HEDONIC PRICE INDEXES AND HEDONIC QUALITY ADJUSTMENTS

Jack E. Triplett
Brookings Institution
Washington, D.C., USA

Note: This document is chapter III in the author's forthcoming OECD publication:

Triplett, Jack E. Forthcoming. Handbook on hedonic indexes and quality adjustments in price indexes. Directorate for Science, Technology and Industry. Paris: Organization for Economic Co-operation and Development.

It is being posted as a background document (only) for:

SSHRC Conference on Index Number Theory and the Measurement of Prices and Productivity

Vancouver, B.C.

June 30 - July 3, 2004

Final June 5, 2004

CHAPTER III

HEDONIC PRICE INDEXES AND HEDONIC QUALITY ADJUSTMENTS

Conventional, matched-model quality adjustment methods are described in Chapter II. Hedonic methods offer an alternative methodology for producing quality-adjusted price indexes.

The first hedonic price index was estimated by Andrew Court (1939).[1] Their modern standing stems from the work of Zvi Griliches (1961), who like Court estimated hedonic price indexes for automobiles. Hedonic indexes have been implemented in the statistical systems of a number of OECD countries, mainly for "high-tech" electronic goods and housing.

A hedonic price index is ***any*** *price index that makes use of a hedonic function*. A *hedonic function* is a relation between the prices of different varieties of a product, such as the various models of personal computers, and the quantities of *characteristics* in them. Section A of this chapter contains an overview of hedonic functions and characteristics, and they are discussed more fully in Chapter V and the theoretical appendix.

As the definition of a hedonic price index implies, hedonic indexes may be computed in a number of ways. Alternative implementations are described in section C of this chapter. Research studies have typically estimated hedonic indexes from a regression by means of the "dummy variable" method (described in section III.C.1), but it is a serious misconception to take the dummy variable method as an essential or necessary part of a hedonic price index. The original Griliches (1961) paper contained several hedonic indexes, not just indexes produced by the dummy variable method. Moreover, hedonic computer equipment indexes that are estimated by OECD country statistical agencies typically do not use the dummy variable method. The definition of a hedonic index in the preceding paragraph makes explicit that a hedonic index is not necessarily a regression price index—and, conversely, a regression price index is not necessarily a hedonic index, if it does not make use of a hedonic function.

Whatever the implementation, estimating a hedonic function is the first step in computing a hedonic price index. Putting the exposition of hedonic indexes (this chapter and the following one) before discussing procedures for estimating hedonic functions violates the order in which a research project will be carried out. It is valuable, however, to consider the use of hedonic functions to compute hedonic price indexes before exploring the technical problems that arise in doing research on hedonic functions. Additionally, the order of the chapters facilitates the comparison of hedonic methods and conventional matched-model methods, a topic on which there has been much confusion.

---

[1] In recent years, historians of economic thought have discovered earlier researchers who developed something that resembles a *hedonic function* (though of course they did not use that term). See the historical note at the end of this chapter.

A. HEDONIC FUNCTIONS: A BRIEF OVERVIEW

The first government hedonic price index for computers was based on research by Dulberger (1989). Cartwright (1986) provides the details of implementation by the Bureau of Economic Analysis in the U.S. national accounts.[2] Most subsequent work on computers, including research on personal computers (PCs), follows an approach that is basically an extension of Dulberger's work, so hedonic functions that resemble hers describe most existing hedonic research on computer equipment (see Chapter V). For these reasons, Dulberger's computer hedonic function provides the illustration for this chapter.

Dulberger's (1989) computer hedonic function (her equation 2•) is:

(3.1) $$\mathrm{Ln}P_{it} = a_0 + 0.783 \ln(\text{speed})_i + 0.219 \ln(\text{memory})_i + \text{“technology” variables} + \varepsilon_{it}$$

Equation (3.1) says that the logarithm of the price for any computer model, i, at time t depends on the logarithm of its speed and the logarithm of the amount of "main" memory included in the machine, measured in megabytes. In Dulberger's study, speed was measured in MIPS, millions of instructions per second, a speed measure that has been widely used in the computer industry. Her equation also included a set of technology variables, which need not concern us at this point, nor do we need to discuss at this point the intercept term, a_0.[3] Finally, ε_{it}, the regression residual, indicates whether the price of the particular computer, model i, is close to the regression line at time t.

In common with many other computer hedonic functions, Dulberger's is in "double log" form—logarithms of all the continuous variables are used in the equation. Other functional forms have been used in hedonic research, including research on computers. These other forms are discussed in Chapter VI, but most of what follows in this chapter applies to hedonic indexes that are based on any of the most frequently-encountered functional forms. Equation (3.1) was estimated by ordinary least squares (OLS) regression, using a database of mainframe computer characteristics and prices that covered IBM mainframe computers and computers from certain other manufacturers that were constructed on similar computer architectural principles (to assure compatibility in the speed measure).

In the economics literature on hedonic functions, variables such as speed and memory size are called *characteristics*. The theory of hedonic indexes is built on the proposition that the characteristics are the variables that the buyers of the product want, and that the characteristics of the product also are costly to produce (see the theoretical appendix). That is, computer buyers want more speed, other things equal, and faster computers are more costly to produce, with a given production technology. An implication of the theory is that variables that do not have both "user value" and "resource cost" interpretations do not belong in hedonic functions. This is considered further in Chapter V.

The *regression coefficients* (0.783 and 0.219 in equation 3.1) are the major interest for this chapter. The regression coefficients value the characteristics. They are often called *implicit prices* or *characteristics prices*, because they indicate the prices charged and paid for an increment of one unit of (respectively) speed and memory. Implicit prices are much like other prices, they are influenced by demand and by supply—

---

[2] However, the first hedonic index used in the U.S. national accounts was a hedonic price index for new houses, introduced in 1974, extending back to 1968 **(***Survey of Current Business*, August, 1974, pp.18-27**).** Lowe (1999) indicates that the first Canadian government hedonic index was also for new houses, beginning in 1974. For earlier computer hedonic indexes, see the historical note to this chapter (Appendix A). A similar French index is described in Laferrère (2003).

[3] The technology variables were dummy variables indicating the technology embodied in the semiconductor chip used in the computer. This is considered in chapter V.

which means, for example, that they do not measure uniquely user value.[4] To avoid confusion, one should note that in a double log hedonic function such as equation (3.1), the regression coefficient is not technically itself the price of the characteristic, but rather the logarithm of the characteristic's price.

Arriving at a satisfactory hedonic function is the major part of the work in constructing hedonic price indexes. For the exposition of this chapter, I assume that a satisfactory hedonic function has already been estimated, following the principles developed in Chapters V and VI.

B. USING THE HEDONIC FUNCTION TO ESTIMATE A PRICE FOR A COMPUTER

For the hedonic price indexes discussed in section III.C, it is essential to explain how one can use a hedonic function to estimate, under certain conditions, the price of a computer (or any other product). For example, a new 50 MIPS computer replaces the 45 MIPS computer that was previously in the price index sample. In the language of Chapter II, this is an item replacement. One wants either (a) to estimate the price of one or both of the two computers for the period in which it was not available (to permit comparing prices of computers having the same specification), or (b) to estimate the value of the 5 MIPS increase in computer speed in one or both of two periods (to serve as a quality adjustment).

To illustrate this diagrammatically, suppose a hypothetical computer hedonic function that has only one characteristic, speed. This would resemble equation (3.1) with only speed in the regression, so it contains only the logarithm of a computer's speed and the logarithm of its price. The graph of this one-variable computer hedonic function (for some period, t) is shown in Figure 3.1. The estimation discussed in the following paragraphs is similar if there are more variables in the computer hedonic function, but it cannot conveniently be illustrated graphically.

1. Estimating Prices for Computers That Were Available and for Those That Were Not

Suppose some computer (call it model r) was sold in period t, but that model r was not in the price index sample of computers. This computer was available on the market, so its price could have been collected, had it been in the sample. Its speed was S_r.

Under the condition that computer r was available, one can use the hedonic function to estimate a price for computer r in period t. This estimated (or predicted) price for computer r is shown on Figure 3.1—if computer r operates at speed S_r, its price is estimated as P_r, the value predicted from the hedonic function for period t. P_r is the best estimate of the price for computer r, in the sense that it is the mean price for a computer that was sold in period t and that has speed S_r.

In this example, we do not know the actual price for computer r. It was not observed, or was not collected for the sample. The actual price for computer r might have been greater or lesser than the estimated price. If the actual price for computer r was lower than the mean price for a computer with speed S_r, then it is a "bargain." Bargain computers lie below the hedonic function—they have negative residuals from the hedonic function, shown on the graph in Figure 3.1 as a solid "dot." Conversely, computer r might have been "overpriced" relative to its speed. In this case, its price will lie above the hedonic function, so it will have a positive regression residual, as shown in Figure 3.1 by an open "dot."

Now consider estimating the price of a new computer—computer model n—that was not available at all in period t. This seemingly is a similar problem. On Figure 3.1, computer n has a speed rating of S_n. *If*

[4] See the Theoretical Appendix.

the introduction of computer n does not change the existing price regime for computers, then the best estimate of the price for computer n is also shown in Figure 3.1: From the speed of computer n, S_n, its predicted price is given by the hedonic regression line as P_n, by a procedure parallel to the one for estimating the price for computer r.

However, estimating a price for computer n from the hedonic regression for period t involves a subtly different economic and statistical problem from estimating the price for computer r. Computer r existed in period t, computer n did not exist. For this reason, the estimates P_r and P_n depend on different assumptions.

Computer r existed in period t, so its influence was already incorporated into the computer market pricing structure, and therefore into the period t hedonic function. No special assumption is required to estimate the price of computer r, *even if computer r was not included in the data from which the hedonic function was estimated*.[5] The question is simply: What is our best estimate of the price that computer r did *actually bring* at time t? Of course, this estimated price is estimated with error, as with all estimates.

But a seemingly similar question concerning the price of computer n is not identical. For computer n, the question is not: What did it sell for? Instead the question is: What *would* computer n have sold for *if it were* on the market in period t? If it were on the market in period t, it might have altered the demands for all computers, so that the hedonic function itself would have been different from the one that actually existed in period t.

The estimate P_n (which is based on the hedonic function for period t) is valid only if we can assume that the introduction of computer n in period t would not have changed the hedonic computer regression line in period t. This might be the case, for example, if computer n was introduced at a price that was exactly on the regression line, and not above or below it. This is quite an old point, but it has often been overlooked.[6] Its importance has been emphasised recently by Pakes (2003): Producers of differentiated products try to find unfilled niches in the product spectrum of varieties, and filling them changes the demands for existing varieties, resulting in changes in the characteristics prices.

Estimating the price of a new computer is most hazardous for computers that are outside the characteristics range that existed in period t. PC's of 1000 megahertz (MHz) speed became available in early 2000, and subsequently substantially higher speeds were offered.[7] Price index compilers may need to consider the question: What would a 1000 megahertz or faster PC have sold for in some earlier period, when machines so fast were not available? There may be no sensible answer to that question.

Presumably, building the 1000 megahertz machine was not possible with the technology of the earlier period. The first IBM PC, introduced in 1981, ran at 4.77 MHz, and the first 400 megahertz PC was introduced in 1998. It was certainly not feasible to produce a 1000MHz PC in 1981, and it is reasonable to presume that 400 MHz represented the upper limit to PC speed with the technology of 1998. As shown in Figure 3.2, the hedonic function for 1998 (designated h(1998)) has a vertical section at the technological upper limit of computer speed, 400 MHz. By 2000, technological change has shifted the hedonic function downward (computers of all speeds became cheaper) and also extended the hedonic function's domain into

---

5 This statement is not intended to imply that there are no econometric problems that might arise when computer r is missing from the sample on which the hedonic function was computed.

6 For example, Committee on National Statistics panel's (Shultze and Mackie, eds, 2002, page 124) review of hedonic indexes says: "The basic idea behind hedonic techniques is that one can use a hedonic equation to calculate the expected price of a particular variety—which may not in fact be offered for sale in the period being considered…." That requires special assumptions, as noted above, which the committee also noted at another point.

7 MHz, sometimes referred to as "clock speed," is another measure of computer speed that is widely used, especially for PC's. Alternative speed measures are reviewed in Chapter V.

regions of the commodity space that were not technologically feasible earlier—the vertical portion of the hedonic function now occurs at 1000MHz.

One cannot use the hedonic function for 1998 to predict what the 1000 megahertz machine would have sold for in 1998, for the correct answer to that question might be infinity, at least from the production side.[8] On the other hand, one can almost always answer the question: What would the typical computer sold in period t cost in period t+1? The hedonic function for 2000 (Figure 3.2) can be used to estimate the price for a 400MHz machine in 2000. Only if the two periods are far apart will the initial period's computer be obsolete. The 400 MHz computer was still available in 2000, but the 4.77 MHz PC was not.

Making quality adjustments in price indexes sometimes requires estimating a price in period t for a new computer that first became available in period t+1. One cannot ignore the fact that the new computer was not actually available in period t, and the possible reasons why it was not available. The special assumptions necessary to validate "backcasting" a price for a machine that was not in fact produced should be kept in mind.

2. Estimating Price Premiums for Improved Computers

For price indexes, another employment of the hedonic function will prove useful. Suppose computer m was in the sample in the previous period, but it was replaced in the sample by a new and faster computer, computer n. Rather than estimating the price for computer n in period t, we can ask instead: What premium should computer n sell for, compared with computer m? The expected premium provides a quality adjustment when computer n is an item replacement for computer m in a price index sample.

Using the hedonic function to estimate a price premium for the improved computer is really just a transformation of the hedonic function estimate considered in the previous section. Suppose computer m, with 400MHz, was in the sample (so its price was known), but it was replaced by computer n, with 500MHz. Suppose additionally for simplicity that the 400 MHz machine was priced to lie on the regression line. This is indicated by the open "dot" in Figure 3.3. The value of the extra 100MHz in computer n is given by the difference ((est P_n) – (P_m)), the logarithms of which can be read from the slope of the regression line in Figure 3.3. Estimating the increment to the price of the extra 100MHz implies exactly the same estimator as estimating the price of the 500MHz computer: The points marked "est $\ln P_n$" in Figures 3.1 and 3.3 are the same estimate, just arrived at in a different way.

Empirical estimates show that the premium for an additional 100 MHz of computer speed declines over time. For this reason, there are two possible valuations for speed in a price index that covers periods t and t+1—the initial period's value in period t and the comparison period's value in period t+1. One must decide which period's valuation is appropriate. Because the same choice is necessary for some other quality adjustment methods (production cost adjustments, for example), this is not a unique problem for hedonic indexes.

Using the hedonic function to estimate the price increment of an additional 100MHz does not evade the backcasting problem discussed in the previous subsection. If the replacement computer is really new and was not on the market in period t, then its introduction might have changed the price increment associated with an additional 100 MHz. Moreover, if the extra 100MHz was beyond the technological frontier of period t (as it is for the 1998 function in Figure 3.2), then the hedonic function is undefined for this increment, and so is the estimated price premium for the faster computer. Notice that even if the premium is

[8] It is common to think about goods that were not available in an earlier period using the demand side: "What would a buyer be willing to pay for a 1000 megahertz PC in 1998?" The suggestion is due to Hicks (1940). See Hausman (1997) for an empirical estimate. The willingness–to-pay estimate is clearly lower than infinity. However, the theory underlying hedonic functions requires that production be possible (see the theoretical Appendix).

undefined in period t, it may be estimated without difficulty for period t+1 (refer to Figure 3.2). On the other hand, if computer n was in fact available in period t—that is, it was missing from the price index sample but not from the market—no special assumption is required to estimate the price increment associated with its increment to speed. This is parallel with the discussion in the previous section. Because statistical agencies frequently choose item replacements that are similar to the computer in the old sample, the replacement is often within the range of characteristics available in the previous period.

3. Residuals

Hedonic functions have residuals, the ε_{it} term in equation (3.1), and the solid and open "dots" in Figure 3.1. For what follows in this chapter, we need to consider why these residuals exist, and what economic and statistical interpretations to give them.

Of course, "noise" always exists in actual data. Statistical noise is one interpretation for the regression residuals. That is another way of saying that the residuals may not mean anything economically, they are just random observational errors of some sort that are hopefully (but not probably) uncorrelated with anything of interest.

But the residuals may also have an economic interpretation. If the prices are transactions prices, negative residuals (the points marked with solid "dots" on Figure 3.1) are "bargains," as already noted: These computers cost less than one would expect from the quantities of characteristics they contain. Conversely, positive residuals (the points marked o on Figure 3.1) cost "too much" for what they provide in characteristics. Actual prices include manufacturer's pricing errors, so the residuals have an economic interpretation, not just a statistical one.

Griliches (1961) long ago noted that if the hedonic function is correctly specified, then the residuals should predict changes in market shares. Bargains should experience increasing market shares. Cowling and Cubbin (1971) explored this suggestion on the market for automobiles in the U.K. Knight (1966) used residuals from an early hedonic function for computers to predict whether a particular computer would find a market, considering its performance relative to its price.[9] Waugh (1928) analysed residuals from his hedonic functions for vegetables.

"Bargains" and "good buys" may exist alongside overpriced computers because people are not perfectly efficient shoppers. Some people know what to buy, others make mistakes. Alternatively, some people find it congenial to shop for the best price and best value, others do not want to invest their time, and buy without searching out the best price. For these reasons, the overpriced computers do not immediately disappear from the market. Retailing and marketing studies persistently show that the "law of one price" does not hold across retail outlets and especially across urban areas and among cities (one recent example of a large literature is O'Connell and Wei, 2002).

It has often been said that the hedonic function requires perfect markets, or that hedonic price indexes require market equilibrium. Such statements are partly germane and partly (depending on what is meant) misunderstandings. Consumer ignorance, in the sense described above, exists, and markets are not "perfect;" this just implies that hedonic functions have residuals. Regression residuals present no particular difficulties for the estimation and use of hedonic functions. Under and over-pricing errors are random for interpreting the hedonic function (though certainly not for analysing buyer behaviour toward market price differentials). It has also often been said that publications such as *Consumer Reports* in the U.S. and *Which?*

---

[9] Knight did not call his function a hedonic function, but it was one. His consulting use of his hedonic function to predict market success of new computers worked quite well (personal conversation with Kenneth Knight). Knight's work was converted into a price index for computers in Triplett (1989).

in the U.K. show there is no correspondence between price and "quality;" that is also a misconception. What these magazines show, in effect, is that hedonic functions have residuals, and they try to point buyers toward negative residuals and away from positive ones.

On the other hand, hedonic residuals could have a less benign interpretation. Suppose an important characteristic that consumers value is left out of the hedonic equation. For example, perhaps the computers marked "o" in Figure 3.1 include packages of software that buyers value and that are not offered on the other computers. If we added data on the missing software characteristics to the hedonic regression we might find that prices of these computers (including software characteristics) would be at the regression line, or even below it. Thus, when bundled software or other characteristics are omitted from the hedonic function, the residuals may reflect specification error in the hedonic function, rather than overpricing.

The residuals may also record other measurement errors that are distributed irregularly across the observations. If list prices are used as the left-hand side variable in the regression, the residuals may measure the error between list and transactions prices, where markups or markdowns differ among the various computers. If the prices are "unit values"—averages aggregated across sellers, as is typical in scanner data—residuals may reflect differences in retail amenities or services that affect the price but are in effect missing variables for a (retail price) hedonic function. Or perhaps we did not record the prices exactly right, or speed or some other variable is measured with error, and the errors are not uniform across the observations.[10] As another alternative, if we inadvertently mis-specify the functional form for the hedonic function (for example, estimate a linear hedonic function, when the true hedonic function is logarithmic), the measured residuals from the fitted hedonic function will not be the true residuals. In these cases, regression residuals may not measure bargains or overpriced machines; instead, they indicate errors in the hedonic function specification.

Missing or omitted or mismeasured variables and other specification errors can cause serious problems for hedonic functions and hedonic price indexes. I discuss in chapter V how to search for and minimise problems that arise from omitted characteristics in hedonic functions, and consider hedonic functional forms and other estimation issues in Chapter VI. For the rest of this chapter, it is appropriate to put these interpretative, procedural and econometric problems aside, for consideration later. We assume, for the exposition of this chapter, that the hedonic function is correctly estimated and its variables are correctly measured.

C. HEDONIC PRICE INDEXES

With these preliminaries, we are now ready to discuss hedonic price indexes. As defined at the beginning of this chapter, a hedonic price index is a price index that uses a hedonic function *in some way*. Four major methods for calculating hedonic price indexes have been developed and used to estimate computer and ICT price indexes.

Each of these four hedonic price index methods uses a different kind of information from the hedonic function. The first two described below (the time dummy variable method and the characteristics price index method) have sometimes been referred to as "direct" methods, because all their price information comes from the hedonic function; no prices come from an alternative source. Direct methods require that a hedonic function be estimated for each period for which a price index is needed.

---

[10] An extensive econometric literature considers the consequences of mismeasured variables for biases in the estimated regression coefficients. See any standard econometrics textbook, such as Gujarati (1995), and Berndt (1991), pp. 128-129, for application to hedonic indexes.

The second two hedonic price index methods (the hedonic price imputation method and the hedonic quality adjustment method) have been described as “indirect” or “composite” methods. They are often called “imputation” methods, because the hedonic function is used only to impute prices or to adjust for quality changes in the sample of computers in cases where matched comparisons break down. The rest of the index is computed according to conventional matched-model methods, using the prices that are collected in the statistical agency’s usual sample.

Indirect methods imply merging two sources of price information—the full cross-sections of computer prices and characteristics, to estimate a hedonic function, and the price index samples that statistical agencies normally collect. Indirect methods also imply that the hedonic function can be estimated from a data source that is different from the one used for calculating the price index, which in turn implies that the hedonic function can be estimated less frequently than the price index publication schedule. For example, if the price index is monthly, use of an indirect method permits estimating the hedonic function quarterly or semi-annually or even annually, according to whatever frequency is required to keep it up to date; the indirect method does not require estimating hedonic functions monthly in order to produce an index that is computed monthly, as the direct methods demand.

This “direct” and “indirect” language is not particularly felicitous. “Indirect” has sometimes been taken as suggesting that the index is somehow not a “real” hedonic index or that it is in some manner inferior to a “full” (“direct”) hedonic index. That is wholly a misconception, as the following sections show. Perhaps the misconception is not entirely attributable to language, but the “direct-indirect” language seems to have fostered it. Nevertheless, I follow what has become common, if somewhat recent, usage because it has attained such currency.

For expositional simplicity, I assume that only speed and memory size matter to computer buyers and sellers, and not some other characteristics of computers that have been left out of the analysis and not considered. That is, there are no omitted variables. I also assume that we have measured speed, memory, and price correctly, and that we have already estimated a hedonic function that is satisfactory, according to the principles discussed in Chapters V and VI. Variables are correctly measured and the hedonic functional form has been determined, empirically. Under these conditions, regression residuals measure over or underpricing of particular computers, relative to their included content of characteristics. Except where noted, I utilize the double log form of equation (3.1) in the exposition, not because a double-log form is necessarily best, but because it has been used so frequently in empirical studies on computers. The calculations are very similar, and in some cases identical, for other functional forms.

