An Experimental Component Index for the CPI: From Annual Computer Data to Monthly Data on Other Goods.

December 18, 2009

Abstract

Until recently the Consumer Price Index consisted solely of "matched model" component indices. The latter are constructed by BLS personnel who visit stores and compare prices of goods with the same set of characteristics over successive periods. This procedure is subject to a selection bias. Goods that were not on the shelves in the second period were discarded and hence never contributed price comparisons. The discarded goods were disproportionately goods which were being obsoleted and had falling prices. Pakes (2003) provided an analytic framework for analyzing this selection effect and showed both that it could be partially corrected using a particular hedonic technique and that the correction for his personal computer example was substantial. The BLS staff has recently increased the rate at which they incorporate techniques to correct for selection effects in their component indices. However recent work shows *very little* difference between hedonic and matched model indices for non computer components of the CPI. This paper explores why.

We look carefully at the data on the component index for TV's and show that differences between the TV and computer markets imply that to obtain an effective selection correction we need to use a more general hedonic procedure than has been used to date. The computer market is special in having well defined cardinal measures of the major product characteristics. In markets where such measures are absent we may need to allow for selection on unmeasured, as well as measured, characteristics. We develop a hedonic selection correction that accounts for unmeasured characteristics, apply it to TVs, and show that it yields a much larger selection correction than the standard hedonic. In particular we find that matched model techniques underestimate the rate of price decline by over 20%.

1 Introduction

This paper develops new hedonic methods for producing price indices. Price indices are often constructed using matched model techniques. Matched model indices are built from prices obtained by data collectors who visit stores and record the prices of goods along with their characteristics. The index is constructed as a weighted average of the ratios of two prices for a good with the same characteristics; one obtained in a base and one in a comparison period (or as a weighted average of the logs of such ratios). When a good which was sampled in the base period was not on the shelves in the comparison period that ratio is not available, so the index is formed as a weighted average of the ratios for the remaining goods. Goods that exit, and hence whose price changes are not included in the index, are disproportionately goods whose characteristics have been obsoleted, and hence whose prices have declined. So omitting these goods tends to remove price changes from the left tail of the distribution of price changes, causing an upward bias in the estimate of the average price increase. Hedonic techniques can be used to correct for this bias.

A simple hedonic procedure consists of estimating a price for the exiting good in the comparison period and substituting the ratio of that price to the observed base period price for the missing price ratio when forming the index. To obtain the estimated prices a hedonic regression of price on characteristics is run. The predicted price is formed as that regression's prediction for the exiting good's characteristics. Pakes (2003) used a model of a differentiated-product market as a framework for clarifying the conditions under which hedonic techniques provide a bound on the transfer needed to compensate consumers for changes in their choice sets (for the "compensating variation"), and showed that use of a hedonic technique that satisfied those conditions caused a dramatic change in the price index for computers.

One of the required conditions is that period-specific hedonic regressions be used to form the estimated prices. This because hedonic regressions do not identify either utility or cost parameters; they do provide a prediction of the market price for a given set of characteristics, but that prediction typically changes markedly from period to period with changes in market conditions.

Until very recently hedonic predictions that were based on regression functions that were updated every period were difficult, if not impossible, to do within the BLS's time constraints. The BLS recently modernized its data gathering procedures by giving its data gatherers hand held computers and enabling them to download the data they gather onto the BLS's central computer at the end of every work day. This has made it possible for the BLS to use hedonics in the construction of the component indices that underlie the CPI. However when the BLS's hedonic procedures were tried on different component groups most of the resultant indices were not much different than the matched model indices for those groups (see the results reported in Table 5 of the survey by Johnson, et. al, 2006).

We begin by comparing the price indices obtained from standard hedonic procedures to those obtained from matched model procedures that do no correction for selection bias for one of the BLS's component indices (TV's; throughout all the data used is data gatherred by the BLS for the construction of the CPI). The findings replicate the earlier results; the hedonic and matched model indices have similar values. We then show that the reason this occurs is not an absence selection bias, rather it is because standard hedonic procedures do little to correct for this bias. The TV market is different from the computer market in that it does not have sharp cardinal measures of most of the characteristics that consumers value. Instead most of our TV characteristics are dummy variables indicating the presence or absence of advanced features. Moreover exit is disproportionately of high priced goods that have most of these features. They exit because they are obsoleted by newer high priced goods with higher quality versions of the same features, and we do not have good quality indices for those features. As a result in the TV market, and we suspect in many other markets, selection is partly based on characteristics the analysts cannot condition on, i.e., on what an econometrician would call "unobservables".

Standard hedonic predictions do not account for the price differences generated by characteristics the analyst does not condition on. One alternative is to augment the standard hedonic with a good-specific "fixed effect" to account for the unobserved characteristics of the good, and then use the coefficients from a regression for the differences of prices of continuing goods on observed characteristics to predict the change in price of the exiting goods. We show that though this procedure does move the index in the expected direction, it only corrects for a small part of the problem. This should not be surprising. We expect the hedonic evaluations of different characteristics to vary across periods. Since the residual summarizes the effects of many unobserved characteristics each of whose value is changing over time, the value of the residual should change over time thus invalidating the fixed effect procedure.

The goal of this paper is to develop hedonic procedures which; at least partially account for the contribution of unobserved characteristics, are robust to the properties of data sets in ways we will make precise, and can be implemented within the BLS's time constraints. The next section of the paper explains how different indices are constructed and compares the standard hedonic to the matched model index for our TV sample. In section 3 we provide evidence that this finding occurs despite the fact that their *is* a selection problem in the TV sample. In particular the exit rate in this sample is about the same as it was in computers and their is ample evidence indicating that the prices of exiting goods are falling at a faster rate then those of continuing goods. Section 4 introduces our hedonic regressions and presents evidence indicating that the residual from that regression, our measure of the value of a good not captured by observable characteristics, is an important determinant of exit.

Section 5 develops hedonic formulae for price predictions which partially account for the role of unobserved characteristics. We provide two of these; one requires data which will only be available in sufficient quantity for particular commodity groups, but is likely to provide a tighter index when available. This section concludes with a third alternative for price prediction; one that has intuitive appeal but which uses assumptions which are hard to justify. If all our assumptions are correct the average (over periods) of the third alternative should lie between the averages of the first two, so we use the third index as a robustness check on our procedures.

Section 6 provides the empirical results from using the various procedures to obtain price indices for our primary sample period. They indicate that standard hedonics overstate inflation by over 20% and are striking in their consistency with our theoretical predictions.

These results were obtained after considerable experimentation; experimentation which would not be possible if the index had to be constructed to meet BLS's deadlines. They were also obtained for a component index with a relatively large number of price quotes. Section 7 of the paper considers two new samples; a TV sample from a later period and a digital camera sample. The later TV sample is used to see whether our index for TV's could be used in "production mode". The digital camera sample allows us to explore what can be done when a smaller number of price quotes are available and enables a constructive comparison to the procedures currently in use at the BLS.

We apply the procedures used for the early TV sample essentially *without change* to the later TV sample and compare the indices that result. Though the later sample is different in important ways from our earlier one, our "out of sample" prediction produced results that strongly reinforced our earlier conclusions. The digital camera sample is a subsample of the sample used by the BLS to construct its photographic equipment index, and has a relativley small number of price quotes, too small to use our more detailed hedonic techniques on (indeed current BLS procedures do not make *any* use of hedonics in their photographic equipment index). The results for the indices we can compute for digital cameras show smaller corrections to the matched model than those obtained for the two TV samples but an identical orderring of indices.

Current BLS procedures for constructing the TV and photgraphic equipment indices do make a partial correction for the selection induced by exit. Price relatives for some of the cameras which exit are dropped, but when a "close" substitute (defined by the good's observed characteristics) is found in the comparison period it is used to form a price ratio, and when a substitute is found but it is not deemed close enough a price ratio is imputed as the average of the price ratios for the goods that exited and had a close substitute. The procedure is similar for TV's except that in TV's about half of the exiting products, products for which there were substitutes but they were not deemed sufficiently close to warrant using them directly, used a variant of hedonic techniques to construct a price ratio.

Section 8 constructs an index we call the CPI mimic; it uses exactly the same price ratios for the goods that exited as the CPI used for those goods, and a matched model price ratio for the goods that do not. The CPI mimic for all three samples shows a noticeably lower rate of deflation than our indices do. Indeed the CPI mimic for the camera sample, the mimic that does not use hedonics at all, produces an index which is lower than the pure matched model index and the difference is both economically and statistically significant. This section also shows that there is an index, closely related to the intutive index we used for robustness checks, that is easier to construct than the CPI mimic and falls at a faster pace than the matched model in all three samples, significantly so in two of them. A short concluding section closes the paper.

2 The Selection Problem and Hedonic Indices.

We will only consider indices that are weighted averages of the logs of "price relatives", where a price relative is defined to be the ratio of a price for the good in the comparison to that in the base period[1]. These indices are referred to as log "geomean" indices, and if we let $\tilde{y}_{qt}$ be the log of an estimate the of ratio of good q's price in period t to that in period $t-1$, they can be written as

$$G_t = \sum_{q \in S_{t-1}} w_{q,t-1} \, \tilde{y}_{qt} \tag{1}$$

where $w_{q,t-1}$ is a (base period) weight and S_{t-1} is a subset of the price quotes available for period $t-1$. The weights $w_{q,t-1}$ we use are taken from the expenditure share weights computed by the BLS[2].

[1] The linear in logs assumptions makes the indices linear in the regression error from the logarithmic hedonic regressions we and others have used to estimate the relationship between price and product characteristics; a fact that makes the relationship of our results to the underlying data transparent. We intend to come back to more complex indices that work directly with this regression error at a later date, as they have a larger role to play in other indices. In particular to construct the Laspeyre's index we need to exponentiate the logs and hence exponentiate the hedonic regression error, and the Laspeyre's index is the only index with a direct interpretation as a bound on compensating variation.

[2] For all but the digital camera sample, the weights we use are obtained by dividing each period-$(t-1)$ regional expenditure-share equally among all the quotes for that region and then renormalizing them so that $\sum_{q \in S_{t-1}} w_{q,t-1} =$

Denoting an actual log price-relative by $y_{qt} = log(p_{qt}/p_{q,t-1})$ and an estimate of a log price relative obtained from the hedonic regressions introduced presently as $\hat{y}_{qt}$, hedonic and matched-model indices can be written as

$$G_t^{hed} = \sum_{q \in A_{t-1}} w_{q,t-1}^{hed} \hat{y}_{qt} \tag{2}$$

$$G_t^{mm} = \sum_{q \in C_{t-1}} w_{q,t-1}^{mm} y_{qt}, \tag{3}$$

where A_{t-1} is the set of quotes for which prices were successfully collected in period $t-1$, and $C_{t-1} = A_{t-1} \cap A_t$. That is matched model indices average *actual* price relatives for goods for which price information was collected in *both* periods, while the hedonic averages *predicted* price relatives for *all* goods whose prices were collected in period $t-1$[3]. It is the fact that the matched model index omits the price changes for goods which had price quotes in period $t-1$ but not in period t that generates its selection bias.

