A Reconsideration of Hedonic Price Indexes with an Application to PC's.

Ariel Pakes, Harvard University and the N.B.E.R. [1]

Revised, August 2003.

[1]My interest in hedonics, and in price indexes 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 the comments of three referees and from discussions with Steve Berry, Ernie Berndt, Gary Chamberlain, Iain Cockburn, Erwin Diewert, Shane Greenstein, Alan Krueger, Minjae Song, and Jack Triplett. Extensive interactions with members of the Bureau of Labor Statistic's staff, specially with Tim Erickson, have been extremely helpful in clarifying concepts and understanding the environment they work in. 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.

Abstract

This paper compares hedonic to matched model indexes. Matched model indexes are averages of the price changes of goods that remain on sampled stores' shelves. Since goods that disappear tend to have falling market values, matched model indexes select from the right tail of price changes. The BLS can construct hedonic indexes that correct for this selection and are justified by standard arguments. In an empirical study of PC's hedonics produce sharp price declines while matched model indexes are near zero. Also, though there are modifications to hedonics that seem desirable, they are not those in current use.

This paper considers the use of hedonics to ameliorate new goods biases in price indexes. It starts with a brief review of the new goods problem and then considers the theory underlying hedonic indexes and the data needed to construct them. We conclude with an empirical analysis of hedonic and matched model price indexes for personal computers (PC's).

The hedonic pricing function is derived from an Industrial Organization model of differentiated product markets. This clarifies the properties of that function, explains "anomalies" in previous hedonic estimates, and provides a framework which can be used to analyze both hedonic and matched model indexes. In particular in gathering data for its indexes the Bureau of Labor Statistics (the BLS) only obtains price changes for those goods in the base period sample that are still sold at the sampled outlet in the comparison period. Goods which disappear from the outlet's shelves in the comparison period tend to be goods which were obsoleted by new products; i.e. goods which have characteristics whose market values have fallen. The matched model index is constructed by averaging the price changes of the goods that *do not* disappear; so it selects disproportionately from the right tail of the distribution of price changes. In the PC market the annual attrition rate averages 85 percent, so the positive selection bias in the matched model index can be large. In contrast the hedonic uses characteristic data to form an estimate of the comparison period prices of *all* goods in the base period samples, and averages the price changes implicit in those estimates.

The paper provides conditions which insure that hedonic indexes provide an upper bound to the compensating variation, and considers detection of situations when it might not. We also explain why the current data generation process enables the BLS to construct hedonic indexes in a timely manner. The hedonic index does have estimation (as well as sampling) variance and hence can be less precise than the matched model index. However when it is there is a hybrid hedonic index (our "complete hybrid") which has the variance reduction properties of the matched model index and the selection correction properties of the hedonic. Thus there are good *a priori* reasons for the BLS's recent (and rather controversial) push to-wards using hedonics (see Paul Liegey, (2001))[1].

Our study of PC's compares hedonic to matched model indexes and considers modified hedonic indexes including the "incomplete hybrid" currently used by the BLS. The hedonic index for PC's shows a rather *sharp decline* in prices over the 1995-99 period ($15-19$ percent per annum). Differences in the rate of decline across years reflect differences in the characteristics of the products marketed and do not change with the details of how the hedonic is constructed. In striking contrast the matched model indexes' price changes are *close to zero* for the whole period.

The BLS's move to hedonics has been to a modified hedonic motivated

[1]The PC component of the CPI began using a hedonic like adjustment in 1998, while the television component began in 1999 (interestingly the sample attrition rate in the BLS's TV sample is similar to that in PC's; see Brent Moulton, Timothy Lafleur, and Karin Moses (2000)). 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.

by variance reduction and timeliness considerations. There are theoretical reasons to worry about the BLS's hedonic and they prove telling in the PC example. We show that other modifications are more likely to help.

1 The New Goods Problem.

The price quotes used to construct most components of the consumer price index (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, these "price relatives" would be averaged to obtain the group's "matched model" index. 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 is instructed to make a "forced substitution" of another good.

The sample rotation and forced substitution processes enable new goods to enter and old goods to exit 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 detail see Jack Triplett (2003)).

The important point, and the source of the new good bias in price indexes, 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 change for this commodity group (say $\dot{p}(t)/p(t)$) will be a weighted average of the price relatives of the new and old good ($\dot{p}^n(t)/p^n(t)$) and $\dot{p}^o(t)/p^o(t)$) with the weight ($w(t)$) being the fraction of price quotes that are obtained from the new good, that is

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

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. When $w(t)$ is increasing there are 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 and have (by revealed preference) increased their utility as a result of the switch. This utility gain is not captured by matched model indexes so it is a "new goods" bias. Depending on introductory pricing patterns it can cause the introduction of a new good to increase the computed price index at the same time as it decreases the true cost of living (for examples see Ariel Pakes (2002)).

More generally the new goods bias in the matched model index is an implication of how goods enter and exit the data gatherer's sample. We turn now to an explanation of how hedonics can be used to ameliorate this bias.

2 Economic Analysis and "Hedonic" Indexes.

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 Andrew Court (1939) and revived by Zvi Griliches (1961) as a way of ameliorating the new goods problems discussed above. These authors argued that since newer models often had more desirable characteristics, the difference between the prices of the newer and the older models should not be entirely attributed to inflation. Court and Griliches suggest estimating a surface which relates prices to characteristics and time and then using the estimated surface to obtain estimates of price changes for products with *constant* characteristics.

A series of papers followed. Those that were theoretical focused on the relationship between characteristics and prices generated by the equilibria of differentiated product markets (Sherwin Rosen (1974), Dennis Epple (1987), Simon Anderson, Andre De Palma, and Jaques-Francois Thisse (1989), Robert Feenstra (1995), and Steven Berry, James Levinsohn and Ariel Pakes (1995), henceforth BLP)[2]. These models 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, residence,....). As a result aggregate demand will depend on both the characteristics of the goods marketed and the distribution of consumer attributes.

- Cost functions which typically include characteristics of the good, input prices, the scale of production, and "productivity" as its arguments.

- An equilibrium assumption. This determines prices (and quantities) given demand and costs.

For our purposes the 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 do this 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].

[2] A notable exception is the work of Robert Pollak, which is summarized in his book (Pollak (1989)), and which is similar in spirit to the analysis in section 3 below.

[3] See Berry, Levinsohn and Pakes (forthcoming). Pakes, Berry and Levinsohn (1993)

We show below that hedonic techniques 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 Hedonic Regression.

