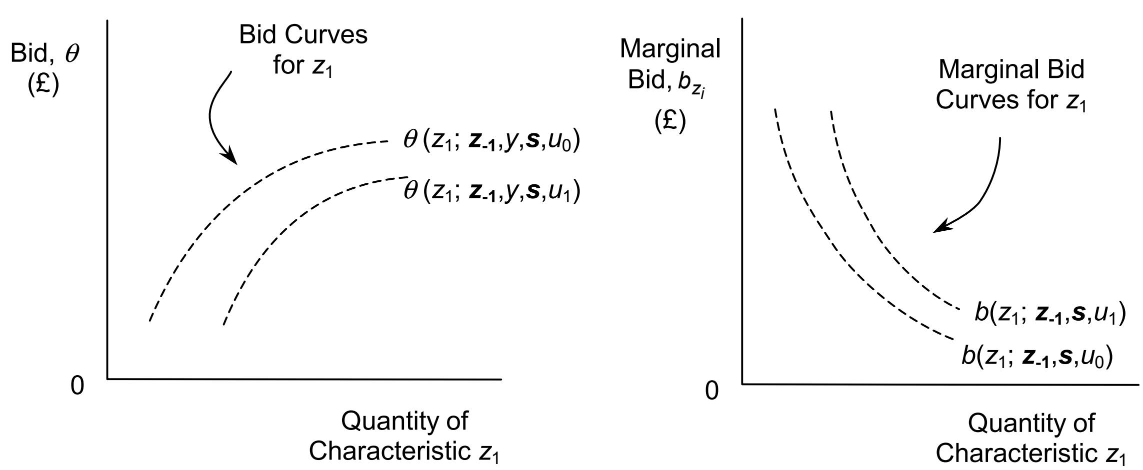

As we would expect, the marginal bid curves in the right hand panel of Figure 18 slope down from left to right. The household is prepared to pay less for each successive unit of attribute z_1. Though not shown in the figure, at some level of z_1 the marginal bid curves will intercept the horizontal axis. 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

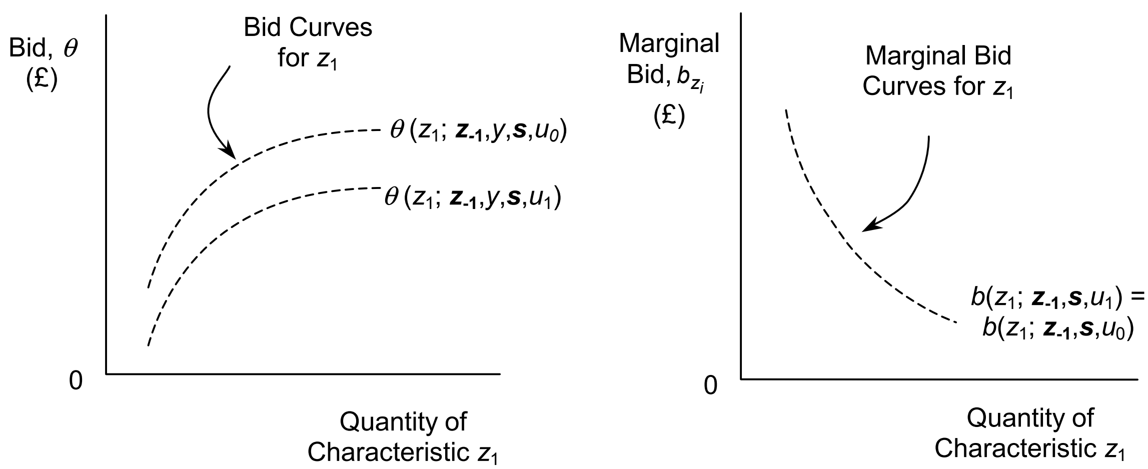

In Chapter 1 we showed how the household's choice of property characteristics could be illustrated using bid functions and the hedonic price function. As shown in the left hand panel of Figure 20, the household chooses the bundle of housing attributes that positions them on the bid curve providing the 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;

$$p_{z_i}(z_i; \mathbf{z}_{-i}) = \frac{\partial P(\mathbf{z})}{\partial z_i} \qquad (4)$$

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, $\mathbf{z}_{-1}$, 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

