

# Cross-validation of quality-adjustment methods for price indexes

*Price index programs may use several quality-adjustment methods, but they often lack guidance on when to use a given method. As part of an effort to improve quality adjustment for network switches in the Producer Price Index, we compare some standard quality-adjustment methods, using cross-validation. For our sample, seemingly reasonable hedonic specifications impute out-of-sample prices less accurately than other, traditional quality-adjustment methods.*

Many price indexes attempt to track price changes for the same set of products over time. Discontinued products complicate this task, because they lack observable prices. Often, a new product is added to the sample to replace a discontinued product. The price difference between the old product in its last period and the replacement product in the next period combines two effects: the quality differences of the two products and the overall price movement.

Several quality-adjustment methods of varying complexity are available to decompose the two effects.<sup>1</sup> Hedonic approaches impute prices with regressions. Other methods can be as simple as asking the price setter about production cost differences or assuming that one of the two effects is zero. Compilers of price indexes want to use the most accurate method, but econometric theory provides little practical advice on which method is best for any particular product category.

This article explores a novel use of out-of-sample cross-validation to empirically compare quality-adjustment methods. Our specific application focuses on network switches, a product category included in producer price indexes, but cross-validation approaches may be applied to any category. We compile observed prices and characteristics for 592 product models. We use quality-adjustment methods to impute one of the prices as if the true price were unknown. The difference between this imputed price and the actual, observed price measures a method's accuracy. We repeat this exercise for all product models.



## Brian Adams

[adams.brian@bls.gov](mailto:adams.brian@bls.gov)

Brian Adams is a research economist in the Office of Prices and Living Conditions, U.S. Bureau of Labor Statistics.

## Alexander Klayman

[klayman.alexander@bls.gov](mailto:klayman.alexander@bls.gov)

Alexander Klayman is an economist in the Office of Prices and Living Conditions, U.S. Bureau of Labor Statistics.

Earlier works have used cross-validation in comparisons of hedonic specifications.<sup>2</sup> This article evaluates hedonic and nonhedonic quality-adjustment methods. We compare the accuracy of the link-to-cell-relative method, the direct comparison method, and several hedonic specifications. We find that, for network switches, some simple hedonic specifications outperform more flexible specifications at imputing out-of-sample prices. However, in our work, none of the hedonic models are more accurate than a simple benchmark of the link-to-cell-relative method, which adjusts a product's previous price by the average percent change in the prices of the other products.

## Industry background and data

Network switches are a type of telecommunications equipment that directs data transfers between computers and other network devices. While similar to routers, network switches are optimized for dynamically establishing dedicated connections among devices within a local network. Network switches, together with telephone and wireline data networking equipment, represent 0.12 percent of final demand and 0.19 percent of processed goods for intermediate demand in the U.S. Producer Price Index (PPI).<sup>3</sup> The comparisons in this article are part of a larger investigation into PPI's potential adoption of hedonic quality adjustments for network switches. This effort is due to frequent difficulties in obtaining direct estimates of quality-adjustment values from respondents.

Data for this study were gathered quarterly from a prominent retailer's online catalog, covering the period December 2016–September 2017. (These data are not those collected for calculating the PPI.) The catalog includes retail prices and, usually, detailed product specifications. The sample includes 592 product listings, with between 480 and 514 listings available in a particular quarter and 408 listings appearing in all four quarterly cross sections. The turnover rate for specific listings is higher than that for network switches generally. Product cycles for switches can last for almost a decade.<sup>4</sup>

Most of the listings are for managed switches, a type of network switch with settings that can be continually adjusted. The average price of these switches is \$4,396, and most of their buyers are large organizations or data centers.

