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

Timothy Erickson
Bureau of Labor Statistics
and
Ariel Pakes*
Harvard University and NBER

September 10, 2008

Abstract

Until recently the Consumer Price Index consisted solely of "matched model" component indexes. 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 indexes. 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%.

*We would like to thank Teague Ruder and Aylin Kumcu for excellent research assistance, Paul Liegey for his patience in explaining the BLS's algorithm for constructing the TV component index to us, Steve Berry, Dale Jorgenson, Emi Nakamura and Jack Triplett for helpful comments, and the BLS NSF and Toulouse Network for for support. All errors and omissions are those of the authors and not of any of these institutions.

1 Introduction

This paper provides hedonic techniques which enable the construction of better price indexes. Along the way we explain why the "biases" in both matched model and in prior hedonic indices seem to differ; (i) across component indexes and (ii) with the time interval between successive price observations.

Pakes (2003) used a model of a differentiated-product market as a framework for clarifying the role of hedonic regressions in the construction of price indexes. The regressions do not identify either utility or cost parameters. Nevertheless under the conditions supplied in that article they can be used to bound the transfer needed to compensate consumers for changes in their choice sets (for the "compensating variation"). The bound is not tight because it does not account for either the inframarginal rents to consumers who would have purchased the good at the highest observed price, or for the substitution possibilities caused by the changes in prices over time. However the bound *is* typically tighter than that given by the matched-model index. This is because it takes partial account of the selection bias in the matched model indices caused by the exit of 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 indexes partially correct for this bias by using the price prediction from a hedonic regression for the exiting period prices of the goods that exit. Since the relationship between the prices of goods and their 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. Subject to these requirements, any sufficiently rich functional form can be used.

The relationship between the hedonic prediction for the price of exiting goods, and that implicit in the matched model index, clarifies the difference between the indexes. The matched-model index implicitly imputes its own value, which is an average of the values for all continuing goods, as the predicted price relative for every exiting good. The hedonic prediction weights more heavily the predicted prices of continuing goods that have characteristics more similar to those of the exiting good. So a matched-model index takes the index weight intended for an exiting good and redistributes it to the continuing goods proportional to their index weights, whereas the hedonic index implicitly redistributes an exit's weight more towards those continuing goods with similar characteristics. In Pakes' (2003) computer application the use of the hedonic rather than the matched model prediction changed the index rather dramatically. This because the value of the observed tuples of characteristics that were similar to those of exiting computers fell rather dramatically – largely in response to the entry of newer machines that obsoleted them.

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 modernized its data gathering procedures by providing their data gatherers with hand held computers and instructing them to download their data nightly onto a central BLS data management system has changed what the BLS can do. It now is possible for the BLS to implement a version of our hedonic index. Before doing so, of course, the index would have to be shown to satisfy all the rigorous monthly production requirements of the CPI.

However when standard hedonic procedures were tried on some of the BLS's component groups most of the resultant indexes were not much different from the relevant matched model indexes (see the results reported in Table 5 of the survey by Johnson, et. al, 2006). Take our TV example. The sample gathered for the TV index has 20% turnover over the two month sampling interval (an almost identical rate to that in computers), and as we show below there is ample evidence indicating that the goods that exit have prices that are falling disproportionately. 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, for more detail on the current CPI TV component index, and Pakes, 2003, for a comparison of the CPI's hedonic adjustment and that obtained from the more standard hedonic procedure used here).

Table 1 presents the index we obtained using hedonic predictions from a linear in logs hedonic regression that was done separately in every period and a set of twenty four characteristics that are similar to the set of characteristics used by the BLS, and compares it to a matched model index based on the same data (see below for details on the data). The hedonic generates an index which is 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 these nine characteristics 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. 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 Appendix 2). 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 indexes for those features. As a result in the TV market, and we suspect in many other markets, selection is partly based

Table 1: **Matched Model and Standard Hedonic Indices.**[2]

Index Calculated	matched model	hedonic[1]
hedonic uses S24[1]	-10.11	-10.20
s.d. (across months)	5.35	7.53
S24 % l.t. mm[1]		.50
hedonic uses S10	-10.11	-8.82
s.d. (across months)	5.35	7.05
S10[1] % l.t. mm		.40

1.*Definitions*: S24 refers to a 24 regressor specification derived from the 22 variables in Table 12 in Appendix 2. Screen size is logged, and two additional regressors are obtained by multiplying log-screensize times itself and the dummy indicating projection TV. S10 refers to a 10 regressor specification consisting of the first 7 variables in Table 12 (with screensize logged) plus the two second-order regressors from S24 and a dummy indicating when an observation is from either NYC, Chicago, or Los Angeles, the only cities sampled on a monthly basis. The dependent variable in both specifications is the log of price.
2.*Entries:* 30 monthly indexes are computed for each of the matched-model and two hedonic specifications, spanning the 31-month interval from June 2000 to Dec 2002. 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).