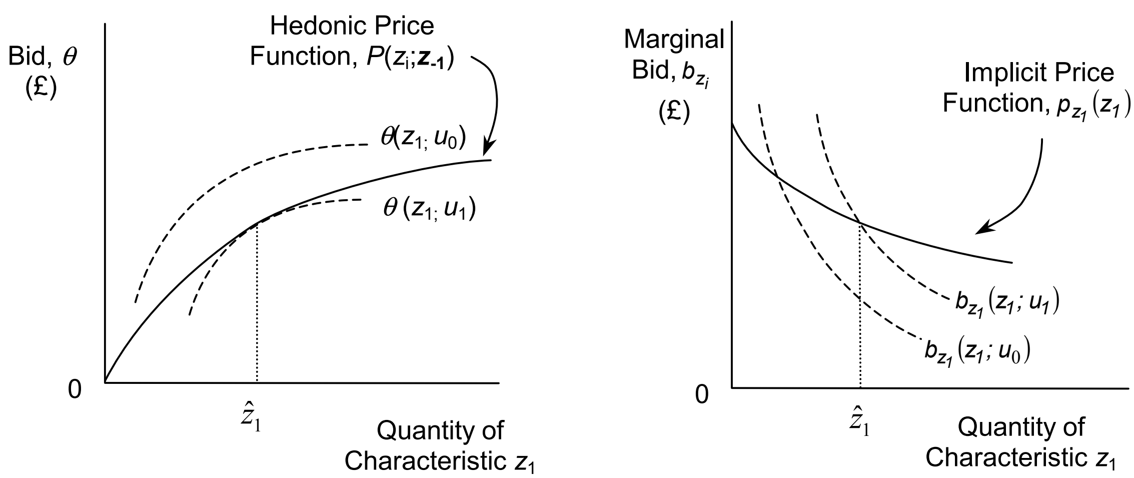

To establish the choice of attribute levels in the marginal analysis one must know in advance the maximised level of utility, u_1. Then the optimal bundle can be found by moving down the marginal bid curve corresponding to u_1 until the household's marginal willingness to pay for extra z_1 is identical to the marginal price of z_1 in the market[1]. 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

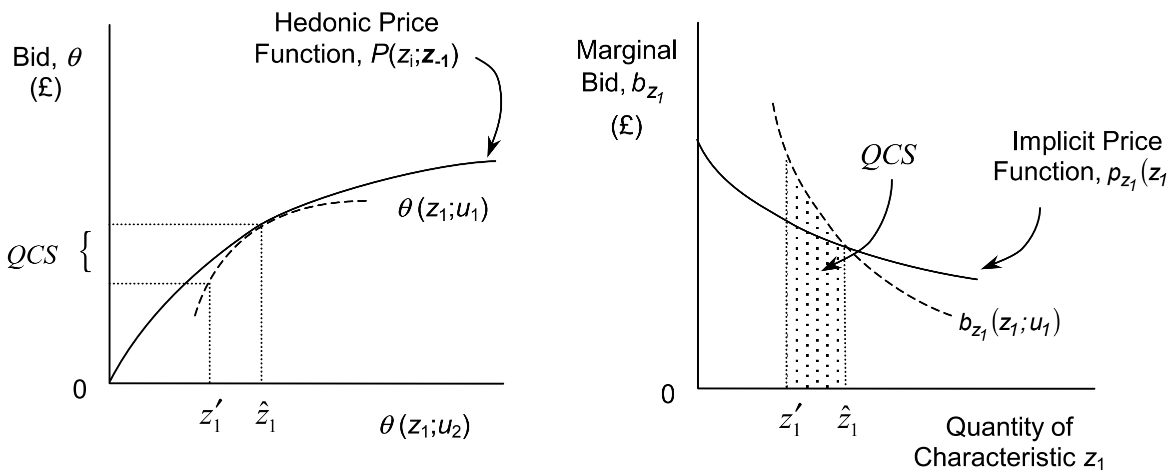

The quantity compensating surplus (*QCS*) defined in Chapter 2 can also be illustrated using marginal bid functions. Imagine a household whose optimal residential location has a level of attribute z_1 given by $\hat{z}_1$. An exogenous change decreases the level of z_1 enjoyed at this location to z_1'. 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 households choice of property attributes[2].