The retailer's catalog includes dozens of product characteristics for most, but not all, listings. The regressions in this article use a smaller set of characteristics that are given in every listing:

- *Brand and series indicators*—indicator variables representing the brand or manufacturer of the device and its series.
- *Layer 2+, layer 3, layer 3 lite, layer 4, layer 4–7*—indicator variables for the Open Systems Interconnection model, which characterizes communication functions implemented by a switch. A layer indicates a switch's level of sophistication and its capability to route network traffic on the basis of that traffic's destination.
- *Ports*—the total number of ports in the switch device.
- *MAC addresses*—variables representing the greatest number of routing entries that the switch can hold in its memory and indicating the upper limit of the size of the Ethernet network that the device can support.
- *Switching capacity*—variable, measured in gigabytes per second, that determines the highest available aggregate bandwidth capacity possible for the switch and represents the degree of dedication of connection lines at full speeds for all port pairs attempting to communicate at the same time.

## Quality-adjustment approaches

PPI, like other price index programs, employs various quality-adjustment methods. The two preferred approaches for network switches are the explicit (respondent-provided) quality adjustment and the overlap method. When a discontinued product is replaced by a product from the same manufacturer, the PPI economist asks the price reporter at the manufacturer, “What is the production cost difference marked up to the net selling price?” If the survey respondent provides an answer, his or her estimate is almost always used. Respondents for network switches usually are unable or unwilling to provide an estimate. When this occurs, the first recourse is the overlap method, which uses the difference in prices in a period when both products are present. If a replacement product is new or if its price in previous periods cannot be collected, the overlap method is not an option.

When these two methods are unavailable, index compilers often use the product category’s average price change as a proxy for the unobservable price movement. PPI economists call this method “link to cell relative.” It is equivalent to dropping the product for a period and reallocating its sample weight to similar products. If the price movements of discontinued products differ from those of surviving products, this method results in a biased index. This bias can result from any of several factors. Diminishing demand and falling prices might cause sellers to discontinue products. Pricing decisions and margins might be adjusted at the switchover phase of a product cycle. In some product categories, this selection bias is thought to be large.<sup>5</sup>

In a direct comparison, the entire price difference between a retired product’s last period and the replacement product’s first period is regarded as a price change. The method assumes no quality difference between the two products. PPI economists use direct comparison when they consider, often with knowledge gained through contact with survey respondents, the products’ differences to be negligible. The opposite extreme would be to assume that the entire price difference is due to a change in quality and none is due to real price movement. Incorrectly assuming away price movement biases an index toward zero.

Hedonic approaches estimate a relationship between prices and product characteristics. This relationship can then be used to impute a product’s price after it is discontinued or to estimate the difference in prices between a discontinued product and its replacement (even if their availability dates do not overlap).<sup>6</sup> Let  $p_{i,t}$  be the price for product  $i$  in period  $t$  and let  $X_i$  be a vector of its observable characteristics. A simple regression model would be

$$p_{i,t} = X_i\beta_t + \varepsilon_{i,t},$$

where  $\varepsilon_{i,t}$  is the econometric error term, which includes the price effect of unobserved characteristics and the idiosyncratic portion of the price treated as random. An imputation of  $p_{i,t}$  is given by  $X_i\hat{\beta}_t$ , where  $\hat{\beta}_t$  is a vector of the coefficients estimated by the regression. The error of such an imputation is

$$p_{i,t} - X_i\hat{\beta}_t = X_i(\beta_t - \hat{\beta}_t) + \varepsilon_{i,t}.$$

Let  $r$  be product  $i$ ’s replacement, so that  $p_{r,t}$  is the price of the replacement. The traditional approach used by the PPI program approximates product  $i$ ’s price as  $p_{r,t} + (X_i - X_r)\hat{\beta}_t$ . The error for the hedonic quality adjustment on the replacement is

$$p_{i,t} - [p_{r,t} + (X_i - X_r)\hat{\beta}_t] = (X_i - X_r)(\beta_t - \hat{\beta}_t) + \varepsilon_{i,t} - \varepsilon_{r,t}.$$

