

“Quality” Adjustment for the High Frequency Price Indexes Used in the CPI: An Experimental Index for TV’s.

Timothy Erickson
Bureau of Labor Statistics

Ariel Pakes*
Harvard University and NBER

December 9, 2006

Abstract

The Consumer Price Index is a weighted average of primarily “matched model” component indexes. The latter are constructed by data gatherers who visit stores and compare prices of goods with the same set characteristics over successive periods. In a previous paper Pakes (2003) suggested that this procedure was subject to a selection bias. Goods that were not on the shelves in the second period, and hence whose price comparisons were discarded, were disproportionately goods which were obsoleted over the period, and consequently represented goods whose prices were falling. Computers (both hardware and software) were the example par excellence of this phenomena, and using computer hardware data Pakes demonstrated that the selection effect could both be very large, and could be corrected using hedonic techniques.

Subsequent work on BLS data for other component indices *did not* show large differences between hedonic and matched model indices; a fact which lessened the incentives to use hedonic procedures in constructing the CPI. This paper explores why. In particular we look carefully at the component index for TV’s and show that differences between the TV and computer markets, together with the fact that the BLS data are high frequency, make it necessary to use a more general hedonic correction for TV’s. The computer market is special in having both well defined cardinal measures of the major product characteristics and exiting goods with relatively low values for them. In markets where such measures are absent and where turnover can be at the high quality end we need to allow for selection on unmeasured, as well as measured, characteristics. Also in high frequency data we need to correct for differential “sticky price” rates among different goods. We develop a hedonic selection correction which accounts for these phenomena and show that when applied to TVs it yields results which are similar to those that were found for computers. Moreover our index can be constructed within the time constraints of the BLS and hence is a “real time” alternative to the current CPI production process.

*We would like to thank the Bureau of Labor Statistics and the Toulouse Network for Information Technology for the support which allowed us to undertake this project. All errors and omissions are those of the authors and not of either of these institutions.

1 Introduction

This paper reports recent progress on improving hedonic quality-adjustment procedures for price indexes. It extends the approach demonstrated in Pakes (2003) for annual data on desktop computers to the monthly and bimonthly data on televisions actually used by the CPI. We show that this requires an extension of the hedonic adjustment techniques used in the earlier paper. Along the way we explain why the “biases” in both matched model indices, and in standard 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 use of hedonic regressions. It showed that such regressions do not identify either utility or cost parameters, but can nevertheless be used to bound the compensating variation which would compensate consumers for changes in the choice set. The bound is typically tighter than that given by the matched-model index 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 removes price changes from the left tail of the distribution of price changes, and this causes an upward bias in the estimate of the average price increase.

We can partially correct for this bias by using hedonic predictions for the exiting year price of the good that exits. In order to reflect the current relationship between prices and characteristics, the regressions used to predict the exiting good’s price must include all relevant characteristics, be updated every period, and have no cross-period constraints. Subject to these requirements, any sufficiently rich functional form can be used.

It is useful to consider the relationship between the hedonic prediction method for the price of exiting goods, and that implicit in the matched model index. 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 closer to those of the exiting good. That is 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 our indexes implicitly redistribute 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. The reason was that 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 which 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's data gatherers now record the data they gather on hand held computers whose contents are downloaded nightly onto a central BLS data management system has, at least in principle, changed this situation. We could now substitute a hedonic for a matched model index and still meet the BLS's production schedule.

However as we shall show below, and as is reported by the National Academy of Sciences (2004), when standard hedonic procedures are tried on the BLS data the resultant indexes are not much different than more traditional matched model indexes. This paper explores the reasons for this phenomena in the context of the component index for TV's. Though there is little difference between the matched model and the standard hedonic index for TV's, on average over 20% of the TV sample turns over during the sampling interval, and there is ample evidence indicating that the goods that are exit have prices that are falling disproportionately.¹

The problem is not that there is no selection bias, rather it is that standard hedonic procedures do not correct this bias. This for two reasons. First the TV market is different from the computer

¹This is despite the fact that some of the attrition from the sample is *temporary*. That is some the goods that are not available in the current period are expected to be back (and in many cases are back) on the shelf in a future period. A good that is temporarily off the shelf may not be off the shelf because it has been obsoleted but rather because of a stock-out caused by unexpectedly high demand. Temporary exits are, however, less than 3% of exits in our sample.

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 3). 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. As a result in the TV market, and we suspect in many other markets, selection is partly based on characteristics the analysts can not condition on, i.e. on what an econometrician would call “unobservables”.

Standard hedonic predictions for the prices of goods which exit do not account for the price differences generated by characteristics the analyst does not condition on. Moreover we will show that because the dynamics of the market insure that the hedonic regression function itself changes over time, the contribution of the unobserved characteristics of a given good to price changes over time. As a result though use of predictions based on hedonic regressions for (log) price differences, regressions which would difference out the impact of unmeasured characteristics whose contributions to price do not change over time, does move the resultant price index in the predicted direction, it only partially controls for the impact of the unobservables (and accordingly generates an index which only improves slightly on standard hedonic techniques).

The second reason why standard hedonic procedures do not adequately control for selection has to do with the relatively high frequency of BLS data (where prices are usually resampled at two month intervals)². At this frequency prices are often “sticky”, i.e. they do not change between successive readings. In fact on average only 40% of the prices change over a two-month interval. If all prices were equally sticky we would not have to adjust our predictions for this fact. However, as we shall see below, prices of goods that are about to exit are, perhaps not surprisingly given that they are in a changing part of the market, systematically less sticky than most.

This paper develops hedonic indexes which partially account for these two phenomena. The

²About 75% of the BLS TV sample comes from alternating bimonthly subsamples. The bimonthly samples we use are obtained by apportioning the CPI monthly subsample of prices between the CPI's two bimonthly samples. The annual rates implied by monthly indexes differ slightly due to the splicing operations used to construct a monthly index from bimonthly indexes.

procedure we adopt abides by the general conditions discussed in Pakes (2003), but differs in several important ways. We already noted that we allow both for selection which is partially based on unobservables, and for differential sticky price rates. Our methodology also enables the analyst to use only a small number of "easy-to-clean" product characteristics in the hedonic regression. Given the recent computerization of the data gathering process, labor-intensive cleaning of a large number of characteristics is now the main reason BLS currently runs only one TV hedonic regression a year. As noted above a hedonic index based on hedonic prediction functions which are not updated every period do not provide a bound to the compensating variation. Since the combination of computerization and our method should enable the BLS to compute a hedonic regression almost instantaneously, it should enable them to construct a hedonic index which bounds compensating variation with in the BLS's time constraint.

Our hedonic makes three changes to standard hedonic procedures. First our predictions for the prices of goods which leave the sample makes an explicit correction for the expected value of the disturbance from the hedonic regression for those goods. The prediction provides an upper bound to the expected contribution of the disturbance, and hence is not "tight". It still, however, makes a substantial difference to the index. Second, we write the fitted regression as a function of the rate of price "stickiness" so that we can condition our predictions on estimates of this rate for different subsets of TVs.

Third, we use local-linear nonparametric regression to insure that our price predictions are based on the price movements of products that have similar characteristics to the characteristics of the product that exited. The price variation in the TV market is enormous; from \$66 to over \$10,000, reflecting the large differences in products that the BLS includes in this commodity group. The entry and exit of particular TVs in this market, as in many markets, tends to disproportionately influence, and be disproportionately influenced by, prices of close competitors. Use of the local-linear nonparametric estimator insures that the hedonic predictions for one good are not overly sensitive to goods which are in very different parts of the product space (indeed movements in

price of very different types of goods are not highly correlated). The local linear methods also help reduce the number of characteristics needed for adequate predictions, a fact which, as noted, is important to enabling use of hedonic techniques in the BLS's production setting.