Goods that exit, and hence whose price changes are not included in the index, are disproportionately goods whose characteristics have been obsoleted, and hence whose prices have declined. So omitting these goods tends to remove price changes from the left tail of the distribution of price changes, causing an upward bias in the estimate of the average price increase. Hedonic indices partially correct for this bias. Formally the matched model index implicitly imputes the average of all continuing goods price relatives of the price relative for the exiting good while the hedonic index weights the price relatives of continuing goods that have observed characteristics more similar to the exiting goods disproprotionately.

As emphasised in Pakes (2003), since the relationship between the prices of goods and their

1. The regional expenditure shares are obtained from the CPI data base. These weights were not available for our camera sample, so we present unweighted indices for it (see below).

[3]An early version of this paper also computed the hybrid indices introduced in Pakes (2003). These impute relatives only for TVs that exit between $t-1$ and t, and use actual price relatives for goods that were available in both periods, i.e. $G_t^{hyb} = \sum_{q \in C_{t-1}} w_{q,t-1}^{hed} y_{qt} + \sum_{q \in A_{t-1}-C_{t-1}} w_{q,t-1}^{hed} \hat{y}_{qt}$. The theoretical attraction of hybrids is that they have no estimation error in their price relatives for the continuing goods, and they eliminate much of the selection bias in the matched model index by using hedonic predictions for the goods that exit. On the other hand they treat the error from the hedonic regression differently for the two types of goods, and this can cause a (different) selection bias. We decided that the tradeoff between bias and variance was not an issue we wanted to deal with in this paper and hence omitted the hybrids. We do consider a different hybrid below, one that serves both as a robustness check on our results, and as an index that the BLS can produce in minimal time.

characteristics changes with almost any change in market conditions (entry, exit, shifts in demand and/or cost,...), for the hedonic prediction to bound the needed compensating variation the hedonic regression on which it is based must be done separately in every period. That regression should include all relevant characteristics and should not be constrained in any way.

The results presented here are based on a seperate linear regression model for the log of price in every period.[4] If Z_t is the $n_t \times K$ matrix of characteristics of all TVs for which prices were collected in period t, and p_t is the corresponding $n_t \times 1$ vector of log prices, then a period-t hedonic regression coefficient is given by

$$d_t = \left(Z_t' Z_t\right)^{-1} Z_t' p_t, \tag{4}$$

and the prediction for log price is $\hat{p}_t = Z_t d_t$. As required by the theory we do not place any restrictions on the coefficients per se or on their relationships across different periods.

Until very recently hedonic predictions that were based on regression functions that were updated every period were difficult, if not impossible, to do within the BLS's monthly time constraints. The fact that the BLS has provided their data gatherers with hand held computers and instructed them to download their data nightly onto a central BLS data management system has changed this situation. Now an index based on monthly hedonic regressions can be used as long as the commodity analyst can clean the downloaded data within the BLS's time constraints, and an appropriate index construction algorithm is programed into the BLS's computers.

However when the BLS's hedonic procedures were applied to some of their commodity groups most of the resultant indices were not much different from the relevant matched model indices (see the results reported in Table 5 of the survey by Johnson, et. al, 2006). To illustrate consider the TV component index that we analyze below. The sample gathered for the TV index has 20% turnover over the two month sampling interval in the period we analyze (and a higher rate in the later period used for "out of sample" testing in section 7 below). This is almost identical to the rate of turnover

[4] An earlier version of the paper also presented results based on a local linear non-parametric kernel hedonic regression for the four and ten characteristic data sets. The non-parametric results did not differ in any substantive way from the results reported below.

in computers where the use of hedonics had a dramatic impact on the component index. Moroever as we show in the next section there is ample evidence indicating that the goods that exit the TV sample have prices that are falling disproportionately. Not surprisingly then the BLS now does use a hedonic regression in constructing the TV component index. Their hedonic regression has a large set of explanatory characteristics and is run only once a year (see Moulton, Lafleur and Moses, 1998, and section 7 of this paper for more detail on the current CPI TV component index).

Table 1 compares the matched model to hedonic indices for our sample. The hedonic predictions were constructed from hedonic regression coefficients obtained from a separate regression for every period of the log of price on a set of twenty four characteristics that are similar to the set of characteristics used by the BLS. The hedonic index has about the same value as that produced by the matched model procedure (and the hedonic is more variant across months). Moreover were the BLS to use a characteristic set this large the amount of data cleaning needed would imply that they could not produce an index which used a new hedonic regression every period and still abide by their time constraints. Since the hedonic can not be justified in terms of a bound on the compensating variation unless the regression underlying it is done separately for every period, we also computed a hedonic index from regressions based on a ten variable characteristic set which does not require extensive cleaning and could be used in a production setting (we come back to these variables below). When we use standard hedonic procedures with this data we get a hedonic index which falls at a *much slower pace* then the matched model index.

The reason that the standard hedonic produces results which are similar to the matched model index is not that there is no selection bias, rather it is because standard hedonic procedures do little to correct for this bias in this market. The TV market is different from the computer market in that it does not have sharp cardinal measures of most of the characteristics that consumers value. Instead most of our TV characteristics are dummy variables indicating the presence or absence of advanced features (see the Appendix). Moreover exit is disproportionately of high priced goods that have most of these features. They exit because they are obsoleted by newer high priced goods

Table 1: **Matched Model and Standard Hedonic Indices.***

Index Calculated	matched model	hedonic
hedonic uses S24	-10.11	-10.20
s.d. (across months)	5.35	7.53
S24 % l.t. mm		.50
hedonic uses S10	-10.11	-8.82
s.d. (across months)	5.35	7.05
S10 % l.t. mm		.40

*30 monthly indices are computed for each of the matched-model and two hedonic specifications, spanning the 31-month interval from June 2000 to Dec 2002. The dependent variable in both hedonic specifications is the log of price. The right hand side variables are the characteristics in S24 and S10 described in the text just prior to Table 4. The rates in the table are the implied annual percentage inflation rate. The last row gives the fraction of the 30 months where a hedonic index is less than the matched-model index (% l.t. mm).

with higher quality versions of the same features, and we do not have good quality indices for those features. As a result in the TV market, and we suspect in many other markets, selection is partly based on characteristics the analysts cannot condition on, i.e., on what an econometrician would call "unobservables".

Standard hedonic predictions do not account for the price differences generated by characteristics the analyst does not condition on. One alternative is to augment the standard hedonic with a good-specific "fixed effect" to account for the unobserved characteristics of the good, and then use the coefficients from a regression for the differences of prices of continuing goods on observed characteristics to predict the change in price of the exiting goods. We show that though this procedure does move the index in the expected direction, it only corrects for a small part of the problem. This should not be surprising. We expect the hedonic evaluations of different characteristics to vary across periods. Since the residual summarizes the effects of many unobserved characteristics each of whose value is changing over time, the value of the residual should change over time thus invalidating the fixed effect procedure.

We begin with a description of the data and a set of reduced form results which illustrate just what the problem with standard hedonic indices is.

3 Background: Properties of the Data and Biases in the Index.

Our primary data set consists of CPI price quotes for the 35 months between March 2000 and January 2003 and a matching characteristic data set built up from the characteristic set used by the BLS's TV commodity analyst for the procedure currently used to construct hedonic adjustments for the TV component index[5]. The average monthly sample contains prices and characteristics from 234 observations. We note that prices in this data ranged from \$66 to over \$10,000, reflecting the rather extreme differences in products that the BLS includes in this commodity group.[6]

Just over three quarters of the CPI price quotes are collected at 2-month intervals from odd and even numbered month subsamples each of which are regionally defined. The other one quarter of the quotes are from New York, Los Angeles and Chicago and are collected at one month intervals. As a result we focus on price relatives, exits, etc. over two month periods (our sampling interval), though all the sample observations available for the two months period are used (whether from the one month or one of the two month subsamples).

On average, 22.5% of the TVs present in any period are not present in the following period, with 19.7% being permanent exits. The non-permanent exits are goods that were available in the prior period, are not available in the current period, but returned to the shelf in a later period. Similarly, 24.0% of TVs in the current period were not present in the prior period, with 17.0% being substitutes (the good that was to be sampled for comparison period prices was not present at the outlet so another good had to be substituted for it) and 4.1% being scheduled additions to the sample (goods that were scheduled to be rotated out of the sample). An average of 2.9% of the

[5] A "cleaned characteristics" subset of each period's July and August data was prepared by the CPI commodity analyst for use in their current hedonic procedure. We assigned the cleaned characteristics to all months by matching model numbers. The resulting 35-month data set contains 8,195 prices, or 79.9% of all prices. On average the months have mean, median, minimum, and maximum prices equal to \$725, \$366, \$81, and \$7836 respectively. Comparing, where possible, statistics for the full and cleaned data sets shows that the latter data is very similar to the full data. Noteworthy departures are slightly lower entry and exit rates (making our problem harder) and a mean price that is about \$40 higher than that for the full data.

[6] As in most markets, the entry and exit of particular TVs tends to disproportionately influence, and be disproportionately influenced by, prices of close competitors. To insure that the hedonic predictions for one good were not overly sensitive to goods which were in very different parts of the product space, an early version of this paper included hedonic indices based on local-linear nonparametric hedonic regressions. We have omitted those results because they did not differ substantially from the results based on the log-linear approximations given below.

exits are temporary, while 2.9% of entering TVs are returning from temporary absence.[7]

Price relatives from different subsets of the data. Price relatives for different subsets of the data play a key role in this paper. We partition the price relatives between any two periods into three groups; a group for which there is a price relative in the prior period but which exit before the next period (our "about-to-exit" or "a-exit " price relatives), a group for which there is a price relative in the following period but not in the preceding period (our recently new or "r-new" price relatives), and a group for which we have price relatives in both adjacent periods (our "other" price relatives). The "full sample" of price relatives refers to the union of these three groups. On average (over months) there were 183.45 price relatives between any two periods and of these; 40.2 (22.4%) are about to exit relatives, 46.0 (25.5%) are recently new relatives, and the remainder (97.25 or 52.1%) are the other price relatives[8].

A good which exits on the day prior to the last period's sampling date will not have a price relative for the current two periods while a good which exits the day after the sampling date will, and the sampling dates are spread across the two month sampling interval. Consequently the behavior of the price relatives for about to exit goods should be more like that of goods which do in fact exit than that of a randomly drawn price relative (direct evidence supporting this assertion is provided below). As a result we will use the a-exit price relatives for clues as to the unobserved price relatives for goods that were in the sample in the first period but exited before the second.

