

Oliner, Stephen D.; Sichel, Daniel E.; Byrne, David M.

Working Paper

How fast are semiconductor prices falling?

AEI Economics Working Paper, No. 2014-06

Provided in Cooperation with:

American Enterprise Institute (AEI), Washington, DC

Suggested Citation: Oliner, Stephen D.; Sichel, Daniel E.; Byrne, David M. (2015) : How fast are semiconductor prices falling?, AEI Economics Working Paper, No. 2014-06, American Enterprise Institute (AEI), Washington, DC

This Version is available at:

<https://hdl.handle.net/10419/280518>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



AMERICAN ENTERPRISE INSTITUTE

How Fast are Semiconductor Prices Falling?

David M. Byrne
Federal Reserve Board

Stephen D. Oliner
American Enterprise Institute and UCLA

Daniel E. Sichel
Wellesley College and NBER

AEI Economic Policy Working Paper 2014-06
November 2015

©2015 by the American Enterprise Institute for Public Policy Research. All rights reserved.

The American Enterprise Institute for Public Policy Research (AEI) is a nonpartisan, nonprofit, 501(c)(3) educational organization and does not take institutional positions on any issues. The views expressed here are those of the author(s).

How Fast Are Semiconductor Prices Falling?

David M. Byrne, Federal Reserve Board, david.m.byrne@frb.gov

Stephen D. Oliner, American Enterprise Institute and UCLA, stephen.oliner@aei.org

Daniel E. Sichel, Wellesley College and NBER, dsichel@wellesley.edu

September 2015

Abstract

The Producer Price Index (PPI) for the United States 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. Roughly coincident with the shift to slower price declines in the PPI, Intel — the leading producer of MPUs — substantially changed its pricing behavior for these chips. As a result of this change, we argue that the matched-model methodology used in the PPI for MPUs likely started to be biased in the mid-2000s and that hedonic indexes can provide a more accurate measure of price change since then. Our preferred hedonic index of MPU prices tracks the PPI closely through 2004. However, from 2004 to 2008, our preferred index fell faster than the PPI, and from 2008 to 2013 the gap widened further, with our preferred index falling at an average annual rate of 43 percent, while the PPI declined at only an 8 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.

KEY WORDS: measurement, hedonic price index, quality adjustment, technological change, microprocessor

We received helpful comments from 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.

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 pace of innovation in that sector.¹

A stalling out of innovation in this sector likely would have broader implications for the economy, as semiconductors are an important general-purpose technology lying behind advances of the digital revolution, including machine learning, robotics, big data, and massive connectivity.² Given this linkage, adverse developments in the semiconductor sector could damp the growth potential of the overall economy.³ On the other hand, if technological progress and attendant price declines were to continue at a rapid pace, powerful incentives would be in place for continued development and diffusion of new applications of this general-purpose technology. Such applications could both enhance the economy's growth potential and push forward the ongoing automation that has generated concerns about job displacement.⁴

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

¹ See Aizcorbe, Oliner, and Sichel (2008) for a discussion of the relationship between price change and innovation for semiconductors.

² Brynjolfsson and McAfee (2014) and Baily, Manyika, and Gupta (2013) highlight key innovations in many sectors of the economy.

³ For discussions of the sources of a possible slowdown in the economy's potential growth rate, see Cowen (2011), Fernald (2014), Gordon (2012, 2014a, and 2014b), Hall (2014), and Jorgenson, Ho, and Samuels (2014).

⁴ For a recent discussion of these employment issues, see Hamilton Project (2015) for the transcript of a wide-ranging Policy Forum.

resolving the puzzle may reside in another development in the semiconductor industry. Roughly coincident with the shift to a slower pace of price decline in the PPI, Intel — the leading MPU producer — dramatically changed its pricing behavior. Prior to the mid-2000s, Intel generally lowered the list prices of existing chips as they aged and as new, higher-performance models were introduced. However, after the mid-2000s, Intel shifted to a new paradigm in which it largely kept the list prices of existing chips unchanged.

This change in Intel's behavior could resolve the puzzling disconnect between continuing improvements in MPU performance and the shift to slower price declines. Namely, we argue that the matched-model methodology used by BLS is not well suited to capturing price trends in a pricing regime like the one that has prevailed since the mid-2000s. In such a pricing regime, if performance is improving over time, then matched-model price indexes for MPUs such as the PPI likely are biased. We argue that hedonic indexes are better suited to capturing price trends in these circumstances 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 in recent years, identifying the correct set of technical characteristics (and likely changes in that set) could be challenging.

Our preferred index of MPU prices tracks the PPI closely from 2000 to 2004 with both indexes falling at an average pace of about 50 percent a year. However, from 2004 to 2008 —

roughly the period of transition to Intel's new pricing regime — our preferred hedonic index fell faster than the PPI. And from 2008 to 2013, the gap widened dramatically, with our preferred index falling at an average annual rate of 43 percent, while the PPI declined at only an 8 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.⁵ 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

⁵ 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.

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.⁶

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) and Berndt and Griliches (1993) called for analysis of computer prices with more refined controls for performance.⁷ 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.⁸ Our research takes the same approach to meeting the concerns raised by Triplett and by Berndt and Griliches.

Another strand of the literature has measured semiconductor prices with matched-model indexes. Aizcorbe et al. (2003) found that under certain circumstances (highly granular data on model prices and high-frequency observations), a matched-model index may produce similar results to a hedonic index. Such highly granular data were used by the Federal Reserve Board for its MPU price index through 2006, which was constructed as a matched-model index. The

⁶ The market for chips going into smartphones and tablets is changing rapidly, and it is challenging to obtain consistent time series on prices and performance. That being said, the production processes and technologies are similar enough to those for MPUs that developments in MPUs likely provide at least a rough gauge of developments in these other types of chips.

⁷ See pages 147 and 91, respectively.

⁸ 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.

source data used by the Fed were discontinued after 2006. However, even if those data were still available, the issues explored in this paper suggest that the matched-model approach would be problematic for more recent years.

Our findings for MPU prices have important implications for gauging the pace of innovation in the semiconductor sector and, more broadly, for the tech sector as a whole. At the same time, they have limited direct implications for the measurement of real GDP or output per hour. Semiconductors mostly are intermediate inputs and so are not counted directly in GDP. Imports and exports are the exception, but trade quantities are small enough on net that, even if our preferred price indexes for MPUs were adopted, measures of real GDP growth would change by only a small amount.

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. This section also discusses the change in the pricing pattern of Intel MPUs that occurred in the mid-2000s. 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) since that change. 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 the plausibility of this extreme slowdown from two different perspectives. The first perspective is whether the slowdown meshes with the trends in technological advance for MPU chips. The

second perspective focuses on the changes that occurred in Intel's pricing patterns for MPUs, developments that, as we show below, have important implications for price index measurement.

2.1 Technology Cycles and Chip Performance⁹

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 (0.7×0.7) 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.¹⁰

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

⁹ The text in this section draws heavily on material in Byrne, Oliner, and Sichel (2013).

¹⁰ 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 be two years based on the actual experience to that point (Intel Corporation, 2005). For a discussion of the outlook for Moore's Law, see Bauer, Veira, and Weig (2013).

years.¹¹ 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) showed this development with performance scores from the Standard Performance Evaluation Corporation (SPEC).¹² 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.¹³ How did this strategy shift affect the rate of increase in performance for end users? Figure 1 extends Pillai's analysis, which ends in 2008, through 2013, plotting average SPEC performance measures for MPUs introduced each year.¹⁴ Starting in 2004, the figure includes a second SPEC performance measure (the "rate" score) that incorporates more fully the

¹¹ See Don Clark, "Intel Rechisels the Tablet on Moore's Law," *The Wall Street Journal*, July 16, 2015.

¹² SPEC is a non-profit corporation that establishes performance benchmark tests for computing equipment and publishes test results submitted by member organizations.

¹³ 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).

¹⁴ 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.

performance gain from using multiple cores than does the standard ("speed") measure.¹⁵ The use of multiple cores enables greater parallelization of processing, which enhances performance.¹⁶

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 2004-13, 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.¹⁷

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 1993-97, and the PPI beginning in 1998, when BLS adopted their current methodology in response to research at the Federal Reserve Board and the Bureau of Economic Analysis. Over this 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 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 accelerated in recent years. If that were the case, we might expect to see a similar price pattern for other types of

¹⁵ See section 4 for further discussion of these alternative SPEC tests.

¹⁶ Thompson (2015) describes how firms able to exploit the parallelization possible on computers with multicore chips realized greater productivity gains than firms that were not able to exploit these innovations.

¹⁷ 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.

semiconductor chips. No such pattern, though, is evident for memory chips (DRAMs).¹⁸ As shown in figure 3, DRAM prices have been quite volatile from year to year, with no clear trend toward slower declines. At least for DRAMs, then, the translation of the engineering frontier to prices does not appear to have changed.

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. We next consider the possibility that changes in the properties of Intel's posted prices after 2006 could have distorted the measurement of price trends in the PPI.

2.2 Intel's Pricing Behavior and Market Position

Between 2003 and 2006, the properties of Intel's posted prices for MPU chips changed dramatically.¹⁹ Prior to 2003, the price of a specific Intel MPU model tended to drop fairly rapidly in the year or two following its introduction, especially once a new, higher performance model became available. By 2006, this pattern had completely changed; the posted price of a specific model tended to remain constant, even after a new, higher performance model became available at a similar price.

Figure 4 illustrates this shift in pricing behavior. The upper panel shows Intel's wholesale list prices for all desktop models during 2000-01, while the lower panel shows the analogous list prices for all desktop models during 2011-12. The difference between the two panels is stark. In the early period, prices fell steeply over a model's lifecycle. However, by

¹⁸ The most important driver of DRAM performance is the size of the components on the chip, which has continued to shrink rapidly. See Flamm (1993) for a discussion of characteristics of DRAM chips.

¹⁹ While the BLS does not indicate which companies' prices are included in the PPI for MPUs, it is only reasonable to assume that Intel is included given its dominant market position. Moreover, it appears that BLS has been making at least some use of Intel's posted prices for MPUs (see Holdway, 2001).

2011-12, the price paths are flat or nearly so, with only a few instances of sizable price declines. Figure 5 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.²⁰ 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 roughly 10 percent for the 2012 cohort.

This shift took place against the backdrop of a changing competitive environment vis-a-vis AMD, Intel's primary challenger in the MPU market.²¹ Figure 6 plots Intel's annual revenue from all products (not only MPUs) as a share of Intel's and AMD's combined revenue. Over this 20-year period, Intel's revenue share ranged from just under 80 percent to slightly above 90 percent. In 2002 and 2003, Intel's share stood at the top of this range but then eroded significantly through 2006. Nosko (2010) also documented this share erosion, noting that from 2002 to 2006, AMD "consistently released products whose price/performance characteristics were similar to or beat Intel's" (p. 8). Documents from legal actions against Intel for alleged antitrust violations also provide evidence that Intel was concerned about its losses of market share to AMD during this period (see, e.g., *State of New York v. Intel Corporation* (2009)).

However, in 2006, Intel began to turn the corner by offering new MPUs that dominated those sold by AMD (see Nosko (2010) for details) and subsequently pulled ahead of AMD in a decisive way, consistent with the recovery in Intel's overall revenue share. By 2013, Intel's dominance had reached the point that AMD effectively had been relegated to the bottom end of

²⁰ The price data in the figure are described in section 4.

²¹ See Goettler and Gordon (2011) for a thorough analysis of competition between Intel and AMD over 1993-2004 and the attendant effects on Intel's rate of innovation. For a further discussion of Intel's competitive and legal environment, see Scherer (2011).

the MPU market. Indeed, as shown in figure 7, AMD received more than 85 percent of its MPU revenue from chips that sold for less than \$75, the low-end part of the market where Intel derived only about 15 percent of its revenue. Conversely, Intel received almost 30 percent of its MPU revenue from chips that sold for \$200 or more, a segment in which AMD earned no revenue.²²

This evidence indicates that Intel has enhanced its position in the MPU market since 2006. With less competition from AMD, Intel has greater scope to manage its price structure.²³ By keeping list prices essentially fixed over a chip's life cycle, Intel may be attempting to extract more revenue from price-insensitive buyers while offering discounts on a case-by-case basis to other customers. Such price discrimination could reduce the information content of its posted list prices, potentially biasing the quality-adjusted indexes generated from these prices. The next section investigates the properties of matched-model and hedonic indexes in the post-2006 pricing environment.

