An overview of valuation techniques for ecosystem accounting

Prof Brett Day

University of East Anglia

Introduction

The purpose of this paper is to provide an overview of the key challenges that will arise in attempts to measure the benefits derived from non-market environmental goods for incorporation in an extended set of national accounts.

The defining feature of such goods is that, as a result of their public goods characteristics, they are not traded in their own markets and do not command an independent price. That is not to say they do not generate benefits for producers and consumers in the economy, far from it. Indeed, for several decades, environmental economists have been developing techniques of non-market valuation for the purposes of quantifying those benefits in terms of welfare changes or, equivalently, changes in surplus. The techniques for valuing the benefits of non-market goods in production are well-established and the application of those methods to generate values compatible with the conventions of national accounting provide no major conceptual challenges (see, for example, Nordhaus, 2006; and SEEA, 2012). Accordingly, this paper focuses on evaluating the benefits of environmental goods as realised by consumers.

A central issue in the application of the techniques of non-market valuation developed by environmental economists in this pursuit, is that those techniques estimate measures of welfare (surplus). In contrast, a central tenet of national accounting is that economic activity should be evaluated at market prices. Indeed, “when market prices are not observable, valuation according to market price equivalents should be used to provide an approximation to market prices.” (SNA, 2008, p. 51). A key purpose of this paper is to explore the extent to which techniques of non-market valuation might be adapted to fit with that requirement.

What to value?

Before commencing on that task we must first establish exactly what we mean by environmental goods. While there has been considerable confusion in this regard, the papers by Fisher and Turner (2008) and Boyd and Krupnick (2009) have firmly established that those environmental goods have to have the features of final goods. In the words of Boyd and Krupnick (2009) “they are commodities that require little or no subsequent biophysical translation in order to make their relevance to utility clear. [For example,] dissolved oxygen is not an endpoint because its role in utility requires understanding of its role in production of subsequent biophysical commodities.”

When considered in these terms the list of utility-relevant outcomes that require valuation actually shrinks quite considerably. For example, many very complex environmental processes are involved in regulating water flow through the hydrological system, but all that really matters to consumers is the degree to which those processes deliver on the final environmental good of reductions in the risk of damage from flooding. Likewise many complex environmental pathways contribute to the quality of the water we drink and the air we breathe, but for consumers the final environmental good is simply reductions in the risk of ill-health or death. Likewise, while habitats and landscapes are incredibly complex, from the consumer's point of view all that matters are that subset of environmental features that meaningfully and impact on the quality of recreational experiences or the amenity values drawn from landscapes. Somewhat more difficult and varied are the non-use values that are derived from the environment. In principle such values could be held for any number of different environmental entities and processes.

Exchange, Markets and Prices

In the System of National Accounts (2008) market prices are defined as "amounts of money that willing buyers pay to acquire something from willing sellers. The exchanges should be made between independent parties on the basis of commercial considerations only, sometimes called 'at arm's length'". Without wishing to bore knowledgeable readers, it is worth beginning by briefly reviewing the basic theory of exchange as envisaged by neoclassical economists. That theory begins by positing economic agents with wants and needs for goods and services (from now on just goods). The strength of those wants and needs can be expressed in terms of a willingness to pay (WTP): that is, the maximum amount of money that an economic agent would willingly give up to acquire a unit of some good.

Often the good in question is owned by, or produced by, another economic agent (though, as we shall see, in the case of many goods emanating from the environment, that is not the case). The terms under which that second agent might participate in exchange depend on the compensation they are being offered in order to give up that good. Again, the required compensation can be measured in money terms by their willingness to accept (WTA): that is, the minimum amount of money that they would accept for giving up a unit of the good. In situations where that economic agent is producing the good then WTA should equate to the marginal costs of production.

Figure 1, illustrates the WTP and WTA of two agents for some good. In the example, the WTP of one economic agent exceeds the WTA of the other such that the possibility exists for those two agents to affect a mutually advantageous exchange in which the good is transferred from the latter to the former in return for a money payment, p. The size of that payment must be such that $WTA < p < WTP$, though importantly observe that in this case exchange could proceed at any price in that range.

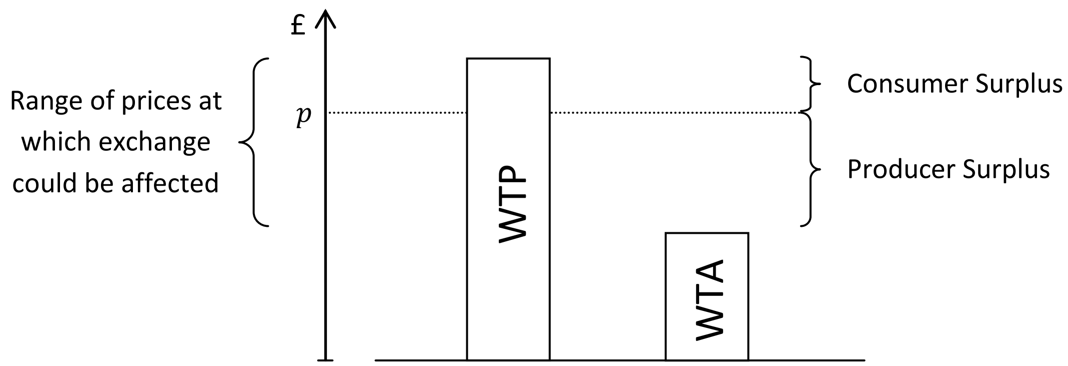

Figure 1: Basic elements of exchange – WTP and WTA

A central concern in economics, particularly welfare economics, is to evaluate the benefits realised by these two agents from participating in the exchange. In essence (and ignoring some complexities concerning ‘exact’ measures), those benefits can be measured by the buyer’s *consumer surplus* (that is, the difference between their WTP and the price) and the seller’s *producer surplus* (that is, the difference between the price and their WTA).