Table 2 provides some summary statistics on the price relatives for the full sample and for these three subsets of the data. We begin with the data on the about to exit goods. The first point to note is that their price relatives show a faster rate of price decline then the other groups of goods.

[7] A good that is temporarily off the shelf may be absent for quite different reasons than goods that have permanently exited. In particular temporary exits may be caused by stock-outs, while permanent exits are more likely to be caused by obsolescence. However the number of temporary exits was too small to cause any noticeable differences in the results reported below. We note that the numbers above come from slightly different series. The exit rates are computed on a series that excludes the last 4 months from each bimonthly subsample, the deleted months used to determine which exits eventually return. Computation of the entry rates exclude the earliest months from each subsample for analogous reasons.

[8] Since we need to be able to check for the existence of a prior price relative to determine whether a price relative is r-new and for the existence of a following price relative to check whether the price relative is a-exit, this table is based on a data set which drops out the first and the last two months from our data series.

The about to exit goods prices decline at about *twice* the average rate of decline and the difference is highly significant (with a t-ratio of about six). If goods that are about to exit have prices declines that are more similar to those of goods that do exit, then these numbers reinforce the belief that, by throwing out the goods that exit, matched model procedures overestimate inflation.

The last panel of this table uses the data from the quarter of the sample with monthly observations. By calculating the first month price relative decline rates of the goods that exit in the second month of the two month sampling interval we can provide direct evidence on the price relatives for exiting goods in the first month of the two-month period in which they actually do exit. We then compare these rates to the rates for the same goods in the period in which they are classified as about to exit.

On average the two month rate of decline of the price relatives for the sample with monthly observations is similar to that of the overall sample, as is that sample's average two month price relative for about to exit goods. 52% of the monthly observations that exit over the two month sampling intervals have observed prices after the first month. The average price relative of these goods for the *first month* is .9756, which is noticeably lower than the average *two month* price relative for the full sample of monthly data (.9835). The average two month rate price relative for *about to exit* goods in the monthly sample is .9679, which is lower than the average one-month price relative of the goods that exit in the second month of the sampling interval. However were we to assume that the rates of price decline were lower in the one-month period of exit then in the one-month period before exit, say because the periods in which goods exit tend to be periods in which goods are under increased price pressure, then we would know that the goods that exited in the second month of the two month period had two month price relatives lower than $(.9756)^2 = .9518$. Below we weaken the assumption that the rate of price decline in the month of exit is greater than in the preceding month and then use the monthly data to get lower bounds to the rate of price decline for exiting goods under the weaker assumptions. However the "back of the envelope results" presented here are illustrative of the more conservative results below.

Interestingly recently introduced goods also have price relatives that on average fall at a faster pace than the other goods, though the difference is not nearly as striking as it is for about to exit goods (it is only 1/4 to 1/5 the differential for a-exit goods, and for r-new goods the difference with other goods is not statistically significant). Still this finding has implications for price index construction procedures. As noted by Pakes (2003) introducing new goods earlier into the index will only ameliorate new goods biases if prices fall in their introductory periods. It seems that early introduction of new goods would indeed ameliorate new goods biases in TVs.

Finally we note that the results in Table 2 go a long way towards explaining the difference in results for matched model indices based on different intervals of time. Compare, for example, the average of the matched model indices with a two month sampling interval with that from a four month sampling interval. The latter omits price changes of two types of goods that are included in the two month interval data; (i) goods that are "about to exit" in the first two month interval, and (ii) goods that are "recently new" in the second two month interval. Both these subgroups of goods have prices that fall at a faster rate then a randomly drawn continuing good. So the four month interval index misses two groups of prices changes whose prices are falling disproportionately.

The fact that the longer sampling interval data omits price changes of about to exit goods accentuates the selection bias we study here. The fact that it omits initial price changes of recently entered goods, will, in markets where initial prices fall disproportionately, accentuate a bias we do not attempt to correct for in this paper. This is the bias caused by the fact that the index does not attempt to capture the inframarginal rents which accrue to individuals who would have bought the new good at a price higher than the highest price at which the new good entered the index (see Pakes, 2003, for further discussion). To get some indication of how these biases increases with the length of the sampling interval, we used our data to calculate the matched model indices when we assumed two, four, and twelve month sampling intervals. The annualized rate of deflation for the three intervals were, respectively, -10.59%, -8.99%, and -6.48%.[9] So going from a two month to an

[9]These indices used equally weighted price relatives (since we did not have expenditure weights) and cover a twenty four month period (so we could compare indices over exactly the same period).

annual interval increases the matched model's estimate of inflation by about 40%. More generally the longer the sampling interval the less accurate the matched model index's measure of inflation is likely to be.

Table 2: **Price Relatives.**

Variable	Full Sample.	a-exit	r-new	other	exit-other	new-other
mean	.9849	.9729	.9844	.9881	-.0152	-.0037
(s.d. of mean)	(.0010)	(.0024)	(.0019)	(.0014)	(.0028)	(.0023)
cross-section s.d.	.0677	.0778	.0606	.0646	n.r.	n.r.

Using the Subsample with Monthly Price Quotes				
variable	All Monthly Data 2-month	a-exit 2-month	late exits (exit after month 1 but before month 2), 1-month	(month 1 exit by month 2)^2
mean price relative	.9835	.9679	.9756	.9518
(s.d. of mean)	(.0016)	(.0036)	(.0068)	(.0136)
# of obs.	1428	334	207	207

Prices of Entering and Exiting Goods. The next table summarizes information on the prices of entering and about to exit goods which will help with an understanding of the role of selection in this market. It has coefficients and t-values from regressions of log prices on a constant and two dummies, one for the goods that just entered and one for goods that are about to exit. The regressions are done differently for odd and even numbered periods as the BLS samples different cities in those periods.

The point made by this table is that both the newly entering goods and the about to exit goods have prices that are *higher* than those of continuing goods. This is not surprising for newly entering goods as new goods enter at the high quality end of spectrum in many markets. What is somewhat surprising is that this is also true for goods that are about to exit. This differentiates the TV market from the market for computers where almost all exits are from the low end of the quality spectrum in the period before they exit. I.e. like in computers, in TV's most new good enter at the high end of the price spectrum. However at least in this period the TV's exitors are also typically high end goods (presumably displaced by the high end entrants). The "low-end" products in the

Table 3: **Characteristics of Entering and Exiting goods.**

Specification	Constrained OLS		Minimum Distance	
	exit	new	exit	new
1. S0 (Odd)	.107	.161	.076	.146
	(2.66)	(4.14)	(1.94)	(3.86)
2. S0 (Even)	.121	.133	.097	.130
	(3.17)	(3.53)	(2.61)	(3.51)

S0 has a constant and two dummies, one for goods about to exit and one for goods that just entered. Odd and Even number periods are done separately as they represent samples from different regions. The constrained OLS estimator is an estimator which allows for different constants in the different months, but constrains the dummy for about to exit (or new) goods to be the same across months. The minimum distance estimator begins by estimating different dummies for each month and then forms the estimates in the table as a weighted average of the monthly estimates with weights inversely proportional to their estimated variances.

TV market do not turn over nearly as much.

Though our characteristic data are rich enough to enable our hedonic function to capture price differences between high and low quality TV's, they are not well suited to distinguish between two high quality TV's one of which is based on older technology and hence has been obsoleted. For example we know which TV's are flat-screen CRT display, but we do not have a good measure of the improvements that have occurred in sharpness of that display over time. This is a second feature which differentiates the TV market from the computer market. In the computer market the major characteristics that are improving over time have natural cardinal measures which make them easy to compare across products (e.g., speed, RAM, hard drive capacity, ...).

4 The Coefficients and Residuals of Hedonic Regressions.

This section provides our hedonic regressions of price on characteristics and examines their properties with an eye to the construction of price indices. All regressions were done separately every month, and used the log of price as the dependent variable. We tried three sets of regressors, denoted by $S5$, $S10$, and $S24$, and defined as follows.

- $S5$: log of screensize in inches, a dummy indicator for projection TVs, the interaction between these two variables, the square of log-screensize, and a dummy variable for whether the observation comes from the monthly subsample[10],

- $S10$: the variables in $S5$ plus dummy indicators for picture-in-picture, flat-screen CRT display, HDTV-ready, a high-quality reputation Brand A, and a low-quality reputation Brand Z,
- $S24$: the variables in $S10$ plus the additional variables listed in Table 13 in the Appendix.

The values for the variables in $S5$ and $S10$ can be verified with minimal effort on the part of CPI staff, and therefore can be used to fit an up-to-date hedonic regression in the time interval available to the BLS when producing their index. This is not so for the additional variables in $S24$. The current hedonic procedure incorporated into the CPI index for TVs uses a list of regressors which is similar to that in S24. Most of the regressors in S24 but not in S10 have values that are difficult to verify in the short period of time between when the BLS obtains the new price quotes and when it has to have produced the index. This is the reason why the current hedonic method used by the BLS fits a regression no more than once a year.

The first three rows of Table 4 show that any of the three sets of characteristics does quite a good job of accounting for variance in the traditional dependent variable of hedonic regression, log-price. Even S5 has R^2's between .87 and .91. There is a noticeable improvement in fit in moving from $S5$ to $S10$, but not much further improvement in adding the 14 characteristics needed for $S24$.[11] Part of the reason for the closeness of the S10 and S24 measures of fit is that the TV's with S10 features have most of the S24 features. Still the striking fact illustrated by table 4 is just how well we do in predicting price; the S10 regressor set gives us R^2's between .94 and .97.

It is not unusual to get high R^2's for hedonic regressions on a set of prices obtained from a given differentiated product market; indeed it is a major reason for the increased use of characteristic

[10] As noted in Pakes, 2003, the hedonic regression can differ with the characteristics of the population in which the goods are marketed as well as with the characteristics of the goods per se. We found that the only population distinction which helped to predict price in our analysis was the dummy variable indicating the observation was one of the three big cities which define the monthly subsample

[11] The improvement in fit in going from S10 to S24 is very close to the improvement we got in moving from the linear regression in S10 to the non-parametric local linear kernel regression in those variables.

based models in demand estimation. However these R^2's are higher than usual. This is partly a result of the quality of the BLS data, and partly a result of the large variance in price of the goods sampled in the TV component index combined with the ability of our characteristic data to differentiate between the "low and high end" of that market.

Recall that we use hedonics to predict the *change in price* of exiting goods. Price changes over a two month period are small, and much more difficult to predict (the regression of price changes on characteristics leads to an average R^2 of only .024). Our worry about using standard hedonics is that these price changes might be largely generated by a change in the value of the characteristics the hedonic does not condition on. This worry is accentuated by the fact that, other than screen size, all the characteristics in S10 are dummies for the presence or absence of advanced features. In particular we do not have a *measure of the quality* of those features and the fact that turnover is concentrated in high end products is indicative of a process of obsolescence in those qualities. So we now turn to a more careful look at the role of unobserved characteristics.