In this notation, the direct comparison method uses  $p_{r,t}$  as an imputation of  $p_{i,t}$ , yielding an error of

$$p_{i,t} - p_{r,t} = (X_i - X_r)\beta_t + \varepsilon_{i,t} - \varepsilon_{r,t}.$$

Without unrealistic assumptions, none of these errors can be shown to be smaller or greater than any other error, and their expected values and variances cannot be ordered. In idealized settings, the econometric error term is uncorrelated with observed product characteristics. If the hedonic regression model satisfies the conditions for the Gauss–Markov theorem, then  $\hat{\beta}_t$  is an unbiased estimate of  $\beta_t$  and the econometric error terms would have an expected value of zero. Hedonic imputation and hedonic quality adjustment on the replacement would have zero expected error. Network switches are differentiated products with some market power, so in this setting hedonic regression estimates may be neither unbiased nor consistent.<sup>7</sup> If the regression bias is large, direct comparison may give more accurate predictions than hedonic imputation. Conversely, if the replacement product is not a close match, hedonic imputation would give a smaller error than direct comparison. However, no theorem exists to determine which method is more accurate in a particular setting.

## Hedonics estimation

To empirically compare hedonic methods with other quality-adjustment methods, we first estimate the hedonic models. Table 1 presents coefficient estimates for four regression specifications using the March 2017 cross section; other periods have slightly different coefficient estimates. Specifications 1–3 (reported in columns 1–3) use price as a dependent variable, and specification 4 uses the natural logarithm of price. The product characteristics (components of the vector  $X_i$ ) include indicator variables for layer types (layer 3, layer 3 lite, layer 4; layer 2 is the omitted characteristic), indicator variables for the product series, the number of MAC addresses, the number of ports, and the switching capacity (measured in gigabytes per second).

**Table 1. Hedonic regression coefficients for March 2017 cross section**

Variable	(1) Price	(2) Price	(3) Price	(4) Ln(price)
Layer 3	1,764.2*** (386.8)	1,602.2*** (347.3)	1,270.6*** (344)	0.661*** (0.0797)
Layer 3 lite	2,188.2 (1,441.1)	1,490.6 (1,294.6)	662.3 (1,258.6)	1.372*** (0.297)
Layer 4	1,746.9 (962.1)	2,355.1** (865.1)	2,171.7* (849.8)	0.782*** (0.199)
MAC addresses	0.0390*** (0.00336)	0.00617 (0.00433)	0.0355** (0.0136)	0.00000410*** (0.000000995)
Ports	41.59*** (10.29)	29.17** (9.308)	4.357 (47.79)	0.0223*** (0.00214)
Switching capacity (gigabytes per second)	—	6.072*** (0.577)	10.29*** (1.587)	0.000774*** (0.000132)
MAC addresses squared (millions)	—	—	-0.036 (0.0463)	—
Ports squared	—	—	0.667 (0.739)	—
Switching capacity squared	—	—	-0.901 (0.506)	—
Switching capacity x ports (thousands)	—	—	-23.01 (26.95)	—
Switching capacity x MAC addresses (thousands)	—	—	-0.00725 (0.0071)	—

See footnotes at end of table.

**Table 1. Hedonic regression coefficients for March 2017 cross section**

Variable	(1) Price	(2) Price	(3) Price	(4) Ln(price)
Ports x MAC addresses (thousands)	—	—	-0.447* (0.221)	—
Series indicators	Yes	Yes	Yes	Yes
Observations	480	480	480	480
$R^2$	0.681	0.744	0.765	0.714
Adjusted $R^2$	0.662	0.728	0.748	0.697
<p>*<math>p &lt; 0.05</math>, **<math>p &lt; 0.01</math>, ***<math>p &lt; 0.001</math></p> <p>Note: Standard errors are shown in parentheses.</p> <p>Source: Authors' calculations based on online retailer data.</p>				

Specification 1 represents the simplest regression. It shows that each additional port is associated with a price increase of \$41.59 and each additional MAC address is associated with a price increase of \$0.04.