Of course, within a real economy there will potentially be very many buyers and sellers. One way of summarising the needs and wants of the buyers is through a *demand curve* which (again ignoring some complexities) might be thought of as the graph of WTP amounts of buyers in the market ordered from highest to lowest from left to right (see left panel of Figure 2). Likewise the compensations required by sellers can be described by a *supply curve*: the graph of WTA amounts this time ordered from lowest to highest from left to right (see middle panel of Figure 2). As shown in the third panel of Figure 2, when placed on the same graph, the intersection of demand and supply curves reveals the quantity of goods that could potentially change hands through a process of mutually advantageous exchange, $\hat{q}$: the sellers of each of those $\hat{q}$ units could be paired with a buyer whose WTP exceeds their WTA. Indeed, the quantity of surplus that might be realised by agents in the economy through the exchange of this good is the sum of differences between those WTP and WTA amounts, a quantity that on the diagram is shown as the area between the demand and supply curves up to $\hat{q}$. While highly stylised, the supply-demand diagram encapsulates the underpinning economic forces which drive non-coercive exchange in the economy.

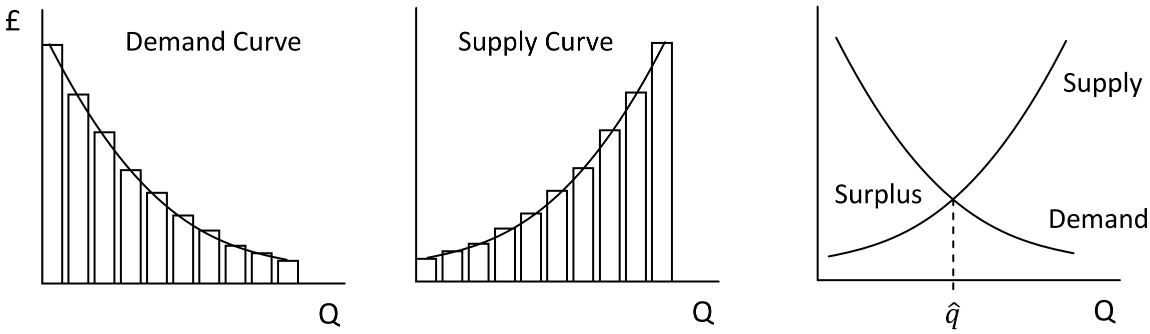

Figure 2: Demand and Supply Curves

Now with those (simplified) basics in place, let us consider how exchange might progress in a real economy. When the economy consists of very many buyers and sellers with perfect information and where none of those buyers or sellers is a sufficiently ‘big player’ to independently manipulate the terms of exchange, then a *perfectly competitive market* may evolve as the institutional setting within which exchange is affected. Perhaps the defining feature of such an institution is that exchange takes place at one particular price, determined by the intersection of the demand and supply curves.

The *law of one price* arises from competitive pressures in the market: attempting to sell above the market price is impossible since others are prepared to sell at the market price and selling below the price is irrational since it will be possible to enjoy greater surplus by selling at the market price. The perfectly competitive market institution is illustrated in the left hand panel of Figure 3. Observe that a perfectly competitive market maximises the surplus generated by exchange and affects a particular division of the gains from exchange between buyers and sellers as shown by the quantities labelled CS (Consumer Surplus) and PS (Producer Surplus). In making those exchanges a certain amount of money, $p\hat{q}$, is passed from buyers to sellers. In providing a monetary measure of this economic activity in the National Accounts, it is this quantity $p\hat{q}$, that is recorded.

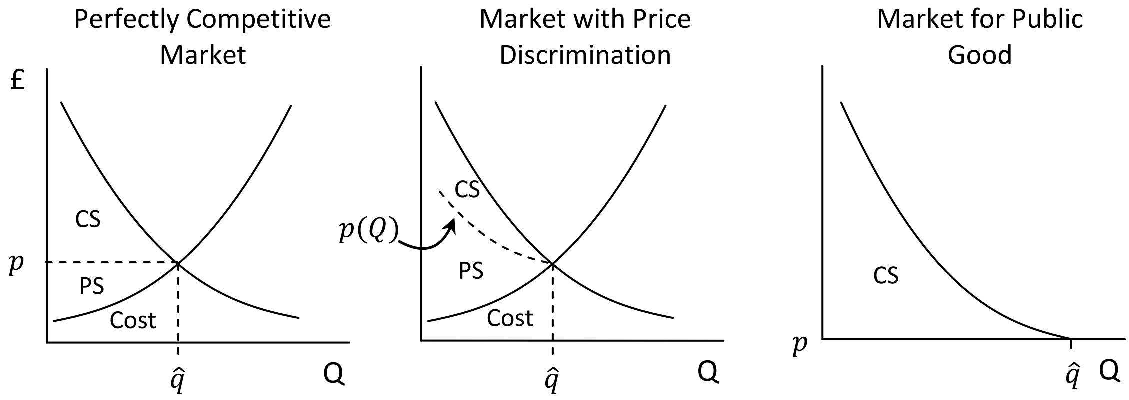

Figure 3: Prices and Market Allocation

Of course, there is no reason to assume that all exchanges are conducted within perfectly competitive markets: very many other institutional settings might arise in the context of a real economy. For example, with incomplete information or through some form of market power sellers (or a seller) may practice price discrimination. To fix thoughts, imagine goods sold through an auction in which non-colluding buyers pay what they bid. The possible outcome of such an exchange mechanism is illustrated in the second panel of Figure 3.