Our bound uses a hedonic regression so we begin with that regression's properties. The structure of the relationship between prices and characteristics depends on the appropriate equilibrium assumption. For specificity we take that to be Nash in prices (or Bertrand). Let (x_i, p_i) denote the characteristics and the price of good i and (x_{-i}, p_{-i}) denote the characteristics and prices of the other goods marketed. Then the demand for good i becomes

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

where A indexes the distribution of consumer attributes that determine consumers' preferences over characteristics.

If all firms are single product firms and and marginal costs are given by $mc(\cdot)$, then prices are

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

where the second term, $\frac{D_i(\cdot)}{|\partial D_i(\cdot)/\partial p|}$, is the mark-up which varies inversely with the elasticity of demand at the point. The hedonic function, say $h(x)$, is the

and Aviv Nevo (2003) provide examples of price indexes build in this way. Though it is not currently practical to use these techniques to compute official price indexes, they could be used to form "exploratory" indexes which help evaluate current procedures.

expectation of price conditional on x. From (2)

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

where the expectation integrates over randomness in the processes generating the characteristics of competing products, input prices, and productivity.

The hedonic function is the expectation of marginal costs plus that of the mark-up conditional on "own-product" characteristics. Further since marginal cost depends on the price of purchased inputs, its conditional expectation depends on the expected markup in input prices. From (3) markups are a complex function of the characteristics of the competing products, the distribution of consumer preferences, and had we allowed for multi-product firms, of the structure of ownership. So if markups in the product of interest (or in its inputs) are significant we should expect the hedonic function to change when the number, the characteristics, or the ownership pattern of competing products (or inputs) changes. For example semiconductor chips are a significant component of the costs of PC's, so the fact that the price of chips varies dramatically with the characteristics of the competing chips that are available (see Minjae Song (2003)) causes the hedonic function in the PC market to change when new chips are introduced (see below).

Given our interest in markets with a flow of new goods I would argue that we should proceed under the assumption that there are large markups (either in the goods studied or in their purchased inputs). This because new goods require prior investments in development and marketing and the markups are necessary to justify those investments[4]. Moreover once we allow for the markups it is clear that the hedonic regression is a “reduced form”; i.e. its coefficients have no obvious interpretation in terms of economic primitives.

Below we provide conditions which insure that this reduced form is all we need to form a price index which bounds the true compensating variation. That is the reduced form coefficients are the objects of interest, so provided we use sufficiently rich regression techniques we need not worry about “coefficient biases”. On the other hand because it is a reduced form there are no *a priori* restrictions on hedonic regressions. Thus the hedonic functions for the same products may differ in markets with different distributions of consumer attributes. More importantly we should expect that function to change with product introductions; new products are typically directed at high markup parts of the characteristic space and cause those markups to fall (e.g. the introduction of PC’s with a new fast chip causes a decrease in

[4]Provided the new goods are not subject to sharply increasing costs (a situation thought unlikely in the early years of a product’s life), this reasoning is independent of the appropriate equilibrium cost or demand assumptions. However with more complex equilibria the argument is more detailed. For example 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 assumption in (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 the current marginal costs and mark-up. Then the mark-ups which justify product development will be garnered at a later stage in the product’s life cycle. For more detail on the theory underlying the markups and entry/exit decisions that underlie our discussion, see the cites in Ariel Pakes (2000).

the mark-up on the machines that had been the speed leaders). Also there is no formal reason to expect an increasing reduced form relationship between the markup and characteristics that we generally think of as desirable[5].

3 Hedonic Price Indexes.

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

[5] I should note that if we had products that had identical values for all but one characteristic, and that characteristic was vertical in the sense that all consumers preferred more of it, then among these products price there should be monotone increasing in the value of the vertical characteristic. This is because when we have products that are identical except in a single characteristic utility theory suffices to determine the sign of the difference in their prices (for a proof see Patrick Bajari and Lanier Benkard (2003)). When there are real costs to developing new products we do not expect to observe a product space that is dense in all characteristics, as is required for this argument.

still have the same amount of income left over to buy other goods. It 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. 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 required for this 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 Alexander Konus (1924) and remains the predominant justification for the indexes currently in use. Not surprisingly then, the problems that arise when using the two indexes are the same; what the use of characteristic space does is enable different solutions[6].

[6]Below we focus on the problems that arise when $x_1 \notin C_2$ but this statement is more

3.1 Hedonic vs. Matched Model: Conceptual Issues.

We begin by assuming that hedonic and matched model indexes can be computed from the same data, and that measured characteristics fully determine the utility a consumer derives from the good. These assumptions are discussed in the next subsection.

Each of the base period goods are either available in the comparison period or not. Consider first the estimated price changes when $x_1 \in C_2$. Then the hedonic bound argument given above applies, but so does Konüs' (1924) 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 between the hedonic and the matched model's price changes is their estimates of $h^t(x)$, a topic I return to below.

The major conceptual difference in the two indexes 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 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 characteristics whose values have fallen

generally true. In particular problems that might arise because utility functions are not separable across commodity groups are common across the two indexes.

more than the average (often due to the entrance of products with superior performance). As a result

the matched model index construction procedure tends to throw out the left tail of the distribution of price changes and produce a price index which is biased up-wards.

The word "tends" here deserves emphasis; given the current data generating process an $x_1 \in C_1$ will not appear in C_2 if it does not appear on the shelves of the outlet at which the good was sampled in the first period. Other reasons this could happen include stock outs at the sampled store and changes in cost of production or ownership structure[7].

In 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).

[7]Note also that the bias occurs even if there are *no* unobservable characteristics. This is the difference between the selection problem here and the econometric analysis of selection; see James Heckman's (1974) classic study of labor supply and the (otherwise closely related) study by Charles Hulten and Frank Wykoff (1981) of selection in the analysis of depreciation rates. Both these studies were concerned with bias in coefficient estimates and there is no coefficient bias if selection is based on observables.

For $h^2(x_1) - h^1(x_1)$ to be an upper bound to the compensating variation when $x_1 \notin C_2$ we require conditions which insure that the consumer will be better off with $h^2(x_1) - h^1(x_1)$ added to its income and C_2 then with its original income and C_1. We provide two sufficient conditions, and then a more general discussion of when the bound might be problematic.