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

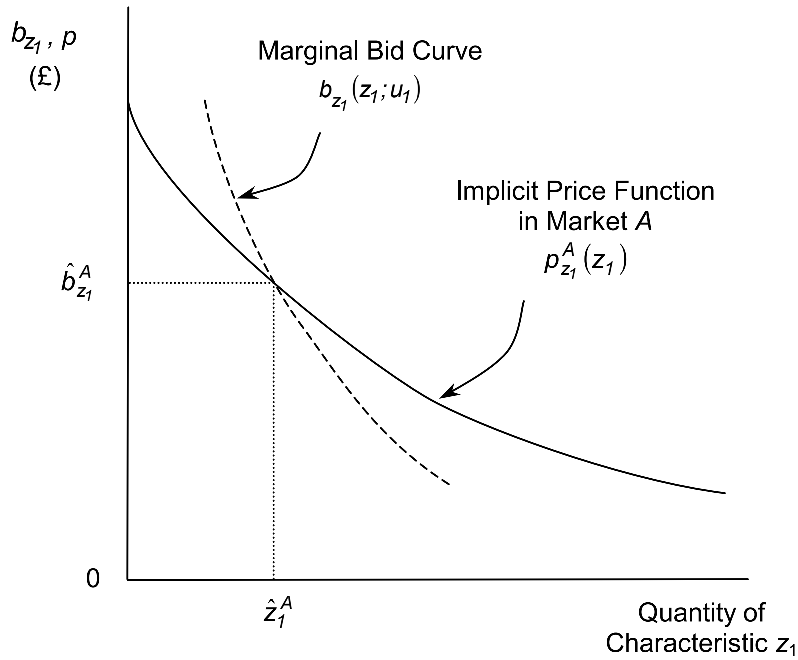

To identify the marginal bid function we would require further information. Specifically, we would need to know the household's marginal bids at alternative levels of z_1 that kept the household at a level of utility u_1.

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

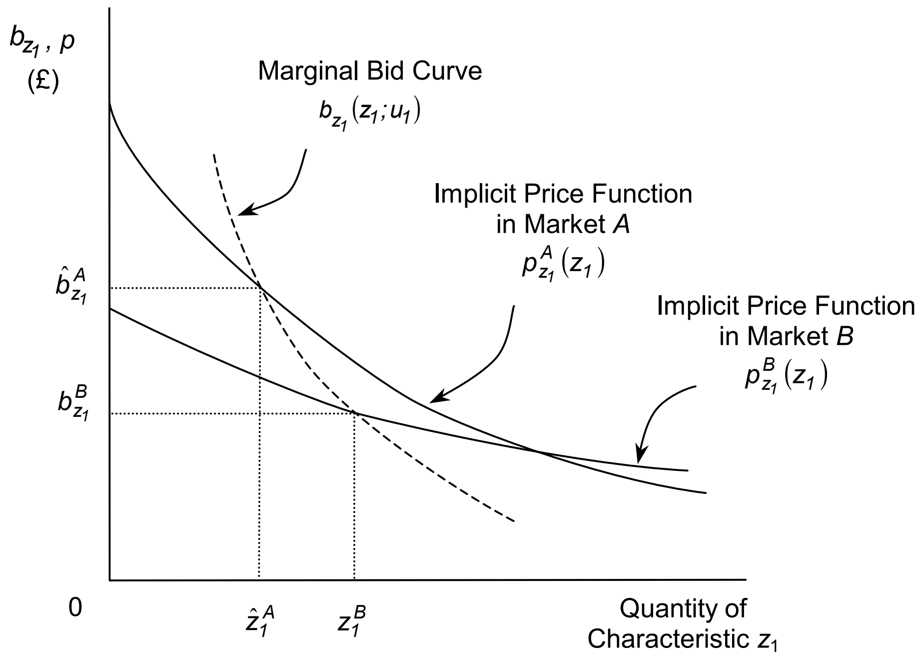

Unfortunately, this is not the case. Since the hedonic price function is different in the second market, the second household's optimal choice of residential location may not afford the same level of utility. For example, if prices are generally lower, then the household's maximised level of utility 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 $p^2_{z_i}$, 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 households[4] 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 $b_{z_1}(z_1;u_1)$, $\left(\hat{z}_1^A, p_{z_1}^A\right)$.

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 $\left(\hat{z}_1^B, p_{z_1}^B\right)$.