The empirical results show that our procedures make a substantial difference to the TV component index. The matched model's estimate of the average annual rate of price change in the TV market is -10.0%. If we apply a standard linear in log's hedonic prediction technique for the exiting goods prices using a set of twenty four characteristics we get a number which is almost the same (we get -10.1%). However were the BLS to use a characteristic set this large the amount of data cleaning needed before they could produce a hedonic index would imply that they could not produce a new hedonic regression every period. So we also consider a nine variable characteristic set which does not require extensive cleaning and hence could be used in a production setting. When we use standard hedonic procedures with these nine characteristics we get a hedonic index falls at a *slower pace* than the matched model index.

The remainder of the analysis was done with the nine variables characteristic set. First we moved to predictions based on hedonic regressions for the log of price differences, regressions which difference out the impacts of unobservables whose effects on price are constant over time. The index then jumps back up to -10.7%, now surpassing the matched model index by a noticeable amount. When we implement our correction for unobserved characteristics but make no adjustment for the differential sticky price rate of exiting goods, the index jumps further to -11.1%.

We now move to the adjustment for sticky prices. The sticky price rate to use in this adjustment is not obvious. We expect the sticky price rates for the goods that exit to be less than the sticky price rate for goods in the year prior to the year they exit (which is about 54%), as the exit year is likely to be a year where large changes occur in the valuations of products with characteristics similar to those of the exiting goods. So if we use the sticky price rate of goods just prior to exit it should give us a conservative index. When we use that sticky price rate the index jumps to -11.5%. If we assume lower sticky price rates for exiting goods the index increases monotonically going

all the way to -15% when we assume that none of the goods that exit are goods with sticky prices.

Though the interval from -11.5% to 15% is quite large, even the smallest number in it implies we can correct the rate of inflation estimated by the matched model index by about 15%. This is a number which is large enough to have a significant impact on the budget deficit, but just how large an impact is on our research “to do” list. We would however like to provide both a robustness check on it, and some guidance as to whether we can limit the range further.

Our hedonic indexes differ systematically from matched-model index because the characteristics (both observed and unobserved) of exiting TVs, and their sticky price rates, do not look like those of a random sample from the distribution of the characteristics of continuing TVs. The exiting goods characteristics *do*, however, look much more similar to the characteristics of a subset of the continuing TVs: those that will exit in the next period (see Appendix 3). So an alternative index we could construct is to use the price changes for the about to exit TV’s for the price change of TV’s that do exit. Though this is not possible in a production setting (since in a production setting we would not know which good would exit in the following period when the index is constructed), an alternative that could be used and has a similar justification, is to use the price falls of the goods that do exit in the period prior to exiting as an estimate of the price falls of those goods in the year they do exit. Though there may be shocks to the price surface in a period which makes the price change in the period preceding the period the good exits different than in the period the good in does exit, we would expect that on average the price changes in the period before a good exits to be similar, though somewhat smaller (in absolute value), than in the period they do exit. “Similar” because actual data gathering dates vary over the two month interval for different goods in an arbitrary way, and smaller because the exiting period was selected as the period when the price drop warranted exit (though we can not prove either assertion).

Substituting the prior year’s price change of the good for the price change in the year they exit leads to an average annual index of -12%.³ This is similar to, but slightly larger (in absolute

³A small fraction of the goods that exit do not have a price change in the period prior to exit, because they are also recent entrants (see below). For these goods we use the prediction developed in the paper for the price change.

value) then, our own hedonic index. Indeed it gives an index with about the same value as one constructed from a sticky price rate of about .45 (instead of .54) per cent. So one might sharpen the interval to -12% to -15% implying adjustments of twenty to fifty percent to the index – though if one believed in the similarity of pre-exit to exit price changes one would stay with the lower part of the interval.

The remainder of this paper is organized as follows. Section 2 describes our data, Section 3 describes the formulas for our research indexes given a set of price prediction equations. Section 4 explains and discusses our methods for predicting price. Section 5 reports our empirical results. Section 6 compares our approach to a research index that mimics the hedonic method currently used in the CPI. Section 7 describes the next steps of our research plan.

2 Background: Characteristics of the Data.

CPI price quotes from March 2000 to January 2003 are used. A "cleaned characteristics" subset of each year'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.⁴ These range from \$66 to \$10,079. The average monthly sample contains 234 prices and has mean, median, minimum, and maximum prices equal to \$725, \$366, \$81, and \$7836 respectively.

CPI collects data at 2-month intervals from an odd-numbered month subsample and an even-numbered subsample, so we can only measure attrition rates over a two month period⁵. On average, 22.5% of the TVs present in any period $t - 2$ are not present in t , with 19.7% of these being permanent exits. Similarly, 24.0% of TVs in t are not present in $t - 2$, with 17.0% of these 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

⁴Comparing, 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.

⁵CPI also collects 1-month interval data, which accounts for 26.8% of the entire data set. For experimental purposes we have converted the 1-month data into 2-month data and merged it with the larger 2-month data sets.

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.⁶

Price relatives for different subsets of the data play a key role in this paper. For any two periods $t - 2$ and t there are TVs in our sample in both periods, and for these we have relatives. Between July 2000 and November 2002 the average number of $t - 2$ to t relatives was 183.45. Some of their characteristics are listed in Table 1.⁷ The average number of $t - 2$ to t relatives that were for TVs that would exit before $t + 2$ was 40.21 (which is 22.37% of those relatives). We say these TVs are "about to exit" and denote them by "a-exit." We expect the behavior of their price relatives to be more like that of goods which are similar to the goods that exit the sample between t and $t + 2$ than that of the randomly drawn price relative. As a result we will use this subsample for clues as to the unobserved price relatives for goods that are similar to those that exit. The average number of $t - 2$ to t relatives that were for TVs that *entered* in $t - 2$ was 46.03 (which is 25.52% of the relatives.) These TVs have "recently entered" from the standpoint of t , and we denote them by "r-new." Table 1 provides some summary statistics on the price relatives for these subsets. We note that 61.55% of all $t - 2$ to t relatives equal 1; that is there are a lot of "sticky" prices. Thus we provide summary statistics for the subsample of non-sticky prices as well as for the overall sample.

A few points about the data on the goods that are about to exit are important for what follows. First the about to exit goods have a faster rate of price decline from $t - 2$ to t than continuing goods, indeed they decline at about twice the rate. Moreover this difference is highly significant (with a t -ratio of about six). Second the about to exit goods have a lower fraction of sticky prices. Moreover if we look just among non-sticky prices the absolute difference between the mean price relative for goods about to exit and the goods that continue is even more pronounced. That is among prices that do change, the prices of the goods that exit fall substantially more than a randomly chosen

⁶These numbers 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.

⁷This 29-month span is derived from our full 35-month data set as follows: One of the 18 odd months is used up to make relatives. Two more of the 18 are used up to determine those relatives that will exit before $t + 2$ and those that entered in $t - 2$. Three of the 17 even months are used up in the same way.

price change. 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 price. These numbers reinforce the belief that, by throwing out the goods that exit, matched model procedures overestimate inflation.