Let utility be a function of x and income available for other goods (or $y - h(x)$). If there are close substitutes to x_1 in C_1 then the first order condition, $\partial U(\cdot)/\partial x \approx \partial U(\cdot)/\partial y \times \partial h^1(\cdot)/\partial x$ will be (approximately) satisfied at the point$(x_1, y - h^1(x_1))$. This enables us to evaluate utility differences in a region around x_1, say $N(x_1)$, with gradients which can be estimated up to a positive constant by the slope of the hedonic surface. The appendix shows that this implies that if

$$max_{\, x \in N(x_1) \cap C^2} \left[\frac{\partial h^1(x)}{\partial x}|_{x_1} - \frac{\partial h^2(x)}{\partial x}|_{x_1} \right] [x - x_1] \geq 0,$$

a condition which can be evaluated empirically, then $h^2(x_1) - h^1(x_1)$ is an upper bound to the compensating variation [i.e. there is an $x \in C_2 \cap N(x_1)$ such that $(x, y + h^2(x_1) - h^1(x_1) - h^2(x))$ is prefered to $(x_1, y - h^1(x_1))$].

Alternatively if the important characteristics of the product are vertical, so that all consumers prefer more of them (as is the case for our PC example), or if the only non-vertical difference between the product that disappears and a product which *is* in C_2 is the outlet of sale and we could bound compensation for those differences, then 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.

More generally for a good which exits to cause a problem for our bound it must have gone from profitable to unprofitable despite still being preferred by a significant share of consumers (at price $h^t(x)$). This is an unlikely event and when it does occur we should be able to correct for it unless; (i) the good has major "horizontal" characteristics (characteristics which some consuming units prefers more of while others prefer less)[8], *and* (ii) the good does not have close substitutes in one of the two periods. The BLS's analyst should check for such situations. A simple check for substitutes is to see whether the exiting good's characteristics are inside the range of characteristics of the goods that are available in C_2. I.e. if $x_i = [x_{i,j}]$ where j indexes the various attributes of the good, $\overline{x}_j = max_{i \in C_2} x_{i,j}$, and $\underline{x}_j = min_{i \in C_2} x_{i,j}$, check if

$$x_i \in \{x : \underline{x}_j \leq x_{i,j} \leq \overline{x}_j, \text{ for } j = 1, \ldots, J\} \equiv \overline{C} \in \mathcal{R}^J. \tag{5}$$

In this spirit our study of PC's will present two hedonic indexes. One assumes the modification needed for the upper bound when goods exit is

[8]For the good to pose a problem we would expect it to also have a large fixed costs of production since then a non-negligible fraction of the population could prefer the old good to the new goods without making the old good profitable to market. An example might be an old automobile (say 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. Note that neither the conditions for the bound, nor the conditions which are likely to make it suspect, have anything to do with how the hedonic function would have shifted had the good that exited not exited.

satisfied for all the goods that exited. The second assumes that the modification is only satisfied when $x_i \in \overline{C}$. For $x_i \notin \overline{C}$ we do exactly what the matched model index does; drop them before calculating the index.

3.2 Hedonic vs Matched Model: Practical Issues.

The practical questions associated with constructing the indexes are closely related to data generation issues. This paper assumes that the BLS continues to use their current data generation procedure. That procedure was designed for the matched model index, but as we now explain it can also be used for constructing hedonic indexes[9].

Matched model indexes 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). The data gathering process generates two lists of product characteristics. 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, and then, after sampling an item from this set, the data collector writes down the values for a second (often more detailed) set

[9]Alternative data generation procedures are now under consideration. The most dramatic (as well as probably the most cost effective) of the suggested changes would be to use explicit quantity or sales weights to form the weighted average of price ratios that becomes the index for a given commodity group (rather than relying on the weights implicit in the current sampling process). This and other changes are likely to make hedonics both easier to implement and more precise; see Pakes (2002).

of characteristics for the entry level item actually sampled. The second list allows a (possibly different) data collector to recall the "same" good, or at least a good with the same listed characteristics, when the outlet is revisited in the next period. If a good with the listed characteristics is found when the outlet is revisited, 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.

Note that this implies that matched model indexes are formed by comparing prices of characteristic tuples, not by comparing the prices of "goods" per se, just as hedonic indexes are. Moreover the BLS analysts have been choosing the characteristics whose prices are compared for years, and it has never been a contentious issue. The lists themselves are typically at least as detailed as those used in hedonic studies and include the name and type of the outlet where the product was sold[10]. A hedonic which conditions on all

[10]For PC's the product characteristic list includes: brand, model number, RAM, hard drive capacity, CPU, speed (megahertz), and dummies for; CD-ROM, CD-writer, zip-drive, monitor (brand, model, and size), keyboard, sound-card, speakers (description), extra software (office, ...), video card (premium or not), floppy drive, and modem (fax capabilities or not). Outlet types are coded (factory outlet, department store, specialty shop....) and the name of the outlet can be used for sales at chain stores. Outlet characteristics are important because the characteristics of the "sale" might matter to the consumer, and though the outlet of purchase does not fix all of those characteristics (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 array of other goods available and their prices), it might proxy

these characteristics will be conditioning on the same variables that matched model comparisons condition on.

There still may be product characteristics that consumers care about but are not listed, and this may be one source of the disturbance in the hedonic relationship. The impact of such unobservables, however, can be quantified. If the data collector looking for a comparison period match can only condition on the listed $x's$ then the disturbance in the hedonic function for the comparison period's price of a good will be independent of that good's base period disturbance. This implies that the expectation of the price relative for products which continue, i.e. for the products that are in both the hedonic and matched model samples, are the same no matter which index is used[11]. So the difference in expectations of the matched model and hedonic indexes equals the sales share of the goods that drop out times the difference between the averages of the price relatives for the goods that continue and the hedonically estimated price relatives for the goods that drop out.

3.2.1 Variance and Hybrid Indexes.

In evaluating estimators we typically consider their variances as well as their means. Both the matched model and the hybrid index have sampling variance

for them.

[11] More formally $E[\frac{h^{t+1}(x)+\epsilon_{t+1}}{h^t(x)+\epsilon_t}|x,\epsilon_t] = \frac{h^{t+1}(x)}{p_t(x)}$, where the equality follows from the fact that ϵ_{t+1} is mean independent of x because $h^{t+1}(\cdot)$ is the expectation of price conditional on x *and* the good being sold in period $t+1$, and is independent of ϵ_t because the data collector only conditions the comparison period draw on price on the observed x's in the base period.

(variance caused by the difference between the sample and the population average of the price relatives used), while the hedonic also has variance due to estimation error. With the sample sizes used at the BLS (see below), both the estimation and sampling variances are likely to be noticeable.

Assuming independent sampling the sampling variance in each index is the variance of a given draw on the index's price relative divided by sample size. The variance of a single price relative differs in the two indexes (since one samples from a selected subsample). Further though the hedonic has estimation error, it also has a larger sample size (as it does not drop goods). So the difference in the variance of the indexes cannot be signed *a priori*.

