

## HOW FAST ARE SEMICONDUCTOR PRICES FALLING?

BY DAVID M. BYRNE

*Federal Reserve Board*

STEPHEN D. OLINER

*American Enterprise Institute and UCLA*

AND

DANIEL E. SICHEL

*Wellesley College and NBER*

The Producer Price Index (PPI) for the U.S. suggests that semiconductor prices have barely been falling in recent years, a dramatic contrast to the rapid declines reported from the mid-1980s to the early 2000s. This slowdown in the rate of decline is puzzling in light of evidence that the performance of microprocessor units (MPUs) has continued to improve at a rapid pace. Over the course of the 2000s, the MPU prices posted by Intel, the dominant producer of MPUs, became much stickier over the chips' life cycle. As a result of this change, we argue that the matched-model methodology used in the PPI for MPUs likely started to be biased after the early 2000s and that hedonic indexes can provide a more accurate measure of price change since then. MPU prices fell rapidly through 2004 on every price measure we present, with the PPI declining at an even quicker pace than the hedonic indexes. However, from 2004 to 2009, our preferred hedonic index fell faster than the PPI, and from 2009 to 2013 the gap widened further, with our preferred index falling at an average annual rate of 42 percent, while the PPI declined at only a 6 percent rate. Given that MPUs currently represent about half of U.S. shipments of semiconductors, this difference has important implications for gauging the rate of innovation in the semiconductor sector.

**JEL Codes:** E01, E31, L63

**Keywords:** hedonic price index, measurement, microprocessor, quality adjustment, technological change

### 1. INTRODUCTION

How fast are semiconductor prices falling? Data from the Producer Price Index (PPI) published by the U.S. Bureau of Labor Statistics indicate that prices of microprocessor units (MPUs) have barely been falling in recent years. This very slow rate of price decline stands in sharp contrast to the rapid declines in MPU prices reported from the mid-1980s up to the early 2000s and the exceptionally rapid declines in the latter half of the 1990s. If accurate, the apparent slowdown in MPU price declines in recent years would be troubling, given the long-run relationship between rates of price decline of semiconductors and the

\*Correspondence to: David Byrne, Federal Reserve Board of Governors 21st and Constitution Ave., NW Washington, DC 20551, USA (david.m.byrne@frb.gov).

pace of innovation in that sector, which has served as an engine of growth throughout the economy.<sup>1</sup>

The apparent slowdown in the rate of price decline is puzzling given evidence that the performance of MPUs continued to improve at a rapid pace after the mid-2000s. The key to resolving the puzzle may reside in another industry development. Over the course of the 2000s, the properties of the MPU prices posted by Intel, the dominant producer, changed dramatically. In the early part of the decade, Intel almost always lowered the posted prices of existing chips as they aged and as new, higher-performance models were introduced. However, by the mid-2000s, the posted prices had become much stickier over the chips' life cycle. The share of MPU chips with a price cut within four quarters of introduction fell from 100 percent for the 2000 and 2001 vintages of chips to about 60 percent for the 2004 to 2009 vintages and to an average about 20 percent for later vintages.

The reason for this change in the life-cycle pattern of posted prices is a matter of speculation. It is possible that Intel actually changed its life-cycle pricing strategy to extract more revenue from older models, with the posted prices reflecting this change. Alternatively, Intel may not have changed its pricing strategy at all, but simply began to post prices that had less connection to transaction prices. Neither we nor anyone else outside Intel has the information to distinguish between these two alternatives, as Intel does not release chip-level transaction prices. Given this lack of information, the challenge is to construct price indices that are robust to potential measurement issues with the available price data.

We argue that the matched-model methodology used by BLS fails this challenge after the early 2000s. In an environment with flat price profiles as models age and with increasing quality over time, matched-model price indexes for MPUs such as the PPI will understate the amount of price decline.<sup>2</sup> We make the case that hedonic indexes are better suited to capturing price trends in this situation and develop new hedonic indexes for quality-adjusted prices using data for Intel MPUs from 2000 to 2013. We also argue that hedonic indexes for MPUs should utilize measures of end-user performance rather than technical variables capturing physical characteristics (such as feature size) and engineering specifications (such as clock speed). Performance measures provide a superior control for quality change when constructing price indexes for MPUs because they gauge the actual output obtained by a user rather than the input characteristics used to produce that output. Moreover, with rapid changes in MPU architecture, identifying the correct set of technical characteristics (and likely changes in that set) can be challenging.

Every index we present shows that MPU prices fell rapidly through 2004, with the PPI declining at an even quicker pace than the hedonic indexes. However, from 2004 to 2009, our preferred hedonic index fell faster than the PPI. And from

<sup>1</sup>See Aizcorbe *et al.* (2008) for a discussion of the relationship between price change and innovation for semiconductors. Brynjolfsson and McAfee (2014) and Baily *et al.* (2013) highlight key innovations across the economy, many of which have been driven by the revolution in computing power. For analyses of the contribution of information technology to economic growth, see Oliner and Sichel (2000), Oliner *et al.* (2007), and Byrne *et al.* (2013).

<sup>2</sup>Although BLS does not indicate which prices are included in the PPI for MPUs, we are able to replicate the trend in the PPI with a matched-model index that uses only Intel's posted prices; see Section 4 below. Hence, any use of other MPU prices in the PPI has not had a material effect on the price trend it shows.

2009 to 2013, the gap widened dramatically, with our preferred index falling at an average annual rate of 42 percent, while the PPI declined at only a 6 percent rate. Thus, our results imply that the PPI vastly overstates the slowdown in price declines for MPUs.

To gauge whether using a measure of end-user performance matters empirically, we also estimate parallel hedonic regressions with technical characteristics as the controls for chip quality, along the lines of the analysis in Flamm (2015). The price indexes generated by the two approaches are strikingly different after 2004. The declines in the indexes based on technical characteristics slowed sharply from that point forward, much like the PPI. After 2004, clock speed—a key technical characteristic which had been highly correlated with user performance—stopped rising in response to problems with heat generation, but Intel continued to boost performance in other ways. The hedonic regressions we estimated with technical characteristics (including clock speed) evidently cannot capture the ongoing gains in performance, which translates into slower declines in constant-quality prices. This result highlights the importance of using direct measures of end-user performance when estimating hedonic price regressions.

We chose to focus on MPUs rather than a broader set of semiconductor products for several reasons.<sup>3</sup> First, MPUs are a large segment of the semiconductor sector; in 2014, they represented about half of U.S. shipments (the scope of the PPI). Second, price series for MPUs extend back to the mid-1980s, allowing for comparisons of price trends over time. Given that price trends in this sector often are used to infer rates of technical progress, this historical comparability is important. Finally, we believe that developments in MPU technology likely provide a rough guide to developments in other parts of the semiconductor sector, such as the chips that are used in smartphones and tablets.

Our work on MPU prices builds on important earlier research. Notable studies that have constructed hedonic price indexes for semiconductors (or computing equipment that embeds these chips) include Cole *et al.* (1986), Grimm (1998), Flamm (2007, 2015), and Song (2010). All of these studies relied on technical characteristics of the chips to control for quality. Recognizing that these characteristics may not fully capture the capabilities of the processor for end users, Triplett (1989, p. 147) and Berndt and Griliches (1993, p. 91) called for analysis of computer prices with more refined controls for performance. In this vein, Chwelos (2000, 2003) and Benkard and Bajari (2005) constructed hedonic price indexes for PCs using performance benchmarks for MPUs rather than characteristics of the MPUs.<sup>4</sup> Our research takes the same approach to meeting the concerns raised by Triplett and by Berndt and Griliches.

<sup>3</sup>Among all MPUs, this paper analyzes prices of the MPUs used in desktop personal computers (PCs), for which data are the most readily available. In other work, we are developing indexes for MPUs going into servers, the types of machines that would support cloud computing in server farms and other processor-intensive applications. In addition, Sichel's undergraduate thesis student, Sophie (Liyang) Sun (2014) developed price indexes for MPUs used in laptop computers.

<sup>4</sup>In addition, Holdway (2001) examined unit value MPU price indexes using benchmark test scores as a measure of quality-adjusted units of computing power. Grimm (1998) developed specifications that controlled for millions of instructions per second (MIPS); this measure of performance, however, has limited ability to account for differences across MPUs in the translation of instructions into program execution.

Another strand of the literature has focused on the choice between matched-model and hedonic measures of semiconductor prices. Aizcorbe *et al.* (2003) showed that under certain conditions (highly granular data on model prices and high-frequency observations), a matched-model index can produce similar results to a hedonic index. However, these conditions may be difficult to meet in practice, and Silver and Heravi (2005) analyzed the potential biases in matched-model indices when they are not met. For our analysis of MPU prices, the limitations of the matched-model approach go beyond the issues highlighted by Silver and Heravi (2005), which relate to the influence of entering and exiting models. In our case, the problem reflects inherent deficiencies in the data for all models after the early years of our sample period.