Interestingly goods that are recently introduced also have price relatives that on average fall at a faster pace than continuing goods, though the difference is not nearly as striking (it is only 1/4 to 1/5 the differential rate of decrease of goods that are on the verge of exit, and the difference with continuing goods is not statistically significant). Still 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. Moreover the tendency for new goods prices to fall is more pronounced among new goods whose prices do change.

Finally we note that the results in Table 1 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 from t to $t + 2$ and from $t + 2$ to $t + 4$ to the matched model index from t to $t + 4$. The latter does not contain the price changes of the goods that are on the “verge of exit” in period t , and are “recently new” in period $t + 2$. Both these subgroups of goods have prices that fall at a faster rate than a randomly drawn continuing good. That is the t to $t + 4$ index misses two groups of prices changes whose prices are falling disproportionately. So matched model indices for longer interval data will provide an upwardly biased rate of inflation, and we should expect the bias to increase the longer the interval. This explains the deterioration in performance for longer interval matched model indices highlighted in the introduction.

The next table provides a characterization 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; (i) log prices on a constant and two dummies, one for the goods that just entered

Table 1: **Price Relatives.**

Variable	Full Sample.	v-exit	r-new	contin.	exit-cont	new-cont
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.
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

and one for goods that are about to exit, and (ii) log price on a constant, the two dummies, and year specific coefficients on the nine variables that we will eventually use in our hedonic regressions (more detail on these below). The first regression provides an indication of the level of prices, and hence the “type” of goods, that just entered and/or are about to exit. The second regression asks whether our characteristics are sufficient to control for the price-type of goods that just enter and/or are about to exit. The regressions are done differently for odd and even numbered years as the BLS samples different cities in those years.

There are two points that this table makes clearly. First both the newly entering goods and the about to exit goods have prices that are *higher* than those of continuing goods. This is not surprising for newly entering goods as it simply means that 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 improvements have been at the high end, but in TV’s the exitors that are displaced by the new entrants are also typically high end goods. The “low-end” products in the TV market do not turnover nearly as much.

Table 2: **Characteristics of Entering and Exiting goods.**

<i>Specification</i>	Constrained OLS		Minimum Distance	
	exit	new	exit	new
1. S0 (Odd)	.106 (2.66)	.161 (4.14)	.075 (1.94)	.146 (3.86)
2. S0 (Even)	.121 (3.17)	.133 (3.53)	.097 (2.61)	.130 (3.51)
3. S9 (Odd)	-.004 (-.456)	.0108 (1.2973)	-.005 (-.558)	.009 (1.136)
4. S9 (Even)	.0003 (.0359)	.0101 (1.241)	-.003 (-.367)	.006 (.732)

S0 has a constant and two dummies, one for goods about to exit and one for goods that just entered, while S9 has these two dummies plus the easy to clean characteristics used in the analysis below (see the notes to table 3). Odd and Even number years done separately as they sample different cities. The constrained OLS regressions allow free coefficients for all the characteristics in each year but constrain the dummy coefficients to be the same across years. t-values in parenthesis below coefficients. The minimum distance estimates are similar but weight with the covariance matrix across years.

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 have liquid crystal display, but we do not have a good measure of the improvements that have occurred in sharpness of the liquid 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 (e.g., speed, RAM, hardrive capacity,) have natural cardinal measures which make them easy to compare across products .

The second point this table makes is that once we control for our (easy to clean) characteristics, the fact that a good is about to exit no longer seems to predict price. I.e. the characteristics seem to “control for” the differential prices of the goods that are about to exit. A similar statement is true for new entrants though here the t-values average about one (while for exits they average under .5).

3 Index Formulas

To begin as simply as possible we start with indexes that are linear in the logs of price relatives. The linearity in the logs of observed price relatives make the regression error in the hedonic regressions easy to deal with, and make the relationship of our results to “quality change” bias transparent. We intend to come back to more complex indexes that work directly with this regression error at a later date. For example to construct the Laspeyre’s index we need to exponentiate the logs and hence exponentiate the 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 disturbance as more appropriate.

3.1 Bimonthly indexes

All indexes are versions of

$$G_t = \sum_{q \in S_{t-2}} w_{q,t-2} y_{qt} \quad (1)$$

where q denotes a quote, w_{qt} is period- t weight, $y_{qt} = \log(p_{qt}/p_{q,t-2})$ is an actual or imputed log-relative, and S_{t-2} is a subset of all quotes active in period $t - 2$.⁸ This is the log of a geometric mean index, which approximates the average proportionate change in prices. The weights $w_{q,t-2}$ are obtained by dividing each period- $(t - 2)$ regional TV expenditure-share equally among all the quotes for that region and then renormalizing them so that $\sum_{q \in S_{t-2}} w_{q,t-2} = 1$. The regional expenditure shares are estimates from the CPI data base.⁹

Explicitly denoting an estimate of a log-relative as \hat{y}_{qt} , hedonic and matched-model indexes can be written as

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

$$G_t^{mm} = \sum_{q \in C_{t-2}} w_{q,t-2}^{mm} y_{qt}, \quad (3)$$

⁸The use of t and $t - 2$ is a result of the bimonthly sampling procedure which implies the basic indexes are bimonthly. Monthly indexes are derived from these by a linear splicing procedure.

⁹Past 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-2} is the set of quotes for which prices were successfully collected in period $t - 2$, and $C_{t-2} = A_{t-2} \cap A_t$. That is, the 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 $t - 2$. The hybrid indexes introduced in Pakes (2003) impute relatives only for TVs that exit between $t - 2$ and t , and use actual price relatives for goods that were available in both periods,¹⁰

$$G_t^{hyb} = \sum_{q \in C_{t-2}} w_{q,t-2}^{hed} y_{qt} + \sum_{q \in A_{t-2} - C_{t-2}} w_{q,t-2}^{hed} \hat{y}_{qt}. \quad (4)$$

4 Hedonic Predictions.

For simplicity we begin with a linear regression model for the (log) price levels of goods in a given period (we present non-parametric results directly thereafter). 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

$$d_t = (Z_t' Z_t)^{-1} Z_t' p_t, \quad (5)$$

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 $S4$, $S9$, and $S24$ are:

- $S4$: log of screensize in inches, a dummy indicator for projection TVs, the interaction between these two variables, and the square of log-screensize.

¹⁰The original rationale for these indices is that by using the actual price relatives for the continuing goods they act like the matched model index in not having any estimation error in those relatives, and by using the hedonic estimates of the price relatives for the exiting goods they eliminate much of the bias in the standard matched model index caused by the omission of price relatives for exiting goods.

- *S9*: the variables in *S4* 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 *s9* plus the additional variables listed in the notes to Table A1 at the end of the paper.

The values for the variables in *S4* and *S9* 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 at the time each index is prepared in a production setting. 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, most of which have values that are difficult to verify in the short period of time during which each index’s production. This is why the current method fits a regression no more than once a year.¹¹

The first three rows of Table 3 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 *S4* 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.

Note that there is a noticeable improvement in fit in moving from *S4* to *S9* but not much further improvement in adding the 15 characteristics needed for *S24*. The fourth panel of the table provides fits from a non-parametric estimate of the hedonic surface. The method used is local linear kernel regression with a cross validated bandwidth. Appendix 1 provides the formulae used in the non-parametric analysis. The small improvement in the fit of the *S24* regression relative to the *S9* regression is similar to the improvement obtained when we substitute the non-parametric (NP)regression for the linear regression; and the NP regressions use only the same 7 characteristics of the *S9* specification. Moreover the NP regression is also easy to compute in a production setting,

¹¹Indeed when we move to out-of-sample predictions, see our “to do” list at the end of this report, we will assume that we only have the *S9* variables to work with.

as it involves only running a pre-programmed algorithm.