3. CONSTRUCTING MPU PRICE INDEXES

As noted, the changes in the pattern of Intel's posted MPU prices raise important questions for price measurement. Can a matched-model index correctly capture price trends given Intel's new pricing regime? If not, what approach to price measurement is likely to deliver estimates that are most robust to possible measurement error in prices? Finally, in an environment with rapid

²² Figure 7 slightly overstates the market separation between Intel and AMD for the following reason. The source data used to construct the figure show the average selling prices for distinct product families (Intel's Core i5 Duo, for example) rather than prices for individual chips. Within each product family, there is some amount of unobserved price variation. The average selling price for AMD's most expensive product family, the FX-8, was less than \$200 in 2013, but two chips in this family (the FX-9590 and the FX-9370) were priced between \$200 and \$500 at introduction in July 2013 (see http://en.wikipedia.org/wiki/List_of_AMD_FX_microprocessors). These isolated cases do not affect the general point that Intel dominates the high end of the MPU market.

²³ For models of pricing and innovation behavior in the MPU industry, see Aizcorbe and Kortum (2005), Nosko (2010), and Pillai (2013).

gains in MPU performance and changes in chip architecture, what are the best ways to measure quality change for MPUs?

In this section, we investigate these questions and make the case that — given the new pricing regime for MPUs that emerged in the mid-2000s — hedonic indexes are preferable to matched-model indexes. In addition, we argue that, given possible measurement error and current limitations on data availability, hedonic regressions that rely only on prices in the period of a model's introduction are likely to be more robust than hedonic regressions that include prices over a chip's 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 Matched-model or Hedonic?

To fix ideas, we briefly describe the matched-model and hedonic approaches to price measurement (see Aizcorbe(2014) for a recent, in-depth treatment of this topic). Matched-model indexes control for quality change between two periods, say period 0 and period 1, by averaging price changes for models that were in the market in both periods. The underlying idea is to measure “pure” (or quality-adjusted) price change by linking period 0 and period 1 prices for identical products, whose quality did not change between periods by definition.

Hedonic price indexes identify “pure” price change by explicitly controlling for quality change with measures of product characteristics or performance. The hedonic approach regresses prices on these measures, with the price change left after controlling for quality taken to be the quality-adjusted price change. Many alternative specifications have been proposed for hedonic regressions. Perhaps the simplest is to include time dummy variables (time fixed

effects) in the hedonic regression and to read off quality-adjusted price change from the coefficients on the time dummies. (We discuss in section 5 our preferred specification for the new price indexes in this paper.)

To illustrate why we prefer hedonic over matched-model indexes for MPUs, we use stylized examples of Intel's pricing behavior to highlight the implications of the shift in behavior. Panel A of figure 8 illustrates the pricing pattern for Intel MPUs prior to the mid-2000s in which the price of each MPU model falls as it ages; this panel is a stylized version of the actual price cascade for 2000-01 shown previously. When a new, higher-quality model is introduced, the price of the old model falls so that price *per unit of quality* is at least roughly equalized across models. In this case, the gap between model prices in periods when two models are in the market would represent the value of improved quality, and a matched-model index would correctly adjust for this quality change.

Panel B shows a scenario in which prices of incumbent models do not fall when new, higher-quality models are introduced — a stylized version of the actual price cascade we presented for 2011-12. Given the upward trend in chip quality, price per unit of quality in panel B is falling as new models are introduced. However, a matched-model index would indicate that quality-adjusted prices are constant because the price of each model stays the same as it ages. Thus, a matched-model index would give a biased read of trends in quality-adjusted prices. This scenario illustrates why we argue against the use of matched-model indexes in the period after the shift in Intel's pricing behavior.

Switching to a hedonic index does not automatically solve the bias problem. As we discuss below, care must be taken to estimate the hedonic index 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 after the shift in Intel's pricing behavior 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 (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 — specifically for the types of tasks actually undertaken by users — are available. Moreover, specifying an appropriate hedonic regression and finding suitable data can be difficult for a rapidly evolving product, and using directly measured performance helps to address this problem.

Accordingly, in this paper, we develop hedonic indexes with performance measures from SPEC that avoid 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, Gordon (1987), Cole et al. (1986), and Fisher et al. (1983) 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 a similar set of concerns apply to MPUs. To explain why, we consider two possible explanations for what lies behind the pricing pattern adopted by Intel in the mid-2000s. Panel C of figure 8 shows a stylized scenario in which prices of older models are measured with error. Specifically, posted list prices in this scenario do not represent true transactions prices because Intel offers “age-related discounts” to selected purchasers as models age.²⁴ In this scenario, new models enter the market at the same list price that exiting models leave the market, just as in panel B. We now add the dashed lines, which show actual transaction prices. These lines slope downward, as progressively larger discounts are offered on older models as newer, higher performance models become available.

In the age-related discounting scenario, a full-sample hedonic index based on observable *posted* prices would go astray because the posted prices are measured with error. The actual

²⁴ 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.

transaction prices of each model are falling over time but a full-sample hedonic index would not account for this measurement error, leading to biased estimates of quality-adjusted price change.²⁵

In contrast, the introduction-period hedonic index could correctly capture trends in quality-adjusted prices if age-related discounts were offered as in panel C.²⁶ In particular, the introduction-period hedonic would omit observations in which prices were measured with error, and the performance variables in the regression would control for improvements in quality in successive periods.

Panel D highlights another scenario (“Drop-off in volume”) that also could create challenges for a full-sample hedonic index. In this scenario, posted prices do not change, but in contrast to panel C, we assume they represent actual transaction prices. With this pricing assumption, the quantity purchased of a model drops off as it ages and faces competition from newer, more powerful models with lower prices per unit of quality. The solid, heavier line represents the early, high-quantity part of each model’s life cycle and the thinner line represents the later part when sales have dropped. If model-level data on shipments or sales were available, a shipments- or sales-weighted index would account for the declining importance of the older model. However, model-level quantity data are not readily available. Accordingly, we (and

²⁵ 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 scenario, 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.

²⁶ The introduction-period index would capture price trends even if list prices at the time of introduction do not represent actual transactions prices provided that discounts at the time of introduction do not vary systematically over time or across models.

other researchers) are forced to consider price indexes that put equal weight on every observation.

In the scenario illustrated in panel D, 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.

For the reasons described here, we emphasize introduction-period hedonic indexes, although we also report full-sample indexes. Our preference does not reflect a belief that introduction-period indexes are inherently better, but rather the view that introduction-period indexes are likely to be more robust to the data limitations that hamper efforts to estimate recent trends in quality-adjusted MPU prices. In particular, we believe that an introduction-period hedonic will be more robust to measurement error in list prices (the age-related discount scenario) and to a 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 2013.²⁷ Intel announces wholesale list prices several times a year for MPUs sold in multiples of

²⁷ 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.

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 take the geometric mean of the overall scores for integer and floating-point tasks.²⁸

SPEC scores are widely used to compare the performance of alternative MPUs or computers and also are used as a standard by computer engineers.²⁹ For example, Hennessy and Patterson (2012, p. 38) — the standard text on computer design and architecture — notes that

²⁸ Although the integer and floating-point scores could be included as separate variables in the hedonic regressions described below, little would be gained relative to the equal-weighted geometric mean, as the correlation between the two scores is greater than 0.98.

²⁹ Goettler and Gordon (2011) also use benchmark scores to evaluate chip performance. And, one of our employers uses SPEC scores in making decisions about purchases of computing equipment.

“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, fig. 1.1)).³⁰

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 184 MPU models from our price data, or 49 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

³⁰ See figure 1.1 on page 3.

of chip characteristics both for chips with SPEC scores and for those without scores.³¹ As shown in Panel A of the table, average entry prices for chips with SPEC scores are considerably higher than for chips without SPEC scores. Perhaps not surprising given the difference in average prices, chips with SPEC scores tend to have technical features associated with higher performance. Clock speed is somewhat greater for chips with SPEC scores; the shift to multicore architecture occurred more quickly; and the amount of power used by the chip (thermal design power³²) is greater. For the most part, these differences in technical features taken individually are not very large, and there is no consistent difference in the final characteristic shown in the panel — the size of the smallest feature on the chip (lithography).

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, over the full period from 2000 to 2013, the average rate of decline in these PPI-like price indexes is very similar across the chips with and without SPEC scores (29 percent versus 26 percent, respectively). The two sets of chips also display similar price trends within the shorter time periods shown in the table. The similarity of the price trends suggests that our results are not biased 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

³¹ The information on characteristics is collected from Intel's product information database (<http://ark.intel.com>).

³² 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.

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 very similar 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:

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

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 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 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.³³ 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.

³³ See Aizcorbe's (2014) discussion and the references there.

As a compromise, we focus on adjacent-period (in our case, adjacent-year) hedonic regressions.³⁴ Specifically, we estimate the following regression for each two-year overlapping period:

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

where $P_{i,t}$ is the price of chip i in year t . We measure $P_{i,t}$ as the average of the observed prices of chip i within the year.³⁵ The dummy variable D_2 equals 1 if the price observation is in the second 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 the previous section, SPEC updated its suite of performance tests in 2006. We use the older (SPEC 2000) benchmarks for the adjacent-year regressions through 2005-06 and the newer (SPEC 2006) benchmarks for the adjacent-year regressions beginning with 2006-07.

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

³⁴ See Triplett (2006) for a discussion of adjacent-period hedonic regressions.

³⁵ We also estimated regressions at the quarterly frequency for each two-year period and included a time trend spanning the eight quarters of data included in each regression. That alternative yielded price trends similar to those reported here.

number of threads.³⁶ 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 2000-06, and table 4 shows estimates for 2006-13. The upper panel in each table presents estimates that rely only on prices in each MPU model's introduction period, 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 close to 0.50 across the full set of regressions shown in tables 3 and 4.

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 5 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 well, with an average adjusted R^2 of about 0.75. Some of the estimated coefficients are in line with expectations. In

³⁶ 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.

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, number of cores, 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.³⁷

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 2000-01 regression.³⁸ We then do the same for the 2001-02 regression, and so on.

We summarize our results in table 5 and figure 9. The table reports average rates of price change over 2000-04, 2004-08, and 2008-13 from six different measures: the hedonic index based on SPEC performance and introduction-period prices (our preferred index), three other hedonic indexes including those using chip characteristics, the MPU price index constructed by

³⁷ 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 less than half of the regressions and insignificant in the rest. The coefficients on the characteristics, though broadly similar to the characteristics-only regression, displayed more frequent instances of significant changes in sign.

³⁸ 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 standard adjustment based on the estimated variance of the regression error; see Aizcorbe (2014) for details. This adjustment had very little effect on estimated price trends over 2000-04, but over 2004-08 and 2008-13 it slowed the estimated annual rate of price decline by as much as 10 percentage points relative to the unadjusted estimates.

the Federal Reserve Board, and the PPI.³⁹ 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 before the mid-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, all of the MPU price indexes fell at an average annual rate of 39 percent or more over 2000-04. In addition, the rate of decline in the PPI (and the Fed’s index) is very similar to that of our preferred hedonic index.

However, the trends in the indexes diverge after 2004. Although our preferred hedonic index fell somewhat less rapidly after 2004 than before, the decline in the PPI slowed much more sharply. Indeed, the PPI fell at an average annual rate of only 8 percent from 2008 to 2013 and barely declined at all in 2012 and 2013. For the reasons highlighted earlier in the paper, we believe that this divergence points to likely bias in the PPI for MPUs and suggests that the PPI could be providing a deeply misleading picture of price trends for MPUs in recent years.⁴⁰

The full-sample hedonic index based on SPEC performance declined less rapidly than the introduction-period index in each period, with the largest difference occurring after 2008. As indicated by the analysis in section 3, the wider gap after the mid-2000s is what would be expected under the “age-related discounting” and “drop-off in volume” scenarios if Intel were keeping posted prices for older models fixed to a greater degree than previously. Because

³⁹ The Federal Reserve series is a matched-model index from 1992 to 2006 linked to an introduction-period hedonic index after 2006. The hedonic index is constructed from an earlier version of the empirical work presented here.