In cases where the variance of the hedonic is larger than the variance of the matched model index there is a hybrid hedonic index available which should dominate the matched model index. It is constructed by averaging the observed price relatives for the goods which are found in the comparison period with the hedonic estimate of the price relatives for the goods which are not.

I will call this hybrid index a "complete hybrid" (to distinguish it from hybrid actually used by the BLS, see below). It does not incur the additional estimation variance of the hedonic for the goods which are sold in the comparison period, and does not incur the selection bias of the matched model index for the goods that do not survive[12].

[12] As pointed out to me by Ellen Dulberger and Jack Triplett the index used in Dulberger (1989) is similar to our complete hybrid. One final point. If there are some characteristics which the data gatherer can condition on but the econometrician cannot (e.g. more de-

A Warning.

We have shown that hedonics can be used to produce indexes which both eliminate much of the selection bias in the matched model index and are likely to have lower variance. Thus I hope it is clear that hedonic price indexes deserve increased attention. Still it is important to conclude this section by stressing that the distance between the upper bound provided by the hedonic index and the least upper bound we are after – say the income transfer that would just insure that the average consuming unit would be at least as well off in the comparison period – may be quite large.

What the hedonic does is allow us to compute price relatives for goods that exit, thus correcting for the bias in the matched model index generated by selection on the observed characteristics of the goods that survive. Neither index makes any adjustment at all for the infra-marginal rents that go to consumers who have reservation prices higher than the highest sampled price ever listed for the new good; for those who value the increased speed of a new PC more than its increased price, or those who valued the portability advantages of laptops more than the difference between their prices and PC prices. Nor does the hedonic register any gains for the increase in the *variety* of the goods offered in a specific characteristic range. To capture these effects we need estimates of a model of household utility.

tailed outlet characteristics), then the measured price relatives of the goods that survive will have different expectations in the matched model and hedonic indexes (the former contains $E(\epsilon_{t+1}|\epsilon_t, survival)$. This same term will appear in the expectation of the observed, but not of the estimated, price relatives in the complete hybrid.

4 PC Price Indexes; Data and Regressions.

We now consider alternative price indexes for PC's. Following the BLS we focus on indexes that are sales weighted averages of estimated price relatives, i.e. if I_t is the index

$$I_t = \frac{\sum_{i \in C_{t-1}} h^t(x_i^{t-1}) q_i^{t-1}}{\sum_{i \in C_{t-1}} p_i^{t-1} q_i^{t-1}} \equiv \sum_{i \in C_{t-1}} \left(\frac{h^t(x_i^{t-1})}{p_i^{t-1}} \times \frac{p_i^{t-1} q_i^{t-1}}{\sum_{i \in C_{t-1}} p_i^{t-1} q_i^{t-1}} \right), \qquad (6)$$

where, as above, $h^t(\cdot)$ is the t^{th} period estimate of the hedonic function.

We begin with a description of our data and then consider issues of functional form and of right hand side variables (for more detail see Pakes (2002)).

4.1 The Data.

We acquired quarterly data from IDC on quantity and mean price by; company, brand name, chip type (generation), and processor speed. We then matched more detailed characteristics and price data from various PC magazines to this sample[13]. 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 characteristic and

[13]For more on IDC see www.IDC.com. 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 (PC World, PC Magazine, Computer Shopper, PC Computing, ...). The additional sources were especially important for 1999.

price data we gathered and a weight computed by splitting the IDC category quantity equally among the observations that fit that 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 their characteristics, 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 larger than the sample size generally 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 underlying their hybrid index (for this they download samples of between 800 and 1200 observations from the internet, see below). Consequently we expect *less* sampling variance in our estimates but *more* variance due to prediction error in our estimates than in the BLS's hybrid.

The table lists the major (but not the only) characteristics used in the hedonic function. Note that they are all "vertical" (we expect every consumer to prefer more of them), and the min, mean, and max of each of them increases over time (illustrating the "quality" improvement that occurs in this market). The rate of change in these characteristics, however, varied significantly between periods; there was a striking increase in quality in 1998 (the year the pentium II became dominant in the market).

Very few machines marketed in the base period are also marketed in the comparison period (between 8 and 23 percent). As a result when we constructed the matched model index we had to drop 84 percent of the observations. This leaves ample room for selection biases.

4.2 Choice of $h^t(\cdot)$.

We care about the mean and the variance of the price index the hedonic regression generates, so our choice of $h^t(\cdot)$ should be determined by these magnitudes (and not by the precision, stability, or sign of any subset of coefficients). The mean square error of the price predictions from a given regression is closely related to its adjusted R^2's, so Table 2 uses this statistic to compare alternative hedonic specifications.

All regressions include; speed and speed squared ("spd" in megahertz), ram and maxram ("rm" and "maxrm" in megabytes), and hard-drive capacity ("hd" in gigabytes). They also include a set of (mostly dummy) variables for "add-ons" (see the notes to the table for a list), 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 table).

Price itself is always the dependent variable. We compared a specification which was linear in the right hand side variables to one which used them to a power. The power function always produced at least as high an adjusted R^2 in all periods so we use it. The difference in adjusted R^2 between our "base" and our "fully loaded" specification (the last row which contains the type dummies and more interactions), or indeed between either of these and any intermediate specification, is quite small. As a result we use *both* the base and the fully loaded specifications below.

Table 3 presents coefficient estimates 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 is positive and sometimes it is negative and the sign of the squared term is always opposite to the sign of the linear term. Other coefficients also vary quite a bit, especially in the fully loaded specification, and sometimes are the "wrong" sign (though never significantly so).

As noted instability and reversals of signs over time are neither unusual results for hedonic regressions nor a particular source of worry. The large variances of some of the parameter estimates could be more troubling. There are at least two ways of trying to reduce that variance; adding data (we consider this below), and constraining coefficients across years. Though there is no reason to expect coefficients to be similar across periods, if they are and we enforce equality we reduce the mean square error of our predictions.

Table 4 presents formal tests for the stability of the coefficients over time. We ran two sets of tests; one constrains only the coefficients of the variables appearing in Table 3, and the other constrains all variables (including the additional included variables listed at the bottom of Table 2). The χ^2 test statistics indicate rejection of just about any hypothesis on the constancy of the coefficients (and the power of the rejections just goes up when we also constrain the additional variables)[14]. The one marginally significant

[14]To obtain the test statistics we first 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

test statistic occurs when we use the fully loaded specification and test for stability only of the Table 3 coefficients and then only over the first three years. However when we constrain those coefficients in the fourth year to be the same as those for the first three all test statistics explode (three out of four have values over 10,000). The fourth year is the year of the mass introduction of the Pentium II; this changed the PC hedonic function markedly.