Our findings for MPU prices have important implications for gauging the pace of innovation in the semiconductor sector. In addition, given the use of MPUs as an input in other industries, changes in MPU prices can affect the allocation of value-added across industries. At the same time, our results have limited direct implications for the measurement of real GDP or output per hour for the economy as a whole. Semiconductors mostly are intermediate inputs and so are not counted directly in GDP. Imports and exports of semiconductors are the exception, with semiconductor exports net of imports counting as GDP. However, the trade quantities are small enough on net that the adoption of our preferred price index for MPUs would leave real GDP growth little changed.

The next section highlights the puzzle raised by the very slow rate of decline in the PPI during recent years at the same time that the engineering frontier for MPUs continued to move out rapidly. Section 3 presents our argument that hedonic price indexes are likely to better capture price trends than are matched-model indexes (such as the PPI) after the early 2000s. Section 4 reviews our data, which cover the period from 2000 to 2013. In Section 5, we describe the hedonic regressions used to obtain measures of quality-adjusted prices. Section 6 presents our results, and Section 7 concludes.

## 2. THE PUZZLE

As noted, the PPI for MPUs has fallen very slowly in recent years. This section explores whether this extreme slowdown meshes with the trends in technological advance for MPU chips, drawing heavily on material in Byrne *et al.* (2013).

The standard definition of a semiconductor technology cycle is the amount of time required to achieve a 30 percent reduction in the width of the smallest feature on a chip. Because chips are rectangular, a 30 percent reduction in both directions implies about a 50 percent reduction ( $0.7 \times 0.7$ ) in the area required for the smallest chip component. As documented in Byrne *et al.* (2013), the semiconductor industry has achieved massive reductions in scaling over time. Indeed, the area occupied by a chip component in 2014 was roughly 400,000 times smaller than in 1969, roughly in line with the prediction made by Moore (1965) 50 years ago in his eponymous Law.<sup>5</sup>

<sup>5</sup>Moore's Law states that the number of components on leading-edge chips will double every two years. Moore's original formulation (Moore, 1965) pegged the doubling time at only one year, but in 1975 he revised the period to two years based on actual experience to that point (Intel Corporation, 2005). For a discussion of the outlook for Moore's Law, see Bauer *et al.* (2013).

TABLE 1  
TECHNOLOGY CYCLES FOR INTEL MPU CHIPS (YEARS NEEDED  
FOR 30 PERCENT REDUCTION IN LINEAR SCALING)

Period	Years
1971–1994	2.9
1994–2014	1.9
1994–2004	1.9
2004–2014	1.9

*Source:* Byrne *et al.* (2013) with update to 2014 based on data posted at <http://ark.intel.com>.

There is a broad consensus that the pace of technical advance in the semiconductor industry sped up in the mid-1990s, a development first brought to the attention of economists by Jorgenson (2001). Table 1 reports the average length of the technology cycle for Intel MPU chips (as defined above) for various periods. As shown, Intel's technology cycle averaged about three years until 1994 and then dropped to about two years from 1994 to 2014. Thus, for the period covered by our empirical work, there had been no pullback from the two-year cycle. Recently, however, Intel's CEO acknowledged that technical challenges have made it impossible to remain on the two-year cycle and that the company is now operating with a cycle closer to 2½ years.<sup>6</sup> Although this shift occurred after our sample period, in the future it could temper the price declines for MPUs.

Until the early to mid-2000s, each new generation of MPU technology allowed for an increase in the number of basic calculations performed per second (clock speed) for a given chip design, thereby boosting performance. However, as speed continued to increase, dissipating the generated heat became problematic and performance gains slowed. Pillai (2013) highlighted this slowdown with performance scores from the Standard Performance Evaluation Corporation (SPEC), a non-profit organization that establishes performance benchmark tests for computing equipment and publishes test results submitted by member organizations. These scores are based on standard tasks designed to reflect the needs of computer users. After having risen 60 percent per year from 1990 to 2000, SPEC performance rose about 41 percent per year on average from 2000 to 2004.

In response to the heat dissipation problem, Intel shifted away from increases in clock speed and boosted performance instead by placing multiple copies of the core architecture on each chip—a change enabled by smaller feature size—and by improving the design of those cores.<sup>7</sup> How did this strategy shift affect the rate of increase in performance for end users? Figure 1 extends Pillai's analysis through 2013 (Pillai's results end in 2008), plotting average SPEC performance measures for MPUs introduced each year.<sup>8</sup> Starting in 2004, the figure includes a second

<sup>6</sup>See Clark (2015).

<sup>7</sup>See Shenoy and Daniel (2006). Also, Hennessy and Patterson (2012) document that increases in clock speed stalled out during this period (Figure 1.11, p. 24).

<sup>8</sup>The data through 1999 were kindly provided by Unni Pillai. We linked these data with SPEC performance results for 1999–2013 that we obtained directly from SPEC. We accessed the SPEC data on May 20, 2014 and used the benchmark suites SPEC® CPU2006 and SPEC® CPU2000. The SPEC score for each year is the average over the Intel desktop MPU chips introduced in that year. Details are provided in Section 4 and in Appendix A.

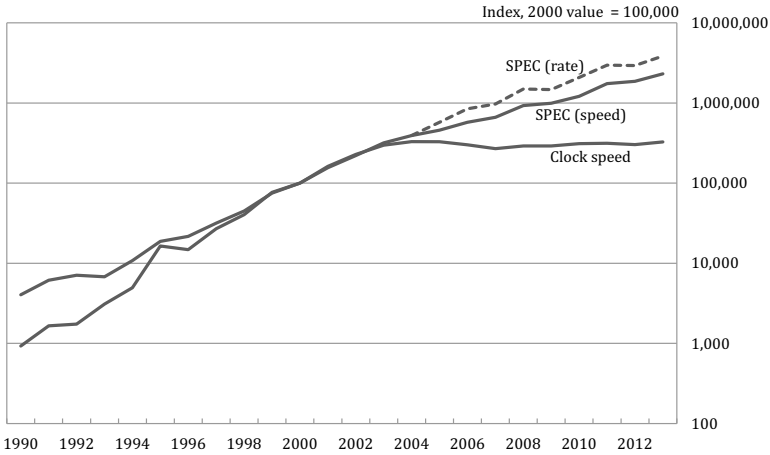


Figure 1. Desktop MPU Performance Measures

Source: Authors’ calculations using data provided by Unni Pillai and performance information from Intel price lists and System Performance Evaluation Corporation.

SPEC performance measure (the “rate” score) that incorporates more fully the performance gain from using multiple cores than does the standard (“speed”) measure.<sup>9</sup> The use of multiple cores enables greater parallelization of processing, which enhances performance.<sup>10</sup>

Pillai noted that the new design approach (adding cores) was not as effective at translating miniaturization into performance as the old design approach (boosting clock speed). Figure 1 bears out this observation. Despite the introduction of multi-core MPUs in 2005, performance gains slowed further to an average rate of 29 percent over 2004–2013, using the SPEC rate measure that accounts more completely for the effects of parallel processing. Even with this downshift, however, the continued increases in MPU performance after 2004 contrast sharply with the stalling out of clock speed.<sup>11</sup>

Historically, improvements in the engineering frontier have translated into steep declines in MPU prices. Figure 2 shows the annual price declines back to 1986, splicing together estimates from Grimm (1998) through 1992, the Federal Reserve Board for 1993–1997, and the PPI beginning in 1998, when BLS adopted their current methodology, described by Holdway (1997), in response to research at the Federal Reserve Board and the Bureau of Economic Analysis. During the full period, MPU prices fell at an average rate of nearly 30 percent per year, with especially sharp drops in the second half of the 1990s. However, reported price declines have slowed dramatically over the past several years. Indeed, the declines

<sup>9</sup>See Section 4 for further discussion of these alternative SPEC scores.

<sup>10</sup>Thompson (2015) describes how firms that were able to exploit the parallelization possible on computers with multicore chips realized greater productivity gains than did firms that were not able to exploit these innovations.

<sup>11</sup>During the first half of the 1990s, performance also increased more rapidly than did clock speed. Hennessy and Patterson (2012) discuss this pattern; see the note to their Figure 1.11.

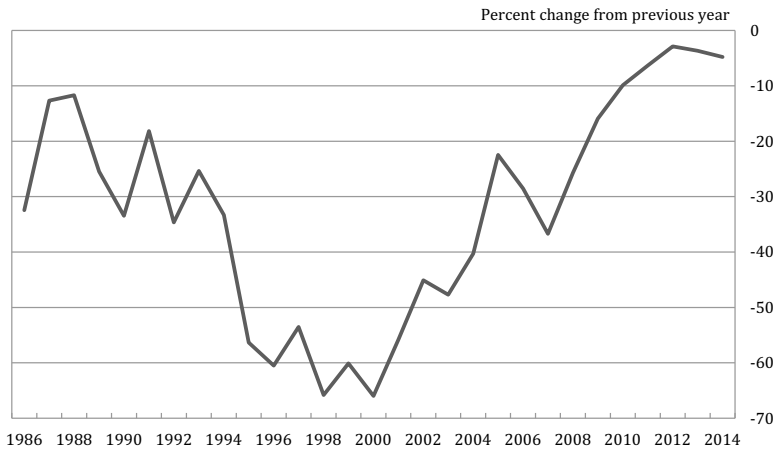


Figure 2. Reported MPU Prices

Sources: Grimm (1998, Table 12) for 1986–92, Federal Reserve Board for 1993–97, and Bureau of Labor Statistics (<http://www.bls.gov>) for 1998–2014.

in each year since 2010 were smaller than in any prior year back to 1986, breaking the link with the continued engineering improvements.