⁴⁰ 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.

discounts from posted prices are unobservable, the full-sample index will understate price declines to the extent that Intel is discounting older models; this bias will be larger than when posted prices were more flexible. Moreover, to the extent that Intel is allowing price-performance ratios on older models to rise above those for new models, sales of the older models likely are falling off more substantially than when posted prices fell over a chip's life cycle. Given that we do not have model-level shipments data, the index based on the unweighted full-sample index would overweight the older models whose prices are not changing. Under either scenario (or a combination of them), the full-sample hedonic index will understate the rate of price decline, supporting our preference for the introduction-period index.⁴¹

The post-2004 slowing in estimated price declines is much more pronounced in the regressions that use chip characteristics than in those using SPEC performance. The results reported for hedonic regressions with chip characteristics are broadly similar to those in Flamm (2015). As shown in table 5, during 2004-08 and 2008-13, the price indexes obtained from the characteristics-based regressions fell at average annual rates ranging from only 1 percent to 21 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 flat path for chip quality implied by the characteristics regression (which we demonstrate below). With no quality improvement, the characteristics regression "thinks" that constant-quality prices are no longer falling more rapidly than observed list prices.

⁴¹As a robustness check for the introduction-period SPEC regression, we estimated an alternative version that included not only chip prices in the quarter of introduction but also prices in the three quarters following the chip's introduction. The price index obtained from this regression fell at a (bias-adjusted) annual rate of 38 percent, 34 percent, and 39 percent, respectively, over 2000-04, 2004-08, and 2008-13. Thus, MPU price declines are estimated to have remained rapid whether we use only the introduction-period price or broaden the sample to include prices observed within a year of introduction.

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 $\sum_k \beta_k (\bar{X}_{k,t_2} - \bar{X}_{k,t_1})$.⁴² 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 10 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 β coefficients and the average chip characteristics are both taken from the introduction-period version of the regressions, and thus represent the price effect of quality at the frontier. As can be seen, the SPEC-based series rises every year, implying a sustained increase in chip quality; this result is consistent with the direct measure of SPEC performance that was shown previously in figure 1. In contrast, the series based on chip characteristics in

⁴² 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.

figure 10 only edges up on net after 2003, implying little effect on prices of the improvement in chip quality from that year forward, punctuated by a sizable decline in 2008. Interestingly, the characteristics-based series in figure 10 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 was essentially stagnant for the decade after 2003 or that the market placed no value on rising performance over that decade. The second condition seems implausible to us 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 in the pricing regime that Intel — the dominant manufacturer — adopted in the mid-2000s. We argue that a hedonic approach based on introduction-period prices is the

preferred way to measure quality-adjusted MPU prices in the current pricing environment. 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 continued to fall rapidly after the mid-2000s, contrary to the picture from the PPI. Our results have important implications for understanding the rate of technical progress in the semiconductor sector and, arguably, for the broader debate about the pace of innovation and growth in the U.S. economy. Notably, concerns that the semiconductor sector had faded as an engine of growth over the period covered by our analysis appear to be unwarranted. Nonetheless, Intel's recent announcement that it had moved to a longer technology cycle could presage some slowdown in MPU price declines, with less rapid advances in technologies enabled by semiconductors.

As a final point, we note that our results raise a new puzzle. In recent years, computer prices in the National Income and Product Accounts have fallen quite slowly by historical standards. If MPU prices have, in fact, continued to decline rapidly, why have prices for computers — which rely on MPUs for their performance — not followed suit?⁴³ We believe there is a reasonable chance that the official price indexes for computers have been measured with error, reflecting in part the challenge of identifying the relevant and likely changing set of physical characteristics used to measure quality. The end-user performance measure employed in this paper could help break this knot for computer prices, and we are investigating this possibility in further work.

⁴³ Byrne and Pinto (2015) discuss in some detail several measurement issues that cast doubt on the plausibility of the slowing decline in the NIPA price index for computer investment.

REFERENCES

- Aizcorbe, Ana, "Why Did Semiconductor Price Indexes Fall So Fast in the 1990s? A Decomposition," *Economic Inquiry*, 44(3), 485-96, 2006.
- _____, *A Practical Guide to Price Index and Hedonic Techniques*, Oxford University Press, 2014.
- _____, Carol Corrado, and Mark 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>
- _____, and Samuel Kortum, "Moore's Law and the Semiconductor Industry: A Vintage Model," *Scandinavian Journal of Economics*, 107(4), 603-30, 2005.
- _____, Stephen D. Oliner, and Daniel E. Sichel, "Shifting Trends in Semiconductor Prices and the Pace of Technological Progress," *Business Economics*, 43(3), 23-39, 2008.
- Baily, Martin Neil, James Manyika, and Shalabh Gupta, "U.S. Productivity Growth: An Optimistic Perspective," *International Productivity Monitor*, 25, 3-12, 2013. <http://www.csls.ca/ipm/ipm25.asp>
- Bauer, Harald, Jan Viera, and Florian 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
- Benkard, Lanier C., and Patrick Bajari, "Hedonic Price Indexes with Unobserved Product Characteristics, and Application to Personal Computers," *Journal of Business and Economic Statistics*, 23(1), 61-75, 2005.
- Berndt, Ernst R., and Zvi Griliches, "Price Indexes for Microcomputers: An Exploratory Study," in Murray F. Foss, Marilyn E. Manser, and Allan 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, vol. 57, 1993. <http://www.nber.org/chapters/c7798>
- Brynjolfsson, Erik, and Andrew McAfee, *The Second Machine Age*, W.W. Norton and Company, 2014.
- Byrne, David M., Stephen D. Oliner, and Daniel E. Sichel, "Is the Information Technology Revolution Over?" *International Productivity Monitor*, 25, 20-36, 2013. <http://www.csls.ca/ipm/ipm25.asp>
- Byrne, David, and Eugenio Pinto, "The recent slowdown in high-tech equipment price declines and some implications for business investment and labor productivity," Federal Reserve Board, FEDS Notes, March 26, 2015. <http://www.federalreserve.gov/econresdata/notes/feds->

[notes/default.html](https://circle.ubc.ca/handle/2429/11401)

Chwelos, Paul, “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(3), 199–224, 2003.

Cole, Rosanne, Y.C. Chen, Joan A. Barquin-Stolleman, Ellen Dulberger, Nurhan Halvacian, and James H. Hodge, “Quality-Adjusted Price Indexes for Computer Processors and Selected Peripheral Equipment,” *Survey of Current Business*, 66(1), 41–50, 1986.

Cowen, Tyler, *The Great Stagnation: How America Ate All the Long-Hanging Fruit of Modern History, Got Sick, and Will (Eventually) Feel Better Again*, Dutton, 2011

Erickson, Tim and Ariel Pakes, “An Experimental Component Index for the CPI: From Annual Computer Data to Monthly Data on Other Goods,” *American Economic Review* 101 (August), 1707-1738, 2011.

Fernald, John G., “Productivity and Potential Output Before, During, and After the Great Recession,” in *NBER Macroeconomics Annual*, ed. by Jonathan Parker and Michael Woodford, 2014.

Fisher, Franklin M., John J. McGowan, and Joen E. Greenwood, *Folded, Spindled, and Mutilated: Economic Analysis and U.S. v. IBM*, MIT Press, 1983.

Flamm, Kenneth, “Measurement of DRAM Prices: Technology and Market Structure,” in Murray F. Foss, Marilyn E. Manser, and Allan H. Young, eds., *Price Measurements and Their Uses*, 157-97, University of Chicago Press and National Bureau of Economic Research, Studies in Income and Wealth, vol. 57, 1993. <http://www.nber.org/chapters/c7802.pdf>

_____, “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 mimeo, 2015.

Goettler, Ronald L., and Brett R. Gordon, “Does AMD Spur Intel to Innovate More?” *Journal of Political Economy*, 119(6), 1141-1200, 2011.

Gordon, Robert J., “The Postwar Evolution of Computer Prices,” NBER Working Paper 2227, April, 1987.

_____, "Is U.S. Economic Growth Over? Faltering Innovation Confronts the Six Headwinds," NBER Working Paper 18315, August, 2012.

_____, "The Demise of U.S. Economic Growth: Restatement, Rebuttal, and Reflections." NBER Working Paper 19895, February, 2014(a).

_____, "A New Method of Estimating Potential Real GDP Growth: Implications for the Labor Market and the Debt/GDP Ratio," NBER Working Paper 20423, August, 2014(b).

Grimm, Bruce T., "Price Indexes for Selected Semiconductors, 1974-96," *Survey of Current Business*, 78(2), 8–24, 1998.

Hall, Robert E. "Quantifying the Lasting Harm to the U.S. Economy from the Financial Crisis," in *NBER Macroeconomics Annual*, ed. by Jonathan Parker and Michael Woodford, 2014.

Hamilton Project, "The Future of Work in the Age of the Machine," Brookings Institution, February 19, 2015.

http://www.hamiltonproject.org/files/downloads_and_links/2015_02_24_THP_Future_of_Work_in_Machine_Age_transcript_unedited.pdf.

Hennessy, John L. and David A. Patterson, *Computer Architecture: A Quantitative Approach* (5th ed.), Elsevier Science, 2012.

Holdway, Mike, "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

Jorgenson, Dale W., "Information Technology and the U.S. Economy," *American Economic Review*, 91(1), 1–32, 2001.

Jorgenson, Dale W., Mun S. Ho, and Jon D. Samuels, "What Will Revive U.S. Economic Growth? Lessons from a Prototype Industry-Level Production Account for the United States," *Journal of Policy Modeling*, 36, 674-91, 2014.

Moore, Gordon E. "Cramming More Components onto Integrated Circuits." *Electronics*, 38(8), April 19, 1965.

Nosko, Christopher, "Competition and Quality Choice in the CPU Market," manuscript, Harvard University, 2010. http://faculty.chicagobooth.edu/chris.nosko/research/nosko_cpu.pdf

Pakes, Ariel, "A Reconsideration of Hedonic Price Indexes with an Application to PC's," *American Economic Review*, 93(5), 1578–96, 2003.

Pillai, Unni, "A Model of Technological Progress in the Microprocessor Industry," *Journal of Industrial Economics*, 61(4), 877–912, 2013.

Scherer, Frederic M., "Abuse of Dominance By High Technology Enterprises: A Comparison of U.S. and E.C. Approaches," *Journal of Industrial and Business Economics*, 38(1), 39-62, 2011.

Shenoy, Sunil R, and Akhilesh Daniel, "Intel Architecture and Silicon Cadence: The Catalyst for Industry Innovation," *Technology at Intel Magazine*, 2006.

Song, Minjae, "The Quality Adjusted Price Index in the Pure Characteristics Demand Model," *Journal of Business and Economic Statistics*, 28(1), 190-9, 2010.

State of New York v. Intel Corporation, No. 1:09-cv-00827 (D. Del), Nov. 4, 2009.

Sun, Sophie, "What are We Paying For? A Quality-adjusted Price Index for Laptop Microprocessors," Wellesley College senior thesis, 2014.

Thompson, Neil, "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, Jack 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, Cambridge, MA, 1989.

_____, "Handbook on Hedonic Indexes and Quality Adjustments: Special Application to Information Technology Products," Organisation for Economic Co-operation and Development, Paris and Washington, D.C., 2006.