1. The Time Dummy Variable Method

Most research hedonic price indexes have used the time dummy variable method. For example, the hedonic price index for PC’s in Berndt, Griliches and Rappaport (1995) is a dummy variable index. The dummy variable method has also sometimes been called *the* direct method, in the sense that the index number is estimated directly from the regression, without another intervening calculation. Because there can be more than one “direct” method, I retain the historical “time dummy variable method” terminology.

a. The Method Explained

To illustrate, suppose one wants a computer price index for three periods. For each computer model, i, one must gather data on the characteristics of the machine, which are—using equation (3.1) for illustration—its speed (measured in MIPS or megahertz, MHz) and its memory size (measured in megabytes). The characteristics of each model of computer do not change between the periods: Indeed, a computer model is defined by its characteristics, so if one of the characteristics changes, this is a new model for our purposes, regardless of the manufacturer's designation.[11] The price of a computer model may of course change between periods, so we must have a price for each machine for each of the three periods.

To discuss the time dummy variable method, I need to change the time notation. Let the three periods for which the price index is calculated be designated τ, $\tau+1$, and $\tau+2$. Then, in the hedonic regression (equation 3.2, below), $t = \tau$ designates the first period, $t = \tau+1$ the second, and $t = \tau+2$ the third.

In one form of the dummy variable method, all three periods' data are combined into regression equation (3.2):

$$\text{(3.2)} \qquad \ln P_{it} = a_o + a_1 \ln(\text{speed})_i + a_2 \ln(\text{memory})_i + b_1(D_{\tau+1}) + b_2(D_{\tau+2}) + \varepsilon_{it}$$

where $t = \tau$, $\tau+1$, and $\tau+2$, respectively, for the three periods. Where equation (3.1) implied a hedonic function using data for only one period, equation (3.2) implies a single hedonic function that covers three periods' data, τ, $\tau+1$, and $\tau+2$. This approach is sometimes called a "pooled" regression, or a "multi-period pooled" regression, because data for several periods are "pooled."

On the left-hand side, the equation has the logarithm of the price of computer i in year t (which is τ, $\tau+1$, or $\tau+2$). A computer that is in the sample for all three years appears in the regression for all three years, with its appropriate price for each year. But some computers appear only in certain periods: If computer m replaces computer n (recall the example in Chapter II), each of these computers is in the regression only for the periods in which each is sold—τ and $\tau+1$, for the old computer m, and $\tau+2$, only, for the new computer n. Using notation adapted from that of Chapter II, equation (3.2) contains three arrays of computer prices: $P(M)_\tau$, $P(M)_{\tau+1}$, and $P(N)_{\tau+2}$.

The right-hand side has the same variables as in equation (3.1)—the intercept term, a_0 and for each computer, i, its speed and its memory. Speed and memory have regression coefficients a_1 and a_2, as in equation (3.1), except in this case the coefficients are constrained to be the same over all periods covered by the regression. They are in effect an average of the coefficients for each of the three periods. Constraining the coefficients is controversial; further discussion appears in section III.C.2.b, which also contains an empirical evaluation.

As with equation (3.1), the term ε_{it} in equation (3.2) records whether the price of computer i in year t is above or below the regression line. Note that in principle computer i could lie above the regression line in one of the three years, and below it in another, and will if its seller lowered its price more or less than the price changes for rival computers.

In addition to the variables in equation (3.1), equation (3.2) contains variables that record the year in which each price is collected: The first "dummy variable," $D_{\tau+1}$, takes on the value of 1 for each computer, i, when its price pertains to period $\tau+1$ (and zero otherwise). The second dummy variable, $D_{\tau+2}$, takes on the value 1 when computer i's price pertains to period $\tau+2$, and zero otherwise. The regression contains no dummy variable for period τ because τ is the base from which price change is computed.

The regression coefficient of the dummy variable, b_1 (strictly speaking, the antilogarithm of it[12]) shows the percentage change in computer prices between period τ and period $\tau+1$, holding constant the

[11] In practice manufacturers sometimes change their model codes when a characteristic of the product changes. However, sometimes a change might be introduced without a change in the model code and sometimes the model code may change without any change in the product's characteristics. Some research on computers has proceeded as if the manufacturer's model code makes getting data on the characteristics unnecessary; this is incorrect.

characteristics of the computer. It represents the price change (sometimes called the "pure" price change) between period τ and period τ+1 that is not associated with changes in computers' speed and memory.

Similarly, the regression coefficient b_2 shows the constant-quality computer price change between period τ and period τ+2. One could also code the data so the dummy variable coefficient b_2 gives the change between τ+1 and τ+2, instead of between τ and τ+2: In this case, the dummy variables are coded $D_{\tau+1} = 1$ for *both* periods τ+1 and τ+2 (that is, for all periods other than τ), and $D_{\tau+2} = 1$ for *only* period τ+2 (as before). The estimates are equivalent in either case. Because the latter—that is, where b_2 provides an estimate of the one-period change between periods τ+1 and τ+2—has a simpler interpretation, I discuss this case in the following.

The dummy variable price index is illustrated in Figure 3.4. Like the preceding figures, Figure 3.4 depicts a single variable hedonic function (so it has only a "speed" coefficient, a_1), with two time dummy variables, $D_{\tau+1}$ and $D_{\tau+2}$.

As Figure 3.4 suggests, with this technique there are really three regression lines, not one. They all have the same slope, a_1 (because the regression coefficients are constrained by the methodology to be the same in all three periods). Coefficients on the dummy variables act as alternative regression intercept terms: The time dummy variables allow the regression for each period to lie above or below the others.

If the coefficient on the dummy variable $D_{\tau+1}$ is negative—that is, $b_1 < 0$—then b_1 measures the rate at which the hedonic line or surface falls, and therefore the rate of decline in computer prices, computer characteristics (and also valuations of the characteristics) held constant. This is shown in Fig. 3.4 as the distance between the solid line ($t = \tau$) and the dashed one ($t = \tau+1$). The coefficient b_2 has a similar interpretation, with respect to τ+2—in Figure 3.4, it is the distance between the dashed and dotted lines.

When the estimated b_1 and b_2 coefficients are positive, computer prices are rising. When prices are rising, the dashed and dotted lines will lie above the solid one. The rest of the discussion applies to both rising and falling price situations.

This technique can be extended to more than three periods by "pooling" more periods of data and adding additional dummy variables as necessary. However, multi-period pooled regressions are not the best way to construct a time series of hedonic indexes.

A preferable alternative is the "adjacent period" approach. Using the same data as for equation 3.2, first combine data for periods τ and τ+1, and drop the variable $D_{\tau+2}$. This adjacent-period regression provides an alternative estimate of the coefficient b_1.[13] Second, combine data for periods τ+1 and τ+2 (discarding the data for period τ), and drop the variable $D_{\tau+1}$. This equation provides an alternative estimate of b_2. The three period price index for τ+2 / τ is obtained by multiplying together (antilogs of) the two adjacent year

[12] It is well established in statistics that the antilog of the OLS regression estimate of b_1 is not an unbiased estimate of b_1, which means that price indexes estimated by the dummy variable method are biased. A standard bias correction (Goldberger, 1968; also Kennedy, 1981, and Teekens and Koerts, 1972) is to add one-half the coefficient's squared standard error to the estimated coefficient. For hedonic indexes, this correction is usually quite small. I corrected nearly all of the computer equipment price indexes reviewed in Triplett (1989). The bias correction made little difference to reported annual rates of price change—because most hedonic functions for computers have very tight "fits," the dummy variable coefficient has a small standard error, and (typically) a rather large coefficient (prices are falling rapidly). This finding does not seem to have informed the recent hedonic price index literature.

[13] Or rather, of the population parameter that b_1 estimates. The statistical estimate of b_1 from equation (3.1) is not the same as the statistical estimate of b_1 from the alternative regression. The same statements apply to b_2. I retain the same notation for simplicity.

estimates b_1 and b_2. The adjacent-period alternative is still a pooled regression, but its pooling is the minimal amount necessary to implement the dummy variable method.

Any pooled regression holds fixed the hedonic coefficients, and holding the coefficients fixed has been criticised. However, the adjacent-period dummy variable approach holds the a_1 and a_2 coefficients constant for only two periods, where the multi-period pooled regression holds them constant for a larger number of periods, as indicated in equation (3.2). The adjacent-period estimator is a more benign constraint on the hedonic coefficients because coefficients usually change less between two adjacent periods than over more extended intervals (see section III.C.2.b). The adjacent-period form of the dummy variable method is preferred for most cases, and indeed should be thought of as "best practice" among dummy variable indexes. For this reason, in the following subsections, I use examples that are adjacent-period dummy variable indexes: As appropriate, some examples refer to an adjacent-period regression covering periods τ and $\tau+1$, others to an adjacent-period regression covering periods $\tau+1$ and $\tau+2$.

b. The Index Number Formula for the Dummy Variable Index

The coefficient of the $D_{\tau+1}$ dummy variable in equation (3.2) is an estimate of the logarithm of the price index between period τ and period $\tau+1$. Griliches (1971) remarked that the dummy variable method gave a price index that "is not well articulated with the rest of the index number literature." Traditionally, price indexes are computed by "formulas"—Laspeyres, Paasche, Fisher, and so forth. What price index formula corresponds to the time dummy variable estimate? We need the answer to compare analytically the dummy variable price index with a conventional price index.

Triplett and McDonald (1977, section IV) noted that the index number formula implied by the dummy variable method can be derived from the expression for the regression coefficient for the time variable, denoted as b_1 in the following. The index formula depends on the functional form for the hedonic function.

For a hedonic function with a logarithmic dependent (price) variable, rearranging the expression for the ordinary least squares (OLS) estimate of b_1 (for the dummy variable $D_{\tau+1}$ in equation (3.2)) yields the price index:

$$\text{(3.3a)}\qquad \text{index } \{(\tau+1) / \tau\} = \exp(b_1)$$

$$= [\ \Pi\ (P_{i,\tau+1})^{1/n} /\ \ \Pi\ (P_{i\tau})^{1/m}\ ] \div [\text{hedonic quality adjustment}]$$

In words, the dummy variable index equals the ratio of unweighted geometric means of computer prices in periods τ and $\tau+1$, divided by a hedonic quality adjustment. In the usual case, hedonic regressions are run on unbalanced samples, so the number of observations may differ in the two periods, as indicated by the subscripts m and n.

The hedonic quality adjustment in equation (3.3a) also depends on the form of the hedonic function. For a logarithmic hedonic function, the hedonic quality adjustment is given by:

$$\text{(3.3b)}\qquad \text{hedonic quality adjustment} = \exp\ [\Sigma\ a_j\ ((\Sigma\ X_{ij\tau+1} /n) - (\Sigma X_{ij\tau} / m))]$$

In equation (3.3b), I designate characteristics by the subscript "j" because in most cases there will be more characteristics in the hedonic function than the two in equation (3.1).

Equation (3.3b) is itself an index number—it is a quantity index that measures the change in characteristics of computers sold in periods τ and $\tau+1$. It is not a standard index number formula, partly because of the unbalanced samples in it. More exactly, the term in square brackets is the mean change in computer characteristics, e.g., speed and memory, between periods τ and $\tau+1$; the changes in computer characteristics are valued by their implicit prices, which are the a_j coefficients from the hedonic function (equation 3.2). The hedonic quality adjustment is the exponential of the square-bracketed term.

If the hedonic function in equation (3.2) were estimated with a sales-weighted regression, equation (3.3a) would be a ratio of weighted geometric means, rather than the equally-weighted geometric mean written there. The quality adjustment in equation (3.3b) would also involve means of characteristics that are weighted by sales, rather than the equally-weighted means in equation (3.3b). Similar statements apply to the following examples, but need not be repeated for every case.

For a linear hedonic function, the formula for the dummy variable index becomes an unweighted arithmetic mean, and the quality adjustment term is a linear quantity index of characteristics:

(3.3c) $$\text{index } \{(\tau+1) / \tau\} = (b_1)$$
$$= [3\ (P_{i,\tau+1})/n \ / \ 3\ (P_{i\tau})/m\,] - [\text{hedonic quality adjustment}]$$

(3.3d) $$\text{hedonic quality adjustment} = [\Sigma\ a_j\ ((\Sigma\ X_{ij\tau+1} / n) - (\Sigma X_{ij\tau} / m))].$$

Other hedonic functional forms imply dummy variable indexes that have index formulas that depend similarly on the expressions for the coefficients of the dummy variables, when the regressions are estimated by OLS.

Triplett and McDonald (1977) computed a price index for refrigerators by the dummy variable method, and then compared it with the conventional (Laspeyres) index number formula used in the U.S. PPI index for refrigerators. Using equation (3.3a) they decomposed the difference into an index number formula effect (geometric, compared with the arithmetic mean formula used in the PPI) and a quality adjustment effect (hedonic, compared with matched model method used in the PPI; the quality adjustment effect was estimated as a residual). The index number formula effect accounted for about 10 per cent of the difference between the dummy variable hedonic index and the PPI.

It is a bit surprising that little subsequent hedonic research has considered index number formula effects. It is well known that geometric and arithmetic mean formulas for price index components give different measures. Sometimes these formula effects are quite large—see, for example, the calculations for the Canadian CPI in Schultz (1994). For this reason, a great amount of attention has been given in recent years to the formula to be used for price index basic components. Diewert (1995) reviews index number theory for basic components; Balk (1999) and Dalén (1999) add more recent developments. This work on index number formulas for price index basic components seems not to have touched research on hedonic indexes, so that the typical hedonic study has remained poorly "articulated" into the price index literature, as Griliches put it. More seriously, few studies have determined what part of the difference between research indexes for PC's and a statistical agency's PC index lies in the use of geometric compared with arithmetic mean index number formulas.[14]

[14] As an example, the U.S. PPI for computers is a quality adjusted (by hedonic methods) arithmetic mean. Some part of the difference between the BLS index and the dummy variable index in Berndt, Griliches and Rappaport (1995) is caused by the difference between arithmetic and geometric means.

c. Comparing the Dummy Variable Index and the Matched Model Index: No Item Replacement

An enduring research issue concerns whether hedonic indexes and matched model indexes give similar or different results. To address this issue, whether empirically or analytically, one needs expressions for the alternative estimators that correspond to matched model and hedonic indexes.[15]

To facilitate comparison, consider exactly the same data that were discussed in chapter II. That is: (a) there are three periods for which we want a price index; (b) a matched set of m computer models exists in period τ and period $\tau+1$, and (c) computer n replaces computer m in period $\tau+2$, with the remaining m-1 models unchanged. In the notation employed in this chapter (see the discussion of equation (3.2), above), we have three arrays of computer prices: $P(M)_{\tau}$, $P(M)_{\tau+1}$, and $P(N)_{\tau+2}$, where the notation $P(N)_{\tau+2}$ indicates that computer n has replaced computer m in the sample in period $\tau+2$. As in Chapter II, this implies that $i = 1, \ldots, n$, but the number of observations in each two-period regression is 2m (because one observation is missing for each of the years).[16] We calculate equations (3.3a and 3.3b).

Suppose the statistical agency uses a geometric mean formula for its matched model index. The hedonic index in equation (3.2) implies a price index that is also a geometric mean, as shown in equation (3.3.a). Thus, no formula effects are present in the comparison of matched model and hedonic indexes.

No item replacement occurs between periods τ and $\tau+1$. The computer models in the price arrays $P(M)_{\tau}$ and $P(M)_{\tau+1}$ are only matched models. In equation (3.3b), each $X_{ij\tau}$ equals the corresponding $X_{ij\tau+1}$, and the number of items is also the same (m=n), so the terms inside the square bracket sum to zero. Thus, *the hedonic quality adjustment in equation (3.3b) is unity*, because there is no quality change in any of computers in the sample.[17]

Accordingly, the price index number formula implied by a dummy variable (logarithmic) regression *run on matched models* is a ratio of equally-weighted geometric means. This is the same unweighted geometric mean formula that has come to be preferred by many statistical agencies for detailed components of consumer price indexes.[18]

One would seldom want to compute a price index from the regression dummy variable when all of the computer models in both periods are matched models. Equation (3.3a) shows why: The dummy variable index yields exactly the matched model index, because (a) matching already holds the quality of computers constant, and (b) the index number formula that is implied by the regression coefficient on the time dummy variable is the same geometric mean that is commonly used by statistical agencies.

For completeness, it should be noted that statistical agencies also use formulas other than the geometric mean index to estimate basic component indexes—a ratio of arithmetic means, an unweighted arithmetic mean of price relatives, and in some cases weighted forms of these estimators. Obviously, the index $\exp(b_1)$ in equation (3.3a) will differ from these other basic component index formulas. For example, the estimate from equation (3.3c) would be consistent with an arithmetic mean index formula.

---

[15] An empirical review of matched model and hedonic indexes is contained in Chapter IV.

[16] The dummy variable method does not require particular assumptions about data availability. I use the data from Chapter II in order to make a clear contrast between different methods *applied to the same data*. A broader consideration of the elements held constant here is contained in Chapter IV.

[17] Note that $e^{(0)} = 1$.

[18] See, for example, Eurostat (1999); United States Bureau of Labor Statistics (1998).

d. Comparing the Dummy Variable Index and the Matched Model Index: with Item Replacements

In our example, periods $\tau+1$, and $\tau+2$ involve an item replacement in the index—computer n replaces computer m in the sample. For this $\tau+2$ / $\tau+1$index, the formula of the dummy variable hedonic index is parallel to the one before, so for an adjacent-period regression involving data for periods $\tau+1$, and $\tau+2$:

(3.4) price index $(\tau+2) / \tau+1 = \exp(b_2)$

$$= [\, A\, (P_{i,\tau+2})^{1/m} / \, A\, (P_{i\tau+1})^{1/m}\,] \div [\text{hedonic quality adjustment}]$$

As before, the hedonic quality adjustment is given by (3.3b), only noting that the subscripts in these expressions involve $\tau+1$ and $\tau+2$. Note that, because of the replacement, there are m computers in each of the two periods; this is a property of the example, which builds on the item replacement problem from Chapter II, in order to provide an explicit comparison of matched model and hedonic methods. Balanced samples are not necessary for hedonic indexes, and might not be typical of practical applications—for example, there might be entering computers in the hedonic index without any corresponding exits, or the number of entering and exiting computers might differ.

When an item replacement takes place, the ratio of geometric means in equation (3.4) is not matched. The geometric mean in the numerator applies to a different set of observations from the geometric mean in the denominator (the former includes computer m, the latter computer n). The dummy variable hedonic index in this case is not equal to the matched model index, even when the matched model index is an unweighted geometric mean index. The matched model index is computed on only the m-1 models that can be matched. In contrast, the hedonic index is computed over all the m models that are in the sample in each of the periods, because the hedonic index provides a quality adjustment, or imputation, for any price change associated with the replacement of computer m by computer n.

The item replacement implies the hedonic quality adjustment in equation (3.3b). The adjustment depends on the change in computer characteristics between models m and n and on the implicit prices of these characteristics, that is to say, on the coefficients of the hedonic function.

However, as explained in Chapter II, conventional index procedures for handling item replacements also imply a quality adjustment for an item replacement. Suppose the deletion (IP-IQ) method is used in the conventional index number methodology (refer to Chapter II). The *implicit* quality adjustment from the deletion (IP-IQ) method (A_3 from equation 2.3b from Chapter II) is reproduced below with the time subscripts changed to fit notation for this section. This implicit quality adjustment can be compared with the *explicit* hedonic quality adjustment (A_h) given by equation (3.3b).

(2.3b) $A_3 = (P_{n,\tau+2}/P_{m\tau+1}) / (A_i\, (P_{i,\tau+2}/P_{i,\tau+1})^{wi}),\ i \neq m, n$

(3.3b) $A_h = \exp\, [\Sigma\, a_j\, ((\Sigma\, X_{ij\tau+2}/n) - (\Sigma X_{ij\tau+1}/m))]$

It is evident that the two quality adjustments differ algebraically, but in a way that is not easy to summarize compactly. It is also evident that all the terms in the two questions can be quantified, with information from price index files and from an estimated hedonic function. It is possible, therefore, to calculate the sizes of the quality adjustments that are made with the two methods.

Suppose some other quality adjustment method is used in the conventional index, for example, manufacturer's production cost. Triplett (1990) and Schultz (2001) present information on manufacturers' cost quality adjustments for automobiles in the U.S. and Canadian CPIs. Such quality adjustment data from agency price index files could be compared quantitatively with the dummy variable hedonic quality adjustment from equation (3.3b).

Despite the quantifiable nature of these comparisons of quality adjustment methods, few empirical comparisons have been performed. Dalén (2002) and Ribe (2002) present implicit quality indexes for certain products and countries in Europe, but at this point no hedonic adjustments exist with which to compare them. Moulton and Moses (1997) present implied quality adjustments from a variety of BLS procedures for handling quality change. Schultze (2002) tabulates comparisons of hedonic and conventional quality adjustments for a number of products in the U.S. CPI, based on data from BLS (but the BLS hedonic indexes are not dummy variable indexes). Triplett and McDonald's similar calculations have already been discussed.

Although it is probably not feasible for outside researchers to produce much information on the adjustments made inside a statistical agency's price index sample, too few have considered the matter, despite the good example set by Griliches' (1961) thoughtful examination of the U.S. automobile indexes. As a result of this research gap, it is often not entirely clear, quantitatively, why research dummy variable hedonic price indexes differ from those of statistical agencies.

e. Concluding Remarks on the Dummy Variable Method[19]

For this chapter, I have used an example in which dummy variable regressions are run on a dataset consisting of the same computer models that were collected for the matched-model price index discussed in Chapter II. The example was chosen to isolate the differences between hedonic quality adjustments and conventional quality adjustments, holding constant other aspects of index number construction, as well as any differences in samples and data. The equations in this section indicate the effect of alternative quality adjustments on a price index for computers, or for other products, and can be used to quantify comparisons of conventional and hedonic quality adjustments.

By using this example, I do not necessarily mean to imply that one would want to run a hedonic regression on the small number of computers in a typical price index sample. Indeed, including a large number of computer models in the regression is important for econometric reasons, essentially to get more precise coefficient estimates. Moulton, LaFleur, and Moses (1999) estimate a hedonic function for television sets, using the TV prices in the U.S. CPI sample, roughly 300 price quotations per month. For most electronic and computer products, the number of prices collected monthly is no doubt much smaller.

In practice dummy variable hedonic indexes often differ from conventional statistical agency indexes because the hedonic indexes are based on a more comprehensive sample. Dummy variable indexes from the research literature have often been based on all the computers that exist in periods τ, $\tau+1$, and $\tau+2$; statistical agency indexes use a sample of these computers, perhaps a sample that remains fixed for all three periods, except for forced replacements. It is well established that fixed samples often give an erroneous estimate of price change when used for technological products where the models or varieties change rapidly (see the discussion in Chapter IV). Thus, dummy variable indexes (or any hedonic index that is based on a more comprehensive sample) may correct for some outside-the-sample changes that are missed in the matched-model, fixed-sample methodology. Outside-the-sample issues are discussed in Chapter IV.

2. The Characteristics Price Index Method

[19] See also the comparison of dummy variable and characteristics price indexes at the end of section III.C.2.

