


A new approach for quality-adjusting PPI microprocessors

The Producer Price Index (PPI) for microprocessors has declined more slowly since 2010 than it did previously. This shift in microprocessors inflation occurred at the same time that a major manufacturer changed its pricing behavior. With these changes, we must explore a different approach to the matched-model methodology that has been used for microprocessors. Hedonic quality adjustment can account for changes in the quality (characteristics) of products that cannot be captured with the use of a matched model. We look at the implementation of a time dummy hedonic model in a recent article and evaluate its suitability to PPI microprocessors. We then develop our own time dummy hedonic model for microprocessors. The choice of characteristics to include in a model is crucial, because the characteristics help determine the inflation rate the model estimates. We turn to statistical learning techniques to select characteristics for our model. We use our model to construct counterfactual PPI indexes for 2009–17 to determine what the effect of using our model would have been.

Measuring the price changes of “high-tech” goods can be challenging. These types of goods see rapid technological improvements, which must be valued so that inflation is accurately estimated. Between 2010 and 2014, the Producer Price Index (PPI) for microprocessors has seen a change in its rate of decline. From 2000 to 2009, the index fell on average 38.14 percent a year.^[1] From 2010 to 2014, the index decreased only 4.21 percent a year. This change in index behavior has spurred interest in alternative methods that account for technological change in microprocessors.

The PPI measures only “pure” price change on the basis of market factors. It must exclude any price change or portion of a price change that is due to a change in the characteristics of a product. A change in the characteristics of a product is also called a quality change. The PPI must separate this pure price change from the quality change.



Steven D. Sawyer

sawyer.steven@bls.gov

Steven D. Sawyer is an economist in the Office of Prices and Living Conditions, U.S. Bureau of Labor Statistics.

Alvin So

alvin.k.so@gmail.com

Alvin So was formerly an economist of the U.S. Bureau of Labor Statistics.

For a product, such as a table, one can easily measure and account for the quality change from switching to oak from maple construction by removing the value of the switch from the price change. This procedure is known as quality adjustment. However, quality adjusting for microprocessor characteristics is more difficult because of their technological complexity. This challenge in properly accounting for the quality change has led to debate about whether the PPI for microprocessors is biased because it is not properly accounting for quality change.

This debate has guided our efforts to develop new quality-adjustment methods for microprocessors.^[2] Some of the quality-adjustment methods proposed by other researchers have been helpful in guiding the general direction of our approach to quality adjustment, which we then refine and develop to suit the needs of the PPI for microprocessors. These methods involve using a type of statistical model for quality adjustment that has never been deployed in a PPI index before. In addition, we use statistical learning techniques in the development of our quality-adjustment method, which have never been used in an official price index before.^[3]

In this article, we will (1) discuss the technical reasons for the debate over the possible magnitude of the bias of the PPI for microprocessors, (2) reexamine some of the results from a study claiming substantial bias in the microprocessors index, and (3) present a quality-adjustment method that both uses new methods and is developed with new methods.

The PPI for microprocessors is a matched-model index. Theoretically, a matched-model index tracks price changes from period to period for the same set of products. However, producers typically stop producing products after a certain amount of time and introduce new products with different characteristics; this scenario creates a problem of price comparison between different products. Before 2009, when new microprocessors were introduced, the prices of existing microprocessors would decline. In this case, a matched-model index would still capture price changes caused by factors such as technological improvements. The market is effectively valuing the technological change. After 2009, however, when new products were introduced, the prices of existing products would usually remain unchanged. This slower rate of decline could be interpreted as a slowdown of technological innovation in the semiconductor industry and as a signal for the end of Moore's Law.

Moore's Law states that the number of transistors on a chip doubles every 2 years, which is substantial because it has driven "feature-size reduction (scaling) that leads to better performance and cost reduction."^[4] This combination of decreasing cost per transistor and increasing performance was responsible for the large declines in the microprocessors index before 2009.^[5] However, at very small feature sizes of microprocessors, which companies began to manufacture in the 2005–07 timeframe, power usage for a given area of a microprocessor began to increase.^[6] Increasing power usage per area makes increasing performance more difficult. While improvements in fabrication technology can address the challenge of increasing power usage per area, at least to some extent, it is clear that the feature-size reduction no longer delivers the same decreases in cost and increases in performance as it did previously.^[7]

Although the slowdown of Moore's Law is debatable, without question, microprocessors have changed substantially since 2009, from both a technological and pricing viewpoint.^[8] These changes have led to discussions about the PPI matched-model methodology for the microprocessors index and whether the index is showing a slower rate of quality increase than a hedonic approach would show. Researchers have used hedonic modeling methods to generate estimates of quality-adjusted price declines for computers.^[9] Since computers and

microprocessors are both high-tech goods and share many similarities, the hedonic methods used for computers are readily applicable to microprocessors.

In a 2018 article, Byrne, Oliner, and Sichel, henceforth referred to as BOS, use a hedonic regression applied to microprocessors to construct a quality-adjusted index of price change for Intel desktop microprocessors.^[10] BOS examine a number of different model specifications that include both physical characteristics and performance benchmarks. Their preferred hedonic index shows a 42-percent price decline per year between 2009 and 2013.

Data and methods

To understand the 42-percent price change, we reconstruct the model presented in the BOS article.^[11] We use Intel processors and prices to recreate a comparable dataset to do this analysis.^[12] BOS uses a time-dummy hedonic model in their analysis of microprocessor prices. A time-dummy hedonic model uses a panel dataset to track price change over time. The BOS article uses 2 adjacent overlapping years for each of their panels. For instance, two of their panels are 2009–10 and 2010–11. Their preferred specification uses a time dummy and one other regressor, the log of a performance benchmark (SPEC speed [Standard Performance Evaluation Corporation speed]), with log price as the dependent variable.^[13] The coefficient on the time-dummy variable shows the price change between the two periods in the panel that is not explained by the other independent variables. The time-dummy coefficient can be used to calculate an annual inflation rate. The SPEC speed measures the performance of a microprocessor by calculating how long a microprocessor takes to run a suite of software. (See appendix A for more information on the SPEC benchmarks.)

In putting together their dataset, BOS identified two possible problems.^[14] The first possible problem is that Intel’s “posted prices do not represent true transactions prices because Intel offers progressively larger discounts to selected purchasers as models age.” The second possible problem is that Intel’s prices are unweighted, which “would put too much weight on price observations for which there were few transactions,” especially for older microprocessors. BOS contend that a newly introduced microprocessor will generate much more revenue than a microprocessor that is several years old, but an unweighted dataset will give them equal importance. Their solution for both of these possible problems was to “use the first four quarterly prices for each model (or fewer prices if the model is in the market for less than a year) and refer to this as the ‘early-price’ hedonic regression.”^[15] They aggregate these quarterly prices into yearly panels. For example, a microprocessor introduced in the third quarter of 2009 would have two observations for 2009 (third and fourth quarters) and two for 2010 (first and second quarters).

After recreating the BOS results, we explore the use of additional characteristics. The data we use in our following research are derived from Intel’s publicly available price sheets. Details on microprocessor characteristics were obtained from Intel’s ARK website.^[16] In addition to SPEC “speed” (and SPEC “rate,” which we discuss in the next section), we look at

- cores—a hardware term that describes the number of independent central processing units (CPUs) on a single computing component (die or chip);
- threads—a software term for the basic ordered sequence of instructions that can be passed through or processed by a single CPU core;

- thermal design power (TDP)—the average power, in watts, that the microprocessor dissipates when operating at base frequency with all cores active under an Intel-defined high-complexity workload;
- base frequency—the rate at which the microprocessor’s transistors open and close (The microprocessor base frequency is the operating point at which TDP is defined. Frequency is measured in gigahertz, or billion cycles per second.);
- turbo frequency—the maximum single-core frequency at which the microprocessor is capable of operating using Intel Turbo Boost Technology;
- cache—an area of fast memory located on the microprocessor (Intel’s Smart Cache refers to the architecture that allows all cores to dynamically share access to the last level cache); and
- graphics—microprocessors that have an integrated graphics processing unit (GPU).[\[17\]](#)

All regressions in this article have log price as the dependent variable and include a time-dummy variable.

Criticisms of PPI microprocessors

The 42-percent decline per year in microprocessor inflation from the preferred BOS model is the result of using only a single regressor, the log of the SPEC speed performance benchmark. When additional characteristics are added to the regression, the rate of price decline becomes much smaller.

As already noted, BOS extensively examined regressions that included characteristics of microprocessors. Arguing that performance, as measured by benchmarks, is the most important aspect of a microprocessor, they focused on regressions that used only performance benchmarks.

PPI is interested in finding the best way to measure the value of the changes in microprocessors. We see nothing wrong in considering performance benchmarks to do this, but we believe that the regressions using these benchmarks must be able to at least approach the model evaluation criteria of other possible specifications, such as those that use characteristics. If regressions that use performance benchmarks exclusively cannot do this, then we cannot confidently use them for quality adjustment in the official PPI.

We replicated the BOS result using publicly available Intel data from 2009 to 2013 and applying their specification on early prices (mentioned earlier). We obtain an average annual price decline of 45.11 percent with an average adjusted R^2 of 0.6516. This result is comparable to the result obtained by BOS. It serves as a check on data compatibility and is shown in table 1.

Table 1. Log performance (SPEC speed) model on first four quarterly early prices, 2009–13

Variable	2009–10	2010–11	2011–12	2012–13
Year dummy	–0.5779*	–1.0819*	–0.4745*	–0.2651*
Standard error	–0.1163	–0.0987	–0.0604	–0.0396
Log performance (speed)	2.9175*	2.8908*	2.7619*	2.9234*
Standard error	–0.2586	–0.2304	–0.1497	–0.0956
Observations (year 1, year 2)	81 (22, 59)	159 (59, 100)	194 (100, 94)	204 (94, 110)
Adjusted R^2	0.6132	0.5247	0.6772	0.8604

See footnotes at end of table.

*Significant at the 5-percent level.

Note: SPEC = Standard Performance Evaluation Corporation.

Source: U.S. Bureau of Labor Statistics.