Not surprisingly the data indicate coefficients which change over time. Recall, however, that we are only worried about instability in coefficients to the extent that it causes instability in our index. The next section shows that this is not the case.

5 Price Indexes for Desktops.

We begin with a hedonic indexes and then compare them to other indexes which either have been, or could be, used by the BLS. Table 5 presents estimates of a number of hedonic indexes and their standard errors. Where there are two estimates of the standard error in brackets, the first is an estimate of the overall standard error and the second estimates what the

the difference between the estimated and constrained coefficients weighted by the inverse of the estimated variance-covariance of the parameter estimates. The χ^2 statistic the estimated value of that quadratic form. Many of the "additional" variables are dummies for characteristics appearing for the first time during the sample period. Interestingly their coefficients start out with large coefficients which become smaller as the characteristics become standard offerings.

standard error would be if there were only sampling (no estimation) error[15].

The top and bottom panel of the table use the "base" and the "fully loaded" set of variables for the hedonic regression (see table 3). In the top sub panel of each panel we sum over the estimated price relatives for all goods marketed in the base period, while in the lower one we drop 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). Each sub panel contains two indexes; the second uses the NLLS coefficients from Table 3 to form the numerator of the price relatives. The first expresses the power function preferred by the data in standard log-log form and uses OLS to estimate the hedonic coefficients. This latter regression predicts the log of price, so when we exponentiate to form predicted price we have to account for the expectation of $exp[\epsilon]$ which we assume equals $exp[.5\hat{Var}(\epsilon)]$ (i.e. the disturbance in the log form is assumed i.i.d. normal)[16].

[15] 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, which, together with the bootstrapped sample of prices and characteristics, are used to compute a hedonic price index. The estimate of total variance is the variance in this index over alternative bootstrapped samples. The estimate of the variance due to sampling error is obtained similarly except that we use the original hedonic regressions to calculate the numerator of the price relatives for all samples (i.e. we do not allow for variance in the estimated coefficients).

[16] This correction does matter. Without it all the numbers are about 20 percent lower.

We computed twice as many indexes as presented. The indexes not presented used $h^t(x)$ instead of $p_t(x)$ in the denominator of all price relatives. They were virtually identical to those presented in the table[17]. Indeed the striking empirical result is just how *similar* all the computed indexes are. It does not matter much whether we; use all the price relatives or drop out the price relatives for goods in $t-1$ whose characteristics are outside $\overline{C}$ (compare the top half of each panel to the bottom half); use the log-log linearization and the "normal" adjustment for the variance in the disturbance or use the nlls estimates directly (so the assumption underlying the adjustments seem reasonable), or use the base or the fully loaded specification for the set of included variables (compare the bottom and top panels of the table).

Importantly this implies that the inclusion of marginally significant characteristics does not effect either the indexes *or* their estimated variances. This is in spite of the fact that their inclusion does change the regression coefficients in Table 3. "High end" machines tend to be better in all dimensions. This produces *both* high variance in coefficients and invariance of the price predictions to the variables included in the regression. Note also that if we can use the log-log regression, the hedonic coefficients can be computed

[17]Whether we should be using $p_t(x)$ or $h^t(x)$ in the denominator of the price relatives of the hedonic index depends on how we interpret the BLS procedure. The interpretation that favors $h^t(x)$ is that they draw random x tuples and then compare the average base to the average comparison period prices for each of them. Alternatively the disturbance could be viewed as representing omitted characteristics that we want to condition on; then we would use p_t in the denominator.

from an analytic formula, and the BLS can fully automate the procedure which constructs the index.

All the indexes agree on three substantive empirical results. First the average price change over the period was a large negative number; the estimates range between 15.6 percent and 17.6 percent. Second, there are significant differences in the estimated indexes across years. 1998, the year of the pentium II, saw a rate of price fall of about double the rate in the earlier years and 1999 is somewhere in the middle. Finally about half of the variance around the estimates seems to be due to sampling error and about half from prediction error (so with sample sizes smaller than ours and better characteristic data we would expect the variance in a hedonic produced from BLS data to be mostly due to sampling error).

5.1 Comparisons to Alternative Indexes.

Table 6 begins by comparing our hedonic (its first panel) to two matched model indexes (its second panel). The Tornquist matched model 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 Anna Aizcorbe, Carol Corrado and Mark Doms (2000), and the literature they cite). The Laspeyres is an arithmetic weighted average using base period sales weights.

The row labeled " percent matched" indicates that only 15 percent of the base-period observations are matched to a comparison period product.

Though this is indicative of a large selection bias in the matched model indexes, it is startling just how large a role selection seems to play. Even the Laspeyre's index, the matched model index with the largest price decline, had a rate of decline less than a fifth of that of the hedonic indexes[18]. In most years the positive effect of selection just about offset the negative effects of technological change on the matched model indexes and both were close to zero. However 1998 was different. In 1998 the pentium II obsoleted 92 percent of the older machines and the 1997 products that did continue were often early pentium II models; models which increased both their prices and sales. The Tornquist index, which weighs these products more heavily, was positive in that year while the Laspeyres was negative (interestingly this produces a correlation between the two matched model indexes of minus one).

The monthly match rate for the BLS's PC sample translates into an annual match rate which is *virtually identical* to that in Table 6. Thus if the selection bias in matched model indexes is linear in the fraction of goods dropping out, the *annual* selection biases implicit in the BLS's monthly matched model indexes would be the same as the bias from our data. It should not be surprising then that in 1998 the BLS analysts decided to

[18]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 indexes, and for two of the years they were smaller.

replace the standard matched model index for PC's with their own hybrid hedonic index.

The BLS's Hybrid and the "Complete" Hybrid.

The BLS's hybrid partitions 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 comparison period (full matches), goods for which they are able to match firm brand and all but a small number of listed characteristics to a base period product (partial matches), and goods that cannot be matched at all. The full, partial, and "no" matches are, respectively, about 85 percent, 12 percent, and 3 percent, of the monthly observations (the implied annual match rates are 15 percent, 55 percent and 30 percent). To construct their hybrid they drop the "no match" cases and compute a sales weighted average of; (i) the actual price relatives from the full matches and (ii) a hedonically adjusted price relatives for the partial matches.