Perhaps the cost of achieving these engineering advances has risen so much as to leave constant-quality MPU prices about unchanged in recent years. If that were the case, the posted prices for newly-introduced MPUs, which do not include any adjustment for quality change, would have to be rising. As a logical matter, this is the only way that constant-quality prices could be flat given the advances in MPU performance. The data, however, show the opposite—that posted prices for newly-introduced MPUs have been trending down. Using the MPU chips in Figure 1 that can be matched to Intel’s price lists, the average posted price for newly-introduced chips fell at an annual pace of 11 percent over 2000–2013, with no evidence of a slowdown late in the sample period. Because these declines pertain to chips whose performance is improving over time, the downtrend in prices after controlling for quality change necessarily will be even steeper.

All in all, the shift to much slower price declines for microprocessors in the PPI is a puzzle in light of the continued substantial improvements on the engineering front and the downtrend in posted prices prior to any quality adjustment.

### 3. CONSTRUCTING MPU PRICE INDEXES

This section addresses several issues that arise in constructing price indexes for MPUs. We begin with a consideration of matched-model versus hedonic indexes in light of the dramatic change that took place in the life-cycle profiles of Intel’s posted MPU prices between the early years and the later years of our sample period. We show that this change has distorted the measurement of price trends in matched-model indexes like the PPI, and we argue that hedonic indexes are preferred for measuring trends in MPU prices. In addition, we argue that,

given possible measurement error and current limitations on data availability, hedonic regressions that rely only on prices in the early part of a model's life cycle are likely to be more robust to mismeasurement than hedonic regressions that include prices over the entire life cycle. We also argue that hedonic indexes using actual measures of performance have important advantages over those using variables capturing physical characteristics of MPUs.

### 3.1. *Life-cycle Prices: Implications for Matched-Model versus Hedonic Indexes*

Figure 3 illustrates the sharp change over the course of the 2000s in the life-cycle properties of Intel's posted prices for MPU chips. The upper panel displays Intel's posted prices for all desktop MPUs during 2000–2001, while the lower panel shows the analogous posted prices during 2011–2012. The difference between the two panels is stark. In the early period, prices fell steeply over a model's life cycle. However, by 2011–2012, price paths are flat or nearly so, with only a few instances of sizable price declines.

Figure 4 summarizes this change in life-cycle pricing over the entire 2000–12 period by showing the share of all Intel desktop MPU models introduced in each year that experienced a price decline within four quarters of introduction.<sup>12</sup> As can be seen, every model introduced in 2000 and 2001 had at least one such price cut. But the share then dropped to about 60 percent for the 2004 to 2009 cohorts and took another leg down thereafter, falling to an average of about 20 percent for the 2010–2012 cohorts.

The flat life-cycle profiles pose difficulties for the matched-model approach. Matched-model indexes measure the change in price from period  $t$  to period  $t+1$  for models with unchanged characteristics that are sold in both periods. As such, these indexes represent an average of model-specific price changes along price profiles of the type shown in Figure 3. If those profiles are flat, a matched-model index will show that constant-quality prices are unchanged, similar to the PPI in recent years. However, this will be a biased indicator of the true change in constant-quality prices if advances in MPU technology enable new models to be brought to market at lower constant-quality prices.<sup>13</sup>

Hedonic price indexes are less susceptible to this bias. The hedonic approach pools a set of models in the market, and then regresses the models' prices on measures of product characteristics or performance to control for differences in quality. In contrast to matched-model indexes, hedonic regressions use information across models to help identify constant-quality prices in a given period and the changes over time. Thus, even in the presence of flat life-cycle price profiles, the introduction of new models at lower constant-quality prices will cause the hedonic index to decline over time.

That said, merely using the hedonic approach need not fully solve the bias problem. As we discuss below, care must be taken to estimate the hedonic index

<sup>12</sup>The price data in the figure are described in Section 4.

<sup>13</sup>Hobijn (2001) and Silver and Heravi (2005) argue that matched-model indexes can be biased even when the observed prices represent true transaction prices. This bias can arise when entering or exiting models (which are omitted from matched-model indexes in their first and final periods of existence) have quality-adjusted prices that differ systematically from those for continuing models.



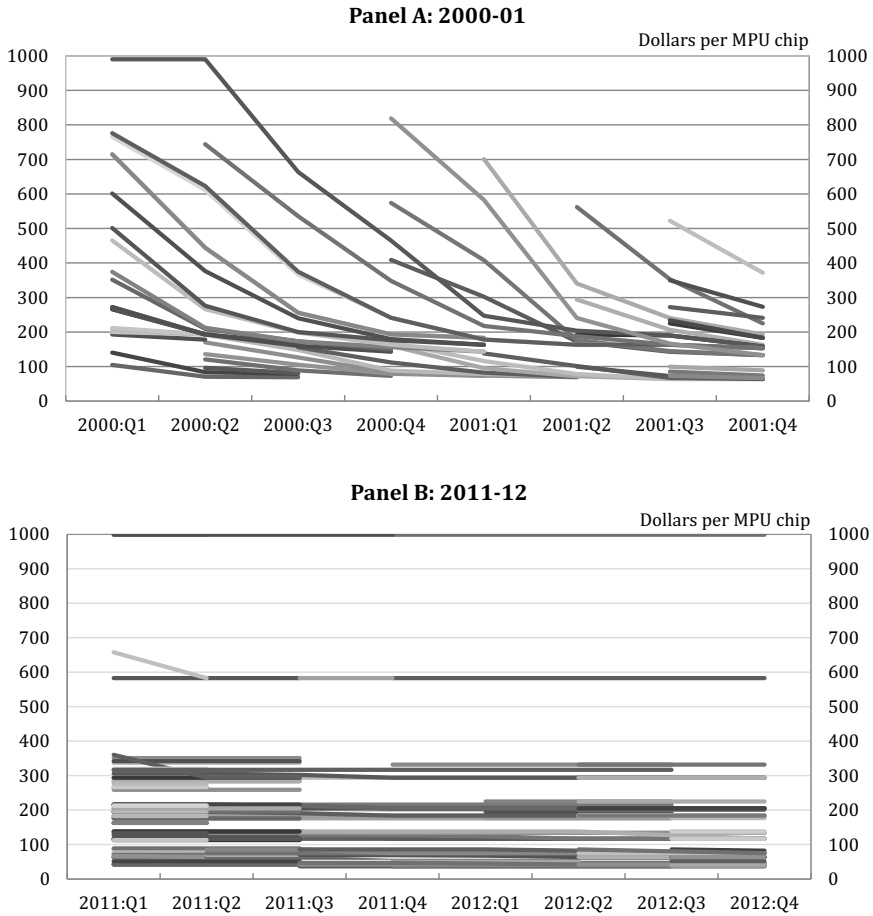


Figure 3. Intel List Prices

Note: Prices for 1,000 unit “trays” of MPUs.

Source: Authors’ calculations from Intel price lists.

in a way that minimizes the risk of bias, and we propose a way to do this. The key point is that a hedonic index can provide an accurate measure of price change over the entire sample period while a matched-model index is doomed to fail.

### 3.2. End-user Performance Measures or Technical Characteristics in Hedonic Regressions?

A long-recognized challenge facing hedonic methods has been the choice of variables to include in the regression to control for quality. The typical approach has been to include measures of key technical characteristics of each model, in the belief that these characteristics will serve as a proxy for what purchasers value. This approach has raised concerns in the past, and researchers have called for the inclusion of actual performance measures rather than technical characteristics

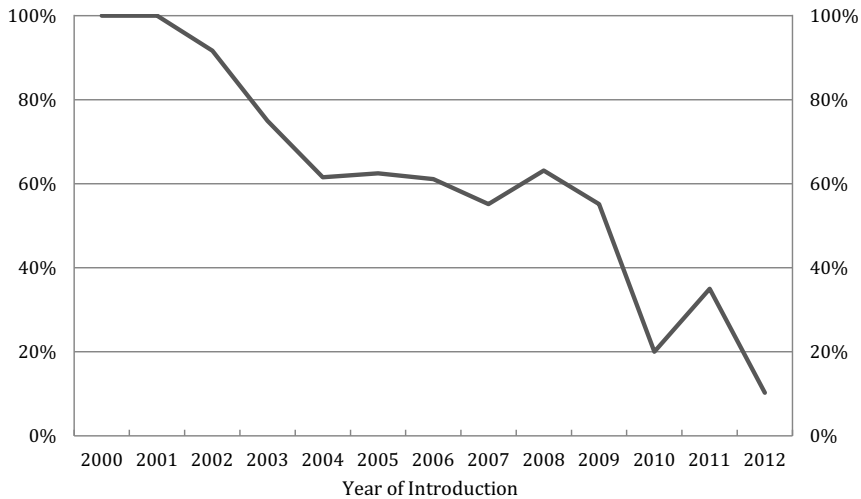


Figure 4. Share of Intel Desktop MPUs with List Price Decline within Four Quarters of Introduction

Source: Authors' calculations from Intel price lists.