In this case, the same quantity of goods is exchanged and the quantity of surplus generated by those exchanges remains the same as in the perfectly competitive market. Importantly, however, the quantity of money changing hands is greater ($\int_0^{\hat{q}} p(Q)dQ$) and the division of surplus also changes in favour of the seller. It follows, somewhat perversely, that while the quantity of goods being exchanged and the total welfare generated by those exchanges remains unchanged, the method of quantifying this economic activity in the SNA would record a greater value $\int_0^{\hat{q}} p(Q)dQ > p\hat{q}$ than would arise if the exchange had taken place under conditions of perfect competition.

The key message from these examples is that prices are not some unchanging feature of market exchange, but reflect as much on the institutional setting through which exchange takes place as on the underlying structure of preferences. Indeed, as with the example in Figure 1 any price function bounded from below by the supply curve, from above by the demand curve and passing through the intersection of those two curves, is compatible with a level of $\hat{q}$ units of exchange and the maximisation of surplus. It follows that summing the prices at which units are exchanged does not provide an invariant measure of economic activity. Using that measure, the same level of economic activity (here the exchange of $\hat{q}$ units of a good) can record different values according to the price function that arises in some particular market institution.

While the SNA's measure of the value of economic activity is not invariant to market institution and provides, at best, only an approximate measure of welfare it is worth noting the aspects of economic activity that are accurately recorded by those accounting methods;

- First, in measuring $\hat{q}$ for all market goods, the accounts record the quantities of goods that are exchanged in market transactions within an accounting period.

- Second, in summing $p\hat{q}$ (or, more generally, quantities multiplied prices) across all market goods, the accounts record the quantity of money circulating through the economy in the form of market exchanges over the accounting period.

In considering how to extend national accounting methods to include consumption of non-market goods, it would be desirable to maintain those two clear interpretations of the figures reported in the national accounts. As we shall see, that proves somewhat problematic.

Environmental Goods and 'Exchange Prices'

One particular market institution with which we are chiefly concerned in this paper arises in the context of environmental goods and services that, on account of issues of non-excludability, have the characteristics of public goods. Those goods and services are enjoyed by economic agents and provide real welfare benefits that as per our previous examples can be depicted as a demand curve (see final panel of Figure 3). In this case, however, since the goods and services are unowned they are not traded independently in their own market.

Notice that as a result of the open-access nature of the environmental good in the final panel of Figure 3, supply is unconstrained and consumers make use of as many units of the good as satisfies their wants and needs: in this case, $\hat{q}$ units[1]. We often describe such environmental goods as being 'non-market' since no formal exchange takes place and individuals come to benefit from the provision of the environmental good without having to part with any money. An alternative and exactly equivalent characterisation would be to consider this as a market exchange but one in which the 'exchange price' is zero. That zero price is indicated on Figure 3.

The central question of this paper then is to consider how the very real benefits that come from the consumption of these environmental goods should be reflected in a modified set of national accounts.

There are at least three ways to progress:

- *Price at Zero:* A strict application of the national accounting convention would be to value consumption of these environmental goods at their implicit marginal price of zero. Following that convention has some advantages. In particular, the accounts would continue to accurately report the quantities of goods and money 'exchanged' in the economy over the accounting period. More problematically, such a practice would see the very real benefits that come from consumption of this environmental good being recorded in the national accounts as zero: a very poor approximation of the welfare generated from environmental goods.
- *Representative Marginal Price*: An alternative would be to mimic a perfectly competitive market and attempt to choose some representative marginal price, $\tilde{p}$, presumably by selecting a value from the range of marginal WTP values that form the environmental good's demand curve. That value could, for example, be the average of the marginal WTP amount between 0

[1] In reality, things may be somewhat more complex. The quantity or quality of the environmental good is often constrained by nature such that not all wants and needs can be satisfied. Moreover, since the environmental good is open-access and not allocated through the rigour of a market, it is unlikely (except in the case of a pure public good) to be the case that the available units of that good are allocated to those with highest WTP.

and $\hat{q}$, though the choice is essentially arbitrary. Multiplying $\tilde{p}$ by $\hat{q}$ would give a value, calculated in the conventional manner as the product of a price and a quantity, to be entered in the national accounts: those accounts would then include an element that (very approximately) captured the benefits derived from non-market environmental goods. Unfortunately, that calculation introduces a logical inconsistency. If the marginal price really were $\tilde{p}$ then consumers would demand some quantity of the environmental good smaller than $\hat{q}$. If a functioning market for the environmental good could be established that resulted in a price $\tilde{p}$, the quantity of money $\tilde{p}\hat{q}$ would not change hands in a process of exchange for this environmental good. Moreover, introducing the quantity $\tilde{p}\hat{q}$ into the national accounts would also mean that those accounts would no longer precisely record the quantity of money circulating through the economy as a result of exchange processes over the accounting period. Of course, there are precedents for 'inventing' money to record economic activity in the SNA for example in the imputed values placed on owner-occupied housing.

- *Representative Discriminatory Prices*: Alternatively, one could choose any set of discriminatory prices that form a function over the range 0 through to $\hat{q}$ that remains below the demand curve and passes through $\hat{q}$ on the horizontal axis. As with the selection of a constant marginal price, valuing the environmental goods consumed in the economy using a set of discriminatory prices would be 'inventing' money. On the other hand, such a set of prices would maintain logical consistency in so much as $\hat{q}$ units of the environmental good would be exchanged if such a set of prices existed in a market. In addition, one might justify such a procedure on the basis of the possibility that market institutions could arise in the real world that returned similar forms of discriminatory prices perfectly compatible the SNA's definition of a market price. Which particular price function to choose is essentially arbitrary though one possibility would simply be to use a marginal price function that perfectly followed the demand curve. In that case, the measure of values derived from this environmental good would amount to the welfare measure given by consumer surplus.[2]

The central message of this discussion is that the market prices used for the purposes of valuing economic activity in the national accounts do not provide an accurate measure of the welfare generated by that activity. Indeed, it is possible for the exact same set of market transactions to be recorded as delivering different values due to differences in the market institution and the nature of the prices that evolve in that market. Given those insights, there is nothing logically inconsistent with the conventions for pricing used in the national accounts to value the benefits derived from public goods by the surplus derived from their consumption.