Table 4: **Hedonic Regressions: Dependent Variable is Log-Price**

Regressors	mean R^2	mean adj R^2	min R^2	min adj R^2	max R^2	max adj R^2
S5	.896	.894	.873	.870	.913	.911
S10	.956	.954	.942	.937	.967	.965
S24	.971	.967	.959	.953	.978	.975

R^2 statistics from log-price regressions run on each month from March 2000 to January 2003.

4.1 Unobserved Characteristics and Hedonic Bounds.

Under standard assumptions on consumer behavior the prices of two goods with identical characteristics should be the same. So if we observed all relevant product characteristics we should be able to predict the prices of goods that exit the sample from the prices of goods with similar characteristics that remain in sample[12]. This prediction problem, however, gets more complicated when

[12] For a statement of this property, and a demand estimation algorithm that makes intensive use of it, see Bajari and Benkard (2005). They require a choice set that fills up a subset of characteristic space. For justification of hedonic indices when the choice set is not this rich see Pakes (2003).

there are characteristics of the goods that consumers value but Econometricians do not observe (and hence can not condition on). So we begin by asking whether there is a need to pay attention to unobserved product characteristics in predicting exiting goods prices.

We look first to the properties of the hedonic regression function per se. Part of the impact of the unobserved product characteristics on price in that regression will be captured by the relationship between unobserved and observed characteristics, but the rest will appear as the residual. If the relationship of the residual to the observed characteristic were no different for exiting goods than for a randomly drawn good we could obtain an unbiased estimate for the price of a good that exited the sample between two periods from the hedonic regression coefficients in the second period and the characteristics of the good that exited (even though this regression only uses observations on the continuing and entering goods). However if unobserved characteristics are important determinants of whether a good exits, then simple economic arguments should lead us to believe that the regression function for goods that exit differs from that for continuing goods, and that this difference will impact its price predictions in systematic ways.

To see this assume, for simplicity, that the true hedonic function is linear and let η measure the contribution of unobserved characteristics to price, so that

$$p = z\beta + \eta, \tag{5}$$

where we have normalized the coefficient of η to be one. Our hedonic equation is obtained from a regression of p on z. To analyze its properties we need the properties of the regression of η on z. Letting $j = \{x, c, n\}$ denote, respectively, exiting continuing and entering goods, then

$$E[\eta|z] = \Sigma_{j=\{c,x,n\}} P\{j|z\} E[\eta|z, j].$$

Though the theory that tells us that goods with the same characteristic should sell for the same prices implies the coefficients on z in equation (5) should not differ between entering, exiting and continuing goods, it says nothing about whether $E[\eta|z, j]$ differs by j. Moreover a standard selection argument would lead us to believe this regression function does differ by j.

If higher priced goods are more profitable (say, because they had higher markups), and more profitable goods are less likely to exit, then if we compare goods with the same observed characteristics those with lower η values will be more likely to exit. Indeed a good whose observed characteristics are associated with high prices will only exit if it has a low value of η, while a good whose observed characteristics are associated with low prices will exit even if it has a relatively high value of η. Before considering the implications of this line of reasoning for the construction of price indices, we look to see if the data support it.

First we ask whether there are significant differences in the relationship between z and η for exiting, continuing, and newly entered goods. To this end we estimated hedonic regressions for each period which allowed each of the three groups of goods to have different z-coefficients. Using the $S10$ regressor set of the last subsection, we then tested whether these coefficients differed from each other. Table 5 presents the fraction of the months for which we could reject the equality assumption based on different significance levels. As pointed out by a referee, since some goods are in the sample for multiple periods and the hedonic errors for a given good tend to be serially correlated, the tests for the different months are not independent of one another. As a result we present results for the for an assortment of significance levels. We can insure that the probability of rejecting the null of equal coefficients when it is true is less than .05 by rejecting only when there is a month that' has a "p-value" less than $.05/28 \approx .0017$. Over 40% of the months have p-values less than that.

Table 5: **Testing for Exit and New Good Interaction Terms.**

p-level	$p = .05$	$p = .005$	$p = .002$	$p = .001$
Fraction of Months Significant At Different α Levels				
$j = x$; Wald-test	.786	.536	.429	.429
$j = n$; Wald-test	.714	.571	.536	.500

Test uses heteroscedastic consistent covariance matrix and a χ^2-statistic for critical values.

Second, if unobserved characteristics are an important determinant of whether a good exits and the marginal costs of the goods are relatively constant over time, then we would expect a negative correlation between the change in the hedonic prediction for price of continuing goods and the values of the residual for those same goods. This because if costs are constant then goods whose observed characteristics increase disproportionately in value will continue even if their η value decreases disproportionately, while goods whose observed characteristics' values decrease disproportionately will only continue if their η values increase disproportionately. The correlation of the change in the observed and unobserved components of price for the continuing goods in our sample was **-.53**, just as a model where selection is partly based on the unobservable would predict.

We can also look for evidence on selection in the difference between the disturbances from the hedonic regressions for about to exit and continuing goods. In particular we can both; (i) compare their values in the period prior to exit, and (ii) compare the change in their values between the two periods prior to the exiting period.

The change in η results also throw light on the appropriateness of a simple procedure for correcting for selection bias; one based on the assumption that the contribution of unobserved characteristics to price does not change over time. If selection was based only on observed characteristics and a *time invariant* unobserved characteristic, or a "fixed effect", the average of $\eta_{t+1} - \eta_t$ should not differ between exiting and continuing goods. So under the fixed effect assumption we can form an unbiased prediction for exiting goods prices by regressing the difference of the logs of the prices of the *continuing* goods onto their characteristics, and then using that regression function to predict the price change for the exiting goods from the change in the valuation of their observed characteristics. Note the arbitrary difference in the way the fixed effect assumption treats the unobserved and the observed characteristics; it calculates its price change predictions from differences in the contribution of the observed characteristics to price over time but assumes, *a priori*, that the contribution of unobserved characteristics to price never changes.

The average difference between the η's of exiting goods and those of continuing goods in the period prior to exit was negative, but only slightly so, and the difference was not statistically significant[13]. Table 6 presents the results from splitting the data into three groups – the goods that are about to exit (they exit during the next sampling interval), those recently new (they enter in period t), and the remaining goods – and then calculating the average change in the residual for each. About to exit goods have an average change in residual which is significantly negative (with a t-value above five). Moreover, though the average change in residual of all goods which continue is also negative, it is less than a fifth the absolute value of the average of the change for about to exit goods. The value of the unmeasured characteristics of goods that are about to exit are falling at a dramatically faster pace than those of continuing goods.

There are a number of implications of this table that are worth noting. First since the contribution of the unobserved characteristic to price is falling at a striking rate just prior to exit, it is likely to be falling during the exiting period (and probably at a faster rate, as the exiting period is the period in which the changes in the environment actually caused exit). This reinforces our worry about the impact of selection on unobserved characteristics on price indices, and rules out corrections based on the assumption that the unobserved characteristic's contribution to price is constant over time. Second the fact that average change in the residual of all the continuing goods is also negative indicates that the new goods that enter have unobserved characteristics that, on average, are more valuable then those of continuing goods (which, given the above discussion, should not be surprising). This fact also has the implication that we would improve on standard hedonic correction for selection, a correction which ignores unobserved characteristics, by making an adjustment for the change in the market's value of the unobserved characteristics of the exiting goods equal to the measured change in the evaluation of the unobserved characteristics of the *continuing* goods. The next section develops this idea further.

[13] We note that this appears to be largely a result of the limits we had to place on the richness of our specification for the hedonic regression. When we considered the difference in residuals between exiting and continuing goods that were "twins" in the sense that they had exactly the same value for all the S10 characteristics the results were stronger; the average difference for the 1103 twins had a t-value of -1.85. Interestingly when we defined twins more liberally and required only that they have the same value for the S5 characteristics, the number of twins goes up to 1359 and the t-value moves up disproportionately to -3.13. That is when we are able to compare goods with the same obseved characteristics (which is what we would do if we could estimate a "perfect" regression function) then when we omit observable variables we increase the difference between the average residuals of continuing and exiting goods, just as we would expect when selection is partly based on unobservables.

Table 6: **Hedonic Disturbances for About to Exit, Recently Entered, Goods.**

Variable	All Continuing	a-Exit	r-New	Remaining Goods.
Using the S10 Specification for the Hedonic Regression[1].				
mean	-.0027	-.0157	-.0049	-.0022
s.d. of mean	.0010	.0026	.0021	.0014
s.d.(across months)	.0090	.0150	.0130	.0130
percent < 0	.6207	.8966	.5517	.5172

[1] See the description of the S10 specification.

5 Hedonic Bounds With Unobserved Characteristics.

Adding an i subscript to differentiate goods and a t subscript to differentiate time periods, our hedonic equation (equation 5) becomes

$$p_{i,t} = z_i\beta_t + \eta_{i,t}. \tag{6}$$

Note that the observed characteristics of the good (the z_i) are constant over time, though the market's evaluation of those characteristics (the β_t) changes with changes with market conditions. The additive disturbance is now not the contribution of unobserved characteristics per se but rather the residual from regressing their contribution to price onto the observed characteristics.

Our problem is that we do not observe the value of $p_{i,t+1}$ for the goods that exit between t and $t+1$. This section provides predictions for $p_{i,t+1} - p_{i,t}$ conditional on $(z_i, \eta_{i,t})$ which, given our assumptions, maintains the hedonic bound in the sense that the resultant predictor for $p_{i,t+1} - p_{i,t}$ will have an expectation which is larger than the expectation of $p_{i,t+1} - p_{i,t}$ conditional on z_i and $\eta_{i,t}$. We develop two bounds. The first only uses the information in the bimonthly sample. The second adds information from the monthly sample on prices at the end of the first month of the bimonthly sampling period. We conclude by noting that there is another computation which, at least intuitively, should provide a lower bound to the average price adjustment over the entire sample period, and it can be used to check the robustness of our results.

5.1 Hedonic Bounds From the Bimonthly Data.

From equation (6) we have

$$E[p_{i,t+1} - p_{i,t}|z_i, \eta_{i,t}] = z_i[\beta_{t+1} - \beta_t] + E[\eta_{i,t+1} - \eta_{i,t}|z_i, \eta_{i,t}], \tag{7}$$

where

$$E[\eta_{i,t+1} - \eta_{i,t}|z_i, \eta_{i,t}] = \sum_{j_{i,t}} E[\eta_{i,t+1}|j_{i,t}, z_i, \eta_{i,t}]Pr\{j_{i,t}|z_i, \eta_{i,t}\} - \eta_{i,t}.$$

We can estimate the probability of continuing and the expected change in η *when the good continues*; i.e. when $j_{i,t} = c$. However to get an upper bound for the price index we also need an upper bound for this conditional expectation when $j_{i,t} = x$. To obtain this bound we need a rule for when the good exits. The rule we use is contained in the following assumption.