The SPEC speed benchmark is designed to measure the speed of a core. According to Cisco,

This benchmark runs the integer or floating-point workloads end to end in a serial fashion, calculating a score based on the amount of time needed to complete the test. This test is meant to represent a single-threaded application or an application that is not designed to run in a multicore system.[\[18\]](#)

The SPEC rate benchmark is designed to measure the

throughput of a machine running simultaneous tasks over time. This benchmark runs several instances of the workloads at once and calculates a score based on how much work was done over a certain amount of time. This test represents a multithreaded application designed to run on a modern multicore system.[\[19\]](#)

The difference between the speed and rate tests can be illustrated by an example. Consider the three microprocessors (i3-6100, i7-4790K, and i7-5960X) in table 2, all with speed scores in the low- to mid-70s. As we can see, the SPEC speed benchmarks are similar for these three microprocessors, while the SPEC rate score better reflects the increase in performance with the increase in price from the i3-6100 to the i7-5960X. The disparity in rate scores is due mostly to different numbers of physical cores, which are two, four, and eight, respectively. The cache is another factor, increasing at even greater rate than the core count: 3, 8, and 20 megabytes. Intel, of course, charges for these features, as the large price spread illustrates. Running a regression with only the log of the SPEC rate benchmark drops the average annual price decline to 29.92 percent, with an average adjusted R^2 of 0.8356. This lower rate of price decline is because SPEC rate controls characteristics of microprocessors better than SPEC speed does. The results are shown in table 3.

Table 2. SPEC speed benchmark and SPEC rate benchmark and price comparisons of three microprocessor models

Attribute	i3-6100	i7-4790K	i7-5960X
SPEC speed benchmark	73	71	72
SPEC rate benchmark	132	183	328
Price	\$117	\$339	\$999

Note: SPEC = Standard Performance Evaluation Corporation.
Sources: SPEC.org and Intel ARK.

Table 3. Log performance (SPEC rate) model on first four quarterly early prices, 2009–13

Variable	2009–10	2010–11	2011–12	2012–13
Year dummy	–0.2745*	–0.7207*	–0.3149*	–0.1122*
Standard error	–0.0846	–0.0548	–0.0381	–0.0251
Log performance (rate)	1.5321*	1.7151*	1.6758*	1.6216*
Standard error	–0.0934	–0.0725	–0.0509	–0.0322

See footnotes at end of table.

Table 3. Log performance (SPEC rate) model on first four quarterly early prices, 2009–13

Variable	2009–10	2010–11	2011–12	2012–13
Observations (year 1, year 2)	81 (22, 59)	159 (59, 100)	194 (100, 94)	204 (94, 110)
Adjusted R^2	0.7714	0.7921	0.8528	0.9261
*Significant at the 5-percent level.				
Note: SPEC = Standard Performance Evaluation Corporation.				
Source: U.S. Bureau of Labor Statistics.				

However, there are reasons to not solely rely on SPEC benchmarks. SPEC has not registered the graphical improvements to many Intel desktop microprocessors in the last decade. In more recent years, Intel has been integrating a GPU onto their microprocessors. For this reason, a hedonic regression examining Intel CPUs should include controls for graphics. The performance of a GPU can be gauged by the number of execution units it has. The regressor we use for graphics is the log of the number of execution units; if it does not have an onboard GPU, we assign a zero.

Additional controls can also be added to distinguish microprocessors with equivalent benchmark performances but differing in operating frequencies, power consumption, thread counts, and cache. The nine-regressor model seen in table 4 includes log cores, log threads, log base frequency, log turbo frequency, log cache per core, log TDP, log graphical statistic, and the single and multicore SPEC benchmarks.

Table 4. Nine-regressor model on first four quarterly early prices, 2009–13

Variable	2009–10	2010–11	2011–12	2012–13
Year dummy	–0.1234*	–0.1089	–0.1259*	–0.0091
Standard error	–0.055	–0.0743	–0.0213	–0.0145
Log performance (speed)	–1.4408*	–2.0935*	–1.8317*	–2.4293*
Standard error	–0.5983	–0.4075	–0.2712	–0.2625
Log performance (rate)	2.1371*	1.9746*	1.4987*	2.1898*
Standard error	–0.6567	–0.416	–0.2591	–0.2238
Log cores	0.0857	0.4079*	0.3011*	0.0412
Standard error	–0.3277	–0.1831	–0.127	–0.1093
Log threads	–0.084	–0.0087	0.2366*	0.1648*
Standard error	–0.1874	–0.1117	–0.0544	–0.0459
Log base frequency	5.104*	2.023*	0.8159*	0.8216*
Standard error	–0.5632	–0.2824	–0.1728	–0.1433
Log turbo frequency	–0.1071	0.8572*	0.9692*	0.9996*
Standard error	–0.5408	–0.213	–0.1871	–0.1569
Log (cache/cores)	0.0077	0.4539*	0.5465*	0.487*
Standard error	–0.2224	–0.134	–0.0805	–0.0682
Log thermal design power	–0.5135*	–0.9042*	–0.448*	–0.4396*
Standard error	–0.2238	–0.0844	–0.0485	–0.0415
Log graphics	–0.2262*	–0.1552*	–0.1538*	–0.0302*
Standard error	–0.0343	–0.0229	–0.0162	–0.0152
Observations (year 1, year 2)	81 (22, 59)	159 (59, 100)	194 (100, 94)	204 (94, 110)
Adjusted R^2	0.947	0.937	0.977	0.9809

See footnotes at end of table.

*Significant at the 5-percent level.

Source: U.S. Bureau of Labor Statistics.

With the nine-regressor specification, the average annual price decline slips to 8.77 percent and the average adjusted R^2 rises to 0.9605. This rate of price decline is near the 12-percent to 14-percent declines BOS calculated with similar specifications.

We perform a standard F-test (where β = the coefficient on a variable and ε = the error term in an equation) on the two nested models, the restricted model with just the time dummy and SPEC speed and the unrestricted model, which include all our additional characteristics:

$$\log \text{ price} = \beta_0 + \beta_1 \text{ Year Dummy} + \beta_2 \log \text{ Performance ("speed")} + \varepsilon$$

$$\log \text{ price} = \beta_0 + \beta_1 \text{ Year Dummy} + \beta_2 \log \text{ Performance ("speed")} + \beta_3 \log \text{ Performance ("rate")}$$

$$+ \beta_4 \log \text{ cores} + \beta_5 \log \text{ Threads} + \beta_6 \log \text{ Base Frequency} + \beta_7 \log \text{ Turbo Frequency}$$

$$+ \beta_8 \log \text{ cache/cores} + \beta_9 \log \text{ TDP} + \beta_{10} \log \text{ Graphics} + \varepsilon$$

For every panel from 2009–10 through 2012–13, we test the null hypothesis (H_0) that our restricted parameters are equal to zero against the alternative hypothesis (H_A) that at least one parameter does not equal zero:

$$H_0: \beta_3 = \beta_4 = \beta_5 = \beta_6 = \beta_7 = \beta_8 = \beta_9 = \beta_{10} = 0$$

$$H_A: \text{at least one } \beta \neq 0$$

From the results of the equation (see text table), the null hypothesis is rejected in all the periods even at the 1-percent level. This result confirms our regression outputs in table 4, in which most of our additional characteristics are significant at the 5-percent level from 2009 to 2013.


F-statistic, 2009–13

Value	2009–10	2010–11	2011–12	2012–13
F-statistic	62.43	128.67	343.84	210.09

Given the statistical significance of variables excluded by BOS, their model may be subject to omitted variable bias. The BOS estimate of inflation is not robust to changes in specification. Indeed, inflation varies greatly with specification changes. This finding raises the question of how different model specifications perform with respect to one another and with respect to the time-dummy coefficient that represents the annual price decline for microprocessors. To provide perspective on this question, we create every possible subset of the nine regressors, which gives 512 different linear regression specifications. We then take these 512 specifications and estimate them for all six of the 2-year panels. Finally, for every one of the 512 specifications, we calculate its average annual rate of inflation and average adjusted R^2 across the six 2-year panels. We get an average annual price decline of 14.35 percent. This rate of inflation is much lower than the model with only the log of SPEC speed as a regressor (the BOS specification), which has a 45.11-percent annual rate of decline, but it is greater than the official PPI for microprocessors, which had an annual rate of decline of 6.6 percent.

The bottom section of figure 1 arrays the 512 models by average adjusted R^2 and average annual price decline. Most cluster above the average adjusted R^2 of 0.6516 of the BOS specification. More important is the range of differences in average annual price declines relative to the BOS result, which appears to be an outlier. This divergent result is even more clearly seen in the top section histogram of figure 1. The blue dashed line shows the average price decline of 14.35 percent for all 512 models. The dotted red line shows the 45.11-percent price decline using the BOS specification, which is a larger rate of decline than that for every other possible model. This result lends evidence that the BOS specification does not include important variables and that it may be an outlier.


Figure 1. Average annual inflation from all 512 models for overlapping 2-year time dummy models, 2009–13



Hover over chart to view data.

Note: Dotted red line at -45.11 percent represents preferred Byrne, Oliner, and Sichel specification. Dashed blue line at -14.35 percent represents average percent change for all models.

Source: U.S. Bureau of Labor Statistics.



Hover over chart to view data.

Note: Dotted red line at -45.11 percent represents preferred Byrne, Oliner, and Sichel specification. Dashed blue line at -14.35 percent represents average percent change for all models.

Source: U.S. Bureau of Labor Statistics.

The annual deflation results obtained from our data are in line with the anecdotal evidence reported by industry observers. For instance, an article on the AnandTech website, detailing sixth-generation microprocessors, states that the increase was around 25 percent, *in total*, in performance between 2011 and 2015.^[20]

From our findings and the results of real-world performance testing from the AnandTech website, we find little empirical support that the 42-percent annual price decline from the preferred BOS model accurately represents the microprocessor industry between 2009 and 2013. However, the BOS criticism of the PPI microprocessors has sparked an evaluation of the PPI semiconductor index methodology and an examination of alternative quality-adjustment methods that PPI could implement.