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: (i) at least partially account for the contribution of unobserved characteristics; (ii) maintain the Konus-Laspeyre's bound on compensating variation; (iii) are robust to the properties of data sets in ways we will make precise; and (iv) can be implemented within the BLS's time constraints. We note that though our assumptions insure that our predictions for the exiting goods price are an upper bound to their expected price, the bounds we suggest are not tight. On the other hand they are relatively easy to implement within the BLS's time constraints and seem to not be sensitive to estimation error. Moreover despite the non-tightness of our bounds, the fact that our hedonic techniques do partially control for the role of unmeasured characteristics generates indexes which falls at a rate over twenty per cent faster than does the standard hedonic index.

We begin with a description of the characteristics of our primary sample (which is an extract from the data the BLS uses to compute their TV component index). This section of the paper also provides evidence that the prices of exiting goods are falling at a faster rate then those of continuing goods and illustrates the features of the data that underlie why our new hedonic techniques are likely to do a better job than standard hedonic techniques. We then turn briefly to the formulae for alternative price indices conditional on estimates of price changes.

This brings us to section 4 which begins by presenting evidence on the importance of accounting for selection on unobserved characteristics in obtaining price predictions. Section 4 then goes on to develop hedonic formulae for price predictions which partially account for the role of unobserved characteristics. We provide two of these. The second uses more of the variables in the data than the first and the extra variables should enable us to produce a tighter index; but since those variables are only available for a subsample the second is also likely to have larger estimation error. Section 4 concludes with a third alternative for price prediction; one that has intuitive appeal but which uses assumptions which are harder to justify (at least as a period-by-period index). 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 5 provides the empirical results from using the various hedonic procedures to obtain price indexes for our primary sample period (from May 2000 to January 2003). They indicate that standard hedonics overstate inflation by over 20%. Moreover they are striking in their consistency with the theoretical arguments given in the paper.

We wondered whether our prior experimentation with the data influenced the nature of these results, and wanted to know how direct application of our procedures would fare in alternative environments. So we drew a new sample (from February 2005 to November 2006) and applied our procedures directly, i.e. without any prior experimentation, to this data. Section 6 of the paper notes that the new sample is different in important ways from our primary sample, and then reports the price indices obtained from our "out of sample" experiment. They reinforce our earlier conclusions.

There is a short concluding section which contains a summary and a reminder that the issues of selection and unobserved characteristics dealt with in this paper are not the only issues with the component indexes that underlie the CPI and that different issues are likely to be more important in constructing different component indexes.

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

Our data 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 CPI industry analyst in the procedure currently used to construct hedonic adjustments for the TV component index[1]. The average monthly sample contains prices and characteristics from 234 observations, and 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.[2]

[1] A "cleaned characteristics" subset of each period's July and August data was prepared by the CPI industry 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.

[2] 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

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 next 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 exits are temporary, while 2.9% of entering TVs are returning from temporary absence.[3]

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

included hedonic indexes 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.

[3]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 a stock-outs, while permanent exit is 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.

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[4].

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. We show below that what evidence exists lends strong support to this belief. 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. Note that 61.55% of all price relatives relatives equal 1; that is there are a lot of "sticky" prices. Since we use this fact below, the table provides summary statistics for the subsample of non-sticky prices as well as for the overall sample.

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 group of goods. 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). Second the about to exit goods have a significantly lower fraction of sticky prices (the standard errors for these fractions vary from .006 to .015). Moreover if we look just among non-sticky prices the absolute difference between the mean price relative for goods about to exit and the other group is even more pronounced. That is among prices that do change, the prices of the goods that are about to exit fall substantially more than a randomly chosen price change[5]. If goods that are about to exit have prices relatives that behave more similarly to the prices of

[4]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.

[5]The about to exit goods also have higher price variance than other goods, though most (though not all) of this increased variance is because they have a larger fraction of non-sticky prices.

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 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). Similarly the sticky price rate in the first month for the goods that exit in the second month is .627 which is lower than the two month sticky price rate for the goods that continue (.6569). 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). Sill this finding has interesting 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 initially prices fall, 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 indexes 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%.[6] So going from a two month to an annual interval

[6] Unlike the rest of the indexes in this paper these use equally weighted price relatives (instead of expenditure relatives) and cover only a twenty four month period.