A second "direct" hedonic method uses the implicit characteristics prices (the regression coefficients from the hedonic function) in a conventional weighted index number formula. Griliches (1971) called this method a "price-of-characteristics index," and Triplett (1989 and earlier) and Dulberger (1989) speak of a "characteristics price index" or a "price index for characteristics." More recently, other names have been introduced. Schultze and Mackie (2002), apparently under the impression that the method is new, call it the "direct characteristics method." Moulton, LaFleur, and Moses (1999) call it "an alternative direct measure of the price change." Okamoto and Sato (2001) call it "the single period method." This recent proliferation of names has caused a certain amount of confusion and has obscured the commonality in several empirical studies reviewed in section C.2.b, below. In this handbook, I use the simpler and more descriptive—and earlier—name "characteristics price index."

The motivation for the characteristics price index method comes from the interpretation of hedonic function coefficients, such as those in equations (3.1-3.3). The coefficients estimate implicit prices for characteristics—the price for a one unit increment to the quantities of speed, memory, and other characteristics embodied in the computer (see the discussion in section III.A and in Chapter V). When estimating hedonic price indexes, researchers have mostly ignored these implicit characteristics prices, or else treated them as merely a step toward calculating something else. For example, in the dummy variable method, only the coefficient on the time dummy is used, though the researcher usually inspects the other coefficients to reassure that the hedonic function is reasonable. Pakes (2003) goes so far as recommending that researchers ignore the coefficients of the characteristics.

If characteristics prices are indeed prices, it is not a great step to construct an explicit price index from them. This price index for characteristics prices looks like any price index, except that the quantity weights are quantities of characteristics, rather than quantities of goods. A price index for computer characteristics is a natural consequence of the idea that consumers, when they buy a computer, purchase so many units of speed, memory and so forth. Computer buyers want the speed and memory contained in a computer box, and not the box itself. The quantities of speed and memory are also what producers of computers produce. This interpretation of the characteristics price index method only requires that the quantities of characteristics are the variables that buyers want and use and that sellers produce and market, and, additionally, that the hedonic function estimates the prices of these variables.[20]

a. The Index

The price index for characteristics is easier to illustrate if we assume that the hedonic function is linear. The Bureau of Labor Statistics uses a linear hedonic function for quality adjustments in its price indexes, partly for ease in interpreting coefficients that are expressed in monetary terms. However, there is nothing in the notion of a price index for characteristics that requires a linear hedonic function. Dulberger (1989) calculated a characteristics price index from a double log hedonic function. Moulton, LaFleur, and Moses (1999) and Okamoto and Sato (2001) calculated it from a semilog hedonic function.

We need hedonic functions for two periods, which following the conventions for this chapter will be designated as τ+1 and τ+2:

$$(3.5)\qquad P_{i,\tau+1} = c_{o,\tau+1} + c_{1,\tau+1}\,(\text{speed})_i + c_{2,\tau+1}\,(\text{memory})_i + \varepsilon_{i,\tau+1}$$

$$P_{i,\tau+2} = c_{o,\tau+2} + c_{1,\tau+2}\,(\text{speed})_i + c_{2,\tau+2}\,(\text{memory})_i + \varepsilon_{i,\tau+2}$$

[20] It does not require, for example, that characteristics prices equal the marginal costs of producers.

In equations (3.5) I have designated the regression coefficients as "c" to avoid confusion with coefficients from nonlinear hedonic regressions discussed earlier: The coefficient c_1 gives the price of speed in direct monetary terms (dollars or euros), and similarly for c_2, the price of memory. For example, a hedonic function estimated on the BLS data indicates that the price of an incremental 1 MB of memory cost \$1.699 in October, 2000 (see the regressions reported in Chapter V).

Using equations (3.5) but ignoring the intercept term, a Laspeyres price index for characteristics can be written as:

(3.6) $\quad \text{Index} = \Sigma_j \, c_{j,\tau+2} \, q_{j,\tau+1} \, / \, \Sigma \, c_{j,\tau+1} \, q_{j,\tau+1}$

As before, I designate the characteristics with the subscript "j" so the characteristics prices in period $\tau+1$ are designated $c_{j,\tau+1}$.

To construct any price index, we need weights, which are designated q_j in equation (3.6). Weights for a characteristics price index are quantities of characteristics. Suppose that this price index applies to the average computer purchased by the CPI population.[21] Then, the numerator of equation (3.6) is constructed from the mean initial-period ($\tau+1$) computer characteristics ($q_{j,\tau+1}$) valued by the second period's hedonic function, that is, by the characteristics prices, $c_{j,\tau+2}$. The denominator is the initial period mean computer characteristics, valued by the initial period's hedonic function, which is just the mean computer price in the initial period.

More generally, the initial period quantities of characteristics can be thought of as the total quantity of characteristics purchased by the index population. In the CPI one might think of the weights as the quantities of computer characteristics purchased by consumers in the initial period: For speed, the quantity q_j is the number of computers of each type times the speed of each type. Suppose for illustration that there are two computer types, with 400 MHz and 500 MHz, respectively. Suppose further that consumers bought 300 of the first type and 200 of the second. Then, consumers bought 220,000 MHz of speed (300 x 400 + 200 x 500).

Alternatively, one might use as weights the characteristics quantities in the typical or average computer consumers purchase. Taking the data in the above example, the mean quantity of speed bought by consumers is 440 MHz. Similar statements apply to the characteristic memory--for example, the weight is the total number of megabytes of memory the index population purchased in the initial period. This formulation of the price index for characteristics is more nearly attuned to the usual formulation of price indexes, and would result if we used "plutocratic" CPI population weights to determine the total quantities of characteristics in equation (3.6).

In a PPI context, equation (3.6) could be thought of as the total number of megabytes of memory produced by the industry and the total number of units of speed produced, in the initial period.[22] The numerator of the index values the initial period's output of characteristics with characteristics prices of the second period: It is the industry's hypothetical revenue if it sold the initial period's production of characteristics at the characteristics prices that prevailed in the second period. The denominator is the actual industry revenue in the initial period.

---

[21] This "average" machine might not exist; it is merely the average quantity of characteristics purchased. One might instead use the "median" computer purchased for calculating the index. In the example, this is the 400 MHz computer.

[22] An anecdote from the mainframe computer days: IBM engineers reportedly spoke of monthly computer production in terms of "number of MIPS shipped." MIPS was the speed measure used by Dulberger (1989). This engineering jargon meshes with the quantity units in the characteristics price index formulation.

The characteristics price index falls as the prices of computer characteristics fall. As an example, Dulberger (1989, table 2.5) reports that the price of one MIP fell from \$1.8 million to \$220 thousand between 1972 and 1984 (an annual rate of decline of 16 per cent), while the price of a megabyte of memory fell from \$497 thousand to \$25 thousand over the same interval (-22 per cent per year). Her characteristics price index fell at the rate of 17 per cent per year.

Some convention must be adopted for handling the intercept term, c_0. It can be interpreted as the price of a computer box with no performance at all. One can treat it simply as another characteristic, or as the group of characteristics that are not included in the regression (software, for example), which always has one unit per box in the index, and put its price in the index (equation 3.6).

As for any other price index, there are two periods and two possible sets of weights for a characteristics price index. Equation (3.6) corresponds to a Laspeyres characteristics price index because its weights were drawn from the initial period. One could also value the typical second-period computer with the hedonic function of the initial-period, or the second period's production of characteristics (in the PPI case) with the hedonic function of the initial period. That corresponds to a Paasche price index for characteristics, because it has characteristics quantity weights from the second period, that is, the weights are $q_{j,\tau+2}$.

A Fisher characteristics price index has the same relation to Laspeyres and Paasche characteristics price indexes as it does to the corresponding "goods-space" indexes. The Fisher index is the geometric mean of the Laspeyres and Paasche price indexes for characteristics. Still presuming a linear hedonic function (equation 3.5), the Fisher price index for characteristics is:

$$(3.6a)\qquad \text{Index} = \{[\Sigma\, c_{j,\tau+2}\, q_{j,\tau+1} / \Sigma\, c_{j,\tau+1}\, q_{j,\tau+1}]\, [\Sigma\, c_{j,\tau+2}\, q_{j,\tau+2} / \Sigma\, c_{j,\tau+1}\, q_{j,\tau+2}]\}^{1/2}$$

In equation (3.6a), the first term is the Laspeyres characteristics price index from equation (3.6), and the second term is the Paasche characteristics price index.

In general, however, the hedonic function is not linear. Using the same notation for characteristics j, and designating the hedonic function for period $\tau+2$ as $h_{\tau+2}$ (and for period $\tau+1$ as $h_{\tau+1}$), in the general case the Fisher price index for characteristics can be written:

$$(3.6b)\qquad \text{Index} = \{[h_{\tau+2}(q_{\tau+1}) / h_{\tau+1}(q_{\tau+1})]\, [h_{\tau+2}(q_{\tau+2}) / h_{\tau+1}(q_{\tau+2})]\}^{1/2}$$

In equation (3.6b), $q_{\tau+1}$ and $q_{\tau+2}$ designate two vectors of characteristics that are evaluated with the hedonic functions indicated.

In words, a Fisher price index for characteristics values two computer characteristics bundles, for example, the mean computer characteristics bought in periods $\tau+1$ and $\tau+2$, having characteristics $q_{\tau+1}$ and $q_{\tau+2}$, respectively. Those two characteristics bundles (weights) define Laspeyres (weights $q_{\tau+1}$) and Paasche (weights $q_{\tau+2}$) index numbers. In turn, the two characteristics bundles are valued by the hedonic function of period $\tau+1$ and by the hedonic function of period $\tau+2$. In the first term on the right hand side of equation (3.6b)—the Laspeyres half—the initial period's characteristics quantities ($q_{\tau+1}$) are valued by the hedonic function in the initial period ($h_{\tau+1}$) in the denominator, and the hedonic function for the comparison period ($h_{\tau+2}$), in the numerator. The Paasche half, the second term in equation (3.6b), uses the characteristics of the comparison period, $\tau+2$, valuing them again by $h_{\tau+2}$ in the numerator and $h_{\tau+1}$ in the denominator. The hedonic function gives the characteristics prices for both indexes; the prices will generally differ between the two periods, as in any price index.

A complication arises when the hedonic price surface is nonlinear, because a characteristic does not have a single market price. Different buyers pay different prices for characteristics, depending on their preferences for characteristics and where they are "located" on the hedonic function (see the theoretical appendix). In equation (3.6b), the hedonic function $h_{\tau+2}$ is used to evaluate computers with characteristics $q_{\tau+1}$ (for the numerator of the Laspeyres index component) and for the characteristics $q_{\tau+2}$ (for the Paasche index component). The corresponding characteristics *prices* will not be the same in these two calculations, because the two machines being evaluated lie on different portions of the hedonic function, and the characteristics prices depend on the point on the (nonlinear) hedonic function that is being evaluated. The same point applies when the same two machines are evaluated with hedonic function $h_{\tau+1}$ (both denominators). Because of this complication, when the nonlinear hedonic function is used to computer the characteristics price index, weights and prices vary in price and quantity indexes for characteristics in ways that they do not in normal price indexes for goods.[23] This is not a difficult computational problem. I have added this explanatory paragraph to avoid confusion.

b. Applications

The price index for characteristics has many applications, calculating it for the typical or average computer in the CPI was just an example. One could calculate separate characteristics price indexes for each model, or for each buyer. In a study of television sets, Moulton, LaFleur, and Moses (1999) estimate price indexes for characteristics for every set in the CPI sample, and calculate Laspeyres and Fisher price indexes for characteristics. That is, in equation (3.6b), they use (for the Laspeyres) each $q_{ij,\tau+1}$ corresponding to each item i in the CPI sample for period τ+1, and each $q_{ij,\tau+2}$ (for the Paasche) in the CPI sample for period τ+2. The Fisher index is calculated in the usual way.

Using the characteristics price index method, it is particularly easy to compute price indexes for individual households, or for household groups. To do this, one sets the characteristics of the computer purchased by household k ($q_{kj,\tau+1}$ and $q_{kj,\tau+2}$), or by household group g, into expressions (3.6), (3.6a), or (3.6b). One might calculate, for example, a computer index for computers bought by the middle deciles of the income distribution, or for any other specified household group. Because they buy different computers, different consumers or consumer groups experience different price changes; the price index for characteristics approach permits calculating different computer price indexes for individuals or for household groups.

As another example, Moch and Triplett (2004) estimate a Fisher-type characteristics price index to obtain a place-to-place (purchasing power parity, or PPP) price index to determine if computer prices differ in France and Germany. This PPP uses characteristics quantities corresponding to mean computer purchases in each of the two countries as weights to calculate the PPP, and uses as well characteristics prices for each country that were computed from hedonic functions for France and Germany.

One can also calculate characteristics quantity indexes. A Laspeyres characteristics quantity index is:

---

[23] This matter does not come up in "good space" indexes, which are the normal kind. For example, the Laspeyres price index for goods corresponds to a price surface for goods that is often called the "budget constraint." It is linear by the *assumptions of the consumer demand model*. Thus, using the standard model implies the assumption that the consumer always pays the same price, no matter what quantity the consumer purchases. This is not the case for implicit characteristics prices, because the corresponding price surface (for characteristics prices) can be nonlinear. See the theoretical appendix.

(3.7) Q-index = $\Sigma\, q_{j,\tau+2}\, c_{j,\tau+1} \,/\, \Sigma\, q_{j,\tau+1}\, c_{j,\tau+1}$

For example, Griliches (1961) computed the change in the characteristics of a group of automobiles (those that were then included in the U.S. CPI), valued by the base-period hedonic function's estimated prices for characteristics. Griliches called this a "quality index" for automobiles, but it was also a quantity index of characteristics, for it measured the changes in characteristics quantities, valued by characteristics prices. Griliches then divided the change in average list price for these CPI cars by the quantity index in equation (3.7) to get a quality-adjusted price index. This is much like national accounts practice, where the change in expenditures on automobiles is divided by a quantity index to produce an implicit price index for automobiles.

As with the characteristics price index, the characteristics quantity index can be computed with Laspeyres or Paasche or Fisher (and other) formulas. When the hedonic function is nonlinear, as it generally is, the Fisher characteristics quantity index is:

(3.7a) Q-index = $\{[\, h_{\tau+1}\,(q_{j,\tau+2}) \,/\, h_{\tau+1}\,(q_{j,\tau+1})]\; [h_{\tau+2}\,(q_{j,\tau+2}) \,/\, h_{\tau+2}\,(q_{j,\tau+1})]\}^{1/2}$

In equation (3.7a) all the symbols correspond to those of the characteristics price index in equation (3.6b). Thus, the Laspeyres component of the quantity index (the first term in the brackets on the right hand side of equation (3.7a)) has period τ+1 characteristics price weights (from the period τ+1 hedonic function), the Paasche component has period τ+2 characteristics price weights (from the period τ+2 hedonic function).

As noted in section III.C.1, the hedonic quality adjustment in the dummy variable method is also a characteristics quantity index, but it is not the same index. Compare equation (3.7a) with the quantity index in equation (3.3b).

c. Comparing Time Dummy Variable and Characteristics Price Index Methods

The characteristics price index method has several advantages over the dummy variable method.

1. Index number formula. The index number formula for the characteristics price index is an entirely separate matter from the form of the hedonic function. As explained in the theoretical appendix, the form of the hedonic function depends only the empirical relation between the prices of varieties of the good and their characteristics. However, in the time dummy variable method, the form of the hedonic function also dictates the index number formula (see section III.C.1); this may not provide the desired price index number formula, because the index number formula depends on the usual theoretical conditions from index number theory, not on the statistical relation established empirically by the hedonic function.

The characteristics price index method permits the two functional forms (hedonic function and index number) to be determined independently. For example, one can compute characteristics price indexes using a superlative index number formula (Diewert, 1976) without imposing a specific superlative aggregator function on the hedonic surface. In recent literature, Silver and Heravi (2001b) emphasise particularly the possibility of estimating superlative indexes of characteristics prices.

The price index for characteristics permits breaking the connection between hedonic functional form and index number functional form. This is a theoretical as well as a practical advantage.

A similar point can be made about weights. A weighted index number is always preferred to an unweighted one, if weights are available. One might also prefer a hedonic function that is estimated from a weighted regression, but that is not always the case (there are econometric arguments for either option—see

Chapter VI). With the time dummy variable method, one cannot have the one without the other. However, a weighted characteristics price index can readily be constructed using coefficients from an unweighted hedonic function, if desired.

2. Constraining the regression coefficients. The dummy variable method requires that hedonic coefficients be constrained to be unchanged over all the time periods included in the regression. This has always been criticized, because it amounts to maintaining the hypothesis of constant (characteristics) prices—a kind of parallel to the "fixed weight" assumption that produces bias in Laspeyres price indexes for goods. The usual "Chow test" (see Berndt, 1991, for the use of Chow tests for hedonic functions) can be applied to determine if hedonic coefficients differ statistically across the periods of a study. Among the studies of high tech goods that have reported the results of such empirical tests, most record rejection of the hypothesis that the coefficients are unchanged, even between adjacent periods.

The characteristics price index method is not subject to this criticism. Because the method uses one hedonic function for each of the periods included in the price index, not a single hedonic function for both periods, all the characteristics prices can change over the periods of the price index calculation.

Constraining the characteristics prices is a difficulty with the dummy variable method, but possibly too much has been made of this matter recently. For example, Schultze and Mackie (2002) urge the U.S. BLS (and by extension, other statistical agencies) to avoid putting resources into the dummy variable method, but to conduct research on the characteristics price index method (which, as noted earlier, the panel thought was new). Is there a conceptual or empirical basis for this recommendation? The following section explores the issues.

3. Evaluation: constraining regression coefficients in the dummy variable method. Empirical evidence indicates that hedonic function coefficients frequently differ in adjacent periods, $\tau+1$ and $\tau+2$, even when the periods are close together. The coefficients are even more likely to differ when periods are far apart, which will be the case when multi-period pooled hedonic regressions are fitted, so constraining them to equality over multiple periods is even more likely to be rejected empirically on normal statistical tests. As noted in the previous section, this is the basis for the recent criticism of the dummy variable method.

However, what matters is the price index, not the coefficients themselves. So, the essential question is: Does constraining the coefficients make any substantial difference to the dummy variable price index? Surprisingly, none of the recent criticisms of the dummy variable method—including Schultze and Mackie (2002)—cites any quantitative support whatever. I review the empirical evidence in the next subsection.

4. Empirical studies. If constraining the hedonic coefficients matters empirically, a dummy variable index will differ from a characteristics price index, which does not constrain the coefficients. A few studies have compared dummy variable and characteristics hedonic price indexes computed from the same data. They are summarized in Table 3.1.

In most cases, I evaluate adjacent-period dummy variable estimates, not multi-period pooled regressions. The objection to constraining the coefficients is, after all, valid. Best practice implementations of the dummy variable method should constrain the coefficients as little as possible. Adjacent-period dummy variable indexes provide the minimum constraint on characteristics prices that is compatible with the dummy variable method. Multi-period hedonic regressions should normally be shunned, unless it can be shown empirically that coefficients have not changed over the periods for which they are held constant in the hedonic regression.

Dulberger (1989) estimated several hedonic indexes for mainframe computers, as noted already. Her dummy variable index and her characteristics price index differed about two percentage points per year, which is not negligible (Table 3.1). However, Dulberger's characteristics price index differs computationally from the methods outlined in the section on that topic (section III.C.2), and it is unclear how much this may matter. She also reported that she could not reject the hypothesis of constant coefficients, so holding them constant in a dummy variable index is not so problematic as it would be when the coefficients change.

Okamoto and Sato (2001) computed dummy variable and characteristics price indexes (they called them "two period" and "one period" estimates) for Japanese PC's, television sets, and digital cameras. In all three cases the two indexes coincided very closely (Table 3.1): In TV's, the two indexes agree to the tenth (10.4 percent per year decline in four years), and the same exact correspondence emerged for digital cameras, over a shorter period (21.9 percent decline). In PC's, the difference between the two hedonic indexes was quantitatively greater, but not relatively so (0.6 percentage points on a 47 percent per year decline). Chow tests on equality of adjacent month coefficients were not published in the paper; however, coefficient equality was rejected for PC's, was not rejected in the case of colour TV sets, and was ambiguous for digital cameras.[24]

Berndt and Rappoport (2002), in research still in progress, compute dummy variable and characteristics price indexes for laptop and desktop PC's. They report rejection of the hypothesis that coefficients were equal across periods. Their index comparisons yield results that are consistent with the other studies: The differences between dummy variable and characteristics price indexes are small. Moreover, they do not always go in the same direction. For example, the laptop indexes differ by one percentage point over the 1996-2001 interval (on a forty percent annual rate of decline), and by about the same amount for 1991-96, but in the opposite direction. For desktops, the differences are slightly larger both absolutely and relatively (1.4 percentage points for each period), but again the sign of the difference reverses between periods.

Finally, Silver and Heravi (2002, 2003) computed dummy variable indexes and characteristics price indexes for U.K. appliances over the twelve months of 1998. In both cases, equality of the adjacent-month coefficients was rejected. For washing machines, the two index number calculations differed by only 0.2 percentage point.[25] For TV's, the difference is slightly larger, 0.4 percentage points, but the TV comparison was done on a dummy variable index for a pooled 12-month regression, so the difference would undoubtedly be smaller for an adjacent-month regression.[26] In neither case is the difference between the two hedonic indexes large.

The studies in Table 3.1 amount to eight author/product pairs, three of them for personal computers. Taken together, they suggest negligible empirical difference between the dummy variable method, which

---

[24] Private communication with Masato Okamoto, January 26, 2003.

[25] The published version of their research compared a dummy variable index from a twelve-month pooled regression with a characteristics price index computed from regression coefficients from monthly regressions, and then chained. This comparison gave a difference of 1.6 percentage points, but the twelve-month pooled regression constrains the coefficients more than is necessary to estimate a dummy variable index. Saeed Heravi has kindly recomputed the dummy variable price index for this study so that it is based on adjacent-month regressions, with the results chained to give a 12-month index (private communication, January 25, 2003). This calculation is the one displayed in table 3.1. The same communication conveyed results of Chow tests on adjacent-month coefficient equality. Notice that 12-month pooled and adjacent-period regressions yield dummy variable indexes that are farther apart than are adjacent-month and characteristics price indexes. Constraining hedonic coefficients matters. However, the relevant comparison is between methods that do not constrain the coefficients at all and a form of the dummy variable method that puts minimal constraint on the coefficients that is compatible with the dummy variable method.

[26] See the previous footnote.

constrains hedonic coefficients to equality across periods, and the characteristics price index method, which does not constrain the coefficients. The only caveat is that the result applies only to best practice dummy variable estimates, not to multi-period pooling, which constrains the coefficients more than is necessary.

Eight studies are hardly the last word on this subject. This suggests a second question: Should we expect constraining the coefficients to make a difference in the price index? How likely are additional studies to find index number differences from constraining hedonic coefficients that are larger than those that have been found so far?

5. Analysis. The estimator for the time dummy variable suggests that constraining the regression coefficients usually will not have a large effect on the index, at least in adjacent-period regressions.

Recall from section III.C.1.a that the dummy variable index can be written as the ratio of two geometric means (in the logarithmic regression case) multiplied by a hedonic quality adjustment. For the logarithmic case, the hedonic quality adjustment is a quantity index of characteristics, having the following form (this is taken from equation (3.3.b), above):

$$\text{hedonic quality adjustment} = \exp\,[\Sigma\, a_j\, ((\Sigma\, X_{ij\tau+2} /n) - (\Sigma X_{ij\tau+1} / m))]$$

As noted earlier, this quality adjustment index is normally computed from an unbalanced panel (m is not necessarily equal to n). But the unbalanced panel poses no problem because the ratio of geometric means that is being adjusted is also computed on an unbalanced panel.