Table 3: **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
S4	.8927	.8908	.8698	.8676	.9127	.9111
S9	.9552	.9533	.9413	.9369	.9663	.9649
S24	.9707	.9672	.9585	.9530	.9777	.9753
NP	.9641	.9626	.9236	.9179	.9730	.9719

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

4.1 Unobserved Characteristics and Hedonic Regression Functions.

Under standard assumptions on consumer behavior the prices of two goods with identical characteristics should be the same¹². Thus 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. This prediction problem, however, gets more complicated when there are characteristics of the goods that consumers value but Econometricians do not observe (and hence can not condition on). Recall that in the TV market exit is largely a result of high quality goods obsoleting older high quality goods, and that we do not have good cardinal measures which capture the differences between the different generations of high quality goods. As a result in correcting for the selection problems induced by exit one might want to pay particular attention to unobserved product characteristics.

Part of the impact of the unobserved product characteristics on price will be captured by their relationship to observed characteristics, but the rest will appear as the residual from the hedonic regression function. 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 t and $t + 1$ from the hedonic

¹²For a statement of this property, and a demand estimation algorithm that makes intensive use of it, see Bajari and Benkard, 2004

regression coefficients in $t + 1$ and the characteristics of the good that exited. However as we now show there are solid economic arguments to lead us to believe that the relationship between the unobserved and observed characteristics is different in the selected sample of exiting goods. Moreover the predictions these arguments make are borne out by the data.

For simplicity assume the true hedonic regression function is linear and let η measure the contribution of unobserved characteristics to price. Then the hedonic regression function is

$$E[p|z, \eta] = z\beta + \eta, \tag{6}$$

where we have normalized the coefficient of η to be one. The regression function we estimate is the regression of p only on z . To analyze its properties we need the properties of the regression of η on z .

If we let x denote exiting goods, n denote new entrants, and c denote continuing goods, then

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

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 (6) 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 is not so.

To see this we need a model for which goods exit. Temporarily assume that a product exits if its price falls below $\underline{p}(z)$. Then

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

In particular when the good's observed characteristics lead to a small $\underline{p}(z) - z\beta$ then the goods that continue will all have values of η which are very high, while if $\underline{p}(z) - z\beta$ is large, goods will continue even if they have low values 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 $S9$ regressor set of the last subsection, we then tested whether these coefficients differed from each other.

The results are presented in Table 4. They clearly reject the null that the new and exiting good interactions are all zero. We also calculated the average difference between the residuals of goods which continued and those that exited. This should approximate the difference in the average value of the unobserved characteristic's contribution to price between exiting and continuing goods (at least that part of the unobserved characteristics not accounted for by observed characteristics). It was $-.017$ with a standard error of $.007$, clearly negative and significant. That is a hedonic prediction for price which did not account for the contribution of unobserved characteristics would significantly overstate the price of exiting goods.

Table 4: **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 Significant At Different α Levels				
$\alpha = .01$.14	.11	.50	.54
$\alpha = .05$.29	.21	.71	.71
$\alpha = .10$.46	.29	.79	.75

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

A second possibility for forming a prediction for the prices of goods that exit is to assume that the contribution of the unobserved characteristic to price does not change over time. Now regress the log of the price differences (or price relatives) for the continuing goods onto their characteristics. Under the assumption that the contribution of the unobserved characteristics to price are constant over time these regressions difference out the contribution of the unobserved characteristics and provide an unbiased estimate of the changes in prices caused by the market's re-evaluation of

characteristics that results from entry of new goods and other changes in the environment. The predicted price change for the exiting goods would then be formed from these estimated differences in evaluations and the good's characteristics.

This procedure depends critically on the assumption that the value of the unobserved characteristic does not change over time.¹³ To see whether this is likely to be appropriate for our data, we split all continuing goods into three groups; the goods about to exit (they exit between $t + 1$ and $t + 2$), those recently new (they enter in year t), and the remaining goods. We then calculated the average change in the residual for each group in each period.

The results are presented in table 5. Interestingly the average change in the residual of all the continuing goods is slightly negative, though it is only marginally significant. This indicates that the new goods that enter in period $t + 1$ do better on the unmeasured characteristics than do the continuing goods, which should not be surprising. About to exit goods, that is goods that will exit between period $t + 1$ and $t + 2$, have an average change in residual which is very negative, more than five times the absolute value of the change in the residual for the average continuing goods. Further there is little doubt that it is significantly negative, as its t-value is also more than five. Since the contribution of the unobserved characteristic to price is falling just prior to exit, we expect it to continue falling during the exiting period. That is the assumption that the unobserved characteristic's contribution to price is constant over time is inconsistent with the data.

The results in this section also reinforce the selection argument that lead us to worry about matched model indices in the first place. I.e. we have shown that the market's evaluation of the unobserved characteristics of products that are about to exit are more negative than those of other goods and falling at a disproportionate rate in the periods just prior to exit.

¹³Though this might be likely if there were only one unobserved characteristic, if there are more than one of them their relative values will be a function of the choice set, which changes over time. This would be one source of the change in the contribution of the unobserved characteristics to price. The value of the unobserved characteristic could also change because consumer's perceptions of a value of one of these unobserved characteristics has changed over time.

Table 5: **Hedonic Disturbances for About to Exit, Recently Entered, Goods.**

<i>Variable</i>	All Continuing	a-Exit	r-New	Remaining Goods.
Using the S9 Specification for the Hedonic Regression ¹ .				
mean	-.0028	-.0150	-.0050	-.0021
s.d. of mean	.0017	.0028	.0025	.0021
s.d.(across months)	.0091	.0151	.0132	.0113
percent < 0	.6207	.8621	.5517	.6552
Using a Local Linear Kernel Regression for the Hedonic ¹ .				
mean	-.0023	-.0133	-.0026	-.0025
s.d. of mean	.0015	.0023	.0024	.0017
s.d.(across months)	.0081	.0126	.0130	.0093
percent < 0	.6897	.7931	.6552	.6552

¹ See the description of the S9 specification and the local linear regression in the text.

4.2 Hedonic Bounds in the Presence of Unobserved Characteristics.

We have hedonic equation for evaluating the observed characteristics of all goods marketed in each period. For continuity of notation we call the residual in that equation η (this is a slight abuse of notation as η is now not the unobserved characteristic per se, but rather it is the residual from projecting the unobserved characteristic onto the observed characteristics). Our problem is that we do not observed the value of η_{t+1} for goods that exit between t and $t + 1$. This section will introduce a prediction for η_{t+1} which maintains the hedonic bound in the sense that the resultant predictor for $p_{t+1} - p_t$ will have an expectation which is larger than the expectation of $p_{t+1} - p_t$ conditional on both (z_t, η_t) and the fact that the good exited between $t + 1$ and t .