Notice, however, that $\left(\hat{z}_1^B, p_{z_1}^B\right)$ 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

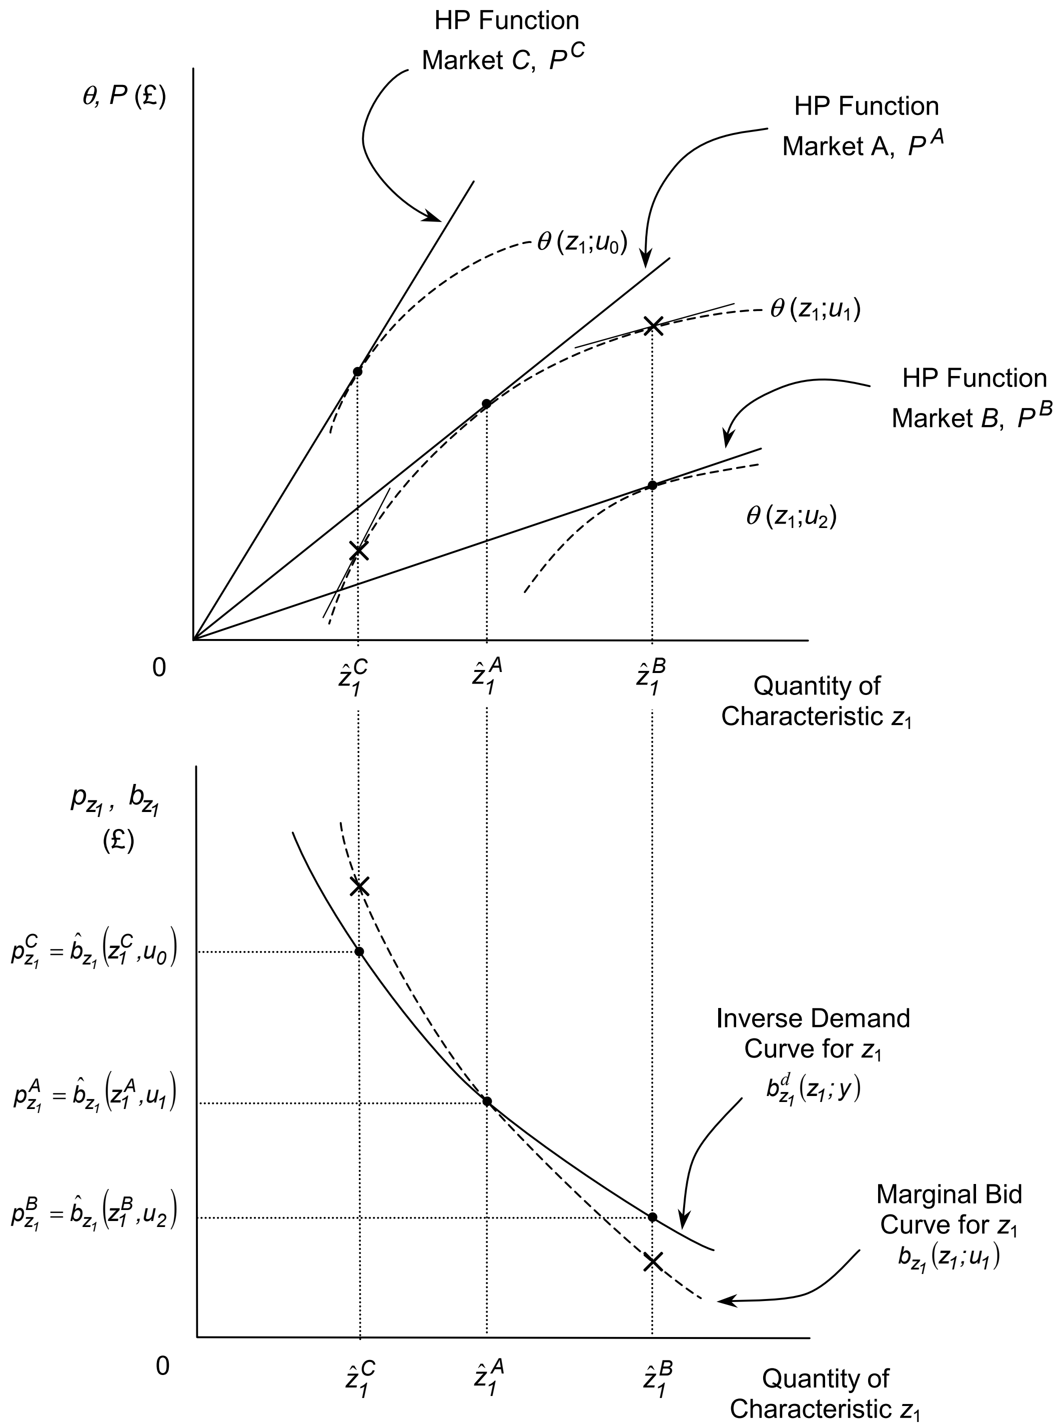

Finally, observe the choice made by the household in market *C*. Here the implicit price of z_1 is the constant $p_{z_1}^C$. Since this is higher than that observed in either of the other markets, the household in market *C* must make do with a lower level of utility. Indeed, the utility maximising choice 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 $\left(\hat{z}_1^C, p_{z_1}^C\right)$. 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;

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

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;

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

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 estimated 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; $\left(\hat{z}_1^A, p_{z_1}^A\right)$. 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)