(see Triplett, 1989, and Berndt and Griliches, 1993). The underlying logic behind using performance measures is to focus on the output received by users rather than the input characteristics used to generate that output. For many products, end-user measures of performance are unavailable so there is little choice but to use characteristics. For MPUs, however, measures of performance are available for tasks actually undertaken by users.

Accordingly, in this paper, we develop hedonic indexes with performance measures from SPEC. Using these performance measures avoids the difficulties entailed in trying to capture quality for MPUs with rapidly changing architecture by identifying relevant technical characteristics. Our preferred performance metrics are described in detail in the next section.

### 3.3. Introduction-period or Full-sample Hedonic Regressions?

A natural starting point would be to assume that hedonic regressions should be estimated using all available data. However, under certain circumstances, a full-sample regression can lead to biased estimates of price change. For example, in the 1980s, Fisher *et al.* (1983), Cole *et al.* (1986), and Gordon (1987) advocated constructing price indexes for mainframe computers using only introduction-period prices for each model. This choice was made for a number of reasons, including concerns that the market for mainframes was not in equilibrium (meaning old and new models were simultaneously in the market and old models were not re-priced to equalize price-performance ratios) and that available IBM list prices for older models might overlook discounts and therefore not be actual transaction prices.

We believe similar concerns apply to MPUs. To explain why, we consider two alternative scenarios. In the first scenario, posted prices do not represent true transactions prices because Intel offers progressively larger discounts to selected purchasers as models age.<sup>14</sup> The unobserved transaction prices of each model are falling over time but a full-sample hedonic index based on posted prices would not account for this measurement error, and thus would understate the rate of quality-adjusted price decline.<sup>15</sup> In contrast, the introduction-period hedonic index could correctly capture trends in quality-adjusted prices. This index would omit observations in which prices were measured with substantial error, and the performance variables in the regression would control for improvements in quality in successive periods. The introduction-period index would be unbiased even if there are unobserved discounts at the time of introduction provided that these discounts do not vary systematically over time or across models.

In addition, the full-sample hedonic index can be biased even if the posted life-cycle prices represent actual transaction prices, as shown in the second scenario. For this scenario, we assume that chips actually sell over their life cycle at the flat posted prices we observe in the latter part of our sample period. In the face of innovation that reduces constant-quality prices at the frontier, the absence of price cuts for existing MPUs implies that these chips become progressively more expensive as they age relative to newly-introduced models. With increasingly unattractive pricing, demand would wane as models age.

If model-level data on shipments or sales were available, a shipments- or sales-weighted hedonic index would account for the declining importance of older models and thus would provide an unbiased picture of pricing trends. However, model-level quantity data are not readily available. Accordingly, we (and other researchers) are forced to consider price indexes that put equal weight on every observation. An unweighted full-sample hedonic index would put too much weight on price observations for which there were few transactions. In contrast, an unweighted introduction-period hedonic index likely would do a better job of capturing the trend in quality-adjusted prices. By focusing on prices at the beginning of each model's life cycle, a regression that applies equal weights to all observations avoids over-weighting models whose quantities have dropped off.

Before estimating an introduction-period regression, one must decide whether to use only the very first price observed for a new model or to include some additional prices that extend further into the model's life cycle. Using only

<sup>14</sup>It is not possible to determine whether Intel did or did not follow this pattern of discounting as the discounts are unobservable. That being said, Scherer (2011) provides an interesting discussion of the reasons why Intel, as an oligopolist in competition with AMD, might want to discount (p. 49–54), and he cites a report indicating that Intel's rebates to Dell during 2003–2007 amounted to \$4.3 billion.

<sup>15</sup>Flamm (2015) expresses skepticism about the unobserved discount hypothesis. Among other concerns, he finds no significant break after 2006 in the relationship between Intel's posted MPU prices and a sample of MPU prices he collected from a retail web site. (The retail market is where small purchasers or small computer manufacturers can buy MPUs. These prices are readily available from online sources). Flamm's finding, however, does not necessarily shed any light on Intel's discounting practices. In particular, Intel could be providing unobserved discounts to larger purchasers, while charging the posted prices to smaller purchasers, who could also buy the chips online in the retail market. In that situation, Flamm's results would imply that there had been no change in the pricing relationship between the two sources of small-volume MPU purchases but would provide no information on pricing for large-volume buyers.

the first price offers the greatest protection against age-related measurement error, but it does so at the cost of greatly reducing the sample size. There is also the possibility that the purchasers of brand-new models have preferences that are not representative of the broader market. Using prices beyond the first observed price helps address both of these concerns, but it increases the risk that age-related measurement error will affect the results. There is no clearly correct way to proceed. Our judgment is that the severe reduction in sample size and the potential for unrepresentative buyers weighs against using only the first observed price. Instead, we 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.

To sum up, we emphasize early-price hedonic indexes over full-sample indexes, although we also report the latter. Our preference does not reflect a belief that the early-price indexes are inherently better, but rather the view that they are likely to be more robust to measurement error in posted prices (the age-related discount scenario) and the lack of model-level shipments data that prevents the use of weighted regressions (the drop-off in volume scenario).

## 4. DATA

### 4.1. *Prices and SPEC Scores*

Our MPU prices are collected from publicly available Intel price lists for the period from 1999 to 2013.<sup>16</sup> Intel announces wholesale list prices several times a year for MPUs sold in multiples of 1,000. Unlike single units sold in retail channels, these “trays” of MPUs do not include a cooling system and carry a shorter warranty. Models are identified by family (e.g., Core i7, Pentium, Core 2 Duo), model ID (e.g. i7-4960X), and selected technical characteristics (for example, amount of cache memory or clock speed). We merged these price lists to create price data at a quarterly frequency. We restrict our attention to the 373 MPU models for desktop computer systems introduced between 2000 and 2013.

To measure the relative quality of different chips we use the end-user performance scores from SPEC that were mentioned in Section 2. (These benchmark tests are described in detail in Appendix A). Briefly, SPEC scores evaluate performance of an MPU on individual tasks that rely heavily on integer computation (such as word processing) and on tasks that rely heavily on floating-point computation (such as speech recognition). Scores for individual tasks are measured in seconds, although SPEC rescales these scores so that higher scores indicate better performance and the units are no longer in seconds. SPEC provides an overall score both for integer and for floating-point computation, which are calculated as geometric means of scores for 12 integer computation tasks and 17 floating-point computation tasks, respectively. To construct a single measure of performance, we

<sup>16</sup>Price lists for the period from April 1999 to December 2006 were collected from an archived version of a website devoted to computer hardware. Price lists for later dates were obtained directly from Intel’s website. On dates when both sources were available, we confirmed that the website prices matched Intel price lists. A list of MPU models included in our dataset is provided in Appendix B.

take the geometric mean of the overall scores for integer and floating-point tasks.<sup>17</sup>

SPEC scores are widely used to compare the performance of alternative MPUs or computers and also are used as a standard by computer engineers. For example, Hennessy and Patterson (2012, p. 38)—the standard text on computer design and architecture—notes that “One of the most successful attempts to create standardized benchmark application suites has been the SPEC (Standard Performance Evaluation Corporation) . . .”. This textbook relies on SPEC scores to measure the growth in processor performance in recent decades (see Hennessy and Patterson (2012, p. 3, Figure 1.1)).<sup>18</sup>

SPEC benchmarks provide several ways to measure MPU performance. The performance of a single task is measured by the “speed” score and the simultaneous performance of multiple tasks is measured by the “rate” score. The “speed” and “rate” scores differ in their use of parallel processing, an important consideration after the introduction of multi-core MPUs in the mid-2000s. In the “speed” test, a single task may be broken into component calculations to be run on different processing cores on the MPU. In the “rate” test, multiple instances of the same task may be run simultaneously to more fully exploit the potential of the chip. We use the “speed” score as our base case, but the results using the “rate” score (available from the authors) are very similar.

SPEC notes that both tests are designed to limit the influence on performance of other computer components such as networking, the operating system, graphics, or the I/O system. Accordingly, SPEC scores should provide a relatively clean read on MPU performance independent of most other changes in PC characteristics.