Constructing a hedonic model for PPI microprocessors

PPI has studied quality-adjusting microprocessors in the past, most notably in an article by Michael Holdway, who looked at using hedonic quality adjustment for microprocessors.^[21] The article concluded that a hedonic model might be infeasible “if technology redefines characteristics over time” because “implicit prices for characteristics may be difficult to interpret.”^[22] The article then devised a method of valuing quality change using performance benchmarks. The main drawback of this method was that estimates of price change were sensitive to the choice of replacement microprocessors.

We find that the BOS article lays out a very useful framework for developing a hedonic model appropriate for use in the PPI. The BOS time-dummy approach solves some of the problems highlighted by the Holdway article because it does not require using the coefficients on the characteristics variables in the model or selecting replacement microprocessors. We believe the preferred BOS specification is not ideal for the needs of the PPI, and we explore methods for picking specifications that can include microprocessors characteristic variables.

Before we look at estimating a model for use in the PPI for microprocessors, we need to create a different dataset than the one BOS used. Because we are using a different dataset, note that any estimates of microprocessor price decline we calculated in this section are not comparable with estimates of price decline calculated in the previous section, in which we looked at the BOS model. We look at microprocessors from 2009 to 2017 and estimate models with overlapping two-period quarterly data. PPI indexes are published monthly, which means models need to be updated more often than annually. Since we are using quarterly data, we cannot use BOS early prices. Instead, we only include microprocessors introduced within a certain interval from a given quarter. In determining what that interval should be, looking at the reasoning behind the BOS early prices is helpful.

As stated earlier, BOS use early prices to address two possible problems. The first possible problem of posted prices not equaling transaction prices for older microprocessors has been examined by Kenneth Flamm.^[23] For the retail microprocessors market, which is 20 percent of Intel’s microprocessors sales by volume, Flamm found “no evidence to support the suggestion that there was some structural change after 2006 in the relationship between observed Intel list price and observed retail market prices.”^[24] Thus, *at least* some of Intel’s sales for older microprocessors are being accurately represented by its posted price list. BOS note, though, that the posted price list does not give any information on discounts off of the list price that Intel may be giving its largest customers, which make up most of its sales. However, we believe Flamm’s findings give us some flexibility on how long a reasonable interval we can set for including older microprocessors in our dataset.

The second problem that early prices address is that after a microprocessor is introduced, its sales peak and then decline. But important to remember is that Intel will continue to sell a microprocessor even after introducing a newer, more technologically advanced version of it is introduced. As mentioned earlier, the PPI tracks the entire production of companies, not just newly introduced products. Furthermore, when Intel introduces a microprocessor, shipments typically start off low and increase for several months before peaking.

For some microprocessors geared toward businesses, Intel guarantees that these microprocessors will be available and supported for 15 months.^[25] If Intel is actively supporting a microprocessor, one can reasonably assume that it should still be selling at a substantial volume. For instance, when a business replaces employees' computers, it usually does so over several months. This extended rollout means that the business will continue to purchase computers with the same microprocessor, even if newer microprocessors are introduced, in order to have a stable configuration. We use this 15-month interval when building our dataset. For example, for first quarter 2015, we include all microprocessors that were introduced from fourth quarter 2013 or later and that are still being sold in the first quarter 2015. This interval is similar to the first four quarterly prices that BOS used.

Several benchmarks are available for testing the performance of microprocessors. Some benchmarks specifically test the performance of rendering and encoding and others test the overall performance of a system. Instead of using the SPEC performance benchmarks, we use the PassMark CPU performance benchmark. (See appendix A for more information on PassMark.) We choose PassMark mainly because of the ease of obtaining the data from its website, which helps in the production nature of the PPI. PassMark also is a comprehensive benchmark that is able to capture the performance gain from multicore microprocessors, as table 5 shows.

Table 5. SPEC speed, SPEC rate, and PassMark benchmarks and price comparisons of microprocessor models

Variable	i3-6100	i7-4790K	i7-5960X
SPEC speed benchmark	73	71	72
SPEC rate benchmark	132	183	328
PassMark benchmark	5454	11,184	15,972
Price	\$117	\$339	\$999


Note: SPEC = Standard Performance Evaluation Corporation.

Sources: SPEC.org, PassMark.com, Intel ARK.

In the first section of this article, we questioned the BOS model because it had a different inflation rate than most of the models with a much higher adjusted R^2 . A first step in evaluating our quarterly data is reproducing figure 1 with them. Note that we constrain the log PassMark variable to be in every model we estimate with our quarterly data because we know for certain that microprocessor performance is increasing over time. We also think that the PassMark benchmark may be capable of accounting for improvements to microprocessors that are not associated with changes in any of the characteristics variables. Microprocessors are complex products with many attributes, so estimating a model that includes their every aspect is not possible. Since the PassMark benchmark shows total microprocessor performance, it may show quality improvements to microprocessors caused by changes to the characteristics we cannot include in our models.

The blue dashed line in figure 2 shows the average annual price decline of 12.74 percent. Just as we saw earlier, a wide variety of possible inflation rates exist. This finding presents us with the problem of selecting our preferred model.

Figure 2. Average annual inflation and average adjusted R^2 from all 128 models for overlapping two-quarter time dummy models, 2009–17



While no definitive method of specification selection exists, the field of statistical learning offers techniques for objectively evaluating the performance of different model specifications. In the following excerpt, Hastie and colleagues give a basic overview behind statistical learning methods:

The *generalization* performance of a learning method relates to its prediction capability on independent test data. Assessment of this performance is extremely important in practice, since it guides the choice of learning method or model, and gives us a measure of quality of the ultimately chosen model.[26]

It is important to emphasize that although statistical learning methods look at the predictive power of the model on the dependent variable, they are evaluating the overall quality of the model, including the model specification. We are using these methods to perform model selection by “estimating the performance of different models in order to choose the best one.”[27]

Models that are overfitted can have a high adjusted R^2 while having poor out-of-sample predictive performance. One way to test for overfitting is to use a procedure known as validation. This validation involves splitting a dataset, estimating the model on one part of the dataset (the training set), and then predicting the dependent variable on the other part of the dataset (the validation set) with the use of the coefficients from the model estimated on the training set. The validation-set predictions of the dependent variable are then compared with the actual validation-set values of the dependent variable to give the prediction errors; the prediction errors are squared and then averaged, which gives the mean squared error (MSE). The lower the MSE is, the better the predictive performance of the model.

We use two different techniques to select model specifications. If we can estimate similar rates of inflation from specifications chosen using different techniques, this consistency would be a sign that we are selecting a reasonable model. With both techniques, we can pick different specifications for every panel. These techniques address Holdway’s criticism that because the characteristics of microprocessors change in their importance over time, a model with a fixed specification may become misspecified over time.[28]

The first method we use is the Bayesian information criterion (BIC). Information criteria, in general, can be defined as “a goodness-of-fit term plus a penalty to control overfitting.”[29] BIC is a traditional technique for evaluating model specifications, but it can also be used for estimating the validation-set MSE of a model.[30] It does this by treating the entire dataset as the training set, and the goodness-of-fit measure calculated for the model estimated on the entire dataset is penalized in order to estimate what the MSE would be on a validation set. The main risk of using BIC for model selection is underfitting, which would be selecting fewer characteristics than the best model that can be supported by the data.[31] Because many characteristics of microprocessors are correlated, there is a tendency for models with different numbers of characteristics to have similar measures of performance. When this happens, BIC tends to pick the model with the least number of characteristics. When we calculate the BIC for all possible models for every two-quarter overlapping panel and then select the model with the minimum BIC for every period, we get an average annual price decline of 10.69 percent between the first quarter 2009 and the third quarter 2017.

Tables 6 and 7 show the models chosen by the minimum BIC for each period (see appendixes B and C for additional information on the BIC and the MSE, respectively). The “Selected variables” in the tables do not include log PassMark or the time dummy since every model contains both of those variables.

Table 6. Selected models by first quarter 2009 to third quarter 2017 with use of minimum BIC

Time of model	Inflation	Selected variables*	Adjusted R^2	Observations
09Q1–09Q2	–0.0355	6	0.9816	45
09Q2–09Q3	–0.0618	6	0.9880	39
09Q3–09Q4	–0.0659	5	0.9672	41
09Q4–10Q1	–0.1596	1	0.8997	53
10Q1–10Q2	–0.0247	5	0.9599	61
10Q2–10Q3	–0.0747	5	0.9591	66
10Q3–10Q4	–0.0719	5	0.9520	72
10Q4–11Q1	–0.0617	4	0.9278	87
11Q1–11Q2	–0.0073	4	0.9225	95
11Q2–11Q3	–0.0026	4	0.9240	81
11Q3–11Q4	–0.0349	4	0.9361	84
11Q4–12Q1	–0.0085	4	0.9425	83
12Q1–12Q2	–0.0049	3	0.9588	70
12Q2–12Q3	–0.0760	2	0.9472	79
12Q3–12Q4	–0.0394	5	0.9634	93
12Q4–13Q1	–0.0070	6	0.9762	105
13Q1–13Q2	–0.0048	6	0.9733	92
13Q2–13Q3	0.0276	2	0.9480	92
13Q3–13Q4	–0.0174	3	0.9495	132
13Q4–14Q1	–0.0105	3	0.9565	139
14Q1–14Q2	–0.0208	4	0.9638	144
14Q2–14Q3	–0.0049	5	0.9590	164
14Q3–14Q4	–0.0044	5	0.9564	169
14Q4–15Q1	–0.0381	5	0.9612	132
15Q1–15Q2	0.0059	5	0.9677	97
15Q2–15Q3	0.0238	5	0.9576	104
15Q3–15Q4	0.0084	4	0.9445	97
15Q4–16Q1	–0.0541	4	0.9299	85
16Q1–16Q2	0	5	0.9186	82
16Q2–16Q3	–0.0005	5	0.9399	78
16Q3–16Q4	–0.0463	5	0.9595	71
16Q4–17Q1	–0.0467	5	0.9642	71
17Q1–17Q2	–0.0121	5	0.9560	69
17Q2–17Q3	0.0073	5	0.9480	66

*Variables do not include log PassMark or the time dummy since every model contains both of those variables.

