NBER WORKING PAPER SERIES

A RECONSIDERATION OF HEDONIC PRICE INDICES WITH AN APPLICATION TO PC'S

Ariel Pakes

Working Paper 8715
http://www.nber.org/papers/w8715

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2002

My interest in hedonics, and in price indices more generally, developed largely through my interactions with Zvi Griliches and the intellectual environment he created. I very much missed the fact that I could not discuss this paper with him. I have benefitted from discussions with Steve Berry, Ernie Berndt, Gary Chamberlain, Ian Cockburn, Erwin Diewert, Shane Greenstein, Alan Krueger, Minjae Song, Jack Triplett, and participants in a number of seminars (most notably a seminar at the BLS). Minjae Song provided excellent research assistance. This research was partially funded from grants from the Bureau of Labor Statistics and the NSF, and I am grateful to those institutions for their support. The views expressed herein are those of the author and not necessarily those of the National Bureau of Economic Research.

A Reconsideration of Hedonic Price Indices with an Application to PC's
Ariel Pakes
NBER Working Paper No. 8715
January 2002
JEL No. C8, L0

ABSTRACT

This paper provides a justification for hedonic price indices and details the properties of hedonic price functions. The analysis is done in a market setting in which a finite number of goods, each defined by its characteristics, interact. We note that proper hedonic indices can be constructed from the same data currently used to construct matched model indices. Since the matched model index does not incorporate price changes for goods which exit, and the goods that exited tend to be those goods whose prices fall, the matched model index has a selection problem which biases it upwards. The hedonic index does not have this problem. We illustrate with a new study of price indices for PC's. The hedonic index shows steep price declines in every year. On average, the matched model indices indicate no price fall at all and one commonly used matched model index is negatively correlated with the hedonic. We also construct and compare alternative price indices used either in research or by the federal statistical agencies. Of these the one that seems to work well is a Pasche style hedonic. Its advantage is that since it does not require computation ofthe current period's hedonic function, it is easier to use when monthly timetables need to be met.

Ariel Pakes
Department of Economics
Littauer 117
Harvard University
Cambridge, MA 02138
and NBER
apakes@kuznets.fas.harvard.edu

This paper has two parts. The first reconsiders the economic model underlying hedonic price indices. It focuses on when hedonic indices are justified and the possiblities for constructing hedonic indices when they are. It then considers how hedonic indices should be constructed when they can be used. A comparison is made to the standard matched model index, the problems that arise in constructing it, and the implications of those problems for the interpretation of matched model indices. The second part of the paper is designed to illustrate the empirical importance of the issues raised in the first part. It contains a new study of price indices for desktop computing machines (PC's).

The empirical study considers different ways of constructing hedonic indices for PC's between 1995 and 1999 and then compares the results to alternative price indices. The alternatives include; two versions of a matched model index, hybrid hedonic indices that mimic the method the BLS has been using to construct their price index for PC's since 1998, a hedonic-like adjustment frequently used in research, and two indices introduced here. One of the latter two is a "complete" hybrid hedonic designed to overcome a problem in the BLS's hybrid, and the other is a Pasche like hedonic designed to enable the BLS to produce a hedonic like index in a more timely manner then they could using a proper hedonic index.

The paper begins with a brief explanation of the new goods biases in price indices and of the Bureau of Labor Statistics' attempts to ameliorate them (section 1). This includes revising sampling procedures in ways that make new goods enter the sample faster and a somewhat controversial push towards using hedonics adjustments in an assortment of commodity groups which began in 1998 (see Liegey, 2001[1]). Section 2 provides an economic model that rationalizes hedonic indices, and then considers that model's implications on the interpretation of the coefficients in hedonic regressions.

The model provides a justification for the use of hedonic price indices and helps clarify many of the apparent anomalies that have been associated with hedonic regression coefficients in the past. It also provides guidance for how a "proper" hedonic index should be constructed. Here a proper index is defined

[1]The PC component of the CPI began using a hedonic like adjustment in 1998, while the television component began in 1999. In 2000 the list of products using a hedonic like adjustment was expanded dramatically to include; twelve audio equipment products and video cameras, VCR's, DVD's, refrigerators, microwave ovens, college textbooks, washing machines and clothes dryers. For the use of hedonics in other statistical agencies, see Moulten, 2001.

as any index that estimates an upper bound to the compensating variation needed to make consumers indifferent between base and comparison period utility. Section 3 considers both the conceptual and practical problems that arise in constructing a hedonic index, and compares these problems to the problems that arise in matched model indices. There are special cases where the rational for hedonic indices is suspect, but they seem to be relatively unimportant.

This section also makes it clear that the data currently produced by the BLS can be used to estimate a proper hedonic index which conditions on the same characteristics that the BLS is implicitly conditioning on in producing their matched model index. Moreover, there is a selection bias in the traditional matched model index and, to a lesser extent, in the hybrid hedonic index the BLS has been using for PC's since 1998. Since the goods that exit are disproportionately goods which embody characteristics whose value has fallen over the period, any index which omits price changes for those goods without incorporating a correction for them is positively biased. This selection bias is *not* present in the proper hedonic index, and is likely to be particularly severe in commodity groups with a high rate of technical change. The section also compares the likely variances of the matched model, hedonic, and hybrid hedonic indices. Though there is no *a priori* ordering of the variances of the three indices, there are good reasons to expect the variance of the hybrid to be the *largest*. Section 3 concludes with a warning. Though hedonic procedures mitigate the selection bias in matched model indices neither index makes any attempt at all to measure the value created by the entry of new goods, and hence even the hedonic may still be quite distant from an "ideal" index.

Section 4 of the paper considers alternative ways of constructing hedonic indices, section 5 introduces our PC data, and section 6 computes and compares alternative price indices for PC's for the period 1995 to 2000. The proper hedonic index for PC's shows a rather sharp decline in PC prices (about 17%). The rate of decline is quite different in different years and these differences are consistent with the changes in the characteristics of the products marketed in those years. Provided we stay within the guidelines of the theoretical discussion, the details on how the hedonic is constructed do not seem to have much impact on the index itself.

On average the matched model indices indicate no change in PC prices over this period, and the Tornquist version of this index actually shows both a positive average price change and a strong negative correlation with the proper hedonic index. Our attempt to mimic the BLS hybrid produces an index which is more similar to the hedonic than the matched model index, but is somewhat erratic and, as expected, has a large variance. The actual hybrid reported in BLS publications differs from both our attempt to mimic BLS procedures and our proper hedonic index. In partciular the BLS estimate is over a third larger (in absolute value) than our proper hedonic in the two years for which the comparison can be made (in both years this difference was over a standard deviation higher than our proper hedonic).

We also introduce and compute; (i) a complete hybrid which should perform better than the BLS's hybrid, (ii) a dummy variable index which is often used in research, and (iii) a Pasche style hedonic which would enable the BLS to produce a hedonic index in a more timely manner. The complete hybrid uses the actual price change for goods which survive to the reference period, and the hedonic estimate of the price change for the goods that do not. Its promise is that though it still has some bias, it mitigates the impact of estimation variance. If it produced an index with lower variance than that of the proper hedonic then one could make an argument for using it. Our calculations indicate that this does not occur in our data, and is even less likely to occur in the BLS's data.

Both the research dummy variable index, and especially the Pasche hedonic index, do better than all the other alternatives to the proper hedonic. Unfortunately the research hedonic that we use requires data on future prices, so the BLS could not use it in constructing their index. Moreover if we would have had the data to construct an analogous index on the basis of current and past data it would have been likely to have much poorer properties. The performance of the Pasche like hedonic index is a surprise and since it only requires estimation of the base (and not the reference period) hedonic function it would be much easier for the BLS to use it to do the initial calculations for their monthly publications. One possibility is to use the Pasche like hedonic in the BLS's initial publications, and then revert to the proper hedonic in their standard revisions.

1 The New Goods Problem in Price Indices.

The price quotes used to construct most of the components of the CPI are obtained by data collectors who make repeated visits to the same outlet and form the ratio of the outlet's prices for the same good over adjacent periods.

In a world where the goods marketed in a given commodity group did not change over time, the "price relatives" obtained in this fashion would simply be averaged to obtain a "matched model" index which would become the CPI component for a given commodity group. Since the goods marketed do change over time and there is a desire to maintain representativeness of the index, the data collector is instructed to rotate a certain per cent of the goods in the index out in every period, and when a good that is not scheduled to be rotated out is no longer sold at the outlet the data collector must make a "forced substitution" of another good.

The sample rotation and forced substitution processes enable new goods to enter the index, and the new goods biases are a result of the mechanics of how these processes work. When new goods are brought into the index they are generally entered either through an overlap pricing procedure or a linking process. In overlap pricing we sample both the good being replaced and the new good in the rotation period, and then use the old good's price for the comparison to the period preceding the rotation period, and the new good's price for the comparison to the period following the rotation period. When a good is linked in it is assumed that the "quality adjusted" price of the new good in the linking period is either the same as the quality adjusted price of the old good in the prior period, or that the difference in the quality adjusted price is equal to some average of the differences in prices of the goods that were available in both periods (for more details on these and other ways of adding new goods to the index, see Moulton and Moses, 1997, and the literature cited their).

The important point, and the source of the new good bias in price indices, is that these mechanisms do not make any adjustment for differences between the "utility per dollar" of the new good and the good(s) it replaces. An example will illustrate. Assume there is a single good in the commodity group and then a new good enters which gradually obsoletes it. The computed price index for this commodity group (say $\Delta p(t)$) will be a weighted average of the price relatives of the new and old good ($\Delta p^n(t)$ and $\Delta p^o(t)$) with the weight ($w(t)$)being the fraction of price quotes that are obtained from the new good, that is

$$\Delta p(t) = w(t)\Delta p^n(t) + (1 - w(t))p^o(t). \tag{1}$$

The formula in (1) would be the traditional upper bound to the average of the income needed to compensate consumers for the price rise (to the compensating variation) if a fraction $w(t)$ of consumers purchased the new good in both periods and a fraction $(1 - w(t))$ purchased the old good. However throughout the adjustment period (the period in which $w(t)$ goes from zero to one) there is a fraction of consumers who switch from the old to the new good. The consumers who do switch have price changes not equal to the change in prices of either the new or the old good. Moreover since the consumers who did switch to the new good could have chosen the old good in the second period, revealed preference theory implies that these consumers have increased their "utility per dollar", or experienced a smaller price change, then the price change on the old good. This is the source of the "new goods" bias in price indices.

Of course if the new good really does obsolete the old good, any reasonable sampling procedure will have $w(t)$ go to one. After $w(t) = 1$ the index in (1) will indeed pick up the correct price increases. That is the measure of price *changes* will only be biased during the adjustment period. However since there is never any correction for the mistakes made during the adjustment, the *level* of the price index will be biased forever. As a result it is easy to find cases where the introduction of the new good will at one and the same time; (i) increase the computed price index, and (ii) decrease the true cost of living.

To see this we add detail to the example above. Say the new good has a smaller cost of production then the old good and that it would preferred by most (though not all) consumers were the prices of the new and old goods equal. Also assume that the producer of the new good launches the new good at a low price (say at its production cost) so as to induce consumers to experiment with the new good and spread information about it. One quarter of the consumers become aware of the new good in each of the first four periods and 90% of those that become aware prefer the new good and purchase it. As more consumers become aware of the new good the reason for the low introductory price disappears. So the producer of the new good gradually increases its price. Say the price of the new good raises by twenty-five percent in each of the first four periods. The old good's optimal response to the entry of the new good is to cede most consumers to the new good and continue charging the same mark-up to the consumers who are either unaware of the new good or who have a preference for the old one and are not very price sensitive. After four periods the prices stabilize. The old good's price remains as before the introduction of the new good, but since the new good has a smaller cost of production it sells at a price which is one-half of that.

Note that no consumer ever faces a higher price than the price the consumer faced before the entry of the new good. Moreover by the end of four periods ninety percent of the consumers consume a good with *both* a significantly lower price *and* a higher value to them then the good they initially consumed. As a result we know that the true average price index *fell* by more than 45% over this period (90% of the consumers consume a better good at half the price). Now assume the weights used in the calculation of the CPI for this commodity group are based on consumption shares in the beginning of every period. Then the CPI for this commodity group will *increase* $25\% \times .9 \times (.25 + .50 + .75) = 33.5\%$ over the period in which the true price index *fell by over* 45%.

Though this example gets at the heart of the new good's problem in the CPI, the way that problem manifests itself in the actual index is considerably more complex. The additional complexity is a result of the fact that there are many different goods in a given commodity group. Indeed it is more realistic to think of the matched model index assigning a weight of either one or zero to each of the different goods. As discussed in detail below, the new goods bias then becomes embedded in the determinants of which of the "one's" of last period remain in the sample during the current period.

Ameliorating New Goods Biases.

Following remarks made by Alan Greenspan and a senate inquiry into new goods biases in price indexes which culminated in the Boskin commission report, there has been both renewed interest in these biases and an increased effort on the part of the BLS to correct for them. The corrections have been primarily based on two procedures

1. revising sampling procedures for matched samples in ways that make new goods enter the sample faster.

2. using hedonic regression to augment the traditional matched sample procedures in a way that allows for some correction for "quality" change.