Assumption 1 (Exit Rule.)

$$j_{i,t} = x \quad \Leftrightarrow \quad \eta_{i,t+1} \leq \underline{\eta}_{t+1}(z_i). \qquad \spadesuit$$

Note that there are no restrictions on $\underline{\eta}_{t+1}(z_i)$. It is a free function of the z_i in every period and that function can change over time[14].

If we were in a world in which goods were taken out of the market when the price they could be sold at was lower than their marginal cost, then Assumption 1 would be equivalent to the assumption that when we look across goods which could have been marketed in a given period, the price they could have been successfully marketed at increases in η more than their marginal costs does. If differences in η represent differences in unmeasured quality characteristics, then an

[14] However we are assuming that there is a single index of unobserved quality which determines its impact on the exit decision. The obvious generalization is to allow for multiple indices of unobserved characteristics and relate their values to the exit decision. If there were two indices we would need assumptions which produced a curve in two dimensional space with the feature that goods with values for the couple of indices which are above the curve would continue while those below would exit (see for e.g. Bajari and Benkard, 2005a). We do not pursue this further because it would lead to an index correction procedure which would both require more assumptions and be more difficult for the BLS to implement.

increase in profitability (measured as the difference between price and marginal cost) as η increases should be expected, as it is this increased profitability which justifies the sunk cost of developing higher quality goods.

If assumption 1 does hold then

$$E[\eta_{i,t+1} - \eta_{i,t} \mid j_{i,t} = x, z_i, \eta_{i,t}] = E[\eta_{i,t+1} - \eta_{i,t} \mid \eta_{i,t+1} \leq \underline{\eta}_{i,t+1}(z_i), \eta_{i,t}, z_i] \leq \tag{8}$$

$$E[\eta_{i,t+1} - \eta_{i,t} \mid \eta_{i,t+1} \geq \underline{\eta}_{t+1}(z_i), \eta_{i,t}, z_i] = E[\eta_{i,t+1} - \eta_{i,t} \mid j_{i,t} = c, z_i, \eta_{i,t}] \equiv gb(z_i, \eta_{i,t+1}).$$

That is the assumption guarantees that the conditional expectation of $\eta_{i,t+1} - \eta_{i,t}$ for continuing goods, an expectation *we can estimate*, provides an upper bound for the unobserved conditional expectation of $\eta_{i,t+1} - \eta_{i,t}$ for exiting goods.

There are cases where Assumption 1 *might* not hold. One example is a situation where the sunk cost of producing a good does not increase in η but the marginal cost does, and it does so at a faster pace then consumer's willingness to pay for η increases. However our bounds only require the inequality in equation (8), an inequality on the change in η (rather than its level), so for a good which exited to violate the bound it must have been profitably produced in a prior period. This should make it possible to detect cases where the bound is in doubt.

What the inequality states is that once we condition on z goods which exit will, on average, have a lower value for the change in η then goods which continue. If we believe that about to exit goods behave more like exiting goods than a randomly chosen good, we can check whether the inequality in equation (8) has been satisfied in the past for any commodity group by comparing the average of the change in η for about to exit goods to that for continuing goods (recall that in our TV sample the rate of η decline of about to exit goods was five to six times that for continuing goods). This does not rule out changes in current costs or market structure (say a merger) which cause goods with η's that people are willing to pay more for and were profitable to produce to no longer be profitable. However the BLS assigns a commodity analyst to each component index and the analyst should be able to spot changes in industry structure or in cost conditions that could have caused such a problem and correct for them when it does.

The change in η regression for continuing goods. To see whether the inequality in equation (8) is likely to help us tighten our bound for the price index we computed the R^2's from the regression of $\eta_{t+1} - \eta_t$ on a polynomial in η and z for continuing goods on the bimonthly data (see Table 7). The η's used here are the residuals from the cross sectional hedonic regressions done separately for each period. As a result the η's from the full sample are uncorrelated with the z's by construction. So if selection were not partially based on the "unobserved" characteristics (our η), we would expect the adjusted R^2's in the second column of table 7 to be zero. The fact that they are highly significant is evidence that the selection into continuing goods is at least partly based on our unobservables. More importantly for present purposes, the fact that the adjusted R^2's go up significantly when we add η_t to the regressions (see column 4 of table 7) indicates that using the information on the η_t of exiting goods together with the conditional expectations of $\eta_{t+1} - \eta_t$ for continuing goods will help bound the change in η for the exiting goods[15].

Of course the bound in (8) may not be very "tight". In particular the last subsection showed that the unobservables for about to exit goods had: (i) systematically lower values of $\eta_{i,t}$ and (ii) systematically lower values of $\eta_{i,t+1} - \eta_{i,t}$ given $\eta_{i,t}$. The bound from equation (8) will make a correction for the lower values of $\eta_{i,t}$ of exiting goods, and for the negative trend in $\eta_{i,t+1} - \eta_{i,t}$ of continuing goods (table 6). However it does not account for the fact that the $\eta_{i,t+1} - \eta_{i,t}$ for exiting goods tends to be less than that for continuing goods[16]. This lack of tightness of the bound is the

[15]Formally we can replace the expectations in equation (8) with expectations conditional on anything in the information set in period t that help predict $\eta_{t+1} - \eta_t$ and the inequality still holds. We noted earlier that the regression function for $\eta_{i,t+1} - \eta_{i,t}$ for newly entered goods might be different than that for other continuing goods, and when we did the $\eta_{i,t+1} - \eta_{i,t}$ regression once using a dummy for newly entered goods we got a small improvement in fit. This explains the last two columns in table 7 and we use predictions that allow for this dummy in what follows (though we get very similar results when predictions without this dummy are used).

[16]We show in an earlier version of this paper (Erickson and Pakes, 2007) that this source of upward bias in the bound in equation (8) can, at least in principal, be corrected if we are willing to make one more assumption; that the stochastic process generating η is Markov and *independent* of z. Recall that each period's $\eta_{i,t}$ is uncorrelated with z_i by construction. The additional assumption we need for the tighter bound corresponds to the movement from zero correlation to full independence. The Markov assumption enables us to use a procedure analogous to that used to correct for selection in production functions by Olley and Pakes (1994) to tighten our bound. However when we tried to implement this procedure we found that the estimates we obtained were quite unstable. There are two possible reasons. First the additional assumption could be inappropriate. Second the tighter bounds, even if appropriate, are quite sensitive to estimation error. Since our intention is to produce a bound which is both robust and can be automated for use by BLS analysts, we did not pursue this further.

Table 7: **Predicting $\eta_{t+1} - \eta_t$ for Continuing Goods in the Bimonthly Sample**.

Condition on	z		(z, η_t)		(z, η_t), r-New.	
Goods/Mean	R^2	Adj. R^2	R^2	Adj. R^2	R^2	Adj. R^2
all continuing	.15	.10	.28	.19	.30	.19
nonsticky-only	.18	.03	.45	.21	.49	.22

motivation for turning to the information in the monthly data in the next subsection.

A point on implementation. An earlier version of this paper distinguished between continuing goods whose prices changed over the sampling interval and those that did not; fully 62% did not change. This plus the fact that the fit of the regression for $\eta_{t+1} - \eta_t$ for the continuing goods *whose prices did change* was noticeably better than the fit of that same regression for *all* continuing goods, leads to a useful way of estimating $gb(z_i, \eta_{i,t}) \equiv E[\eta_{i,t+1} - \eta_{i,t} \mid z_i, \eta_{i,t}, j_{i,t} = c]$.

Let $q \in \{\Delta, s\}$ indicate whether a good's price changes ($q = \Delta$) or stays the same ($q = s$). Noting that the value of $\eta_{t+1} - \eta_t$ for the goods whose prices do not change is, by definition, $p_{i,t} - z_i\beta_{t+1} - \eta_{i,t}$, we have

$$gb(z_i, \eta_{i,t}) = \sum_{q \in \{\Delta, s\}} E[\eta_{i,t+1} - \eta_{i,t} \mid q, j_{i,t} = c, z_i, \eta_{i,t}] Pr\{q \mid j_{i,t} = c, z_i, \eta_{i,t}\} = \tag{9}$$

$$[p_{i,t} - z_i\beta_{t+1} - \eta_{i,t}] P_q\{s \mid j_{i,t} = c, z_i, \eta_{i,t}\} + E[\eta_{i,t+1} - \eta_{i,t} \mid q = \Delta, j_{i,t} = c, z_i, \eta_{i,t}][1 - P_q\{s \mid j_{i,t} = c, z_i, \eta_{i,t}\}],$$

where $P_q\{s \mid j_{i,t} = c, z_i, \eta_{i,t}\}$ is the conditional probability that $q = s$. Our estimated bounds are found by substituting nonparametric estimates of $E[\eta_{i,t+1} - \eta_{i,t} \mid q = \Delta, j_{i,t} = c, z_i, \eta_{i,t}]$ and of $P_q\{s \mid j_{i,t} = c, z_i, \eta_{i,t}\}$ for their true values in equation (9)[17].

5.2 Hedonic Bounds That Use The Monthly Data.

The monthly data contain the price changes in the first month of the two month sampling period for about half of the exiting goods in that subsample; the half that exited during the second

[17] The nonparametric estimate of $P_q\{s \mid j_{i,t} = c, z_i, \eta_{i,t}\}$ is obtained from a "probit" model with a polynomial in $(z_i, \eta_{i,t})$ as the regressors, and the nonparametric estimate of the expectation of $\eta_{i,t+1} - \eta_{i,t}$ is obtained from an OLS regerssion with a polynomial in $(z_i, \eta_{i,t})$ as regressors.

month. The average one month price relative for goods which exited in the second month was .976, a number which was substantially lower than the average one-month price relative for the goods which continued (.996). Indeed the two months price relative for the goods that continued was .985, substantially higher than the one-month relative for the goods that exited after the first month. If we were to assume that the two month price relative for the goods that exit in the second month was less than the square of their price relative in the first month, we would estimate an average annual rate of deflation of about 28% for those goods, almost three times larger than the rates from either the standard hedonic or the matched model reported in Table 1. The assumption that the two month price relative for exiting goods is bounded by the square of the one-month price fall is hard to justify, but we can use assumptions similar to those used in the last subsection to justify a bound which uses the monthly data and should be tighter than the bound of the last subsection.