Note: BIC = Bayesian information criterion, Q1 = first quarter, Q2 = second quarter, Q3 = third quarter, and Q4 = fourth quarter.

Source: U.S. Bureau of Labor Statistics.

Table 7. Selected models by first quarter 2009 to third quarter 2017 with use of MSE

Time of model	Inflation	Selected variables*	Adjusted R^2	Observations
09Q1–09Q2	–0.03546	6	0.9816	45
09Q2–09Q3	–0.0618	6	0.9880	39
09Q3–09Q4	–0.0808	3	0.9567	41

See footnotes at end of table.

Table 7. Selected models by first quarter 2009 to third quarter 2017 with use of MSE

Time of model	Inflation	Selected variables*	Adjusted R^2	Observations
09Q4–10Q1	−0.1596	1	0.8997	53
10Q1–10Q2	−0.0247	5	0.9599	61
10Q2–10Q3	−0.0747	5	0.9591	66
10Q3–10Q4	−0.0697	7	0.9542	72
10Q4–11Q1	−0.0651	6	0.9332	87
11Q1–11Q2	−0.0126	6	0.9273	95
11Q2–11Q3	−0.0066	6	0.9278	81
11Q3–11Q4	−0.0349	4	0.9361	84
11Q4–12Q1	−0.0339	6	0.9476	83
12Q1–12Q2	−0.0049	3	0.9588	70
12Q2–12Q3	−0.0399	5	0.9550	79
12Q3–12Q4	−0.0429	7	0.9642	93
12Q4–13Q1	−0.0070	6	0.9762	105
13Q1–13Q2	−0.0048	6	0.9733	92
13Q2–13Q3	0.0253	4	0.9521	92
13Q3–13Q4	−0.0149	1	0.9449	132
13Q4–14Q1	−0.0105	3	0.9565	139
14Q1–14Q2	−0.0232	3	0.9623	144
14Q2–14Q3	0.0046	4	0.9571	164
14Q3–14Q4	−0.0044	5	0.9564	169
14Q4–15Q1	−0.0381	5	0.9612	132
15Q1–15Q2	0.0059	5	0.9677	97
15Q2–15Q3	0.0263	6	0.9589	104
15Q3–15Q4	0.0084	4	0.9445	97
15Q4–16Q1	−0.0436	6	0.9353	85
16Q1–16Q2	0	6	0.9333	82
16Q2–16Q3	0.0027	6	0.9408	78
16Q3–16Q4	−0.0433	6	0.9599	71
16Q4–17Q1	−0.0620	7	0.9660	71
17Q1–17Q2	−0.0078	7	0.9591	69
17Q2–17Q3	−0.0146	2	0.9153	66

*Variables do not include log PassMark or the time dummy since every model contains both of those variables.

Note: MSE = mean squared error, Q1 = first quarter, Q2 = second quarter, Q3 = third quarter, and Q4 = fourth quarter.

Source: U.S. Bureau of Labor Statistics.

Most of the selected models have at least four variables. There is also a decreasing rate of price decline, especially starting in 2011.

Using BIC to select a model specification has several advantages for the PPI. Because BIC is a widely known technique for model evaluation, it makes the process of model selection easy to explain to users of PPI data. These users will have more confidence and trust in our data when our methods are clear and transparent. Using BIC for model specification selection is also quick and easy to implement. This ease of implementation is important if a model will be used in the PPI. A PPI industry analyst would need to re-estimate the model quarterly. A hedonic model can only be used operationally in the PPI if the model can be developed efficiently.

The second model specification selection technique we use is k -fold cross-validation. Just as with validation, which we explained previously, this method objectively evaluates a model by assessing its out-of-sample predictive performance. The disadvantage of validation is that part of the dataset is not used to estimate the model, which increases the variability of the estimated model parameters. With k -fold cross-validation, the dataset is split into k parts. Each of the k parts is held out in turn (the validation set), and the model is estimated on the remaining data (the training set). Then, just as with validation, cross-validation calculates the MSE by making predictions on the validation set by using the model estimated on the training set. This procedure yields k MSEs, which are averaged to produce an overall MSE.

Performing k -fold cross-validation on every possible model is computationally intensive. To reduce the number of models to evaluate, we can first prescreen the dataset. It is important to note that the prescreening is only performed on the training set and not on the full dataset. With our dataset, we have seven possible regressors to choose from (the time dummy and log PassMark are always included). We start by calculating the residual sum of squares (RSS) for every model that contains just one regressor (plus the time dummy and log PassMark). The one-regressor model with the lowest RSS makes it through the prescreening. We then repeat this procedure for every model that contains two regressors (plus the time dummy and log PassMark). The two-regressor model with the lowest RSS makes it through the prescreening. We continue this process, increasing the number of regressors by one each time, until we have prescreened a model that contains six regressors (plus the time dummy and log PassMark). Because we have seven regressors, only one seven-regressor model exists. At the end of the prescreening procedure, we have seven models, which contain from one to seven regressors. We then use cross-validation to calculate the MSE for each of the seven models.^[32] (See appendix D for more details.) Since we use 10-fold cross-validation, we repeat this procedure 10 times.

Splitting the dataset into 10 folds is done randomly. To average the random variation from splitting the dataset, we perform the procedure just discussed 500 times and take the average MSE for each of the seven models.^[33]

Because we are performing this procedure 500 times (which means 5,000 different training sets), different specifications can possibly be selected for models of the same size during the prescreening procedure.^[34] That is, at the end of the procedure, we know the MSE for models with from one to seven regressors, but we do not know which regressors are included in those models. To select the specific regressors, we find which number of regressors in the prescreening step had the lowest MSE and then find the RSS for every model with that number of regressors in the whole dataset. The model with the lowest RSS is then selected.

Only 16 of the 34 models had the same specifications selected by both methods. Despite this finding, the average annual price decline is very close across specifications, with the price change obtained from the cross-validation method declining 10.95 percent and the minimum BIC method declining 10.69 percent.

Displaying the results from the minimum BIC method and the minimum MSE method, table 8 summarizes the inflation rates and the corresponding indexes for the models chosen.

Table 8. Minimum microprocessor inflation rates and indexes, first quarter 2009 to third quarter 2017

Minimum inflation rate & index: 09Q1–10Q4	09Q1–09Q2	09Q2–09Q3	09Q3–09Q4	09Q4–10Q1	10Q1–10Q2	10Q2–10Q3	10Q3–10Q4
---	-----------	-----------	-----------	-----------	-----------	-----------	-----------

See footnotes at end of table.

Table 8. Minimum microprocessor inflation rates and indexes, first quarter 2009 to third quarter 2017


Minimum inflation rate & index: 09Q1–10Q4	09Q1–09Q2	09Q2–09Q3	09Q3–09Q4	09Q4–10Q1	10Q1–10Q2	10Q2–10Q3	10Q3–10Q4
BIC inflation	–0.0355	–0.0618	–0.0659	–0.1596	–0.0247	–0.0747	–0.0719
MSE inflation	–0.0355	–0.0618	–0.0808	–0.1596	–0.0247	–0.0747	–0.0697
BIC index*	96.45	90.49	84.53	71.03	69.28	64.11	59.5
MSE index*	96.45	90.49	83.18	69.9	68.17	63.08	58.69
Minimum inflation rate & index: 10Q4–12Q3	10Q4–11Q1	11Q1–11Q2	11Q2–11Q3	11Q3–11Q4	11Q4–12Q1	12Q1–12Q2	12Q2–12Q3
BIC inflation	–0.0617	–0.0073	–0.0026	–0.0349	–0.0085	–0.0049	–0.0760
Minimum MSE inflation	–0.0651	–0.0126	–0.0066	–0.0349	–0.0339	–0.0049	–0.0399
BIC index*	55.83	55.42	55.28	53.35	52.9	52.64	48.64
MSE index*	54.87	54.18	53.82	51.94	50.18	49.94	47.95
Minimum inflation rate & index: 12Q3–14Q2	12Q3–12Q4	12Q4–13Q1	13Q1–13Q2	13Q2–13Q3	13Q3–13Q4	13Q4–14Q1	14Q1–14Q2
BIC inflation	–0.0394	–0.007	–0.0048	0.0276	–0.0174	–0.0105	–0.0208
MSE inflation	–0.0429	–0.007	–0.0048	0.0253	–0.0149	–0.0105	–0.0232
BIC index*	46.72	46.4	46.18	47.45	46.62	46.13	45.17
MSE index*	45.89	45.57	45.35	46.5	45.81	45.33	44.27
Minimum inflation rate & index: 14Q2–16Q1	14Q2–14Q3	14Q3–14Q4	14Q4–15Q1	15Q1–15Q2	15Q2–15Q3	15Q3–15Q4	15Q4–16Q1
BIC inflation	–0.0049	–0.0044	–0.0381	0.0059	0.0238	0.0084	–0.0541
MSE inflation	0.0046	–0.0044	–0.0381	0.0059	0.0263	0.0084	–0.0436
BIC index*	44.95	44.75	43.05	43.31	44.33	44.71	42.29
MSE index*	44.48	44.28	42.6	42.85	43.98	44.35	42.41
Minimum inflation rate & index: 14Q2–16Q1	16Q1–16Q2	16Q2–16Q3	16Q3–16Q4	16Q4–17Q1	17Q1–17Q2	17Q2–17Q3	—
BIC inflation	0	–0.0005	–0.0463	–0.0467	–0.0121	0.0073	—
MSE inflation	0	0.0027	–0.0433	–0.062	–0.0078	–0.0146	—
BIC index	42.29	42.27	40.31	38.43	37.97	38.24	—
MSE index	42.41	42.53	40.69	38.16	37.87	37.31	—

*Indexes start at 100 in 08Q4–09Q1.

Note: BIC = Bayesian information criterion, MSE = minimum squared error, Q1 = first quarter, Q2 = second quarter, Q3 = third quarter, and Q4 = fourth quarter.


Source: U.S. Bureau of Labor Statistics.

The two different methods of selecting specifications produce indexes that track closely over time. In figure 3, we compare the actual PPI for microprocessors with counterfactual hedonic indexes that adjust desktop microprocessor items with the inflation rates in the previous tables. Note that the PPI for microprocessors stopped publishing in March 2015 because it no longer met publication criteria.