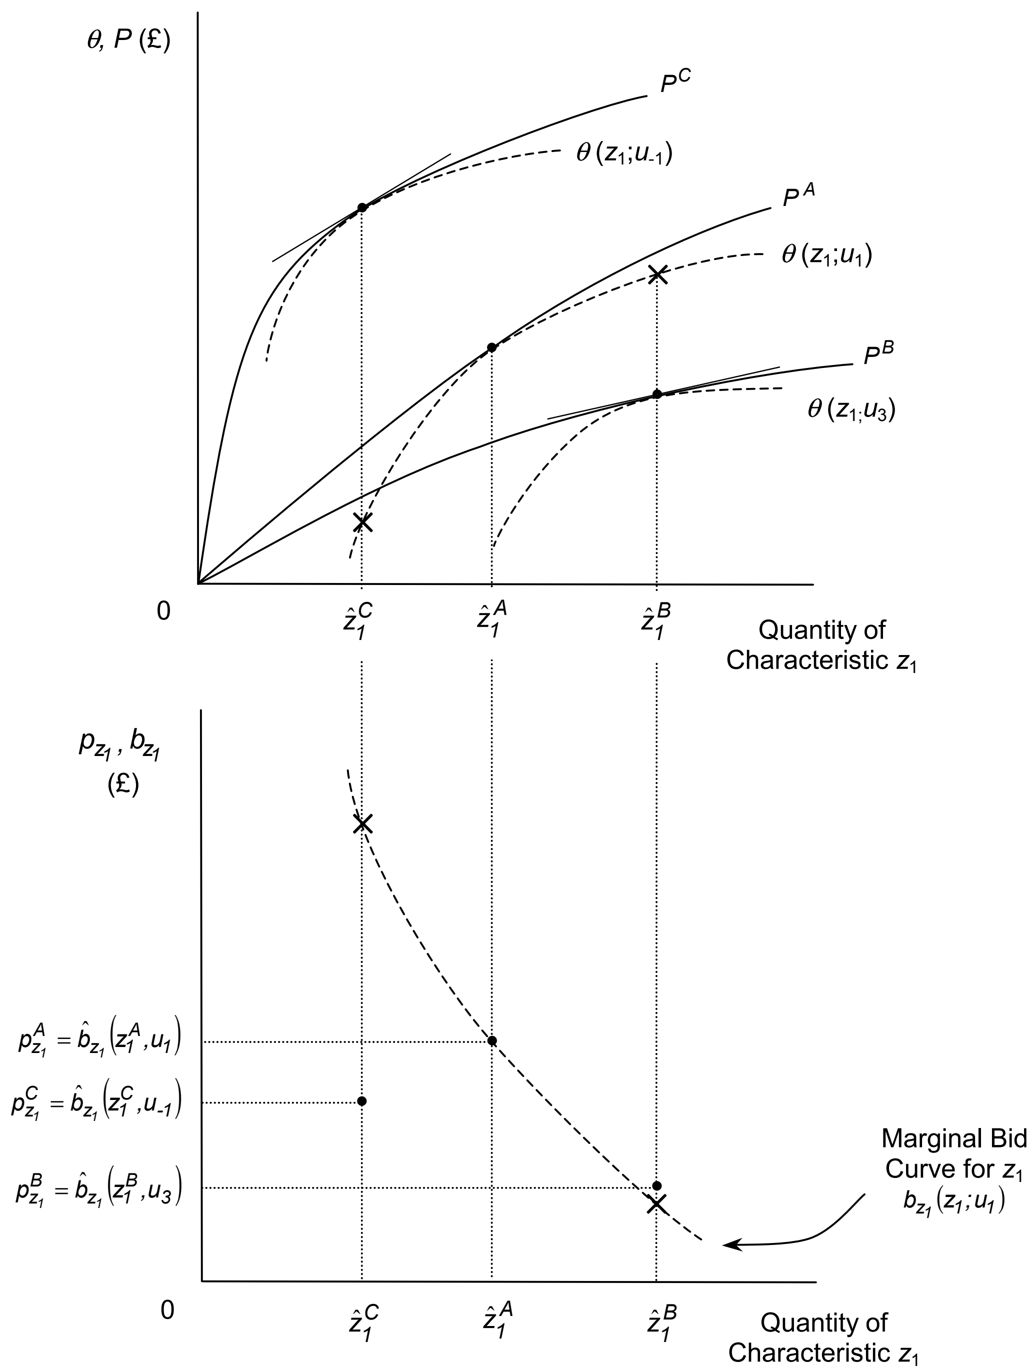

Clearly, with non-linear hedonic price functions and preferences that are not quasilinear, observed choices do not plot out a nice downward sloping inverse ordinary demand curve[8].

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 curve[9]. 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 enjoyed[10].

---

[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)