The constrained hedonic coefficients are, of course, the a_j terms, which the adjacent period dummy variable method requires to be the same in both periods. The true coefficient vectors (the implicit prices) vary by period, that is, they are $a_{j\tau+1}$ and $a_{j\tau+2}$. Empirically, the adjacent period regression that combines data for periods τ+1 and τ+2 often yields coefficients that are approximately the average of coefficients obtained from a separate regression for τ+1 and another for τ+2. Thus,

$$a_j \,.\, (a_{j\tau+1} + a_{j\tau+2}) / 2$$

This is not a necessary condition, from the econometrics of the regressions, but it is one of the earliest and most widely documented empirical regularities in hedonic functions. For example, this condition was found in Griliches' (1961) pioneering study of automobiles (see the extract in Table 3.2): Every coefficient in the pooled, adjacent-period regression lies midway between coefficients of the corresponding single-year regressions.

Substituting the approximation back into equation (3.3b) yields:

$$\text{modified hedonic quality adjustment} \,.\, \exp\,[\Sigma\, ((a_{j\tau+1} + a_{j\tau+2}) / 2)\, ((\Sigma\, X_{ij\tau+2} /n) - (\Sigma X_{ij\tau+1} / m))]$$

The formula for the modified quality adjustment corresponds to no standard index number formula, but its weights resemble those in a Tornqvist index (though the second term is a difference, rather than the Tornqvist index ratio). Is it likely to differ from the (unmodified) quality adjustment term implied by the usual dummy variable formulation, in which the coefficients are constrained? The algebra suggests: not much.

A number of writers have denounced the dummy variable method. Most of them seemingly have not thought through fully the index number implications of holding constant the characteristics prices. Although I slightly prefer other hedonic methods over the dummy variable method (for other reasons), the "fixed coefficients" matter has not nearly so much empirical importance as some critics have presumed.

Holding coefficients fixed in a multi-period pooled regression will no doubt create more serious problems, but multi-period pooled regressions can easily be avoided in favour of a sequence of adjacent-period regressions. In the rare case where small samples force multi-period regressions, the researcher should bear in mind the potential problems created by constraining the coefficients to equality over too lengthy a period, but no practical advice can be provided in such situations, solutions must depend on the researchers' good judgements.

d. Concluding Remarks on the Characteristics Price Index Method

This method is not a new idea. The U.S. new house price index, the first hedonic price index in any country's statistics (it dates from 1968), is a form of the characteristics price index. In the hedonic computer price literature, Dulberger (1989) introduced the price index for characteristics. It was also discussed in the review of computer price research in Triplett (1989), and actually dates back to Griliches (1961), and in quantity index form to Court (1939). Some regrettable confusion has arisen because some recent researchers have not recognized the lengthy intellectual antecedents of the characteristics price index approach.

The characteristics price index method has a number of useful attributes. First, the characteristics price index provides an alternative way to think about hedonic price indexes and quality change, and a more explicit or obvious way to link hedonic indexes into the conventional price index literature. A theoretical characteristics price index was the subject of Triplett's (1983) translation of price index theory from "goods space" to "characteristics space," in order to analyse quality change. Fixler and Zieschang (1992) have fruitfully extended the same idea, as have some other following writers. The theory of hedonic price indexes (see the Theoretical Appendix) is really a theory of the characteristics price index, so estimating a hedonic price index for characteristics meshes more obviously with the theory than is the case for some other estimation methods, where the connection is less immediately apparent.[27] When the characteristics price index is computed using a superlative index number formula, it has the natural interpretation (in the consumer price case) of an approximation to a COLI "subindex" on computer characteristics.[28]

The characteristics price index concept also suggests alternative hedonic calculations that can be quite useful and enlightening in certain contexts. Because hedonic prices differ across household units (see the theoretical appendix), so also will the quality adjustment. This is not a new point, but it is usefully emphasized in the Committee on National Statistics report (Schultze and Mackie, 2002). The weights that individual households give to characteristics will also differ. For example, households who need graphics capabilities value graphics and video cards more than households who use their computers only to calculate, or for word processing. The characteristics price index offers a way to take into account these differences in households' tastes, incomes, the characteristics prices they face in consequence, and the weights that should apply to the characteristics in calculating the index.

One theoretical problem becomes more apparent with a characteristics price index. The price index for characteristics implies that *only* the characteristics matter to buyers. This is a very special assumption. Two 400 MHz computers will not do what one 800 MHz computer will do, so how the characteristics are packaged into the computer box matters to consumers. See the discussion of this point in the theoretical

[27] I suspect that some of the "hedonic indexes have no theory" criticism that is still heard, long after considerable theory has been developed, might have less currency if the price index for characteristics were better known.

[28] A COLI subindex is defined in Pollak (1975).

appendix, especially the contribution of Pollak (1983). One gets around this problem to an extent by including the price of the box (the intercept term) in the price index, but it is an end run, not a confrontation or a solution.

Both time dummy and characteristics price index methods use information from the hedonic function to compute price indexes, without recourse to the prices for the individual goods themselves. They both require that the hedonic function be estimated for each period. The price index for May, for example, requires a hedonic function for May, as well as one for the previous or initial period. This requirement poses severe problems in practice.

Both methods also imply that the price index is computed on a large dataset, because estimating a hedonic function requires an extensive cross-section of prices and characteristics. Thus, if not a universe of computers, these two methods require at least a larger sample than is usual for statistical agency methodology, and (as noted previously) this larger sample must be available monthly, for a monthly index.[29]

The timeliness and database requirements of the dummy variable and characteristics price index methods have greatly inhibited their adoption in official statistics. These problems have led to the development of alternative methods for implementing hedonic indexes (considered in subsections 3 and 4) that have less restrictive practical implications.

3. The Hedonic Price Imputation Method

In this subsection and in the next, I consider two hedonic methods that have been called “indirect” or “composite” methods because they superimpose a hedonic imputation or hedonic quality adjustment onto the traditional matched-model index. This language is a sense a misnomer: Section C.1.b shows that a matched model index and a dummy variable hedonic index are algebraically the same when there is no quality change and index number formulas (and databases) coincide. When there is quality change, the dummy variable hedonic index is equivalent to a matched model index with a hedonic quality adjustment. From this, using an explicit hedonic imputation (this section) or an explicit hedonic adjustment (the next section) in a matched model index produces a hedonic index that is no less “direct” and no more “indirect” that any other method, because the observations that are matched are handled equivalently in both “direct” and “indirect” methods.

a. Motivation

To motivate the hedonic price imputation method, note that three arrays of computer prices, $P(M)_\tau$, $P(M)_{\tau+1}$, $P(N)_{\tau+2}$, were used in the regression of equation (3.2). The models in the first two arrays—P(M) in both period τ and period τ+1—are all matched, so their characteristics are already held constant by the fundamental index number matching methodology described in Chapter II.

Somewhat the same point can be made about the second interval (periods τ+1 and τ+2). Although an item replacement occurred, only two differences exist between arrays $P(M)_{\tau+1}$ and $P(N)_{\tau+2}$: One computer (computer m) disappeared and one new one (computer n) appeared. All of the other m-1 models were unchanged.

---

[29] The hedonic functions for TV sets estimated by Moulton, LaFleur and Moses (1999) used the dataset actually collected for the U.S. CPI. It contained around 300 observations monthly. The sample for many high tech products is likely to be smaller, even in the U.S. Additionally, as the authors and Schultze (2002) point out, the CPI replacement rules restrict the hedonic sample, so that new items, or items with characteristics differing from the items that exit the sample, are less likely to be included.

Statistical agencies may be reluctant to change calculation methods for the m-1 computers where the traditional approach appears adequate. Pakes (2003) emphasizes the variance of the index: Matched observations contain sampling variance,[30] but hedonic imputations contain both sampling variance and estimation variance. Another reason is the omitted variables that are associated with the retailer, and not necessarily with the product. Matching holds those mostly unobservable variables constant, so the agencies are properly reluctant to surrender control over them in order to improve their methods for handling a relatively small number of sample exits and entries.

For both operational and statistical reasons, it is thus natural to use traditional matched model price comparisons for the m-1 unchanged computers, and to direct attention to devising an imputation for the missing computer prices. The hedonic imputation method permits exactly that. Where matched model comparisons are possible, they are used. Where they are not possible, a hedonic imputation is made for the item replacement. Hedonic imputation methods make maximum use of observed data, and minimum use of imputation, thereby minimizing estimation variance. The hedonic imputation method was employed in the hedonic computer indexes introduced into the U.S. national accounts in 1985 (Cartwright, 1986).

b. The Imputation and the Index

Section III.B.1 already discussed using the hedonic function to estimate one, or both, of the "missing" prices, which are $P_{m,\tau+2}$ and $P_{n,\tau+1}$. One imputation involves computer n, an entering machine that first appeared in the sample in period τ+2. Because statistical agencies will nearly always link over in some fashion to follow computer n in future periods, it is natural to impute the price of computer n in the period (τ+1) before it was introduced, and compare the imputed price of computer n in period τ+1with its actual price in period τ+2.

We begin from a regression similar to equation (3.2), but specified for only one period:

$$(3.8) \qquad \ln P_{it} = a_o + a_1 \ln(\text{speed})_i + a_2 \ln(\text{memory})_i + \varepsilon_{it}$$

The imputation for computer n (which first appeared in period τ+2) requires that equation (3.8) be run on data for the previous period, period τ+1. Thus, in equation (3.8), t = τ+1. From an operational point of view, this imputation is convenient because the hedonic function for period τ+1 can be prepared in advance and the agency can have it ready to impute the price when needed.

The imputed price for computer n is calculated from equation (3.8) as:

$$(3.9) \qquad \text{est } P_{n,\tau+1} = \exp\ \{a_{0\tau+1} + a_{1\tau+1}\ (\text{speed})_n + a_{2\tau+1}(\text{memory})_n\ )\}$$

where the "est" designates an estimated or imputed price, and the subscripts τ+1 indicate coefficients from the hedonic function for period τ+1. The values of speed and memory in equation (3.9) correspond to the characteristics of computer n, as the subscripts indicate. Thus, we estimate the price of new computer n in the period prior to its introduction by valuing its characteristics—its speed and memory—at the implicit prices for speed and memory that are estimated from the regression that applies to period τ+1. Because

[30] If a probability sample; a variance is not defined for non-probability samples.

equation (3.9) is logarithmic, the regression prediction is biased as an estimate of the predicted price; the usual adjustment involves adding one-half of the regression variance estimate as a bias correction.[31]

The imputed price for the new computer (model n), from equation (3.9), can be used to estimate a matched model price relative $(1 + \Delta)$ for computer n, or:

(3.10a) $\text{est } (1 + \Delta)_{n,\tau+2,\,\tau+1} = P_{n,\tau+2} / \text{ est } P_{n,\tau+1}$

The notation for the price relative term $(1 + \Delta)$ follows the convention adopted in Chapter II. The numerator is the actual observed price for the new computer, the denominator is its imputed price in the previous period.

Alternatively, one could impute the price in period $\tau+2$ of the machine that disappeared or exited (computer m). The price imputation for exiting computer m in period $\tau+2$ is determined analogously by substituting data for the appropriate time period $(\tau+2)$ in equation (3.8) and data for the exiting computer (model m) in equation (3.9). Then, the estimated price relative for computer m is:

(3.10b) $\text{est } (1 + \Delta)_{m,\tau+2,\,\tau+1} = (\text{est } P_{m,\tau+2}) / P_{m,\tau+1}$

The denominator in equation (3.10b) is the actual price observed for the old computer in period $\tau+1$, and the numerator is its estimated price when computer m is no longer observed. For exits, one must have the hedonic regression for period $\tau+2$, which poses operational problems because the price index for period $\tau+2$ cannot be published until the hedonic regression for period $\tau+2$ has been estimated and analysed. On the other hand, there is asymmetry in imputing for entries and not for exits (or the other way around).

To compute the price index, either the estimated price or the estimated price relative, est $(1 + \Delta)$, can be used in an ordinary unweighted geometric or arithmetic mean formula, as in equation (2.1), along with the matched prices for the other m-1 computers in the sample. Taking as an example the geometric mean formula for basic components and using the imputation for the new computer from equation (3.10a) gives as the hedonic imputation price index:

(3.11) $I_{\tau+2,\,\tau+1} = \{ \prod_i (P_{i,\tau+2} / P_{i,\tau+1})^{1/m} \} = \prod_i (P_{1,\tau+2}/P_{1,\tau+1},\ P_{2,\,\tau+2}/P_{2,\tau+1},\ \ldots\ P_{m-1,\tau+2}/P_{m-1,\tau+1},\ P_{n,\tau+2}/ \text{ est } P_{n,\tau+1})$
1/m

As written, equation (3.11) imputes for the new computer, model n. If one were to impute for the exiting computer, model m, as in equation (3.10b), then the last term in equation (3.11) would be equation (3.10b), rather than equation (3.10a). If one were to impute a price relative for both entering and exiting computers, then both equations (3.10 and 3.10b) appear in equation (3.11), but one must then split the weight between them to avoid over-weighting the computer observation that changed.

---

[31] The imputed price P_{it} incorporates the bias to the estimate of a_0 that was discussed in the preceding section, plus what Teekens and Koerts (1972) refer to as a "transformation bias." Wooldridge (1999, page 202) explains the bias to the predicted price and the reasons for adjusting the estimate (by half the squared stardard error of the regression) for the bias; he illustrates predicting the price of a house from a logarithmic hedonic house price function. Note that there is a close relation between the adjustment for prediction bias and the adjustment for regression coefficient bias—the former uses half the squared standard error of the regression (standard error of estimate), the latter uses half the squared standard error of the estimated regression coefficient.

Compare equation (3.11) with the matched model indexes in Chapter II, for example, equation (2.1). Compare also the imputed price relatives in equations (3.10a and 3.10b) with the actual or imputed price relatives for the conventional quality adjustments shown in table 2.1.

I use the device of the "residual diagram" (Figure 3.5) to illustrate and explain hedonic price imputations. Suppose, solely to simplify the diagram, that the hedonic function is the same in both periods (the examples suggest instead that it must have declined).

Consider first that we impute a price for exiting computer m and that computer m's last pre-exit price was above the hedonic line, as shown by the open "dot" in Figure 3.5—that is, computer m was over-priced, relative to its speed. Any imputed price, of course, lies on the hedonic line (see section III.B, above), so computer m's imputed price in period τ+2 lies on the hedonic surface. Thus, in this case the hedonic price imputation method results in a price decrease (refer to the estimated price relative in equation (3.10b)). This imputation seems intuitively reasonable: The exit of an overpriced machine lowers the average (quality-adjusted) price of computers.

Had computer m been a "bargain" (price lying below the regression line in Figure 3.5), hedonic imputation would have resulted in a rising price index, because the imputed price in period τ+2 (on the regression line) would have been higher than computer m's actual price in period τ+1 (below the line). The exit of a bargain priced computer raises the average (quality-adjusted) price of computers, just as the exit of an overpriced machine lowers the average price level.[32]

In practice, the hedonic regression line will not remain unchanged. Thus, the hedonic imputation will reflect two forces—whether the hedonic surface is lower (or higher) in τ+2 than in τ+1, and whether the exiting computer was over- or under-priced (or neither) in the period before it existed.[33]

Now apply the same exercise to the imputation of computer n, which is shown as introduced in period τ+2 at a price that lies below the regression line in Figure 3.5. However, the new computer's imputed price in period τ+1 is on the regression line. Hedonic price imputation in this case results in a falling price estimate. This imputation makes intuitive sense because the introduction of a new model at a better quality-adjusted price should be recorded as a price decrease. Had the new computer been introduced at a price above the regression line, the hedonic imputation would have recorded an increase. The residual analysis for the new computer, model n, is parallel to the case for the exiting computer, model m.

The Boskin Commission (1996) emphasized the entry of cheaper (relative to their quality) varieties. In the computer context, this has often happened. Quality-adjusted price decline associated with the introduction of new computers was implied by Fisher, McGowan, and Greenwood's (1983) discussion of equilibrium and computer hedonic functions; it motivated the empirical methods employed by Dulberger (1989).

However, newly introduced products might have introductory prices that are more expensive relative to their quality if "newness" itself is desired by some buyers, or if manufacturers take the opportunity presented by new model introductions to incorporate price increases at the same time. Silver and Heravi (2002) provide examples of new introductions that raise the quality-adjusted price. Whether new introductions are associated with rising or falling prices is thus an empirical question, which is discussed further in Chapter IV.

---

[32] Recall the interpretation of the imputation, from section III.B: The regression line gives the mean price of a computer with speed S_m.

[33] Note that if imputations were made for all the computers sold in period τ+2, this would just recover the hedonic regression line in period τ+2, because regression residuals (over and under-pricing) sum to 0 for all the observations.

Conversely, Pakes (2002) suggests that products that exit from the price index sample have price changes that fall, relative to products that survive.[34] This might be so, but it is an empirical proposition. Evidence is reviewed in Chapter IV.

In general, it is a very good idea to estimate both missing prices when possible. Price change can occur from the entry into the commodity spectrum of new models that are cheaper, relative to their quality, than the models that existed before, or from the entry of new models that are more expensive, relative to their quality. It can also occur because some relatively good buys disappear, or because "market shakeout" leads to the exit of models that offer poor price/performance values. Imputing prices for new introductions as well as for disappearances provides estimates of both kinds of price changes.[35]

When indexes are estimated with the hedonic price imputation method, researchers and statistical agencies should beware of imputing from an outdated hedonic function. Because of resource constraints, it is tempting to estimate a hedonic function for one period, say τ, and use it to impute for later periods, say $\tau+1$ or $\tau+2$. Computer prices fall very rapidly. The estimated price based on equation (3.9) will usually be lower in period $\tau+1$ than in period τ, and lower yet in period $\tau+2$; using coefficients for period τ in equation (3.9)—rather than the coefficients for period $\tau+1$—results in an estimated denominator in equation (3.10a) that is too high, and this causes the index to fall too fast (because the price in the numerator is so much lower than the overstated denominator).

Dulberger (1989) introduced imputation hedonic indexes into the hedonic literature on computer prices, using the estimations in equations (3.10a) and (3.10b). Indeed, all the IBM-BEA hedonic computer equipment price indexes (Cartwright, 1986) were estimated by the hedonic price imputation method. BEA called its imputation indexes "composite" indexes, to emphasize that the hedonic estimate was only used for calculating some price relatives, the others being calculated by matching models. Pakes (2003) uses the term "hybrid" index for any form of the hedonic imputation index of equation (3.11), a term that has also been applied by others.

Silver and Heravi (2002) refer to hedonic price imputation indexes (and also to the hedonic quality adjustment method discussed in the following section) by the colourful, if slightly pejorative, term "patching," because they view them as repairing the normal statistical agency fixed sample for "holes"—when items in the sample disappear. However, hedonic imputation methods remain relevant even if the price index is computed on the full universe of computers. If the agency collected the universe of computer prices, it might still prefer to calculate the m-1 price relatives directly, for the statistical reasons noted above, and to make hedonic imputations or hedonic quality adjustments (or "patches") only for exits and new introductions in the universe. Changes in the range of computers available are great, even between periods that are not far apart, so quality change remains a serious problem no matter how large is the sample.

Note, particularly, that hedonic imputation can be useful for outside-the-sample quality change. A statistical agency may know that its fixed sample misses some price change, for the reasons discussed in Chapter IV. It may react by jettisoning the fixed sample, and aggressively augmenting the sample to bring in the new varieties promptly, or as soon after their introduction as possible. However, augmenting the sample will not be fully effective if the new computers are simply linked into the index without taking account of

[34] "Goods which disappear…tend to be goods which were obsoleted by new products, i.e., goods which have characteristics whose market values have fallen. The matched model index is constructed by averaging the price changes of the goods that do not disappear, so it selects disproportionately from the right tail of the distribution of price changes" (Pakes, 2003, page 1).

[35] On the other hand, new introductions often occur outside the sample, so the adjustments outlined in this section, though appropriate, may not be adequate. See Chapter IV.

any price/quality advantages they may have over existing machines. Bringing in the new computers by using hedonic imputations for their previous prices permits a fuller measure of the price impacts of entering computers. Accordingly, handling quality change by hedonic price imputation "patching" remains a viable option even if a universe of prices is available, though patching is by no means the only option.

c. A Double Imputation Proposal

In the hedonic price imputation method, *all observed prices are used*. The price for an entering computer is observed in the period when it enters, τ+2; its observed price is not replaced with an imputation. Similarly, the observed price of an exiting computer in the period before it exits, period τ+1, is not imputed. As explained in the preceding section, an imputation appears in the index *only* when a price is unavailable.

Some researchers have proposed discarding the observed prices for entering (or exiting) computers and imputing both prices. That is, for entering computers one carries out the double imputation:

$$(1 + \Delta)_{n,\tau+2,\ \tau+1} = \text{est } P_{n,\tau+2} / \text{ est } P_{n,\tau+1}$$

Pakes (2003) proposes the same thing for sample exits (forced replacements): His "complete hybrid" index calls for discarding the observed period τ+1 price for exits (computer m) and for imputing the price relative for periods τ+1 and τ+2 from the hedonic function, giving:[36]

$$(1 + \Delta)_{m,\tau+2,\ \tau+1} = \text{est } P_{m,\tau+2} / \text{ est } P_{m,\tau+1}$$

Alternative notations for the preceding two expressions are, respectively: $h_{\tau+2}(n)$ / $h_{\tau+1}(n)$ and $h_{\tau+2}(m)$ / $h_{\tau+1}(m)$, where $h_{\tau+2}(n)$ indicates that the characteristics of computer n are valued with the hedonic function of period τ+2.

Why would one want to discard a "hard" price that has been collected for a price index and replace it with an imputation? Silver and Heravi (2001a, page 7) present one argument for this, Pakes (2003) another. The residual diagram device can be used to explain and evaluate (Figure 3.5).[37]

Consider an entering computer (computer n) that corresponds to the solid point in Figure 3.5: The price of computer n lies below the hedonic regression line, when it is introduced. With single imputation, the index falls, as explained previously, because the actual price for period τ+2 lies below the regression line, while the imputed price for period τ+1 lies on the line—see est $P_{n,\tau+1}$ in Figure 3.5. The index will record an additional price change if the prices of continuing computers fall in reaction and push the whole hedonic function down further, as explained in the previous section.

---

[36] Pakes (2003, page 19): The complete hybrid "is constructed by averaging the observed price relatives for the goods [present in the initial period's sample] which are found in the comparison period with the hedonic estimate of the price relatives for goods which are not." He also suggests what he refers to as a "Paasche-like hedonic index" that would replace *all* of the observed prices in the initial period with hedonic imputations, which would be compared with observed prices in the comparison period, thus, for all observations: $P_{i,\tau+2}$ / est $P_{i,\tau+1}$. This index would obviously include the new computer, n, but not the exiting computer, m.

[37] As before, I suppose for simplicity in drawing the diagram that the same hedonic function existed in two periods. This is inconsistent with the parts of the example that imply that the hedonic function should fall, or rise, but this is initially neglected for clarity. The matter is considered in later in this section.