Since we do not observe the η_{t+1} of the exiting goods in the second period, we will have to base our prediction on the η_{t+1} of the continuing goods. That is the regression function we can estimate is

$$E[\eta_{t+1} - \eta_t | z, \eta_t, j = c]. \tag{7}$$

To evaluate this expression we need models for the evolution of η_t and for exit. We will assume

that

$$\mathbf{F}_\eta = \{F(\eta_{t+1} | \eta_t); \eta_t \in \mathcal{R}\}, \quad \text{and that} \quad j = c \Leftrightarrow \eta_{t+1} \geq \underline{\eta}_{t+1}(z). \quad (8)$$

There are two assumptions in equation (8). First it is assumed that the stochastic process generating η is Markov and independent of z . Since our prediction for price conditional on z is a regression function, each period's η is mean independent of z by construction. So the content of the first part of the assumption is only to move us from mean independence to full independence. Note, however, that we do not restrict the functional form of the Markov process in any way.

The second assumption provides the exit rule; a good with observed characteristics z exits only if $\eta_{t+1} \leq \underline{\eta}_{t+1}(z)$. We want to emphasize that we place no assumptions on the function $\underline{\eta}_{t+1}(z)$, so this rule is consistent with virtually any model for exit in each period. In this sense our procedure is similar to the procedure used to correct for selection in production functions in Olley and Pakes (1994). Indeed, as is shown in Appendix 2, we could push the Olley and Pakes (1994) analogy further and generate a three step estimator for the η_{t+1} which, modulus estimation error, should generate a sharper bound than the one we use here. However as we explain in that appendix the tighter bounds are quite sensitive to estimation error, and since our intention is to produce a bound which is both robust and can be automated for use by the the BLS analysts, we shall ignore the tighter bound in the text.

Given these assumptions

$$\begin{aligned} E[\eta_{t+1} - \eta_t | z, \eta_t, j = x] &= E[\eta_{t+1} - \eta_t | \eta_{t+1} \leq \underline{\eta}_{t+1}(z), \eta_t, z] \leq E[\eta_{t+1} - \eta_t | \eta_{t+1} \geq \underline{\eta}_{t+1}(z), \eta_t, z] \\ &= E[\eta_{t+1} - \eta_t | z, \eta_t, j = c] = \frac{\int_{\underline{\eta}_{t+1}(z)} [\eta_{t+1} - \eta_t] dF(\eta_{t+1} | \eta_t)}{1 - F(\underline{\eta}_{t+1}(z), \eta_t)} \equiv g(\underline{\eta}_{t+1}(z), \eta_t). \end{aligned} \quad (9)$$

It is important to note that the inequality in equation (9) makes the expression

$$\hat{p}_{t+1} = z\beta_{t+1} + g(\underline{\eta}_{t+1}(z), \eta_t) + \eta_t \quad (10)$$

greater than expected price for *both* continuing and exiting goods.

We estimate β_{t+1} and $g(\cdot)$, substitute the estimates into equation (10), and use the resulting equation for our hedonic price predictions. The results of the last section indicated that the unobservables for exiting goods had two properties: (i) lower values of η_t and (ii) lower values of $\eta_{t+1} - \eta_t$ given η_t . The prediction from equation (10) will make a correction for the lower values of η_t of exiting goods, but it does not take into account the lower values of $\eta_{t+1} - \eta_t$ conditional on η_t . This is the source of the upward bias in the bound in equation (10) relative to the bound in appendix 2.

Partly as a result we will also introduce an alternative prediction for the prices of exiting goods. The alternative simply assumes that the change in the exiting good's price in the period in which it exits is at least as negative as it was in the period prior to it exiting. In this case we simply use the price change between periods t and $t - 1$ for the 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 always use the inequality in equation (10).

Table 6 presents the R^2 's from regressing $\eta_{t+1} - \eta_t$ on a polynomial in η and z for continuing goods; that is it provides the R^2 's from regressions used to form $g(\cdot)$ in equation (10). Recall that since the η 's are from the hedonic regression, the full sample of η 's are mean independent of the z 's by the properties of that regression function. So if there were no selection problem we would expect the first set of regressions to have adjusted R^2 's of zero. In fact they are highly significant which is evidence that the selection into continuing goods is at least partly based on the unobservables, as is modelled in equation (8). The prediction for $\eta_{t+1} - \eta_t$ improves noticeably when we use η_t as a predictor, so our ability to condition on the η_t prior to exit will be helpful in forming our predictions for goods similar to those that exit.

Two smaller points about this table are worth noting. First the regressions are meant to control for selection, and the appropriate selection correction for newly entered goods might be different

Table 6: **Predicting $\eta_{t+1} - \eta_t$ for Continuing Goods.**

r.h.s.	z		(z, η_t)		$(z, \eta_t), \text{r-New.}$	
Goods/Mean	R^2	Adj. R^2	R^2	Adj. R^2	R^2	Adj. R^2
all continuing	.15	.10	.27	.18	.28	.19
nonsticky-only	.16	.04	.43	.20	.47	.21

than for other continuing goods. So we did this once using a dummy for newly entered goods and once not. We get a small improvement in fit with the dummy, and hence will use the $\hat{g}(\cdot)$ that includes this dummy in what follows (though we get very similar results when the prediction without this dummy are used). The second point to note is that the selection correction fit is better for goods in the sample of goods that actually change their prices than in the overall sample (i.e. the sample which includes goods with sticky prices). We come back to this presently.

We conclude by summarizing the steps used in producing an index based on the predictions in equation (10).

- First we use all of the data to estimate unrestricted regression functions for each period.
- In the second step we use the residuals from those regressions to estimate equation (9) non-parametrically. This can only be done for the continuing goods and it gives us our estimate of $g(\cdot)$.
- There are now two possible estimates of β . One is from a continuing goods regression for the change in p minus the regression function for the change in η on the characteristics. The other is from the level regressions. We have tried both and though they need not be the same, they delivered two indices which are virtually identical.
- Finally use the estimated β and $g(\cdot)$ to construct the price predictions in equation (10).

4.3 Conditioning on the sticky-price rate

We noted that many of the prices are sticky, i.e. they do not change in value across two periods. As a result the hedonic price relative equation (and consequently the hedonic “ η -relative” equation), are quite different for the sticky and non-sticky price goods. As seen in Table 1 the prices for about to exit goods are noticeably less sticky than for a randomly drawn good. Moreover we expect goods which are similar to the exiting goods to have more reasons to change their prices in the year they exit than in the prior year. Consequently we now consider an adjustment to $g(\cdot)$ which accounts for possible differences in sticky price rates among the two groups of good.

Consider a polynomial regression function, and partition the continuing observations so that $\eta = (\eta'_1, \eta'_2)'$ where η_1 is the $n_1 \times 1$ vector associated with nonsticky prices, while η_2 is the $n_2 \times 1$ vector for the sticky. Let Z be the polynomial terms formed from (z, η) . If we sort and partition Z so that it is conformable to η , we can write the estimated coefficients for $g(\cdot)$ as

$$d = [Z'_1 Z_1 + Z'_2 Z_2]^{-1} [Z'_1 Z_1 d_1 + Z'_2 Z_2 d_2], \quad (11)$$

where

$$d_j = (Z'_j Z_j)^{-1} Z'_j \eta_j, \quad \text{for } d = 1, 2. \quad (12)$$

Equation (11) implies that the coefficient vector used to build $g(\cdot)$, i.e. our d , is a matrix-weighted average of d_1 and d_2 .