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.
Fraction of Subsample With Relatives						
Equal 1 (or "sticky")	.6155	.5390	.6203	.6380	-.0990	-.0176
Greater than 1	.1166	.1097	.1142	.1213	n.r.	n.r.
Less than 1	.2679	.3513	.2655	.2407	n.r.	n.r.
# of obs.	5320	1167	1335	2818	n.r.	n.r.
Among Price Relatives Not Equal to 1 (i.e. not "sticky").						
mean	.9622	.9460	.9608	.9682	-.0222	-.0074
(s.d. of mean)	(.0024)	(.0056)	(.0049)	(.0034)	(.0063)	(.0058)
cross-section s.d.	.1039	.1083	.0920	.1024	.0059	-.0104
# of obs.	2017	549	514	1067	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)
sticky price rate	.6569	.5776	.6270	.3931
# 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

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

Specification	Constrained OLS		Minimum Distance	
	exit	new	exit	new
1. S0 (Odd)	.107	.128	.074	.121
	(2.63)	(3.23)	(1.88)	(3.12)
2. S0 (Even)	.114	.121	.090	.117
	(2.89)	(3.05)	(2.33)	(3.00)

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 and minimum distance estimates differ in that the latter weights with the covariance matrix across periods.

entering goods as new goods typically enter at the high quality end of spectrum. 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. 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 TV market do not turnover nearly as much.

We will see that though our characteristics can differentiate between high and low quality TV's, they have more difficulty with distinguishing 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 projection, but we do not have a good measure of the improvements that have occurred in sharpness of their 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, ...).

3 Inputs For the Indexes.

We require; (i) a formula which enables us to calculate the price index from a given set of price relatives, and (ii) the hedonic regressions that underlie our hedonic price predictions. We start with the index formula.

3.1 Index Formulae.

We begin as simply as possible and use indexes that are linear in the logs of price relatives. This makes the indexes linear in the regression error from the logarithmic hedonic regressions we and others have used, and this in turn makes the relationship of our results to the underlying data transparent[7].

Letting t index our two month sampling period, all indexes we present are versions of

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

where q denotes a quote, w_{qt} is period-t weight, $\tilde{y}_{qt}$ is an actual or imputed log-relative, and S_{t-1} is a subset of all quotes available for period $t-1$[8]. This is the log of a geometric mean index. The weights $w_{q,t-1}$ are obtained by dividing each period-$(t-1)$ regional TV expenditure-share equally among all the quotes for that region and then renormalizing them so that $\sum_{q \in S_{t-1}} w_{q,t-1} = 1$. The regional expenditure shares are estimates from the CPI data base.[9]

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

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

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

[7] We intend to come back to more complex indexes that work directly with this regression error at a later date, as they have a larger role to play in other indexes. In particular to construct the Laspeyre's index we need to exponentiate the logs and hence exponentiate the hedonic regression error. Since the Laspeyre's index is the only index that has an interpretation in terms of a bound on compensating variation, there are good reasons for thinking the indices that deal directly with the regression error might be important.

[8] This mimics the BLS's bimonthly sampling procedure. Their monthly indexes are derived from these by a linear splicing procedure.

[9] Past values of the CPI subindex for TVs are used in making the estimate for any period t. We take these estimates as *given;* we do not prepare our own estimates based on past values of any of our research indexes.

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 indexes average the 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 $t-1$[10].

3.2 Hedonic Regressions.

The results presented here are based on a linear regression model for the (log) price levels of goods in a given period.[11] Let Z_t be the $n \times K$ matrix of characteristics of those TVs for which prices were collected in period t and p_t be the corresponding $n \times 1$ vector of log prices. Then a typical 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 noted above there are no restrictions on these coefficients and there is no necessary relationship between the coefficient vectors estimated in different periods.

We fit this regression to every month in each bimonthly sample, using each of three different sets of regressors for Z, all of which include a column of ones. The three sets of regressors, to be denoted by $S5$, $S10$, and $S24$ are

- $S5$: 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[12],

[10] An early version of this paper also computed the hybrid indexes 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 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. The actual values of the hybrids we calculated for our sampling period are, however, available on request from the authors.

[11] 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.

[12] As noted in Pakes, 2003, the hedonic regression can differ with the characteristics of the population in which

- $S10$: 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 the notes to Table A1 at the end of the paper.

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 of the CPI index for TVs uses a different but similarly lengthy list of regressors as our S24, most of which 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 very high R^2's. It is not unusual to get high R^2's in hedonic regressions on differentiated product markets, indeed it is a major reason for the increased use of characteristic models in demand estimation. However these R^2's are higher than usual, which probably attests to the quality of the BLS data.

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 $S24$.[13] 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. So perhaps the more striking fact illustrated by table 4 is just how well we do in predicting price; the residual from either the S24 or the S10 regression accounts for only a relatively small fraction

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

[13]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.

of the price variance (and there may be sources of measurement error).

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

Table gives summary statistics from log-price regressions run on each of the 35 months from March 2000 to January 2003.