With double imputation, prices in both periods are estimated from the regression line. Thus when the hedonic function does not change: est $P_{n,\tau+2}$ = est $P_{n,\tau+1}$ and $(1 + \Delta)_{n,\tau+2,\, \tau+1} = 1$. Refer to Figure 3.5. The double imputation procedure, in effect, holds the regression residual constant in the imputations that are used in estimating the index. With double imputation, the price index cannot change from the effects of a new product that is priced above or below the regression line, nor from an exit of an overpriced or bargain product variety.

Of course, even if no other price changes, the hedonic function will not remain unchanged. The entry of computer n will push the regression line down, by an amount that depends on the size of the regression residual in Figure 3.5 and computer n's market share. Suppose the new regression line is the dotted line in Figure 3.6 (where for simplicity, I assume that the regression slope remains unchanged). Double imputation means that prices for computer n are estimated from the regression line in both periods; this in turn means that we record the amount b as an imputed price for computer n, and not the amount Δ_n, the amount we would record with the normal single imputation discussed above.

As the price decline of computer n, do we want to record the amount Δ_n in Figure (3.6)? Or b? That is the same thing as asking whether we want to hold the residual constant in estimating the price relatives in equations (3.10a) and (3.10b).

In one sense, the answer is straightforward from the example. The quantity b is the *average price change in the sample*. In the example, the average price change is composed of "none" for the m-1 matched machines and Δ_n for the machine that changed. Putting in "none" for the m-1 matched machines and "b" for the machine that changed clearly understates the price decline for the new machine (by the difference $d = \Delta_n - b$), and it also understates the average price change (by $w_n d$, where w_n is the weight for the machine that changed).

The same error occurs, incidentally, if one were to impute the price change of computer n from the time dummy coefficient in an adjacent period regression covering periods τ+1 and τ+2, which has also been proposed.[38] As explained in section III.C.1, b equals the coefficient on a time dummy variable, so double imputation approximates using the time dummy coefficient as the imputation for the price change implied by the introduction of computer n. Either way, the price change of the changed or new machine is greatly under-estimated, and *the price index is biased toward the change in the matched model index*. Double imputation forces the imputation hedonic index to agree with the matched model index because the price change that is imputed to the new model is forced to be similar to the price changes for continuing models. Double imputation yields a "hedonic" price index that approximates the IP-IQ (deletion) method of Chapter II, because in both cases the price change for the model that changed is estimated by the price changes for the models that are matched.

An indication of the magnitudes comes from the computer price indexes of Van Mulligen (2003), which are discussed in chapter IV. Van Mulligen estimated both single imputation and double imputation indexes for PC's, notebooks, and servers (in addition to dummy variable and matched model indexes—see the summary in Table 4.7 and the footnote to the table). In every case, the double imputation index lies midway between the single imputation index and the matched model index. For example, for PC's matched model and single imputation hedonic indexes yielded price changes of, respectively, -21.9 percent, and –26.2 percent average annual rates of change. The double imputation hedonic index recorded -24.3 percent per year over the same period, with the same data.

1. Silver and Heravi's proposal. Another part of the answer depends on one's view about what the residual is measuring. The exposition of this chapter has been built on the specification that hedonic

38 Using the dummy variable coefficient for imputation has been explored by Van Mulligen (2003).

residuals are measures of under or over-pricing. This interpretation of the residuals depends on having a well-specified hedonic function (see section III.B.3). In these circumstances, there is no good reason to assume, *contrary to observation*, that computer n was not introduced at a "bargain" price below the regression line. Its actual price at introduction is $P_{n,\tau+2}$; it is not est $P_{n,\tau+2}$.

Perhaps, however, the hedonic function is not well specified. Suppose the hedonic function has missing characteristics, such as the omission of a variable to measure the amount of software that is bundled into the transaction. A negative residual for the new computer might just reflect its containing less software (or a positive residual more software) than the average for other machines in the sample. Holding the residual constant might be interpreted as controlling for the amount of bundled software in the new machine—holding constant, that is, the unobserved omitted variables.

This is one justification for Silver and Heravi's double imputation proposal (which they call "predicted versus predicted"): In their example, the residual, Δ_n is associated with a "missing" brand dummy. If a brand dummy for the seller of computer n were included in a hedonic function, it might be associated with a premium or discount in both periods, but the brand effect would be missed if brand is not included as an explanatory variable. In this circumstance, using the price predicted from the regression for only one period introduces an imputation error, because the brand effect is not held constant between the two periods. [39]

However, the mis-specified hedonic function argument is not a rationale for general application of double imputation. The treatment Silver and Heravi propose relies on specific information on the nature of omitted variables that is not incorporated into the hedonic function, *especially, information that these unobserved variables have remained constant*. Generally, we do not know that. Generally, therefore, use of double imputation introduces error.

Alternatively, perhaps the hedonic residual associated with an exiting computer represents unobserved store amenities; in this case, a replacement computer *in the same outlet* might be associated with the same unobserved store amenities. Holding the regression residuals constant might be interpreted as a mechanism for controlling for unobservable retailing variables, in the same way that the matched model index relies on matching by outlet for controlling unobserved characteristics of the transaction (see Chapter II). Again, however, this is not a rationale for general application of double imputation. It is a rationale that depends on specific additional information on the outlet, information that was not, for some reason, introduced into the hedonic function—so if the replacement computer was in the same outlet we can assume that the unobserved outlet characteristics are constant, and can assume, additionally, that the entire regression residual reflects outlet variables.

Imputing the price of computer n (or m) from the regression line for the period when computer n (or m) is not observed rests on the statistical principle that the predicted value from the regression line is the best estimate for a computer that has the characteristics of computer n (or m), in the period in which this computer is not observed. [40] However, the regression line does not provide the best estimate of the price when the price of computer n (or m) is actually observed. The observed price is always more accurate than the imputed price.

In general, the double imputation procedure introduces additional estimation variance into the index, and very possibly introduces bias. For these reasons it should be avoided, unless the investigator has reason to believe that omitted variables, or store amenities, or some related reasons, account for the residual,

[39] This paragraph draws on conversations with Mick Silver.
[40] See the discussion of prediction in section III.B and in Gujarati (1995), pages 137-139.

and that the omitted variables have not changed. This is a demanding condition that is not, for the most part, met in the data sets normally used for estimating hedonic price indexes.

However, it might be justified for specific applications of hedonic imputations within an agency's price index. For example, the pricing agent may have information about specific changes within the retail outlet, or the agency's commodity analyst may know that the brand of the entering computer is associated with a smaller amount of unmeasured software than is present for other computers, and may know that this accounts for the negative residuals shown in Figures 3.5 and 3.6. One always wants to take into account all the information available on quality change, and to resist mechanical application of any statistical procedure. This same principle means that general application of double imputation should be avoided, it should be utilized only when specific information that is not incorporated into the hedonic function applies to a specific case.

Double imputation assures that entries and exits from the sample carry with them no price changes. We know that they often do.

2. Pakes' proposal. Pakes (2003) has independently made a similar double imputation proposal, though he proposes to impute for sample exits, only. His reasoning begins from a different point.

Considering the buyer of discontinued computer m, Pakes seeks the amount that would compensate that buyer for the disappearance of the computer bought in the initial period. He reasons that a bound on the compensation amount is the change in the hedonic function between periods $\tau+1$ and $\tau+2$, evaluated at the bundle of characteristics embodied in the exiting computer, model m.[41] That gives est $P_{m,\tau+2}$/ est $P_{m,\tau+1}$, as noted above. If the hedonic function is unchanged, no compensation is necessary and the price relative is unity (see Figure 3.5); if the hedonic function changes, this compensation is the amount b in Figure (3.6).

Pakes does not discuss the residuals. Based on seminar presentations of his paper, the text of the revised paper, and his earlier drafts, I think Pakes would emphasize two hypotheses about the residuals. First, they may measure omitted characteristics. This point has already been addressed. It is valid but should be confronted by improving the hedonic specification where possible, and does not make a case for *routine* application of double imputation.

Second, Pakes points to differences in preferences among buyers: Some individuals value a particular bundle of characteristics more than others. Under or overpricing, relative to the average valuations incorporated into the hedonic function, reflects pricing strategies of oligopoly sellers who exploit market niches that arise because of these differences in buyers' preferences. This is the right way to look at the theory (see the theoretical appendix), and Pakes is right to stress markups.

Taking these two points together, what appears to be an overpriced computer (a positive hedonic residual), is not necessarily judged overpriced by the buyers who choose it. Even if there are no unobserved characteristics that induce buyers to choose computer m, their valuation of its characteristics may influence them to choose it because their tastes differ from buyers who choose a similar, but not identical specification.

But this does not yield a clear prescription for how to handle the price change that occurs when computer m exits. If computer m was a bargain in period $\tau+1$, compensation by the shift in the hedonic function will be insufficient.

[41] He goes on to note the standard cost-of-living index result that this compensation is too large if consumer substitution is feasible, but I ignore the COL index question for present purposes.

Suppose computer m's price lay above the hedonic regression line in period τ+1, because it served a market niche that valued its characteristics bundle more highly than did other buyers. Imputing computer m's price from the new hedonic function for period τ+2 obviously ignores whatever it was that made computer m unique for its period τ+1 buyers. If computer m's uniqueness lay in its constellation of *observed* characteristics, and if it was truly not available in period τ+2, then an imputation from the regression line does not compensate. If computer m contained unobserved characteristics that accounted for its hedonic residual, then we cannot be sure that imputation from comparable a point on the period τ+2 hedonic function implies the same unobserved characteristics.

3. Summary: Double imputation. This matter is not settled, and depends, as noted in the introductory paragraph in this section, on one's interpretation of hedonic residuals. If the hedonic function is correctly specified, then it seems incontrovertible that double imputation creates error. If the hedonic function is not correctly specified—if the residuals measure the effects of omitted variables (Silver and Heravi) or differential markups because sellers of unique varieties have different amounts of market power (Pakes)—then one could justify double imputation in circumstances where the researcher knows something about the omitted variables for individual cases that are to be imputed. It seems harder to justify routine application of double imputation. The remedy for incorrectly specified hedonic functions is improving the specification, and it is not clear that routine application of double imputation can correct for specification errors.

It is worth noting that Griliches (1961) pointed to the interpretation of hedonic residuals, so this question, also, is not a new one in the hedonic literature. He suggested research on residuals that, though not proposed in the context of the present discussion, would be useful in evaluating the double imputation proposals. Griliches suggested that if residuals were evidence of under or overpricing, then they should predict changes in market shares. If they do, then this is evidence that the hedonic function is correctly specified, that the residuals are measures of under and over pricing, and double imputation is therefore inappropriate. On the other hand, if residuals reflect omitted characteristics, as the proposals for double imputation suggest, then they should not be associated with changes in market shares. Research on hedonic residuals and changes in market shares would clarify the issues. If the residuals are associated with differential markups, then the methods pioneered by Berry, Levinsohn and Pakes (1995) may be used to analyse the economic interpretation of hedonic residuals, and accordingly their appropriate treatment in price indexes. This is a complicated research effort that has hardly begun.

4. The Hedonic Quality Adjustment Method

An alternative imputation method takes the form of estimating a quality hedonic adjustment, rather than imputing the missing price. The hedonic quality adjustment method has properties that distinguish it from other hedonic methods.

Dummy variable and characteristics price index methods imply, as noted in previous sections, that the database for the hedonic function and the database for the price index must be the same, and that the hedonic function must be estimated with the same timeliness as the index. As it has actually been implemented in work on computers, the hedonic price imputation method also uses the same database, because otherwise the imputations in equations (3.10a and 3.10b) might differ in some manner from the actual observed prices in the index—in the discounts, for example. Dulberger (1989) constructed an imputation index from the same data she used to estimate the hedonic function.

The hedonic quality adjustment method makes it possible to estimate the hedonic function from a database that is different from the database for the index itself. In this, it is unique among hedonic methods. The hedonic function database may be larger and it might be drawn from a completely different source. It may also refer to a different period from the month or quarter for which the index is published. These are practical advantages in statistical agency environments, and are the reasons why the hedonic quality adjustment method is the most widely used hedonic method within statistical agencies—as examples, the U.S. CPI (Fixler, Fortuna, Greenlees, and Lane, 1999), PPI (Holdway, 2001), and the French national accounts computer deflator (Lequillier, 2001), and the British PC index (Ball, et al., 2002).

a. The Method Explained

The estimator for the hedonic quality adjustment method has already been presented in the section "estimating price premiums for improved computers" (section III.B.2). As an example, suppose that computer n was 10 per cent faster than the computer it replaced (computer m) and that it had 15 per cent more memory. The hedonic coefficients (see equation (3.8) or (3.1)) can be used to estimate a hedonic quality adjustment, (A(h)), which is:

$$(3.12)\qquad A(h) = \exp\ \{(a_1[(\text{speed})_n / (\text{speed})_m) + (a_2(\text{memory})_n / (\text{memory})_m)\}$$

Taking for illustration the regression coefficients of equation (3.1), the additional speed adds approximately 9 per cent to the price of a computer (antilog .783 x 1.10) and the additional memory adds about 3 per cent (antilog .219 x 1.15). Computer n is thus worth approximately 12 per cent (1.09 x 1.03) more than computer m. Suppose that computer m cost €4000; the quality adjustment in monetary terms is then €490 (4000 x 0.12227). This implies that the estimated price for computer n in the period before it was introduced (period τ+1) is €4490. Thus, we have:

$$(3.13)\qquad \text{est } P_{n,\tau+1} = P_{m,\,\tau+1}\,(A(h)),\ \text{and}$$

$$\text{est } (1+\Delta)_{n,\tau+2,\,\tau+1} = P_{n,\tau+1} / \text{est } P_{n,\tau+1}$$

For nonlinear hedonic functional forms, including the double log form, the value of the hedonic quality adjustment depends on which point on the hedonic function is chosen for evaluation.

Equation (3.12) provides a hedonic quality adjustment—A(h)—that makes one of the computers equivalent to the other. It does not really matter which one, but it is customary in many statistical agencies to reduce the price of the new computer to make it comparable to the old one rather than the other way around. This requires that equation (3.12) be computed from coefficients from a regression using period τ+1 data, the period before computer n was introduced. For example, Bascher and LaCroix (1999, their section 1.1.1) write the hedonic quality adjustment method as applied in France as (translating into the notation used in this section):

$$(3.12a)\qquad \text{est } P_{n,\tau+1} = P_{m,\tau+1} + \{\exp\,[a_1S_n + a_2M_n] - \exp\,[a_1S_m + a_2M_m]\}$$

Thus, they estimate a previous period price (they say base period) for the new computer by quality adjusting the price of the old computer for the difference between the characteristics of the new and old computers.

Applying the hedonic quality adjustment to the price of the new machine, the adjusted price becomes the estimated price for computer n in period τ+1(in the example, €4490). It is used as the estimated price in the denominator in the price relative of equation (3.13), that is, est $(1 + \Delta)_n = P_{n,t+2}/$ est $P_{n,t+1}$. Compare equation (3.13) with the corresponding equation (3.10e) of the preceding section.

The price index is then calculated conventionally, by using the quality-adjusted price relative (3.13) in the matched-model formula of equation (2.1). That is:

(3.14) $I_{\tau+2,\,\tau+1} = \{ \prod_i (P_{i,\tau+2} / P_{i,\tau+1})^{1/m} \} = \prod_i (P_{1,\tau+2}/P_{1,\tau+1},\ P_{2,\,\tau+2}/P_{2,\tau+1},\ \ldots\ P_{m-1,\tau+2}/P_{m-1,\tau+1},\ P_{n,\tau+2}/\ \text{est}\ P_{n,\tau+1})^{1/m}$

Equation (3.14) is identical to equation (3.11) from the previous section. They differ only in the way the denominator of the final term is estimated, that is whether est $P_{n,\tau+1}$ is imputed directly from the hedonic function for period τ+1 (equation 3.11) or is imputed by applying a hedonic quality adjustment derived from the hedonic function for period τ+1 to the price of the exiting machine, computer m (equation 3.14).

The hedonic quality adjustment method and the hedonic imputation method are equivalent *when they are applied to the same data*. Though I return to this in a later section, the point needs emphasis here because it has often been overlooked in recent criticism of the hedonic quality adjustment method, for example, Schultze and Mackie (2002) and Pakes (2003).

One might apply the hedonic quality adjustment to the price of the exiting computer, computer model m, by adjusting its price for its lower level of characteristics to make it comparable to the price of computer n in period τ+2. Hoven (1999, section 3.2) suggests that adjusting the old price is the standard Dutch practice. This is exactly parallel to the calculations described in equations (3.12-3.14), except that the adjusted price est $P_{m,\tau+2}$ becomes the numerator in the price relative of equation (3.13), the actual price $P_{m,\tau+1}$ being the denominator. The price index is calculated as before, by using the estimated price relative for computer m in equation (3.14).

To quality adjust the exiting computer, one needs coefficients from a regression for period τ+2 in equation (3.12). This often poses operational difficulties because the regression may not be available in time to calculate the quality adjustment that pertains to period τ+2. In contrast, estimating the previous period's quality-adjusted price for the entering computer (that is, the denominator of equation 3.13) implies having the period τ+1 regression available before period τ+2, which is inherently feasible. Thus, quality adjusting the price of the new computer not only fits in with normal agency practice it also has operational advantages in providing time for preparing the hedonic function.

b. Diagrammatic Illustration

The hedonic quality adjustment is shown diagramatically in Figure 3.7, where, as before, a one-variable hedonic function is shown, in order to draw a two-dimensional diagram.

As an initial example, suppose that computer m was overpriced relative to its speed, as shown by the open "dot" in the figure. Suppose, additionally for simplicity, that computer n was introduced at the price indicated by the solid "dot" in Figure 3.7 (its price lies on the hedonic line).

The value of the speed difference between computers m and n, as determined by the hedonic function for period τ+1, is designated as A(h) in the figure. Then, the quality-adjusted difference between the prices for computer m and computer n consists of an adjustment component (the term A(h)), plus the price decrease shown as the difference between the adjusted price for computer n and the actual price for computer m, marked "B" in the figure. The total price change between computer m and computer n (which is $\ln p_n - \ln p_m$) can thus be decomposed into (-B + A(h)). Of this, +A(h) is the quality adjustment for the improvement in computer n's speed, and –B is the price change that goes into the index.

If computer m's price had been below the hedonic line, the B term would have been a price increase (because a bargain exited from the market). This situation is depicted in Figure 3.8. In this case, the difference between P_n and P_m equals +A(h) +B. The quality adjustment term A(h) is removed from the total change between the prices of computers m and n, leaving B as the price change (an increase in this case).

In these two examples, computer n's introductory price lies on the hedonic line. Suppose, however, that computer n had a lower introductory price (below the regression line, for example). Then, as shown in Figure 3.9, the total price decline becomes –A-B-C, where C is the amount that computer n lies below the regression line, and A and B are defined as before. In this case, there are three parts to the price decline: The exit of an over-priced machine (B), the entry of a bargain-priced machine (C) and the value of the improvement between the two machines' specifications (A). Obviously, there are corresponding cases where computer n is introduced at a higher price than the regression line, computer m exited as a bargain, and so forth, which are not illustrated but follow the same principles.

Comment. The U.S. Bureau of Labor Statistics hedonic indexes for computers (Holdway, 2001) use the hedonic quality adjustment method. They combine a linear form of equations (3.8) and (3.12) to obtain the quality adjustment and estimated price relative of equation (3.13). In the case of the Producer Price Index, the index number formula is the average of price relatives (AR), rather than the geometric mean formula of equation (3.14), but otherwise the procedures in this section apply to the BLS computer indexes, for the most part. In the BLS implementation, hedonic functions are re-estimated roughly three times a year because the coefficients change rapidly. The U.S. CPI also uses the hedonic quality adjustment method for most of its hedonic indexes (see Fixler, Fortuna, Greenlees and Lane, 1999, and the summary in Schultze, 2002). Similar employments of the method are the PC price indexes computed by Statistics Canada (Barzyk and MacDonald, 2001), INSEE, and ONS (Ball, et al., 2002).

Ideally, equations (3.8) and (3.12) should pertain to period τ+1 or to period τ+2. It is common practice to estimate the hedonic regression for some previous period (perhaps τ or τ-1) and use it for making quality adjustments for several periods before re-estimating the equation. When this is done, the potential for error from using an outdated regression may be considerable. An extra 100MHz of speed, for example, cost far more in 1998 than it does presently, so making a quality adjustment from an old regression over-adjusts for the current value of the quality change, and biases the index downward. The same point was made above in the discussion of the imputation index. The Schultze and Mackie (2002) report criticized the BLS CPI hedonic indexes for appliances because BLS had insufficient plans for updating the coefficients.

The hedonic quality adjustment method requires that coefficients be estimated precisely. The method is subject to bias from omitted variable problems, and to mis-specification of the hedonic function. In these cases, the coefficients used in equation (3.12) may be biased and will give incorrect quality adjustments. Omitted variable and mis-specification problems, however, affect all hedonic methods, and not uniquely the hedonic quality adjustment method (as has sometimes erroneously been suggested). These matters are discussed in Chapter V.

c. Empirical Comparisons with Conventional Methods

The hedonic quality adjustment method permits ready comparison with other quality adjustments that might be employed in the index, which is a major advantage for analytical purposes, and yields quantitative assessments. For example, Schultze (2002) tabulates and discusses BLS data that compare hedonic and conventional quality adjustments in a number of CPI index components, and Moulton, Lefleur and Moses (1999) performed a comprehensive comparison for TV sets. In the U.S. CPI, quality changes are mostly handled by direct comparison or by the deletion (IP-IQ) method.[42] The BLS studies summarized by

[42] Actually by the "class mean" variant of the deletion (IP-IQ) method. See Chapter II.

Schultze replaced direct comparisons or IP-IQ treatments with hedonic quality adjustments, and he compared the resulting indexes. In some cases, the hedonic-adjusted index shows less price increase, in some more, but in most the differences were not great. Moulton, Lefleur and Moses (1999) reported that their simulated CPI TV index that used hedonic quality adjustments for sample replacements declined by 13.5 percent, over the 1993-1997 interval, which was very close to their "simulated" CPI comparison index (which declined 13.2 percent).

The Triplett and McDonald (1977) study has already been mentioned: They applied a hedonic quality adjustment corresponding to equation (3.12) to each quality change that occurred in the U.S. PPI refrigerator index, holding everything else in the price index compilation the same, and compared the result with the quality adjustments that had actually been made in the published index. In the U.S. PPI, quality changes are handled either by direct comparison (when the difference in quality between new and old machines was judged small) or by link-to-show-no-change.

Silver and Heravi (2003) carried out a similar study, only using scanner data. They calculated a benchmark index from the scanner data that replicated normal agency practice for replacements; this matched model index fell by 9.1 percent over the period they studied. They compared this matched model index with several different hedonic indexes computed from the same data, including a version of the hedonic imputation method, and one using the hedonic quality adjustment method. The hedonic indexes mostly fell somewhat less than the matched model index—8.6 percent for the imputation index, 8.8 percent for the hedonic quality adjusted index, but all the hedonic indexes were relatively close together.[43] Silver and Heravi (2003) are the only authors, to my knowledge, who have compared hedonic imputation and hedonic quality adjustment methods on the same data—indeed, they compare all four forms of hedonic methods with alternative forms of matched model indexes, all using the same database. None of their hedonic indexes are very far apart (see table 3.3), and all fall somewhat more slowly than a matched model index computed from the same data.

d. Criticism of the Hedonic Quality Adjustment Method and Comparison with Hedonic Imputation Method

Schultze and Mackie (2002) have criticized the hedonic quality adjustment method, as has Pakes (2003). The issues need sorting.