Using d for prediction of η changes is appropriate for imputing the relative of a good that is randomly selected from the distribution of all continuing TVs. If instead we know that a good is randomly selected from the distribution for continuing TVs that have non-sticking prices, then we would want impute its relative using d_1 . There is, however, a question of what should we do if we know only that the good is randomly selected from the distribution of exiting TVs. As we showed before, the probability of a good which is about to exit changing prices over the two periods just prior to exit is higher than that probability for a randomly chosen good. As a result we expect goods that are similar to the goods that exit to have a higher probability of price change than

randomly chosen goods, but just how much higher is uncertain. Accordingly we would like to make price predictions for exiting goods with some estimate intermediate “between” d_1 and d_2 .

So we introduce a family of coefficient estimators indexed by s , a “sticky-price” rate, defined as

$$d(s) = \left[(1-s) \left(\frac{Z_1' Z_1}{n_1} \right) + s \left(\frac{Z_2' Z_2}{n_2} \right) \right]^{-1} \left[(1-s) \left(\frac{Z_1' Z_1}{n_1} \right) d_1 + s \left(\frac{Z_2' Z_2}{n_2} \right) d_2 \right] \quad (13)$$

where s ranges from 0 to

$$s_0 = \frac{n_2}{n_1 + n_2},$$

the sticky-price rate for TVs that continue from $t-2$ to t . Note that $d(s_0) = d_1$ and $d(s=0) = d_2$.

CPI personnel do not observe a sticky price rate for the exiting goods, so there is no obvious way to choose an s . Consequently our choices for s in the empirical section below will be pragmatic and conservative. For example, we know from Table 1 that the sticky-price rate for “about-to-exit” TVs is on average lower than s_0 , and since the date for dividing months is arbitrary, exiting TVs should much more closely resemble about-to-exit TVs than they do other continuing TVs. Indeed, as noted, we expect the prices of goods with characteristics similar to the goods that exit to change more rapidly than those about to exit. This suggests using the about-to-exit sticky-price rate as an upper bound for s . We will provide the index that corresponds to s_0 , the about-to-exit s , and rates lower than that. This will give us a chance to see how sensitive our results are to the choice of s .

5 Geomean Indexes for TVs: Empirical Results

Table 7 gives the results of fitting alternative indexes to the TV data. It is divided into panels, each of which corresponds to a different procedure for calculating the index. Panel A uses a traditional hedonic regression in the log of price levels to construct the hedonic indices it displays. Panel B uses the change in log levels, or the price relative regression, to construct its indices. Panel C uses a combination of the prior year log price changes where available and our prediction (i.e equation 10) where not. Finally Panel D uses our prediction equation with adjustments for the sticky price

rate as explained in section 4.3. The panels differ somewhat in structure, as is needed to convey what we have actually learned from the data.

Throughout *hed9* and *hyb9* denote hedonic and hybrid indexes using linear regression with the *S9* regressor set. *hedNP* and *hybNP* use local-linear nonparametric regression with the same seven product characteristics as in *S9*. All reported indices are the mean of 31 monthly indexes, multiplied by 1200 to give the implied percentage annual inflation rate.

Consider first the matched model comparison to the index based on the traditional hedonic regression in log price levels in Panel A. When we use the *S9* regressor set the matched-model index registers more deflation than any of the regression-based indexes. However, this is largely a result of the fact that to do a reasonable job the level regression requires the use of the *S24* regressor set. When we use the *S24* regressor set the matched model and hedonic indices are about the same. This is similar to the results obtained by the BLS staff from the hedonic analysis on most component indices other than the index for computers. We note that this was the only panel in which the *S24* regressor set led to noticeably different results than from the *S9* regressor set. As a result of this, and the fact that the *S24* regressor set can not be used to produce an online hedonic index for the reasons discussed above, we only present the results for the *S24* regressor set in what follows.

The second panel reports an index based on a hedonic function estimated from the change in log prices, or the price relatives. This is the regression that would be appropriate if the contribution of unobserved characteristics to price were constant over time. The alternative indexes register small but appreciable reductions in inflation vis a vis both the matched model and the traditional hedonic indices. Notice also that the hedonic indices exhibit only slightly more variance across months than does the matched model indices. As expected (see Pakes,2003) the hybrid indices are in the middle in this regard. We have not yet investigated whether the increased variance is “real” or is an artifact of estimation variance, but whatever the source it is not large enough to worry too much about. All indexes now have a .7 or greater probability of being less than the matched-model.

Table 7: **Alternative Monthly Indexes for TV¹**

<i>SummaryStat</i>	matched model	hed9	hyb9	hedNP	hybNP
Panel A: Using Log-Price Regression Fit to All Observations					
mean	-10.01	-8.85	-9.68	-9.08	-9.74
(Use S24) ²	(-10.01)	(-10.09)	(-10.10)	(n.c.)	(n.c.)
s.d.	5.80	8.36	5.72	7.88	6.10
% l.t. mm		.39	.36	.39	.42
(S24 %)		(.52)	(.52)	(n.c.)	(n.c.)
Panel B: Using Log-Relative Regression Fit to Continuing TVs Only					
mean	-10.01	-10.63	-10.25	-10.68	-10.30
s.d.	5.80	6.27	5.92	6.39	5.98
% l.t. mm		.77	.68	.74	.65
Panel C: Using A-exit Price Changes Where Available, and Equation (10) Where Not ³ .					
mean	-10.01	-11.97	-11.48	n.c.	n.c.
use pre-exit s	-10.01	-12.02	-11.53	n.c.	n.c.
s.d.	5.80	6.21	6.08	n.c.	n.c.
% l.t. mm		.87	.68	n.c.	n.c.
Panel D: Using Equation (10) With Different Sticky-Price Rates s^4					
$s = s_0$	-10.01	-11.09	-10.59	n.c.	n.c.
s.d.	5.80	6.48	6.21	n.c.	n.c.
%l.t.mm		.78	.74	n.c.	n.c.
pre-exit	-10.01	-11.47	-10.98	n.c.	n.c.
avg pre-exit ratio	-10.01	-11.44	-10.94	n.c.	n.c.
$s = 0.5$	-10.01	-11.62	-11.13	n.c.	n.c.
$s = 0.4$	-10.01	-12.17	-11.68	n.c.	n.c.
$s = 0.3$	-10.01	-12.74	-12.25	n.c.	n.c.
$s = 0.2$	-10.01	-13.35	-12.86	n.c.	n.c.
$s = 0.1$	-10.01	-14.05	-13.56	n.c.	n.c.
$s = 0.0$	-10.01	-15.05	-14.66	n.c.	n.c.

Notes:

1. Above values are implied rates of percent annual change, obtained by multiplying the average monthly index by 1200. Averages are over 31 monthly indexes covering the period from June 2000 to January 2003.
2. S24 refers to the S24 regressor set. All other indices in this table are based on the S9 regressor set. n.c. means not calculated because there were too few observations to use non-parametric regression with this many regressors.
3. 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).
4. s_0 is the current period's sticky price rate. "pre-exit" sets s equal to the sticky-price rate for the preceding period's about-to-exit TVs (the current period's exited TVs). "avg pre-exit ratio" sets s equal to s_0 times the long-run average value (from the relevant bimonthly subsample) of sx_t/s_t , where s_t and sx_t are the unconditional and conditional-on-about-to-exit sticky price rates for period t . 4. The bandwidth used for the upper right entry in the third panel differs from that used for the bottom right entry in the bottom panel.

In sum, moving to the price relative regression, that is accounting for constant differences in the contributions of unobservables to price, moves us in the expected direction. However it is not clear whether it moves us far enough in that direction.