<http://browse.oecdbookshop.org/oecd/pdfs/free/9306081e.pdf>

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, Oliner, and Sichel (2013) with update to 2014 based on data posted at <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

	2000-04	2005-08	2009-13
Number of Intel desktop MPU models			
With SPEC available	38	60	86
No SPEC available	47	45	97
Universe	85	105	183
Entry price (\$)			
With SPEC available	512	416	243
No SPEC available	181	197	168
Clock speed (ghz)			
With SPEC available	2.07	2.86	3.00
No SPEC available	1.95	2.80	2.76
Number of cores			
With SPEC available	1.0	2.25	3.09
No SPEC available	1.0	1.44	2.76
Maximum thermal design power (watts)			
With SPEC available	60.3	92.3	75.1
No SPEC available	55.2	78.8	59.1
Lithography size (nanometers)			
With SPEC available	147.9	65.2	31.3
No SPEC available	134.7	73.9	30.9

Panel B: Price changes (percent, average annual rate)

	2000-04	2004-08	2008-13	2000-13
PPI-like matched-model price index (Intel MPUs)				
With SPEC available	-56	-19	-7	-29
No SPEC available	-48	-24	-6	-26
All desktops	-53	-22	-6	-28
PPI	-48	-29	-8	-28

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

Table 3. Regression Results for 2000-06**Panel A: Introduction period only**

	2000-01	2001-02	2002-03	2003-04	2004-05	2005-06
<i>Year dummy</i>	-.895** (.177)	-.363 (.209)	-.995** (.287)	-.929* (.281)	-.263 (.277)	-.819* (.311)
<i>ln Performance</i>	.63* (.27)	1.12** (.29)	2.98** (.69)	3.99** (.80)	4.39** (1.01)	2.95** (.74)
Number of Obs.	20	18	14	11	23	28
Adjusted R ²	.56	.48	.56	.70	.44	.34

Panel B: Full sample

	2000-01	2001-02	2002-03	2003-04	2004-05	2005-06
<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 ²	.46	.44	.47	.62	.43	.35

Note: The dependent variable is ln(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 2006-13**Panel A: Introduction period only**

	2006-07	2007-08	2008-09	2009-10	2010-11	2011-12	2012-13
<i>Year dummy</i>	-.726 (.388)	-.784* (.226)	-.754 (.369)	-.335 (.244)	-1.229** (.226)	-.481** (.110)	-.533** (.091)
<i>ln Performance</i>	1.87* (.76)	2.39** (.82)	2.86** (.70)	2.79** (.65)	3.31** (.50)	2.88** (.31)	2.88** (.20)
Observations	24	26	27	25	35	50	31
Adjusted R ²	.16	.28	.41	.41	.58	.65	.87

Panel B: Full sample

	2006-07	2007-08	2008-09	2009-10	2010-11	2011-12	2012-13
<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 ²	.41	.41	.46	.57	.35	.29	.63

Note: The dependent variable is ln(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 prices¹
(Average annual percent change over periods shown)

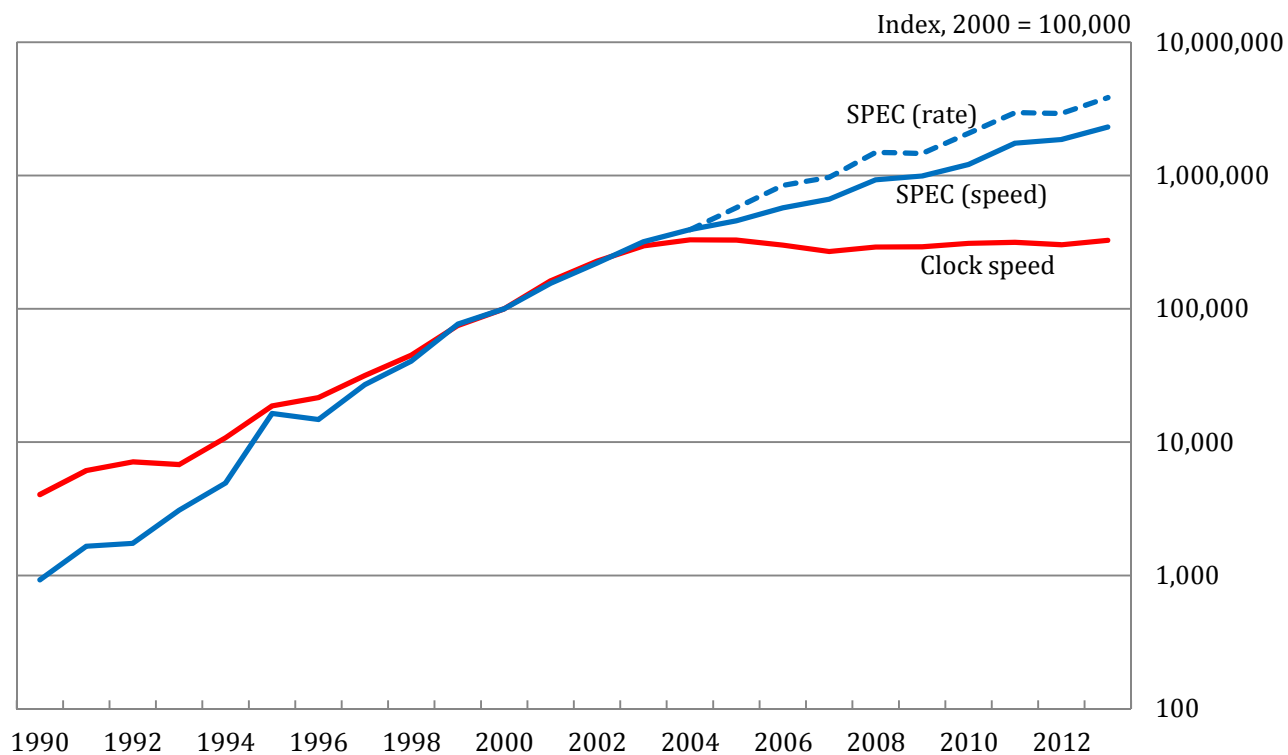
	2000-04	2004-08	2008-13	2000-13
Hedonic, SPEC performance				
Introduction period	-52	-37	-43	-44
Full sample	-40	-33	-22	-31
Hedonic, chip characteristics				
Introduction period	-43	-1	-21	-23
Full sample	-39	-16	-13	-23
Federal Reserve Board ²	-51	-29	-31	-38
PPI	-48	-29	-8	-28

1. 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.

2. Constructed as a matched-model index through 2006 and as an introduction-period hedonic index after 2006. The dataset and hedonic regression specification reflect an earlier version of the empirical work presented in this paper.

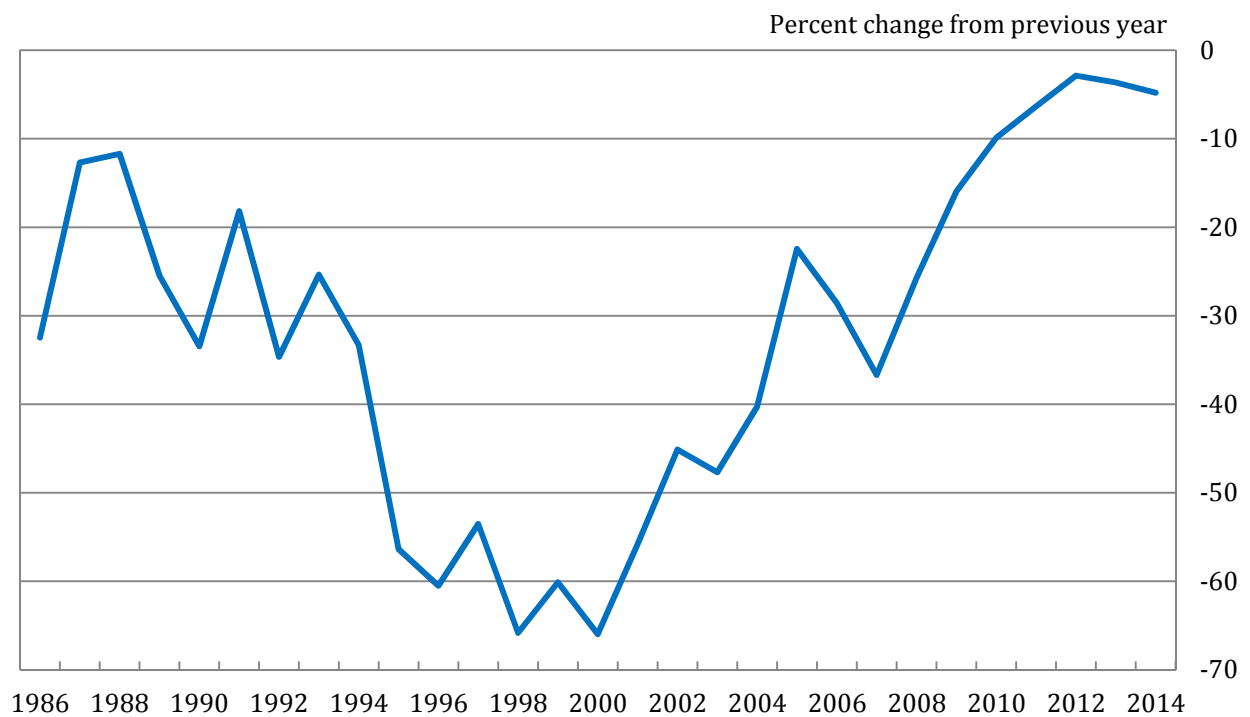
Source. Authors' calculations and Bureau of Labor Statistics.

Figure 1: Desktop MPU Performance Measures



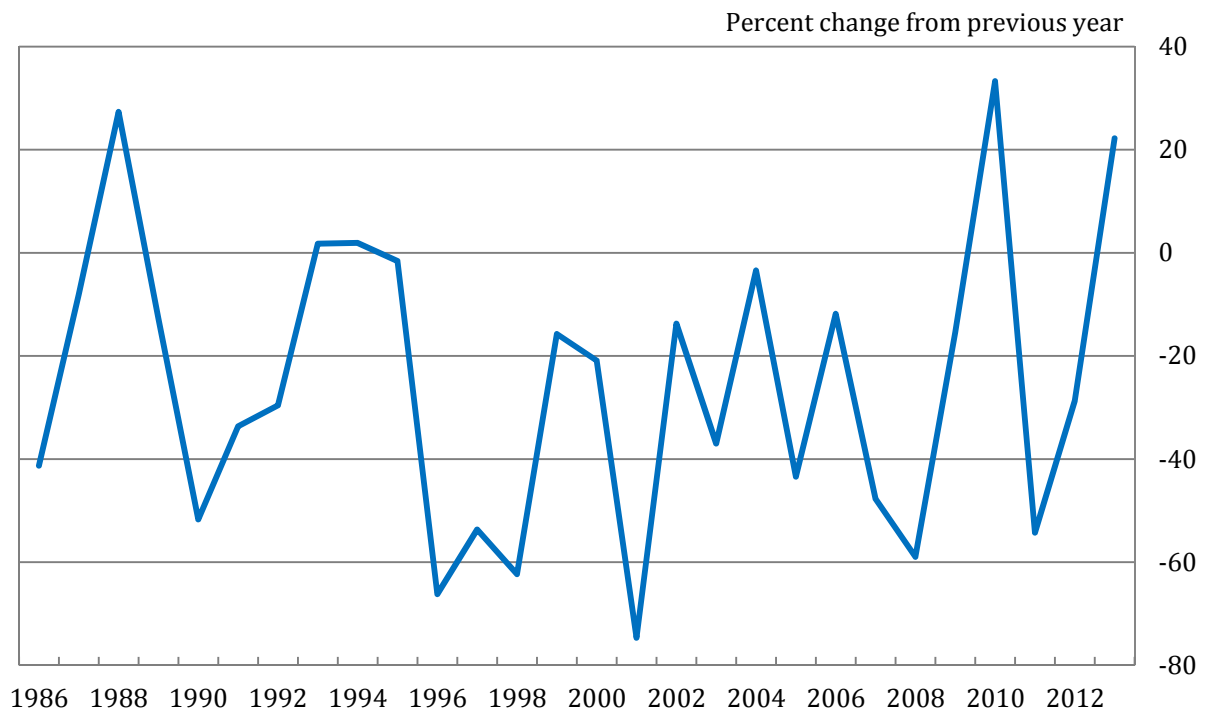
Source. Authors' calculations using data provided by Unni Pillai and performance information from Intel price lists and SPEC corporation.

Figure 2: Reported MPU Prices



Sources. Grimm (1998, table 12) for 1986-92, Federal Reserve Board for 1993-97, and Bureau of Labor Statistics for 1998-2014.