Specification 2 adds an important variable, switching capacity. With that addition, the adjusted  $R^2$  increases from 0.662 to 0.728. An additional port is associated with a price increase of \$29.17, and an additional gigabyte per second of switching capacity is associated with a price increase of \$6.07.

Specification 3 adds higher order and interaction terms for all characteristics, except the layer and series indicators. This addition further increases the adjusted  $R^2$ , but only to 0.748. Despite 480 observations (and, thus, 445 degrees of freedom), adding these terms causes overfitting and, as will be shown in a later section, reduces out-of-sample accuracy.

Specification 4 uses the same explanatory variables as those used in specification 2, but its dependent variable is the logarithm of price. Here, an additional port is associated with a price increase of roughly 0.02 percent, an additional MAC address with a price increase of 0.000004 percent, and an additional gigabyte per second of switching capacity with a price increase of 0.00077 percent.

Our specifications are typical of the hand-selected hedonic models used in price indexes. They fit respectably well, use important product characteristics, and have few unexpected signs. Like all hedonic models, they introduce some estimation error into quality adjustment. How their fit compares with that of entirely different quality-adjustment approaches has rarely been explored.

## Comparing quality-adjustment methods

If the price of a discontinued product could be observed, the accuracy of quality-adjustment methods could be assessed directly. The methods' imputed prices could be compared with the actual price. Because discontinued products do not have observable prices, there is a need to impute prices and make quality adjustments. So instead of imputing prices for discontinued products, we impute prices for continuing products. We then compare the known prices of continuing products with the various imputations that quality-adjustment methods would have provided if these prices were unknown. Our cross-validation exercise cannot detect the selection bias that might afflict the link-to-cell-relative method or a direct comparison, but it can detect poor fit and inaccurate imputation in

continuing products. A model or method that cannot accurately impute the prices of continuing products likely cannot accurately impute the prices of discontinued products either.

The first step in the cross-validation exercise is to select one of the continuing products and pretend its new price is unknown. We then impute that price in several ways. We use all the other products available in the new period to re-estimate the hedonic model and then multiply the coefficient estimates by the selected product's characteristics to find the hedonic imputation value. Next, we observe the selected product's old price and multiply it by the average price increase for all other products, to simulate the link to cell relative, or by one, to simulate the assumption of no price change. We then match the product with a simulated replacement and use that replacement's price as the value from direct comparison. Finally, we find the difference in hedonic imputations for the selected and replacement products and add that difference to the replacement product's price in the new period to calculate the price implied by hedonic quality adjustment on the replacement.

Usually, the choice of a replacement product in the PPI is influenced by information provided by a survey respondent. In this exercise, however, it is selected through an algorithm designed to simulate that process. When possible, the replacement product is of the same brand, layer type, and series as the selected product. If more than one product has the same brand, layer type, and series, the replacement product is the one nearest to the selected product in switching fabric capacity. If there are multiple products of the same switching fabric capacity, the difference in the number of MAC addresses serves as a first tiebreaker and the difference in the number of ports serves as a second tiebreaker.

Table 2 presents an example illustrating the imputation calculations. One of the products in the online retailer's catalog had an observed price of \$13,767 in both March and June. The exercise is to impute the June price as if it were unobserved. The simplest imputation is to assume no price change and use the last period's price. In this example, the last period's price was \$13,767, which, because there was no change, perfectly imputes the June price. The geometric average price change for all other products was -0.38 percent, so adjusting the March price by the average price change yields a price of  $\$13,767 \times (1 - 0.0038) = \$13,714$ , which underestimates the actual June price by \$53.