Methods of Non-Market Valuation and 'Exchange Prices'

Our primary concern in this paper is with an environmental good that is not traded in its own independent market: that is to say, a consumer cannot simply purchase units of the environmental good in a market at some price. Instead the environmental good is provided in some fixed quantity or quality though, as we shall see that quantity or quality may vary across space or activity.

[2] A concern here, however, is that for certain forms of environmental good (say, drinking water) the demand curve may tend towards infinity as quantity or quality approaches zero. While that may be a good reflection on the welfare value of an environmental good, it would make a nonsense of a system of national accounts primarily designed to value levels of economic activity.

In order to think about the possibilities for valuing the consumption of environmental goods, it is essential to be more explicit about how the environmental good enters a consumer's choice problem. Economists recognise several ways in which this might happen, and each of those different forms of choice problem suggest a different approach to the task of non-market valuation.

Here we briefly review the different forms before exploring methods for valuing the different cases in more detail;

- *Attribute of Heterogeneous Market Good*: Quantities of an environmental good form an attribute of some other good that can be purchased in a market. The standard example is property and the environmental quality of its location.
- *Attribute of a Waged Job*: While researchers have more usually studied how wages reflect occupational risks of illness, accidents or death, it is certainly possible that wages may also capture aspects of the environmental quality at the location of employment.
- *Substitute Good in Household Production Function*: In a number of circumstances, market goods can act as substitutes for environmental goods such that expenditure on those substitutes informs on the value households place on the environmental good.
- *Complementary Good in Household Production Function*: In the same way, environmental quality may act as a complement to marketed goods. The environmental good and marketed goods must be combined in order to derive utility such that expenditure on the market good informs on the value place on the environmental good.
- *Pure Public Good*: In some circumstances, consumers derive value directly from environmental goods as a final good or service and cannot influence the level of provision of that environmental good through the choices they make in markets. Such values are derived in complete isolation from marketed goods and, as such, their value to consumers cannot be inferred from market behaviour.

Attribute of Heterogeneous Market Good: The Hedonic Price Method

While it may not be possible to independently buy units of an environmental good, under certain circumstances quantities of that good may form an attribute of some other good that can be purchased in a market.

The standard example is property, in which the location of a house determines levels of local environmental quality (for example, levels of noise pollution, air pollution, views of and proximity to natural areas). Indeed, property is an example of a *heterogeneous good*, where the set of units that are traded in a market differ in terms of the levels of a number of attributes. In the case of property, that list of attributes would include not only environmental quality, but also the physical characteristics and quality of the building and the proximity of the property to local amenities. Other examples of heterogeneous goods include cars, computers and breakfast cereals.

Economists' theoretical understanding of the functioning of markets for heterogeneous goods was first outlined by Rosen (1974). Without going into detail (see Day, 2001), Rosen explored the outcome of a competitive market characterised by a supply of units of the heterogeneous good endowed with different levels of attributes and a set of consumers with some distribution of preferences for those attributes. Rosen showed how that market would be bought into equilibrium through the emergence of a *hedonic price function* describing the price at which a unit of the heterogeneous good with particular attributes will sell for in the market.

Since households prefer better environmental quality to worse, this hedonic price function will tend to be increasing in environmental quality. In other words, the marginal price of environmental quality, or what in the hedonics literature is called its *implicit price*, is positive. Importantly, however, since the heterogeneous good cannot be unbundled and its constituent attributes sold independently, the implicit price of environmental quality (or any other attribute for that matter) will not necessarily be constant. How much more must be paid for one extra unit of environmental quality will depend in a complicated way on the availability of goods with more quality.[3]

A simplified presentation of the household's choice problem in such circumstances is as follows;

$$Utility\ Function{:}\quad u\left(x_1, \{q, z\}_{x_2}\right)$$

$$Budget\ Constraint{:}\quad M \leq x_1 p_1 + p_2(q, z)$$

Here a household must choose the quantities of a standard market good, x_1, and the attribute levels of a heterogeneous good, x_2, where those attributes consist of levels of environmental quality, q, and, for the sake of simplicity, one other attribute, z. Each unit of the standard market good can be purchased at its unit price of p_1 while the cost of purchasing different varieties of the heterogeneous good are determined by hedonic price function $p_2(q, z)$.

Basic economic theory implies that optimising consumers will choose levels of x_1, q and z that equate their marginal WTP (MWTP) for those goods and attributes with the marginal price of each. For the standard market good that implies $MWTP_{x_1} = p_1$. For environmental quality, on the other hand, that implies $MWTP_q = \partial p_2(q, z)/\partial q$, where $\partial p_2(q, z)/\partial q$ is the implicit price of environmental quality, and since implicit prices are not constant, consumers of the heterogeneous good choosing different levels of environmental quality also face a different marginal price. That outcome is shown in the left hand panel of Figure 3 where, faced by a non-constant implicit price function household a selects a marginal price of $\hat{p}_a$ and household b a marginal price of $\hat{p}_b$.