The PPI microprocessors index fell, in total, 30.88 percent, whereas, at the same time, the minimum BIC model fell 38.97 percent and the minimum MSE model fell 39.09 percent. The PPI microprocessors index begins diverging substantially from the counterfactual indexes in the beginning of 2010. This result supports the hypothesis that after 2009, the matched-model methodology was no longer fully accounting for quality change for PPI microprocessors. The hedonic indexes are also very similar, which suggests that our hedonic estimates of price change are consistent, even when selected with the use of two different techniques. Most likely, if the notebook and server microprocessors were also adjusted with hedonic time-dummy models, the rate of decline for the hedonic indexes would be greater. This assumption is based on the inference that the time-dummy coefficients for the hypothetical notebook and server microprocessors models would be similar to the desktop models. Also important to note is that as an item declines in price, its relative revenue weight in an index declines, which lessens the items effect on the index over time.^[35]

In figure 4, we compare the semiconductors index for primary products (PPI code PCU334413334413P),^[36] with counterfactual hedonic indexes that adjust desktop microprocessor items with the inflation rates in table 8. PPI semiconductors primary products has published for the entire time that we calculated the counterfactual indexes. PPI semiconductor primary products aggregates together indexes for different semiconductors products, including microprocessors.



The PPI semiconductors primary products index fell, in total, 29.37 percent while the minimum BIC model index and the minimum MSE model index both fell 30.92 percent. The hedonic indexes again show an overall greater decline than the PPI semiconductors primary products index, though it is minimal. This decline is the result of two factors. First, desktop microprocessors are a smaller share of PPI semiconductor primary products than they are of PPI microprocessors. Second, as an index declines, its relative weight also declines. As the PPI microprocessors index declines faster than other PPI semiconductor indexes, its relative weight in PPI semiconductor primary products declines.^[37]

Starting in January 2018, PPI began using the minimum MSE model to estimate prices for desktop microprocessor items in the official semiconductors index. The minimum MSE model was selected because statistical learning is a developing field, and we hope to refine our methods continually over time.

Conclusion

Challenges in recent years of continuing the historic trend of performance improvements of semiconductors have led to a debate about the rate of microprocessor price declines. Part of this debate, the BOS article, has a preferred model that has a very large rate of price decline in 2009–13. By estimating models with more characteristics, we have shown that the rate of price decline decreases and the goodness-of-fit measures for the models increase.

Drawing on past work done in hedonics, particularly the BOS article, we have devised an alternative to the matched-model methodology used for microprocessors in PPI semiconductors. Our models use overlapping adjacent quarters, a time dummy, and specifications selected with the statistical learning technique of repeated cross-validation. Our microprocessors models are the first time-dummy models used in the PPI and are the first to use a statistical learning technique in any official price index.

The counterfactual indexes we created using our models show that the official PPI for microprocessors was likely understating price declines, but not dramatically so. With the implementation of time-dummy models with a specification selected using statistical learning, the PPI may better represent the price movements of desktop microprocessors.

Acknowledgments

We first and foremost would like to thank Tim Erickson. He spent countless hours with us discussing economics, statistics, statistical learning, and the technical details of microprocessors. The knowledge we gained from these conversations made this article possible. We would also like to thank Kelly McConville, Deanna Bathgate, David Byrne, Stephen Oliner, Daniel Sichel, Ralph Bradley, Peter Uimonen, Brian Adams, and Vincent Russo for their assistance and comments.

Appendix A: SPEC and PassMark benchmarks

Our dataset includes the SPEC (Standard Performance Evaluation Corporation) central processing unit (CPU) 2006 benchmark score for desktop microprocessors.[\[38\]](#) We obtained SPEC benchmarks from the SPEC website in January 2016. Both SPEC speed and SPEC rate benchmarks were collected; the speed metrics are used for measuring the capability of a computer to complete a single task. The SPEC speed metric measures the performance of the system using a single core of the system microprocessor, which is greatly affected by the clock speed of the microprocessor and its cache size. The SPEC rate metric measures the throughput of a machine performing several simultaneous tasks. This metric provides a good overall measure of performance of modern multicore microprocessors. SPEC rate metrics are typically most affected by the number of microprocessor cores on a system. We include the SPEC rate benchmark to our evaluation because a key price-determining characteristic of today's microprocessors is based on the number of cores it has.[\[39\]](#) The SPEC benchmark suite is used primarily by server and workstation manufacturers to performance test their products. Many low-end desktop CPUs do not have SPEC benchmark scores. To better reflect all the desktop CPUs that Intel produces, we turn to PassMark.

PassMark CPU benchmark results are gathered from users' submissions to the PassMark website as well as from PassMark's own internal testing. Its benchmarking software PerformanceTest is available for purchase on the PassMark website. This software conducts eight different tests and then averages the results together to determine the CPUMark rating (a performance measure of the microprocessor) for a system. PassMark runs one simultaneous CPU test for every logical CPU, physical CPU core, or physical CPU package. The eight tests include an Integer Maths test, a Compression Test, a Prime Number Test, an Encryption Test, a Floating-Point Math Test, an Extended Instructions Test, a String Sorting test, and a Physics Test.[\[40\]](#)

Appendix B: Selected quarterly models of minimum Bayesian information criterion

Table B-1. Minimum BIC selected models, first quarter 2009 to fourth quarter 2010

Variable	09Q1–09Q2	09Q2–09Q3	09Q3–09Q4	09Q4–10Q1	10Q1–10Q2	10Q2–10Q3	10Q3–10Q4
Quarter time dummy	–0.0361	–0.0638*	–0.0682	–0.1739*	–0.0250	–0.0776	–0.0746
	–0.0304	–0.0267	–0.0503	–0.0736	–0.0427	–0.0421	–0.0408
Log PassMark	5.6791*	6.3611*	8.2641*	1.1733*	–0.2040	–0.3234	–0.2659
	–1.2062	–0.9380	–0.6484	–0.1223	–0.2432	–0.2445	–0.2000
Log cores	–9.3942*	–9.4896*	–4.7662*	—	1.0789*	0.9855*	1.1182*
	–0.5407	–0.6929	–0.4859	—	–0.1760	–0.1801	–0.1671
Log threads	6.1792*	5.7827*	—	—	0.2816	0.4226*	0.5758*
	–0.9694	–0.6473	—	—	–0.1597	–0.1397	–0.0937
Log base frequency	50.2257*	47.7293*	4.9030*	—	1.5867*	1.6701*	1.8986*
	–4.4902	–3.6284	–0.9733	—	–0.5289	–0.4259	–0.3350
Log turbo frequency	–54.1835*	–52.4618*	–10.6041*	1.4998*	2.6945*	2.6959*	1.7294*
	–4.0868	–3.807	–1.3089	–0.3227	–0.3367	–0.3377	–0.3080
Log (cache/cores)	–0.4955*	–0.5179*	–0.8236*	—	—	—	0.2214*
	–0.2209	–0.1287	–0.1350	—	—	—	–0.0859
Log thermal design power	–2.4443*	–2.8018*	–4.0385*	—	0.6121*	0.6825*	—
	–0.2469	–0.2191	–0.4619	—	–0.2566	–0.2743	—
Log graphics	—	—	—	—	—	—	—
Observations	45	39	41	53	61	66	72
Adjusted R^2	0.9816	0.9880	0.9672	0.8997	0.9599	0.9591	0.9520
BIC	–38.8	–44.26	–3.39	28.85	–10.02	–11.98	–8.81

*Significant at the 5-percent level.

Note: BIC = Bayesian information criterion, Q1 = first quarter, Q2 = second quarter, Q3 = third quarter, and Q4 = fourth quarter.

Source: U.S. Bureau of Labor Statistics.

Table B-2. Minimum BIC selected models, fourth quarter 2010 to third quarter 2012

Variable	10Q4–11Q1	11Q1–11Q2	11Q2–11Q3	11Q3–11Q4	11Q4–12Q1	12Q1–12Q2	12Q2–12Q3
Quarter time dummy	–0.0637	–0.0073	–0.0026	–0.0355	–0.0085	–0.0049	–0.0790*
	–0.0438	–0.0406	–0.0451	–0.0383	–0.0403	–0.036	–0.0388
Log PassMark	–0.0686	–0.1726	–0.1300	–0.1971	–0.2227	1.7098*	1.1021*
	–0.1152	–0.0891	–0.0971	–0.1234	–0.1506	–0.1048	–0.0546
Log cores	0.5841*	0.5441*	0.4741*	0.5298*	0.6028*	–0.4930*	—
	–0.0996	–0.0845	–0.0906	–0.1009	–0.1050	–0.0841	—
Log threads	0.6276*	0.7343*	0.7519*	0.8064*	0.7705*	—	—
	–0.0901	–0.0844	–0.0821	–0.0813	–0.0891	—	—
Log base frequency	—	—	—	—	—	–1.6054*	–0.8222*
	—	—	—	—	—	–0.1860	–0.2057
Log turbo frequency	1.6939*	1.5227*	1.2316*	0.7053*	0.6420*	—	—
	–0.3749	–0.3265	–0.2765	–0.2529	–0.2347	—	—
Log (cache/cores)	0.3651*	0.3938*	0.3992*	0.4276*	0.4629*	0.4447*	0.8269*
	–0.0763	–0.0700	–0.0760	–0.0932	–0.1161	–0.2111	–0.1731
Log thermal design power	—	—	—	—	—	—	—

See footnotes at end of table.

Table B-2. Minimum BIC selected models, fourth quarter 2010 to third quarter 2012

Variable	10Q4–11Q1	11Q1–11Q2	11Q2–11Q3	11Q3–11Q4	11Q4–12Q1	12Q1–12Q2	12Q2–12Q3
Log graphics	—	—	—	—	—	—	—
Observations	87	95	81	84	83	70	80
Adjusted R^2	0.9278	0.9225	0.9240	0.9361	0.9425	0.9588	0.9472
BIC	2.69	–2.18	–1.41	–13.07	–17.54	–36.54	–26.41

*Significant at the 5-percent level.