Panel C uses the price change in the period prior to the period in which the goods exit as a measure of the price change of goods similar to the goods that exit in the period in which they exit when that is available (it is available for about 83% of the exits) and our prediction, or equation (10), when not. The logic is that the price change for goods like the good that exits should be, if anything, more negative in the year they exit than in prior years. These results are strikingly different and indicate that there is a bias of about 20% in the matched model indices. Again using the estimated index results in only a slight increase in variance across months over that in the matched model index.

Panel D uses our prediction, equation (10), for all predicted price relatives. The first three rows present the results without any correction for the fact that the prices of goods that are similar to the goods that exit are likely to be less sticky than those of a randomly chosen good. These results are in between the results in panels B and C, with standard deviations across months similar to those in Panel C. This should have been expected since table 1 showed that; (i) about to exit goods have prices that change more often than those of a randomly chosen good, and (ii) when an about to exit good does have a price change that price change is, on average, more negative than that of a randomly chosen good whose price changes.

The next two rows present estimates which corrects for the first of these two phenomena. It assumes the sticky price rate for exiting goods is the same as it is for the goods that are about to exit (which is about .54). We do this twice; once for the actual pre-exit rate of change and once for the average of the pre-exit rates, but the results are virtually identical. These estimates show a fall in prices that is about .5% less than the estimates which use the about to exit price changes directly, with the difference likely to be attributable to (ii) above. Indeed we would have to decrease the stickiness parameter for exiting goods to between .4 and .5 for the predicted prices

from our method to match those obtained when we use the about to exit price changes directly. Note that if we assume all goods which are similar to those that exit have a price change, then we move our estimate from just below 12% to just over 15% and over 90% of our monthly estimates are lower than those from the matched model.

The movement of the price index from 10% to just below 12% is a movement of about 20%. The additional movement to 15% would give us a change in the price index of 50%. Even the 20%, which is a figure which is consistent with our different ways of accounting for the price falls of the goods that are like our exiting goods, is quite a sharp decrease in the index. Moreover throughout we have used assumptions which we would view as leading us to a lower bound. We note that the difference might well be different for different component indices, and one topic for future research is to get a rough idea of what fraction of the index, when corrected for exiting goods biases, are likely to result in changes of twenty per cent or more.

6 Comparison to the current CPI method of imputing relatives

The CPI subindex for TV is more elaborate than the simple matched model index used here. Our contribution is in the way to deal with exiting goods. Indeed we could accommodate all the other sources of complexity substituting only our treatment for exiting goods, and do a comparison of the resulting indexes. Though we are thinking of doing that, for now what we do is to adapt our matched model index to their treatment of exiting goods, and then compare the result to our index.

There are a number of ways in which the current CPI treats exiting goods. First it ignores exits for which the data gatherer does not provide an immediate substitute. This is equivalent to the treatment in our matched model index. On average about a quarter of exits are so ignored, and this is one source of bias.

When there is a substitute, either of two methods are used. Most often the price of the substitute TV obtained in period t is used as the numerator of the imputed relative. The denominator is the collected price of the period- $(t - 2)$ exit plus $b(z_q^s - z_q)$, where z_q^s is the characteristic vector for the

substitute, z_q is that for the exit, and b is a coefficient vector. The coefficient vector is obtained from a hedonic regression that is only updated annually. The coefficients are obtained from an OLS regression after prior restrictions are placed on coefficient signs.

We computed a "CPI-like" index using this method. The coefficient vector b was obtained by regressing log-price on our $S24$ variables and then setting to zero all negative coefficient estimates for product characteristics that, in our opinion, all consumers would like. A regression is run on the first month of the year in each of our bimonthly subsamples, and then re-fit at 12-month intervals. The spliced monthly index series generated an average annual rate of inflation of -9.99 percent, somewhat *lower* in absolute value than the rate of price decline calculated from our basic matched-model index (which was -10.01). If the sign restrictions are not imposed on the annual regression, we get an annual rate of -9.58 . A version of this method where the regressions are updated bimonthly is even farther from the matched-model index, at -9.88 . If sign restrictions are not imposed on these regressions, then -9.73 results.

The second CPI method for imputing relatives at a substitution is not explicitly hedonic. It is used when the substitute differs from the exit in an important "price-determining" way that is not represented by an element in x . First, the imputed relatives for all those substitutions for which hedonic imputation was possible are gathered.¹⁴ Then a geometric mean price index is constructed from these relatives only, and the index value is used as the desired imputation. We have not yet attempted to mimic this approach in a research index.

7 Future Research

There are a number of small details we need to examine and a few important tests we need to run before we make the suggestion that the BLS follow our, rather than prior, procedures for

¹⁴This is done at the regional level. The CPI computes TV indexes for 38 regions and aggregates them into a national index. We note that there are also a small number of substitutions that are judged to be essentially the same as the exiting TV. CPI relatives for these are merely the price of the substitute divided by the price of the exit. For the imputation method being described, they are included with the quality-adjusted relatives. For our preferred indexes, we predict their price just as we do other exits. Interestingly, the price dynamics for these "no-difference" substitutes differ from other "continuing TVs." They have a much lower sticky-price rate for example.

constructing this component index. First among the important tests is the need to do an out of sample prediction and compare the result to what the BLS obtained for the out of sample period. So we will write a program which automates the procedure for constructing the *S9* and *NP* indices and then apply it the data that has been gathered by the BLS since August 2001.

Second, we will extend our approach to regional geomean and Laspeyres indexes. The CPI computes 38 regional TV subindexes, and then treats each one as a price relative to be plugged into a national Laspeyres index. Currently the CPI uses geomean regional indexes for TVs, and there are good theoretical reasons for thinking the true index may differ across regions (see Pakes, 2003). We will construct a simple research index that mimics this two-stage procedure.

This will require exponentiating regional log-geomean indexes based on predicted log-relatives. Exponentiating log-geomean indexes converts a zero mean random prediction error in log levels into a residual which is not zero mean and hence must be accounted for in computing the index. This can done with a standard method if prediction errors are homoscedastic; see Pakes (2003). We have generalized this method to heteroscedastic errors, with initially promising results. One alternative here is to use local-linear (or some other non-parametric method) on levels rather than log-levels. This does away with the need for exponentiating an error, and should be rich enough to give a reasonably accurate picture of the hedonic surface (this would not be the case for the *S9* regressor set and linear in levels regressions).

We will also construct Laspeyres price indexes, which, if based on log level regressions also requires exponentiating individual predicted log-relatives. Our Laspeyres experimentation is based on the argument in Pakes (2003) that the upper bound it affords for a true COLI has a quality-change induced upward bias that can be reduced or eliminated by hedonic methods. We will also apply these arguments to a computation of the the first-order effect of quality change on a geomean index. We can do so because a Taylor expansion of a geomean index yields a Laspeyres index as its first-order term. The Laspeyres term has weights that depend only on the expenditure-share weights of the geomean and the vector of relatives around which the expansion is made, freeing us

from the difficult problem of calculating true Laspeyres weights.¹⁵

Finally we want to study the problem of how to choose the sticky-price parameter s more carefully. As noted we do have credible data-based values for a lower bound for this parameter, but we worry that they may be too conservative, resulting in a misleadingly insufficient correction of the quality-change bias of the matched-model index.

References

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

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.

Erickson, T., and P. Langohr. "Equilibrium Market Dynamics and the Measurement of Inflation and Quality Change, *mimeo*, 2006, Bureau of Labor Statistics.