**Table 2. Imputation and calculation example A**

Category	Selected model	Replacement model
Characteristics:		
Ports	40	48
Switching fabric capacity (gigabytes per second)	1,475	496
MAC addresses	294,912	65,536
Layer type	Layer 3	Layer 3
Power over Ethernet	No	Yes
June 2017 price	\$13,767	\$5,910
Imputed prices:		
No change (March 2017 price)	\$13,767	—
Average change	\$13,714	—
Hedonic imputation (using specification 2)	\$10,412	\$5,623
Hedonic imputation (using specification 3)	\$22,008	\$4,575
Direct comparison	\$5,910	—

See footnotes at end of table.

**Table 2. Imputation and calculation example A**

Category	Selected model	Replacement model
Hedonic quality adjustment on replacement (using specification 2)	\$10,699	—
Hedonic quality adjustment on replacement (using specification 3)	\$23,343	—
Errors:		
No change	\$0	—
Average change	-\$53	—
Hedonic imputation (using specification 2)	-\$3,355	—
Hedonic imputation (using specification 3)	\$8,241	—
Direct comparison	-\$7,857	—
Hedonic quality adjustment on replacement (using specification 2)	-\$3,068	—
Hedonic quality adjustment on replacement (using specification 3)	\$9,576	—

Source: Authors' calculations based on online retailer data.

Because this exercise treats the June price for one product as unobserved, the hedonic regression is estimated anew by excluding the selected product and using the other 486 products available in June. Coefficient estimates differ slightly from those in table 1, because the period differs and because one product is omitted. The estimated coefficients multiply the observed characteristics to give a fitted value for price, and this number is the hedonic imputation. If specification 2 (presented in column 2 of table 1) is used, the hedonic imputation is \$10,412, which underestimates the observed price by \$3,355. If specification 3 is used, the hedonic imputation is \$22,008, which is \$8,241 more than the actual price.

The algorithm that finds a replacement product matches the selected switch to a closely related product, with a product number that differs by only one digit. Both products have the same layer type, switching fabric capacity, and number of ports. The selected product handles more MAC addresses than its replacement, and the two models differ in other characteristics, such as the capability to be powered over Ethernet. A direct comparison would use the replacement's price of \$5,910 as a proxy for the selected product's June price. In this example, the selected product's actual price is \$7,857 more than the price of the replacement, which serves as a poor proxy.

Because the two product models have different characteristics, hedonic regressions impute different prices. The difference in imputed prices is the estimated quality adjustment. One regression specification expects the selected product to command \$4,789 more than its simulated replacement, the other specification \$17,433 more. Added to the replacement product's price of \$5,910, these quality adjustments give estimates of \$10,699 and \$23,343 for the June price of the selected product. The actual June price of \$13,767 is \$3,068 higher and \$9,576 lower than these estimates. The quality adjustment is too small in the former case and too big in the latter.

Table 3 presents a second example. Between March and June, the price of the product selected for this exercise increased from \$5,729 to \$6,227. If the June price were unobserved, the March price would be an imperfect estimate of it. An adjustment assuming no price change generates an error of \$498, and an adjustment assuming the average price change generates an error of \$521. Hedonic imputation errs in the other direction, estimating a price of \$7,096 (specification 2) or \$8,220 (specification 3).

**Table 3. Imputation and calculation example B**

Category	Selected model	Replacement model
Characteristics:		
Ports	48	48
Switching fabric capacity (gigabytes per second)	176	176
MAC addresses	32,000	32,000
Layer type	Layer 3	Layer 3
Image type	LAN base	Standard
June 2017 price	\$6,227	\$6,750
Imputed prices:		
No change (March 2017 price)	\$5,729	—
Average change	\$5,706	—
Hedonic imputation (using specification 2)	\$7,096	\$7,096
Hedonic imputation (using specification 3)	\$8,220	\$8,220
Direct comparison	\$6,750	—
Hedonic quality adjustment on replacement	\$6,750	—
Errors:		
No change	-\$498	—
Average change	-\$521	—
Hedonic imputation (using specification 2)	\$869	—
Hedonic imputation (using specification 3)	\$1,993	—
Direct comparison	\$523	—
Hedonic quality adjustment on replacement	\$523	—
Source: Authors' calculations based on online retailer data.		