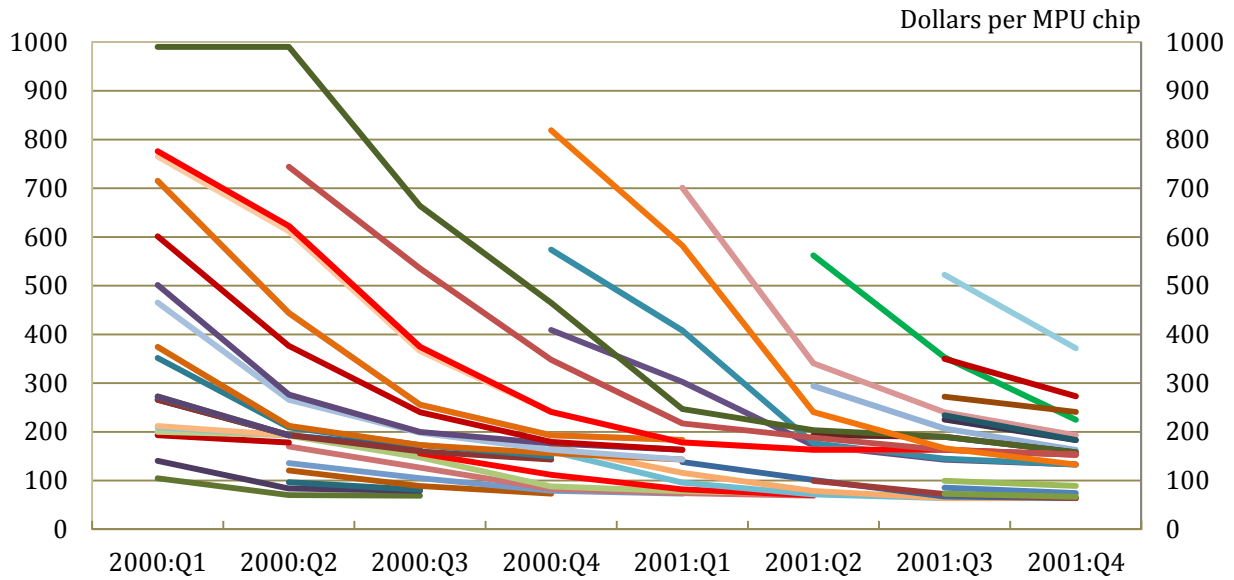
Figure 3: DRAM Prices



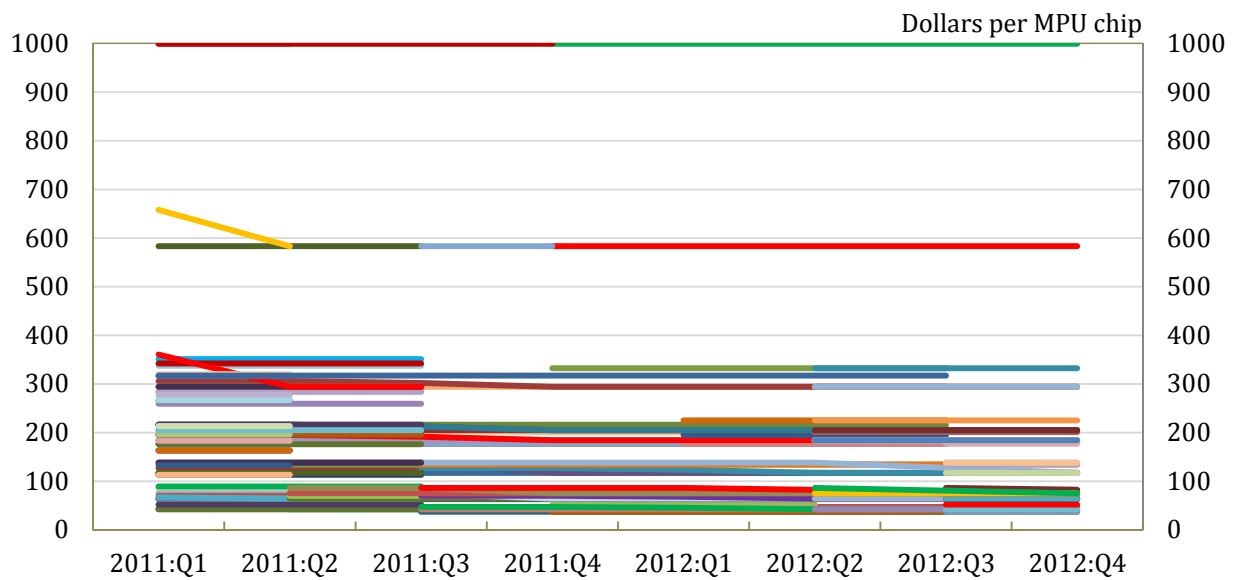
Sources. Grimm (1988, table 5) for 1986-92 and Federal Reserve Board for 1993-2013.

Figure 4: Intel List Prices

Panel A: 2000-01

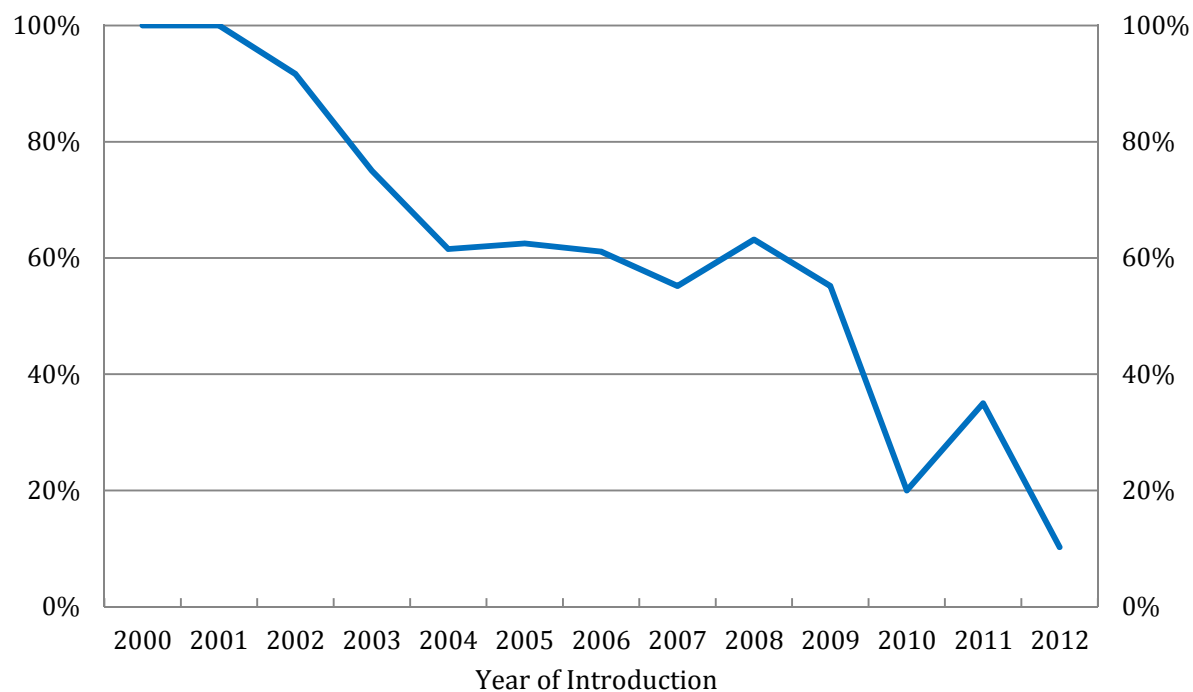


Panel B: 2011-12



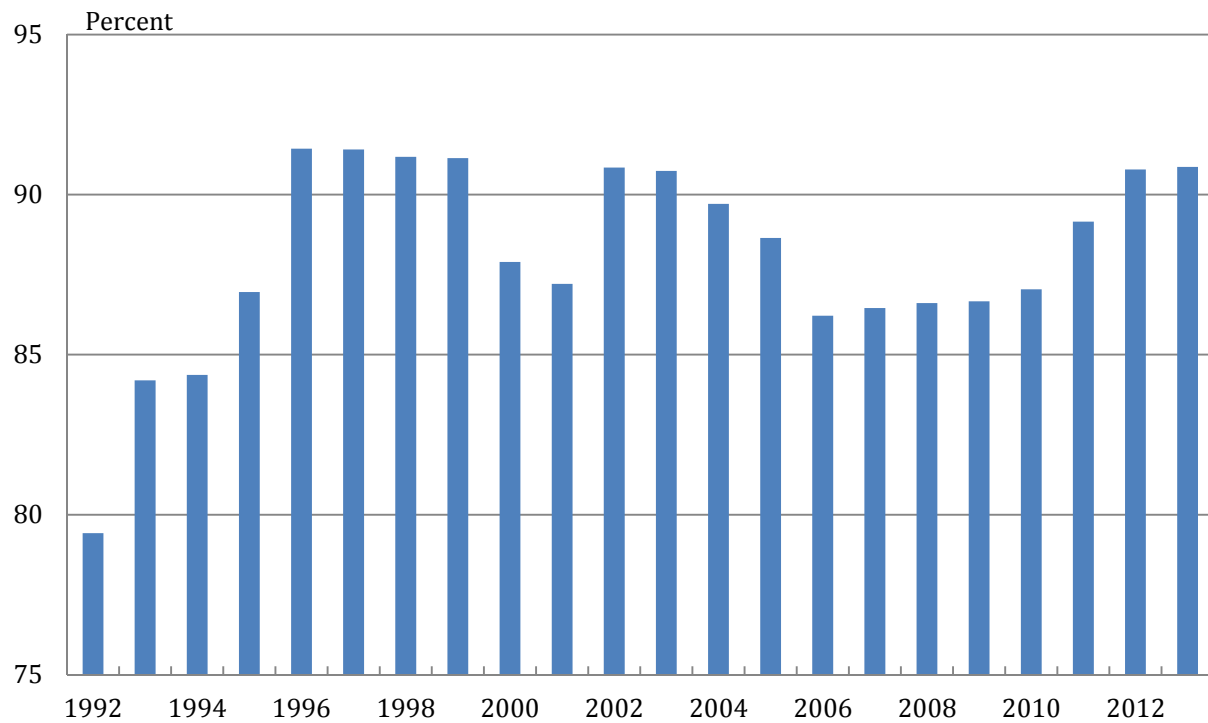
Note. Prices for 1,000 unit "trays" of MPUs.
Source. Authors' calculations from Intel price lists.

Figure 5: Share of Intel Desktop MPUs with List Price Decline within Four Quarters of Introduction



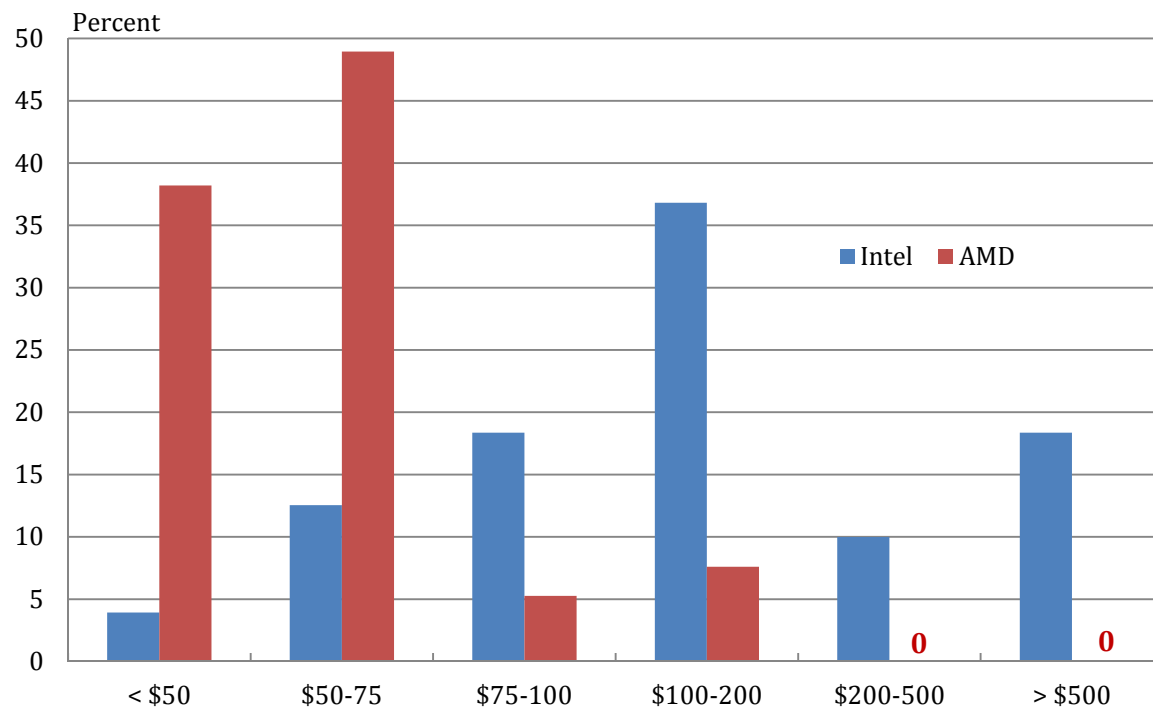
Source. Authors' calculations from Intel price lists.

Figure 6: Intel Share of Combined Intel and AMD Revenue



Source. Authors' calculations based on data from financial reports filed with the Securities and Exchange Commission.