Fan, J., and I. Gijbels. *Local Polynomial Modelling and Its Applications*, Chapman and Hall/CRC, 1996, Boca Raton, Florida.

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.

Newey, W.

¹⁵BLS interprets the regional geometric indexes as estimates of the true COLI for a representative consumer with a Cobb-Douglas utility function, which implies expenditure shares that do not depend on price and therefore remain constant over time. In contrast, Laspeyres expenditure-share weights must be updated every period to reflect price change.

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.

Pakes, A. and P. McGuire. "Computing Markov Perfect Equilibria; Numerical Implications of a Dynamic Differentiated Product Model" *RAND journal of Economics*, 25, , pp. 555-589.

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

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

Appendix 1: Local-linear Nonparametric Regression.

Local linear estimation predicts each individual value of the dependent variable with a separate weighted least squares regression. Specifically, the local linear WLS estimator for predicting the log-relative for the q^{th} value of the dependent variable is

$$\hat{\delta}(h, q) = (Z'\Omega(h, q)Z)^{-1} Z'\Omega(h, q)y, \quad (14)$$

where y is the vector of dependent variables, Z is the matrix of independent variables $\Omega(h, q)$ is a diagonal matrix whose diagonal elements are the weights assigned to each observation for the prediction $\hat{y}_{qt} = z_q\hat{\delta}(h, q)$, where z_q is the q^{th} row of Z . Note that it is indexed by the bandwidth parameter h as well as by q .

The i -th diagonal element of $\Omega(h, q)$ is a decreasing function of the distance between z_q and the z_i of the i -th observation. Specifically,

$$\Omega_{ii}(h, q) \propto \prod_{j=2}^K \exp \left\{ -\frac{1}{2} \left(\frac{z_{q,j} - z_{i,j}}{h \times s_j} \right)^2 \right\}, \quad (15)$$

where j indexes the columns of Z , and $s_j^2 = \sum_{i=1}^n (z_{i,j} - \bar{z}_j)^2 / (n - 1)$ is the sample variance of column j , the column mean being $\bar{z}_j = \sum_{i=1}^n z_{i,j} / n$.

The bandwidth h determines the rate at which the weights decrease with distance. We let the data select h by cross validation,

$$h = \arg \min \sum_{i=1}^n \left(y_i - z_i \hat{\delta}_{cv}(h, i) \right)^2,$$

where $\hat{\delta}_{cv}(h, i)$ is defined for each i by deleting z_i from Z and y_i from y and then evaluating (14) with the remaining $n - 1$ data points. We use the same h for all q . See Fan and Gijbels (1996) for details of this estimator.

Local-linear regression conditional on the sticky-price rate

Linearity means there is a sticky-price version of (14) also:

$$\begin{aligned} \hat{\delta}(h, q, s) &= \left[(1-s) \left(\frac{Z_1' \Omega_1(h, q) Z_1}{n_1} \right) + s \left(\frac{Z_2' \Omega_2(h, q) Z_2}{n_2} \right) \right]^{-1} \times \\ &\left[(1-s) \left(\frac{Z_1' \Omega_1(h, q) Z_1}{n_1} \right) \hat{\delta}_1(h, q) + s \left(\frac{Z_2' \Omega_2(h, q) Z_2}{n_s} \right) \hat{\delta}_s(h, q) \right] \end{aligned} \quad (16)$$

where s is defined as in the text

$$\Omega(h, q) = \begin{pmatrix} \Omega_1(h, q) & \mathbf{0} \\ \mathbf{0} & \Omega_2(h, q) \end{pmatrix},$$

and

$$\hat{\delta}_j(h, q) = (Z_j' \Omega_j(h, q) Z_j)^{-1} Z_j' \Omega_j(h, q) y_j. \quad (17)$$

Corresponding hybrid and hedonic indexes that are conditional on s are

$$G_t^{hyb}(s, h) = \sum_{q \in C_{t-2}} w_{q,t-2}^{hed} y_{qt} + \sum_{q \in A_{t-2} - C_{t-2}} w_{q,t-2}^{hed} x_q \hat{\delta}(h, q, s) \quad (18)$$

and

$$G_t^{hed}(s, h) = \sum_{q \in C_{t-2}} w_{q,t-2}^{hed} x_q \hat{\delta}(h, q, s) + \sum_{q \in A_{t-2} - C_{t-2}} w_{q,t-2}^{hed} x_q \hat{\delta}(h, q, s) \quad (19)$$

Appendix 2 “Tighter” Hedonic Bounds.

Recall that section 4.2 equation (9) gave us the expectation of $\eta_{t+1} - \eta_t$ conditional on survival as

$$E[\eta_{t+1} - \eta_t \mid z, \eta_t, j = c] = \frac{\int_{\underline{\eta}_{t+1}(z)} [\eta_{t+1} - \eta_t] dF(\eta_{t+1} \mid \eta_t)}{F(\underline{\eta}_{t+1}(z), \eta_t)} \equiv gx(\underline{\eta}_{t+1}(z), \eta_t).$$

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

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

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

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

$$\underline{\eta}_{t+1}(z) = \mathcal{F}_\eta^{-1}[\mathcal{F}(\underline{\eta}_{t+1}(z), \eta_t)],$$

and substitute it into equation (9) to obtain

$$E[\eta_{t+1} - \eta_t \mid z, \eta_t, x = c] = g(\mathcal{F}_\eta^{-1}[\mathcal{F}(\underline{\eta}_{t+1}(z), \eta_t)], \eta_t) \equiv h(\mathcal{F}_t, \eta_t), \quad (20)$$

where $\mathcal{F}_t \equiv \mathcal{F}(\underline{\eta}_{t+1}(z), \eta_t)$.

Both \mathcal{F}_t and η_t can be estimated, and hence, if we temporarily ignore estimation error, can be treated as observable. So we can substitute equation (20) into equation (7) to obtain

$$E[p_{t+1} - p_t \mid z, \eta_t, j = c] = z(\beta_{t+1} - \beta_t) + h(\mathcal{F}_t, \eta_t). \quad (21)$$

This is the equation we take to data, and it allows us to estimate both the function $h(\cdot)$, and

$(\beta_{t+1} - \beta_t)$.¹⁶

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

$$0 = E[\eta_{t+1} - \eta_t \mid z] \equiv \mathcal{F}_t(z)E[\eta_{t+1} - \eta_t \mid \eta_{t+1} \geq \underline{\eta}_{t+1}(z), z] + [1 - \mathcal{F}_t(z)]E[\eta_{t+1} - \eta_t \mid \eta_{t+1} \leq \underline{\eta}_{t+1}(z), z].$$

Consequently

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

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

$$E[p_{t+1} - p_t \mid z, \eta_t, j = x] = z(\beta_{t+1} - \beta_t) - \frac{\mathcal{F}_t h(\mathcal{F}_t, \eta_t)}{[1 - \mathcal{F}_t]}. \quad (22)$$

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. The likely reason is that the denominator in the latter equation is estimated and if the true value of \mathcal{F}_t is near one even a small amount of estimation error will cause very imprecise estimates of the truncated expectation. Accordingly we advise against using this estimator, unless of course the underlying data set is very large.

Appendix 3: 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.

¹⁶Formally 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 econometric details see the review of semiparametric techniques by Newey () and the literature he cites.

Table 8: **Average Characteristic Vectors for Subsets of TVs.**

<i>characteristic</i>	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 sur	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 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" refers to TVs present in t but not present in $t-2$.