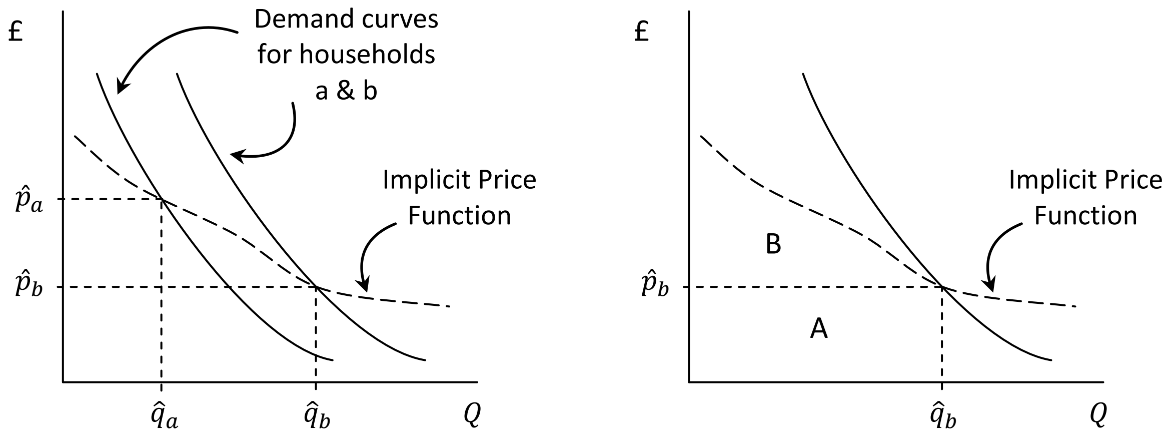

Figure 3: Marginal Price and Quantity Choices for a Heterogeneous Good

[3] For example, imagine that the environmental quality attribute could take one of three equally spaced levels: low, medium and high. Imagine further that the market contains lots of units in the low and medium category but only one in the high category. Under those circumstances, the extra paid for a unit with the medium level over the low level will tend to be much smaller than that paid for the high level over the medium.

While the environmental economics literature has focused on the very complex problem of how to recover estimates of the underlying preference function from information on implicit prices and households' market choices, the great attraction of the hedonic price mechanism for the purposes of the national accounts is that it returns market-based marginal prices of environmental goods akin to those used for standard market goods.

Of course, the range of environmental goods for which implicit prices can be derived from hedonic property price methods is constrained to those that exhibit variation over the spatial extent of a property market. Accordingly, applications to differences in noise pollution, or access to urban green space or views of natural areas may well be successful. The hedonic price method will not be appropriate for environmental goods that do not exhibit such spatial variation: for example, the existence value placed on some remote habitat.

As illustrated in the left hand panel of Figure 3, a further complexity is that within any particular market the hedonic method identifies a continuum of implicit prices. A strict application of the marginal price convention in national accounting would require this implicit price to be calculated for each household.

Even more troublesome is that, with a non-linear implicit price function, the straight multiplication of a household's marginal price by the quantity of the environmental good they have acquired through purchase of the heterogeneous good , does not accurately reflect their actual expenditure on the environmental good. Observe the right hand panel of Figure 3. Here the conventional measure of expenditure using just the marginal price: that is to say $\hat{p}_b\hat{q}_b$, is shown by area B. In this case, that measure underestimates the household's actual implicit expenditure on the environmental good given by the area A + B: that is to say $\int_0^{\hat{q}_b} \partial p_2(q,z)/\partial q \, dq$ (where the zero level of the environmental good is taken to be the lowest level observed in the market).

A further complication arises from the fact that the particular continuum of implicit prices calculated from the application of the hedonic price method in one property market is specific to the conditions of demand and supply in that market. It would not be theoretically appropriate to transfer those implicit prices to another property market characterised by different conditions of supply and demand.

Finally, it should be noted that the measures of expenditure on environmental goods that can be derived from hedonic pricing methods are expenditures already captured in the national accounts through the recording of property transactions. What the method offers then is a means of partitioning those transactions so as to identify that part of the expenditure that can be attributed to environmental goods.

Attribute of Waged Job: The Hedonic Wage Method

A qualitatively different situation in which an environmental good may constitute an attribute in a consumer's choice problem occurs in labour markets. In particular, consumers may select employment from an array of different jobs where those jobs differ according to a variety of attributes. Of particular interest are attributes of environmental quality and of safety in the workplace.

As in the hedonic pricing method, economic theory suggests that the interaction of firms, supplying jobs with different attributes, and consumers, with different preferences for those job attributes, will lead to the establishment of a *hedonic wage function* that clears the labour market. Since consumers place positive value on environmental quality and workplace safety, that wage function should be decreasing in those attributes. Though, because job attributes cannot be unbundled, there is once again no reason to assume that the hedonic wage function will be linear in these attributes.

A simplified presentation of the consumer's choice problem in such circumstances is as follows;

$$Utility\ Function: \quad u(x_1, q, z)$$

$$Budget\ Constraint: \quad \bar{M} + w(q, z) \leq x_1 p_1$$

Once again the consumer can choose quantities of a standard market good, x_1, with constant marginal price, p_1. In addition, the consumer can select their levels of exposure to environmental quality, q, and, for the sake of simplicity, one other job attribute, z, through their selection of employment. The remuneration the consumer enjoys from choosing a job with particular levels of chasing different varieties of q and z is determined by the hedonic wage function $w(q, z)$.

The consumers optimal choice of q and z will once again be such that the their marginal value for environmental quality will be equal to the marginal trade off between wages and environmental quality in the work place; $\partial w(q, z)/\partial q$. Once again, $\partial w(q, z)/\partial q$ can be interpreted as an estimate of the marginal price of environmental quality; it records how much money a consumer would have to forfeit in lost wages in order to receive one more unit of environmental quality.

The use of marginal prices derived from the hedonic wage model to estimate expenditures on environmental goods raises the same set of issues as discussed previously for the hedonic price model.

As discussed by Bockstael and McConnell (2007), a number of other empirical and theoretical issues are of concern in transferring marginal prices form hedonic wage models to value changes in risks of accidents or death resulting from environmental change. Most important, perhaps, is the fact that the major beneficiaries of such changes are often thought to be children and the elderly while the preferences recorded in marginal prices from hedonic wage models tend to be those of males in their 30s and 40s.