Figure 7: 2013 Distribution of Desktop MPU Revenue, by Price of Chip

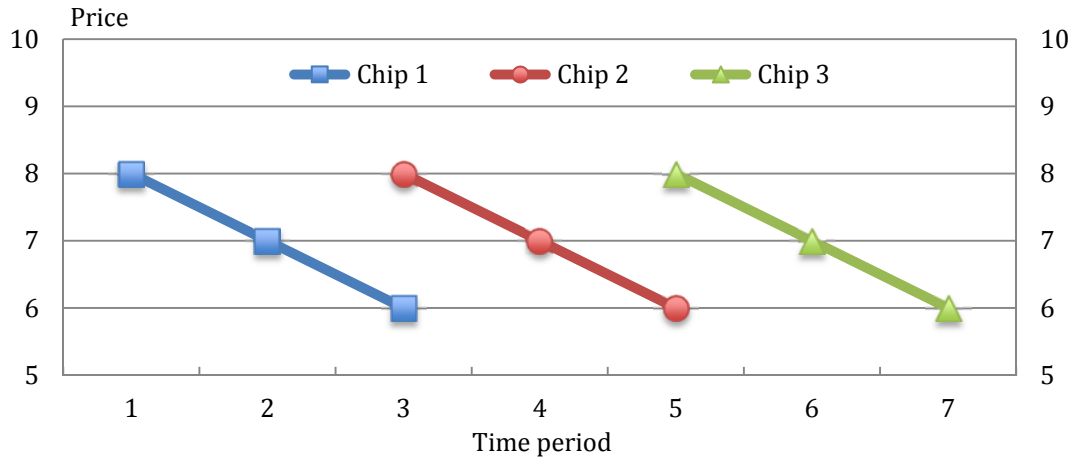


Note. Data cover 2013:Q1 through 2013:Q3.

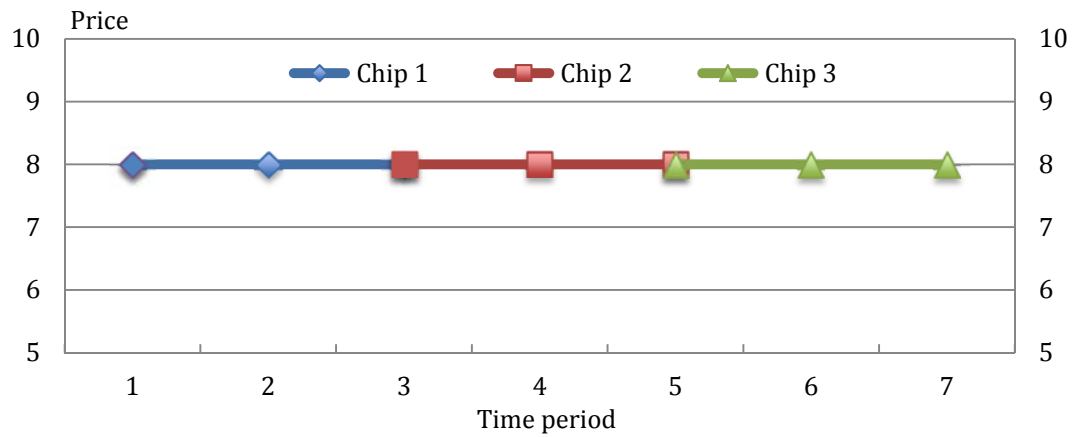
Source. Authors' calculations based on data from IDC Research, Inc.

Figure 8: Alternative Scenarios

Panel A: Pricing Pattern Before the Mid-2000s



Panel B: Pricing Pattern After the Mid-2000s



Panel C: After the Mid-2000s: Age-related Discounts

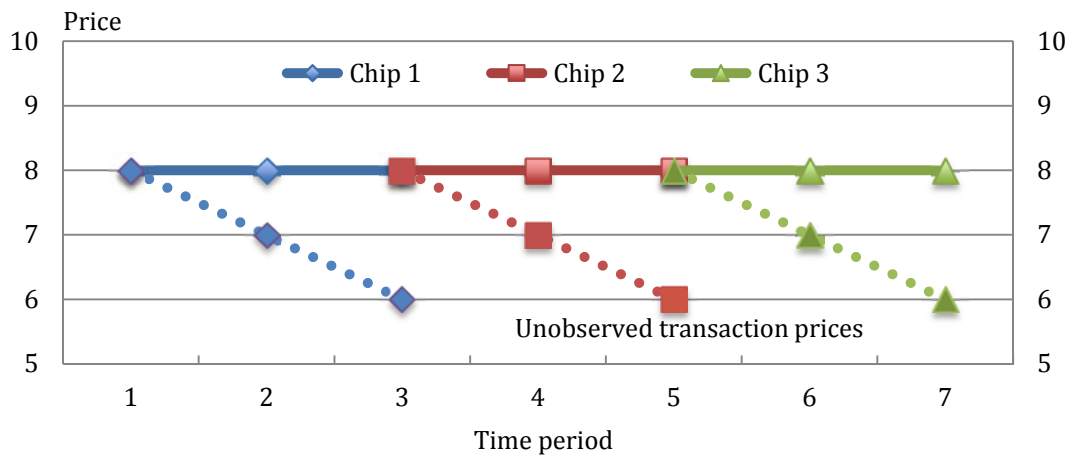


Figure 8: Alternative Scenarios

Panel D: After the Mid-2000s: Drop-off in Volume

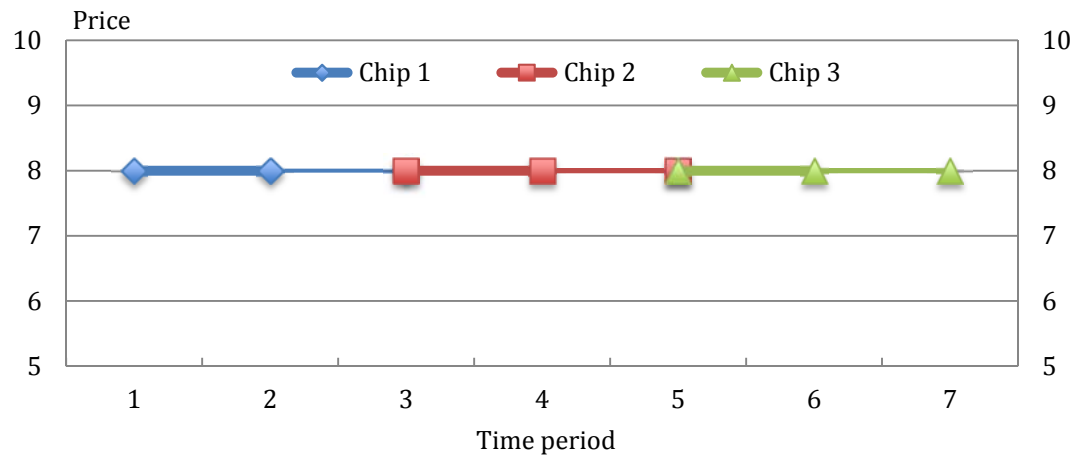
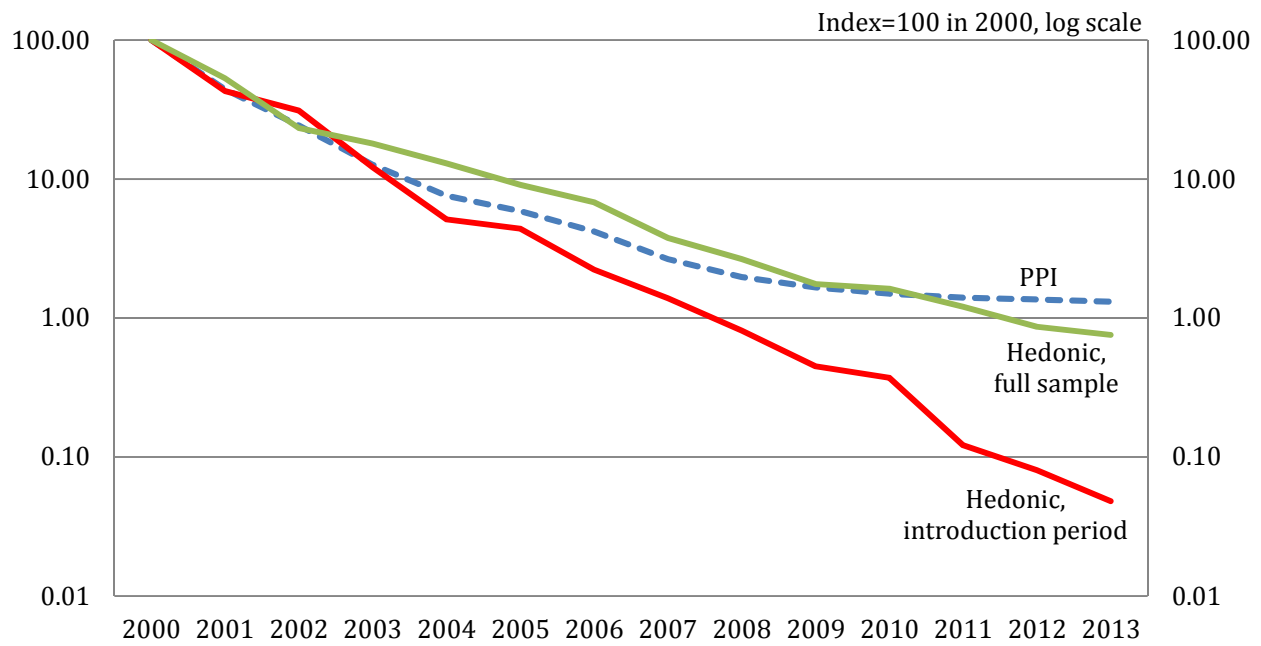
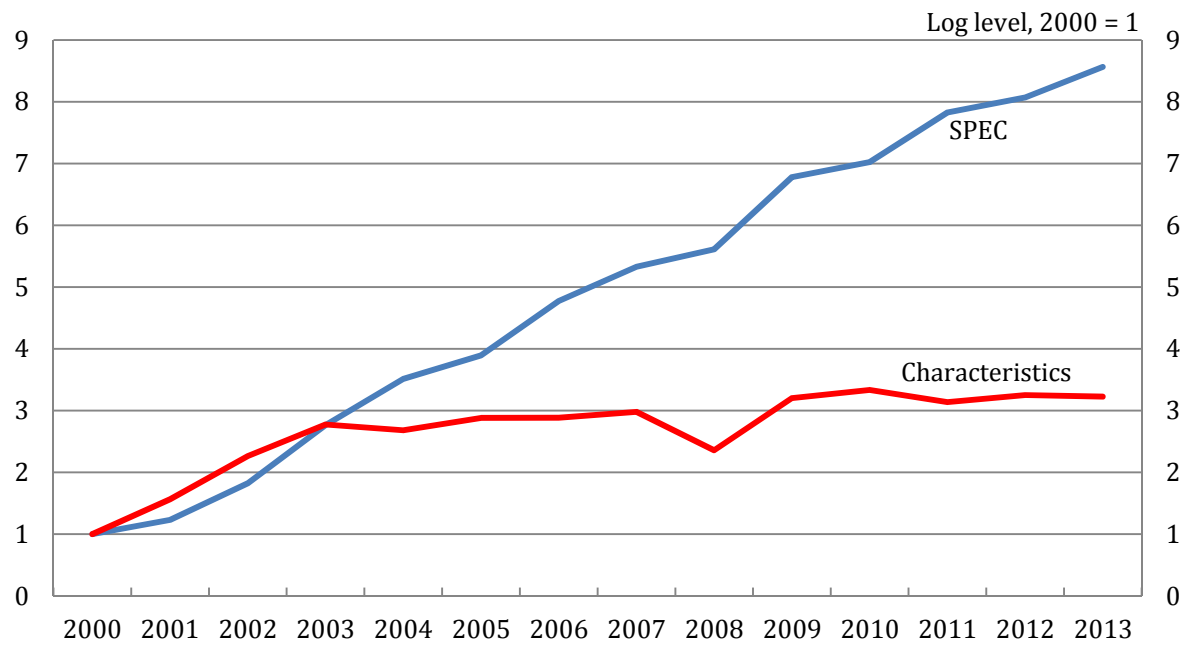


Figure 9: MPU Price Levels



Note. Hedonic indexes are based on regressions that use SPEC performance to control for quality.
Source. Bureau of Labor Statistics and authors' calculations.