Before moving on to the focus of this paper (which is (2)), a word clarifying what we can expect from the revised sampling procedures in (1) is in order. Provided we maintain that the new goods bias in the computed price index is positive, the investments made by the BLS in more timely incorporation of new goods into the price index will ameliorate or exacerbate the new goods bias according as the more timely price index rises at a slower or faster rate than the alternative index. Whether it does or not depends on the details of the structure of the market in which the new good is sold, and is therefore likely to vary from market to market.

For example take a market which, like the personal computer or "PC" market we study below, sees frequent introductions of new products. The new products are "better" than the old products, and therefore gradually obsolete them. At some point it no longer pays the producer of an old product to market it, and the product exits. For clarity say the life cycle of a single product is three periods and that between the first and the second period each product introduced in the first period experiences a percentage price change of g_1 (which may be negative) while between the second and the third the price change is g_2. Now compare two price indices. A "timely" index that picks up the new goods at the date they are introduced, and an index which picks up the new good only after they had been sold for a period. The more timely index records a price change of $wg_1 + (1-w)g_2$, where w is the weight the sampling procedure gives to goods which are new in the base period. The less timely index records a price change of g_2. Thus the difference between the new and old indices is $(g_1 - g_2)w$. So the more timely index would exacerbate or ameliorate the bias according as $g_1 - g_2$ is greater or less than zero. In the home market for PC's we tend to get sharper price falls later on in the life cycle of the product indicating that faster incorporation of the good into a matched model index would exacerbate the bias problem, while research by Cockburn and Griliches (1994) indicates that in pharmaceuticals the opposite is in fact true.

Clearly we want to include all new goods in the index eventually. If we did not we would run the risk of having an index of the prices of goods that we do not care about. However, given that we are using the matched model procedure, it may be appropriate not to introduce the new goods very early on in their life cycle, but rather to wait until after market prices have settled down. Whether we want to wait depends on the institutional details that set pricing patterns in the market of interest. Given this, it is somewhat distressing that for most all commodity groups those details are neither known nor a focus of current research (at the BLS or anywhere else I am aware of), and that despite this we are allocating a large fraction of the resources available for improving our price indices to decreasing the time lag before incorporating new products.

2 Economic Analysis and "Hedonic" Indices.

To begin we consider the economics that underlie hedonic price functions, as this provides the rational for hedonic price indices and the appropriate method of constructing them. Hedonic price functions are empirical summaries of the relationship between the prices and the characteristics of goods sold in differentiated product markets. They were introduced by Court (1939) and revived by Griliches(1961) as a way of ameliorating the new goods problems discussed above.

These authors reasoned that since newer models of goods often had more desirable characteristics, the difference between the prices of the newer and the older models should not be entirely attributed to inflation. On the other hand if we build our price indices entirely from inter period price comparisons of goods sold in both periods, that is if we never compare "old" to "new" goods directly, we will never capture the effect of switching to new goods on welfare. Since a revealed preference argument indicated that the new price-characteristic tuple is preferred to the old, omitting the comparison will bias measured welfare increments downward, and price indices upward. Court and Griliches suggest estimating a surface which relates prices to characteristics and then using the estimated surface to obtain estimates of "quality adjusted" price changes for products with given sets of characteristics.

A series of empirical and theoretical papers using hedonics followed. As now seems natural, the theoretical focus was on examining the relationship between characteristics and prices generated by the equilibria of differentiated product markets (see Rosen,1974, Epple,1987 and, more recently, Anderson, De Palma, and Thisse,1989, Feenstra,1996, and Berry Levinsohn and Pakes,1995, henceforth BLP[2]). The models all require three "primitives":

- Utility functions defined directly on the characteristics of products, rather than on products per se. Typically the preferences for different characteristics are allowed to depend on the attributes of the consuming unit (their income, family size, ...). As a result aggregate demand, obtained as the sum of individual demands, will depend on both the characteristics of the goods marketed and the distribution of the attributes of the consuming units in the market of interest.

[2]A notable exception here is found in the work of Robert Pollak, which is summarized in his book (Pollak, 1989), and which is similar in spirit to the analysis in section 4.1 below.

- Cost functions. Typically these determine marginal costs as a function of the characteristics of the good, factor prices, the scale of production, and, "productivity".
- An equilibrium assumption. This determines prices (and quantities) given demand and costs.

For our purposes the big advantage of these models is that they allow us to compare new goods to old rather directly; we simply compare the value consumers attach to the characteristics of the old good to those of the new. The problem is that to use this comparison in the construction of our price indices we need estimates of the distribution of preferences over characteristics. Obtaining those estimates is a complicated task that, at least with current resources, requires a number of a priori assumptions[3].

As is now explained, the advantage of hedonics is that, conditional on our framework, they can be used to provide a bound on the true cost of living index that is both *independent* of these detailed assumptions, and relatively *easy* to construct.

2.1 The Economics of Hedonic Regressions.

We begin with the relationship between prices and characteristics implied by economic theory. That relationship depends on the equilibrium assumption that is appropriate for the market of interest. For simplicity we consider only two such equilibria (i) a marginal cost pricing equilibrium, and (ii) a Nash in prices, or Bertrand equilibrium. However, as will become clear, any equilibrium that generates mark-ups could replace the Bertrand assumption and similar comments would apply.

To this end, let (x_i, p_i) denote the vector of characteristics and the price of good i, (x_{-i}, p_{-i}) denote the characteristics and prices of the other goods

[3]See Berry Levinsohn and Pakes,2000. For examples of price indices build in this way see Pakes, Berry and Levinsohn, 1993, and Nevo,2000. Though I do not think it is practical to expect the BLS to use these procedures to compute "exact" official price indices, it would be worthwhile to use them to construct a number of "exploratory" indices (e.g. for commodity groups with high rates of technological change and whose major characteristics are easy to measure, like the PC market discussed below). This would give us some indication of; the bias in more traditional indices, the extent to which hedonics bridges the gap, and the data needs and difficulties with using the new techniques.

marketed, and write the demand for good i as

$$D_i(\cdot) = D(x_i, p_i, x_{-i}, p_{-i}; A)$$

where A indexes the attributes which determine the distribution of consumer preferences over characteristics. Also assume that marginal cost is constant and given by $mc(x)$ and that all firms are single product firms.

Then, if we introduce the indicator function $\chi(B)$ which equals one in the Bertrand pricing equilibrium and zero when prices equal marginal cost, the pricing functions from our two different equilibria can be expressed as

$$p_i = mc(x_i) + \chi(B)\frac{D_i(\cdot)}{|\partial D_i(\cdot)/\partial p|}, \tag{2}$$

where the second term, $\frac{D_i(\cdot)}{|\partial D_i(\cdot)/\partial p|}$, is usually referred to as the mark-up; it varies inversely with the elasticity of demand at the point.

The hedonic function, say $h(x)$, is just the expectation of price conditional on x. From (2)

$$h(x_i) \equiv E[p_i|x_i] = mc(x_i) + \chi(B)E\left(\frac{D_i(\cdot)}{|\partial D_i(\cdot)/\partial p|}|x_i\right). \tag{3}$$

That is, in the marginal cost pricing equilibrium the hedonic function *is* the marginal cost function. However in the Bertrand equilibrium the hedonic function is the sum of the marginal cost function and a function that summarizes the the relationship between mark-ups and characteristics.

The marginal cost function is a "primitive" of the problem. Thus we might, for example, have priors on the signs of its coefficients or on their stability over time. The equilibrium mark-up, however, is derived as a complex function of

- the characteristics of the competing products and
- the distribution of consumer preferences.

So the coefficients in the hedonic equation should be expected to change when either of these "primitives" change, and their sign may be different than the sign in the cost function. The existence of a mark-up and the fact that it is a complex function of underlying primitives would result from any equilibrium other than marginal cost pricing (it does not depend upon the "static" Nash in prices assumption above).

Given our interest in new goods, I would argue that we can rule out the assumption of marginal cost pricing as *a priori* irrelevant. The new products we study are largely a result of prior investments in developing and marketing. It is only sensible for a firm that made these investments to price the new good at marginal cost if it is faced with an infinitely elastic demand curve. We do not expect infinitely elastic demand curves in differentiated product markets unless two products with identical characteristics are developed. This, in turn, is an unlikely equilibrium of reasonable models for the development of new products, and when it does happen it probably would not persist (there would be an incentive for firms to modify their product). Moreover the fact that there are no mark-ups in a marginal cost pricing equilibrium implies that in that equilibrium all producers are losing the money that went into the development and marketing of their new products (at least provided the new goods are not subject to sharply increasing costs to production, a situation thought unlikely in the early years of a product's life). As a result we would not expect a market based on marginal cost pricing to persist [4].

Put differently, since mark-ups generate the incentives to develop new products, it is hard to imagine a model designed to analyze new products which doesn't allow for them. This implies that we should expect the regression of price on characteristics to contain the relationship between mark-ups and characteristics. Two direct implications of this are:

1. We should expect hedonic functions to change over time. The changes are likely to be sharper in time periods when there is substantial technical change embodied in new products. This because their will be an attempt to direct new products to parts of the characteristic space where mark-ups had been high, and these products will, when they enter, decrease the mark-ups of the products that had been leading in that part of the characteristic space. In our computer example, the introduction of PC's with a new fast chip cause a noticeable decrease

[4] There are good reasons for producers of new goods to price below marginal cost when production is subject to learning by doing, or when consumption today spreads information which shifts out the demand curve in future years. However in these cases the static Bertrand equilibrium which generates (2) is inappropriate. Instead prices charged today will depend on the impact of current price changes on future costs (or on future demand) as well as on marginal costs and the current mark-up. Then the mark-ups which justify the development of the products will be garnered at a later stage in the product's life cycle. For a discussion of when the static Nash in prices is likely to be irrelevant, see Pakes,2000.

in the mark-up on the machines that had been the speed leaders in the prior period. For similar reasons hedonic functions for the same products may well be different in markets with different distributions of consumer attributes.

2. We should not expect the hedonic function to have the properties we might expect the cost function to have. In particular there is usually no economic reason for the hedonic function to generate a positive relationships between price and the characteristics of products that we generally think of as desirable.

(2) needs some clarification. It is true that if we compare prices of two products whose characteristics are identical except for say the first characteristic, and the first characteristic is a "vertical" characteristic (i.e. a characteristic everybody would like more of), then if our framework is right, the price of the first good should be larger than the price of the second good. The reason this is true is because in the case when we have products that are the same except in one characteristic, we can rely on utility theory to determine the sign of the differences in their prices (for a proof see Bajari and Benkard,2001). However there are two empirically important problems in pushing this line of reasoning.

First in a world where there are real sunk costs to the development of new products, not all parts of the characteristic space are "filled in". Then the reasoning that leads to positive coefficients on vertical characteristics breaks down[5]. Indeed once we acknowledge the fact that there are not a "continuum" of products marketed, then the properties of the hedonic regression function depend on factors that have no direct impact on the utility of *any* consumer; for e.g. which products are owned by which firms.

Second, and perhaps more important, when we use a characteristic model to provide an explanation of behavior we often have to go to a more complete characteristic model that the analyst who estimates hedonic functions cannot use. The more complete model has consumers producing utility from combining product characteristics with their own, individual-specific attributes

[5]In particular when development costs are a function of the characteristics of the products, product development decision often require trading off one vertical characteristic against another, and mark-ups can easily generate hedonic functions with a "wrong" sign. For example, among autos of a given year size and miles per gallon (m.p.g.) are negatively correlated, and a hedonic regression which includes both often produces a negative coefficient for m.p.g. (see, for e.g. Table 3 in BLP).

(see Pollack,1989, and the literature cited their for a deeper discussion). The production function the consumer uses in this process would generally differ with the consumer's attributes and the characteristics that utility is defined over (the "utility" characteristics") would be different than the characteristics embodied in the goods bought (the "input characteristics"). In particular a given vertical utility characteristic might be produced with different inputs by consumers with different attributes. Often the best the analyst can do is find some summary measures (e.g. the average) of the values of the utility characteristic produced by a given "input", and this is used as a "characteristic" of that input in a hedonic analysis of input prices. The reasoning that leads one to expect a positive effect of a vertical "characteristic" on price requires the assumption that we can order the inputs by the amount of the vertical utility characteristic they produce, and that this ordering is the same for *all consumers*. There are many cases when this assumption is *not* satisfied. When it is not there is no reason to expect that a product which, say on average, produces more of the vertical characteristic than a competitor will have a higher price.

An example, taken from Cockburn and Janis' (1998) study of drugs for the treatment of Rheumatoid Arthritis, will clarify this point. Assume people sick with this disease care about their ability to grip objects and their overall health (we include the health index because the drugs that treat Rheumatoid Arthritis can have toxic side effects). The drugs actually marketed are defined by their content of various chemicals, and the transformation from chemicals to "grip ability" and "overall health" varies by individual. We simplify slightly and assume that their are two types of drugs that treat Rheumatoid Arthritis (type A and type B). The NIH performed a series of tests and they showed that for the vast majority of patients drug type A is as effective as drug B and is far less toxic. In particular the "toxicity" rating of drug A (measured as fraction of people who the drug causes serious harm to) is essentially zero. On the other hand Drug B is effective on most of the, say, 5% of the population that drug A does not help, but has a toxicity rating of .7.

Once the results of these tests are made public, the drug companies rush to produce different versions of the type A drug. Indeed companies keep entering into that part of the market until the expected discounted value of profits from marketing such a drug falls below its development costs. The large number of producers of type A drugs generates competition in that drug's market and this forces down mark-ups on drugs of type A. Of course the large number of patients consuming these drugs implies that the firms that produced it are still able to cover their development costs (even given the small mark-ups). The market for drugs of type B is too small to support more than one firm, so that producer sells its product at a "monopoly" price.