1. Database. There is first the difficult question of databases. Database issues have not always been fully considered by critics.

The hedonic quality adjustment method is designed to be used in situations where the database for the index is inadequate for estimating the hedonic function, which therefore must be estimated from a different database. First, normal price index samples are frequently too small to estimate accurate hedonic functions (though there are exceptions). Second, in the normal statistical agency environment, there is not sufficient time to estimate a hedonic function and have one of the "direct" hedonic price indexes ready in time for an ongoing monthly price index program.

The hedonic quality adjustment method was designed to surmount these two difficulties. As implemented by the BLS, for example, one estimates the hedonic function beforehand so coefficients are

[43] I used their "geometric mean" versions of each, as tabulated in their Table 9.7. It is interesting that both Silver and Heravi (2003) and Moulton, Lefleur and Moses(1999) report that the price index for characteristics gave larger differences from the matched model index than the index that applied hedonic quality adjustments to forced replacements. This suggests that hedonic adjustments for forced replacements inside the sample are outweighed in significance by outside-the-sample price changes that are missed by the matched model method. See the discussion in Chapter IV.

available for making a quality adjustment when a forced replacement takes place. The hedonic function is estimated from a different database from the one that is used to calculate the index. Indeed, the hedonic quality adjustment method was designed to permit maximum flexibility in choice of databases.

If an agency were to implement any of the other hedonic methods (including the hedonic imputation method), the database for the hedonic function and for the index must match. This usually implies changing price collection and computational methods. Changing traditional methods should not be an issue that is out of bounds. One might believe that traditional collection methods are flawed, and that alternative price collections (such as scanner data) might be better. One advantage of scanner data, which might weigh more heavily in some judgements, is that scanner data permit more alternatives for implementing hedonic indexes. But it is a separate issue (see the following section).

Critics of the hedonic quality adjustment method, particularly in the U.S., have mostly ignored the database matter, and proceeded instead with econometric and statistical estimation contentions, which are discussed in the next subsection. The statistical properties of alternative implementations of hedonic indexes are important, but they cannot be considered independently of the choice and availability of databases. If the database for estimating the hedonic function must be different from the one used for the index, the hedonic imputation method is not a feasible substitute in practice for the hedonic quality adjustment method.

Thus, it seems useful to proceed in the following way: First, distinguish circumstances where the hedonic imputation method and the hedonic quality adjustment method can be implemented on the same database. An example is when the CPI sample contains a sufficient number of observations to estimate a hedonic function, as in Moulton, Moses and Lefleur's (1999) TV indexes from the U.S. CPI database, or Van Mulligen's (2003) price indexes estimated from the database for the Dutch CPI. This same database case is considered in subsection 2, to follow. I conclude that when the hedonic quality adjustment and hedonic imputation methods are estimated from the same database there is no conceptual or statistical reason for preferring one over the other, contrary to much that has been written recently.

Then I consider cases where that is not possible, or not desirable. In this case, statistical agencies estimate the hedonic function from sellers' website data or market information datasets on computers, and use the results in an index that uses prices that are directly collected for that purpose. This case matches the U.S. PPI, Statistics Canada, and the U.K. price indexes for personal computers. It raises a different set of statistical issues that have not figured very prominently in recent discussions and criticisms of hedonic indexes. These issues are reviewed in subsection 3.

2. Multicollinearity and the precision of the adjustment. All hedonic price indexes are subject to problems arising from multicollinearity, which arises when independent variables in regressions (characteristics in hedonic functions) are highly correlated. Multicollinearity in hedonic functions is defined and discussed in Chapter VI. An extreme case is where multicollinearity, usually in conjunction with measurement errors of some sort, results in negative values for one or more coefficients that ought, on a priori grounds, to have positive coefficients. This problem is well known in the hedonic literature, it is not a new point, either for hedonic price indexes generally or for the hedonic quality adjustment method in particular.

Pakes (2003) criticizes the hedonic quality adjustment method, as implemented in the BLS PPI and CPI computer price indexes on mulicollinearity grounds:

> "…characteristics are typically highly correlated…and this produces negatively correlated regression coefficients. Consequently the weighted sums of coefficients used to predict comparison period prices will be estimated more precisely than the individual coefficients used for the BLS's [hedonic quality adjustment method]" Pakes (2003, page 32)

The criticisms in Schultze and Mackie (2002) are also tied to multicollinearity—see their chapter 4. Both Pakes and the Committee on National Statistics group prefer the hedonic imputation method. Some of the criticisms in both sources pertain to idiosyncratic implementations by the BLS. My main focus in the following is on the method, not particular applications of it.

The econometric point is illustrated with equation (3.9), reproduced below:

$$\text{est } P_{n,\tau+1} = \exp \{a_{0\tau+1} + a_{1\tau+1}\,(\text{speed})_n + a_{2\tau+1}(\text{memory})_n\,)\}$$

Even if the hedonic coefficients a_1 and a_2 are not estimated very precisely, the imputed price, est P, may be estimated with low variance, if the overall fit of the regression is good. On the other hand, using the OLS estimate of either a_1 or a_2 by itself as an adjustment may not produce a very accurate estimate.

However, the hedonic quality adjustment method, correctly applied, uses all the coefficients, not just one coefficient, or a subset of them. Because of this, the hedonic quality adjustment method and the hedonic imputation method yield the same result, *when applied to the same data*. This seems not adequately appreciated by the critics.

Using the definition of the hedonic quality adjustment, A(h) from equation (3.12), one can obtain an estimated price for the replacement variety, computer n, in the period before it was introduced by adjusting the price of the exiting computer (computer m) for the value of the difference in the two machines' characteristics:

$$\text{est } P_{n,\tau+1} = P_{m,\tau+1}\ \{\ A(h)\ \}$$

$$= P_{m,\tau+1}\ \{\exp\,((a_1[(\text{speed})_n / (\text{speed})_m) + (a_2(\text{memory})_n \,/\, (\text{memory})_m))$$

$$= \exp \{a_{0\tau+1} + a_{1\tau+1}\,(\text{speed})_n + a_{2\tau+1}(\text{memory})_n\,)\}$$

$$= \text{equation (3.9)}$$

when $P_{m,\tau+1}$ lies on the regression line (which results in the simplification in the third line). Thus, when the hedonic imputation and hedonic quality adjustment methods employ the same data, the imputations derived from each are computationally and statistically equivalent.

Because the hedonic quality adjustment method uses all the coefficients, it is equivalently a price imputation for the adjusted product variety, as explained earlier. Equation (3.12) is an adjustment that makes the price of computer m equivalent to that of computer n in period τ+1. It is an adjustment for all of the characteristics, not just one or a subset of them.

The point may be clarified with the example similar to the one originally used to illustrate equation (3.12), above. Suppose that the new computer was ten percent faster than the computer it replaced, but that its memory size was the same. Following the illustrative calculations for equation (3.12), the new computer's additional speed adds 9 percent to its price (antilog .219 x 1.15, as before). But memory change is zero for the purpose of calculating the hedonic quality adjustment in equation (3.12), so there is no additional adjustment for memory size.

It might appear that the hedonic quality adjustment in this case uses only the coefficient for the speed variable, so it is subject to the multicollinearity objection. However, the coefficient for memory is still used in the hedonic quality adjustment, it is just that the memory change is zero, so the coefficient contributes nothing to the adjustment. Applying the hedonic quality adjustment through equation (3.12) yields the same imputed price as imputing the value of the new machine directly through equation (3.9).

The general "multicollinearity" criticism of the hedonic quality adjustment method seems to have taken a problem that affects all hedonic methods (including the hedonic imputation method) and directed it at the hedonic quality adjustment method. Correctly applied, the hedonic quality adjustment method is no more subject to the multicollinearity objection than are other hedonic methods.

Comment: BLS implementation. Nevertheless, BLS procedures for hedonic computer price indexes do not escape the multicollinearity criticism, because BLS uses a subset of the hedonic coefficients, not all of them. The BLS hedonic function contains a speed coefficient, but BLS does not use the speed coefficient from its hedonic function to adjust computer prices for speed changes. Instead, it uses information on the cost of the computer microprocessor chip as an adjustment (Holdway, 2001). To the extent that multicollinearity between the ignored speed variable affects some other coefficient that is used in the BLS hedonic quality adjustment process, error might result.

In general, one should not implement the hedonic quality adjustment method by employing single coefficients, or subsets of coefficients when multicollinearity among variables is strong. In these cases, it may be possible to adjust for the joint effect of the collinear variables. For example, if it is known (from the regression diagnostics, for example) that the size of the hard drive and the size of the main memory are highly correlated in the sample, then one should always use the combined coefficients from the two variables as a quality adjustment for the combined effect of the correlated variables. Normally, when confronted with a "wrong" sign on a coefficient, the investigator will not make an adjustment that obviously goes in the wrong direction; the combined values of the coefficients of the collinear variables may provide an acceptable quality adjustment.

A final observation is relevant. It is striking that critics of BLS procedures provide no quantitative evidence. The multicollinearity criticism of the hedonic quality adjustment method depends for its force on the presumption that hedonic functions have high multicollinearity. In Chapter IV, I show that multicollinearity is rather low in the BLS hedonic functions for computers. Secondly, the important question is: How much difference does a departure from best practice make to the index? Does the BLS' particular implementation of the hedonic quality adjustment method yield a different price index from the index that the hedonic imputation method yields? The matter cries out for an empirical assessment, which is wholly absent from recent criticisms, including the Committee on National Statistics Panel's review (Schultze and Mackie, 2002).

3. Combining different databases: conditions and assumptions. The hedonic quality adjustment method normally is implemented by estimating the adjustment from one database and applying it to another. For example, in the U.S. BLS hedonic computer indexes, the database for estimating the hedonic function is compiled from publicly-available data on computer sellers' internet sites. The adjustments are then applied to adjust actual prices collected for the PPI, the CPI and international price indexes for changes in characteristics encountered in computers in the BLS samples. U.K. personal computer price indexes are estimated similarly, as is the Statistics Canada index.

Even though combining databases it the method's strength, it is necessary to consider the conditions under which it is legitimate to estimate a hedonic quality adjustment from one database and transfer it to an index that is computed from another database. The issues are most transparently posed if we assume a linear hedonic function, under which indeed the problems are greatest.

Suppose the linear hedonic function is estimated from manufacturers' list prices or on prices collected from manufacturers' internet sites, both of which are common data sources for cross-section computer prices used in hedonic functions. Using such data in a linear hedonic function yields the price of speed in dollars or euros, so an incremental MHz might be estimated (from the hedonic function) to cost, say, $1.19 (which was the actual BLS estimate in October, 2000). This implicit price does not reflect any discounts that are offered—it is the list price of MHz, not the transaction price of MHz, or it may be the transaction price for direct internet sales or the price to resellers, but it is not necessarily the price for sales through other distribution channels. Even if the price quoted on sellers' websites is a transactions price for single unit internet sales, the PPI may incorporate discounts for 10, 25 or 100 computers sold at the same time. The implicit price of MHz to volume buyers will be lower than to single computer buyers, so the estimate from the hedonic function will be too large to apply to these volume sales.

The class of problems discussed here also reverberates to another question: Can one make quality adjustments to both the PPI and the CPI with estimates from the same hedonic function?[44] The appropriate considerations are similar to those discussed immediately above. If manufacturers' selling prices, the database for estimating the hedonic function, are prices charged to resellers, the $1.19 per MHz price is a wholesale price, not a retail price. Using $1.19 as a quality adjustment for an additional MHz in the CPI understates the retail price of MHz. A similar point applies to hedonic functions estimated from retail prices collected by market information firms: They typically report the average selling price of a computer, aggregated across sellers. This average may differ from the retail price collected for a particular outlet in the CPI (it might be higher or lower). The estimated price per MHz might therefore be too high or too low as a quality adjustment for a replacement computer in the CPI.

These are certainly important and difficult points. By reciting them, I do not mean to imply they have not been considered by statistical agencies that have implemented the hedonic quality adjustment method. For example, in the BLS implementation, the quality adjustments are themselves adjusted to value them at retail and wholesale prices, for the CPI and PPI, respectively (see Holdway, 2001). This revaluation of the hedonic quality adjustments is parallel to revaluation of other quality adjustments. Manufacturers' production cost, for example, is collected by BLS at the manufacturers' level and marked up to retail by BLS.

Revaluation problems are less severe with other hedonic functional forms, under reasonable assumptions. For the double-log form (equation 3.8), the regression coefficient on MHz estimates the elasticity. For the semi-log form, the coefficient estimates its percentage contributions to the price of the computer. Suppose the retail markup on computers does not depend on the *proportion of MHz and MB in individual computers contained in equation (3.12),* or suppose that discounting does not depend on this proportion. Then, the percentage change in MHz and MB estimated from equation (3.8) can be applied to either the CPI or the PPI. Refer to the numerical illustration in section C.3.a, above: If a ten percent increment to speed adds 9 percent to the price of the computer, it adds 9 percent to the wholesale price and 9 percent to the retail price. Similarly, ten percent additional memory adds 3 percent to both wholesale and retail prices.

To put this another way, consider estimating two versions of equation (3.8), one with the left-hand side variable the retail price, the other using the wholesale price. Under the assumption that retail markups and discounting do not depend on the proportions of the characteristics, the two regressions would be the

---

[44] A more tangled issue is discussed in Chapter VI and the Theoretical Appendix: Does the "resource cost" quality adjustment criterion generally accepted for the PPI require a different quality adjustment from the "user value" criterion generally accepted for the CPI? For reasons discussed in the appendix, I believe that generally the same hedonic information can serve for each, subject only to the level of the adjustment (the problem discussed in the present section). The conceptual issues are, however, complicated. I have discussed this class of questions in Triplett (1983 and 1987).

same except for the intercept term, a_0. The difference in intercept terms between the two regressions measures the discounts or markups. However, the intercept term does not appear in the hedonic quality adjustment in equation (3.12). Thus, the quality adjustment would be the same (in percentage terms or elasticity terms) whether the regression was run on wholesale or retail prices, or on list or transactions prices.[45] This is one of the advantages of non-linear functional forms. Additional discussion occurs in Chapter VI.

On the other hand, the crucial assumption might be wrong: Larger computers might be discounted more heavily than smaller ones. If differential discounting twists the price structure in this manner, it would bias the estimate of the hedonic price of MHz (and of MB). If low margin computer sellers concentrate, say, on high-end machines, this will also twist the price structure and potentially bias an adjustment applied to computers sold by sellers of smaller machines who charge closer to list prices.

The only thing that can be said about this class of potential problems is that researchers need to be aware of them and to check the accuracy of their estimates against any other information they may have available. Other methods are not immune to similar difficulties—differential discounting will affect any hedonic index.

4. Residuals. For the hedonic imputation method, the interpretation of residuals is an issue, which was reviewed in an earlier section (III.C.3.c). For the hedonic imputation method, the residual issue involves est $P_{n,\tau+2}$/ est $P_{n,\tau+1}$ (or the comparable estimation for computer m). That is, the interpretation of residuals came up because the imputation required comparing the *same characteristics bundle on two hedonic functions*.

We need to consider whether there is a parallel issue for the hedonic quality adjustment method, though I do not know that anyone has actually raised this issue. For the hedonic quality adjustment method, we compare estimates for two *different characteristics bundles on the same hedonic function*, that is: est $P_{n,\tau+1}$ / est $P_{m,\tau+1}$ (or the equivalent for the hedonic function for period τ+2). The hedonic quality adjustment method is neutral with respect to the residuals.

5. Evaluation of database and computation choices. Although the choice of price index computation and collection methods is generally beyond the scope of this handbook, something must be said.

Many statistical agencies presently have no price indexes for computer equipment, or judge their present indexes to be inadequate. Many are also exploring alternative data sources to the usual direct collection of prices, scanner data, for example (see the volume on this subject by Feenstra and Shapiro, 2003). Others are exploring proprietary datasets, such as from the International Data Corporation (IDC); examples are the Australian Bureau of Statistics (ABS), the French statistical agency (INSEE), Statistics Netherlands, and Eurostat (Konijn, Moch, and Dalén, 2003).

For some agencies, there is a trade-off between proprietary datasets, where the agency has less control over collection and data quality issues, and more expensive direct collection. If in their judgements, their database needs are met by proprietary datasets, the choice among hedonic methods is accordingly larger, because the index dataset and the hedonic function dataset can be the same. For example, ABS has contemplated using proprietary data, and computing a time dummy variable hedonic index from the proprietary data.

---

[45] The dollar or euro value would of course be different, because the percentage adjustment would be applied to a different level of prices. See the numerical example in section III.C.3.

On the other hand, if an agency feels that its price collection methods for computer or ICT equipment are optimal for its needs, but that quality adjustment is the issue, then they typically wish to find a hedonic method that suits their price collection strategy. Obviously, this is again a cost-benefit calculation, which weighs the data quality assurance that traditional collection methods are thought to provide against the relative costs and accuracy of alternative collections. These questions cannot be confronted in this handbook. Statistical agencies have made varying judgements, based on their own experiences, resources, and their assessments of the relative strengths and costs of the alternatives.

On the relative strengths of the alternatives, compilers of proprietary datasets do not always have price index purposes in mind, so comparability of observations over even adjacent periods does not always receive high priority (Evans, 2002, discusses this problem and explains how the French resolved it). Also, matching observations by model numbers or nomenclatures (sometimes the only price information available in proprietary datasets) invites undetected quality change in the matches as sellers upgrade ICT equipment without necessarily changing model numbers. This, of course, can also be a problem for conventional price collection.

D. THE HEDONIC INDEX WHEN THERE ARE NEW CHARACTERISTICS

The four versions of the hedonic index that are reviewed in this chapter all make allowances or adjustments for quality changes that can be represented as increases or decreases in the quantities of the characteristics that were available in both periods. For example, if a CD/RW drive was included in the price of some computers in the period τ+1, as well as in some computers in period τ+2, then one can estimate an implicit price for inclusion of a CD/RW drive in both periods; the implicit price can be used for the hedonic quality adjustment method if the replacement and the machine that exited the sample differed in having or not having a CD/RW drive. Alternatively, one could estimate a dummy variable index in this case, since the CD/RW drive was present in both periods.

Sometimes, however, a new characteristic arrives in the second period. At some point in the past, no computer had a CD/RW drive. If the CD/RW drive was present in period τ+2 but was not present in period τ+1, no period τ+1 implicit price is available to carry out the hedonic quality adjustment method or the hedonic imputation method. The hedonic quality adjustment method requires valuing the characteristics in which computers n and m differ, and so it requires the implicit price of the CD/RW drive in period τ+1. Similarly, the hedonic imputation method for the new computer, model n, requires estimating its price in period τ+1, but that is not possible, unless one ignores the value of the of the CD/RW. Ignoring it effectively ignores the quality change that took place.

The dummy variable method will also be ineffective. If one adds a dummy variable for the CD/RW drive to data for the second period, this dummy variable will be exactly collinear with the time dummy. The method cannot work.

There are a few ways around the problem. One can still estimate the price of the old computer in period τ+2, so long as some computers are still sold that do not have CD/RW drives. Thus, a hedonic imputation for the exiting machine is still practical, even if hedonic imputation for the price of the new machine is not. If one thinks of the Laspeyres index as measuring the prices of the machines in the period τ+1 sample, then in this restricted sense the hedonic imputation method can be applied.[46] One can also use the hedonic quality adjustment method by taking the adjustment from the hedonic function for period τ+2.

[46] Though this is a popular to think about the price index sample in a Laspeyres index, it is probably a misleading way to think about it. See Dalén (1999).

As well, the Laspeyres version of the characteristics price method works The weight of CD/RW in period τ+1 is zero, so the missing period τ+1 characteristics price for the CD/RW drive is not a difficulty (though it is not clear that the Laspeyres version is satisfactory). Neither the Paasche version nor the Fisher version can be computed.

Generally, the arrival of a new characteristic cannot be evaluated satisfactorily with hedonic methods. A truly new characteristic is like a truly new good: Its value can readily be estimated after it is introduced, but its value in the period before it is introduced is problematic. Hausman (2003) correctly notes this problem with hedonic indexes and connects it to the estimation of a "virtual price" for new product varieties (see his well known estimate for new breakfast cereals, Hausman, 1997). The appropriate valuation in a COL index is the minimum price in the previous period that would have resulted in consumers demanding none of the product, or in our example, the price which would have induced computer buyers to forego entirely a CD/RW drive.

Even if one can compute one of these Laspeyres-type indexes described above, the problem remains unresolved: Computer quality changed in the example by inclusion of a new characteristic. One needs to value that characteristic in order to estimate a price index on all the computers. Ignoring the new variety with its improved specification will, like ignoring any new introduction, produce bias in the index, for the reasons discussed previously in this chapter, and in chapter IV.

E. RESEARCH HEDONIC INDEXES

The language and examples in this chapter have been oriented to the statistical agency problem of forced replacements in fixed samples. However, the content of the chapter applies as well to research hedonic indexes, or to cases where a statistical agency might decide to estimate hedonic indexes from some extended database such as scanner or market information databases. A few matters from the research context justify some additional remarks.

As noted earlier, most research hedonic indexes use the dummy variable method, and most of them have been conducted on large databases that include, so far as possible, all the varieties of a product that exist in all the periods of the study. Few statistical agencies have considered this method, but the reasons have much to do with the database used for the price index, the need to produce an index in a timely manner, and so forth, as discussed earlier in this chapter. These constraints do not similarly constrain researchers.

From the research standpoint, the dummy variable method is the simplest (which is the major reason it has been employed so frequently in research). The hedonic imputation method and hedonic quality adjustment methods are the most resource intensive.

The hedonic quality adjustment method will probably not appeal to researchers. First, its advantages stem from the freedom it gives statistical agencies to estimate the quality adjustment from a different database from the one used for the index, and to estimate the hedonic function (from which the adjustment is taken) prior to the publication of the index. Neither is any particular advantage to the researcher, except for cases where the researcher desires to replicate price index practices for some reason.

Relevant research choices, then, appear to be the dummy variable method, the characteristics price index method and the hedonic imputation method.

The dummy variable method has been criticized, seemingly forever, for its holding characteristics prices fixed over the estimation period, and (far less prominently) for its implied index number formula, which might not be the appropriate one for index number theory. Table 3.1 suggests that the empirical importance of these points has been exaggerated, at least if best practice is pursued. However, certainly with respect to the first criticism, researchers ought to abandon the multi-period pooled form of the dummy variable method, in favour of a series of adjacent-period estimates. The logic is simple: Constraining the coefficients is clearly undesirable, and usual statistical tests indicate that coefficients in fact change, so best practice calls for the minimum constraint on the coefficients that is consistent with the method. Minimum constraint is the adjacent-period regression, which ideally should be estimated for relatively short intervals, monthly or quarterly, if data permit, and not (say) annually or longer intervals.