Substitute Good in Household Production Function: The Defensive Expenditure Method

In a number of circumstances, environmental goods are best thought of as inputs that households use in providing themselves with some final service from which they derive utility. The standard example concerns the final service of health. For many types of environmental quality, such as those relating to air and water pollution, it is not the pollution itself that concerns consumers but how that pollution impacts on their health.

While the level of environmental quality is out of their control, consumers can purchase other goods and services that act as substitutes for environmental quality in the production of health end points. For example, items including air filters, sun screen and bottled water have been posited as marketed substitutes for environmental quality in producing health.

A simplified presentation of the consumer's choice problem is given by;

$$Utility\ Function: \quad u\big(x_1, f(x_3, q)\big)$$

$$Budget\ Constraint: \quad M \leq x_1 p_1 + x_3 p_3$$

In addition to quantities of the standard market good, x_1, the consumer can produce some valued output through either exogenously determined environmental quality or through the purchase of the market good, x_3 , at constant marginal prices p_3. The production technology is given by the function, $f(x_3, q)$.

Bartik (1988), has shown that when the market good and environmental good are substitutes a lower bound to the value consumers put on some change in q can be estimated from observations of the changes in expenditure on x_3 following that change. Intuitively, when environmental quality falls, consumers will respond through making defensive expenditures on the substitute market good. The payments they make in that offsetting will never exceed the value to them of returning q to its original level.

Observe that an upper bound on the value of a change in q can be estimated from an understanding of the household production technology, $f(x_3, q)$. To illustrate, it must be the case that the value of a fall in q cannot exceed the extra expenditure on x_3 that is needed to return production back to its original level. In general economists have been sceptical of this measure, since it is quite possible for the costs of the required defensive expenditures to far outweigh the value of returning production to previous levels.

In the environmental economics literature, the focus has been on defining circumstances where restrictions on the production technology allow us to deduce more about the exact welfare change associated with change in q.

The lower bound estimate provided by observed defensive expenditures in reaction to a unit change in provision of an environmental good, provides a value that could reasonably be interpreted as approximating a marginal price. It is an amount that a household would be WTP for the marginal unit of environmental quality and, given the assumption of diminishing marginal product of environmental quality, an amount they would also be prepared to pay for each other unit of environmental quality currently enjoyed.

Notice that in this case, the environmental good has the characteristics of a pure public good: it is provided to consumers without any charge. Accordingly, using the defensive expenditures method to identify a marginal price at which those environmental goods might be valued, adds a new item to the national accounts. The values generated by the environmental good are not captured anywhere else in the accounts.[4]

[4] When an environmental good is a characteristic of a heterogeneous market good (e.g. property) individuals are already implicitly paying for that service flow (e.g. through their housing rent), likewise when environmental goods are complementary to market goods (e.g. in the production of recreational experiences) individuals are already 'paying' for the environmental good through their payments for the market goods (e.g. travel costs). The substitute case is a little different insomuch as individuals enjoy the environmental service flow without having to incur any expenditure on market goods. Of course, if the service flow is somehow damaged then we might observe defensive expenditures on the substitute good and from those we could deduce a lower bound to the value of the environmental good. In other words, to value the service flow currently being experience by individuals in the economy one could draw inference from defensive expenditures that were made elsewhere or at some other time when that environmental service flow was not being enjoyed. Accordingly, applying that figure to the current levels of service flow adds on an "expenditure"that is not actually incurred (it doesn't have to be because the environment is providing that service for free).

As with all non-market valuation techniques, application of the defensive expenditure method is beset by numerous theoretical and empirical difficulties. Perhaps the most problematic is the identification of market goods that constitute substitutes to an environmental good in the way assumed by the theory and the subsequent collection of data detailing changes in defensive expenditures in response to changes in environmental quality.

Likewise, adjustments may need to be made to marginal prices derived from defensive expenditures if the goods bought with those expenditures confer utility through some other route. For example, households may respond to traffic noise pollution by installing double-glazing. According to the assumptions of the defensive expenditure method, the costs of double-glazing should provide a lower bound estimate to the value of the reduced exposure to traffic noise. Of course, double-glazing contributes to household welfare in other ways: perhaps through their insulating properties or perhaps through reducing maintenance efforts.

The double-glazing example illustrates another complicating factor for the defensive expenditure method. In particular, those expenditures may be 'lumpy' and constitute long term investments. Since defensive expenditure is only a lower bound approximation to welfare, the more lumpy the market good the greater the likelihood of significant underestimation of the value of environmental quality[5]. Moreover, with long term investments the expenditure on the substitute good will need to be disaggregated using a suitable discount factor to identify a value that represents the benefits of environmental quality over the duration of one accounting period.

Complementary Good in Household Production Function: The Travel Cost and Associated Methods

An alternative household production relationship that may exist between environmental quality and marketed goods is one of complementarity.

Here the standard example is recreational experiences in natural areas. To enjoy such a recreational experience, consumers must combine the environmental quality of the natural area with a series of complementary market goods, most notably the costs of transport to that area. Since the quality of the natural area cannot be enjoyed without the market purchases, those purchases provide information on the value households place on environmental quality. In the context of valuing the contribution of environmental quality to recreational experiences, this approach is commonly termed the travel cost method.

[5] If the substitute good is "fine-grained", then expenditures on the substitute will closely resemble the benefits being foregone from the loss of environmental service flow. So, if for example tap water quality reduces (say from an outbreak of Giardia or some other nasty) then observing expenditures on bottled water as a substitute would give a pretty accurate reflection of the value of clean tap water. Alternatively, if the substitute good is lumpy then for many people we may not observe any defensive expenditure, even if those people have suffered considerable welfare loss as a result of a reduction in environmental service flow. For example, imagine the construction of a new road that significantly increases the noise exposure of properties along its length. The substitute good here might be double-glazing, but double-glazing is "lumpy" so for many people the cost of installing the new windows may outweigh the benefits that would come from substituting for the lost environmental service flow (i.e. a quiet environment). That means that in any study the expenditure on substitutes is always going to be a worse approximation to the value of the environmental service flow the lumpier the substitute.