The matching algorithm returns another product model closely related to the simulated replacement product. The selected and replacement products differ in some characteristics, but not in the characteristics used as explanatory variables in the hedonic regressions. Thus, the imputed price for the replacement is the same as that for the original. Consequently, the estimated quality adjustment is zero, making the hedonic quality adjustment on the replacement identical to direct comparison. Direct comparison errs only by \$523, as the replacement product is quite similar in characteristics and price to the selected model.

A leave-one-out cross-validation repeats this exercise for all continuing products in all three time intervals. Table 4 presents the average of the imputed price subtracted from the actual price. For a large-enough sample, this difference would measure the imputation method's bias. Assuming no price change has an obvious bias; it systematically underestimates prices in inflationary periods. Yet, because actual prices change so slightly in this sample and period, the assumption produces the third-smallest error of the methods tested. Adjusting the old price by the average price change produces the smallest mean error, -\$2.

**Table 4. Mean imputation error, by quality-adjustment method**

Method	Mean error			
	March	June	September	Overall
Assume no change	-\$67	-\$6	-\$2	-\$24

See footnotes at end of table.

**Table 4. Mean imputation error, by quality-adjustment method**

Method	Mean error			
	March	June	September	Overall
Assume average change	-55	11	34	-2
Hedonic imputation				
Using specification 2	13	71	21	35
Using specification 3	-9	40	-21	3
Using specification 4	-166	-164	-270	-202
Direct comparison	118	98	157	126
Hedonic quality adjustment on replacement				
Using specification 2	73	35	78	62
Using specification 3	46	14	41	34
Using specification 4	-445	-485	-444	-457

Source: Authors' calculations based on online retailer data.

An unbiased imputation method might lower accuracy, making large but offsetting imputation errors. Another measure of accuracy, the mean of the absolute values of the differences between imputed and actual prices, appears in table 5. By this measure, assuming no price change gives the most accurate imputations in every period. Indeed, 59 percent of the listings in the sample have the same March and June prices, and so have no error.

**Table 5. Mean absolute imputation error, by quality-adjustment method**

Method	Mean error			
	March	June	September	Overall
Assume no change	\$202	\$151	\$76	\$141
Assume average change	211	164	105	159
Hedonic imputation				
Using specification 2	1,710	1,737	1,662	1,702
Using specification 3	1,783	1,836	1,739	1,784
Using specification 4	2,579	2,645	2,586	2,603
Direct comparison	1,574	1,566	1,476	1,537
Hedonic quality adjustment on replacement				
Using specification 2	1,568	1,546	1,469	1,526
Using specification 3	1,604	1,649	1,538	1,595
Using specification 4	2,258	2,310	2,203	2,255

Source: Authors' calculations based on online retailer data.

Assuming the average price change is similar to assuming no price change, because the average change is so small for these periods. Using the average change of the continuing goods to adjust the price of discontinued goods is problematic whenever discontinued and continuing goods have different price movements. The present exercise cannot address this bias, as errors can be calculated only for continuing goods.

A hedonic quality adjustment on the replacement using the simpler specification still outperforms direct comparison and (by some measures) the other hedonic methods. Direct comparison has a modest mean absolute error despite having a high mean error. While its imputation errors are small, they tend to move in the same direction, underestimating new prices. Hedonic quality adjustment and direct comparison give the same price imputations and the same errors if the selected and replacement products have the same values for all characteristics used in the regression. Such exact matches occur in at least a third of the sample. Hedonic quality adjustment has less error than direct comparison for most of the products for which specification characteristics differ, and it holds its advantage even in subgroups for which the replacement is a close match.