Note: BIC = Bayesian information criterion, Q1 = first quarter, Q2 = second quarter, Q3 = third quarter, and Q4 = fourth quarter.

Source: U.S. Bureau of Labor Statistics.

Table B-3. Minimum BIC selected models, third quarter 2012 to second quarter 2014

Variable	12Q3–12Q4	12Q4–13Q1	13Q1–13Q2	13Q2–13Q3	13Q3–13Q4	13Q4–14Q1	14Q1–14Q2
Quarter time dummy	–0.0402	–0.0070	–0.0048	0.0272	–0.0176	–0.0106	–0.0210
	–0.0319	–0.0238	–0.0316	–0.0332	–0.0265	–0.0260	–0.0246
Log PassMark	0.5797*	–0.0359	–0.1480	0.9511*	1.0947*	0.6152*	0.0936
	–0.1154	–0.2205	–0.2822	–0.0731	–0.1036	–0.1346	–0.2995
Log cores	—	0.5319*	0.6150*	—	—	0.3212*	0.6307*
	—	–0.138	–0.1735	—	—	–0.0951	–0.1402
Log threads	0.3802*	0.6329*	0.6799*	0.4566*	0.3725*	0.4491*	0.5718*
	–0.0673	–0.0942	–0.1175	–0.0737	–0.0858	–0.0755	–0.1104
Log base frequency	–0.7328*	—	—	—	—	—	—
	–0.2420	—	—	—	—	—	—
Log turbo frequency	0.9030*	1.1779*	1.3342*	—	—	—	0.6335
	–0.3435	–0.3241	–0.3839	—	—	—	–0.3311
Log (cache/cores)	0.6013*	0.6522*	0.6270*	—	—	0.4668*	0.6112*
	–0.1301	–0.0901	–0.1047	—	—	–0.1609	–0.1385
Log thermal design power	—	–0.2494*	–0.2500*	—	–0.1599*	—	—
	—	–0.0682	–0.0800	—	–0.0765	—	—
Log graphics	–0.0701*	–0.0885*	–0.0997*	–0.0650	–0.0637*	—	—
	–0.0277	–0.0322	–0.0326	–0.0331	–0.0278	—	—
Observations	94	105	92	92	132	139	144
Adjusted R^2	0.9634	0.9762	0.9733	0.948	0.9495	0.9565	0.9638
BIC	–53.16	–99.37	–67.51	–52.55	–90.29	–95.41	–111.33

*Significant at the 5-percent level.

Note: BIC = Bayesian information criterion, Q1 = first quarter, Q2 = second quarter, Q3 = third quarter, and Q4 = fourth quarter.

Source: U.S. Bureau of Labor Statistics.

Table B-4. Minimum BIC selected models, second quarter 2014 to first quarter 2016

Variable	14Q2–14Q3	14Q3–14Q4	14Q4–15Q1	15Q1–15Q2	15Q2–15Q3	15Q3–15Q4	15Q4–16Q1
Quarter time dummy	–0.0049	–0.0044	–0.0388	0.0059	0.0235	0.0084	–0.0556
	–0.0219	–0.0223	–0.0270	–0.0257	–0.0273	–0.0330	–0.0352
Log PassMark	–0.0852	0.0537	0.0043	–0.2386	0.4295*	0.2986	0.1016
	–0.2615	–0.2518	–0.2926	–0.2489	–0.1957	–0.1859	–0.1149

See footnotes at end of table.

Table B-4. Minimum BIC selected models, second quarter 2014 to first quarter 2016

Variable	14Q2–14Q3	14Q3–14Q4	14Q4–15Q1	15Q1–15Q2	15Q2–15Q3	15Q3–15Q4	15Q4–16Q1
Log cores	0.7774*	0.7334*	0.7930*	0.9721*	1.0864*	0.6907*	0.3721*
	–0.1214	–0.1141	–0.1240	–0.1019	–0.1246	–0.0706	–0.1301
Log threads	0.5904*	0.5138*	0.4936*	0.5248*	0.2830*	0.5136*	0.6222*
	–0.0909	–0.0836	–0.0934	–0.0930	–0.0810	–0.0979	–0.0678
Log base frequency	—	—	—	—	0.7982*	—	–1.3323*
	—	—	—	—	–0.1556	—	–0.3558
Log turbo frequency	1.0726*	1.0838*	1.1390*	1.2687*	—	0.6776*	2.2484*
	–0.2982	–0.2723	–0.3026	–0.2445	—	–0.2696	–0.4814
Log (cache/cores)	0.5443*	0.4601*	0.5282*	0.7386*	0.4538*	—	—
	–0.1418	–0.1412	–0.1532	–0.1338	–0.1453	—	—
Log thermal design power	–0.1282*	–0.1764*	–0.1997*	–0.2599*	–0.5620*	–0.2187*	—
	–0.0612	–0.0567	–0.0631	–0.0761	–0.0779	–0.0674	—
Log graphics	—	—	—	—	—	—	—
Observations	164	169	132	97	104	100	88
Adjusted R^2	0.9590	0.9564	0.9612	0.9677	0.9576	0.9445	0.9299
BIC	–125.61	–126.85	–95.85	–81.12	–66.69	–46.25	–29.86

*Significant at the 5-percent level.

Note: BIC = Bayesian information criterion, Q1 = first quarter, Q2 = second quarter, Q3 = third quarter, and Q4 = fourth quarter.

Source: U.S. Bureau of Labor Statistics.

Table B-5. Minimum BIC selected models, first quarter 2016 to third quarter 2017

Variable	16Q1–16Q2	16Q2–16Q3	16Q3–16Q4	16Q4–17Q1	17Q1–17Q2	17Q2–17Q3
Quarter time dummy	0	–0.0005	–0.0474	–0.0478	–0.0122	0.0073
	–0.0368	–0.0386	–0.0373	–0.0371	–0.0423	–0.0483
Log PassMark	–0.0901	–0.1573	–0.2331	–0.5198*	–0.5839*	–0.4767
	–0.1342	–0.1437	–0.1604	–0.1189	–0.1303	–0.3027
Log cores	0.5151*	0.7310*	0.8177*	1.0809*	1.2325*	1.9084*
	–0.1171	–0.1448	–0.1496	–0.1617	–0.1720	–0.3073
Log threads	0.5943*	0.6368*	0.6438*	0.5919*	0.5127*	0.3544*
	–0.0648	–0.0648	–0.0634	–0.0635	–0.0661	–0.0695
Log base frequency	–1.4332*	–0.9775*	–0.7959*	–0.7833*	–0.8436*	1.6006*
	–0.2355	–0.3276	–0.3234	–0.2777	–0.2906	–0.5177
Log turbo frequency	2.7389*	2.2187*	2.0403*	2.1470*	2.3450*	—
	–0.3234	–0.4100	–0.3747	–0.2973	–0.3009	—
Log (cache/cores)	0.2229	0.2829*	0.4079*	0.7375*	0.8111*	0.7575*
	–0.1362	–0.1299	–0.1317	–0.0819	–0.0857	–0.1762
Log thermal design power	—	—	—	—	—	–0.4409*
	—	—	—	—	—	–0.1044
Log graphics	—	—	—	—	—	—
Observations	82	78	71	71	69	66
Adjusted R^2	0.9186	0.9399	0.9595	0.9642	0.9560	0.9480
BIC	–33.4	–11.37	–16.5	–23.25	–8.36	7.13

*Significant at the 5-percent level.

Note: BIC = Bayesian information criterion, Q1 = first quarter, Q2 = second quarter, Q3 = third quarter, and Q4 = fourth quarter.

See footnotes at end of table.

Source: U.S. Bureau of Labor Statistics.

Appendix C: Selected quarterly models of minimum mean squared error

Table C-1. Minimum MSE selected models, first quarter 2009 to fourth quarter 2010

Variable	09Q1–09Q2	09Q2–09Q3	09Q3–09Q4	09Q4–10Q1	10Q1–10Q2	10Q2–10Q3	10Q3–10Q4
Quarter time dummy	–0.0361	–0.0638*	–0.0843	–0.1739*	–0.025	–0.0776	–0.0722
	–0.0304	–0.0267	–0.0539	–0.0736	–0.0427	–0.0421	–0.0398
Log PassMark	5.6791*	6.3611*	1.5041*	1.1733*	–0.2040	–0.3234	–0.3767*
	–1.2062	–0.9380	–0.2641	–0.1223	–0.2432	–0.2445	–0.1782
Log cores	–9.3942*	–9.4896*	—	—	1.0789*	0.9855*	1.1359*
	–0.5407	–0.6929	—	—	–0.1760	–0.1801	–0.2275
Log threads	6.1792*	5.7827*	2.0948*	—	0.2816	0.4226*	0.5105*
	–0.9694	–0.6473	–0.2631	—	–0.1597	–0.1397	–0.0859
Log base frequency	50.2257*	47.7293*	1.4333*	—	1.5867*	1.6701*	1.5911*
	–4.4902	–3.6284	–0.4846	—	–0.5289	–0.4259	–0.3495
Log turbo frequency	–54.1835*	–52.4618*	—	1.4998*	2.6945*	2.6959*	2.2455*
	–4.0868	–3.8070	—	–0.3227	–0.3367	–0.3377	–0.4500
Log (cache/cores)	–0.4955*	–0.5179*	—	—	—	—	0.1234
	–0.2209	–0.1287	—	—	—	—	–0.1438
Log thermal design power	–2.4443*	–2.8018*	–4.4002*	—	0.6121*	0.6825*	0.4321
	–0.2469	–0.2191	–0.4044	—	–0.2566	–0.2743	–0.3118
Log graphics	—	—	—	—	—	—	0.0436
	—	—	—	—	—	—	–0.0575
Observations	45	39	41	53	61	66	72
Adjusted R^2	0.9816	0.9880	0.9567	0.8997	0.9599	0.9591	0.9542

*Significant at the 5-percent level.

Note: MSE = minimum squared error, Q1 = first quarter, Q2 = second quarter, Q3 = third quarter, and Q4 = fourth quarter.

Source: U.S. Bureau of Labor Statistics.