Indeed the variance accounted for by the residual is small enough to call into question the need to worry about omitted variables when correcting for selection in matched model indexes. It is true however that; the price variance is large, exit is concentrated in particular parts of the product space, and the variance accounted for by the residual may well be be a larger portion of overall variance in that part of the characteristic space where exit is concentrated. This worry about omitted characteristics 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 the advanced features and the fact that the turnover is concentrated in the high end of the product price spectrum 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.

4 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[14]. This prediction problem, however, gets more complicated when there are characteristics of the goods that consumers value but Econometricians do not observe

[14]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).

(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, then 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 is only done with observations on the continuing goods and the new entrants). However if unobserved characteristics are important determinants of whether a good exits, then simple economic arguments should lead us to believe that; (i) the regression function for goods that exit is different from that for continuing goods, and (ii) that the prediction for the price of the good that exits obtained from this regression function will be systematically biased in a particular direction.

For simplicity assume 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.

If we let $j = x$ denote exiting goods, $j = n$ denote new entrants, and $j = c$ denote continuing 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.

To see this we need a model for which goods exit. Temporarily assume that a product exits if its price falls below its marginal cost, denoted $m(z,\eta)$, and that $\partial m(z,\eta)/\partial\eta < 1$ (everywhere). Our normalization implies that price increases one to one with η, so we are assuming that price increases more than marginal cost when unobserved quality increases (this rationalizes the extra sunk costs usually required to develop higher quality products). Then there is a function $\underline{m}(z)$, such that the good exits if and only if $\eta \leq \underline{m}(z) - z\beta$, and

$$E[\eta|z, j = x] = E[\eta|\eta \leq \underline{m}(z) - z\beta] \leq E[\eta|z].$$

In particular when the good's observed characteristics lead to a small $\underline{m}(z) - z\beta$ then the good will continue even if it has a low value of η, while if $\underline{m}(z) - z\beta$ is large the good will only continue if it has a relatively large value of η. So the distribution of η conditional on z (its support, its mean,....) will be different for the continuing than for the exiting goods.

To see whether such logic leads to a significant differences in the relationship between z and η for exiting, continuing and newly entered goods in our data set, 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. The results are presented in Table 5. They clearly reject the null that the new and exiting good interactions are all zero.

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

Test	$j = x$; F-test	$j = n$; F-test	$j = x$; Wald-test	$j = n$; Wald-test
Fraction of Months Significant At Different α Levels				
$\alpha = .01$	.11	.07	.57	.61
$\alpha = .05$	.25	.21	.79	.71
$\alpha = .10$	.39	.36	.79	.79

F-test assumes homoscedastic variance-covariance,
Wald-test allows for heteroscedastic consistent covariance matrix.

We can go one step further here. If unobserved characteristics are an important determinant of whether a good exits and marginal costs are relatively constant over time, then we would expect there to be 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 $z\beta$ increase disproportionately will continue even if their η value decreases disproportionately, while goods whose $z\beta$ decreases 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.

The data contains at least two more pieces of evidence on the difference between the unobserved characteristics of the exiting and continuing goods. First we can compare the estimates of the unobserved characteristics (i.e. the residuals from the hedonic regression) of exiting goods to those of continuing goods in the period prior to exit. Second we can compare the change in the estimate of the $\eta's$ of the exiting goods over the period immediately preceding the period in which they exit to the change in estimate of the η's of the continuing goods over that period.

The change in η results also throw light on the appropriateness of an alternate 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[15]. Note the arbitrary difference in the way the fixed effect assumption treats the

[15]This because the fixed effect assumption guarantees both that; (i) the contribution of the unobserved characteristics to price does not change over time, and (ii) that the regression of the log price relatives provides an unbiased estimate of the changes in prices caused by the market's re-evaluation of observed characteristics.

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. 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. So though the values of the unobserved characteristic of the exiting goods in the year prior to exit are only marginally lower than that of the continuing goods, the values of η of the exiting goods are falling at a dramatically faster pace than those of the 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 rather 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). That is the assumption that the unobserved characteristic's contribution to price is constant over time seems inconsistent with the data. 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, have larger values then do 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

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.

the *continuing* goods – a point we come back to in a more formal way below.

The results in this section make it doubtful that we can get an adequate correction for the selection bias in exiting goods by simply re-evaluating the *observed* characteristics of those goods in the period in which they exit. So we now consider alternative selection corrections.

4.1 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 in market conditions. The additive disturbance is now not the contribution of unobserved characteristics per se but rather the residual from regressing 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 introduces method for predicting $p_{i,t+1} - p_{i,t}$ 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 for the goods that exited during the second month of that 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, at the very least, it can be used to check the robustness of our results.

4.1.1 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.