The purpose of these quality-adjustment methods, however, is not to estimate individual prices but to produce an accurate price index. Table 6 displays the quarterly price changes as they would be measured if imputed prices were used instead of actual prices. More specifically, each row gives the unweighted geometric mean of the new price estimated by the quality-adjustment method divided by the actual old price. This mean would be the equivalent of the price index if all products were discontinued and the given quality-adjustment approach was used.

**Table 6. Price-change estimates, by quality-adjustment method**

Method	Estimated mean price change (percent)		
	December–March	March–June	June–September
Actual prices	-0.30	-0.50	-1.00
Assume no change	0.00	0.00	0.00
Hedonic imputation			
Using specification 2	16.20	14.80	18.20
Using specification 3	9.20	7.30	10.00
Using specification 4	-1.00	-2.50	-2.60
Direct comparison	-5.30	-4.80	-8.30
Hedonic quality adjustment on replacement			
Using specification 2	-4.20	-5.40	-6.50
Using specification 3	-2.90	-4.30	-5.20
Using specification 4	-4.30	-4.30	-7.50

Source: Authors' calculations based on online retailer data.

Assuming no price change for all products always returns a price-change estimate of zero. From December 2016 to September 2017, the actual price movement was near zero, with quarterly price changes of -0.3 percent (December–March), -0.5 percent (March–June), and -1.0 percent (June–September). An index biased toward zero has little error. In periods with large price movements, assuming no price change would produce large errors.

Comparing price ratios shows the advantage of the log regression specifications. Although these specifications' imputed prices have the highest mean and absolute errors, their imputed price ratios are more accurate than those from linear regression specifications. Using specification 4 in the interval ending in March, hedonic imputation of all products produces a price-change estimate of -1.0 percent, compared with 9.2 percent or 16.2 percent for the linear specifications.

## Conclusion

Out-of-sample cross-validation can be applied to selecting hedonic regression specifications. It may also help compare hedonics with other quality-adjustment methods. The method cannot solve the problems arising from the fact that prices of discontinued items are fundamentally unobservable or that the final price movements of discontinued items might differ from those of continuing goods. However, our cross-validation approach can show how accurately various methods impute prices and correct for quality differences, at least among continuing products. For network switches in early 2017, the average price change of all other products yielded better approximations of actual prices than did any other quality-adjustment method tested. Our best performing hedonic specifications had much higher imputation error than PPI's current imputation methodology for network switches. This finding was one consideration in the program's decision not to extend hedonic quality adjustment to network switches in the PPI. Our results may be sector and time specific, but our approach can be applied widely and repeatedly. In particular, it can be reapplied when data with additional product characteristics become available or when new macroeconomic regimes arise.

### SUGGESTED CITATION

Brian Adams and Alexander Klayman, "Cross-validation of quality-adjustment methods for price indexes," *Monthly Labor Review*, U.S. Bureau of Labor Statistics, June 2018, <https://doi.org/10.21916/mlr.2018.18>

### NOTES

<sup>1</sup> For brief introductions to the quality-adjustment approaches used by the PPI program, see *Handbook of methods*, chapter 14, "Producer prices" (U.S. Bureau of Labor Statistics), <https://www.bls.gov/opub/hom/pdf/homch14.pdf>; and "Quality adjustment in the Producer Price Index" (U.S. Bureau of Labor Statistics, August 2014), <https://www.bls.gov/ppi/qualityadjustment.pdf>.

<sup>2</sup> For an example from real estate economics, see Allen C. Goodman and Thomas G. Thibodeau, "Housing market segmentation and hedonic prediction accuracy," *Journal of Housing Economics*, vol. 12, no. 3, 2003, pp. 181–201.

<sup>3</sup> Telephone and wireline data networking equipment is PPI commodity WPU117601.