Figure 10: Effect of Chip Quality on Price



Source. Authors' calculations. See text for details.

APPENDIX A: PERFORMANCE MEASURES FROM SPEC (SYSTEM PERFORMANCE EVALUATION CORPORATION)

As noted in the text, SPEC is a non-profit corporation that develops performance benchmarks for computers. This appendix describes the performance benchmarks from SPEC that we use in our analysis.

For MPUs, SPEC has developed a suite of benchmark tests that evaluate how quickly an MPU can complete a set of tasks that are developed from real user applications. These benchmark suites are updated periodically to reflect changes in MPU architecture and in relevant tasks. We use the latest benchmark, called CPU2006, as well as CPU2000.

The performance of an MPU will depend on characteristics of the system other than just the MPU; these other elements include memory and the compiler used. SPEC has benchmarks for a standard configuration (“base” metrics) and a configuration in which compilers are tuned for maximum performance (“peak” metrics). We use the base metrics. That said, SPEC notes that the 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.

Further details are provided below on how we use the CPU2006 and CPU2000 benchmarks to construct the performance measure used in our analysis.

CPU2006

CPU2006 was introduced in 2006, and we use results from CPU2006 to measure the performance of chips from 2006 to 2013. This benchmark consists of two suites of tasks. The first suite contains 12 tests that focus on integer calculations, and the second suite contains 17 tests that focus on floating point calculations. For each test, SPEC normalizes the test time by taking the ratio of the test time on a standardized reference machine to the test time for the MPU being tested. By scaling results in this way, shorter test times result in higher performance scores. SPEC then constructs a composite score for calculations by taking the geometric mean of the normalized individual integer scores; SPEC constructs a composite floating-point score in a parallel manner. The composite scores we use are called SPECint_base2006 and SPECfp_base2006.

As noted, these tests cover actual user applications. The integer and floating-point applications include the following.

- ***Integer*** applications include running PERL scripts, file compression, running a C compiler, combinatorial optimization, artificial intelligence (playing the games Go and chess), searching gene sequences, simulating a quantum computer, video compression, discrete event simulation, running path-finding algorithms, and XML processing.
- ***Floating-point*** applications include computations for fluid dynamics, quantum chemistry, quantum chromodynamics, molecular dynamics, general relativity, finite

element analysis, linear programming, image rendering, structural mechanics, computational electromagnetics, weather modeling, and speech recognition.

To construct the performance measure used for most of our analysis, we take the geometric mean of the integer and floating-point composites described above.

With the advent of parallel processing, SPEC began distinguishing between *speed* and *rate* measures of performance. Speed measures focus on how fast a computer completes a single task. Rate measures focus on how many tasks a computer can complete in a given amount of time, taking advantage of available parallel processing.⁴⁴ The integer and floating-point performance suites described above are speed measures. Our analysis primarily focuses on these speed measures as they are available over a longer time span. That said, we also consider rate measures to more fully account for parallel processing and the rise of multicore chips.

CPU2000

CPU2000 was introduced in 1999 and retired in 2007. We use results from CPU2000 to measure the performance of chips from 2000 to 2006. The calculation of CPU2000 is very similar to that of CPU2006 except that the individual performance tests are of a type appropriate to the computing environment in the earlier period. The CPU2000 integer suite includes 12 tests, and the floating-point suite includes 14 tests.

Just as with CPU2006, for CPU2000, our analysis relies on the geometric mean of the integer and floating-point composite metrics, which themselves are geometric means of normalized results of the individual integer and floating-point tests.

Additional Details

We are able to bridge across the SPEC2000 and SPEC2006 benchmarks because results for 2006 were reported on both benchmarks for many chips. This allows us to estimate hedonic regressions for adjacent pairs of years through 2005-06 using SPEC2000 and for adjacent pairs of years starting with 2006-07 using SPEC2006.

An MPU chip often has multiple scores for each of the SPEC benchmark tests. Multiple scores can arise either because more than one computer vendor tested the chip or because a given vendor tested the chip under different conditions. The variation in test conditions can reflect differences in hardware (e.g., the circuit board or amount and type of DRAM) or software (e.g., the operating system or compiler). When multiple scores are available for a specific model, we use the model's median score.

⁴⁴ Speed scores do account for some parallel processing; in particular, speed scores allow for auto-parallelization for a single task, and by 2014 almost all speed scores reported using this feature.

APPENDIX B: INTEL DESKTOP MPUS ON WHOLESALE PRICE LISTS

The table below (see next page) lists every Intel desktop MPU shown on the company's wholesale price lists starting in 1999. The most recent price list is posted at <http://www.intc.com/pricelist.cfm>; price lists for earlier periods were collected from this location and other online sites. Entry year denotes the year that the MPU first appeared on a price list. The MPU description is taken directly from Intel's ARK database, located at <http://ark.intel.com>, which contains a full history of Intel microprocessors; the note at the bottom of the table describes the minor ways we edited the ARK description to save space. The final column shows whether a SPEC score exists for each model. All of the models with a SPEC score are included in our empirical analysis. Of 373 models introduced by Intel during the period, 184 have SPEC scores available.

Entry Year	MPU Description ¹	SPEC Score?
1999	Celeron® 333 MHz, 128K Cache, 66 MHz FSB	No
1999	Celeron® 366 MHz, 128K Cache, 66 MHz FSB	No
1999	Celeron® 400 MHz, 128K Cache, 66 MHz FSB	No
1999	Celeron® 433 MHz, 128K Cache, 66 MHz FSB	No
1999	Celeron® 466 MHz, 128K Cache, 66 MHz FSB	No
1999	Celeron® 500 MHz, 128K Cache, 66 MHz FSB	No
1999	Pentium® II 350 MHz, 512K Cache, 100 MHz FSB	No
1999	Pentium® II 400 MHz, 512K Cache, 100 MHz FSB	No
1999	Pentium® II 450 MHz, 512K Cache, 100 MHz FSB	No
1999	Pentium® III 450 MHz, 512K Cache, 100 MHz FSB	Yes
1999	Pentium® III 500 MHz, 256K Cache, 100 MHz FSB	No
1999	Pentium® III 500 MHz, 512K Cache, 100 MHz FSB	No
1999	Pentium® III 533 MHz, 256K Cache, 133 MHz FSB	No
1999	Pentium® III 533 MHz, 512K Cache, 133 MHz FSB	No
1999	Pentium® III 550 MHz, 256K Cache, 100 MHz FSB	No
1999	Pentium® III 550 MHz, 512K Cache, 100 MHz FSB	No
1999	Pentium® III 600 MHz, 256K Cache, 100 MHz FSB	No
1999	Pentium® III 600 MHz, 256K Cache, 133 MHz FSB	No
1999	Pentium® III 600 MHz, 512K Cache, 100 MHz FSB	No
1999	Pentium® III 600 MHz, 512K Cache, 133 MHz FSB	No
1999	Pentium® III 650 MHz, 256K Cache, 100 MHz FSB	Yes
1999	Pentium® III 667 MHz, 256K Cache, 133 MHz FSB	Yes
1999	Pentium® III 700 MHz, 256K Cache, 100 MHz FSB	Yes
1999	Pentium® III 733 MHz, 256K Cache, 133 MHz FSB	Yes
2000	Celeron® 533 MHz, 128K Cache, 66 MHz FSB	No
2000	Celeron® 566 MHz, 128K Cache, 66 MHz FSB	No
2000	Celeron® 600 MHz, 128K Cache, 66 MHz FSB	No
2000	Celeron® 633 MHz, 128K Cache, 66 MHz FSB	No
2000	Celeron® 667 MHz, 128K Cache, 66 MHz FSB	No
2000	Celeron® 700 MHz, 128K Cache, 66 MHz FSB	No
2000	Celeron® 733 MHz, 128K Cache, 66 MHz FSB	No
2000	Celeron® 766 MHz, 128K Cache, 66 MHz FSB	No
2000	Celeron® 800 MHz, 128K Cache, 100 MHz FSB	No
2000	Pentium® 4 1.30 GHz, 256K Cache, 400 MHz FSB	Yes
2000	Pentium® 4 1.40 GHz, 256K Cache, 400 MHz FSB	Yes
2000	Pentium® 4 1.50 GHz, 256K Cache, 400 MHz FSB	Yes
2000	Pentium® III 1.00 GHz, 256K Cache, 133 MHz FSB	Yes
2000	Pentium® III 750 MHz, 256K Cache, 100 MHz FSB	Yes
2000	Pentium® III 800 MHz, 256K Cache, 100 MHz FSB	No
2000	Pentium® III 800 MHz, 256K Cache, 133 MHz FSB	Yes
2000	Pentium® III 850 MHz, 256K Cache, 100 MHz FSB	Yes
2000	Pentium® III 866 MHz, 256K Cache, 133 MHz FSB	Yes
2000	Pentium® III 933 MHz, 256K Cache, 133 MHz FSB	Yes
2001	Celeron® 1.00 GHz, 128K Cache, 100 MHz FSB	No
2001	Celeron® 1.00 GHz, 256K Cache, 100 MHz FSB	No
2001	Celeron® 1.10 GHz, 128K Cache, 100 MHz FSB	No
2001	Celeron® 1.10 GHz, 256K Cache, 100 MHz FSB	No
2001	Celeron® 1.20 GHz, 256K Cache, 100 MHz FSB	No
2001	Celeron® 1.30 GHz, 256K Cache, 100 MHz FSB	No
2001	Celeron® 850 MHz, 128K Cache, 100 MHz FSB	No
2001	Celeron® 900 MHz, 128K Cache, 100 MHz FSB	No
2001	Celeron® 950 MHz, 128K Cache, 100 MHz FSB	No
2001	Pentium® 4 1.60 GHz, 256K Cache, 400 MHz FSB	Yes
2001	Pentium® 4 1.70 GHz, 256K Cache, 400 MHz FSB	Yes
2001	Pentium® 4 1.80 GHz, 256K Cache, 400 MHz FSB	Yes
2001	Pentium® 4 1.90 GHz, 256K Cache, 400 MHz FSB	Yes
2001	Pentium® 4 2.00 GHz, 256K Cache, 400 MHz FSB	Yes
2001	Pentium® 4 2.00 GHz, 512K Cache, 400 MHz FSB	Yes