The price index for characteristics is straightforward and relatively easy to do. Even though its neglect in the research literature has left some methodological matters relatively unexplored, it should be estimated more frequently in research studies. For one thing, it permits evaluating the weighting and index number formula objections to the dummy variable method. Even if future studies produce similar findings to those displayed in Figure 3.1 (that is, the dummy variable index and the price index for characteristics agree closely), that is a research finding of importance. The price index for characteristics fits more closely the theory of hedonic indexes and ought to receive more attention in empirical studies than it has in the past.

Finally, research studies ought to consider the hedonic imputation method. With modern computers, this is not a taxing computation. Imputing all entering and exiting machines is valuable in itself, for it assures that the price index implications of all changes in the sample are built into the index. But as well, displaying these imputations separately, as done in a few hedonic studies, is essential information to determine why hedonic indexes differ from matched model indexes. The reasoning and existing empirical examples are presented more fully in Chapter IV.

F. CONCLUSIONS: HEDONIC PRICE INDEXES

Hedonic price indexes have sometimes been called "regression price indexes." Although it is true that regressions are used to estimate the hedonic function, hedonic indexes are not necessarily regression *indexes*. A hedonic index may be constructed in at least four ways: The dummy variable method, the characteristics price index method, the imputation method, and the hedonic quality adjustment method. In practice, statistical agencies that have implemented hedonic indexes have mostly used the latter, partly because of the necessity for producing a timely index. The hedonic quality adjustment method can be estimated using a hedonic function from a prior period, where the dummy variable method (and other methods) requires the current period's hedonic function as well. But there is no reason why the dummy variable method should not be employed when it is feasible. Its major liability is the difficulty in introducing weights into the dummy variable index. As discussed in Chapter IV, there is virtue in methods that make use of all the data that can be collected, and the dummy variable method, as well as the characteristics price index method, does that.

APPENDIX A TO CHAPTER III

HISTORICAL NOTE

Although Andrew Court (1939) published the first article on hedonic price indexes, researchers before Court discovered relations that resemble hedonic functions to some extent. Examples are Waugh (1928), who estimated price-characteristics functions on vegetables, and Haas, who even earlier estimated land price-location functions, as discussed in Colwell and Dilmore (1999). Taylor (1916), discussed as another precursor by Ethridge (2002), investigated quality dispersion in the cotton market and associated price differentials, but he did not relate price differentials to the characteristics of cotton in a statistical analysis.[47]

Many relations in economics resemble a hedonic function. A hedonic function is simply a regression on value and explanatory variables that is based on the notion that the transaction that one can observe is a bundle of lower order transactions that determine the value of the bundle—the "hedonic hypothesis." For example, "human capital" wage regressions from the labour economics literature are also hedonic functions—earnings on the left hand side of the equation are regressed against variables such as years of schooling and years of experience, where the characteristics schooling and experience stand for the intrinsic productiveness of the worker. The estimated regression coefficients are interpreted, or used to obtain, the implicit return to investment in education, or the price the employer pays for education in the workforce, which is exactly parallel to using the regression coefficients of a hedonic function on computers to estimate the implicit price of computer characteristics. Indeed, such wage regressions are sometimes referred to as "hedonic wage regressions." The idea of human capital has been traced back to Adam Smith.

Thus, even outside agricultural economics ideas similar to the hedonic hypothesis have been around for a long time. Court (1939) was undoubtedly not the first to discover a relation between value and explanatory variables that one could call "hedonic," and of course he never claimed that. Finding "precursors" to Court in this sense is somewhat beside the point. If one wants to determine who was the first to actually estimate what we would now call a hedonic function, it was undoubtedly an agricultural economist (I say that because in so many empirical research topics in economics, agricultural economists were estimating things, when most other economists were still indulging in speculation—"armchair" reasoning, as it used to be called).

Court does seem to be the first to came up with the idea of applying a value function (which he named "hedonic") to the problems of quality change in price indexes. It is, however, interesting that Court himself gave credit to Sidney W. Wilcox, then chief statistician of BLS, for suggesting the statistical analysis so the proper place to look for a precursor to Court might be in BLS published or unpublished materials somewhere. Court credited Andrew Sachs for suggesting the name hedonic, which was chosen on the presumption hedonic indexes measured "the potential contribution…to the welfare and happiness of its purchaser and the community" (Court, 1939, page 107). Of course, we now know that this "user value" interpretation of hedonic indexes is not necessarily the appropriate one (a question reviewed in the Theoretical Appendix).

---

47 Though Taylor was interested in quality price differentials in the cotton market (as was Waugh for the asparagus market), Taylor did not quite grasp the hedonic *function* idea. Both Taylor and Waugh were trying to provide analysis that would help either the farmers or improve the efficiency of agricultural markets. The two examples show that agricultural economists were working on problems of heterogeneous products when most non-agricultural economists were quite content with the assumption that products were homogeneous.

A second Court—Louis Court (1941)—was the first to explore the idea of constructing a model of economic behaviour toward the characteristics. Louis Court was thus a precursor (as was Gorman, 1980, but written much earlier) of Lancaster (1971) and Ironmonger (1973), who explored consumer demand toward the characteristics of complex goods.

The dummy variable method for estimating hedonic price indexes first appeared in Court (1939). Stone (1956) and Griliches (1961) estimated dummy variable indexes in several specifications, and the dummy variable method was the major method employed for hedonic price indexes in most of the research of the 1960's and 1970's.

Griliches (1961 and 1971) was the first to produce hedonic indexes by methods other than the dummy variable one, and to discuss the advantages of alternative estimation methods for hedonic indexes. He estimated a quality adjusted automobile index that used a method that was closely related to the hedonic imputation indexes described above. In a second example, Griliches calculated the (geometric) mean unadjusted price change for a group of cars and divided it by a "quality index." Griliches' quality index valued the changes in average specifications for the automobiles in his sample, using the coefficients from a hedonic function. Court (1939) suggested an identical procedure, but did not estimate it. See also Griliches (1971), where he discusses a number of difficulties with the dummy variable approach and remarks that using the hedonic function only for estimating the implicit prices of characteristics is the approach "most directly in the usual price index spirit." As noted in the main body of the text, Griliches is also the first to describe the characteristics price index, which he called the "price-of-characteristics" index.[48]

The first government application of the characteristics price method was the New House price index, which has been constructed by the U.S. Census Bureau since 1968. This index was introduced in the U.S. national accounts beginning in 1974, and extended back to 1968. Thus, the hedonic index for new house construction is not only the first hedonic index in any country's economic statistics, it is also the first hedonic index used in any country's national accounts. This index is still published, and can be retrieved at the U.S. Census Bureau website (U.S. Census Bureau, undated).

Considering the controversy that has sometimes surrounded the extension of hedonic indexes to other products, it is odd that the new house hedonic index has been free from controversy (see the review of construction price indexes in Pieper (1990), for example). Even Denison (1989), who denounced the BEA hedonic price indexes for computer equipment, accepted the hedonic price index for new houses without criticism.

The method of applying hedonic quality adjustments to replacement items in a price index was first introduced in Triplett and McDonald (1977). They replaced the actual quality adjustments that had been made in the U.S. Producer Price Index (then called the Wholesale Price Index) for refrigerators with hedonic quality adjustments applied to the same PPI refrigerator sample, and compared the resulting indexes. Most of these PPI refrigerator quality adjustments at that time were link-to-show-no-price-change methods (see Chapter II).

As noted in Chapter V, the earliest computer hedonic functions and price indexes were estimated by computer scientists, and by economists such as Chow (1967) and Knight (1966), mostly with the dummy variable method. The first government price index for computers (Cole, et al., 1986; Cartwright, 1986) used the hedonic price imputation method to estimate prices for computers that were present in one period but not in the other. This IBM-BEA index was constructed by the Paasche price index number formula because it

---

[48] The bizarre title of the paper by Koskimäki and Vartia (2001) shows how much the flavor of Griliches' work has sometimes been misinterpreted. Indeed, in Griliches (1971) he expressed reservations about the dummy variable method.

corresponded to the deflation system then in use in the U.S. national accounts; weights were based on sales data for individual models of (mainframe) computers, and of peripheral equipment. In the research that led up to the final index, several other hedonic indexes were calculated for research and comparison purposes, including a price index for characteristics and a time dummy variable index. All of them coincided closely. These are discussed in Dulberger (1989). In Europe, the first government hedonic price indexes for computers built on the work of Dalén (1999) for Sweden and Moreau (1996) and Bourot (1997) for France.

Although there was, dating from Griliches (1961), a great amount of interest in defining the hedonic index from a theoretical perspective, little of that proceeded as if the index number could be defined in characteristics space. For example, in Adelman and Griliches (1961), which the authors themselves note was assembled from independent manuscripts, the index number theory is not a theory of a characteristics space index, but a very general theory of a goods space index to which the characteristics have been somewhat uncomfortably grafted on. Triplett (1971b) was the first to redefine a price index formula into characteristics space. He pointed out that you could think about a Laspeyres index, such as the CPI, as if the index weights were the quantities of characteristics of the product that were consumed in the base period, and not just the quantities of the products themselves. "Quality adjustment" (an adjustment for changes in characteristics included in the product) could then be given the interpretation that it was merely a device for holding constant the index (characteristics) weights. Turvey (1999) apparently discovered this idea independently. This characteristics-space index number framework is an implication of Lancaster's (1971) demand for characteristics model, where the characteristics are the quantities demanded and consumers respond to changes in the characteristics prices.

A more general price index formulation of characteristics price indexes is Triplett (1983, 1987), who distinguished exact (in the sense of Diewert, 1976) characteristics space input or cost-of-living indexes and exact output price indexes for characteristics. These two correspond to CPI and PPI concepts, respectively. It has long been well known that CPI (cost-of-living) and output price indexes (PPI) have different theoretical properties (Fisher and Shell, 1972; Archibald, 1977); Fisher and Shell show that this difference carries over to the treatment of quality change. Triplett's results are consistent with that tradition. Feenstra (1995) expanded on the idea of exact COLI indexes for characteristics, and suggested (stringent) conditions where they might be produced empirically by hedonic indexes. Fixler and Zieschang (1992) explore a similar exact characteristics price index idea. Pakes (2003) discusses bounds that might be placed on the true characteristics price index by using the hedonic function (he actually talks about bounds on the compensating variation—the amount that would have to be given to an individual to leave the individual on the same indifference curve—but the concepts are derived from the same principle). However, Pollak (1983) had already showed that all characteristics space COLI formulations (including all of the above) are special cases that depend on particular specifications of how characteristics enter the utility function (for COLI indexes). He also showed that there are many special cases. This means that deriving bounds on the COLI indexes, or on the compensating variation, is not so simple as it is in the normal Laspeyres, good-space world in which the usual COLI theory is developed.

APPENDIX B FOR CHAPTER III

HEDONIC QUALITY ADJUSTMENTS AND THE OVERLAPPING LINK METHOD

As explained in chapter II, The overlapping link method takes the actual difference in prices between two computers in some overlapping period as the "ideal" quality adjustment. As also noted in chapter II, the overlapping link method is seldom employed in practice, but it provides the conceptual framework for the actual adjustments made conventionally. The hedonic function can be used to examine the conditions under which the overlapping link method gives the appropriate measure, so this Appendix evaluates the conceptual framework for conventional quality adjustment.

For this Appendix, I use figures that have already been used in this chapter, to avoid replication of additional figures; however, the analysis of this Appendix refers to situations where prices of entering and exiting computers are both observed at the same time, where the earlier use of these figures applied to cases where these prices were only observed for different periods. For this Appendix, the reader must ignore the time subscripts in these figures. Otherwise, they suit the analysis.

1. The Overlapping Link Method if Both Prices Are on the Hedonic Function. In Figure 3.3, computers m and n (discontinued and replacement computers) both lie exactly on the hedonic function (recall that for the overlapping link method, prices for both computers are observed in the overlap period, which is period t+1). Thus, the higher market price for computer n in period t+1 equals exactly the value of its higher level of characteristics, no more and no less.

Under this scenario, the treatment of these two computers under a hedonic price index and under the overlapping link method is exactly the same. From Chapter III, section C.3, the hedonic quality adjustment, shown as A(h) on Figure 3.3, equals the ratio of the prices of computers m and n in period t, or R_t. This R_t is also the adjustment, A_o, implied by the overlapping link method discussed in Chapter II.

It is sometimes said that hedonic indexes and the overlapping link method require the same assumption, namely that the two prices indicate relative valuations of the quality of the two computers. Figure 4.1 shows that this assertion is true only if both the old computer and its replacement lie exactly on the hedonic line.

Note that for this example *no price change takes place either when a model exits or when a new one enters*. However, it will not in general be true that both the discontinued computer and its sample replacement both lie exactly on the hedonic surface. Accordingly, it will not in general be true that the introduction of new computer models and the exits of old ones have no implications for price change. Moreover, it will not in general be true that a hedonic quality adjustment equals the adjustment implied by the overlapping link method. I turn to other cases in the following.

2. The Overlapping Link Method if Prices Are Not Exactly on the Hedonic Function. Suppose that computer m is supplanted in the market by computer n because computer m was overpriced, relative to its characteristics content. Economists usually believe that the overpriced machines will be displaced because buyers will shift to better values, that is, toward more favorable price/quality ratios.

Assume for the moment and for simplicity that the new computer, n, is introduced at a price that lies exactly on the hedonic surface. This case is depicted in Figure 3.8. Computer n's price is just equal to the value of its characteristics, so the introduction of computer n does not change the hedonic function (which is why the assumption was said to simplify the discussion). The old one, computer m, was overpriced, relative to its characteristics content, so that it lies above the hedonic function.

Using the hedonic function to estimate the value of the difference between m and n yields the same hedonic quality adjustment A(h), as before. However, using R_t (the ratio of the two computers' actual prices) as a quality change estimate is much too small. Computer n supplants computer m because for very little increase in price, computer buyers can get a substantial increment in speed. In this case, if A(h) is used as a quality adjustment, the quality adjusted price falls (by the amount B). The overlapping link method would suggest that prices have remained constant, which is clearly incorrect.

Now assume, instead, that the price of computer m, the old computer, lies below the hedonic line. For example, it might have been on sale when prices were collected in period t. This would introduce the opposite error. This error is shown in Figure 3.9.[49]

However, price statistics agencies have become reluctant to permit a price observation to either enter or leave the sample when the price collected is a special sales price, to prevent cases where the return to normal of a sales price is advertently linked out of the index. I think they are right to do this. This suggests an asymmetry in prevailing statistical agency practice: Price changes for computers that are underpriced because they are on sale are censored when they leave the sample. There is no equivalent censoring of overpriced models.

Pakes (2003) has suggested another mechanism that motivates exits: Producers will withdraw machines whose prices are *below* the hedonic surface (bargains), because they are less profitable. This can only happen when producers have market power (otherwise new entrants will provide the bargains), but Pakes also points out that market power is typical of producers in markets that are characterised by rapid technological progress. In Pakes case, figure 3.9 also applies, and market exits produce a downward bias in conventional indexes—the exit of a bargain raises the price, but this will not be detected by the overlapping link method, nor by normal matched model methods.

When the new computer is priced so it is exactly on the hedonic line, the direction of the bias from the overlapping link method (or from the IP-IQ method) depends on whether there is a systematic tendency for the new computer to displace old ones that were overpriced or ones that were underpriced. The seller of overpriced computer, m, for example, ought to attempt to meet the (quality adjusted) price of the new computer, n. There is evidence in the computer market that this does not happen, that more typically, the overpriced computers just lose their market shares and disappear, without ever meeting the price/quality terms that are offered by newer machines. On the other hand, Silver and Heravi (2002) suggest that exits have higher quality-adjusted prices than continuing models in the market they explored; this finding supports Pakes (2003) expectation that producer's search for more profitable models explains market exits. There is need for more research on the impact of new machines and new technologies on existing markets.

Finally, the new computer might be priced above or below the (old) hedonic line. Pricing below the line is the case that empirically emerges in many price index studies of high-tech equipment. This

[49] A reader of the draft version of this handbook has suggested that in a probabilistic sense, the errors suggested by Figures 3.8 and 3.9 should cancel out, so there may be no overall index error. Although this might be true, the diagrams are intended to show the errors incumbent with each case; one suspects that the distribution of the cases themselves are not random, particularly for ICT products. If sellers systematically introduce new products at price/quality ratios that are more favourable than those of continuing products, the introduction of new products implies price reduction. The contrary is true if sellers take the introduction of new products as an opportunity to increase prices, or if price increases are timed to coincide with new model introductions.

casecorresponds to Figure 3.5. For almost any price of the old computer, the ratio of actual prices (the overlap) is an inaccurate measure of the quality difference between the two computers.

Table 3.1

Comparison of Hedonic Index Methods: Adjacent-Period Dummy Variable Method and Characteristics Price Index Method

Study and Product	Dummy Variable	Characteristics Price Index
Dulberger (1989, table 2.6)		
Computers, AARG (1972-84)	- 19.2%	- 17.3%[a]
Okamoto and Sato (2001, charts 2 and 5)		
TV's (AARG, 1995-99)[b]	- 10.4%	- 10.4%
PC's (AARG, 1995-99)[b]	- 45.1%	- 45.7%
Digital Cameras (AARG, Jan 2000-Dec 2001)[c]	- 21.9%	- 21.9%
Silver and Heravi (2002, 2003)		
TV's (total 11 month change)	- 10.5%	- 10.1%
Washing machines (ibid)	- 7.4%[d]	- 7.6%
Berndt and Rappaport (2002, table 2)		
PC desktops (AARG)		
1991-1996	- 37.0%	- 38.4%
1996-2001	- 35.7%	- 34.3%
PC laptops (AARG)		
1991-1996	- 26.9%	- 26.0%
1996-2001	- 39.6%	- 40.6%

Notes:

a. Characteristics prices taken from multiperiod pooled regression (see text).
b. Computed from values for Jan-Mar (first quarter) for the end periods, not from the annual averages given in the charts noted above.
c. New estimate, communication from Masato Okamoto to the author, January 26, 2003
d. New estimate, communication from Saeed Heravi to the author, January 24, 2003

Table 3.2

Hedonic Coefficients, Adjacent Period and Pooled Regressions
Automobiles, 1959 and 1960, Separate and Pooled Regression Coefficients

	1960	1959	1959-60
Horsepower	0.119	0.118	0.114
Weight	0.136	0.238	0.212
Length	0.015	- 0.016	- 0.006
V-8 engine?	-0.039	-.0.070	- 0.059
Hardtop styling?	0.058	0.027	0.040
Auto-transmission?	0.003	0.063	0.040
Power steering?	0.225	0.188	0.206
Compact car?	NA	NA	0.052
R^2	0.951	0.934	0.943

Source: Griliches (1961), tables 3 and 4.

Table 3.3

Matched-model and Hedonic Indexes for UK Washing Machines
(index percent changes, January to December 1998)

Matched Model, Deletion (IP-IQ)	- 9.2 %
Hedonic Imputation	- 8.6 %
Hedonic Quality Adjustment	- 8.8 %
Characteristic Price Index (Fisher)	- 7.6 %
Dummy Variable	- 6.0 % [a]

Note: /a This time dummy index constrains coefficients over 12 periods. Results for constraining only for adjacent periods (a preferred implementation) are discussed in a previous section of this chapter.

Source: Silver and Heravi (2003).
First three lines: Table 9.7, lines titled “Imputation,” “Predicted versus Actual, Geometric Mean,” and “Adjustment via Coefficients, Geometric Mean.”Fourth and fifth lines: Table 9.5, “Exact Hedonic, Fisher” and “Time Dummy Variable: Linear.”