The marginal cost of production is similar for both types of drugs, so the hedonic function will be largely determined by differences in mark-ups. Since mark-ups are higher for the more toxic drug, we should expect the hedonic regression of price against efficacy and toxicity to pick up a strong positive coefficient on toxicity. This is precisely what Cockburn and Anis (1998) find. There is nothing "wrong" with his result, indeed standard economic arguments lead us to expect it. In particular it does not mean that the market prefers more toxic drugs; it just means that profits were such that entry drove down mark-ups more on the less toxic then on the more toxic drugs – not an unreasonable finding at all.

Hedonic regressions have been used in research for some time and they are often found to have coefficients which are "unstable" either over time or across markets, and which clash with the naive intuition that characteristics which are generally thought to be desirable should have positive coefficients. This intuition was formalized in a series of early models whose equilibrium implied that the "marginal willingness to pay for a characteristic equaled its marginal cost of production". I hope this discussion has made it amply clear that these models can be *very misleading*. The derivatives of a hedonic price function should not be interpreted as either willingness to pay derivatives or cost derivatives; rather they are formed from a complex equilibrium process.

3 Hedonic Price Indices.

Of course just because we can not interpret the coefficients from hedonic regressions does not mean that these regressions cannot provide information on how to construct an appropriate price index. We turn to the task of using hedonic regression functions to construct price indices now.

3.1 Bounds From Hedonic Regression Functions.

The theoretical rational underlying the hedonic lower bound to an exact price index is deceptively simple. Let $h^t(x)$ be the hedonic function in period t (as defined in equation 3), and C_t be that period's choice set (the list of the

x tuples that can be purchased in t). We designate the base period to be period one and the comparison period to be period two. The base period (or "Laspeyres" type) hedonic adjustment to the second period's income of a consumer who purchased x_1 in period one is

$$h^2(x_1) - h^1(x_1). \tag{4}$$

This is the change in the base period's income that would allow the consumer to buy the same good in period two that it bought in period one and still have the same amount of income left over to buy other goods. So this compensates the consumer for price changes by insuring the consumer the possibility of buying the same goods in the comparison period as were bought in the base period.

The theoretical point to note is that since the new choice set is generally different from the old, we know that *provided x_1 is available in period two* (i.e. $x_1 \subset C_2$), a consumer with income $y + h^2(x_1) - h^1(x_1)$ facing the second period's choice set and hedonic surface (h^2, C_2) will generally choose a different x than x_1, and any such choice will result in utility greater than the utility from the period one choice (since this is still feasible choice). Thus we have the qualitative result we are after:

The base period hedonic adjustment (or 4) is too generous to the consumer.

It is important to note that the only assumptions we require for this argument is that $x_1 \subset C_2$ and that the marginal utility of income is positive for everyone. Precisely the same assumptions (and proof) show that the traditional Laspeyres index is an upper bound to the compensating variation in models where utility is defined directly on products. This argument dates back to Konus (1924) and remains the predominant justification for the indices currently in use.

I will call any index which provides an upper bound on the compensating variation independent of the form of the utility function a "proper" price index, and will be concerned primarily with proper indices from models where utility is defined on the characteristics of products, rather than on products per se.

3.2 Hedonic vs. Matched Model: Conceptual Issues.

We begin with the "conceptual" differences between the two indices. These are the differences that would arise if the hedonic and the matched model indices were computed from the same base period data, and measured characteristics fully determine the utility a consumer derives from the good. The next section considers "practical" issues in the construction of the two indices. This includes differences in their; data requirements, costs, and likely precision.

In this subsection, then, we compare hedonic and matched model indices constructed from the same base period data and assuming that measured characteristics fully determine the utility a consumer derives from the good. Each of the base period goods are either available in the comparison period or not.

Consider first the estimated price changes for goods that have $x_1 \in C_2$. In these cases the hedonic lower bound argument given above is air tight, but so is Konus'(1924) lower bound argument for the Laspeyre's matched model index. Moreover given enough data the two would lead to precisely the same price change. So for goods that do not exit the only difference in the price change computed by the hedonic and the matched model indices is their estimates of $h^t(x)$, a topic I return to below.

The major conceptual difference in the two indices arises as a result of their treatment of cases where $x_1 \notin C_2$. The matched model index is simply *not defined* in these cases, so what practitioners do is drop the good in question and link in another. This generates a selection problem. The goods that are sold in the first period and are not in the next tend to be goods whose measured characteristics were desirable relative to those of the other products sold in the first period, but were not in the next. Consequently they tend to be the goods which are intensive in characteristic whose values have fallen more than the average (often due to the entrance of products with superior characteristics)[6]. As a result

the matched model index construction procedure tends to throw

[6]There are, of course, other reasons that induce firms to pull a good off the market. These include differences in costs of production and (once we allow for multi-product firms) ownership patterns over time. Note that either of these could generate a situation in which one good is pulled off the market while another good with the same consumer characteristics is not.

out the left tail of the distribution of price changes and produce a price index which is biased upwards.

Note that this bias occurs even if there are *no* unobservable characteristics[7]

In rather stark contrast to the matched model procedure, the hedonic procedure is the *same* whether or not $x_1 \in C_2$. If $x_1 \notin C_2$, the hedonic estimate of the price in the second period, or $h^2(x_1)$, is obtained as a weighted average of the prices of products whose characteristics *are marketed* in the second period with the weights being larger for those products that have characteristic vectors similar to x_1 (the precise form of the weights depends on how one estimates the hedonic function).

Of course to rationalize using $h^2(x_1) - h^1(x_1)$ as an upper bound to the equivalent variation when $x_1 \notin C_2$ we need an additional assumption. The additional assumption that is required is an assumption which guarantees that the consumer will be better off with $h^2(x_1) - h^1(x_1)$ added to its income and the choice set in period two then with the original income and the period one choice set. If there are goods in the choice set in period two that are not too different from x_1 the assumption required is that their exists a direction of movement in characteristic space from x_1 to a product which *is* marketed in period two, where the willingness to pay for the implied changes in characteristics is greater than the slope of the hedonic function. Whether this is likely to be a problematic assumption depends on both the nature of the characteristics, and on whether there are goods available in period two which have characteristics that are close to those of x_1.

A case where it is not problematic is when the important characteristics of the product are vertical, so that all consumers prefer more of them (this is

[7] This is the difference between the selection problem here and the econometric analysis of selection that dates back at least to Heckman's (1974) classic study of selection in estimates of labor supply, and Hulten and Wykoff's (1981) use of selection corrections in their analysis of depreciation rates. Similarly though the issue raised in Olley and Pakes' analysis of the role of exit in studying the impact of a policy or environmental changes on firm performance (firm's who were negatively effected by the change tended to exit), the initial measurement problem was different. In the Olley and Pakes(1995) example the focus was on working out the implications of the selection problem on the estimation of a vector of parameters describing a primitive of the problem (the production function), and then using those estimates to compare the average level of an outcome (productivity) over time. Like in Heckman's (1974) and Hulten and Wykoff's (1981) case, selection induced relationship of unobservables to observables that generated the bias in the production function coefficients and subsequent averages. If selection were only a function of observables, there would have been no bias.

the case for our PC example below). In this case we can modify the hedonic estimate for the second period to give individuals who bought the exiting good the cost of a good which *is* available in the second period and is at least as valuable to the consumer as the good that exited. This insures that our hedonic index is proper in spite of the fact that $x_1 \notin C_2$.

So our assumption is only problematic for goods with important "horizontal" characteristics (goods where one person prefers more of a characteristic and another prefers less of it). Even in this case we would only expect a good to exit if there were a very small number of consumers who prefer it and a relatively high fixed costs of producing it – since otherwise it would be profitable to market the good[8].

I would argue that in a reasonably well working market the impact of this bias is likely to be small, and if there is a bias in a particular price it is most likely to be for a good whose characteristics are not inside the "span" (or the range of characteristics) of the goods available in C_2 (since then there would be no feasible direction of movement for consumers who had a strong preference for the characteristic which the exiting good was very intensive in). More precisely if we let $x_i = [x_{i,j}]$ where j indexes the various attributes of the good, and $\overline{x}_j = max_{i \in C_2} x_{i,j}$ while $\underline{x}_j = min_{i \in C_2} x_{i,j}$, then we might worry most about a bias for goods where

$$x_1 \notin \Pi_j [\underline{x}_j, \overline{x}_j] \equiv \overline{C}. \tag{5}$$

For this reason we will present two hedonic calculations in our empirical study of PC's. One will be a proper hedonic assuming that the additional assumption needed for the lower bound is satisfied for all the goods that exited in period two. The second will assume that the additional assumption is only satisfied when $x_1 \in \overline{C}$. For goods where this condition is not satisfied we will do exactly what the matched model index would do; we drop them and omit them from calculation from the index. Throughout the remainder of the paper, however, I assume that at least one of these indices is a proper upper bound to the appropriate compensating variation[9].

[8] An example might be an old model of an auto (e.g. a model "T" Ford) that antique car buffs might be willing to purchase at its old price (adjusted for inflation), but that would not sell sufficient quantity at that price to cover the sunk costs of producing it.

[9] I should note that since in the PC case most of the relevant characteristics are vertical, the fact that there is not too much difference in the two indices for our case, should not be extrapolated too far.

3.3 Hedonic vs Matched Model: Practical Issues.

The practical questions associated with constructing the indices are intricately tied up with data generation issues. I begin by assuming we use the current data generation process. This was designed for the matched model index, but as will be explained presently, can be used for estimating hedonic indices as well. I then briefly consider how one might alter the data generating process to take fuller advantage of the properties of hedonic indices.

Data Needs.

The most important point to note here is that characteristic data play exactly the same role in obtaining the data needed to construct the matched model index as they do in constructing the hedonic index. In both cases they define the product whose price is being compared over the two periods.

Matched model indices are built from data provided by data collectors who collect data on price changes of a particular good at a particular outlet (both of which are chosen by a sampling process). There are two points where product characteristics come into play. First there is the list of characteristics that define the characteristic tuples (products) that the data collector is allowed to sample from when obtaining an entry level item for a given commodity group (for a discussion of this process, see Triplett,1961). After sampling from an item from this set, the data collector then writes down the values for a second (often more detailed) set of characteristics for the entry level item actually sampled. This second list allows the data collector to recall the "same" good, or at least a good with the same listed characteristics, when the outlet is revisited in the reference period.

If a good with the listed characteristics is found when the data collector revisits the location in the reference period, the price of that "good" is recorded, and the ratio of the second period to the first period price becomes the "price relative" for the good. These price relatives are averaged to obtain the index for the particular commodity group. If no good with the listed characteristics is found in the reference period, that "good" is dropped from the list of goods used to form the index for that commodity group. Clearly then

the matched model indices are formed by comparing prices of characteristic tuples, not by comparing the prices of "goods" *per se, just as hedonic indices are.*

Let x^m be the set of characteristics the data collector writes down. One solution to the problem of determining the characteristic to be used in the hedonic analysis, say x^h, is to set $x^h = x^m$. This would insure that any problem in defining the characteristics for the hedonic index would also be a problem in defining the matched model index, and vica versa. In reality the BLS has been choosing the "relevant" characteristics (the x^m) of products for years and it has never been a terribly contentious issue. At least for PC's the the list is quite detailed and includes: brand, model number, amount of RAM, size of Hardrive, CPU, speed in megahertz, and dummies for; CD-rom, cd-writer, zip-drive, monitor (brand, model, and size when there is one), keyboard, sound-card, speakers (description), extra software (office, ...), video card (premium or not), floppy drive, and modem (fax capabilities or not). Most of us would be happy to use the same set of characteristics for the hedonic study.

There would remain however one difference between x^m and x^h. The BLS also implicitly conditions on the outlet at which the purchase took place. The characteristics of the "sale", as well as the characteristics of the product purchased, might matter to the consumer, and though the outlet of purchase does not fix the "characteristics" of the sale , it might well proxy for them (for e.g. it does not fix; the time required to make the purchase, the time of day and day of the week of the purchase, and the other amenities offered at the outlet at the time of the purchase like the array of other goods and their prices). If we were to use hedonics then we ought probably to add to the list of characteristics of goods the data collector currently collects some of the characteristics of the outlet (is it a factory outlet, department store, specialty store....). We could then condition the hedonic on outlet characteristics, and largely eliminate this source of difference in the two indices. The discussion of this section will assume this one, fairly simple, alteration to the data gathering procedure.

Matched Model, Hedonic, and Hybrid Indices from Current Data.

This section will consider differences in indices when all of them must be constructed from the data currently gathered (which, recall, includes characteristic data). Our comparisons will focus on two issues:

1. Differences between the expectation of the upper bound provided by the index and the true compensating variation (or the "bias" of the indices), and
2. Differences in the variance of the indices which, in turn, consists of
   (a) traditional sampling variance (the population average of price relatives is different from their sample average) and
   (b) the variance associated with using an estimate of the comparison period price.

Recall that the differences between hedonic and matched model indices based on the same data are;

(i) for the goods in the sample in both the base and the reference period, the matched model index would form price relatives from the actual observations, while the pure hedonic would substitute an estimate of its second period price for its actual value, and

(ii) the matched model index would simply drop the comparisons for goods which were present in the base period but not the comparison period, while the hedonic index would use an estimate of the second period value.