Table C-2. Minimum MSE selected models, fourth quarter 2010 to third quarter 2012

Variable	10Q4–11Q1	11Q1–11Q2	11Q2–11Q3	11Q3–11Q4	11Q4–12Q1	12Q1–12Q2	12Q2–12Q3
Quarter time dummy	–0.0673	–0.0127	–0.0066	–0.0355	–0.0345	–0.0049	–0.0407
	–0.0426	–0.0393	–0.0443	–0.0383	–0.037	–0.0360	–0.0386
Log PassMark	–0.2849*	–0.3968*	–0.3075*	–0.1971	0.0755	1.7098*	0.5849*
	–0.1281	–0.1150	–0.1147	–0.1234	–0.1806	–0.1048	–0.1323
Log cores	1.0143*	1.0006*	0.8715*	0.5298*	0.3037*	–0.4930*	—
	–0.1728	–0.1741	–0.1862	–0.1009	–0.1498	–0.0841	—
Log threads	0.6856*	0.7706*	0.7640*	0.8064*	0.6600*	—	0.3704*
	–0.0913	–0.0832	–0.0752	–0.0813	–0.0961	—	–0.0849
Log base frequency	1.1616*	1.0799*	0.8645	—	–0.6130*	–1.6054*	–0.9900*
	–0.4239	–0.4319	–0.4543	—	–0.2781	–0.1860	–0.2135
Log turbo frequency	1.2431*	1.1928*	1.0031*	0.7053*	0.9753*	—	0.9854*
	–0.3588	–0.2999	–0.3150	–0.2529	–0.3056	—	–0.3290
Log (cache/cores)	0.4312*	0.4768*	0.4852*	0.4276*	0.4605*	0.4447*	0.5292*

See footnotes at end of table.

Table C-2. Minimum MSE selected models, fourth quarter 2010 to third quarter 2012

Variable	10Q4–11Q1	11Q1–11Q2	11Q2–11Q3	11Q3–11Q4	11Q4–12Q1	12Q1–12Q2	12Q2–12Q3
	–0.0755	–0.0657	–0.0768	–0.0932	–0.1163	–0.2111	–0.1809
Log thermal design power	–0.3643	–0.4301*	–0.4027*	—	—	—	—
	–0.1915	–0.1549	–0.1419	—	—	—	—
Log graphics	—	—	—	—	–0.0644*	—	–0.0649*
	—	—	—	—	–0.0313	—	–0.0246
Observations	87	95	81	84	83	70	79
Adjusted R^2	0.9332	0.9273	0.9278	0.9361	0.9476	0.9588	0.9550

*Significant at the 5-percent level.

Note: MSE = minimum squared error, Q1 = first quarter, Q2 = second quarter, Q3 = third quarter, and Q4 = fourth quarter.

Source: U.S. Bureau of Labor Statistics.

Table C-3. Minimum MSE selected models, third quarter 2012 to first quarter 2014

Variable	12Q3–12Q4	12Q4–13Q1	13Q1–13Q2	13Q2–13Q3	13Q3–13Q4	13Q4–14Q1	14Q1–14Q2
Quarter time dummy	–0.0438	–0.0070	–0.0048	0.0250	–0.0150	–0.0106	–0.0235
	–0.0316	–0.0238	–0.0316	–0.0308	–0.0277	–0.026	–0.0250
Log PassMark	0.2457	–0.0359	–0.148	1.0946*	0.9158*	0.6152*	0.5108*
	–0.3403	–0.2205	–0.2822	–0.1023	–0.0643	–0.1346	–0.1186
Log cores	0.2626	0.5319*	0.6150*	—	—	0.3212*	0.4510*
	–0.2602	–0.1380	–0.1735	—	—	–0.0951	–0.0782
Log threads	0.4995*	0.6329*	0.6799*	0.3031*	0.4580*	0.4491*	0.4358*
	–0.1339	–0.0942	–0.1175	–0.0966	–0.0719	–0.0755	–0.0726
Log base frequency	–0.4244	—	—	—	—	—	—
	–0.4302	—	—	—	—	—	—
Log turbo frequency	1.1021*	1.1779*	1.3342*	—	—	—	—
	–0.3887	–0.3241	–0.3839	—	—	—	—
Log (cache/cores)	0.5961*	0.6522*	0.6270*	0.2803	—	0.4668*	0.6278*
	–0.1218	–0.0901	–0.1047	–0.2188	—	–0.1609	–0.1360
Log thermal design power	–0.1304	–0.2494*	–0.2500*	–0.1466	—	—	—
	–0.1193	–0.0682	–0.0800	–0.0862	—	—	—
Log graphics	–0.0780*	–0.0885*	–0.0997*	–0.0765*	—	—	—
	–0.0259	–0.0322	–0.0326	–0.0309	—	—	—
Observations	93	105	92	92	132	139	144
Adjusted R^2	0.9642	0.9762	0.9733	0.9521	0.9449	0.9565	0.9623

*Significant at the 5-percent level.

Note: MSE = minimum squared error, Q1 = first quarter, Q2 = second quarter, Q3 = third quarter, and Q4 = fourth quarter.

Source: U.S. Bureau of Labor Statistics.

Table C-4. Minimum MSE selected models, second quarter 2014 to first quarter 2016

Variable	14Q2–14Q3	14Q3–14Q4	14Q4–15Q1	15Q1–15Q2	15Q2–15Q3	15Q3–15Q4	15Q4–16Q1
Quarter time dummy	0.0046	–0.0044	–0.0388	0.0059	0.0260	0.0084	–0.0446
	–0.0224	–0.0223	–0.027	–0.0257	–0.0269	–0.0330	–0.0334
Log PassMark	–0.2458	0.0537	0.0043	–0.2386	0.2787	0.2986	–0.1682

See footnotes at end of table.

Table C-4. Minimum MSE selected models, second quarter 2014 to first quarter 2016

Variable	14Q2–14Q3	14Q3–14Q4	14Q4–15Q1	15Q1–15Q2	15Q2–15Q3	15Q3–15Q4	15Q4–16Q1
	–0.2478	–0.2518	–0.2926	–0.2489	–0.2458	–0.1859	–0.1778
Log cores	0.7996*	0.7334*	0.7930*	0.9721*	1.0173*	0.6907*	0.6594*
	–0.1194	–0.1141	–0.1240	–0.1019	–0.113	–0.0706	–0.1471
Log threads	0.6775*	0.5138*	0.4936*	0.5248*	0.3536*	0.5136*	0.6677*
	–0.0769	–0.0836	–0.0934	–0.0930	–0.0951	–0.0979	–0.0675
Log base frequency	—	—	—	—	0.4691	—	–0.9499*
	—	—	—	—	–0.2673	—	–0.3543
Log turbo frequency	1.0399*	1.0838*	1.1390*	1.2687*	0.5461	0.6776*	1.8802*
	–0.3040	–0.2723	–0.3026	–0.2445	–0.4002	–0.2696	–0.3858
Log (cache/cores)	0.5503*	0.4601*	0.5282*	0.7386*	0.4337*	—	0.2618*
	–0.1427	–0.1412	–0.1532	–0.1338	–0.1323	—	–0.1249
Log thermal design power	—	–0.1764*	–0.1997*	–0.2599*	–0.4771*	–0.2187*	—
	—	–0.0567	–0.0631	–0.0761	–0.0795	–0.0674	—
Log graphics	—	—	—	—	—	—	0.0671
	—	—	—	—	—	—	–0.0389
Observations	164	169	132	97	104	97	85
Adjusted R^2	0.9571	0.9564	0.9612	0.9677	0.9589	0.9445	0.9353

*Significant at the 5-percent level.

Note: MSE = minimum squared error, Q1 = first quarter, Q2 = second quarter, Q3 = third quarter, and Q4 = fourth quarter.

Source: U.S. Bureau of Labor Statistics.

Table C-5. Minimum MSE selected models, first quarter 2016 to third quarter 2017

Variable	16Q1–16Q2	16Q2–16Q3	16Q3–16Q4	16Q4–17Q1	17Q1–17Q2	17Q2–17Q3
Quarter time dummy	0	0.0027	–0.0443	–0.064	–0.0078	–0.0147
	–0.0333	–0.0379	–0.0374	–0.0369	–0.0411	–0.0629
Log PassMark	–0.4185*	–0.2707	–0.3118	–0.2435	–0.2349	–0.0562
	–0.1255	–0.1637	–0.159	–0.1787	–0.1970	–0.1756
Log cores	0.7163*	0.8597*	0.9135*	0.9956*	1.1211*	1.4852*
	–0.1190	–0.1576	–0.1536	–0.1535	–0.1893	–0.1387
Log threads	0.6699*	0.6745*	0.6762*	0.4275*	0.2995*	—
	–0.0538	–0.0706	–0.0710	–0.0973	–0.0993	—
Log base frequency	–0.8992*	–0.6968	–0.5853	–0.7459*	–0.6896	—
	–0.2569	–0.3729	–0.3668	–0.2959	–0.3647	—
Log turbo frequency	2.2450*	1.9030*	1.7867*	2.3522*	2.5434*	—
	–0.3333	–0.4112	–0.4089	–0.3420	–0.3720	—
Log (cache/cores)	0.3113*	0.3498*	0.4521*	0.6533*	0.6685*	0.7848*
	–0.1216	–0.1226	–0.1207	–0.0839	–0.1007	–0.1994
Log thermal design power	—	—	—	–0.1728	–0.2316*	—
	—	—	—	–0.1025	–0.1094	—
Log graphics	0.2048*	0.0431	0.0299	–0.0909*	–0.1324*	—
	–0.0467	–0.0339	–0.0302	–0.0366	–0.0433	—
Observations	82	78	71	71	69	66
Adjusted R^2	0.9333	0.9408	0.9599	0.966	0.9591	0.9153

*Significant at the 5-percent level.

Note: MSE = minimum squared error, Q1 = first quarter, Q2 = second quarter, Q3 = third quarter, and Q4 = fourth quarter.

See footnotes at end of table.

Source: U.S. Bureau of Labor Statistics.