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 then their marginal cost, then Assumption 1 would be equivalent to the assumption that price increases in η more than marginal costs does.[16] Indeed it is the increase in the difference between price and marginal cost as quality increases that is often viewed as the justification for the sunk cost of developing higher quality goods.

Assumption 1 implies that

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

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

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.

We have already shown that the η of about to exit goods fall at a faster pace than other continuing goods (see Table 6). What equation 8 shows us is that Assumption 1 is sufficient to extend this inequality to goods that do exit. To see whether the bound is likely to help 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

[16] There are at least two situations when the conditions underlying Assumption 1 might not hold, though there is an industry analyst for each component index who is able to amend the index in cases where either condition might overturn our bounding argument. First the cost of the unobserved characteristics may increase enough relative to the price it can command to warrant an exit decision. Second a merger or an acquisition could lead to a situation where a good exits even though Assumption 1 is satisfied. Note also that the logic underlying Assumption 1 is based on their being 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 the values of those indices to the exit decision. This would require a significantly more complex set of assumptions and computational algorithms, see for e.g. Bajari and Benkard (2005a), and generate an index whose relationship to the data is less transparent.

the conditional expectations of $\eta_{t+1} - \eta_t$ for continuing goods will help bound the change in η for the exiting goods[17].

Of course the bound in (8) may not be very "tight". In particular the last subsection showed that the unobservables for exiting 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 (see 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.

This source of the 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.[18] Appendix 1 shows that with the Markov assumption a procedure analogous to that used to correct for selection in production functions by Olley and Pakes (1994) can be used to tighten our bound (this without restricting the Markov process in any way). 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, as we explain in appendix 2, 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 ignore the Olley-Pakes bound below. This does however leaves us with a bound which we know is not tight, and this is part of the motivation for turning to the information in the monthly data in the next subsection.

[17]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).

[18]We have done the same analysis using local linear kernel regressions. Since, at least in the limit, these give us non-parametric estimates of conditional expectations, the additional assumption then moves us from mean independence to full independence. The results from using the local linear kernels were virtually identical to those described below.

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

One final point on implementation. When we included η_t the fit of the regression for $\eta_{t+1} - \eta_t$ for the continuing goods *whose prices change* was noticeably better than the fit of that same regression for*all* continuing goods. The value of $\eta_{t+1} - \eta_t$ for the goods whose prices do not change is, by definition, $p_{i,t} - z_i\beta_{t+1} - \eta_{i,t}$. So instead of estimating $gb(\cdot) = E[\eta_{i,t+1} - \eta_{i,t} \mid z_i, \eta_{i,t}, j_{i,t} = c]$ directly, we make use of the information in the price change variable to obtain a more precise estimate of the the conditional expectation of $\eta_{t+1} - \eta_t$. That is we let $q \in \{\Delta, s\}$ indicate whether a good's price changes ($q = \Delta$) or stays the same ($q = s$) and rewrite our bound as

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

$$= \Big(E[\eta_{i,t+1} - \eta_{i,t} \mid q = \Delta, j_{i,t} = c, z_i, \eta_{i,t}] - [p_t - z_i\beta_{t+1} - \eta_{i,t}] \Big) Pr\{q = \Delta \mid j_{i,t} = c, z_i, \eta_{i,t}\} + [p_t - z_i\beta_{t+1} - \eta_{i,t}].$$

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 $Pr\{q = \Delta \mid j_{i,t} = c, z_i, \eta_{i,t}\}$ for their true values in equation (9).

4.1.2 Hedonic Bounds That Use The Information in 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 month. Table 8 is based on all the data in the monthly sample and compares the behavior of prices during the first month of the two month sampling period of goods that exit in the second month to that of goods who continue over the entire two months. It is clear from the table that the first month price fall and fraction non-sticky for goods that exit in the second month are quite a bit larger than the same figures for the goods that continue.

Table 8: **Summary Statistics from The Monthly Sample.**

Data From The First Month of the Sampling Period.		
	goods exiting in second month	continuing goods
2. Fraction With Sticky Price	.584	.756
3. Average Price Relative	.9756	.9957
4. Av. Rel. when $q^- = \Delta$	.933	.974

Indeed comparing these figures to those in table 2 we see that the first month price fall for the goods that exit (.976) is larger than the two month price fall for the goods that continue (.984). Moreover the latter figure is quite close to the square of the one month price relative for the continuing goods ($.993^2 = .986$). If we were to assume that the two month price relative for the goods that exit in the first 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 can not be justified without assumptions which are harder to justify than those used in this paper. On the other hand we can use assumptions very similar to those used in the last subsection to justify a bound which uses the monthly data and should be tighter than the bound which uses only bimonthly data provided in the last subsection.

In particular we use the monthly analogue of Assumption 1 to relate the observed first month price change of goods that exited in the second month to the unobserved first month price changes of goods that exited in the first month. I.e. the assumption 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. Formally 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 we assume that

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

and then develop bounds for the left hand side of this equation.