<sup>4</sup> Product cycles for managed network switches are typically triggered by technological advances in network speeds, microprocessors, and computer memory. Major product cycles have generally occurred gradually, in step with increasing network speeds. Around 1988, when data network switches were introduced, 100-megabit network connections were the fastest. Around 1996, 1-gigabit switches were introduced as a result of technological advancement. Around 2002, 10-gigabit switches entered the market, and circa 2010, 100-gigabit switches began to appear. These major speed improvements for high-end models have also been associated with better cost efficiency for lower-end models with slower connections. Thus, separate upgrade cycles in all sections of the market were triggered. Currently, a major upgrade cycle is progressing, with many companies and organizations upgrading switch models to faster line speeds.

<sup>5</sup> See, for example, Ariel Pakes, "A reconsideration of hedonic price indexes with an application to PCs," *American Economic Review*, vol. 93, no. 5, December 2003, pp. 1578–1596.

<sup>6</sup> Andrew Court first proposed using hedonic regressions in price indexes in 1939; see Court, "Hedonic price indexes with automotive examples," in *The dynamics of automobile demand* (New York: General Motors Corp., 1939), pp. 98–119. Later, Zvi Griliches championed the approach; see, for example, Griliches, "Hedonic price indexes for automobiles: an econometric of quality change," in *The price statistics of the federal government* (Cambridge, MA: National Bureau of Economic Research, 1961), pp. 173–196. The selective use of hedonics began in 1987 for the U.S. Consumer Price Index and in 1991 for the PPI. For a history of hedonic methods in U.S. official statistics, see Brent R. Moulton, "The expanding role of hedonic methods in the official statistics of the United States," BEA Papers 0018 (U.S. Bureau of Economic Analysis, 2001). For current uses of hedonics in the PPI, see "Hedonic models in the Producer Price Index," *Producer Price Indexes* (U.S. Bureau of Labor Statistics, June 2011), <https://www.bls.gov/ppi/ppicomqa.htm>.

For examples of recent developments, see “PPI introduces hedonic quality adjustment for Internet access indexes,” *Producer Price Indexes* (U.S. Bureau of Labor Statistics, February 2017), <https://www.bls.gov/ppi/broadbandhedonicmodel.htm>; and “PPI introduces hedonic price estimation for integrated microcircuits indexes,” *Producer Price Indexes* (U.S. Bureau of Labor Statistics, February 2018), <https://www.bls.gov/ppi/microcircuitshedonics.htm>.

<sup>7</sup> Hedonic regressions model a product’s price as a function of its own characteristics. In models of imperfect competition, equilibrium prices depend on the prices and characteristics of competing goods. In Hotelling-style models, for example, the location of the nearest competitor influences equilibrium prices. The substitutability of competing products likely affects prices and likely varies with a product’s observable characteristics, but it is an omitted variable in the hedonic model, violating assumptions of zero conditional mean for the error term. For settings that can be approximated by perfect competition, the further assumptions needed for interpreting hedonic coefficients as welfare measures are given in Robert C. Feenstra, “Exact hedonic price indexes,” *The Review of Economics and Statistics*, vol. 77, no. 4, 1995, pp. 634–653.

#### RELATED CONTENT

##### Related Articles

[Comparing NAICS-based Producer Price Index industry net output data and International Price Program import data](#), *Monthly Labor Review*, March 2018.

[Crude petroleum prices and retail fuel margins: an empirical examination](#), *Monthly Labor Review*, January 2018.

[New PPI net inputs to industry indexes](#), *Monthly Labor Review*, October 2015.

[Quality adjustment in the Producer Price Index](#), *Producer Price Indexes*, August 2014.

[Analyzing price movements within the Producer Price Index Final Demand–Intermediate Demand aggregation system](#), *Monthly Labor Review*, January 2014.

[Comparing new final-demand producer price indexes with other government price indexes](#), *Monthly Labor Review*, January 2014.

##### Related Subjects

[Producer price index](#) | [Prices and Spending](#) | [Statistical programs and methods](#) | [Technology](#) | [Prices](#)