(i) creates estimation variance (2b above) in the hedonic index which is absent from the matched model index. Letting $\hat{h}(\cdot)$ be our estimate of $h(\cdot)$, the estimation error in the hedonic is $\hat{h}^{t+1}(\cdot) - h^{t+1}(\cdot)$[10]. (ii) creates a positive selection bias in the matched model index (1 above) which is absent from the hedonic index. In addition (ii) implies that the matched model index will use a smaller number of price relatives than does the hedonic index, and this could increase its sampling variance (2a above) relative to that of

[10] There is also a difference between the two comparison period price estimates caused by the disturbance in the regression function (the ϵ). There is a question here of whether we want to consider ϵ as a disturbance about a price we want to measure (say $h(x)$), or to take ϵ as a result of unobserved characteristics of the product we want to condition on in our comparisons. In the former case we would like to use the base period hedonic ($h(x)$) as the denominator in the formula for all price relatives. For simplicity I am simply going to follow the BLS procedure and use the actual base period price in the denominator throughout this paper. Given this choice and the fact that the base and reference period ϵ's are correlated, the inclusion of ϵ in the comparison period price does not necessarily increase the variance in price relatives (see below).

the hedonic index[11]. Thus, perhaps not surprisingly, the tradeoff between the matched model and proper hedonic indices comes down to a tradeoff between the larger bias of the matched model index and the possibility of a larger variance in the hedonic index. Consequently it is hard to discuss this tradeoff without empirical magnitudes in mind, so we come back to it in our discussion of the empirical example below.

For comparison to currently used procedures, I also want to introduce a hybrid index constructed as an average of

- the observed price relatives for goods which continue to be sold in the comparison period, and
- the hedonic estimate of the price relative for the rest of the goods.

The apparent advantage of the hybrid is that it does not incur the additional estimation variance of the hedonic for the goods which are sold in the reference period, and does not incur the selection bias of the matched model index for the goods which are not.

We will call this index the "complete hybrid" index. Use of the term "complete" is to distinguish our hybrid from a "cousin" hybrid which is often used by the BLS when it does produce a "hedonic-like" index. I will consider the BLS's price index for PC's in detail in the empirical section of this paper. For now suffice it to note that the BLS's hybrid differs from ours in two ways; they do not obtain *any* price relatives for some of the goods sold in the base period, and they do not use our hedonic estimate of the reference period price for the goods they do adjust by hedonics.

Can we say anything about the relationship between the means and/or the variances of the complete hybrid and the proper hedonic? Recall that the comparison period price is $p_{t+1} = h^{t+1}(x) + \epsilon_{t+1}$ with $E[\epsilon|x] = 0$. The proper hedonic uses $\hat{h}^{t+1}(x)$ for the comparison period price of all goods. The complete hybrid hedonic uses $h^{t+1}(x) + \epsilon$ for those goods which do survive, and $\hat{h}^{t+1}(x)$ for those who do not. Consequently, if we let χ_i^{t+1} be the indicator function which takes the value of one if the good is sold in period $t+1$, the difference between the complete hybrid and the proper hedonic is

$$\frac{1}{n_t}\sum_{i\in C(t)}[1-\chi_i^{t+1}][h^{t+1}(x_i) - \hat{h}^{t+1}(x_i)] \; + \; \frac{1}{n_t}\sum_{i\in C(t)}\chi_i^{t+1}\frac{\epsilon_{i,t+1}}{p_i^t}, \tag{6}$$

[11] It could also be the case that the sampling variance of the price relatives of the goods selected on survival is different than the overall sampling variance.

where n_t is the number of products sampled in period t.

Assuming that the average estimation error is approximately zero, the expectation of the difference in (6) is

$$E \frac{1}{n_t} \sum_{i \in C(t)} \chi_i^{t+1} \frac{\epsilon_{i,t+1}}{p_i^t} \frac{1}{n_t} \approx E \sum_{i \in C(t)} \chi_i^{t+1} \frac{\epsilon_{i,t+1}}{h^t(x_i)} \left[1 - \frac{\epsilon_{i,t}}{h^t(x_i)} \right],$$

where the approximation sign indicates we have used a first order expansion of the $\epsilon_{t+1}/[h^t(\cdot) + \epsilon_t]$ from the point $\epsilon_{t+1}/h^t(\cdot)$.

To evaluate this expression we need a model for both the evolution of ϵ over time and for selection, i.e. for the event $\chi^{t+1} = 1$. If we thought of ϵ as being determined by unmeasured product characteristics, we might be willing to assume that goods with higher values of ϵ will be more likely to survive. However it is not clear whether survival is more closely related to $\epsilon_{i,t}$ or $\epsilon_{i,t+1}$, and, at least to a first order approximation, they enter the extra bias term in the complete hybrid with different signs. As a result even the sign of this term is not clear. A similar argument shows that despite the fact that their is additional estimation error in the hedonic, the difference in the variance of the proper hedonic and the complete hybrid cannot be signed. We return to the comparison between the complete hybrid and the proper hedonic in the empirical section below.

Alternative Data Sources.

All indices discussed in this paper are weighted averages of price relatives for a given commodity group. Perhaps the most important distinction among data generation procedures is how the weights are obtained. Currently the weights are implicit in the BLS's sampling process. This first samples an outlet and then sends a data collector to that outlet to sample a good from the goods with a set of prespecified characteristics. The data collector then revisits the sampled outlet in the next period and looks for the price of a good with the same characteristics as the good sampled in the first period. The repete visit procedure which generates the price relatives is quite expensive. As a result the indices are typically based on a relatively small number of price quotes; about one hundred and fifty for the PC index discussed below.

One alternative would be to base the weights on explicit quantity or sales data, and then form price relatives from price and characteristic data that can generally be obtained from any of a number of other sources (store catalogues, Internet web sites, scanner data ....). For the commodity groups for which reliable quantity or sales weights are available, the alternative would provide a data gathering procedure which is likely to be quite a bit cheaper than the current data generation procedure[12]. A deeper discussion of these possiblities is, however, beyond the scope of this paper. Here I only consider alternative data generation procedures that keep the current sampling process.

The estimation error in the hedonic index arises from errors in estimating the hedonic function, i.e. from $\hat{h}(x) - h(x)$. We can reduce the variance in $\hat{h}(x) - h(x)$ by using data from alternative sources on price characteristic combinations when we estimate the hedonic function. For e.g. the BLS uses data downloaded from Internet web sites to estimate their hedonic function for PC's, and, as a result, their estimated hedonic is based on eight to twelve hundred observations (in contrast to the 150 price relatives that go into their sample of price relatives). So this is one change that is relatively easy to make and likely to be useful.

A more basic change in data generating procedures is also possible if we were to switch to the proper hedonic. We could let data collectors determine entry level items just as they do now, but not require the data collector to find the price associated with the characteristics of the entry level item in the second period. Instead every year the data collector would collect price data for a set of different entry level items, and the second period price for the current entry level items would be determined by the hedonic regression function. This would cause at least two important changes in the index.

- First it would allow for one-hundred per cent replacement every period, and as a result the implicit weights that go into the index would *fully* reflect base period sales patterns. In contrast the partial rotation and forced substitution methods that define the weights for the current matched model procedure attempt to use a particular average of sales weights over past periods, but in fact use weights that are distorted in a non-random and poorly understood way by forced substitutions.

- There would be a change in the the cost per quote. In particular

    - the BLS would have to obtain permission from store owners to

[12]We would no longer need to send data collectors to outlets. Note that the alternative would also generate the data needed to construct more ambitious experimental indices based on direct estimates of the distribution of preferences, see the discussion in section 3.

sample from new set of stores every period, and this would increase cost,

- but we would not have to send the data collector out to exactly the same outlet in the second period, and the data collector would not have to search for a product with the same listed characteristics as those sampled in the previous period, and this would decrease cost.

There is an open question of whether this alternative, which is only available for hedonic indices, and which has a clear advantage in that it produces an index based on recent weights in every period, is a cost effective way of improving our indices.

A Warning.

We have shown that we can produce a hedonic index which would eliminate the selection bias in the matched model index from the exact same data used now. In addition there are a number of ways of changing the data gathering procedure to make hedonic methods even more attractive. As a result I hope it is clear that hedonic methods of constructing price indices deserve increased attention. Still it is important to conclude this section by stressing that the distance between the upper bond provided by the hedonic index and the least upper bound we are after – say the income transfer that would insure that the average consuming unit would be at least as well off in the second period – may be quite large.

What the hedonic does is allow us to compute price relatives for goods that exit, thus correcting for the selection bias in matched model indices. It makes no adjustment for the value created by the entry of new goods per se. As a result it is possible to point to numerous aspects of the new good problem that increase consumer utility but will never be reflected in either the hedonic or the matched model index. Thus neither index will reflect any gains from increases in the variety of the goods offered in a certain characteristic range. Nor will either index register any gain as a result of the introduction of a good which truly expands the range of characteristics offered, thereby enabling us to use goods in the same commodity group in new and more productive ways. For example when the laptop computer was introduced its closest competitors were desktop machines. The desktops had more speed, storage capability, and reliability, than laptops and a significantly lower price.

However the significant reduction in size and weight in a laptop was highly valued by consumers. None of the indices discussed here could ever pick up the increase in utility that the laptop generated (for a more detailed discussion see Pakes,1995). To capture these effects we need a more complete model of household utility.

That is there is good reason not to expect the upper bound provided by the hedonic to be particularly close to the least upper bound we are after. Moreover the difference between the hedonic and this least upper bound is likely to be especially large in technologically dynamic industries.

4 The Statistics of Constructing An Index.

The next section computes and then compares several different price indices for PC's. These indices are all weighted averages of estimated price relatives. The literature on price indices distinguishes between "plutocratic" weights, which are weights that are proportional to sales, and "democratic" weights, or weights that are proportional to the quantity sold. The weights used by the BLS in their indices are determined by the BLS's sampling procedures which make the probability of sampling different entry level items proportional to the items' fraction of sales in the given period. The BLS therefore produces a plutocratic index. For comparability, we will do the same.

Consequently the alternative hedonic indices we present will all be constructed as

$$I_t = \frac{\sum_{i \in C_{t-1}} h^t(x_i^{t-1}) q_i^{t-1}}{\sum_{i \in C_{t-1}} p_i^{t-1} q_i^{t-1}} \equiv \sum_{i \in C_{t-1}} \left(\frac{h^t(x_i^{t-1})}{p_i^{t-1}} \times \frac{p_i^{t-1} q_i^{t-1}}{\sum_{i \in C_{t-1}} p_i^{t-1} q_i^{t-1}} \right), \tag{7}$$

or,

$$\overline{I}_t = \frac{\sum_{i \in \overline{C}_{t-1}} h^t(x_i^{t-1}) q_i^{t-1}}{\sum_{i \in \overline{C}_{t-1}} p_i^{t-1} q_i^{t-1}},$$

where, as above, $h^t(\cdot)$ is the t^{th} period estimate of the hedonic function, and $\overline{I}$ differs from I in that it does not attempt to compute price relatives for goods which are sampled in the base period but whose characteristics are outside of the range of the choice set in the comparison period.

Many alternatives estimates of $h^t(\cdot)$ have been used in the literature. They differ in functional form, in right hand side variables, and in procedure used to predict price conditional on functional form and right hand side variables. I begin with a brief review of issues that arise in making these choices, and then go directly to a comparison of alternative indices for our PC example.

Functional Forms

Typically linear or log-linear functional forms are used when estimating hedonic functions, and we will follow suit. There are, however, two points which are frequently forgotten in the literature, and which have distinct implications for how the index should be used.

First even if the log-linear functional form works out to be preferable (and it frequently does), we have to keep in mind that our goal is to predict p ; *not* to predict $log[p]$. In particular there are two ways of obtaining a consistent estimate for $h(x)$ for the log-linear specification. One is to use nonlinear least squares and estimate

$$p = A\Pi_j x_j^{\alpha_j} + \epsilon^1.$$

In this case our estimate $\hat{h}(x)$ is given by

$$\hat{h}(x) = \hat{A}\Pi_j x_j^{\hat{\alpha}_j},$$

where a hat over a variable indicates its estimated value.

The other, which is computationally easier and as a result used quite frequently, is to linearize the logarithmic model and use ordinary least squares to estimate

$$log[p] = a + \sum_j \alpha_j log[x_j] + \epsilon^2.$$

In this case the prediction equation should be

$$\hat{h}(x) = exp[\hat{a}]\Pi_j x_j^{\hat{\alpha}_j} Eexp[\epsilon^2]. \tag{8}$$

I.e. if we use the linearization and O.L.S. we must correct the estimated price index for the variance in the disturbance of the model (our $Eexp[\epsilon^2]$). If we do not make that correction we will obtain a biased estimate of the second period prices, and hence a biased price index. This point is also made by Berndt (1991) and can be quantitatively important (see below).

The second point about functional forms was already emphasythed in the theoretical discussion (section 3). There we noted that we *expect* the hedonic function to be *unstable* over time. If, given this knowledge, we still choose to restrict the coefficients in the hedonic equation to be the same in different periods, then our index no longer has an expectation which is a lower bound for the true price index. Of course we may be willing to incur this "bias" if it is more than compensated for by a decrease in variance of our estimated price index (see below). However this variance-bias trade off is the only reason for constraining coefficients over time and should be explicitly evaluated.