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

$$E[\eta_{i,t+1} - \eta_{i,t} \mid j_{i,t}^{+} = x, \eta_{i,t}^{+}, j_{i,t}^{-} = c, z_i, \eta_{i,t}] \tag{11}$$

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

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

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 only conditions on $(z_i, \eta_{i,t})$. 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 equation (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 $\eta_{i,t+1} - \eta_{i,t}$ then becomes the prediction from this equation 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.

4.1.3 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 and a disproportionate fraction those sales were caught in our about to exit sampling period instead of in the period the good exits, or (ii) if there were period effects in the prices of goods due to, say, cost changes (or for some other reason). 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 might 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 $t+1$ to 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.

5 Geomean Indexes for TVs: Empirical Results

Table 1 showed that a standard hedonic index (one 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. Moreover, as noted in the introduction, if the BLS is to use a hedonic index that can be justified in terms of a bound on the compensating variation, it must use a bound based on a characteristic set similar to our ten variable characteristic set.

Table 9 provides the results from computing hedonic indexes 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)[19]. 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, at -10.6% compared to -10.11% for the matched model.

Panel B of the table provides the hedonic index obtained when we use our first suggestion for a non-parametric selection correction, the one based only on the bimonthly data. It corrects for the change in unobserved characteristics with the results from regressing the 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 correction procedures produces an index of -11.2%, which is 11% *lower than* the matched model index. Note also that use of this non-parametric

[19]This index is constructed by first regressing the change in the log 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.

selection correction also 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 on the price change of the goods that exit in the first month of the two month sampling period to the selection correction procedure. 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 the sampling interval. This results in a *further* 16% fall in the index. The index is now 12.88, almost 28% lower than the matched model index.

Table 9: **Alternative Monthly Indexes 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		.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		.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.c.	.83

[1] See also the notes to table 1. n.c.=not calculated, % 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.

The index that uses the monthly data is also more variant across months then the other indexes. 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 indexes.

Tables 1 and 9 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 10 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 indexes 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 10 lie between the estimates in panels B and C of Table 9. 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 indexes 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 both the Table 9 and the Table 10 estimates). We take this as strong evidence in support of our assumptions.

Table 10: **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		.83
Panel B: Correction Using Also Monthly Data Otherwise.		
mean	-10.11	-12.27
standard deviation	5.35	5.91
% l.t. mm		.93

Notes: Also see the notes to Tables 1 and 9. The average (over all months) fractions of goods that are continuing, exiting-with-a-previous-relative, and exiting-without-a-relative are, respectively, $(.793, .171, .036)$.

The indexes in table 10 have two attractive features. First substituting the price changes in the year prior to exit for the missing data on the price changes of the exiting goods would provide an easy way for the BLS to correct for most of the missing data. Moreover, if the results in Table 10 are indicative, this simple substitution would tighten up the matched model bound significantly. Second use of this procedure decreases the standard deviations across months noticeably; the table 10 standard deviations are comparable to the lowest standard errors in table 9. On the other hand recall that there is reason to worry whether the assumptions needed for these indexes to actually bound the compensating variation will always be satisfied, and the index that results will be a mix of the price falls over two periods, rather than over the period of interest.

6 An "Out of Sample" Application.

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. To see whether our index could be used in "production mode" we took a sample from a later period, from February 2005 to November 6, and re-ran all the tables and indexes in this paper with the new data. We made no change to the procedures used in the computation, but we did change the observed characteristic set used in the hedonic regression to reflect changes in the TV's marketed that we were quite sure the BLS's industry analyst for TV's would have picked up on.[20] As it turns out the later period was a period in which the market for TV's evolved differently than the earlier period, so the exercise also provides a test of the robustness our procedures to different underlying market conditions.

The later period data replicated virtually all of the qualitative features of the data from the earlier period that were discussed in the first four sections of this paper,[21] though magnitudes did change in notable ways. In particular the later period was a period where there was *more rapid* rates of turnover in the TV's marketed, with average exit and entry rates over the two month sampling period of 37.6% and 36.1% respectively, about 50% larger than in the earlier period. The turnover was still largely at the high end. Just as in the earlier period, the prices of goods in the period before their exit were both significantly higher and falling at a faster pace then the prices of a randomly chosen good. So the higher turnover rates may well be indicative of a more rapid rate of obsolescence of advanced features in the later than in the earlier period. Further since the average R^2 for the hedonic regressions in the later period was less than that for the earlier period (.90 vs about .95), the unobserved qualities of these advanced features may have also played a larger role

[20]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.