#### 4.2. *Sample Selection*

We matched 177 MPU models from our price data, or 53 percent, to at least one performance score published by SPEC. An important question is whether this selection on the availability of SPEC scores could bias our results. Table 2 provides information on this point, showing a range of characteristics for chips with SPEC scores and those without scores.<sup>19</sup> As shown in Panel A of the table, average entry prices for chips with SPEC scores are higher than for chips without SPEC scores in each of the three subperiods (2000–2004, 2005–2009, and 2010–2013). The changes over time also differ between the two groups of chips, as the average entry prices for the chips with SPEC scores fell substantially from 2000–2004 to 2010–2013, while those for the chips without scores were more stable. In addition, more often than not, the chips with SPEC scores have technical features associated with higher

<sup>17</sup>Because the integer and floating-point scores are highly correlated across MPU models ( $\rho \approx 0.98$ ), the geometric mean provides essentially the same information about performance as the two scores separately. As a robustness check, we re-estimated the SPEC regressions reported below after substituting the two separate scores for the geometric mean score and found that the estimated trends in constant-quality prices were virtually unchanged.

<sup>18</sup>In addition, Goettler and Gordon (2011) use benchmark scores from other sources to evaluate chip performance.

<sup>19</sup>The information on characteristics is collected from Intel’s product information database (<http://ark.intel.com>).

TABLE 2  
SAMPLE CHARACTERISTICS: IS SAMPLE SELECTION A PROBLEM?

Panel A: Number of Intel models, characteristics, and price levels, by year of introduction			
	2000–2004	2005–2009	2010–2013
<b>Number of Intel desktop MPU models</b>			
With SPEC available	32	72	73
No SPEC available	28	48	82
Universe	60	120	155
<b>Entry price (\$)</b>			
With SPEC available	511	415	210
No SPEC available	198	207	170
<b>Clock speed (ghz)</b>			
With SPEC available	2.38	2.86	3.02
No SPEC available	2.61	2.79	2.76
<b>Number of cores</b>			
With SPEC available	1.0	2.50	3.01
No SPEC available	1.0	1.94	2.88
<b>Maximum thermal design power (watts)</b>			
With SPEC available	69.6	94.1	69.9
No SPEC available	75.7	75.3	57.7
<b>Lithography size (nanometers)</b>			
With SPEC available	140.6	61.1	28.8
No SPEC available	106.8	63.1	28.4

Panel B: Price changes (percent, average annual rate)				
	2000–2004	2004–2009	2009–2013	2000–2013
<b>PPI-like matched-model price index (Intel MPUs)</b>				
With SPEC available	–54	–23	–5	–30
No SPEC available	–42	–21	–4	–24
All desktops	–50	–23	–5	–28
<b>PPI</b>	–48	–26	–6	–28

*Source:* Authors' calculations based on data from System Performance Evaluation Corporation, Intel price lists, and data from <http://ark.intel.com>.

performance: faster average clock speed, a more rapid shift to multicore architecture, and greater average power use (as measured by thermal design power).<sup>20</sup> Thus, there are some differences between the chips with SPEC scores and those without scores.

For our purposes, the essential issue is whether price trends differ across the samples of chips with and without SPEC scores. To examine this question, we construct matched-model indexes using the PPI methodology (labeled PPI-like) for both samples. As shown in Panel B of the table, prices of chips with SPEC scores fell somewhat more rapidly over 2000–2004 than did the prices of chips without scores, but overall the two indexes trace out the same pattern: steep

<sup>20</sup>Thermal design power (TDP) measures the amount of heat generated when running typical software for the chip. The amount of heat generated (measured in watts) is closely related to the MPU's power consumption.

declines in the early years of the sample period followed by a marked slowdown in later years. The similarity of the price trends using the PPI methodology suggests that our results are not biased in a material way by the absence of SPEC scores for some chip models.

#### 4.3. Representativeness of Our Sample Relative to the PPI

A second important question is how well price trends in our sample track those in the PPI. Our sample is for Intel desktop chips, while the scope of the PPI includes desktop chips from other manufacturers (notably including AMD), as well as chips for servers and laptops. As can be seen in Panel B, our PPI-like indexes for Intel desktop chips and the PPI for all MPUs display virtually the same price trends. Accordingly, we are comfortable using our results for Intel desktop chips to draw inferences about the performance of the PPI.

### 5. SPECIFICATION OF THE HEDONIC REGRESSIONS

To fix ideas, we first describe a dummy-variable hedonic specification:

$$(1) \quad \ln(P_{i,t}) = \alpha + \sum_k \beta_k X_{k,i,t} + \sum_t \delta_t D_{i,t} + \varepsilon_{i,t}$$

where  $P_{i,t}$  is the price of chip  $i$  in period  $t$ ,  $X_{k,i,t}$  is the value of characteristic  $k$  for chip  $i$  in period  $t$  (measured in logs or levels, as appropriate),  $D_{i,t}$  is a time dummy variable (fixed effect) that equals 1 if chip  $i$  is observed in period  $t$  and zero otherwise, and  $\varepsilon_{i,t}$  is an error term.

A potential shortcoming of equation (1), highlighted by Pakes (2003) and Erickson and Pakes (2011), is that the coefficients on the characteristic or performance variables are constrained to remain constant over the full sample period. One response to that concern (see Aizcorbe, 2014) is to run a cross-section regression for every time period and then to use results from those regressions to build up a price index. Such an approach is appealing because it provides maximum flexibility for estimated coefficients to change over time and allows the results to be used in price index formulas. However, our sample size is too small to run reliable cross-section hedonic regressions for every quarter or even every year.

As a compromise, we focus on adjacent-period (in our case, adjacent-year) hedonic regressions.<sup>21</sup> Specifically, we estimate the following regression for each overlapping two-year period:

$$(2) \quad \ln(P_{i,t}) = \alpha + \sum_k \beta_k X_{k,i,t} + \delta D_2 + \varepsilon_{i,t}$$

where  $P_{i,t}$  is a price observation for chip  $i$  in year  $t$ . Because our price data are quarterly, there will be as many as four price observations for chip  $i$  within the year. The dummy variable  $D_2$  equals 1 if the price observation is in the second

<sup>21</sup>See Triplett (2004) for a discussion of adjacent-period hedonic regressions.

year of the two-year overlapping period and 0 otherwise. To construct a price index from this sequence of regressions, we spliced together the percent changes implied by the estimated coefficients on the  $D_2$  variables. In our main results, we rely on the SPEC variable to capture the performance of each MPU as experienced by users, so that  $\sum_k \beta_k X_{k,i,t}$  reduces to  $\beta \ln(\text{SPEC}_{i,t})$ .

As noted in Section 4, SPEC updated its suite of performance tests in 2006. We use the older (SPEC 2000) benchmarks for the adjacent-year regressions through 2005–2006 and the newer (SPEC 2006) benchmarks for the adjacent-year regressions beginning with 2006–2007.

As a comparison, we also estimate equation (2) with a set of chip characteristics on the right-hand side instead of the SPEC variable. This alternative regression represents the usual approach to hedonic specification in the literature and allows us to gauge the effect on estimated price trends of controlling directly for performance. The characteristics included in the regression are those shown in Table 2—clock speed (in gigahertz), number of cores, maximum thermal design power (in watts), and lithography size (in nanometers)—plus cache memory (in megabytes), a dummy for whether the chip has a separate graphics processing unit, and the number of threads.<sup>22</sup> We use the natural log of clock speed, thermal design power, lithography size, and cache memory; the number of cores and number of threads enter the regression in levels.

## 6. RESULTS

### 6.1. *Estimated Coefficients*

Table 3 shows estimates of equation (2) using SPEC performance for the overlapping two-year periods during 2000–2006, and Table 4 shows estimates for 2006–2013. The upper panel in each table presents estimates that rely only on prices in the first four quarters of each MPU model’s life cycle (the “early price” regression), while the lower panel presents estimates based on the full sample of price observations. The coefficient on the second-year dummy variable in each adjacent-year regression measures the rate of change in quality-adjusted MPU prices from the first year to the second.

Overall, the regressions explain much of the variation in MPU prices. The adjusted  $R^2$  averages about 0.45 across the early-price regressions shown in Tables 3 and 4 as well as for the full-sample regressions.

Although Pakes (2003) cautions against providing structural interpretations of the coefficients, we note that the coefficients on SPEC performance are uniformly positive and significant at the 1 percent level. These coefficients indicate that higher performing MPU models sell for higher prices and suggest that the performance measure captures an important element of the quality differences across MPU models.

Analogous tables for the regressions that replace SPEC performance with chip characteristics are provided in Appendix C. These regressions fit the data

<sup>22</sup>MPUs can perform multiple threads—sequences of related program instructions—either by employing multiple cores or by sharing resources on an individual core. Thus, additional threads provide a form of parallel processing.