The hedonically adjusted price relatives are obtained as the ratio of the actual price of the "matched" comparison period good to an adjusted base period price. The adjusted base period price is the actual base period price plus the hedonic function's "evaluation" of the difference in characteristics between the base and the matched comparison period good (computed as the difference in their characteristics times the estimated coefficients of those characteristics). The price relatives for the PPI are formed in a similar way except that the comparison (instead of the base) period price is adjusted for the differences in characteristics.

The coefficients for the adjustment are obtained from a *constrained* hedonic regression using data downloaded from the internet. Why constrained? Recall that hedonic regressions coefficients need not equal consumers' marginal characteristic valuations so hedonic coefficients of vertical characteristics can have "wrong" signs (or "implausible" magnitudes) and this can make BLS's adjustments seem absurd. If one constrains the relationship between price and characteristics to be linear signs typically make more sense (it is higher order and interaction terms that tend to generate adjustments of the "wrong" sign in portions of the characteristic space). Also since the impact of speed is clearly nonlinear (Table 3), the coefficients from a separate regression of the cost of chips to PC producers on measures of chip performance are used for the speed adjustments (see Michael Holdway (2001); note that this is a cost, not a consumer valuation, based adjustment). Finally if any of the linear coefficients still turn out to be "incorrectly" signed they are set to zero.

The BLS procedure still incurs the selection bias resulting from dropping the goods not matched at all (now 30 percent, instead of 85 percent, of the goods in annual samples). Also the modified Konüs argument we used to justify the hedonic index does not apply to the BLS's hybrids. Finally price relatives based on the BLS's adjustments are likely to have *much* larger variance than those from our hedonic procedure, particularly in "high tech" markets. In these markets vertical characteristics are typically highly correlated across products (our "high end" PC's, which have more speed, memory, storage capacity,...) and this produces negatively correlated regression coefficients. Consequently the weighted sums of coefficients our hedonic uses to predict comparison period prices will be estimated more precisely than the individual coefficients used for the BLS's adjustments.

We divided our sample into fully, partially, and unmatched goods and then mimicked the BLS's hybrid using alternative estimates of the hedonic function[19]. When we used the regression coefficients from Table 3 for the adjustments the indexes varied from plus to minus a thousand percent or more, making it evident why the BLS analysts used constrained regressions. So we used the coefficients from a linear regression of price on the characteristics in Table 3 for the adjustments (the only second order term in this specification is speed squared). These results appear as the $lin(CPI)$ (which applies the adjustment to the base period product) and the $lin(PPI)$ (which applies it to the comparison period product) rows in panel 3 of table 6. They show similar averages but *very* different patterns over time than the hedonic indexes (they show a price rise in 1997/98 and the largest price falls in the earlier two periods). Moreover they have variances which are an *order of magnitude* larger than those of our hedonic indexes.

[19]15 percent of our sample was fully, 40 percent was partially, and 45 percent were not matched. A base period product was labeled 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. 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.

We could not also mimic the BLS's speed adjustment because their data was confidential. However we did find data on the price at which Intel sold their later generation chips to PC producers (from the original pentium) and we used it to estimate a "cost-based" speed adjustment for Intel chips. Using these coefficients for the speed adjustments, and partially adjusting only machines based on Intel chips of a pentium or later vintage, gave us the rows in panel 4 of Table 6. The speed adjustment helps some with the "variance" problem but not with the "timing" problem in the BLS's hybrid[20].

Our *complete* hybrid, say C, should ameliorate the variance, selection, and interpretation problems with the BLS's hybrid, and *may* have less variance than the hedonic. C is a sales weighted average of the Laspeyres matched model index for the base period goods which survive into the comparison period, and a hedonic index for *all* the goods that do not (both the partial and no matches). I.e. $C = wM + (1-w)I(d)$, where w is the fraction of sales in continuing goods, M is the Laspeyres matched model index, and $I(d)$ is the hedonic index for the goods that do not survive.

Note that since our hedonic, or $I = wI(c) + (1-w)I(d)$, where $I(c)$ is the hedonic for the goods that continue, $C - I = w(M - I(c))$. So when the match rate is low C will be similar to I. It is not surprising then that our estimates of C and its variance (panel 5 of table 6), are not significantly different from our estimates of I and its variance. However, since the monthly w's will be

[20]Interestingly the actual BLS hybrid for 97/98 and 98/99 were -.358 and -.265, respectively, both larger (in absolute value) but within two standard deviations of, the figures in Table 6.

larger than those from annual data, this conclusion is likely to vary were we to have access to the BLS's monthly data (also the difference between M and $I(c)$ may not be proportional to the length of the sampling period). Thus though our results indicate that the BLS's hybrid is problematic, we withhold judgment on the choice between the complete hybrid and the hedonic until an experiment is run on monthly data.

Modifications to the Hedonic.

We conclude with two hedonic like indexes that might be easier for the BLS to implement than our hedonic. The first is obtained from a log-log hedonic regression that constrains all coefficients but period specific dummy variables to be the same across periods, and uses the dummy variables' coefficients to form the index. Its virtue is that it can be constructed from a simple procedure that can be fully automated (it is also familiar from economic research). However since hedonic coefficients vary across periods it has no theoretical justification, and since the BLS has to produce their index in real time the index for the current period would have to be obtained from a regression which only uses data from *prior* periods.

Panel 6 of table 6 provides the standard dummy variable index, while panel 7 provides a dummy variable index that the BLS could actually use, but only for the last two periods (the index for 97/98 is obtained from a regression which uses data from 95 to 98, while the 98/99 index is obtained using data from 96 to 99). The index in panel 6 "smooths out" the differences in price changes across years (the absolute difference with the hedonic averages 30 percent and is statistically significant in one year)[21]. So it is not surprising that the index which the BLS could use (panel 7) has the unattractive feature that it is noticeably lower (higher) then the hedonic index when the hedonic indexes for the years immediately preceding the current year are lower (higher) than that year's hedonic.

The "Paasche-like" hedonic index is calculated by using the base period hedonic function to estimate what prices would have been for the comparison period goods in the base period, and then forming price relatives from the ratios of the actual comparison period prices to those predictions; i.e. as $1/n_t[\sum_{i=1}^{n_t}(P_i^t/h^{t-1}(x_i^t))]$. Its advantage is that, since the hedonic function it uses can be estimated *before* the current period's data arrives, it may be computable in a more timely manner than our index. It also asks the (slightly) different question; how much income would make a consumer who purchased the comparison period bundle in the base period at least as well off in the comparison as in the base period?