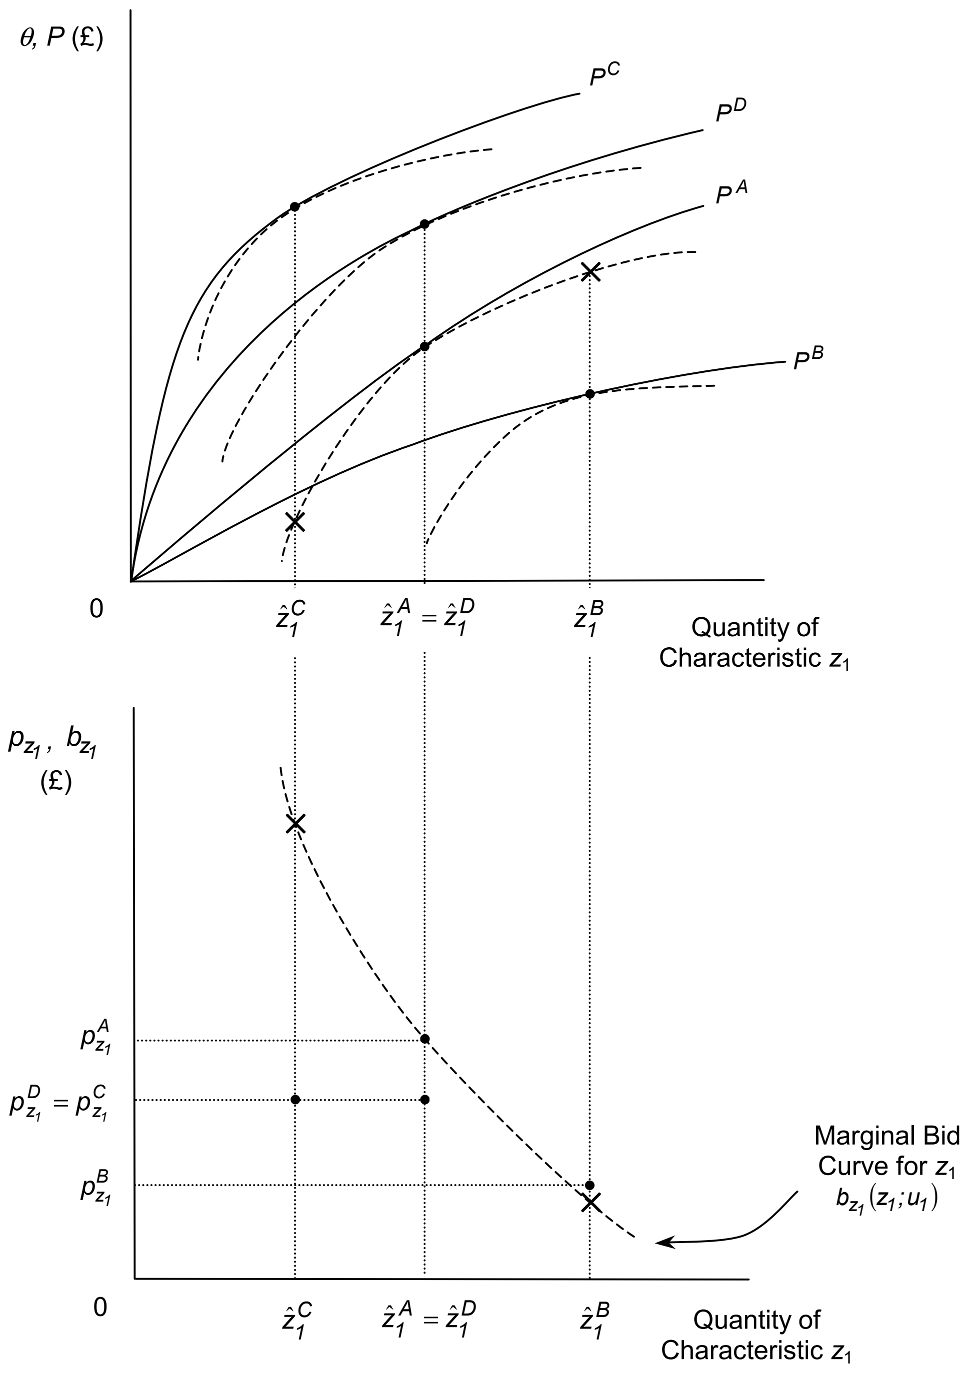

f. Mythical Demand Curves: Linearising the Budget Constraint

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 diagram[11].

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 ya 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_1. For example, if z_1 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