The nature of the consumer's problem is identical to that described for defensive expenditures model except now environmental quality, q, and the market good, x_3, are complements in the production of recreational experiences.

Economists have identified methods for using data on changes in expenditures on x_3 as a result of changes in the environmental quality to estimate exact welfare measures for changes in q.

In the case of recreational use of natural areas, expenditures incurred in taking each trip to a recreational site can be interpreted as the price of accessing the site. The intuition is simple: to enjoy the recreational site I have to pay the travel costs of getting to that site. As such, like any other price, I will choose to take trips to that site up until the point at which the marginal benefit I enjoy from taking a trip to that site falls to the costs of a trip.

Accordingly, for the purposes of national accounting, one possibility is to use travel costs as marginal prices for the recreational benefits of trips to natural areas.

Notice that unlike the substitutes case, the use of this marginal price to value an environmental good does not add a new item to the accounts. Rather, in the complements case, enjoyment of the environmental good must be paid for through purchases of market goods. Accordingly, the outcome is simply to reclassify already identified transactions and ascribe them to the environmental goods.

One complexity recognised by practitioners of the travel cost method is that travelling to a recreational site uses time that a consumer could have employed undertaking other utility-raising activities. Accordingly, in nearly all applications, practitioners will add an element to the travel cost that represents the opportunity cost of time spent travelling. From a welfare point of view this makes perfect sense, though the case for including that time cost element in marginal prices used for national accounting is not so clear. In particular, the national accounts do not cost the time associated with other market activities such that their inclusion for the valuation of recreational sites would introduce inconsistencies.

One limitation of the travel cost method is that it does not extend easily to situations in which consumers are faced by an array of substitute recreational sites. In those circumstances, the consumer's are as concerned with the choice between sites as the choice of the number of trips to take to one particular site. The standard method applied in the case of multiple sites is provided by the random utility model: a discrete choice modelling technique in which consumers are assumed to choose which particular site to visit based on the qualities of and costs of travel to the different substitute sites. A fundamental feature of random utility models is that they predict the probabilities that a household will choose to visit each of the different available recreational sites, as well as the probability that they won't take a recreational trip at all. Accordingly, with such models an 'exchange price' might be arrived at by taking the travel costs to each recreational site and calculating the probability-weighted sum of expenditures on outdoor recreation across the course of a year.

As with many other environmental goods and services, an issue for accounting for recreation benefits is that location is everything. The value an individual derives from a recreational site (perhaps of some particular habitat type) depends crucially on where it is located and what other recreational opportunities are available. So for one community a local woodland may be extremely valuable and receive many of visits, but for another community an identical woodland may be of little value perhaps because there is an abundance of woodland in the region or other recreational opportunities that offer preferable alternatives. Accordingly, it is impossible to imagine any sensible accounting price based on recreational sites of types (e.g. of a particular habitat). More sensible would be to consider constructing a single aggregate account that cuts across habitats. That aggregate quantity could be decomposed into the aggregate value that was derived from recreational sites of each habitat type, but not into some uniform unit price for each site.

Pure Public Good: Stated Preference Methods

In some circumstances, consumers derive value directly from the environment as a final good or service and are unable to influence their level of provision of that environmental good through purchases in markets.

A standard example of such a case might be the pure existence value that a consumer derives from the on-going existence of a species, say the blue whale, or natural habitat, say the Amazon rainforest. Such values are derived in complete isolation from marketed goods and, as such, their value to consumers cannot be inferred from market behaviour.

A simplified presentation of the consumer's choice problem is given by;

$$Utility\ Function: \quad u(x_1, q)$$

$$Budget\ Constraint: \quad M \leq x_1 p_1$$

For these goods, practitioners of non-market valuation have no option but to adopt stated preference methods.

Since stated preference techniques estimate individuals' values through their responses to questions concerning purchases in hypothetical markets, they can potentially return values for any conceivable change in provision of an environmental good. The question is: for the purposes of providing a quantity akin to a marginal price, which particular change in provision should be valued?

Consider the aggregate demand curves for two different public goods illustrated in Figure 4. Since these are public goods the benefits of each unit of those goods are enjoyed simultaneously by everyone. In other words, the demand curves are generated by vertically summing individual demand curves across all individuals; they show the aggregate amount individuals in the economy would pay for each unit of the environmental good. For both goods, provision is constrained at a level of $\bar{q}$.

In the left hand panel, the constrained level of provision ensures that at the margin individuals hold positive values for the environmental good. In these circumstances, an argument could be made to support the contention that the natural counterpart to a market price is this marginal valuation, p. Under that premise, application of stated preference techniques to inform the national accounts should focus on the valuation of the marginal unit of environmental good.

An alternative outcome is illustrated for a second public good shown in the right hand panel of the figure. Here the marginal valuation of the environmental good is zero. From a welfare point of view applying that value as a marginal price in the national accounts makes no sense, particularly as the level of total welfare generated by the two environmental goods (the area under the respective demand curves) is of similar magnitude. On the other hand, one might argue that in the current state of affairs there is no scarcity in the second public good and, as such, a marginal price of zero is appropriate.