The Paasche-like index in panel 8 is quite close to our hedonic; closer (in absolute value) than *any* of the other indexes. It does have larger variances than the Laspeyres hedonic in three of the four years. This because "high

[21]Note that its standard error *is* lower than that of the hedonic. This raises the issue of whether we might constrain some coefficients over some years to reduce the variance in the hedonic index. We tried several alternative ways of doing this (e.g. constraining the coefficients of only the first three years, and using a two-step procedure designed to minimize mean square error) but none of them lead to an index with a (bootstrapped) estimate of mean square error less than that of our unconstrained hedonic.

end" comparison period characteristic bundles are not observed a lot in the base period, and hence have hedonic values that are estimated imprecisely[22]; a fact that has to be weighed against any timing advantages.

6 Concluding Remarks.

This paper shows that there are relatively easy to construct hedonic price indexes that can be justified in terms of providing an upper bound to the compensating variation. These hedonic indexes circumvent a selection problem in matched model indexes that, as our example shows, can be quite severe for commodity groups in which there is a lot of technical change.

There are three disadvantages of hedonic indexes. First the upper bound to the compensating variation they provide need not be particularly close to the least upper bound we are after. Second, in addition to the sampling variance that is inherent in matched model indexes, they also have variance due to prediction error. Finally it may be difficult for the BLS to produce the hedonic and still meet their monthly deadlines.

The fact that the hedonic may be too generous to the consumer is a natural result of looking for an index which can be justified non-parametrically;

[22]Indeed the reason that we did not use the fully loaded specification to calculate the Paasche hedonic is that then we would need to use the base period regression function to evaluate say, types of chips, that were not present in the base period. With related problems in mind we also calculated the Paasche 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.

i.e. of looking for an index whose justification is valid regardless of the precise form of the distribution of utility functions (so this issue also arises when using Laspeyres' indexes in product space). Though some experimentation with price indexes which use estimated distributions of utility functions seems worthwhile (as this might provide an indication of the bias in traditional indexes and the extent to which hedonics bridges the gap), they are not likely to be incorporated in our official indexes for some time.

There are a number of low cost ways for the BLS to augment their samples to reduce estimation variance[23]. However if estimation variance remains a problem it would seem sensible to move to the complete hybrid, rather than the matched model, index. The hybrid eliminates estimation variance to the same extent as the matched model and mitigates the matched model's selection problems.

Finally the BLS's data gatherers now carry hand held computer and download their data daily onto a central computer. As noted there are automated procedures which can compute our hedonic index from that data within seconds of receiving it. If the BLS were uncomfortable with using the automated estimate of the hedonic function (and they might be at least for an initial period), use of a "Paasche-like" hedonic would provide them with a month to scrutinize the estimated hedonic function before using it.

References.

[23]The use of data downloaded from the internet was noted in the text. Alternatively the data collectors could double the number of quotes they get at a given outlet, and use the second set of price quotes only for estimation (not for price relatives).

Aizcorbe, Ana; Corrado, Carol and Doms, Mark. “Constructing Price and Quantity Indexes for High Technology Goods.” *Mimeo*, Federal Reserve Board, Washington D.C., 2000.

Anderson, Simon P.; de Palma, Andre and Thisse, Jacques-Francois. *Discrete choice theory of product differentiation.*, Cambridge and London: MIT Press, 1992.

Bajari, Patrick and Benkard, Lanier. “Demand Estimation with Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach.” Working paper, Stanford University, 2003.

Berry, Steven; Levinsohn, James and Pakes, Ariel. ”Automobile Prices in Market Equilibrium.” *Econometrica*, July 1995, 63(4), pp. 841-890.

Berry, Steven; Levinsohn, James and Pakes, Ariel. ”Differentiated Product Models from a Combination of Micro and Macro Data: The New Car Market”, *Journal of Political Economy,* forthcoming. N.B.E.R. Discussion paper 6481.

Court, Andrew. “Hedonic Price Indexes with Automotive Examples”, in *The dynamics of automobile demand*, General Motors Corporation, 1939, pp. 99-117.

Dulberger, Ellen. “The Application of a Hedonic Model to a Quality-Adjusted Price Index for Computer Processors.” In D. Jorgenson and R. Landau (ed.s) *Technology and capital formation.* Cambridge: MIT Press, 1989, pp. 37-75.

Epple, Dennis. “Hedonic Prices and Implicit Markets: Estimating Demand and Supply Functions for Differentiated Products.” *Journal of Political Economy,* February 1987, 95(1), pp. 59-80.

Feenstra, Robert. “Exact Hedonic Price Indexes.” *Review of Economics and Statistics*, November 1995, 77(4), pp. 634-653.

Griliches, Zvi. “Hedonic Price Indexes for Automobiles: An Econometric Analysis of Quality Change”, in *The price statistics of the federal government*, New York, National Bureau of Economic Research, 1961, pp. 173-196.

Heckman, James. ”Shadow Prices, Market Wages, and Labor Supply.” *Econometrica*, July 1974, 42(4), pp. 679-94.

Holdway, Michael. “Quality-Adjusting Computer Prices in the Producer Price Index: An Overview.” Obtained from the BLS web site, 2001.

Hulten, Charles and Wykoff, Frank. “The Estimation of Economic Depreciation Using Vintage Asset Prices: An Application of the Box-Cox Power Transformation.” *The Journal of Econometrics*, April 1981, 15(3), pp. 367-96.

Konüs, A.A. "The Problem of the True Index of the Cost of Living" *Econometrica*, January 1939, 7(1), pp. 10-29.

Liegey, Paul. "Hedonic Quality Adjustments in the U.S. CPI: A Statistical Agency Perspective", Working paper, Bureau of Labor Statistics, 2001.

Moulton, Brent, LaFleur, Timothy and Moses, Karen. "Research On Improved Quality Adjustment in the CPI: The Case of Televisions." Working paper, the Bureau of Labor Statistics, 1998.

Nevo, Aviv. "New Products, Quality Changes and Welfare Measures Computed form Estimated Demand Systems." *The Review of Economics and Statistics*, May 2003, 85(2), pp. 266-275.

Pakes, Ariel; Berry, Steven and Levinsohn, James. "Applications and Limitations of Some Recent Advances in Empirical I.O.: Price Indexes and Analysis of Environmental Change." *American Economic Review*, May 1993, 83(2), pp. 240-46.

Pakes, Ariel. "A Framework for Applied Dynamic Analysis in Industrial Organization." Working Paper, NBER # 8024, 2000.