Choice of Regressors.

As stressed in the theoretical discussion we have neither any intrinsic interest in, nor any interpretation for, the coefficients obtained when we estimate the hedonic function. Consequently it is a *mistake* to choose regressors on the basis of the precision, stability, or the "sign" of any particular subset of those coefficients. This is true whether the additional regressors are different characteristics than those already included, or just increasingly complex functions (say higher order terms) of the characteristics that are already in the regression.

We do care about the precision of the price index those coefficients generate, and our choice of included right hand side variables could logically be determined by that. It would be sensible, for e.g., to chose regressors by whether they reduce the mean square error (the squared bias plus the variance) of the estimated price index. In the special case where we use a linear hedonic function it is easy to express the price index in (6) as a linear combination of the characteristic coefficients estimated in the hedonic function. As a result it is not difficult to ask whether either including a particular regressor, or constraining coefficients to be the same over time, would reduce the mean square error we are interested in. The problem is that the answer depends on unknown parameters (in particular on the true values of the coefficients and of their variance- covariance matrix).

For example, if we add a regressor we will increase or decrease the mean square error of an estimated price relative according as the "true" $t - value$ of that regressor (the true value of the coefficient of that regressor divided by its true standard error) is greater than one. The true t-value is never observed but the estimated $t - value$ increases when the adjusted R^2 from the regression goes up. This gives us some indication of when we might want to add a regressor to the model (when it increases the adjusted R^2 substantially), but since the estimated t-value is not an unbiased estimate of the true t-value it does not give us a very precise answer to our question.

Moreover once we allow for the log-linear regression function that our data ends up preferring (see below), the mean square error of the price index becomes a rather complicated function of the entire joint distribution of the coefficients.

¿From the point of view of the BLS it may be worthwhile to design a Monte Carlo study which would provide a precise answer to the question of which regressors to include[13]. What we do in the example which follows is provide the adjusted R^2's, and then focus on estimates of the price index for the two specifications that the data indicates are likely to be informative of the range of possible index values.

Variances.

Recall that there are two sources of variance in the hedonic index; sampling variance and estimation variance. Our sample characteristics differ from those of the BLS's sample (see below), and the impact of those differences on the sampling variance is likely to differ from their impact on the estimation variance. Consequently we provide an estimate of both the overall variance in the hedonic index and the variance due just to sampling errors.

All variance estimates are obtained from bootstrap estimation procedures. To derive our estimate of the total variance we draw a random sample with replacement from every period in our original data (with sample size equal to the actual sample size). This bootstrap sample is then used to estimate new hedonic regressions for each period. The predictions from this hedonic regression, together with the bootstrapped sample of prices and characteristics, are then used to compute the hedonic price index as in (7). We do this many times and then compute the variance in the hedonic price indices over the alternative bootstrapped samples. If our sample were large enough, this estimate would converge to the total variance generated by both the sampling and the prediction exercises inherent in the construction of our index.

The second variance we want to estimate is the variance that is just due to sampling error. To calculate this variance we begin by drawing random samples with replacement from the original data just as above. Now however; (i) we assume that the coefficients from the hedonic regression in each year are exactly equal to the coefficients we estimated from the true data (i.e. we never re-estimate the hedonic regression), and (ii) we assume that the prices of each

[13]Similarly a Monte Carlo study could be used to ask whether to constrain the coefficients of those regressors over time.

good in *both* the base and the comparison period are equal to the estimate of the price obtained from the good's characteristics and the estimate of the hedonic in the two periods. So the price relative for the i^{th} vector of characteristics drawn in the period $t-1$ bootstrap is $h^t(x_{t-1,i})/h^{t-1}(x_{t-1,i})$. The price index is calculated as the average of these price relatives over the characteristics in the bootstrapped sample. Our estimate of the sampling variance in our index is the variance in this price index across bootstrapped samples. Note that this estimate assumes both that the hedonic surface is correctly estimated and that $\epsilon \equiv 0$. Thus all of its variance is due to variance in the characteristic tuples sampled.

5 PC Price Indices; Data and Regressions.

This section computes hedonic regression functions for PC's sold in the "home market" between 1995 to 1999. The next section will compute and compare alternative price indices for PC's for this period.

Since we do not have the resources to send data collectors to outlets to sample prices, we will have to make due with somewhat less precise and much less frequent data than the BLS has at their disposal[14]. As a result some aspects of our comparisons across indices are likely to differ quite a bit from what they would have been would we have had the richer data the BLS has at its disposal. The best we can do in this respect is to try and be clear where we expect these differences to matter.

5.1 The Data.

We acquired quarterly data from IDC on quantity and mean price by; company, brand name, chip type (generation), and processor speed. This was our base sample and we matched more detailed characteristics and price data to it. The more detailed data was primarily from the annual January issue of *Datasources*, but we checked and augmented this with data from other issues of *Datasources* as well as other PC magazines (PCworld, PCmagazine, computer shopper, PC computing, ...)[15].

[14] The total cost of the data we use here is about ten thousand dollars, a tiny fraction of the cost of the BLS's PC component of the CPI over this five year period.

[15] The additional sources were especially important for 1999.

The detailed characteristic and price data is for the first quarter of every year. Consequently we produce an annual index based on that data and the quantity information in the first quarter of the IDC data. An "observation" in this data consists of the detailed characteristic and price data we gathered from the sources noted above and a quantity weight obtained by splitting the IDC category quantity equally among the observations that fit into the IDC category. The number of observations per IDC category varied from 2 to 10, and averaged 3.4.

Table 1 provides information on the number of products per year, the distribution of the characteristics of those products, and the fraction of products marketed in period $t-1$ that we were able to match to observations in year t. There are between 150 and 250 observations a year. This is generally larger than the sample size used by BLS to construct the PC component of the CPI (about 150 items), but much smaller than the sample used by the BLS to estimate the hedonic functions used for the hybrid index they have produced since 1998 (this sample usually contains between 800 and 1200 observations). Consequently we would expect *less* sampling variance in our estimates than in the BLS's estimates, but *more* variance due to prediction error in our hedonic estimates than in the hedonic estimates one could produce with their data.

A number of important facts can be read directly from this table. First all the characteristics of PC's listed in this table (and these will end up being the major, but not the only, characteristics that determine the hedonic function) are "vertical" in the sense that we expect every consumer to prefer more of them. Moreover the min, mean, and max of these characteristics all increase over time, illustrating the "quality" improvement that we all know occurs in this market. Note, however that the rate of change in the characteristics varied significantly from period to period. There was a striking increase in the quality of the machines marketed in 1998 (the year the pentium II became dominant in the market).

Note also that very few machines that are marketed in the base period are also marketed in the comparison period. Depending on the year, just under 10 to just over 20 percent of the observations were matched in adjacent years. As a result when we construct a matched model index we will have to drop 80 to 90 percent of the observations. This leaves ample room for the selection biases discussed above. This, however, is also one of those facts that would be likely to differ significantly were we using the BLS's data. The primary reason is that the BLS matches across months rather than across years, and more timely sampling would undoubtedly cause the match rate to go up[16]. It is important to remember, however, that the bias generated by the selection bias in the monthly data would have to be compounded over twelve months to give the bias in the annual index we are after.

5.2 Choice of Regressors.

Table 2 provides the adjusted R^2's from regressions of price against alternative right hand side variables. All regressions include five continuous variables; speed and speed squared ("spd" measured in megahertz), ram and maxram ("rm" and "maxrm" measured in megabytes), and hard-rive capacity ("hd" measured in gigabytes). They also include a set of (mostly dummy) variables for "add-ons" (see the notes to the table for a list of these variables), a dummy which differentiates apple machines, and that dummy interacted with speed. Some regressions also include processor type dummies in each period and additional interaction terms (see the notes to the table).

All regressions used price itself as the left hand side variable. We compared a specification which was linear in the right hand side variables and was estimated by OLS to one which used a power function of the right hand side variables and was estimated by non-linear least squares. The power function always produced at least as high an adjusted R^2 in all periods, and since, as noted above, we have no priors on either the sign or the stability of the coefficients, we will stick with the power function for most of the discussion below.

The difference in adjusted R^2 between our "base" nlls, which does not contain the type dummies nor the interaction terms and our fully loaded nlls (the last row), or indeed between either of these and any intermediate specification, is quite small. Indeed they are so small that we found it difficult to predict which of the regressions would generate a hedonic index with a smaller mean square error (see the discussion in section 5.1). As a result we will produce the indices that emanate from *both* the minimalist "base" specification and the maximilist "fully augmented" specification.

[16]We matched on a different set of characteristics than the BLS does and this could also effect the match rate. Generally the BLS has data on more characteristics than we do, so for any given time horizon they are likely to obtain a lower match rate than we do.

5.3 The Estimated Coefficients.

Table 3 presents the estimated coefficients and (heteroscedastic-consistent) estimates of their standard errors. It is clear that the coefficients jump around quite a lot between periods. For example, though we always need both a linear and a quadratic in speed in the regression, sometimes the linear part of the quadratic is positive and sometimes it is negative and the sign of the squared term is always opposite to the sign of the linear term. The coefficients on ram, max-ram, and hardrive capacity also vary quite a bit across time, especially in the fully augmented specification, and sometimes are the "wrong" sign (though they are never significantly so).

As noted above instability over time and reversals of signs are not unusual results for hedonic regressions, nor should they be a particular source of worry. Somewhat more troublesome is that the variance of at least some the parameter estimates seems to be substantial. There are at least two ways of trying to reduce that variance. One is to use more data. This is indeed what the BLS has done in estimating the hedonic regressions that go into their hybrid index, but neither we, nor most other empirical researchers, have the resources the BLS has at its disposal. Another method frequently used to reduce variance in the estimates of coefficients in hedonic studies is to constrain coefficients across years. As noted there is no particular reason to expect coefficients to be similar across periods, but if they are and we constrain them to be equal we may reduce the mean square error of the resulting index.

To see if constraining coefficients is a reasonable thing to do we computed some formal tests for the stability of the coefficients over time. Table 4 presents our results. For both the base and the augmented specification, we ran two sets of stability tests. One constrains only the "base" variable coefficients to be the same across years (those variables appearing in Table 3). The other constrains all variables (including the additional included variables listed at the bottom of Table 2) to have the same coefficients across years. To obtain the test statistics we proceeded as follows. First we used non linear least squares to estimate separate coefficients for each year and form a heteroscedastic consistent estimate of their covariance matrix. We then find the constrained coefficient vector that minimizes a quadratic form in the difference between the estimated and constrained coefficients weighted by the inverse of the estimated variance-covariance of the parameter estimates. We report as our χ^2 statistic the estimated value of that quadratic form (this is just Chamberlain's, 1983, "Π" matrix test for our problem).

Table 4 indicates that the data rejects just about any hypothesis on the constancy of the base coefficients. The rejections are even stronger when we also constrain the coefficients of the additional variables. Many of these latter variables are dummies for characteristics appearing for the first time during the sample period. The characteristics are initially offered only on a few premium machines, but then become more or less standard offerings. Consequently they tend to get larger coefficients in the earlier years.

The one marginal case for constraining coefficients occurs when we use the fully augmented specification and test for stability only of the base coefficients and then only over the first three years. If this were an independent test statistic we could accept it at a 1% but not at a 5% significance level.

It is instructive to consider what happens when we constrain only the base coefficients in the fourth year to be the same as those estimated for the first three years. All test statistics go up by factors of over a hundred, and three of the four of them have values over 10,000. The fourth year is the year of the mass introduction of the Pentium II, and that changed the mark-ups for all characteristics markedly. As noted in our theoretical discussion, this is just what we should expect to happen in periods of rapid technological change.

On the basis of these results, it is hard to make a case for constraining any of the coefficients over time.

6 Price Indices for Desktops.

We begin by comparing alternative ways of constructing a proper hedonic index. Then we compare these indices to the following alternatives; a matched model index, indices that mimic the hybrid procedure used by the BLS since 1998, the complete hybrid introduced in section 3, an alternative hedonic often used in research, and a comparison year (or Pasche like) hedonic index that would be easier for the BLS to produce in a timely fashion then the base period (or Laspeyres like) hedonic that we focus on.

6.1 "Proper" Hedonic Indices and Their Variances.

Table 5 presents alternative estimates of hedonic indices and their estimated standard errors. Where there are two estimates of the standard error presented, the first is an estimate of the overall standard error and the second estimates the contribution of sampling variance to that standard error. All estimated standard errors are obtained from the bootstrap procedures discussed in section 4.

The top panel of the table uses the "base" specification in Table 3 for our estimates of the hedonic regression function, while the bottom panel uses the fully augmented specification. Within each panel there are two sub panels. The top sub panel uses the estimated price relatives for all goods marketed in the base period, while the bottom drops the price relatives for goods that were marketed in the base period but had characteristics that were outside of the "range" of characteristics marketed in the comparison period (as defined by $\overline{C}$ in equation 5).

Within each of the sub panels there are three estimated hedonic indices for each period. The first two use the standard log-log linearization of the power function for the hedonic regression, and then estimate with OLS (recall that the power function is preferred by the data). The difference between these two rows is that the row labelled $\nu_t = .5\sigma^2$ adjusts the prediction equation from the log-log specification for the variance in the estimated disturbance, as equation(8) tells us to do. Our adjustment assumes that $Eexp[\epsilon^2] = .5\hat{\sigma}^2$ (which would be a consistent estimator of the adjustment if the disturbance distributed log normally). The row labelled $\nu_t = 0$ uses the same estimates of the hedonic function but makes no adjustment for the variance of the disturbance. The last row provides the hedonic price index based on the non linear least squares estimates of the hedonic equation and, as noted in section 5.1, then there is no need to do any adjustments.