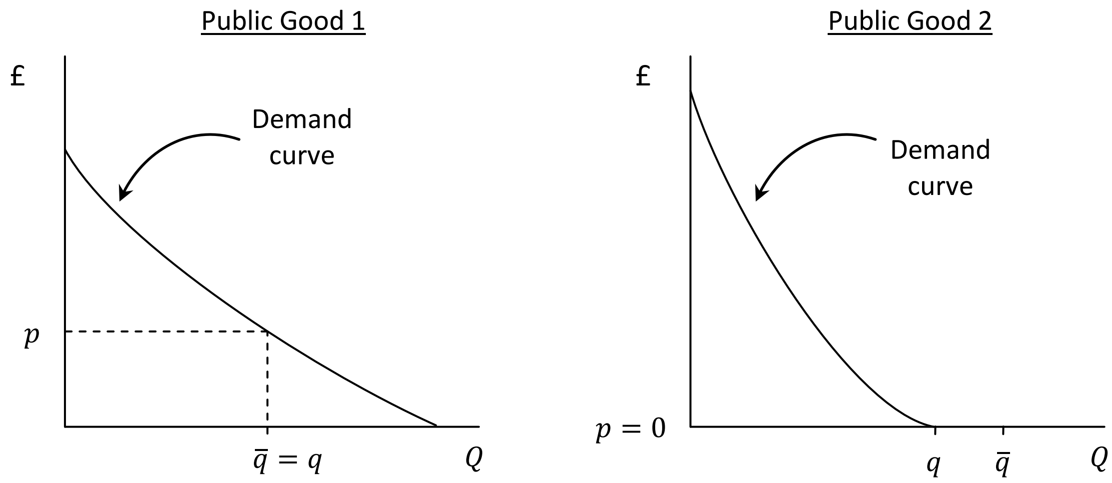

Figure 4: Marginal WTP for two different public goods with an exogenously determined level of provision $\bar{q}$

Of course, as we have seen any positive marginal price function that remains beneath the demand curve and passes through q is compatible with the households' preferences, though the selection of any particular price function is essentially arbitrary.

When to use Stated Preference Methods

Stated preference valuation methods are not limited in their application to the pure public good case discussed above. Rather practitioners could just as well apply stated preference method to the valuation of environmental goods that are implicitly traded in a property or wage market or that enter a household production function as a complement or substitute to some market good. However, the use of values from such stated preference exercises to inform the national accounts, carries considerable risks. In particular, in the cases of implicitly traded environmental goods and environmental goods that form complementary relationships in household production functions, the expenditures associated with those transactions already appear in the national accounts. A stated preference valuation of those environmental goods which ignored that fact would, at best, simply be adding measures of surplus to the accounts and at worst would result in double counting.

Replacement Costs

An alternative to preference based methods of valuing environmental goods is that of quantifying replacement costs. The replacement cost method values an environmental good according to how much it would cost to mitigate the losses that would result from a deterioration in the level of provision of that good. Such methods lack theoretical coherence with the preference-based values used to evaluate other productive activities. Replacement costs could potentially be very much higher than the values derived from those mitigating activities. In general, replacement cost methods should be avoided.

Conclusions

The intention of this note has been to explore and review the theoretical underpinnings of non-market valuation with the intention of providing insights as to how those methods might best be applied in identifying the "exchange prices" required for use in national accounting exercises.

A strong emphasis has been placed on establishing exactly how environmental goods provide utility to consumers. What is clear is that in many circumstances that utility provision can be linked in some way to associated purchases of a market good. That association with purchases of a market good is important for two reasons;

- First it provides a means by which we might infer the value of the environmental good through observations of real world market purchases
- Second it makes clear the circumstances under which the service flows coming from environmental goods have found expression in exchanges that are already recorded in the national accounts.

In the opinion of the author, methods of non-market valuation have the capacity to play a central role in identifying exchange prices that might be used in constructing an extended set of national accounts. At the same time, those methods have generally been applied on a much smaller scale; focusing on projects or localised areas. Scaling those methods up to undertake valuation exercises across a whole nation is a considerable challenge.

Of particular concern is the fact that the value of environmental goods will undoubtedly exhibit considerable spatial heterogeneity. Applications of hedonic pricing methods, for example, may need to be undertaken across a broad range of urban areas rather than the current norm of applying the method in just one location.

An important next step will be to identify a precisely defined set of environmental goods for which values are required. As discussed in the paper, the list of those final environmental good should not be excessively long and might include values for changes of risk of environmental damage to property, changes in risks of environmentally mediated ill-health or death, values for environmental features that significantly impact on the quality of recreational experiences or the amenity values drawn from landscapes.

To support that endeavour, a crucial task will be the on-going construction of high-quality spatially explicit datasets at a national scale that record environmental service flows and relevant market transactions.

Finally, this note has only considered the application of non-market valuation techniques to the valuation of flows of environmental goods. A separate set of issues, particularly those concerned with complexities and nonlinearities in the functioning of environmental systems and with discounting future environmental service flows must be addressed in attempting to value the environmental stocks from which those flows emanate.

References

Bartik, T. J., (1988). "Evaluating the benefits of non-marginal reductions in pollution using information on defensive expenditures", *Journal of Environmental Economics and Management,* 15, pp.111-127.

Bockstael, N. E. and K. E. McConnell (2006). *Environmental and resource valuation with revealed preferences: A theoretical guide to empirical models*, in the economics of non-market goods and services, vol 7. Springer, Dordrecht

Fisher, B. and Turner, K. (2008): Ecosystem services: Classification for valuation. *Biological Conservation* 141: 1167-1169.

Nordhaus, W. D. (2006). "Principles of National Accounting for Nonmarket Accounts," in Dale W. Jorgenson, J. Steven Landefeld, and Nordhaus, William D., *A New Architecture for the U. S. National Accounts*. Chicago: University of Chicago Press for NBER, pp. 143-60

Rosen, S., (1974). "Hedonic prices and implicit markets: Product differentiation in pure competition", *Journal of Political Economy*, 82(1) pp. 34-55.

SEEA (2012). *System of Environmental-Economic Accounting: Central Framework*. http://unstats.un.org/unsd/envaccounting/White_cover.pdf

SNA (2008). System of National Accounts. http://unstats.un.org/unsd/nationalaccount/sna.asp