The monthly analogue of Assumption 1 assumes that, conditional on the observed characteristics (our z_i) and the initial value of the unobserved characteristic ($\eta_{i,t}$), the η change of goods which continue into the second month of the period is greater than or equal to the η change of goods which exited in the first month of the sampling period. I.e. if we let $j^- = x$ ($j^+ = x$) denote the event that the good had exited by the end of the first (second) month of the sampling period

$$E[\eta_{i,t+1} - \eta_{i,t} \mid j_{i,t}^+ = x, j_{i,t}^- = c, z_i, \eta_{i,t}] \geq E[\eta_{i,t+1} - \eta_{i,t} \mid j_{i,t}^+ = x, j_{i,t}^- = x, z_i, \eta_{i,t}]. \tag{10}$$

If the good is in sample at the end of the first month, we know its price and can estimate its value of η, say $\eta_{i,t}^+$, at that time. Adding this $\eta_{i,t}^+$ to the conditioning set and recalculating the expectation on the left hand side of equation (10), that expectation becomes

$$E[\eta_{i,t+1} - \eta_{i,t} \mid j_{i,t}^+ = x, \eta_{i,t}^+, j_{i,t}^- = c, z_i, \eta_{i,t}] = E[\eta_{i,t+1} - \eta_{i,t}^+ | \jmath_{i,t}^+ = x, \eta_{i,t}^+, j_{i,t}^- = c, z_i, \eta_{i,t}] \; + \; [\eta_{i,t}^+ - \eta_{i,t}]$$

$$\leq \; E[\eta_{i,t+1} - \eta_{i,t}^+ | \jmath_{i,t}^+ = c, \eta_{i,t}^+, j_{i,t}^- = c, z_i, \eta_{i,t}] \; + \; [\eta_{i,t}^+ - \eta_{i,t}], \tag{11}$$

where the inequality again follows from our Assumption 1.

Since both $\eta_{i,t}^+$ and $E[\eta_{i,t+1} - \eta_{i,t}^+ \; | \jmath_{i,t}^+ = c, \eta_{i,t}^+, j_{i,t}^- = c, z_i, \eta_{i,t}]$ can be estimated, we can, at least in principal, construct a bound for $E[\eta_{i,t+1} - \eta_{i,t} \mid j_{i,t}^+ = x, \eta_{i,t}^+, j_{i,t}^- = c, z_i, \eta_{i,t}]$ for each value of

$(\eta^+_{i,t}, j^-_{i,t} = c, z_i, \eta_{i,t})$ observed in the monthly sample. We do not observe $\eta^+_{i,t}$ for the three quarters of the price quotes that are sampled at a two month interval, and we need a bound on the η change for the exits from that subsample; i.e. we need a bound which conditions on $(z_i, \eta_{i,t})$ but not on $\eta^+_{i,t}$. To obtain that bound from the estimates of the right hand side of equation (11), we can average over the distribution of $\eta^+_{i,t}$ conditional on $(j^+_{i,t} = x, j^-_{i,t} = c, z_i, \eta_{i,t})$ in the monthly panel. These averages are: (i) only a function of $(z_i, \eta_{i,t})$, and (ii) from equations (8) and (10) are consistent bounds for the average of $E[\eta_{i,t+1} - \eta_{i,t} \, | j_{i,t+1} = x, z_i, \eta_{i,t}]$ across all exiting observations.

There is an empirical problem with implementing this procedure. As noted the subsample with monthly quotes is about 25% of the total sample. Since about 80% of that 25% continues into the next bimonthly sampling period, we can use about 20% of the original sample to estimate $E[\eta_{i,t+1} - \eta^+_{i,t} \, | \jmath^+_{i,t} = c, \eta^+_{i,t}, j^-_{i,t} = c, z_i, \eta_{i,t}]$. However since only 10% of the monthly subsample exit in the first month of the sampling period, we only have about 2.5% of the original sample available to compute the distribution of $\eta^+_{i,t}$ conditional on $(j^+_{i,t} = x, j^-_{i,t} = c, z_i, \eta_{i,t})$. That is not enough information to estimate the needed distribution with sufficient accuracy and when we tried to do so our estimates were not robust to reasonable changes in estimation methodology and conditioning sets. As a result we move to an alternative procedure that is still based on the inequality in (11), but restricts the difference between the regression functions for $\eta_{t+1} - \eta_t$ conditional on (z, η_t) for those who exit and those who do not to be a difference in the constant term.

The steps in the alternative are as follows. We first use equation (11) to obtain a bound for $\eta_{t+1} - \eta_t$ for the late exits from the monthly sample that conditions only on (z_i, η_t). To get this bound we first regress $\eta_{i,t+1} - \eta^+_{i,t}$ for the continuing goods in the monthly sample on $(\eta^+_{i,t}, z_i, \eta_{i,t})$. Under our assumptions the predictions obtained from this equation for the $\eta_{i,t+1} - \eta_{i,t}$ of late exits are upper bounds to their values, which makes them an upper bound for the values of the early exits as well. Next we augment the sample for the regression of $\eta_{i,t+1} - \eta_{i,t}$ for the continuing goods in the bimonthly sample with these predictions and rerun the regression for predicting this change using the augmented sample and adding a dummy variable to the list of regressors which takes the value of one when the observation is a predicted value. The prediction for $\eta_{i,t+1} - \eta_{i,t}$ from these regressions are the prediction used with the dummy variable set to equal one if the good exited, and with the dummy variable set to equal zero if it did not. We implemented this procedure in several different ways, weighting the predicted observations in the $\eta_{t+1} - \eta_t$ regressions differently and using different interactions in that regression. The results were quite stable across the different specifications.

5.3 A Robustness Check.

Partly as a result of the fact that we know that our bounds are not tight, and partly to check the robustness of our procedures, we also consider an alternative bound for the price relatives of exiting goods. The alternative assumes that the change in the evaluation of the exiting good's characteristics in the period in which it exits is, on average, at least as negative as was the price fall for the same good in the period prior to it exiting. We expect exiting goods to be goods which are being obsoleted. The intuition for this bound is that the period in which exit occurs is likely to be the period when the impact of changes in market conditions on the value of the characteristics of the good that exited was particularly sharp.

There is both a conceptual and a practical problem with this bound. Conceptually it is possible for the about-to-exit price fall to be larger than the fall in the value of the goods characteristics in the exiting period. This worry would be accentuated if either; (i) there were "clearance" sales just before exit with price declines that were unlikely to be repeated over two periods, or (ii) if there were period effects in the prices of goods due to changes in overall market conditions. The practical problem is that this index mixes up information on price declines in the current period with information on price declines in a prior period. Though this may not be an important problem for some uses of the CPI, it would be important to a monetary authority that is interested in high frequency movements in price.

If we ignore these problems and accept the reasoning that leads us to expect larger price declines in the period of exit than in the period preceding exit we can simply use the price change between periods t and $t-1$ for the unobserved price change for the goods that exit between $t+1$ and t. Of course we can only do this for the goods that exited between t and $t+1$ but *were present* in period $t-1$. This is about 85% of the goods that exit between t and $t+1$. The other 15% entered between $t-1$ and t and then exited before $t+1$. For this latter group of goods we use one of our other bounds.

6 Geomean Indices for TVs: Empirical Results

Table 1 showed that a standard hedonic index that does not correct for unobserved characteristics has an average value about equal to the matched model index when a set of twenty five characteristics are used in the hedonic regression, but falls at a noticeably slower rate than the matched model index when only the ten variable regressor set is used. The matched model and the standard hedonic differ in two ways. The matched model does not account for the change in value of either the observed or the unobserved characteristics of exiting goods, but it does account for the change in value of the unobserved characteristics of the continuing goods. The standard hedonic, on the other hand, does account for the change in value of the observed characteristics of the exiting goods, but it does not account for the change in value of the unobserved characteristics of either the exiting or the continuing goods. Apparently when only ten characteristics are used the fall in value of the unobserved characteristics of the continuing goods captured by the matched model index more than compensates for the fall in value of the observed characteristics of exiting goods captured by the hedonic. This is distrubing since, as noted above, if the BLS is to use a hedonic index that can be justified in terms of a bound on the compensating variation, practical considerations will require it to use a characteristic set similar to our ten variable characteristic set.

Table 8 provides the results from computing hedonic indices based on the ten characteristic regressor set when we use the procedures developed in the last subsection to *correct for both unobserved characteristics and selection.* Panel A provides the results when it is assumed that the contribution of the unobserved characteristics to price is constant over time (the "fixed effect" assumption)[18]. As argued in the text this is unlikely to be true, but use of the fixed effect assumption does account for some of the impact of unobserved characteristics. The index based on the fixed effects assumption falls at about a 5% *faster pace* than does the matched model index (-10.6% vs. -10.11% per annum).

Panel B of the table provides the hedonic index obtained when we use the non-parametric selection correction that only requires the bimonthly data. It corrects for the change in unobserved characteristics with the results from regressing the change in residuals of the *continuing* goods on their observed characteristics and the estimate of the value of their unobserved characteristic in the initial period. This simple procedure produces an index of -11.2%, which is 11% *lower than* the matched model index. Note also that the non-parametric selection correction lowers the standard deviation of the index across months; it is now less than that of the matched model index.

Panel C adds information from the monthly data. It uses; (i) the fall in value of the unobserved characteristics of the goods that survive the first month but exit in the second to correct for the fall in value of the unobserved characteristics of the exiting goods in the first month, and (ii) the fall in value of the unobserved characteristics of the continuing goods to correct for the fall in value of the unobserved characteristics of exiting goods in the second month of our intervals. This generates a *further* 16% fall in the index, to -12.88, almost 28% lower than the matched model index.

The index that uses the monthly data is also more variant across months then the other indices. The variance across months could either be caused by estimation error, or by true variance in the value of the index. If the increased standard deviation of the index which uses the smaller monthly samples is due to estimation error, and not due to real variance in index values across months that the courser bimonthly sample does not pick up, it is undesirable. However we should keep in mind that the CPI itself is an average over many component indices, and the averaging process should do away with much of the worry about estimation variance in the component indices.

[18] This index is constructed by first regressing the change in the logs of the price levels (i.e. the log of the price relatives) of the continuing goods on their observed characteristics, and then constructing predicted price relatives for both exiting and continuing goods from these regression coefficients.

Table 8: **Alternative Monthly Indices for TV**[1]**.**

Index Calculated	matched model	hedonic
Panel A: Fixed Effects (in logs) Selection Correction.		
mean	-10.11	-10.62
standard deviation	5.35	5.79
% l.t. mm	n.r.	.70
Panel B: Non-Parametric Selection Model Using Only Bimonthly Data.		
mean	-10.11	-11.17
standard deviation	5.35	5.01
%l.t.mm	n.r.	.80
Panel C: Non-Parametric Selection Model Using Also Monthly Data.		
mean	-10.11	-12.88
standard deviation	5.35	8.21
%l.t.mm	n.r.	.83

[1] See also the notes to table 1. n.r.=not relevant, % l.t. mm = percentage less than matched model, standard deviation is the standard deviation of the index across months. All panels use the S10 regressor set.