Twelve indices are calculated. We being with differences in the method of constructing the indices which *do not* have a strong impact on the estimates of either the index per se or its standard error.

1. Since the bottom and top panels of the table are very similar, there is not much difference between the estimated hedonic indices generated by the model that uses the base, and the model that uses the fully augmented, specification for the hedonic regression.

2. Since the top half of each panel is similar to the bottom half, we conclude that it does not matter much whether we drop out the price relatives for goods in $t-1$ whose characteristics were outside of the range of characteristics of products marketed at t.

3. The estimates of the hedonic index which use the log-log linearization and adjust for the variance in the disturbance are very close to the nlls estimates, indicating that the assumption underlying the adjustments are reasonable.

(1) implies that the impact of differences in the characteristic set used in the hedonic analysis on the estimated hedonic price index is small; a fact which may not be that surprising when we remember that both regressions must have residuals which average to zero. Recall that the base and augmented hedonic specifications showed large differences in the estimated coefficients of the variables that appeared in both of them (Table 3). That is when we omit the extra characteristics in the augmented specification we *do* change the coefficients of our "major" variables; but that change is about what is needed to give the same prediction for the index.

Somewhat more surprising is the fact that the variance of the estimated index does not differ much between the base and the fully augmented specification. This has to do with the fact that the variables omitted in the base specification were only marginally significant (Table 2). Apparently provided we have data on the major characteristics that are available, we should not worry too much about the other characteristics used in the hedonic analysis, even if they have an impact on the coefficients of the major variables (recall that these coefficients are of no independent interest).

Similarly (2) and (3) indicate that we should not worry too much about goods which were marketed in the initial period but had characteristics which were outside of the range of the characteristics of goods marketed in the reference period, or about the method of estimation. We do, however, have to adjust the price prediction from the log-log regression for the variance in the disturbance. The estimates from the log-log regression which do not correct for the variance in the disturbance are consistently off the mark. They invariably overestimate the rate of the price fall by about .03, which, on average, is about twenty per cent of the actual index; a matter of some concern. We focus on the nlls estimates in the discussion which follows, but could have equally well focused on the estimates from the log-log linearization *provided* we make the adjustment for the variance.

Table 5 also contains at least four results on the characteristics of hedonic price indices for PC's during this period.

1. First the average price change over the period was a large negative number, sixteen to seventeen percent (and it does not matter much

which particular proper hedonic you use to estimate this number; our lowest estimate is 15.6% and our highest is 17.6%).

2. There is significant variance in the estimated index across years. 1998, the year that saw the mass introduction of the pentium II, saw a rate of price fall which is about double the rate of price fall in the earlier years, and 1999 is somewhere in the middle.

3. There is noticeable estimation variance in estimates; our estimated standard errors are on the order of five per cent for three of the four years, and about nine per cent for the fourth.

4. As noted the variance in the index consists of both sampling and estimation variance. The other striking fact is that much of the variance in the index is caused by sampling variance. In three of the four years the sampling variance was estimated to be over fifty per cent of the total variance in our index, and in one year it was estimated at eighty per cent.

The last point is striking in light of the comparison between the data at the disposal of the BLS and our data. Since our sample or price relatives is larger than the sample the BLS uses, but the sample we used for estimating the regression function is only one fifth of the size of the sample the BLS uses for this purpose, a BLS proper hedonic would have more sampling variance and quite a bit less estimation variance than the indices we produced. So we would expect most of the variance in the BLS's hedonic estimate to be due to sampling. Moreover the extent of sampling variance in the BLS's matched model index should be similar to the sampling variance in the proper hedonic it could produce. It would be surprising, then, if the BLS's matched model index had a much smaller variance than the proper hedonic index it could produce. The next section compares the bias of the matched model index to that of the hedonic index.

6.2 Comparisons to Alternative Indices.

Table 6 compares our proper hedonic indices (first panel) to alternative indices that have been used in either the statistical agencies or in research.

Matched Model Indices.

The second panel of the table provides two matched model indices that are frequently used. The first, a Tornquist index, is a weighted geometric mean of price relatives in the two periods with weights set equal to the average share of the good in sales over the two periods (see, for e.g.., Aizcorbe, Corrado and Doms, 2000, and the literature cited their for a discussion), and the second, a Laspeyres index, is a simple weighted average using base period sales weights. In a world in which all goods are available in both periods, the Tornquist is often thought to provide a closer approximation to an exact price index, while the Laspeyre's would satisfy our conditions for being a proper index.

The row labelled "% matched" indicates that about eighty five percent of the base-period observations cannot be matched in the comparison period, and hence are dropped when we construct the matched model index. This left ample room for the selection biases discussed in section 4.2. Even so, however, it is surprising just how large a role selection seems to play in this data. In particular

1. Both matched model indices show a price change about equal to zero in all years but 1997/98, and in 1997/98 the Tornquist index is *positive* and the Laspeyres is negative (both with absolute values being under a third of that of the proper index for that year).

2. The Tornquist index is *negatively* correlated with the proper hedonic indices across years, while the Laspeyres is positively correlated.

3. Apparently the way one constructs the matched model index matters. The correlation between the Tornquist and Laspeyres indices is minus one.

(1) indicates that the positive selection bias effects in the matched model indices just about exactly compensates for the negative effects of the fall in prices on those indices. Note that this implies that selection has its biggest effects in the years with the largest rate of technical change, as those are the years in which all but the most superior of base year goods are obsoleted by entering products. One can see the reflections of this clearly in the raw data. The hedonic index registers its largest fall in 1998, the year the pentium II obsoleted most older machines. Precisely because of this obsolescence, 1998 was the year with the lowest match rate. The 1997 products which did

survive were the few PC's which had already introduced the pentium II, and their prices did not fall nearly as much as the others.

More generally (2) implies that in the Tornquist index the increase in the positive selection effect that occurred with higher rates of technical change tended to outweigh the negative effects of the technical change itself, producing a negative correlation between the Tornquist and the proper hedonic. (3) is also of interest. Among the goods that survive until the reference period, those who grew disproportionately had prices which actually increased over time, giving the Tornquist index a larger positive bias (so a change which might have improved the index when there was no selection exacerbates the selection bias and ends up producing more harm then good).

There are two more points to make in this subsection. We did not provide estimated variances for the matched model index. This is because the number of matches was so small that we thought those variance estimates were unreliable. We note, however, that the estimated variances for two of the years were larger than the estimated variances of the proper hedonic indices, and for two of the years they were smaller. So there is no indication that the variance in the matched model index is smaller than that of the proper hedonic. I.e. *neither* bias nor variance considerations seem to favor matched model indices.

Finally though it is clear that the BLS's monthly sampling process results in a much higher match rate than that from our annual data, this should not have much of an affect on the *annual* selection bias implicit in the monthly indices. That is were we to compute an annual index from the monthly matched model index we would be compounding the monthly biases in those indices. Moreover the BLS's monthly match rate is about eighty-five per cent. This translates into an annual match of about fourteen per cent, which is *almost identical* to the match rates in our sample, indicating a that selection should have similar effects in the two types of data.

The BLS's Hybrid.

The last section should have made it clearly why PC's were the earliest of the commodity groups in which the industry analysts at the BLS recently decided to move away from use of a standard matched model index. Starting in 1998 the BLS used their own hybrid hedonic procedure to compute the PC component of the CPI. I begin by describing how the "BLS hybrid" is constructed.

The BLS's hybrid partitions the the goods sampled in the base period into three groups

- goods whose characteristics can be matched exactly to the characteristics of a good sampled in the reference period (full matches),
- goods for which they are able to match firm, brand and all but a small number of listed characteristics to a period t product (partial matches), and
- goods that cannot be matched at all.

As noted their monthly samples fully match about eighty-five per cent of products. About twelve per cent of products are partial matches, and three per cent are not matched at all.

The goods that are not matched at all are simply dropped from the sample (just as they are in the standard matched model index). A variant of hedonics is used to obtain price relatives for the partially matched goods, and those price relatives are averaged with the price relatives for the full matches to obtain the index.

The BLS's hedonic adjustment for partially matched goods begins by using data downloaded from the Internet to estimate a "constrained" hedonic regression. The comparison period price for a partially matched good is obtained as the price of the good partially matched to the base period good we are interested in, plus an adjustment for the difference in characteristics between that base period good and the reference period good that it was matched to. The adjustment equals the sum, over characteristics, of the difference in the quantity of the characteristic between the two goods multiplied by the estimated regression coefficient for that characteristic (it is the hedonic functions "evaluation" of the difference in characteristics).

I now come back to the difference between their constrained hedonic regression and the regressions we have been considering. As noted earlier, except in rare instances the coefficients obtained from hedonic regressions *are not* equal to the consumers' marginal valuations of characteristics. As a result they can easily provide a "wrong" sign for an increase in a "vertical" characteristic over some range, and even where the adjustment is the right sign its magnitude might not make any "sense" as an adjustment coefficient.

The BLS analyst are too good at their jobs to allow senseless adjustments to take place so they have adapted their adjustment procedures in ways which seem to "work" empirically. For their hedonic they regress price against a linear (instead of a power) function of all regressors; i.e. they do not include either squared or interaction terms in the characteristics as right hand side variables. The squared and interaction terms tend to generate negative effects of vertical characteristics in certain portions of the characteristic space. Also they set any of the coefficients which turn out to be negative in this regression to zero for the adjustments.

Even then not all of the coefficients they obtain from this regression are used for the adjustments. The major coefficient from these regressions which is *not* used to make adjustments is the coefficient of speed. Speed simply varies too much for the linear term to produce adjustments which seem reasonable to the analysts (speed also picked up significant squared terms in all of our regressions; see Table 3). So instead of using the hedonic coefficient to adjust for speed differences they substitute a separate adjustment coefficient that is obtained from an entirely different regression whose left hand side variable is the cost of chips to PC producers and whose right hand variable is a measures of chip performance (see Holdway,). Note that from the point of view of consumers, their speed adjustment is a cost of production based adjustment, and has nothing to do with consumer valuations.

There are three problems with the BLS's procedure.

1. It still incurs the sample selection bias resulting from the dropping out of the goods that are not matched at all. Their three percent drop out rate per month translates into an annual drop out rate of just over thirty percent, significantly lower than the eighty five percent dropped out when no adjustments are made, but still large enough to be important.

2. We noted in section 3.3 that even a complete hybrid that used the actual regression function ($h(\cdot)$ where $p = h(x) + \epsilon$ and $E[\epsilon|x] = 0$) to predict comparison period prices for goods that exited would be subject to the bias resulting from only using the ϵ for price comparisons of goods which do not exit. The complex procedure the BLS uses to obtain the price relatives of the partially matched data creates additional unknown biases. As a result the BLS's hybrid does not satisfy the conditions needed for the index to be an upper bound for the compensating variation.

3. Last, but certainly not least, the estimates the BLS's hybrid makes for

the comparison period price of the partially matched goods is likely to have larger variance than that from standard hedonic predictions (see below).

(3) is likely to be particularly important in high tech markets, like the PC market, where vertical characteristics are likely to be highly correlated both across products and over time (e.g. we talk of "high end" or top of the line, PC's, which have more speed, memory, storage capacity,...). In this case, the regression function will have difficulty sorting out which of the characteristics are associated with price changes producing individual regression coefficients with large variances even when the regression function itself is well estimated. Put differently the regression coefficients of characteristics will tend to be negatively correlated so that the (weighted) sums of coefficients used for the comparison period price in the proper hedonic will be estimated more precisely than the individual coefficients used for the BLS's adjustments.

The third panel of Table 6 provides estimates of BLS-type hybrids. To obtain these estimates the data in each base year was partitioned into a group of products that were perfectly matched to comparison period products, a group that was partially matched, and one that was not matched at all according to the criteria listed in this footnote[17]. The fraction of our sample that was perfectly matched was about what would have been predicted from the BLS monthly sampling experiments, but we got a slightly lower fraction of partial matches then did the BLS (forty vs. fifty five per cent).

The first hybrids we tried were BLS-type hybrids that used the coefficients from the hedonic regressions provided in Table 3 for the adjustments to the comparison period prices of the partially matched goods. However it quickly became clear that when we use the coefficients from any of our regression functions for the BLS adjustments the price indices we obtained made no sense – the numbers varied from plus to minus a thousand percent or more depending on the year and the exact coefficients. Given their adjustment procedure, it is easy to see how the BLS analysts moved away from estimating an unconstrained hedonic regression.

[17]A base period product was labelled as partially matched to a comparison period product if there existed a comparison period product with the same firm, brand, and processor type as the base product of interest. Among period t products with the same firm, brand, and processor type as the $t-1$ product of interest we chose the product that was "closest" to the base period product as our match; where closeness was defined in terms of the minimum of a weighted sum of the absolute values of differences in characteristics, with weights equal to the characteristic coefficients estimated in our hedonic regression.

