Monthly Labor Review mlr BLS

June 2018

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.

Brian Adams
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

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

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.[1] 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.

Earlier works have used cross-validation in comparisons of hedonic specifications.[2] 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).[3] 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.[4]

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.[5]

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).[6] 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 β_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.[7] 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)

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)
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)	—
Ports x MAC addresses (thousands)	—	—	-0.447* (0.221)	—
Series indicators	Yes	Yes	Yes	Yes
Observations	480	480	480	

See footnotes at end of table.