Tables 1 and 8 taken together make two points. First if we are to base a hedonic adjustment on hedonic regressions with variable sets that enable a different hedonic regression to be used in each period and we do not take account of the value of unobserved characteristics in components where they are important, we are likely to significantly understate the rate of price decline. In the case of the TV component index for our period we would understate it more than a matched model procedure would understate it. Second if we do account for the change in value of unobserved characteristics then the hedonic correction generates price falls that are substantially larger than that obtained from the matched model index. Moreover as we try corrections that theory tells us should get closer to the true price change, we obtain progressively larger falls in the index.

Table 9 asks whether these results are consistent with the estimates obtained from assuming that the rate of fall in the evaluation of the characteristics of exiting goods in the year they exit is at least as large as the price fall those goods experienced in the period before they exit (the "a-exit" price falls). The indices in this table average; (i) hedonic price relatives for continuing goods with, (ii) the price relative in the period prior to the period in which goods exit for exiting goods for which there was an *observed* price relative in the period prior to exit, with (iii) estimates of price relatives from one of our other procedures for exiting goods which did not have an observed price change in the period prior to exit.

Both estimates of the rate of deflation in Table 9 lie between the estimates in panels B and C of Table 8 (and they have lower standard deviations across months than do the estimates based on the monthly data in that table). That is when we use the observed price falls in the period before the good exits as a bound on the fall in the period in which the good exits we obtain rates of price change which are larger (in absolute value) then the indices which only take account of the change in the values of the unobserved characteristics of continuing goods, but lower than the estimates that take account of the actual change in value of the unobserved characteristics in the first month of the goods that exit in the second month of the two month sampling period. This is exactly what we would expect if *all* of our assumptions were correct; the assumptions underlying the estimates in both Table 9 and Table 8. Partly as a result, we come back to indices based on pre-exit data in section 8 of this paper.

Table 9: **Robustness Analysis.**

A-Exit Price Changes If They Exist and Our Correction Otherwise		
Index Calculated	matched model	hedonic
Panel A: Correction Using Bimonthly Data Otherwise.		
mean	-10.11	-12.15
standard deviation	5.35	5.13
% l.t. mm	n.r.	.83
Panel B: Correction Using Also Monthly Data Otherwise.		
mean	-10.11	-12.27
standard deviation	5.35	5.91
% l.t. mm	n.r.	.93

Notes: See the notes to Tables 1 and 8. The average (over all months) fractions of goods that are continuing, exiting-with-a-previous-relative, and exiting-without-a-relative are, respectively, $(.793, .171, .036)$.

7 Out of Sample Predictions and Alternative Sample Sizes.

The results presented in the last section were obtained after considerable experimentation, experimentation which would not be possible if the index would have had to be constructed between the time the BLS receives its data and its deadline for publishing the index. They were also obtained for a component index with a relatively large number of price quotes. This section does two things. First it examines whether our index for TV's could be used in "production mode" by applying it, without change, to a later period of TV data (our "out of sample" prediction). Second we ask whether an analogous index could be used for digital cameras, a product with a relatively low number of price quotes. Table 10 compares the three data sets.

Table 10: **Comparing Data Sets.**

Variable (averaged over sample periods)	TV: May 2000 Jan. 2003	TV: Feb. 2005 Nov. 2006	Digital Cameras May 2007/April 2009
1. # Of Obs.	236.5	215.4	57.7
2. Exit Rate	22%	31%	36%
3. Ave. Adjusted R^2. (Number of Regressors)	.95 (10)	.91 (9)	.86 (8)
4. # Continuing (monthly sample)	50.1	39.4	9.4

7.1 Out of Sample Predictions.

The analysis of the later period TV sample makes no change to the algorithms used in the earlier parts of this paper, so once the data from the later period was inputted we could reproduce all of our tables in seconds. We did change the observed characteristic set used in the hedonic regressions to reflect changes in the TV's marketed that we were quite sure the BLS's commodity analyst for TV's would have picked up on.[19]

The later period data replicated virtually all of the qualitative features of the data from the

[19]The most important change in the characteristic set was replacing the flat screen indicator for a "flat panel" indicator, and interacting that with the screen size variable. Almost all TV's in the later period were flat screen TV's. The flat panel category includes both liquid crystal display and plasma TV's. We also eliminated the brand dummies as the two brands were a smaller (and decreasing) portion of the market in the later period.

earlier period,[20] though magnitudes did change in notable ways. In particular; (i) exit rates were almost 50% *higher* in the later period (a fact that may well be indicative of a more rapid rate of obsolescence), and (ii) the adjusted R^2 from the hedonic regressions were lower (.90 vs. .95; this is probably a result of doing the analysis of the later period without the extensive experimentation with variables and data cleaning we did with the earlier sample).

Table 11 shows that, as one might expect, the more rapid rate of turnover in the later period goes hand in hand with a more rapid rate of price decline, no matter which index one uses to measure price declines. Both the rate of deflation, and the standard deviation in the rate of deflation, almost double in absolute value in the later period. Most of the difference between periods is picked up by the matched model index; i.e. in the later period goods prices were falling at a more rapid rate in the periods *before they exited*. Despite this difference between the periods, the indices we suggest each provide a correction to the matched model index of about the same magnitude in the later as they did in the earlier period (the later period corrections are actually slightly larger in absolute value than they were in the earlier period, but noticeably smaller as a ratio of the matched model index). Equally important from a methodological point of view, the order of the various indices in the later period is precisely the same as it was in the earlier period which, recall, is exactly the order our reasoning predicts. We conclude that our indices seem to generate corrections which abide by our priors when applied directly to new data *without any prior experimentation*; even when the data differs in ways that have significant effects on the magnitude of the indices.

7.2 Digital Cameras.

Digital cameras are not a separate component index, but rather are a part of the "photographic equiptment" index. However hedonic adjustments for cameras would presumably require a camera specific hedonic function, and 97% of the camera price quotes in the last three periods of our sample are digital[21]. The digital camera sample is only about one fourth the size of the TV samples and

[20] The sole exception was that in the later data the rate of fall of price relatives for recently introduced goods was slightly less than that for all continuing goods, while in the earlier data it was slightly more.

[21] About 65% all price quotes in the component index are cameras. Since cameras are not a separate component index we do not have the weights that the BLS would use to construct a camera index and as a result our camera

Table 11: **Comparing Indices In Different Data Sets.**

Index Calculated	TV May 2000 Jan. 2003	TV Feb. 2005 Nov. 2006	Cameras May 2007 April 2009
1. matched model	-10.11	-19.29	-29.89
Standard deviation	5.35	9.31	16.42
Hedonic With Adjustment for Unobservables.			
2. Bimonthly Adj.	-11.17	-20.44	-30.46
Standard deviation	5.01	10.95	18.73
Adj. to mm	-1.06	-.47	-.47
3. Monthly Adj.	-12.88	-23.20	n.a.
Standard deviation	8.21	11.15	n.a.
Adj. to mm	-2.71	-3.91	n.a.
Pre-Exit with Hedonic Adj. When Not Available.			
4. Bimonthly Adj.	-12.15	-22.30	-30.84
Standard deviation	5.13	8.80	13.45
Adj. to mm	-2.04	-3.78	-.95
5. Monthly Adj.	-12.27	-22.69	n.a.
Standard deviation	5.91	9.34	n.a.
Adj. to mm	-2.17	-3.40	n.a.

Notes: The indices in rows 2 and 3 are calculated as are the indices in panels B and C in table 9, and the indices in rows 4 and 5 are calculated as are the indices in panels A and B in table 10. "Standard deviation" is the standard deviation of the index across months and "Adj. to mm" is the difference between the calculated index and the matched model index.

has, on average, only 9.4 continuing price quotes and 2.8 late exits in its monthly subsample. This is not enough to compute indices which require monthly hedonic regressions.

As can be seen from Table 11 digital cameras have the highest exit rate (36%) and lowest average adjusted hedonic R^2 (86%) of any of our samples. Again the higher exit rate seems indicative of a faster rate of obsolescence which is largely picked up in the price falls of goods in the period *prior to* exit. The average rate of decline of the matched model index is almost 30%.

The average value of the bimonthly adjusted hedonic index presented in the table varied somewhat with the specification of the regression of $\eta_{t+1} - \eta_t$ on η_t and z. The results in the table were obtained by constraining that regression to be linear, as this produced both the most conser-

indices are unweighted. As noted our TV indices are weighted. We also computed unweighted TV indices and the differences between the alternative unweighted indices were virtually identical to those we obtained using weights; though the levels of all unweighted indices were, on average, about 5% lower.

vative index (i.e. the index with the slowest rate of price decline), and the least variable index[22]. This generated a price fall which, on average, is about .5 percentage points more then that of the matched model index; an increment that matches that of the later TV sample. The index which uses the pre-exit price falls when they are available and the bimonthly hedonic adjustment when it does not falls at a slower rate than the analogous incremental fall in the two TV samples, but at a rate which is still larger than that of the hedonic with the bimonthly adjustment.

Our hedonic adjustements to the matched model index for the digital camera sample are smaller than those for the TV samples. This is partly due to the lack of enough observations to use the adjustments that require the monthly sample. Nonetheless the adjustments we obtain from the hedonic procedures we can apply to the digital camera sample have precisely the same qualitative properties as did the indices we computed from the two TV samples; they all produce faster rates of price decline than the matched model index, and they are ordered in the predicted order.

8 The CPI Mimic and a Pure Pre-Exit Index.

The BLS no longer uses a pure matched model index for either their phototgraphic or TV component index. Rather they use a "hybrid" index constructed as a weighted average of matched model price relatives when they are available, and one of the four constructs for price relatives described below when it is not[23]. We now consider how the BLS's procedures for treating goods which are not available in both periods effects price indices.

In the BLS's current indices these goods are: (i) dropped from the sample before averaging if a close substitute can not be found, (ii) given a a price relative formed as the ratio of the price of the close substitute to that of the original good when a "very close" substitute is found (very close in terms of observed charateristics), (iii) given a price relative constructed as the ratio of a

[22]When instead we chose the specification for the $\eta_{t+1} - \eta_t$ regression which maximized the adjusted R^2 we obtained an average rate of price decline of -39.52 for our adjusted hedonic. This is a rate of price fall which is about 30% larger than that from the matched model. The standard deviation of that index over periods was about 45% larger than that of the index reported in the text.

[23]The BLS procedure also differs from ours in that it attempts to obtain price relatives for the goods sampled in the comparison rather than for the goods sampled in the base period; a procedure which need not produce a bound on the compensating variation. For more on the BLS's hybrids, see Pakes, 2003.

hedonically adjusted price of a substitute good to that of the original good if a substitute is found but its observed characteristics differ in important ways from those of the original good, or (iv) given a price relative equal to the average of the price relatives for the goods that exit *and* obtain a price relative if a substitute is found but the analyst is not confortable using either (ii) or (iii) above. The hedonic adjustment in (iii) is obtained by multiplying the difference in characteristics between the original and substitute goods by the coefficients of those characteristics obtained from a hedonic regression which is only updated infrequently (usually annually).