TABLE 3  
REGRESSION RESULTS FOR 2000–2006

Panel A: Early Prices						
	2000–2001	2001–2002	2002–2003	2003–2004	2004–2005	2005–2006
<i>Year dummy</i>	−.981** (.147)	−.318** (.098)	−.304* (.129)	−.635** (.198)	−.552** (.147)	−.387** (.120)
<i>ln Performance</i>	1.16** (.22)	.91** (.15)	1.16** (.19)	3.28** (.51)	3.95** (.53)	2.30** (.37)
Number of Obs.	73	73	60	44	77	102
Adjusted R <sup>2</sup>	.37	.34	.39	.50	.41	.27
Panel B: Full sample						
	2000–2001	2001–2002	2002–2003	2003–2004	2004–2005	2005–2006
<i>Year dummy</i>	−.875** (.097)	−.318** (.072)	−.403** (.090)	−.489** (.115)	−.436** (.124)	−.494** (.092)
<i>ln Performance</i>	1.05** (.15)	.87** (.09)	1.13** (.12)	2.81** (.26)	3.15** (.35)	2.52** (.28)
Number of Obs.	100	111	94	78	107	157
Adjusted R <sup>2</sup>	.46	.44	.47	.62	.43	.35

*Note:* The dependent variable is  $\ln(\text{MPU price})$ ; the regression includes a constant, not shown above. Standard errors are in parentheses. \* and \*\* indicate significance at the 5% and 1% levels, respectively.

TABLE 4  
REGRESSION RESULTS FOR 2006–2013

Panel A: Early prices							
	2006–2007	2007–2008	2008–2009	2009–2010	2010–2011	2011–2012	2012–2013
<i>Year dummy</i>	−.683** (.165)	−.721** (.152)	−.830** (.131)	−.438** (.104)	−1.027** (.105)	−.450** (.072)	−.251** (.065)
<i>ln Performance</i>	2.08** (.33)	2.60** (.34)	2.58** (.26)	2.57** (.21)	2.59** (.25)	2.62** (.19)	2.86** (.16)
Observations	71	96	107	104	148	173	118
Adjusted R <sup>2</sup>	.36	.39	.48	.60	.47	.53	.73
Panel B: Full sample							
	2006–2007	2007–2008	2008–2009	2009–2010	2010–2011	2011–2012	2012–2013
<i>Year dummy</i>	−.743** (.140)	−.533** (.112)	−.555** (.082)	−.177** (.064)	−.436** (.069)	−.469** (.068)	−.203** (.054)
<i>ln Performance</i>	2.22** (.28)	2.19** (.22)	1.73** (.14)	1.53** (.09)	1.14** (.09)	1.21** (.11)	2.98** (.16)
Observations	88	141	187	217	281	294	214
Adjusted R <sup>2</sup>	.41	.41	.46	.57	.35	.29	.63

*Note:* The dependent variable is  $\ln(\text{MPU price})$ ; the regression includes a constant, not shown above. Standard errors are in parentheses. \* and \*\* indicate significance at the 5% and 1% levels, respectively.

TABLE 5  
 RATES OF CHANGE IN MPU PRICES<sup>1</sup> (AVERAGE ANNUAL PERCENT CHANGE OVER PERIODS SHOWN)

	2000–2004	2004–2009	2009–2013	2000–2013
Hedonic, SPEC performance				
<b>Early prices</b>	<b>–42</b>	<b>–46</b>	<b>–42</b>	<b>–44</b>
Full sample	–40	–42	–27	–37
Hedonic, chip characteristics				
Early prices	–41	–18	–14	–25
Full sample	–41	–25	–12	–27
PPI	–48	–26	–6	–28

<sup>1</sup>The results for the hedonic regressions have been bias-corrected for the conversion from the natural log of price to the price level. See the text for details.

Source: Authors' calculations and Bureau of Labor Statistics (<http://www.bls.gov>).

well, with an average adjusted  $R^2$  slightly above 0.70. A number of the estimated coefficients are in line with expectations. In particular, the coefficients on clock speed, number of threads, and cache memory are generally positive and significant, with no instances of significant negative coefficients. In contrast, the significant coefficients on thermal design power, lithography, and the graphics processor dummy change sign at least once. In addition, the significant coefficients on the graphics processor dummy are generally negative. Overall, these regressions exhibit the interpretational issues that often attend hedonic regressions estimated with a vector of product characteristics.<sup>23</sup>

## 6.2. Price Indexes

To construct annual price indexes from the SPEC and characteristics regressions, we set the 2000 value of the index to 100, and then move the index forward year by year with the implied percent change from each adjacent-year regression. For example, to calculate the percent change from 2000 to 2001, we exponentiate the coefficient on the year dummy in the 2000–2001 regression.<sup>24</sup> We then do the same for the 2001–2002 regression, and so on.

We summarize our results in Table 5 and Figure 5. The table reports average rates of price change over 2000–2004, 2004–2009, and 2009–2013 from five

<sup>23</sup>For completeness, we also estimated a version of the hedonic regression that included both SPEC performance and chip characteristics as controls for quality. The results were not crisp, as might be expected given the inclusion of both a summary measure of performance and chip features that influence the summary measure. In particular, the coefficient on SPEC performance was positive and significant in a bit less than half of the regressions, insignificant in about half, and negative and significant in two others. The coefficients on the characteristics, though broadly similar to the characteristics-only regression, displayed more frequent instances of significant changes in sign.

<sup>24</sup>Because the exponential function is nonlinear, the translation from the natural log of prices to price levels requires an adjustment in order to be unbiased. We apply the usual adjustment for adjacent-period regressions, which is based on the standard error of the estimated coefficient  $\delta$  on the second-year dummy; see van Dalen and Bode (2004) and Triplett (2004) for details. This adjustment had very little effect on the estimated price trends.

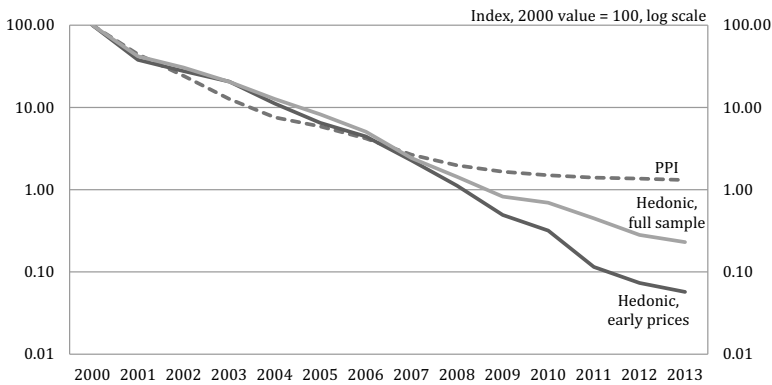


Figure 5. MPU Price Levels

*Note:* Hedonic indexes are based on regressions that use SPEC performance to control for quality.

*Source:* Bureau of Labor Statistics (<http://www.bls.gov/ppi>) and authors' calculations.

different measures: the hedonic index based on SPEC performance and early prices (our preferred index), three other hedonic indexes including those using chip characteristics, and the PPI.<sup>25</sup> The figure plots the levels of the PPI and the two hedonic indexes based on SPEC performance.

From 2000 to 2004, all of the indexes show very rapid declines in MPU prices. As discussed in Section 3, with Intel's ubiquitous downward re-pricing of existing chips through the early 2000s, all of the price indexes—both matched model and hedonic—would be expected to capture the downward trend in quality-adjusted prices. This expectation is borne out by our results. As shown in Table 5, the hedonic indexes fell at an average annual rate between 40 and 42 percent over 2000–2004, while the PPI declined even more rapidly.<sup>26</sup>

However, the trends in the indexes diverge after 2004. Although our preferred hedonic index remained on essentially the same downward trend after 2004 as before, the decline in the PPI slowed sharply. Indeed, the PPI fell at an average annual rate of only 6 percent from 2009 to 2013. For the reasons highlighted earlier in the paper, we believe that this divergence points to likely bias in the PPI for

<sup>25</sup>The Federal Reserve Board also calculates an MPU price index as part of its statistical program to track industrial production and capacity utilization. The Federal Reserve's MPU price series is constructed as a matched-model index through 2006, which is linked to a hedonic index after 2006. Through 2006, the Fed series closely tracks the PPI. For later years, the Fed's series is based on the hedonic indexes in an earlier version of this paper. Because those indexes from the earlier version of this paper have since been revised, the Fed series is not fully in sync with the results in Table 5. In the future, the Fed series will be revised to reflect the hedonic results in this paper.

<sup>26</sup>The faster decline in the PPI compared with the hedonic indexes is consistent with the findings in Silver and Heravi (2005) for several types of consumer durable goods. Silver and Heravi show analytically that this pattern can arise when new models enter the market at prices above the estimated hedonic surface or old models exit the market at prices below the estimated hedonic surface.