When we simplified and used a linear (instead of a power) function *and* stuck to our base specification for included characteristics (recall that then the only second order term included was speed squared) we got numbers that were a bit more reasonable (these are the "lin" numbers in the third panel in table 6). The exact numbers differed quite a bit depending on whether we used the base $(t-1)$ or comparison (t) period hedonic regression coefficients for the adjustments, so we present them both in table 6[18]. Still, given what we know about this market, these numbers do not make a lot of sense. 1997/98, which recall is the year with the most selection and likely the year with the highest rate of improvement in characteristics, shows a price *rise*, and 1996/97 which was a year which registered only a modest price fall in the proper hedonic index, now registers a forty seven per cent fall in prices. It seems that without the separate speed coefficient, even the linear function does not do well.

As a result we tried to also mimic their speed adjustment. We were limited here by the fact that we did not have the data they had to estimate their speed adjustments, and since their estimated coefficients were considered confidential, we could not get access to them. We did however have data on the price at which Intel sold their later generation chips to PC producers (from the original pentium until now), and we used that data to estimate a "cost-based" speed adjustment for Intel chips in every year. Using the speed coefficient from this regression, and partially adjusting only machines based on Intel chips that were of a pentium or later vintage, gives us the indices in the next two rows of Table 6 (these are the rows labelled "adjusted"; note that the percentage matched in 1995/96 here is particularly low, as 1995 still had many machines using 486 chips). This adjustment brings us closer to our proper hedonic, but still leaves an index which seems problematic. The adjusted BLS hybrid hedonic are; (i) negatively correlated with the proper hedonic, (ii) on average quite a bit larger than the proper hedonic, and perhaps most important (iii)they have *two and a half* times the standard errors of our proper hedonic index.

There are also economically important differences between the actual BLS figures for the two years for which they use the hybrid, and our proper hedonic for those two years. The actual BLS hybrid for 97/98 reported a price change

[18] The BLS re-estimates the hedonic function used for their adjustments every three months or so. Thus the numbers which use the comparison period coefficients might be closer to mimicking what the BLS does.

of -.358, and for 98/99 it reported a price change of -.265. Though the actual BLS figures are more similar to our pure hedonic than our attempt to mimic their behavior, the BLS still estimates quite a bit larger price declines than we are estimating[19]. That is they estimate a rate of price fall which, in both years is, in absolute value, about a third larger than the price fall we are estimating (this is over a standard deviation larger in both years. These differences are large enough to be of some economic importance.

Complete Hybrids.

We also computed the complete hybrid introduced in section 3 (panel 4 of table 6). This is a (base period) sales weighted average of price relatives for *all* of the goods marketed in $t-1$. For goods which continue to be marketed in t it uses the ratio of the actual period t to period $t-1$ price for its price relatives. For all other goods it uses the period t hedonic regression to estimate the comparison period price for the good. Thus the only difference between our complete hybrid index and the proper hedonic indices is that now we use the actual prices relative when they are available.

The complete hybrid can be expressed as a weighted average of the Laspeyres matched model index given above and a Laspeyres hedonic index for the goods that are not. The weight given to the matched model index is equal to the share of the continuing goods in base period sales. This is reported as "w(full match)" in the table. The difference between the Laspeyres matched model for full matches and the Laspeyres hedonic for the other goods is reported as "difference" in the table.

The actual value of the indices are quite close to the values of the proper hedonics provided in panel one in all years. This is largely a result of the small weight given to the matched model index (less than ten per cent in all years and only three percent in 1998/99). The difference between the matched model index for the goods that survive and the hedonic index for the goods that do not are, per se, very large and almost equal to the index itself. For the same reason the variance in the complete hedonic is close to the variance in the proper hedonic; if anything the complete hedonic has a slightly larger variance.

[19]This could be a result of many problems, among them using linear projections for the impact of characteristics which seem to want to enter the hedonic function in a nonlinear fashion; see our table (2) above

Given the discussion of section three which indicated that the complete hedonic does not have a transparent interpretation, and the fact that the performance of the complete hedonic seems to be no better than that of the proper hedonic, we see no reason for using the complete hedonic, at least for the PC market.

The Research Dummy Variable Index.

This is the index which is most frequently used in economic research. It begins with a log-log hedonic regression that constrains coefficients on all variables but the constant to be the same across years. The estimated differences in the constants across years is then the hedonic estimate of the rate of price change.

There are two problems with this index.

1. This index is an "improper" hedonic; it does not necessarily provide an unbiased estimate of a lower bound to the compensating variation.

2. This index cannot be used by the BLS because it requires data which is not available when the BLS index must be constructed.

Recall that we may want to accept an improper index if it does not seem terribly different than a proper index and has lower variance. However in constructing the research dummy variable index the dummy for a given period is typically obtained from a regression that uses data from future as well as past periods. We could consider hedonic regressions which just use current and past data but there properties are likely to be different. In particular if there are trends in the values of characteristics over time (and in high tech products there surely will be), we would expect a dummy variable index constructed from current and past data to be both; less precise in the range of products currently being evaluated (this year's products have different, more advanced, characteristic then the products of prior years), and, if characteristic trends result in trends in the predictions of the hedonic functions for parts of the characteristic space (i.e. if new characteristic tuples obsolete old ones), we would expect this index to be biased towards past evaluations of those characteristic tuples.

Unfortunately we do not have a long enough time series to evaluate the dummy variable hedonic based only on current and past data, so we will have to stick to the research dummy hedonic which uses all the data we have[20]. The results in panel 5 of Table 6 indicate that the research dummy variable hedonic is closer to the proper hedonic then either the matched model index or the BLS-style hybrids.

The average absolute difference between the research dummy variable hedonic and the proper hedonic is about 30% of the average absolute value, and in one year the difference is statistically significant at traditional levels. As one might expect from a regression which used a single set of coefficients when coefficients are in fact changing over time (see Table 4), what the dummy-variable index does is smooth out the differences over years. It is lower than the proper hedonic when the proper hedonic is relatively high, and higher than the proper hedonic in years where the proper hedonic is relatively low. Recall that if we were to use only current and past data, as the BLS would have to do, the dummy variable index would have the additional problem of using data whose central tendencies are those of a year in the distant past, and not of a year which is in the middle of the sample, as is the case here.

The standard error of the research dummy variable index *is* lower than that of the proper hedonic[21]. Since the sampling variance from the two procedures should be similar and the variance of the regression disturbance from the research dummy variable regression is larger than those from the individual year regressions, the smaller variance of the dummy variable index is caused by the fact that the dummy variable regression obtains more precise parameter estimates. This difference in variances is likely to be much smaller

[20]For comparison with our other estimates we actually used nlls to estimate the traditional hedonic specification (then the price indices are obtained from year-specific multiplicative constants).

[21]This raises the issue of whether we might constrain some coefficients over some years to reduce the variance in the proper hedonic index even if it does result in an "improper" index. We tried several alternative ways of doing this, but none of them lead to a procedure which decreased our (bootstrapped) estimate of the mean square error of our index. One specification we tried constrained only the first three years of coefficients to be the same. We also tried a two-step procedure designed to produce an approximation to the minimum mean square error estimator. Here we used a linear specification and a two stage estimation procedure. The first stage estimates a "proper" hedonic and is used to get estimates of the variance and bias of every year's coefficients in estimating the true coefficients in every other year. These were then used to find the weighted average of the coefficients in the various years that minimized the (estimate of) the mean square prediction error in each given year.

in a BLS produced dummy variable index for two reasons; (i) as noted the BLS uses much larger samples to estimate their hedonic regression than we do and as a result the variance induced by imprecise coefficients is likely to be a smaller fraction of the overall variance in both indices, (ii) the BLS will have to use current and past data, and in high tech industries, data from the distant past will have less information in them on the hedonic surface in the current range of characteristic values then data from the near future.

"Timely" Pasche Indices

As noted above the BLS cannot produce an index which requires a regression function estimated on future data. The fact is the BLS would find it difficult to produce an index where the regression function is based on current data, as it would be difficult to estimate that function and put out the new index all in the same month. If current data cannot be used we cannot construct our "proper hedonic" index either.

There are at least two possible solutions to this problem. One is to delay publication of the index for PC's by a month. Though this is undesirable, it is done for other component indices, and so could presumably also be done for the PC index. The other is to substitute a Pasche, or reference period, type hedonic index for the current Laspeyres, or base period type hedonic, at least for the period in between the original publication of the index and the revision of the index.

The Pasche-type hedonic would evaluate reference period goods with the base period hedonic function; i.e. it would construct its average price relatives as

$$\frac{1}{n_t}\sum_{i=1}^{n_t}\frac{P_i^t}{h^{t-1}(x_i^t)}$$

where $i = 1, \ldots, n_t$ are the goods sampled in period t. Note that this Pasche-hedonic has a different, and less intuitive, justification then the Laspeyres index. It asks how much income would make a consumer who purchased the reference period bundle in the base period, at least as well off in the reference as in the base period.

The Pasche-hedonic computed from the base specification of our hedonic function (see Table 2) is provided in panel 6 of Table 6. Somewhat surprisingly they are quite close to our Laspeyres hedonic, closer (in absolute value) than *any* of the other indices. They do have higher variances than the Laspeyres hedonic in three of the four years (though no where near the variances of the BLS hybrids). This probably should be expected, since the "high end" reference period characteristic bundles are not prevalent in the base period, and hence the base period's regression function's estimates for those characteristics should have higher variance. Indeed this was the reason that we did not use the fully augmented specification to calculate the Pasche hedonic. If we did we would have had to use the base period regression function to evaluate say, types of chips, that were not present in the base period[22].

On the whole however the performance of the Pasche index is quite impressive. One possibility is to use it in lieu of the Laspeyre's type hedonic in initial publications, and then substitute the Laspeyres hedonic in the standard revision procedure.

7 Concluding Remarks.

This paper shows that there are relatively easy to construct hedonic indices that are "proper" indices in the sense that they provide an upper bound to the compensating variation needed to insure that comparison period utility is at least as large as utility in the base period. These hedonic indices circumvent a selection problem in matched model indices that, as our example shows, is likely to be particularly severe for commodity groups in which there is a lot of technical change.

There are two disadvantages to hedonic indices. First, in addition to the sampling variance that is inherent in matched model indices, they also have variance due to prediction error. Second the upper bound to the compensating variation they provide need not be particularly close to the least upper bound we are after.

The fact that the hedonic may be too generous is a natural result of looking for an index which can be justified "non parametrically"; i.e. of looking for an index whose justification is valid regardless of the functional form for and the distribution of utility functions[23]. Though we can argue about whether we should make more detailed assumptions about the distribution

[22]With related problems in mind we also calculated the Pasche index for goods in $C_{t-1} \bigcap \overline{C}_t$, i.e. omitting goods that were not in the span of characteristics of the goods marketed in $t-1$, but the numbers were very similar to those reported in the table.

[23]As a result the same problem arises in Konus,1924, justification for, and our use of, Laspeyres' indices in product space.

of utility functions, it is unlikely that we will want to use such assumptions in our official indices without quite a bit of prior experimentation with them.

It is hard to assess the importance of the variance in hedonic indices without an empirical example in mind. In the context of our PC example, it seems that much of the variance in the index is sampling variance. This variance is going to be common to all indexes which rely on the BLS's sampling process for the (implicit) weights used to average over alternative price changes. Second the BLS can and does use easily accessible external sources of information on prices and characteristics to decrease the variance in the estimated hedonic coefficients, and this should make the contribution of estimation variance to the total variance in the BLS's index even smaller than its contribution to our index.

Finally we noted that it would be hard for the BLS to construct our proper index in a timely enough manner for it to be used for the initial announcements of their monthly indices. This is because our proper hedonic requires an estimate of the current period hedonic regression function, and it would be hard for the BLS to use the incoming data to estimate the required function and then compute the index in time for the announcements. One can, however, construct a Pasche like hedonic that uses the base period hedonic function to compute the increase in prices of the reference period bundle of goods sold. Strictly speaking the Pasche hedonic is not a proper hedonic since a proper hedonic measures the price rise of the base (in contrast to the reference) period bundle of goods. However our estimates indicate that the two indices are not very different, suggesting it might be reasonable to use the Pasche hedonic index for the initial publications and substitute the proper hedonic when the BLS does its standard revision.

References.

Abraham, K., J. Greenless, and R. Moulton (1998): "Working to Improve the Consumer Price Index", *The Journal of Economic Perspectives*,pp. 27-36.

Aizcorbe, A., C. Corrado, and M. Doms (2000); "Constructing Price and Quantity Indexes for High Quality Goods", *mimeo*, Federal Reserve Board, Washington D.C. 20551.

Anderson, S., A. DePalma, and F.Thisse (1992): *Discrete Choice Theory of Product Differentiation*, the M.I.T. Press, Cambridge and London.

Bajari, P and L. Benkard (2001): "Demand Estiamtion with Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach" *mimeo*, Standord University.

Berry, S., J. Levinsohn, and A.Pakes (1995):"Automobile Prices in Market Equilibrium" *Econometrica*, Vol. 63, No. 4, pp.841-890.

Berry, S., J. Levinsohn, and A.Pakes (1998):"Differentiated Product Models from a Combination of Micro and Macro Data: The New Car Market", N.B.E.R. Discussion paper 6481.

Berndt, E (1991); "The Practice of Econometrics". Reading, Mass: Addison-Wesley.

Boskin,M., E.Dullberger, R.Gordon, A.Griliches, and D.Joregenson (1996):"Toward a More Accurate Measure of the Cost of Living", final report to the *Senate Finance Committee* from the Advisory Commissioin To Study the Consumer Price Index.