Pakes, Ariel. "A Reconsideration of Hedonic Price Indexes with an Application to PC's." Working Paper, NBER # 8715, 2002.

Pollak, Robert. *The theory of the cost-of-living index.* New York: Oxford University Press, 1989.

Rosen, Sherwin. "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition." *Journal of Political Economy*, Jan.-Feb. 1974, 82(1), pp. 34-55.

Song, Minjae. "Investment, Innovation, and Consumer Welfare." Unpublished Ph.D thesis, Harvard University, 2003.

Triplett, Jack. "The OECD Hedonic Handbook." Mimeo, *The Brookings Institution*, forthcoming.

Appendix: A Gradient Inequality.

Let $U(x, y-h(x))$ provide the utility an agent derives from consuming a good with characteristics x when there is $y-h(x)$ income left over to use in any way desired. Here $h(x)$ is the price of x. A person who chose x_t in period t and is given the income $y+[h^{t+1}(x_t)-h^t(x_t)]$ in period $t+1$ will have at least as high a utility in period $t+1$ as in period t if

$$sup_{x \in C(t+1)} \left[U\left(x, y - h^{t+1}(x) + h^{t+1}(x_t) - h^t(x_t)\right) - U(x_t, y - h^t(x_t)) \right] \geq 0.$$

Assuming the utility function is well approximated in a neighborhood of x_t, say $N(x_t)$, by the differential from x_t this condition becomes

$$sup_{\ x\ \in\ N(x_t)\ \cap\ C(t+1)} \left[\frac{\partial U}{\partial x}(x - x_t) + \frac{\partial U}{\partial y}\left(h^{t+1}(x_t) - h^{t+1}(x)\right) \right] \geq 0,$$

where all derivatives are evaluated at the point $(x(t), y - h(x(t)))$.

If $N(x_t) \cap C(t)$ contains a sufficient number of products then the consumer's choice will satisfy $U'_x \approx U'_y \frac{\partial h^t}{\partial x}$. Since we assume $U'_y > 0$, our condition then becomes

$$sup_{x \in N(x(t))\ \cap\ C(t+1)} \left[\frac{\partial h^t}{\partial x}(x - x_t) - \left(h^{t+1}(x) - h^{t+1}(x_t)\right) \right] \geq 0.$$

Finally if $h^{t+1}(\cdot)$ is sufficiently smooth then this can be rewritten as

$$sup_{x\ \in\ N(x(t))\ \cap\ C(t+1)} \left[\frac{\partial h^t}{\partial x}|_{x(t)} - \frac{\partial h^{t+1}}{\partial x}|_{x(t)} \right] (x - x_t) \geq 0,$$

which is the condition used in the text. ♠

Table 1: Characteristics of Data*.

year	95	96	97	98	99
# observations	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 loaded 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

[*1]See 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 percent but not at 5 percent.

[*4] accept at .25*percent* but not at .1 percent

Table 5: Hedonic PC Price Indexes and Their Variance*

		95/96	96/97	97/98	98/99	av.
Specification With Base Variables Only.						
$x \in$						
C_{t-1}	log-log	-.102	-.111	-.292	-.172	-.169
		(.037/.023)	(.059/.037)	(.041/.023)	(.092/.025)	(.09/n.c.)
	NLLS	-.097	-.108	-.295	-.155	-.164
		(.04/.024)	(.063/.037)	(.045/.021)	(.099/.022)	(.09/n.c.)
$x \in$						
$C_{t-1} \bigcap \overline{C}_t$	log-log	-.100	-.115	-.267	-.161	-.161
		(.032)	(.054)	(.038)	(.062)	(.08)
	NLLS	-.094	-.111	-.270	-.150	-.156
		(.039)	(.052)	(.044)	(.054)	(.08)
Specification With Type and Interactions.						
$x \in$						
C_{t-1}	log-log	-.106	-.123	-.277	-.188	-.174
		(.035/.028)	(.066/.041)	(.041/.026)	(.128/.021)	(.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	-.105	-.126	-.269	-.173	-.168
		(.035)	(.057)	(.043)	(.062)	(.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 corresponds to the estimate of the actual variance, and the second corresponds to the estimate of just the sampling variance component. All standard errors are estimated by a bootstrap based on 100 repetitions (higher numbers of repetitions did not change the estimates).

Table 6: Alternative PC Price Indexes*

	Year	95/96	96/97	97/98	98/99	av.
1. Hedonics	base	-.097	-.108	-.295	-.155	-.164
		(.040)	(.063)	(.045)	(.099)	(.091)
(or, I)	f.a.	-.094	-.111	-.270	-.150	-.156
		(.039)	(.052)	(.044)	(.054)	(.079)
2. Matched	Tornquist	.012	.002	.09	.011	.028
model (M)	Laspeyres	-.013	-.002	-.08	-.011	-.027
	percent matched	16.6	22.8	8.0	11.5	14.7
3. BLS	lin(CPI)	-.159	-.419	.008	-.059	-.157
		(.37)	(.11)	(.43)	(.48)	
Hybrids	lin(PPI)	-.167	-.661	.007	-.059	-.220
		(.71)	(.40)	(.15)	(.81)	
	percent matched	43.6	71.7	50.3	45.2	52.7
4. (with speed	lin/spd(CPI)	-.264	-.409	-.193	-.082	-.237
		(.20)	(.12)	(.23)	(.11)	
adjustment)	lin/spd(PPI)	-.345	-.623	-.208	-.082	-.315
		(.28)	(.26)	(.14)	(.24)	
	percent matched	25	64.1	48.7	35.7	43.4
5. Complete	base	-.125	-.135	-.300	-.168	-.182
		(.043)	(.062)	(.046)	(.115)	
Hybrids (C)	f.a.	-.128	-.164	-.292	-.194	-.195
		(.042)	(.063)	(.044)	(.141)	
6. Dummy	base	-.135	-.098	-.160	-.170	-.141
		(.038)	(.035)	(.027)	(.040)	(.032)
Variables	f.a.	-.152	-.122	-.213	-.143	-.158
		(.040)	(.032)	(.041)	(.028)	(.039)
7. BLS Feasible	base	n.r.	n.r.	-.156	-.168	n.r.
Dummies		n.r.	n.r.	(.034)	(.036)	n.r.
8. Paasche	base	-.104	-.167	-.337	-.119	-.181
Hedonic		(.07)	(.049)	(.091)	(.062)	(.070)

*Standard errors appear in brackets below estimate. They are estimated by a bootstrap based on 100 repetitions. "base" refers to base specification and "f.a." refers to fully loaded specification in table 2. "n.r." = not relevant.