Figure 3.1

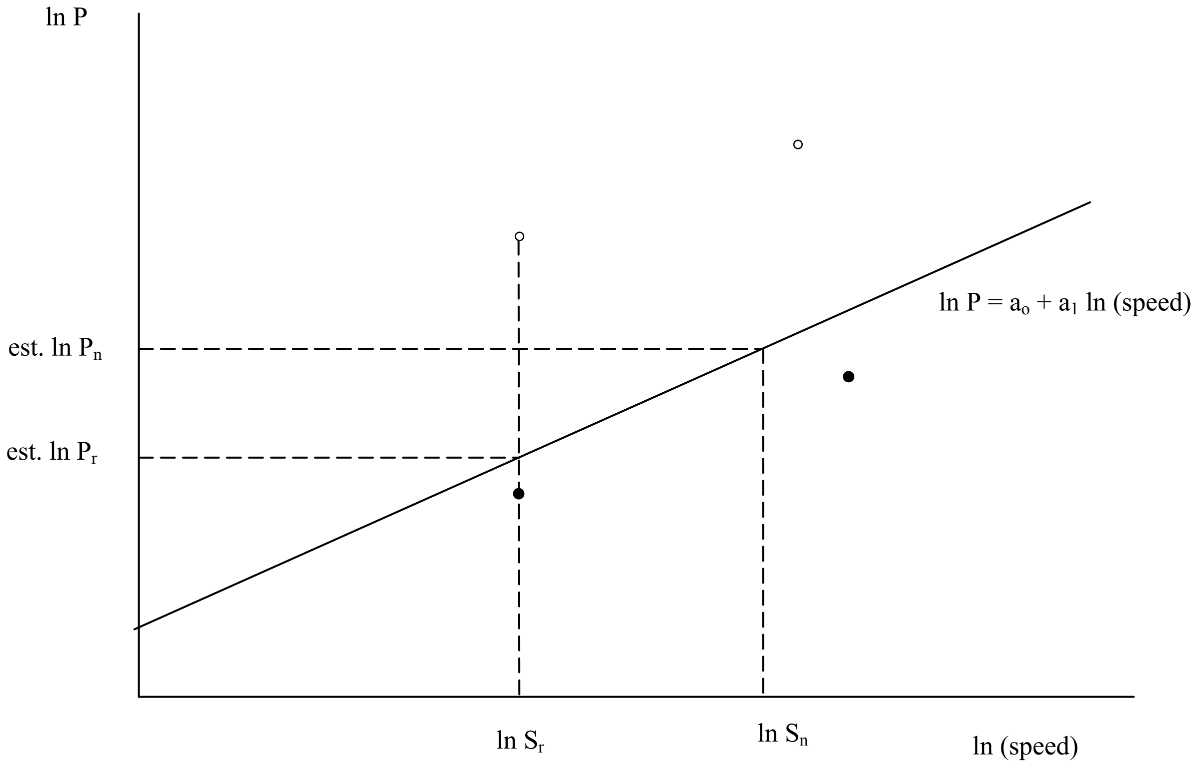

Figure 3.2

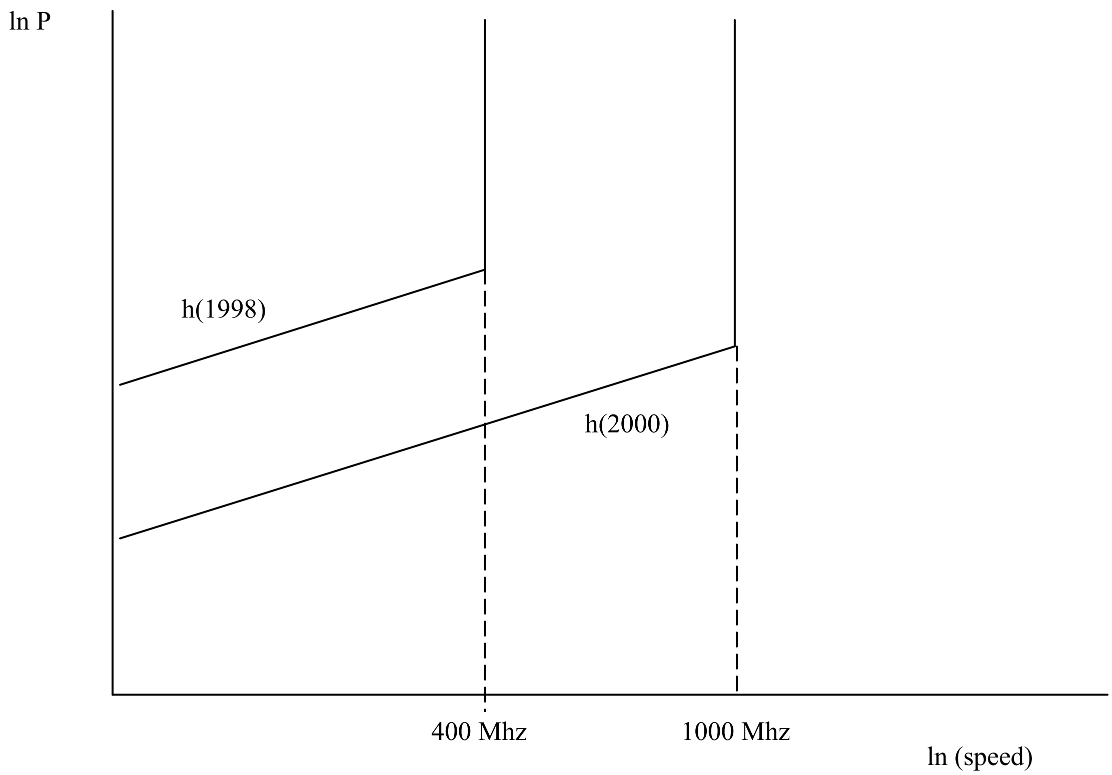

Figure 3.3

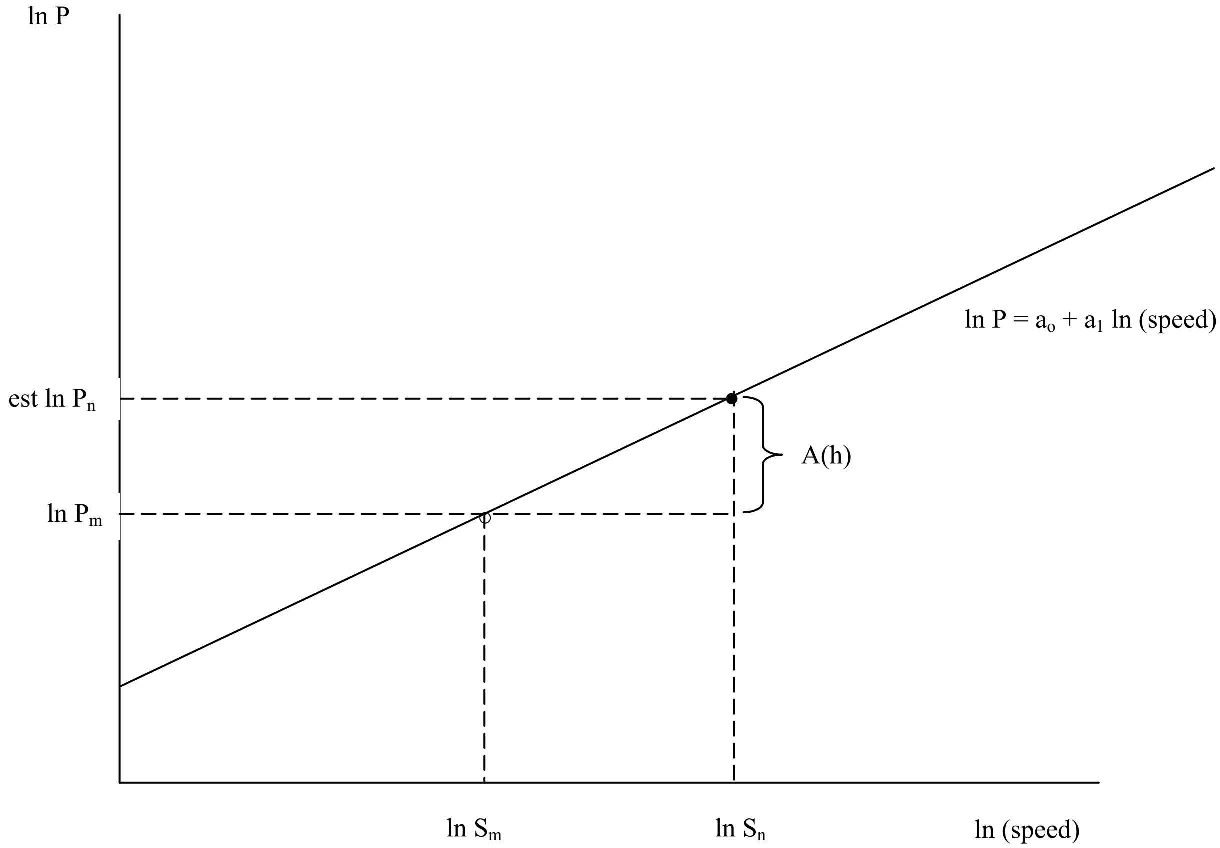

Figure 3.4

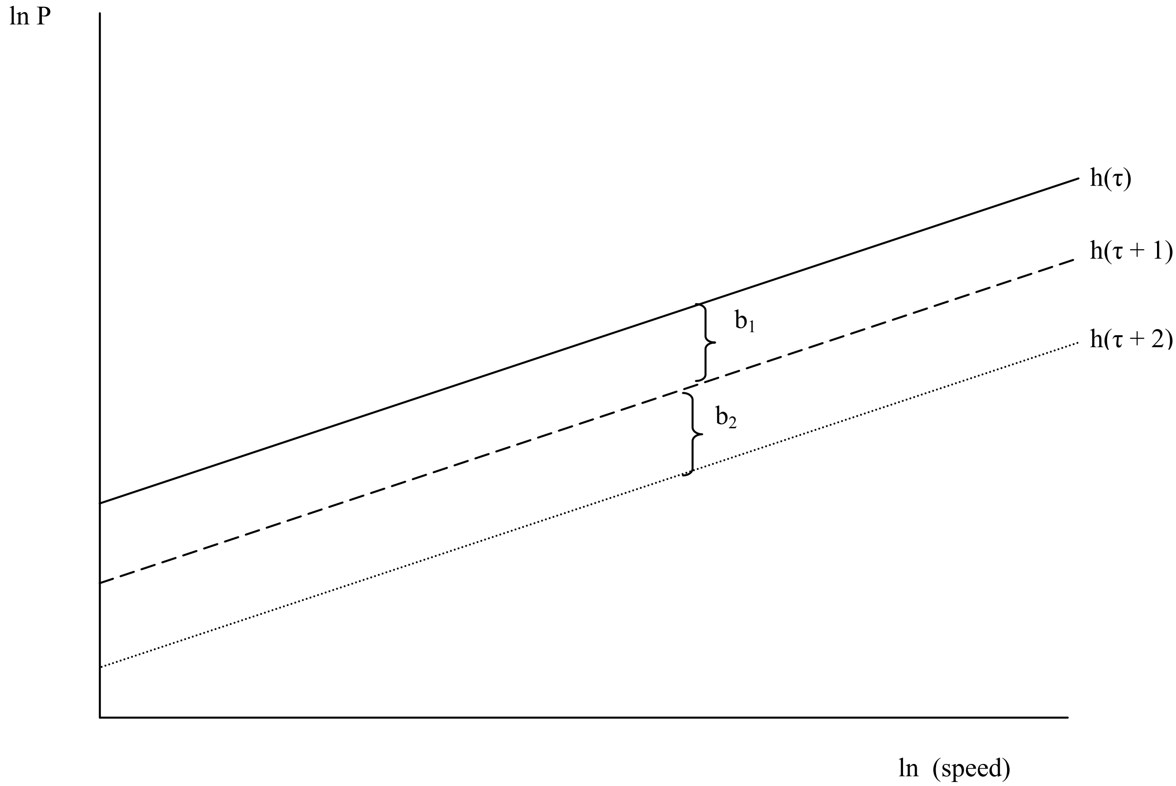

Figure 3.5

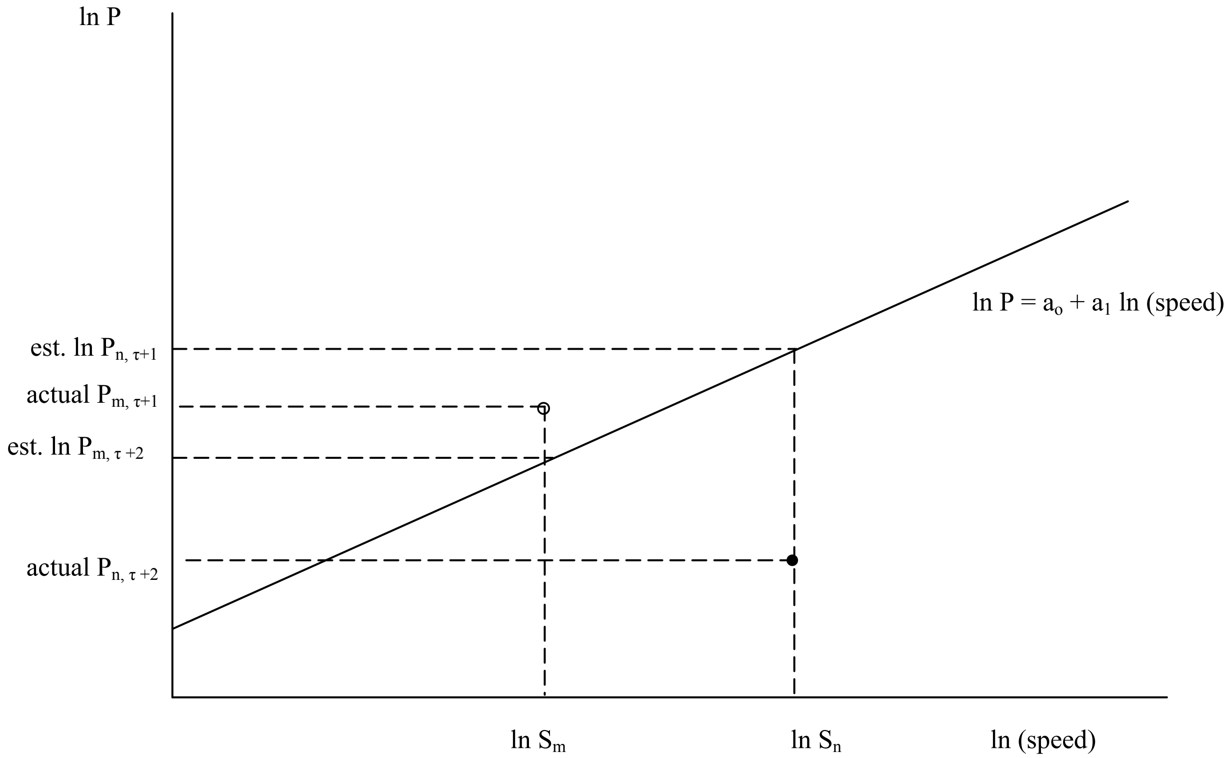
ln P
ln P = a$_o$ + a$_1$ ln (speed)
est. ln P$_{n, \tau+1}$
actual P$_{m, \tau+1}$
est. ln P$_{m, \tau+2}$
actual P$_{n, \tau+2}$
ln S$_m$
ln S$_n$
ln (speed)

Figure 3.6

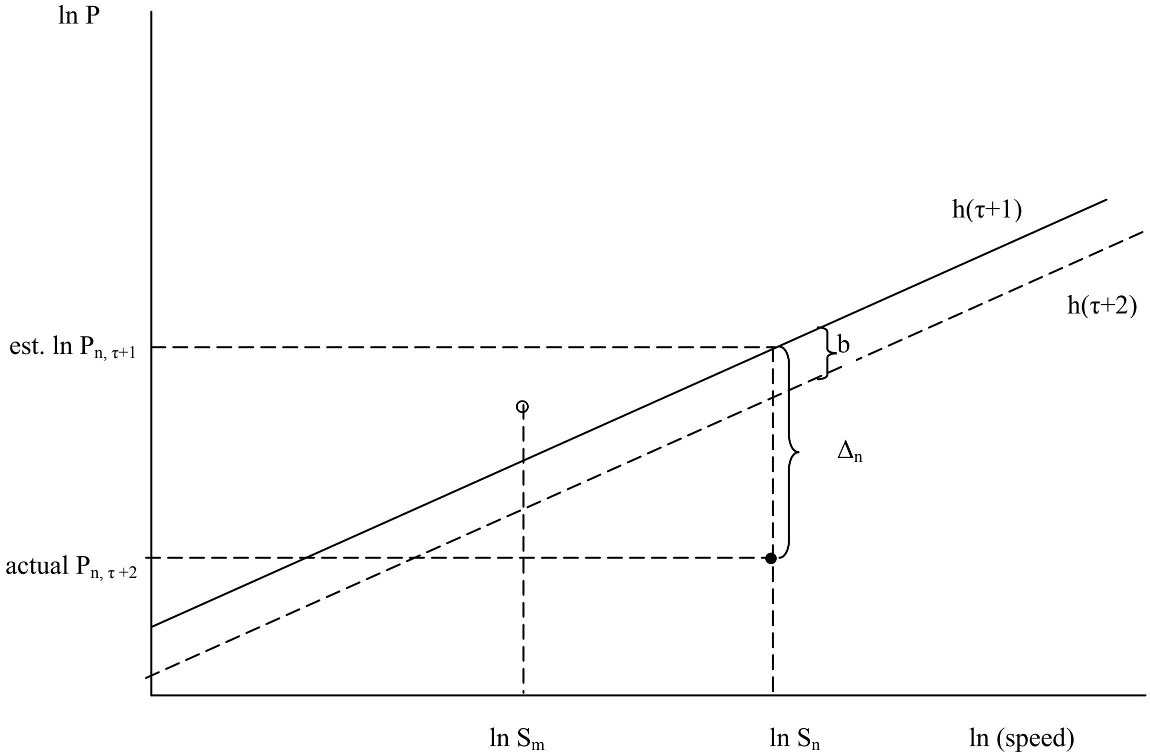
ln P
h(τ+1)
h(τ+2)
est. ln $P_{n,\, \tau+1}$
b
Δ_n
actual $P_{n,\, \tau+2}$
ln S_m
ln S_n
ln (speed)

Figure 3.7

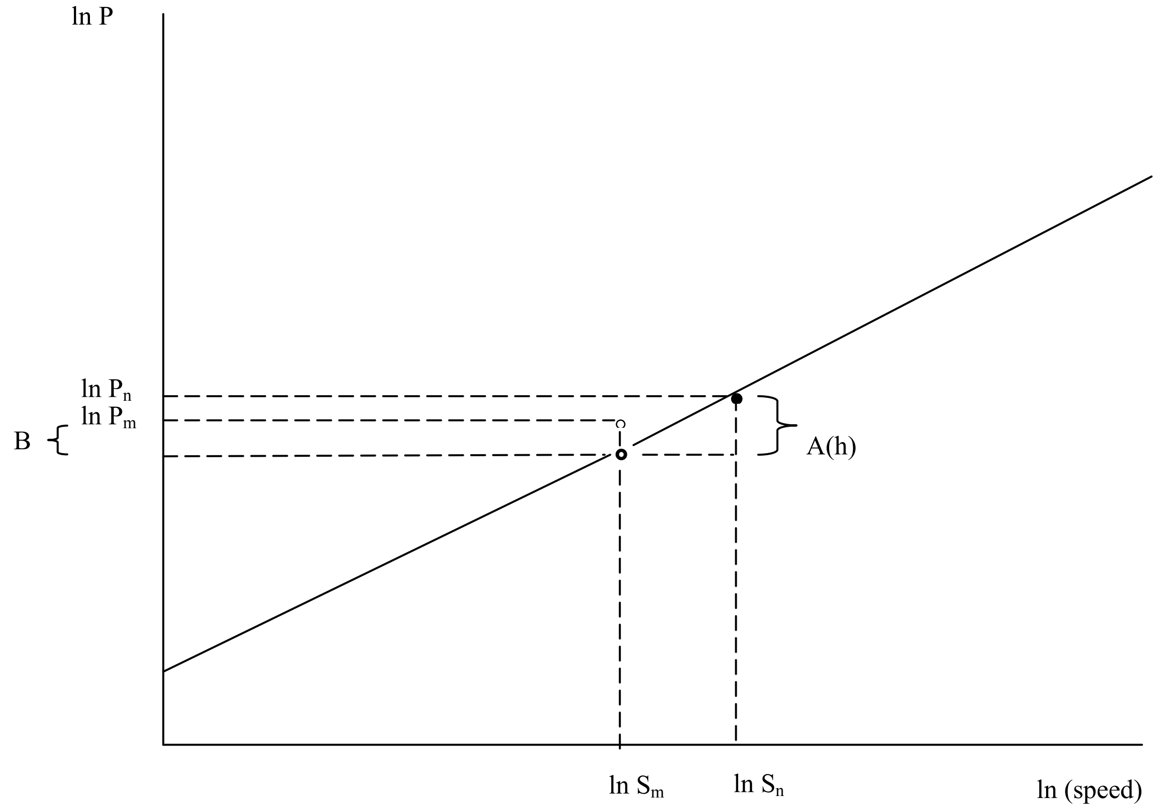

Figure 3.8

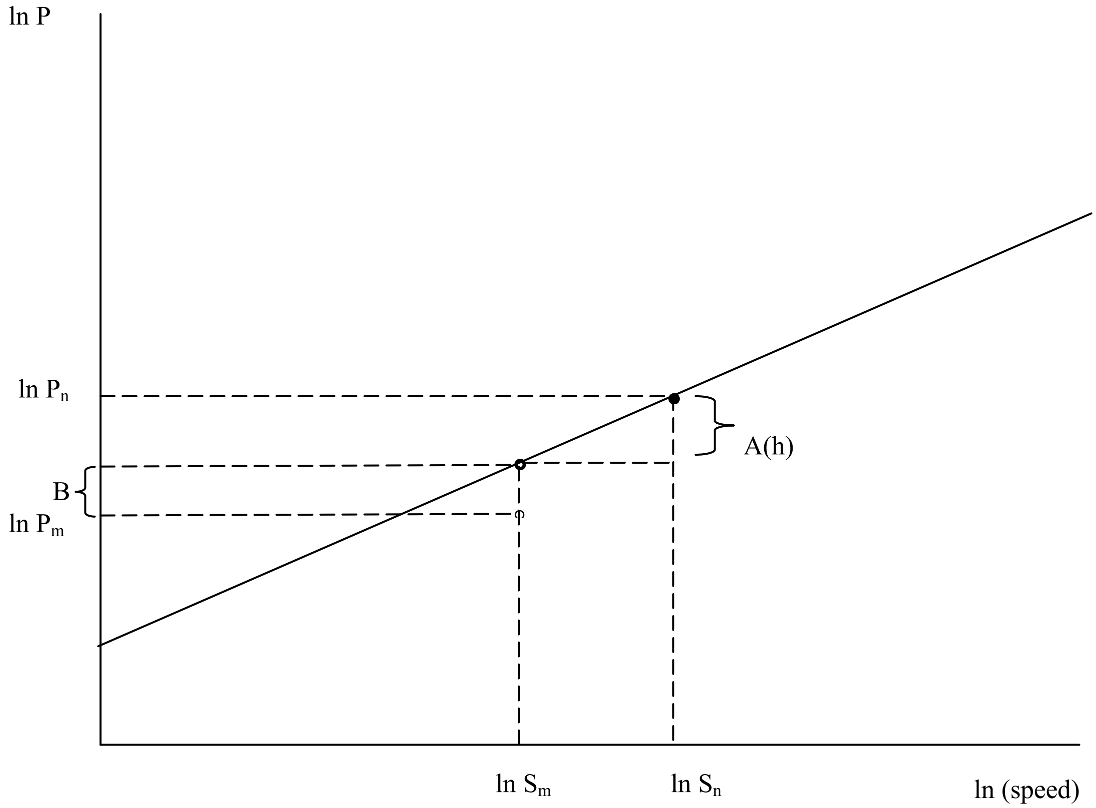

Figure 3.9

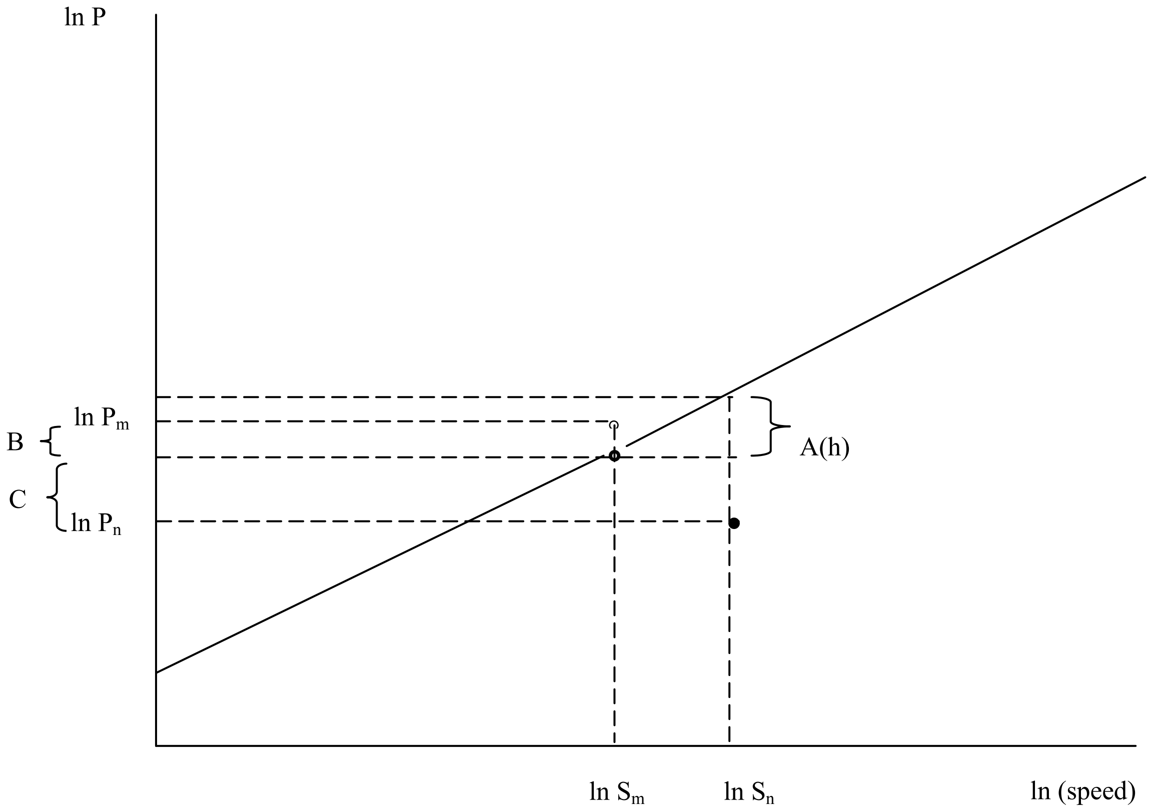

Figure 3.10