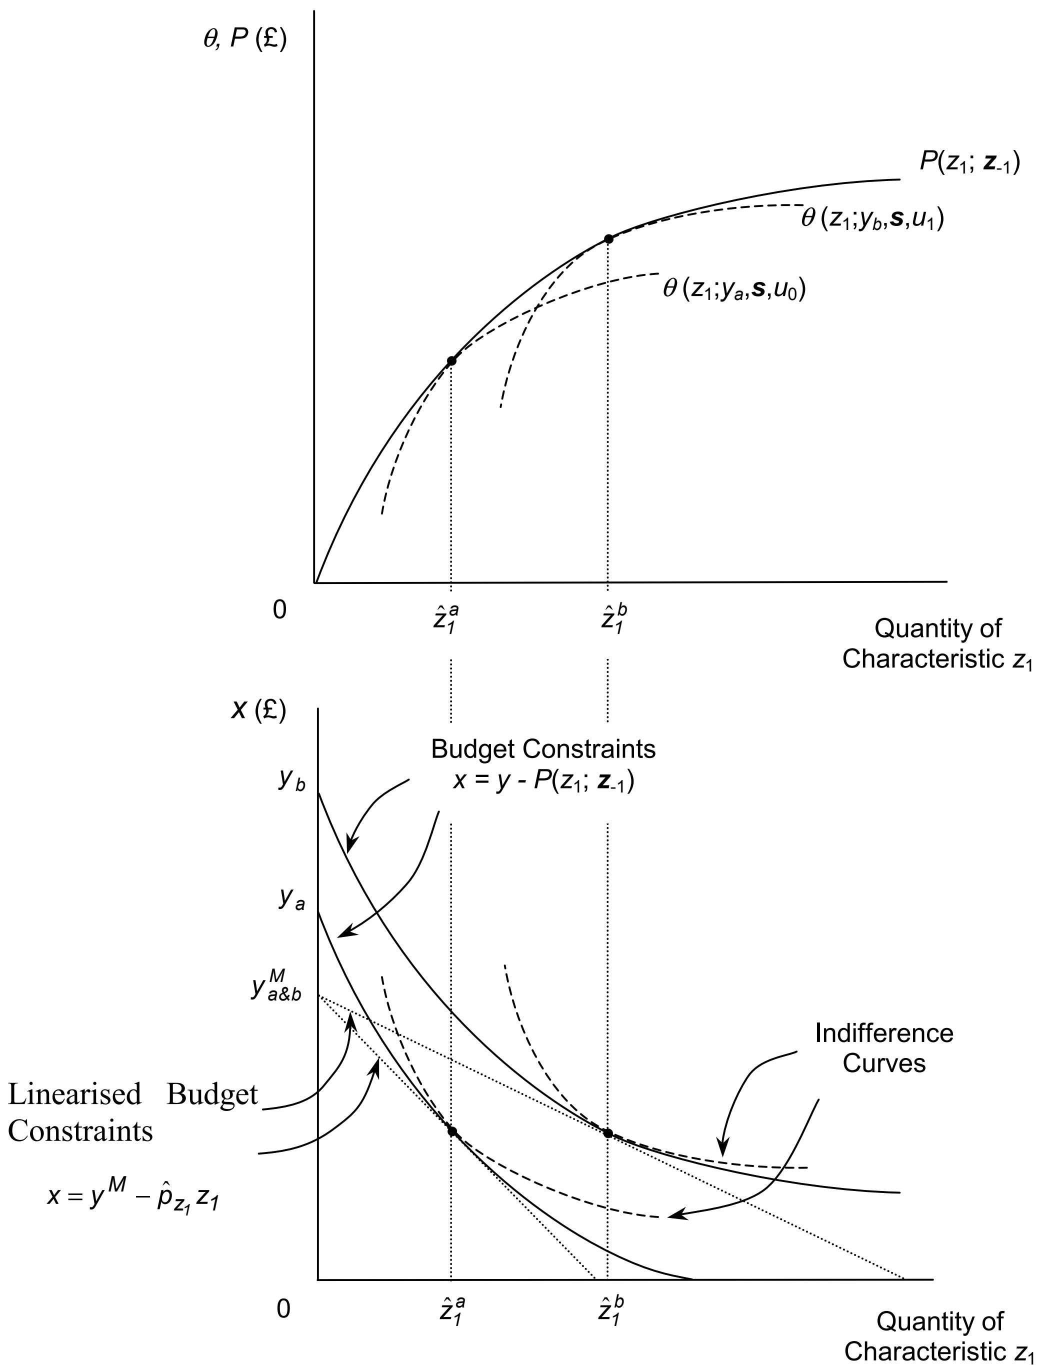

Now consider the choice made by household *b*. Following the same procedure, we can construct a mythical linear budget constraint whose slope is defined by the implicit price of the attribute at household *b*'s optimal choice, $\hat{p}_{z_1}^{b}$. The intercept of this budget constraint with the y-axis 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 function[12],
- Calculate the implicit price for each housing attribute
- Calculate the implied mythical income at these implicit prices according to;

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

- 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;

$$\hat{p}_{z_1} = b_{z_1}^M\left(\hat{z}_1, \hat{\boldsymbol{z}}_{-1}, y^M, \boldsymbol{s}\right) \tag{39}$$

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.

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

where $\hat{\boldsymbol{p}}_{z_{-1}}$ 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

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

Step 4	*Calculate each Household's Mythical Income* Using the implicit prices estimated in step 3 calculate each household's mythical income according to; $$y^M = y - P(\hat{\mathbf{z}}) + \sum_{i=1}^{K} \hat{p}_i \hat{z}_i$$
Step 5	*Calculate instruments for Implicit Prices and Mythical Income* Select instruments for implicit prices. Suitable candidates include; • Socioeconomic characteristics • Implicit prices chosen by similar (demographic traits/spatial proximity) households Select instruments for Mythical Income. Suitable candidates include; • Socioeconomic characteristics Using data from all markets estimate two ancillary equations regressing observed implicit prices and mythical income on instruments Use the regression results to calculate predicted values for implicit prices and mythical income for each household. Call these; $\tilde{y}^M$ and $\tilde{p}_{zi}$
Step 6	*Estimate Mythical Ordinary Demand Function* Using predicted values calculated in step 5 estimate the demand function according to; $$\hat{z}_1 = z_{z_1}^M \left(\tilde{p}_{z_1}, \tilde{\mathbf{p}}_{z_{-1}}, \tilde{y}^M, \mathbf{s} \right)$$
Step 7	*Calculate QCS welfare measures* Integrate under the mythical demand curve between the initial level of the attribute and that following some external change

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 changes[13]. 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