MPUs and suggests that the PPI could be providing a deeply misleading picture of price trends for MPUs in recent years.<sup>27</sup>

The full-sample hedonic index based on SPEC performance declined a bit less rapidly than the early-price index over 2004–2009, and the gap widened substantially after 2009. As indicated by the analysis in Section 3, the widening gap is what would be expected if Intel were keeping posted (but not transaction) prices for older models fixed to a greater degree than previously. In all likelihood, the full-sample regression either uses posted prices that have become increasingly disconnected from transaction prices or, if the transaction prices do in fact mirror the posted prices, overweights the sparse transaction volume for older models at very high quality-adjusted prices. Either way, the full-sample hedonic index would understate the rate of price decline, supporting our preference for the early-price index.

As a robustness check for the early-price SPEC regression, we estimated an alternative SPEC regression that included only chip prices in the quarter of introduction. The price index obtained from this introduction-period regression fell at (bias-adjusted) annual rates of 54 percent, 46 percent, and 47 percent, respectively, over 2000–2004, 2004–2009, and 2009–2013. The rates of decline over 2000–2004 and 2009–2013 are somewhat faster than those implied by the early-price version of the regression. But despite these differences, the central message from the introduction-period regression is the same as from the early-price regression—MPU price declines have remained rapid.<sup>28</sup>

Table 5 also shows price trends from the regressions that use chip characteristics rather than SPEC performance. In these results, which are broadly similar to those in Flamm (2015), the post-2004 slowing in estimated price declines is much more pronounced than in those using SPEC performance. In particular, during 2004–2009 and 2009–2013, the price indexes obtained from the characteristics-based regressions fell at average annual rates ranging from only 12 percent to 25 percent, a far slower rate of decline than in the corresponding SPEC-based indexes.

The wide gap stems from the post-2004 divergence between the continued performance gains indicated by SPEC scores and the nearly flat path for chip quality implied by the characteristics regression (which we demonstrate below). With little quality improvement, the characteristics regression “thinks” that constant-quality prices are no longer falling much more rapidly than observed list prices.

<sup>27</sup>In a closely related study, Sun (2014) estimated hedonic regressions for MPUs used in laptops. Because SPEC scores were not available for a wide enough set of laptop chips, Sun used a variety of other performance benchmarks in the regressions. She found that prices for laptop MPUs trended down at a 20 to 30 percent average annual pace over the past ten years, depending on the regression specification. The rate of price decline slowed after 2010, though not to the extent shown by the PPI.

<sup>28</sup>At the suggestion of one of the referees, we also ran a version of the full-sample SPEC regression that included the chip’s model age as an additional explanatory variable. In the early years of the sample period, the coefficient on model age is negative and significant, indicating that older models sell at a lower price than newer models, all else equal. Controlling for model age results in a slower estimated rate of constant-quality price decline. This finding implies that part of the constant-quality price decline for an MPU from one year to the next is due to the approach of obsolescence. However, in the later years of the sample period, the coefficient on model age flips sign to become positive and significant. The positive coefficient on model age in the later years of the sample is consistent with the hypothesis discussed above that Intel offered discounts on older chips that were not reflected in their posted prices. Under this hypothesis, the posted prices used in the full-sample regression would be increasingly too high relative to the model’s characteristics as the model aged, which would show up as a positive coefficient on age in the later part of the sample.

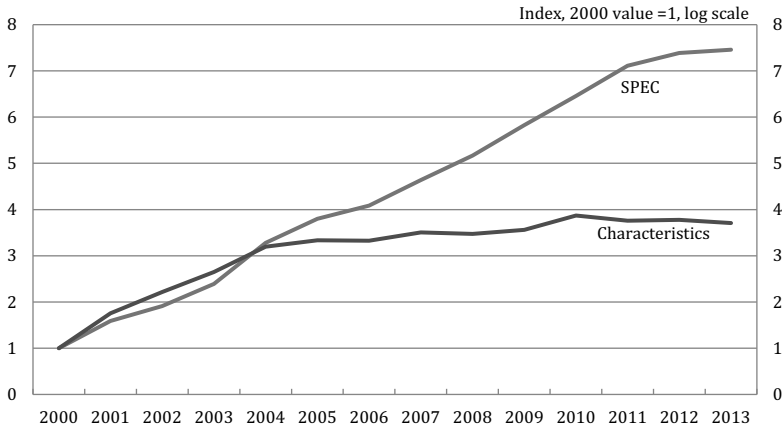


Figure 6. Effect of Chip Quality on Price

Source: Authors' calculations. See text for details.

### 6.3. Implied Chip Quality

To measure the implied change in chip quality over time, recall from equation (2) that the effect of quality on the log of observed chip price in the adjacent-year regression is  $\sum_k \beta_k X_{k,i,t}$ . Letting  $t_1$  and  $t_2$  denote the first and second of the two adjacent years, the average value of the quality effect for the chips that appear in the regression in  $t_1$  is  $\sum_k \beta_k \bar{X}_{k,t_1}$ , with an analogous expression for  $t_2$ . Thus, the change in the quality effect from  $t_1$  to  $t_2$  is  $\sum_k \beta_k (\bar{X}_{k,t_2} - \bar{X}_{k,t_1})$ .<sup>29</sup> In words, this is the change in the average quantity of each characteristic, with each change weighted by the estimated marginal value from the hedonic regression. When SPEC performance is used as the sole control for chip quality, this expression reduces to  $\beta(\bar{X}_{t_2} - \bar{X}_{t_1})$ , where  $X$  is  $\ln(\text{SPEC})$ . It is important to note that  $\sum_k \beta_k (\bar{X}_{k,t_2} - \bar{X}_{k,t_1})$  measures the effect of quality change on prices, not the pure change in quality itself. However, on the reasonable assumption that higher quality is associated on average with higher chip prices, the sign of  $\sum_k \beta_k (\bar{X}_{k,t_2} - \bar{X}_{k,t_1})$  indicates whether chip quality is improving, worsening, or remaining unchanged, which is sufficient for our purpose.

Figure 6 shows the time series for the price effect of chip quality, measured from each adjacent-year regression as  $\sum_k \beta_k (\bar{X}_{k,t_2} - \bar{X}_{k,t_1})$ , with each year-pair then linked together to form the time series. The  $\beta$  coefficients and the average chip characteristics are both taken from the early-price version of the regressions, and thus represent the price effect of quality close to the frontier. As can be seen, the SPEC-based series rises every year, implying a sustained increase in chip quality. Although the rate of increase slows near the end of the sample period, the overall pattern is consistent with the direct measure of SPEC performance that

<sup>29</sup>Aizcorbe (2006) also measures the price effect of changes in chip quality. Although her method involves taking the difference between the changes in observed prices and quality-adjusted prices, algebraically this is very close to what we do.

was shown previously in Figure 1. In contrast, the series based on chip characteristics in Figure 6 rises quite slowly from 2004 to 2010 and is completely flat from 2010 to 2013. Interestingly, the characteristics-based series in Figure 6 bears a striking resemblance to the series for clock speed in Figure 1. Thus, even though the characteristics-based regression includes seven MPU characteristics, clock speed appears to exert a powerful influence on the implied measure of chip quality in the regression.

These results raise questions about the constant-quality MPU price indexes obtained from the regressions that control for quality with chip characteristics. To view those indexes as credible, one must accept either that the quality of MPU chips improved very slowly after 2004 or that the market placed very little value on rising performance over that long period. The second condition seems implausible on its face, while the first implies that the upward march in chip quality indicated by SPEC ratings—the industry standard for performance measurement—is spectacularly wrong. The much more likely conclusion, in our view, is that chip performance has continued to improve and that constant-quality MPU prices have remained on a steep downtrend.

## 7. CONCLUSION

After falling rapidly through the mid-2000s, the PPI for MPUs has declined very slowly by historical standards in recent years. Such a slowdown is puzzling given evidence of ongoing rapid advances in semiconductor technology. To reconcile these observations, this paper demonstrates that the matched-model procedure used for the PPI for MPUs likely is inappropriate given the changes in the properties of Intel's MPU prices after the early 2000s. We argue that a hedonic approach based on prices from the early part of the chip's life cycle is the preferred way to measure quality-adjusted MPU prices. We implement this hedonic approach with an MPU performance measure that addresses longstanding concerns in the literature about the use of product characteristics to proxy for performance.

The results from our preferred hedonic price index indicate that quality-adjusted MPU prices have continued to fall rapidly, contrary to the picture from the PPI. Our results are consistent with other indicators of continued rapid technical progress in the semiconductor sector. Concerns that the semiconductor sector had faded as an engine of growth over the period covered by our analysis appear to be unwarranted.

## ACKNOWLEDGMENTS

We received helpful comments from two anonymous referees, Ana Aizcorbe, Ernie Berndt, Erwin Diewert, Ken Flamm, Charlie Gilbert, Robert Hill, Dale Jorgenson, Kevin Krewell, Catherine Mann, Chris Nosko, Unni Pillai, Paul Thomas, Neil Thompson, Vincent Russo and other analysts in the Producer Price Index program at the Bureau of Labor Statistics, as well as seminar participants at the Federal Reserve Banks of Boston and Cleveland, the 2013

IARIW/UNSW Conference and Workshop, the NBER weekly productivity seminar, the 2014 CRIW/NBER Summer Institute, and the Wellesley Economics Department brown-bag series. Unni Pillai generously shared MPU performance data from SPEC prior to 1999. The views expressed herein are those of the authors and do not necessarily reflect the views of institutions with which we are affiliated.