[21]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.

in the later period.[22]

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 indexes 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 indexes in the later period is precisely the same as it was in the earlier period which, recall, is exactly the order our theoretical reasoning predicts. We conclude that our indexes seem to do generate corrections which abide by our priors when applied directly to new data *without any prior experimentation*, and this even when the new data differs from the old in ways that have significant effects on the magnitude of the indexes.

7 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 indexes correct for the markets 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.

[22]This also points to the desirability of adding variables to the observed characteristic set in the later period. Perhaps the most obvious candidate is an indicator which would let us keep separate dummies for plasma and liquid crystal display. These are now lumped together in the flat panel dummy.

Table 11: **Comparison Between Time Periods.**

Index Calculated	May 2000 January 2003	February 2005 November 2006
1. matched model	-10.11	-19.29
Standard deviation	5.35	9.31
Hedonic With Adjustment for Unobservables.		
2. Bimonthly Adj.	-11.17	-20.44
Standard deviation	5.01	10.95
Adj. to mm	1.06	1.15
3. Monthly Adj.	-12.88	-23.20
Standard deviation	8.21	11.15
Adj. to mm	2.71	3.91
Pre-Exit with Hedonic Adj. When Not Available.		
4. Bimonthly Adj.	-12.15	-22.30
Standard deviation	5.13	8.80
Adj. to mm	2.04	2.68
5. Monthly Adj.	-12.27	-22.69
Standard deviation	5.91	9.34
Adj. to mm	2.17	3.40

Notes: The indexes in rows 2 and 3 are calculated as are the indexes in panels B and C in table 9, and the indexes in rows 4 and 5 are calculated as are the indexes 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.

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 indexes that at least partially correct for both the selection bias induced by exit *and* for the contribution of unmeasured characteristics. We have shown both that, at least in our example, the indexes 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 want to conclude by noting that the problems of selection and of unmeasured characteristics are not the only problems with the component indexes 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 indexes.

We already noted that none of the indexes 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 falls when the season ends. Finally the sampling scheme used to construct the component indexes 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).[23]

References

Andrews, D. "Asymptotics for Semiparametric Econometric Models via Stochastic Equicontinuity", *Econometrica* 62, 1994, pp.43-72.

Aizcorbe, A., B. Bridgman, and J. Nalewaik. "The Implications of Heterogenous Buyers for Mea-

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

suring Quality Change", *in process*, Bureau of Economic Analysis.

Berry, S., Levinsohn, J., and A. Pakes. "Automobile Prices in Market Equilibrium," *Econometrica,* 63, pp. 841-890.

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.

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

Erickson, T. "On Incorrect Signs in Hedonic Regressions", *mimeo,* 2004, Bureau of Labor Statistics.

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.

Konus, A. "The Problem of the True Cost-of-Living Index" translated in *Econometrica 7*, 1924, January, p. 10-29.

Lancaster, K. *Consumer Demand: A New Approach,* Columbia University Press, 1971, New York, New York.

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.

Newey, W., 1994, "The Asymptotic Variance of Semiparametric Estimators" *Econometrica*, 62, pp. 1249-82.

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.

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

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

Appendix 1 "Tighter" Hedonic Bounds.

This appendix assumes, in addition to Assumption 1, that the $\{\eta_t\}$ in equation (8) follow a Markov process. If for simplicity we assume that process is first order, the formal statement of the additional assumption would be that the stochastic process generating $\{\eta_t\}$ is given by the family of probability distributions

$$\mathbf{F}_\eta = \{F(\ \eta_{t+1} \mid \eta_t\ );\ \eta_t \in \mathcal{R}\}. \tag{12}$$

This plus the exit rule in Assumption 1 implies that the expectation of $\eta_{i,t+1} - \eta_{i,t}$ conditional on survival is given by

$$E[\eta_{i,t+1} - \eta_{i,t} \mid z_i, \eta_{i,t}, j_{i,t} = c] = \frac{\int_{\underline{\eta}_{t+1}(z_i)} \Big[\eta_{i,t+1} - \eta_{i,t}\Big] dF(\eta_{i,t+1} \mid \eta_{i,t})}{1 - F\Big(\underline{\eta}_{t+1}(z_i) \mid \eta_{i,t}\Big)} \equiv g\Big(\underline{\eta}_{t+1}(z_i), \eta_{i,t}\Big). \tag{13}$$

We have an estimate of $\eta_{i,t}$ from the hedonic regression that uses *all* of the data. However we need an estimate of $\underline{\eta}_{t+1}(z_i)$.