Chamberlain, G. (1983) ; "Panel Data", chapter 22 in *The Handbook of Econometrics*, Z. Griliches and M.Intriligator ed.s, North Halland, Amsterdam and New York.

Cockburn I. and A. Anis (1998); "Hedonic Analysis of Arthritis Drugs", N.B.E.R. Working paper, 6574, Cambridge, Mass.

Court, A.T. (1939): "Hedonic Price Indexes with Automotive Examples", in *The Dynamics of Automobile Demand*, pp.99-117, General Motors Corporation.

Epple, D. (1987): "Hedonic Prices and Implicit Markets: Estimating Demand and Supply Functions for Differentiated Products", the *Journal of Political Economy,* Vol.95, pp.59-80.

Feenstra, R. (1996): "Exact Hedonic Price Indexes", the *Review of Economics and Statistics*, pp.634-653.

Liegey, P. (2001): "Hedonic Quality Adjustments in the U.S. CPI: A Statistical Agency Perspective", *mimeo*, the Bureau of Labor Statistics.

Griliches, Z. (1961): "Hedonic Price Indexes for Automobiles: An Econometic Analysis of Quality Change", in *The Price Statistics of the Federal Government*, New York, the National Bureau of Economic Research.

Griliches, Z.(1971): *Price Indices and Quality Change*, Harvard University Press.

Griliches Z. and I. Cockburn (1994): "Generics and New Goods in Pharmaceutical Price Indices", *American Economic Review*, Vol.84, No.5, December, pp. 1213-1232.

Heckman, J. (1974):"Shadow Prices, Market Wages, and Labor Supply", *Econometrica*, 42, 679-83.

Holdway, M. (2001): "Quality-Adjusting Computer Prices in the Producer Price Index: An Overview", obtained from the BLS web site.

Hulten C. and F. Wykoff (1979): "The Estimation of Economic Depreciation Using Vintage Asset Prices" *The Journal of Econometrics* pp.367-96.

Konus, A. (1924); "The Problem of the True Cost of Living Index" translated in *Econometrica* 7, (1939), January, pp.10-29.

Moulton, B.(1996): "Bias in the Consumer Price Index: What Is the Evidence?" the *Journal of Economic Perspectives*, Fall, pp159-77.

Moulton, B. and K. Moses (1997): "Adressing the Quality Change Issue in the CPI", *Brookings Papers on Economic Activity*, pp. 305-366.

Moulton, B. (2001): "The Expanding Role of Hedonic Methods in the Official Statistics of the U.S.", paper presented at the Bureau of Ecocnomic Analysis Advisory Committee meeting.

Nevo, A.(2000): "New Products, Quality Changes and Consumer Welfare" *mimeo*, Univeristy of California, Berkeley.

Olley, S. and A. Pakes (1996): "The Dynamics of Productivity in the Telecommunications Equipment Industry", *Econometrica* pp. 1263-1297.

Pakes, A. (1995): "Statement". In "Consumer Price Index: Hearings Before the Committee on Finance, the United States Senate." Senate Hearing 104-69,, Washington, D.C.: U.S. Government Printing Office, 1995, pp.44-48.

Pakes, A., Berry, S., and J. Levinsohn (1993): "Applications of Recent Advances in Empirical I.O.: Price Indices and Quality Change", the *American Economic Review.* pp.240-46.

Pakes, A. (2000); "A Framework for Applied Dynamic Analysis in Industrial Organization", NBER Working Paper, 8024.

Pollak, R. (1989): *The Theory of the Cost-of-Living Index*, Oxford University Press.

Rosen, S. (1974): "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition" the *Journal of Political Economy*, Vol.82, pp.34-55.

Triplett, J. (1971); "Quality Bias in Price Indexes and New Methods of Quality Measurement", pp.180-214 in Z. Griliches, op.cit.

Table 1: Characteristics of Data*.

year	95	96	97	98	99
nobservations	264	237	199	252	154
matched to $t+1$	44	54	16	29	n.r.
characteristics					
speed (Mhz)					
min	25	25	33	140	180
mean	65	102	153	245	370
max	133	200	240	450	550
ram (MB)					
min	2	4	4	8	16
mean	7	12	18	42	73
max	32	64	64	128	128
hard disk (GB)					
min	.1	.1	.2	.9	2
mean	.5	1	1.8	4.5	8.5
max	1.6	4.3	4.3	16.8	25.5

Table 2: Adjusted R^2 Comparisons; .*

Model	95	96	97	98	99
Base lin-lin	.363	.503	.469	.431	.259
Base; nlls	.372	.503	.481	.480	.308
Same + rm*maxrm	.383	.509	.484	.479	.308
Same+ spd*rm	.381	.508	.492	.487	.311
Base + type	.373	.502	.491	.506	.308
Same + spd*ram	.372	.500	.500	.514	.316
Same +type*speed	.404	.498	.514	.520	.316
same+ram*mxrm	.407	.504	.513	.518	.311

* Notes to Table.

- "lin-lin" is a model which is linear in price and in characteristics. "nlls" is linear in price but but has the characteristics to estimated powers times an estimated constant as right hand side variables.

- All regressions include also dummies for; apple, CDRM, Modem, ethernet, DVD. In addition they all include graphic memory (MB), sound card (0,1,2 representing none, normal, 3-dimensional), and apple $\times$ spd.

- "Base" variables are; spd (MH), spd^2, ram(MB), maxrm, hd (GB).

- "Type" is processor type. Types (other than apple types) are: in 1995; pentium, 486: in 1996 and 1997; pentium, pentiumII, 486: in 1998 pentiumII, pentium: in 1999 just pentium II.

Table 3: Major Coefficients and Their Variance.*

Using "base" nlls from Table 1.

x	95	96	97	98	99	panel
const	1.12 (1.46)	12.77 (2.04)	16.84 (2.59)	-39.58 (8.87)	-5.91 (20.12)	n.r.
spd(MH)	2.82 (.70)	-3.24 (.90)	-4.72 (1.13)	16.79 (2.47)	4.34 (6.90)	1.33 (.26)
spd^2	-.325 (.09)	.37 (.10)	.49 (.12)	-1.56 (.22)	-.44 (.59)	-.15 (.03)
ram(MB)	.39 (.06)	.32 (.085)	.20 (.07)	.52 (.12)	.30 (.09)	.34(.06)
maxr	-.00 (.06)	.19 (.05)	.20 (.05)	.09 (.05)	-.04 (.045)	.13 (.03)
hd(GB)	-.05 (.08)	.07 (.07)	.10 (.083)	.02 (.09)	.28 (.07)	.03 (.04)
Using fully augmented specification from Table 1.						
const	6.80 (2.32)	12.73 (2.45)	15.25 (.25)	-12.13 (10.35)	5.37 (28.61)	n.r.
spd(MH)	-.83 (1.19)	-2.70 (1.31)	-2.65 (1.94)	5.12 (4.35)	-2.16 (6.90)	1.17(.44)
spd^2	.15 (.165)	.34 (.16)	.18 (.25)	-.36 (.47)	.29 (1.08)	-.12(.06)
ram(MB)	1.065 (.65)	.18 (.72)	-1.73 (1.59)	2.47 (1.31)	3.58 (2.83)	.28(.21)
maxr	.24 (.16)	-.19 (.24)	.03 (.35)	.10 (.34)	.21 (.65)	.03(.10)
hd(GB)	-.01 (.09)	.06 (.08)	.12 (.07)	.06 (.06)	.31 (.07)	.08(.04)
Obs.	264	237	199	252	154	

* Heteroscedastic consistent estimates of standard errors appear in brackets after coefficient estimates. For other variables included in these regressions see table 1.

Table 4: χ^2 Tests For Constancy of Coefficients.

test	Base Spec.*1——				Fully Aug. Spec.*1 ——			
restricted variables	basic*2		all*2		basic*2		all*2 ——	
	d.f.	χ^2	d.f.	χ^2	d.f.	χ^2	d.f.	χ^2
All Years	20	152	57	450	20	61	73	3,309
First three years	10	50	28	165	10	22*3	38	237
First four vs. first three	5	10,391	14	12,480	5	569	18	60,999
Fourth and fifth year	5	71.7	14	135	5	18.8*4	17	868

*1See tables 1 to 3 for variables included in these two specifications. See the text for how the test statistics are constructed.

*3 accept at 1% but not at 5%.

*4 accept at .25% but not at .1%

Table 5: "Proper" Hedonic PC Price Indices and Their Variance*

		v_t	95/96	96/97	97/98	98/99	av.
Specification With Base Variables Only.							
$x \in C_{t-1}$	log-log						
		$.5\sigma^2$	-.102	-.111	-.292	-.172	-.169
			(.037/.023)	(.059/.037)	(.041/.023)	(.092/.025)	(.09/n.c.)
		0	-.132	-.140	-.318	-.199	-.197
			(.035/.022)	(.056/.035)	(.038/.022)	(.091/.025)	(.09/n.c.)
	NLLS		-.097	-.108	-.295	-.155	-.164
			(.04/.024)	(.063/.035)	(.045/.026)	(.099/.023)	(.09/n.c.)
$x \in C_{t-1} \bigcap \overline{C}_t$	log-log						
		$.5\sigma^2$	-.100	-.115	-.267	-.161	-.161
			(.032)	(.054)	(.038)	(.062)	(.08)
		0	-.130	-.143	-.294	-.189	-.189
			(.031)	(.052)	(.036)	(.061)	(.07)
	NLLS		-.094	-.111	-.270	-.150	-.156
			(.039)	(.052)	(.044)	(.054)	(.08)
Specification With Type and Interactions.							
$x \in C_{t-1}$	log-log						
		$.5\sigma^2$	-.106	-.123	-.277	-.188	-.174
			(.035/.028)	(.066/.041)	(.041/.026)	(.128/.021)	(.08/n.c.)
		0	-.137	-.150	-.301	-.215	-.200
			(.034/.027)	(.064/.039)	(.039/.026)	(.125/.02)	(.08/n.c.)
	NLLS		-.099	-.137	-.292	-.176	-.176
			(.036/.024)	(.063/.035)	(.049/.026)	(.131/.023)	(.08/n.c.)
$x \in C_{t-1} \bigcap \overline{C}_t$	log-log						
		$.5\sigma^2$	-.105	-.126	-.269	-.173	-.168
			(.035)	(.057)	(.043)	(.062)	(.07)
		0	-.135	-.153	-.293	-.201	-.195
			(.033)	(.056)	(.041)	(.061)	(.07)
	NLLS		-.097	-.138	-.280	-.164	-.170
			(.037)	(.057)	(.044)	(.083)	(.08)

*Standard errors appear in brackets below estimate. Where there are two standard error estimates, the first correspondes to the estimate of the actual variance, and the second corresponds to the estimate of just the sampling variance component (see section 4). All standard errors are estimated by a bootstrap based on 100 repetitions (higher numbers of repetitions did not change the index).

Table 6: Alternative PC Price Indices*

	Year	95/96	96/97	97/98	98/99	av.
1. Proper	base	-.097	-.108	-.295	-.155	-.164
		(.040)	(.063)	(.045)	(.099)	(.091)
Hedonics	f.a.	-.094	-.111	-.270	-.150	-.156
		(.039)	(.052)	(.044)	(.054)	(.079)
2. Matched	Tornquist	.012	.002	.09	.011	.028
model	Laspeyres	-.013	-.002	-.08	-.011	-.027
	% matched	16.6	22.8	8.0	11.5	14.7
3. BLS	$\text{lin}(h^t)$	-.469	-.468	.014	-.057	-.245
		(.39)	(.16)	(.19)	(.13)	
Hybrids	$\text{lin}(h^{t-1})$	-.167	-.661	.007	-.059	-.220
		(.71)	(.40)	(.15)	(.81)	
	% matched	43.6	71.7	50.3	45.2	52.7
	$\text{adj}(h^t)$	-.245	-.366	-.203	-.107	-.230
		(.18)	(.15)	(.13)	(.11)	
	$\text{adj}(h^{t-1})$	-.345	-.623	-.208	-.082	-.315
		(.28)	(.26)	(.14)	(.24)	
	% matched	25	64.1	48.7	35.7	43.4
4. Complete	base	-.125	-.135	-.300	-.168	-.182
		(.043)	(.062)	(.046)	(.115)	
Hybrids	f.a.	-.128	-.164	-.292	-.194	-.195
		(.042)	(.063)	(.044)	(.141)	
	w(full match)	.099	.061	.043	.03	
	difference	.124	.143	.230	.162	
5. Research	base	-.135	-.098	-.160	-.170	-.141
Dummy		(.038)	(.035)	(.027)	(.040)	(.032)
Variables	f.a.	-.152	-.122	-.213	-.143	-.158
		(.040)	(.032)	(.041)	(.028)	(.039)
Research	base	.038	-.01	-.135	-.009	abs=.049**
minus		(.041)	(.061)	(.042)	(.094)	
Proper	f.a.	.053	-.015	-.079	-.033	abs=.045**
		(.043)	(.067)	(.055)	(.125)	
7. Pasche	base	-.104	-.167	-.337	-.119	-.181
Hedonic		(.07)	(.049)	(.091)	(.062)	(.070)

*Standard errors appear in brackets after estimate. They are estimated by a bootstrap based on 100 repetitions. "base" refers to base specification and "f.a." refers to fully augmented specification in table 2.