Procedure (i) generates an exit bias identical to the bias that motivates hedonic procedures. (ii) and (iii) make a correction for the re-evaluation of observed, but not unobserved, characteristics, and the correction for observed characteristics need not satisfy the conditions required for a bound on compensating variation. The properties of (iv) depends on those of (ii) and (iii).

We found the price relatives *actually used* by the CPI for the exits from each of our samples, and used them to construct inidices which "mimic" the CPI. Our CPI mimic uses matched model price relatives for continuing goods and, to the extent possible, the procedure actually used in the CPI for the goods that exit. The relative importance of the procedures used to obtain price relatives for our CPI mimic differed by sample. In the TV samples; (i) about 26% of the goods were dropped, (ii) 18% were adjusted using substitutes that were deemed not to need hedonic adjustments, (iii) 46% were adjusted with hedonic adjustments, and (iv) the remaining 10% were given the average price relative of the goods that exited but had price relatives. The camera sampled differed markedly in this respect; hedonics *were not used at all* for cameras. In that sample; (i) 20% of the exiting goods were simply dropped out, (ii) 53% obtained price relatives by using an unadjusted substitute directly, and (iii) 27% were allocated the average of those that were adjusted with the substitute. So in the camera sample goods were either dropped out, or adjusted with a procedure *which does not attempt* to account for differences in *either observed or unobserved* characteristics between the substitute and the original good.

Panel A of Table 12 provides the average difference between CPI Mimic and a matched model index for our three samples, its standard deviation across sample periods, and the "t-value" of the average difference. For comparison we provide the same three numbers for a "pure pre-exit" index in Panel B. The pure pre exit is a hybrid which uses; (i) the matched model price relatives for the continuing goods, (ii) the pre-exit price changes for the exiting goods for which there is a pre-exit price fall observation, and (iii) the average of the price changes of the exiting goods that did have a pre-exit price fall for the exiting goods for which we do not observe a price fall in the period before they exit. The advantage of the pure pre-exit index is that it is easier to use than *any* of the indices considered thus far; its disadvantages are that it does not necessarily form a bound for the compensating variation and it is partly based on price movements in the period prior to the period of interest. The first disadvanatge is shared by the CPI mimic, the second is not but the CPI mimic; (i) has an exiting good bias and (ii) requires an *ad hoc* allocation of goods into groups.

The numbers in Table 12 are striking. For the early TV sample the CPI mimic does adjust the matched model index downward, but the average adjustment is not significant, and its magnitude is only about half that provided by the pure pre-exit index (which is less than the adjustments provided by three of the four indices in table 11). For the later TV sample the CPI mimic adjusts the matched model *upward*; i.e. it produces a *slower* rate of price fall than the matched model (though the average difference is not significant). In contrast the pure pre-exit index adjusts the price index downward, as expected, and the difference is both statistically and economically significant. Finally for cameras, the component index which does *not use hedonics at all*, the CPI mimic adjusts the index *upward* by 33%. This is a rate of price fall which is both statistically and economically *significantly slower* than that of the matched model index. Again the pure pre-exit index adjusts downward as expected, though the average difference is not statistically dignificant.

At the very least Table 12 throws into question the usefullness of indices that do not account for differences in the characteristics of goods that exit. Moreover the figures for cameras indicate that continuing goods that are very close in their observed characteristics to exiting goods need not be close in their unobserved characteristics. We noted earlier that use of pre-exit price relatives has several problems, but it is *easier* to compute and seems to perform far *better* than the CPI mimic.

Table 12: **The CPI Mimic and A Pure Pre-Exit Index.**

Index Calculated	TV May 2000 Jan. 2003	TV Feb. 2005 Nov. 2006	Cameras May 2007 April 2009
Panel A: CPI Mimic Minus Matched Model.			
Average	-.79	1.03	9.96
Standard deviation	2.80	4.0	1.77
t-value	-1.58	1.10	5.65
Panel B: Pure Pre-Exit Minus Matched Model.			
Average	-1.56	-2.41	-1.09
Standard deviation	2.15	4.03	8.28
t-value	-4.04	-2.54	-.62.

Notes: Since we wanted all rows of this table to be comparable, and we do not have regional weights for the cameras sample, the differences are all unweighted average differences of price relatives. The standard error is the standard error of the average difference across months. The CPI mimic is an index which treats each price relative in the manner the actual CPI treats that particular price relative (see the text). The pre-exit index uses the price change between the two periods immediately preceeding exit for the unobserved price fall of the exiting good price when the pre-exit price fall is available, and assigns the average of the pre-exit price changes that do exit for the cases where the pre-exit price fall does not exist.

9 Conclusions and Caveats.

Standard hedonic procedures correct for the market's re-evaluation of the observed characteristics of exiting goods, but do not correct for the re-evaluation of the unmeasured characteristics of either continuing or exiting goods. Matched model indices correct for the market's re-evaluation of the unmeasured characteristics of continuing goods but do not correct for the change in value of either the observed or unobserved characteristics of exiting goods. As a result when there is substantial turnover *and* important unmeasured characteristics both indices are likely to be inadequate, and either index can be larger than the other.

Unmeasured characteristics can arise either because there are no sharp cardinal measures of important characteristics of the good available, or because the measures that are available can not be used without an extensive data cleaning procedure. Extensive data cleaning is inconsistent with the combination of the time constraints of the BLS and the need to compute new hedonic regressions every period in order to insure that the resultant index does in fact abide by the Konus-Laspeyres bound to the compensating variations.

This paper provides several ways of constructing hedonic-like indices that at least partially correct for both the selection bias induced by exit *and* for the contribution of unmeasured characteristics. Along the way we explain why the “biases” in both matched model and in prior hedonic indices seem to differ; (i) across component indices and (ii) with the time interval between successive price observations.

Our empirical results show that, at least in our examples, the indices we suggest can be produced in a timely way. Moreover they produced values which were consistent with the economic arguments that lead to them and were noticeably lower than both the matched model and standard hedonic indices. We also mimic current BLS procedures for correcting for the selection induced by exit and find that they do not do nearly as well.

We want to conclude by noting that the problems of selection and of unmeasured characteristics are not the only problems with the component indices that underlie the CPI. A number of other problems remain and just as it was more important to account for unmeasured characteristics in the TV than in the personal computer component index, the importance of these other problems is likely to vary across component indices.

In particular none of the indices make any adjustment for the inframarginal rents that accrue to consumers that would have bought a new good at the highest price at which the good is observed. Also none of the component indices take account of changes in either the environment, or in the availability of related goods which impact on the utility of the goods in the particular component index of interest. For example the fall in the price of clothing as the season that the clothing was designed for ends is partially a result of the fact that the utility the consumer derives from that clothing changes when the season ends. Finally the sampling scheme used to construct the component indices attempts to measure changes in the price of a (sales weighted) average purchase from the commodity group in question. In fact different consumers purchase at different points of time. At least in markets for goods which are somewhat durable and in which there are seasonal or intra model-year patterns in prices, we might think it more in line with the compensating variation rational for price indices to compute an average of price changes over the intervals at which consumers' purchase the good, rather than an average over the purchases in a given interval (for an application of this idea to automobiles see Ana Aizcorbe et. al, in process).[24]

References

Aizcorbe, A., B. Bridgman, and J. Nalewaik. "The Implications of Heterogenous Buyers for Measuring Quality Change", *in process*, Bureau of Economic Analysis.

Bajari, P. and L. Benkard. "Demand Estimation with Heterogenous Consumers and Unobserved Product Characteristics: A Hedonic Approach "(2005) *Journal of Political Economy.*

Bajari, P. and L. Benkard, "Hedonic Price Indexes with Unobserved Product Characteristics, with an Application to PC's" *Journal of Business and Economic Statistics*, 2005a.

Erickson, T., and A. Pakes: "An Experimental Component Index for the CPI: From Annual Computer Data to Monthly Data on Other Goods" 2009, National Bureau of Economic Research Discussion Paper ....

Johnson, D., S. Reed and K. Stewart. "Price Measurement in the United States: A Decade after the Boskin Report", *The Monthly Labor Review*, 2006, May, pp. 10-19. *Econometrica 7*, 1924, January, p. 10-29.

Moulton B., T. Lafleur, and K. Moses, *Research On Improving Quality Adjustment in the CPI: The Case of Televisions*, Proceeding of the fourth meeting of the International Working Group on Price Indices, U.s. Department of Labor, April 1998, Bureau of Labor Statistics.

[24]This is related to the more general issue of constructing and using household or demographic specific indices, a topic beyond the scope of this paper.

Olley, S. and A. Pakes. "The Dynamics of Productivity in the Telecommunication Equipment Industry", *Econometrica*, 1994, pp.

Pakes, A. "A Reconsideration of Hedonic Price Indices with an Application to PC's." *American Economic Review* 93. 2003. pp. 1578-1596.

Appendix: Characteristic Data.

The next table defines the characteristics we use and gives their average values for different subsets of the data. All variables are 0-1 dummy variables except screen size and the number of DVD player inputs.

Table 13: **Average Characteristic Vectors for Subsets of TVs.**

characteristic	continue	exit	about to exit	enter
screen size (inches)	29.22	30.74	30.84	30.91
picture in picture	0.28	0.32	0.33	0.34
flat screen (not flat panel)	0.096	0.092	0.095	0.136
Projection TV (rear only)	0.148	0.181	0.188	0.185
High-definition ready (no tuner)	0.069	0.070	0.076	0.098
A prominent "quality" brand	0.232	0.202	0.205	0.209
A prominent "value" brand	0.142	0.145	0.149	0.141
1 extra video input	0.282	0.253	0.253	0.240
2 extra video inputs	0.288	0.310	0.304	0.273
3 extra video inputs	0.268	0.283	0.287	0.333
4 extra video inputs	0.046	0.047	0.049	0.069
No. DVD player inputs	0.442	0.481	0.491	0.613
A 3D comb filter	0.148	0.171	0.179	0.192
wide screen (16:9 aspect ratio)	0.023	0.031	0.035	0.037
mtx surround sound	0.394	0.410	0.409	0.427
store 1	0.159	0.155	0.153	0.161
store 2	0.205	0.192	0.191	0.206
store 3	0.118	0.114	0.112	0.112
store 4	0.099	0.063	0.065	0.069
New York City	0.105	0.112	0.115	0.107
Chicago	0.058	0.064	0.068	0.059
LA	0.105	0.092	0.095	0.108

Notes: 1. In the regressions the first characteristic is log-screensize; it is unlogged here. 2. Table is the average of the mean characteristic vectors in period t-2 for each of 29 bimonthly intervals t-2 to t: 15 from the odd-month subsample and 14 from the even-month subsample. "continue" indicates all TVs present in both t-2 and t. "exit" are those present in t-2 but not in t. "about to exit" are present in t-2 and t but not in t+2. "enter" gives the period t characteristic for TVs present in t but not present in t-2.