## REFERENCES

- Aizcorbe, A., "Why Did Semiconductor Price Indexes Fall So Fast in the 1990s? A Decomposition," *Economic Inquiry*, 44, 485–96, 2006.
- , *A Practical Guide to Price Index and Hedonic Techniques*, Oxford University Press, 2014.
- Aizcorbe, A., C. Corrado, and M. Doms, "When Do Matched-Model and Hedonic Techniques Yield Similar Measures?" Federal Reserve Bank of San Francisco Working Paper No. 2003–14, 2003. <http://www.frbsf.org/economic-research/files/wp03-14bk.pdf>
- Aizcorbe, A., S. D. Oliner, and D. E. Sichel, "Shifting Trends in Semiconductor Prices and the Pace of Technological Progress," *Business Economics*, 43, 23–39, 2008.
- Baily, M. N., J. Manyika, and S. Gupta, "U.S. Productivity Growth: An Optimistic Perspective," *International Productivity Monitor*, 25, 3–12, 2013. <http://www.csls.ca/ipm/ipm25.asp>
- Bauer, H., J. Viera, and F. Weig, "Moore's Law: Repeal or Renewal," McKinsey & Company, December, 2013. [http://www.mckinsey.com/Insights/High\\_Tech\\_Telecoms\\_Internet/Moores\\_law\\_Repeal\\_or\\_renewal?cid=other-eml-alt-mip-mck-oth-1312](http://www.mckinsey.com/Insights/High_Tech_Telecoms_Internet/Moores_law_Repeal_or_renewal?cid=other-eml-alt-mip-mck-oth-1312)
- Benkard, L. C. and P. Bajari, "Hedonic Price Indexes with Unobserved Product Characteristics, and Application to Personal Computers," *Journal of Business and Economic Statistics*, 23, 61–75, 2005.
- Berndt, E. R. and Z. Griliches, "Price Indexes for Microcomputers: An Exploratory Study," in M. F. Foss, M. E. Manser, and A. H. Young (eds), *Price Measurements and Their Uses*, 63–93, University of Chicago Press and National Bureau of Economic Research, Studies in Income and Wealth, 57, 1993. <http://www.nber.org/chapters/c7798>
- Brynjolfsson, E. and A. McAfee, *The Second Machine Age*, W.W. Norton and Company, New York, 2014.
- Byrne, D. M., S. D. Oliner, and D. E. Sichel, "Is the Information Technology Revolution Over?" *International Productivity Monitor*, 25, 20–36, 2013. <http://www.csls.ca/ipm/ipm25.asp>
- Chwelos, P., "Hedonic Approaches to Measuring Price and Quality Change in Personal Computer Systems," Ph.D. thesis, University of British Columbia, 2000. <https://circle.ubc.ca/handle/2429/11401>
- , "Approaches to Performance Measurement in Hedonic Analysis: Price Indexes for Laptop Computers in the 1990's," *Economics of Innovation and New Technology*, 12, 199–224, 2003.
- Clark, D., "Intel Rechisels the Tablet on Moore's Law," *Wall Street Journal*, July 16, 2015.
- Cole, R., Y. C. Chen, J. A. Barquin-Stolleman, E. Dulberger, N. Halvacian, and J. H. Hodge, "Quality-Adjusted Price Indexes for Computer Processors and Selected Peripheral Equipment," *Survey of Current Business*, 66, 41–50, 1986.
- van Dalen, Jan, and Ben Bode, "Estimation Biases in Quality-Adjusted Hedonic Price Indices, Working paper, October 1, 2004. <http://www.ipeer.ca/papers/vanDalenBodeOct.1,2004,SSHRC%20Paper17.pdf>
- Erickson, T. and A. Pakes, "An Experimental Component Index for the CPI: From Annual Computer Data to Monthly Data on Other Goods," *American Economic Review*, 101, 1707–38, 2011.
- Fisher, F. M., J. J. McGowan, and J. E. Greenwood, *Folded, Spindled, and Mutilated: Economic Analysis and U.S. v. IBM*, MIT Press, 1983.
- Flamm, K., "The Microeconomics of Microprocessor Innovation," manuscript, University of Texas, Austin, 2007.
- , "Causes and Economic Consequences of Diminishing Rates of Technical Innovation in the Semiconductor and Computer Industries," University of Texas at Austin manuscript, 2015.
- Goettler, R. L., and B. R. Gordon, "Does AMD Spur Intel to Innovate More?" *Journal of Political Economy*, 119, 1141–200, 2011.
- Gordon, R. J., "The Postwar Evolution of Computer Prices," NBER Working Paper 2227, April, 1987.
- Grimm, B. T., "Price Indexes for Selected Semiconductors, 1974–96," *Survey of Current Business*, 78, 8–24, 1998.

- Hennessy, J. L. and D. A. Patterson, *Computer Architecture: A Quantitative Approach* (5th ed.), Elsevier Science, Amsterdam, New York, 2012.
- Hobijn, B., "Is Equipment Price Deflation a Statistical Artifact?," Federal Reserve Bank of New York Staff Report #139, November 2001. [https://www.newyorkfed.org/research/staff\\_reports/sr139.html](https://www.newyorkfed.org/research/staff_reports/sr139.html).
- Holdway, M. "Changes in the PPI for Semiconductors Indexes." PPI Detailed Report, January 1997, 10–11.
- , "An Alternative Methodology: Valuing Quality Change for Microprocessors in the PPI," manuscript, 2001. <http://www.bea.gov/papers/pdf/mpuvqa.pdf>
- Intel Corporation, "Excerpts from a Conversation with Gordon Moore: Moore's Law," 2005. [http://large.stanford.edu/courses/2012/ph250/lee1/docs/Excepts\\_A\\_Conversation\\_with\\_Gordon\\_Moore.pdf](http://large.stanford.edu/courses/2012/ph250/lee1/docs/Excepts_A_Conversation_with_Gordon_Moore.pdf)
- Jorgenson, D. W., "Information Technology and the U.S. Economy," *American Economic Review*, 91, 1–32, 2001.
- Moore, G. E., "Cramming More Components onto Integrated Circuits," *Electronics*, 38, April 19, 1965.
- Oliner, S. D. and D. E. Sichel, "The Resurgence of Growth in the Late 1990s: Is Information Technology the Story?" *Journal of Economic Perspectives*, 14, 3–22, 2000.
- Oliner, S. D., D. E. Sichel, and K. J. Stiroh, "Explaining a Productive Decade," *Brookings Papers on Economic Activity*, no. 1, 81–152, 2007.
- Pakes, A., "A Reconsideration of Hedonic Price Indexes with an Application to PC's," *American Economic Review*, 93, 1578–96, 2003.
- Pillai, U., "A Model of Technological Progress in the Microprocessor Industry," *Journal of Industrial Economics*, 61, 877–912, 2013.
- Scherer, F. M., "Abuse of Dominance by High Technology Enterprises: A Comparison of U.S. and E.C. Approaches," *Journal of Industrial and Business Economics*, 38, 39–62, 2011.
- Shenoy, S. R., and A. Daniel, "Intel Architecture and Silicon Cadence: The Catalyst for Industry Innovation," Intel White Paper, 2006.
- Silver, M., and S. Heravi, "A Failure in the Measurement of Inflation," *Journal of Business and Economic Statistics*, 23, 269–81, 2005.
- Song, M., "The Quality Adjusted Price Index in the Pure Characteristics Demand Model," *Journal of Business and Economic Statistics*, 28, 190–9, 2010.
- Sun, S. (Liyang), "What are We Paying For? A Quality-adjusted Price Index for Laptop Microprocessors," Wellesley College senior thesis, 2014.
- Thompson, N., "Moore's Law goes Multicore: The economic and strategic consequences of a fundamental change in how computers work," Sloan School of Management, MIT, 2015.
- Triplett, J. E., "Price and Technological Change in a Capital Good: A Survey of Research on Computers," in Dale W. Jorgenson and Ralph Landau (eds), *Technology and Capital Formation*, 127–213, MIT Press, 1989.
- , "Handbook on Hedonic Indexes and Quality Adjustments in Price Indexes: Special Application to Information Technology Products," Organisation for Economic Co-operation and Development, STI Working Paper 2004/9, 2004. <http://www.oecd.org/sti/33789552.pdf>

## SUPPORTING INFORMATION

Additional supporting information may be found in the online version of this article at the publisher's web-site:

**Appendix A:** Performance Measures From SPEC (System Performance Evaluation Corporation)

**Appendix B:** Intel Desktop MPUs On Wholesale Price Lists

**Appendix C:** Regression Results Using MPU Characteristics