As in Olley and Pakes (1994) the estimate of $\underline{\eta}_{t+1}(z_i)$ is obtained from the exit equation which is given by

$$Pr\{\eta_{i,t+1} \geq \underline{\eta}_{t+1}(z_i) \mid \eta_{i,t}\} = 1 - F\Big(\underline{\eta}_{t+1}(z_i) \mid \eta_{i,t}\Big) \equiv \mathcal{F}(\underline{\eta}_{t+1}(z_i), \eta_{i,t}).$$

The function $\mathcal{F}(\cdot)$ maps values of $(\underline{\eta}_{t+1}(z_i), \eta_{i,t})$ into the interval $(0,1)$ and, provided $F(\cdot|\eta_{i,t})$ has a density which is positive everywhere, is monotone decreasing in $\underline{\eta}_{t+1}(z_i)$ for any given value of $\eta_{i,t}$. This implies that for any $\eta_{i,t}$ there is an inverse which provides $\underline{\eta}_{t+1}(z_i)$ as a function of the value of $\mathcal{F}(\cdot)$ and $\eta_{i,t}$. Call that inverse $\mathcal{F}_\eta^{-1}$, so that

$$\underline{\eta}_{t+1}(z_i) = \mathcal{F}_\eta^{-1}\Big[\mathcal{F}(\underline{\eta}_{t+1}(z_i), \eta_{i,t})\Big],$$

and substitute it into equation (13) to obtain

$$E[\eta_{i,t+1} - \eta_t \mid z_i, \eta_{i,t}, j_{i,t} = c] = g(\mathcal{F}_\eta^{-1}\Big[\mathcal{F}(\underline{\eta}_{t+1}(z_i), \eta_{i,t})\Big], \eta_{i,t}) \equiv h(\mathcal{F}_{i,t}, \eta_{i,t}), \tag{14}$$

where $\mathcal{F}_{i,t} \equiv \mathcal{F}(\underline{\eta}_{t+1}(z_i), \eta_{i,t})$.

Both $\mathcal{F}_{i,t}$ and $\eta_{i,t}$ can be estimated, and hence, if we temporarily ignore estimation error, can be treated as observable. So, after conditioning on the fact that the good is in sample in period t+1, we can substitute equation (14) into equation (7) to obtain

$$E[p_{i,t+1} - p_{i,t} \mid z_i, \eta_{i,t}, j_{i,t} = c] = z_i(\beta_{t+1} - \beta_t) + h(\mathcal{F}_{i,t}, \eta_{i,t}). \tag{15}$$

This equation can be taken to data to estimate both the function $h(\cdot)$, and $(\beta_{t+1} - \beta_t)$.[24]

We now move to the prediction for *exiting* goods conditional on both observed and unobserved characteristics. First note that

$$0 = E[\eta_{i,t+1} - \eta_{i,t} \mid z_i] \equiv$$

$$\mathcal{F}_{i,t}E[\eta_{i,t+1} - \eta_{i,t} \mid \eta_{i,t+1} \geq \underline{\eta}_{t+1}(z_i), z_i] + [1 - \mathcal{F}_{i,t}]E[\eta_{i,t+1} - \eta_{i,t} \mid \eta_{i,t+1} \leq \underline{\eta}_{t+1}(z_i), z_i].$$

Consequently

$$E[\eta_{i,t+1} - \eta_{i,t} \mid \eta_{i,t+1} \leq \underline{\eta}_{t+1}(z_i), z_i] = -\frac{\mathcal{F}_{i,t}E[\eta_{i,t+1} - \eta_{i,t} \mid \eta_{i,t+1} \geq \underline{\eta}_{t+1}(z_i), z_i]}{[1 - \mathcal{F}_{i,t}]} \equiv -\frac{\mathcal{F}_{i,t}h(\mathcal{F}_{i,t}, \eta_{i,t})}{[1 - \mathcal{F}_{i,t}]}.$$

So the hedonic prediction for the price relatives of exiting goods conditional on both observed and unobserved characteristics could be obtained by

$$E[p_{i,t+1} - p_{i,t} \mid z_i, \eta_{i,t}, j_{i,t} = x] = z_i(\beta_{t+1} - \beta_t) - \frac{\mathcal{F}_{i,t}h(\mathcal{F}_{i,t}, \eta_{i,t})}{[1 - \mathcal{F}_{i,t}]}. \tag{16}$$

We found that the estimates we obtained in this way to be quite imprecise and to vary a great deal with the way one estimates the non-parametric function. There are two possible reasons. First

[24]Formally the estimator is a two-stage semiparametric estimator. The non-parametric components are the functions $\mathcal{F}(\cdot)$ and $h(\cdot)$ and the parametric components are β_{t+1} and β_t. For the econometric details of semiparametric techniques see Andrews (1994) and Newey (1994) and the literature they cite.

the independence assumption in equation (12) might be inappropriate. Second in the empirical work $\mathcal{F}_{i,t}$ must be estimated and if its true value is near one even a small amount of estimation error will cause very imprecise estimates of the truncated expectation.

Appendix 2: 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 12: **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.