Appendix D: An introduction to statistical learning model specification selection

The basic steps that we use for prescreening and cross-validation come from *An introduction to statistical learning*:^[41]

1. Let M_0 denote the null model, which contains no (p) predictors. This model simply predicts the sample mean for each observation.
2. For $k = 1, 2, \dots, p$:
 - a. Fit all (p/k) models that contain k predictors.
 - b. Pick the best among these (p/k) models, and call it M_k . Here, *best* is defined as having the smallest residual sum of squares, or equivalently R^2 .
3. Select a single best model from among M_0, \dots, M_p using cross-validated prediction error, C_p (Akaike information criterion or AIC), Bayesian information criterion, or adjusted R^2 .

For step 3, we only use cross-validation (10 fold). We repeat the steps 500 times, and we calculate the standard errors for each of the 1 through p models. The model with the smallest number of predictors whose standard error was within range of the lowest mean-squared-error value is selected.

The code we use to implement this technique is based on code from page 250 of *An introduction to statistical learning*.^[42] We added code to repeat the process 500 times and to calculate the standard errors.

SUGGESTED CITATION

Steven D. Sawyer and Alvin So, "A new approach for quality-adjusting PPI microprocessors," *Monthly Labor Review*, U.S. Bureau of Labor Statistics, December 2018, <https://doi.org/10.21916/mlr.2018.29>

NOTES

¹ The PPI is one of the nation's principal federal economic indicators that measures the average change over time of the selling prices received by domestic producers of goods and services. This family of indexes is made up of approximately 10,000 PPIs for individual products and groups of products that are published each month, one of which is the PPI for semiconductor and related-device manufacturing. One subcomponent of semiconductor and related-device manufacturing is the index for microprocessors. For more information, see *BLS handbook of methods*, chapter 14, "Producer prices" (U.S. Bureau of Labor Statistics, 2014), <https://www.bls.gov/opub/hom/pdf/homch14.pdf>. The PPI for microprocessors (including microcontrollers) has not been published since March 2015, since it does not meet publication standards for the PPI.

² Bruce T. Grimm, "Price indexes for selected semiconductors, 1974–96," *Survey of Current Business*, February 1998, pp. 78, 8–24; David M. Byrne, Stephen D. Oliner, and Daniel E. Sichel, "How fast are semiconductor prices falling?" *Review of Income and Wealth*, vol. 64, no. 3, April 2017, [Author: I found that this was published in Sep 2018—see <https://onlinelibrary.wiley.com/doi/pdf/10.1111/roiw.12308>] pp. 679–702, <https://doi.org/10.1111/roiw.12308>; Ana Aizcorbe, "Why did semiconductor price indexes fall so fast in the 1990s? A decomposition," *Economic Inquiry*, vol. 44, no. 3, 2006, pp. 485–496 [Author: I found this at <https://onlinelibrary.wiley.com/doi/abs/>

[10.1093/ei/cbj027](#) and it shows March 2007 as publishing date.]; and Liyang Sun, “What are we paying for? A quality-adjusted price Index for laptop microprocessors” (Ph.D. dissertation, Wellesley College, 2014).

[3](#) Allen C. Goodman and Thomas G. Thibodeau, “Housing market segmentation and hedonic prediction accuracy,” *Journal of Housing Economics*, vol. 12, no. 3, February 2003, pp. 181–201. They used statistical learning techniques to compare hedonic specifications for models that predicted housing prices outside of an official price index.

[4](#) Peter Van Zant, *Microchip fabrication: a practical guide to semiconductor processing*, 6th ed. (New York: McGraw Hill, 2014), p. 394.

[5](#) Kenneth Flamm, “Has Moore’s Law been repealed? An economist’s perspective,” *Computing in Science and Engineering*, March–April 2017. He shows in table 1, p. 33, that for Intel, the cost per transistor has continued to decline even as wafer processing costs have increased.

[6](#) For a more indepth description of this phenomenon, see “Understanding Dennard scaling” (Sunnyvale, CA: Rambus, 2016), <https://www.rambus.com/blogs/understanding-dennard-scaling-2/?nabe=4857318206603264:1.6583178454368256:0>.

[7](#) In Van Zant, *Microchip fabrication*, p. 434, Van Zant describes some of the strategies of the methods for increasing performance even as feature size decreases. In addition, performance implicitly includes power usage.

[8](#) Byrne et al., “How fast are semiconductor prices falling?” p. 1–23. The authors discuss the challenges to using a matched-model methodology beginning in the mid-2000s.

[9](#) Ariel Pakes, “A reconsideration of hedonic price indexes with an application to PC’s,” *American Economic Review*, vol. 93, no. 5, August 2003, pp. 1578–1596; Patrick Bajari and C. Lanier Benkard, “Demand estimation with heterogeneous consumers and unobserved product characteristics: a hedonic approach,” *Journal of Political Economy*, vol. 113, no. 6, December 2005, pp. 1,239–1,276; Patrick Bajari and C. Lanier Benkard, “Hedonic price indexes with unobserved product characteristics, and application to personal computers,” *Journal of Business & Economic Statistics*, vol. 23, no. 1, January 2005, pp. 61–75; and Tim Erickson and Ariel Pakes, “An experimental component index for the CPI: from annual computer data to monthly data on other goods,” *American Economic Review*, vol. 101, no. 5, August 2011, pp. 1,707–1,738. A quality-adjusted price also accounts for technological change in a product.

[10](#) Byrne et al., “How fast are semiconductor prices falling?” pp. 1–23.

[11](#) Ibid.

[12](#) Any reference to Intel in this article is for this article’s analysis only.

[13](#) For more information about SPEC, go to <https://www.spec.org/>.

[14](#) Byrne et al., “How fast are semiconductor prices falling?” p. 11.

[15](#) Ibid.

[16](#) For more information, see Intel ARK, “Product specifications,” <https://ark.intel.com/>.

[17](#) All definitions were obtained from Intel’s website <https://ark.intel.com/>. Under product specifications, each category has a link to the definition for that particular specification.

[18](#) “Using SPEC CPU2006 benchmark results to compare the compute performance of servers,” white paper, revision 1.0 (Cisco, June 2010), p. 4, https://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/unified-computing/SPECCPU2006_overview.pdf.

[19](#) Ibid.

[20](#) Ian Cutress, “The Intel 6th Gen Skylake: Core i7-6700K and i5-6600K tested” (AnandTech, August 2015), <http://www.anandtech.com/show/9483/intel-skylake-review-6700k-6600k-ddr4-ddr3-ipc-6th-generation>.

- ²¹ Michael Holdway, “An alternative methodology: valuing quality change for microprocessors in the PPI” (U.S. Bureau of Labor Statistics, 2001), <https://www.semanticscholar.org/paper/An-Alternative-Methodology-%3A-Valuing-Quality-Change-Holdway/6a9b6605d25e3acacd23786bf53ac889f39e6f41>.
- ²² Ibid., p. 24.
- ²³ Flamm, “Has Moore’s Law been repealed? An economist’s perspective.”
- ²⁴ Ibid., p. 26.
- ²⁵ For more information on the Intel Stable Image Platform, see “Intel Stable Image Platform Program [Intel SIPP],” <http://www.intel.com/content/www/us/en/computer-upgrades/pc-upgrades/sipp-intel-stable-image-platform-program.html>.
- ²⁶ Trevor Hastie, Robert Tibshirani, and Jerome Friedman, *The elements of statistical learning: data mining, inference, and prediction*, 2d ed. (New York: Springer, 2009), p. 219, http://statweb.stanford.edu/~tibs/ElemStatLearn/printings/ESLII_print10.pdf.
- ²⁷ Ibid., p. 222.
- ²⁸ Holdway, “An alternative methodology: valuing quality change for microprocessors in the PPI.”
- ²⁹ John J. Dziak, Donna L. Coffman, Stephanie T. Lanza, and Runze Li, “Sensitivity and specificity of information criteria,” technical report 12-119 (The Pennsylvania State University, The Methodology Center, June 2012), p. 2, <https://www.methodology.psu.edu/files/2019/03/12-119-2e90hc6.pdf>.
- ³⁰ Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani, *An introduction to statistical learning: with applications in R*, Springer Texts in Statistics (New York: Springer, 2013), pp. 211–212.
- ³¹ Dziak et al., “Sensitivity and specificity of information criteria,” p. 23.
- ³² James et al., *An introduction to statistical learning*. We used code from this source to implement this procedure.
- ³³ Ibid. To be more precise, to select the best model, calculate the model with the smallest number of variables whose standard error is within range of the lowest MSE value. This procedure is called the “one-standard-error-rule.” See page 214 for more detail.
- ³⁴ Ten-fold cross-validation is repeated 500 times.
- ³⁵ In addition, relative weights for items within an index are reset to their revenue share of an index when the industry is resampled. PPI semiconductors is typically resampled every 5 years.
- ³⁶ For more information on the PPI code, see <https://data.bls.gov/timeseries/PCU334413334413P>.
- ³⁷ Note that relative weights for indexes within an industry are reset to their revenue share of an industry when the industry is resampled. PPI semiconductors are typically resampled every 5 years.
- ³⁸ For more information about SPEC, go to <https://www.spec.org/>.
- ³⁹ HP, “An overview of the SPEC CPU2006 benchmark on HP ProLiant servers and server blades,” October 2007, ftp://ftp.hp.com/pub/c-products/servers/benchmarks/SPEC_CPU2006_Overview_101907.pdf.
- ⁴⁰ For more detail regarding each test, see “PassMark CPU test information,” PassMark Software, CPU Benchmarks, https://www.cpubenchmark.net/cpu_test_info.html.
- ⁴¹ James et al., *An introduction to statistical learning*, p. 205.
- ⁴² Ibid., p. 250.

RELATED CONTENT

Related Articles

[Multifactor productivity slowdown in U.S. manufacturing](#), *Monthly Labor Review*, July 2018.

[Cross-validation of quality-adjustment methods for price indexes](#), *Monthly Labor Review*, June 2018.

[Producer prices, 2017: inflation accelerates across the stages of production](#), *Monthly Labor Review*, June 2018.

Related Subjects

[Benchmarking](#) | [Producer price index](#) | [Statistical programs and methods](#) | [industry](#) | [Prices](#) | [Experimental methodology](#) | [Inflation](#)