Entry Year	MPU Description ¹	SPEC Score?
2001	Pentium® 4 2.20 GHz, 512K Cache, 400 MHz FSB	Yes
2001	Pentium® III 1.00 GHz, 256K Cache, 100 MHz FSB	Yes
2001	Pentium® III 1.10 GHz, 256K Cache, 100 MHz FSB	Yes
2001	Pentium® III 1.13 GHz, 256K Cache, 133 MHz FSB	Yes
2001	Pentium® III 1.20 GHz, 256K Cache, 133 MHz FSB	Yes
2002	Celeron® 1.40 GHz, 256K Cache, 100 MHz FSB	No
2002	Celeron® 2.00 GHz, 128K Cache, 400 MHz FSB	No
2002	Celeron® 2.10 GHz, 128K Cache, 400 MHz FSB	No
2002	Pentium® 4 2.26 GHz, 512K Cache, 533 MHz FSB	Yes
2002	Pentium® 4 2.40 GHz, 512K Cache, 400 MHz FSB	Yes
2002	Pentium® 4 2.40 GHz, 512K Cache, 533 MHz FSB	Yes
2002	Pentium® 4 2.50 GHz, 512K Cache, 400 MHz FSB	No
2002	Pentium® 4 2.53 GHz, 512K Cache, 533 MHz FSB	Yes
2002	Pentium® 4 2.60 GHz, 512K Cache, 400 MHz FSB	No
2002	Pentium® 4 2.66 GHz, 512K Cache, 533 MHz FSB	Yes
2002	Pentium® 4 2.80 GHz, 512K Cache, 533 MHz FSB	Yes
2002	Pentium® 4 supporting HT Tech. 3.06 GHz, 512K Cache, 533 MHz FSB	Yes
2003	Celeron® 2.20 GHz, 128K Cache, 400 MHz FSB	No
2003	Celeron® 2.30 GHz, 128K Cache, 400 MHz FSB	No
2003	Celeron® 2.40 GHz, 128K Cache, 400 MHz FSB	No
2003	Celeron® 2.50 GHz, 128K Cache, 400 MHz FSB	No
2003	Celeron® 2.60 GHz, 128K Cache, 400 MHz FSB	No
2003	Celeron® 2.70 GHz, 128K Cache, 400 MHz FSB	No
2003	Celeron® 2.80 GHz, 128K Cache, 400 MHz FSB	No
2003	Pentium® 4 2.80A GHz, 1M Cache, 533 MHz FSB	No
2003	Pentium® 4 supporting HT Tech. 2.40 GHz, 512K Cache, 800 MHz FSB	Yes
2003	Pentium® 4 supporting HT Tech. 2.60 GHz, 512K Cache, 800 MHz FSB	Yes
2003	Pentium® 4 supporting HT Tech. 2.80 GHz, 512K Cache, 800 MHz FSB	Yes
2003	Pentium® 4 supporting HT Tech. 2.80E GHz, 1M Cache, 800 MHz FSB	No
2003	Pentium® 4 supporting HT Tech. 3.00 GHz, 1M Cache, 800 MHz FSB	No
2003	Pentium® 4 supporting HT Tech. 3.00 GHz, 512K Cache, 800 MHz FSB	No
2003	Pentium® 4 supporting HT Tech. 3.20 GHz, 1M Cache, 800 MHz FSB	No
2003	Pentium® 4 supporting HT Tech. 3.20 GHz, 512K Cache, 800 MHz FSB	Yes
2003	Pentium® 4 supporting HT Tech. 3.40 GHz, 1M Cache, 800 MHz FSB	No
2003	Pentium® 4 supporting HT Tech. 3.40 GHz, 512K Cache, 800 MHz FSB	Yes
2003	Pentium® 4 Extreme Ed. supporting HT Tech. 3.20 GHz, 2M Cache, 800 MHz FSB	Yes
2003	Pentium® 4 Extreme Ed. supporting HT Tech. 3.40 GHz, 2M Cache, 800 MHz FSB	Yes
2004	Celeron® D 320 (256K Cache, 2.40 GHz, 533 MHz FSB)	No
2004	Celeron® D 325 (256K Cache, 2.53 GHz, 533 MHz FSB)	No
2004	Celeron® D 330 (256K Cache, 2.66 GHz, 533 MHz FSB)	No
2004	Celeron® D 335 (256K Cache, 2.80 GHz, 533 MHz FSB)	No
2004	Celeron® 2.40 GHz, 256K Cache, 533 MHz FSB	No
2004	Pentium® 4 2.80 GHz, 1M Cache, 533 MHz FSB	Yes
2004	Pentium® 4 530 supporting HT Tech. (1M Cache, 3.00 GHz, 800 MHz FSB)	Yes
2004	Pentium® 4 540 supporting HT Tech. (1M Cache, 3.20 GHz, 800 MHz FSB)	No
2004	Pentium® 4 540J supporting HT Tech. (1M Cache, 3.20 GHz, 800 MHz FSB)	No
2004	Pentium® 4 550 supporting HT Tech. (1M Cache, 3.40 GHz, 800 MHz FSB)	No
2004	Pentium® 4 550J supporting HT Tech. (1M Cache, 3.40 GHz, 800 MHz FSB)	No
2004	Pentium® 4 560 supporting HT Tech. (1M Cache, 3.60 GHz, 800 MHz FSB)	No
2004	Pentium® 4 560J supporting HT Tech. (1M Cache, 3.60 GHz, 800 MHz FSB)	Yes
2004	Pentium® 4 Extreme Ed. supporting HT Tech. 3.46 GHz, 2M Cache, 1066 MHz FSB	Yes
2005	Celeron® D 326 (256K Cache, 2.53 GHz, 533 MHz FSB)	No
2005	Celeron® D 331 (256K Cache, 2.66 GHz, 533 MHz FSB)	No
2005	Celeron® D 336 (256K Cache, 2.80 GHz, 533 MHz FSB)	No
2005	Celeron® D 340 (256K Cache, 2.93 GHz, 533 MHz FSB)	No
2005	Celeron® D 341 (256K Cache, 2.93 GHz, 533 MHz FSB)	No
2005	Celeron® D 345 (256K Cache, 3.06 GHz, 533 MHz FSB)	No
2005	Celeron® D 346 (256K Cache, 3.06 GHz, 533 MHz FSB)	No

Entry Year	MPU Description ¹	SPEC Score?
2005	Celeron® D 350 (256K Cache, 3.20 GHz, 533 MHz FSB)	No
2005	Celeron® D 351 (256K Cache, 3.20 GHz, 533 MHz FSB)	No
2005	Celeron® D 355 (256K Cache, 3.33 GHz, 533 MHz FSB)	No
2005	Pentium® 4 520J supporting HT Tech. (1M Cache, 2.80 GHz, 800 MHz FSB)	No
2005	Pentium® 4 521 supporting HT Tech. (1M Cache, 2.80 GHz, 800 MHz FSB)	Yes
2005	Pentium® 4 530J supporting HT Tech. (1M Cache, 3.00 GHz, 800 MHz FSB)	No
2005	Pentium® 4 531 supporting HT Tech. (1M Cache, 3.00 GHz, 800 MHz FSB)	Yes
2005	Pentium® 4 541 supporting HT Tech. (1M Cache, 3.20 GHz, 800 MHz FSB)	No
2005	Pentium® 4 551 supporting HT Tech. (1M Cache, 3.40 GHz, 800 MHz FSB)	No
2005	Pentium® 4 561 supporting HT Tech. (1M Cache, 3.60 GHz, 800 MHz FSB)	No
2005	Pentium® 4 570J supporting HT Tech. (1M Cache, 3.80 GHz, 800 MHz FSB)	Yes
2005	Pentium® 4 571 supporting HT Tech. (1M Cache, 3.80 GHz, 800 MHz FSB)	No
2005	Pentium® 4 630 supporting HT Tech. (2M Cache, 3.00 GHz, 800 MHz FSB)	Yes
2005	Pentium® 4 631 supporting HT Tech. (2M Cache, 3.00 GHz, 800 MHz FSB)	Yes
2005	Pentium® 4 640 supporting HT Tech. (2M Cache, 3.20 GHz, 800 MHz FSB)	Yes
2005	Pentium® 4 641 supporting HT Tech. (2M Cache, 3.20 GHz, 800 MHz FSB)	No
2005	Pentium® 4 650 supporting HT Tech. (2M Cache, 3.40 GHz, 800 MHz FSB)	Yes
2005	Pentium® 4 651 supporting HT Tech. (2M Cache, 3.40 GHz, 800 MHz FSB)	Yes
2005	Pentium® 4 660 supporting HT Tech. (2M Cache, 3.60 GHz, 800 MHz FSB)	Yes
2005	Pentium® 4 661 supporting HT Tech. (2M Cache, 3.60 GHz, 800 MHz FSB)	No
2005	Pentium® 4 662 supporting HT Tech. (2M Cache, 3.60 GHz, 800 MHz FSB)	No
2005	Pentium® 4 670 supporting HT Tech. (2M Cache, 3.80 GHz, 800 MHz FSB)	Yes
2005	Pentium® 4 672 supporting HT Tech. (2M Cache, 3.80 GHz, 800 MHz FSB)	No
2005	Pentium® D 820 (2M Cache, 2.80 GHz, 800 MHz FSB)	Yes
2005	Pentium® D 830 (2M Cache, 3.00 GHz, 800 MHz FSB)	Yes
2005	Pentium® D 840 (2M Cache, 3.20 GHz, 800 MHz FSB)	Yes
2005	Pentium® D 920 (4M Cache, 2.80 GHz, 800 MHz FSB)	Yes
2005	Pentium® D 930 (4M Cache, 3.00 GHz, 800 MHz FSB)	Yes
2005	Pentium® D 940 (4M Cache, 3.20 GHz, 800 MHz FSB)	Yes
2005	Pentium® D 950 (4M Cache, 3.40 GHz, 800 MHz FSB)	Yes
2005	Pentium® Extreme Ed. 840 (2M Cache, 3.20 GHz, 800 MHz FSB)	Yes
2005	Pentium® Extreme Ed. 955 (4M Cache, 3.46 GHz, 1066 MHz FSB)	Yes
2005	Pentium® 4 Extreme Ed. supporting HT Tech. 3.73 GHz, 2M Cache, 1066 MHz FSB	Yes
2006	Celeron® D 315 (256K Cache, 2.26 GHz, 533 MHz FSB)	No
2006	Celeron® D 347 (512K Cache, 3.06 GHz, 533 MHz FSB)	No
2006	Celeron® D 352 (512K Cache, 3.20 GHz, 533 MHz FSB)	No
2006	Celeron® D 356 (512K Cache, 3.33 GHz, 533 MHz FSB)	No
2006	Core™2 Duo E6300 (2M Cache, 1.86 GHz, 1066 MHz FSB)	Yes
2006	Core™2 Duo E6400 (2M Cache, 2.13 GHz, 1066 MHz FSB)	Yes
2006	Core™2 Duo E6600 (4M Cache, 2.40 GHz, 1066 MHz FSB)	Yes
2006	Core™2 Duo E6700 (4M Cache, 2.66 GHz, 1066 MHz FSB)	Yes
2006	Core™2 Extreme QX6700 (8M Cache, 2.66 GHz, 1066 MHz FSB)	Yes
2006	Core™2 Extreme X6800 (4M Cache, 2.93 GHz, 1066 MHz FSB)	Yes
2006	Pentium® 4 524 supporting HT Tech. (1M Cache, 3.06 GHz, 533 MHz FSB)	No
2006	Pentium® D 805 (2M Cache, 2.66 GHz, 533 MHz FSB)	Yes
2006	Pentium® D 915 (4M Cache, 2.80 GHz, 800 MHz FSB)	No
2006	Pentium® D 925 (4M Cache, 3.00 GHz, 800 MHz FSB)	Yes
2006	Pentium® D 945 (4M Cache, 3.40 GHz, 800 MHz FSB)	Yes
2006	Pentium® D 960 (4M Cache, 3.60 GHz, 800 MHz FSB)	Yes
2006	Pentium® Extreme Ed. 965 (4M Cache, 3.73 GHz, 1066 MHz FSB)	Yes
2007	Celeron® D 365 (512K Cache, 3.60 GHz, 533 MHz FSB)	No
2007	Celeron® 430 (512K Cache, 1.80 GHz, 800 MHz FSB)	No
2007	Celeron® 440 (512K Cache, 2.00 GHz, 800 MHz FSB)	No
2007	Core™2 Duo E4300 (2M Cache, 1.80 GHz, 800 MHz FSB)	Yes
2007	Core™2 Duo E4400 (2M Cache, 2.00 GHz, 800 MHz FSB)	No
2007	Core™2 Duo E4500 (2M Cache, 2.20 GHz, 800 MHz FSB)	Yes
2007	Core™2 Duo E4600 (2M Cache, 2.40 GHz, 800 MHz FSB)	Yes
2007	Core™2 Duo E6320 (4M Cache, 1.86 GHz, 1066 MHz FSB)	No

Entry Year	MPU Description ¹	SPEC Score?
2007	Core™2 Duo E6420 (4M Cache, 2.13 GHz, 1066 MHz FSB)	No
2007	Core™2 Duo E6550 (4M Cache, 2.33 GHz, 1333 MHz FSB)	No
2007	Core™2 Duo E6750 (4M Cache, 2.66 GHz, 1333 MHz FSB)	Yes
2007	Core™2 Duo E6850 (4M Cache, 3.00 GHz, 1333 MHz FSB)	Yes
2007	Core™2 Duo E8190 (6M Cache, 2.66 GHz, 1333 MHz FSB)	No
2007	Core™2 Duo E8200 (6M Cache, 2.66 GHz, 1333 MHz FSB)	Yes
2007	Core™2 Duo E8400 (6M Cache, 3.00 GHz, 1333 MHz FSB)	Yes
2007	Core™2 Duo E8500 (6M Cache, 3.16 GHz, 1333 MHz FSB)	Yes
2007	Core™2 Extreme QX6800 (8M Cache, 2.93 GHz, 1066 MHz FSB)	Yes
2007	Core™2 Extreme QX6850 (8M Cache, 3.00 GHz, 1333 MHz FSB)	Yes
2007	Core™2 Extreme QX9650 (12M Cache, 3.00 GHz, 1333 MHz FSB)	Yes
2007	Core™2 Quad Q6600 (8M Cache, 2.40 GHz, 1066 MHz FSB)	Yes
2007	Core™2 Quad Q6700 (8M Cache, 2.66 GHz, 1066 MHz FSB)	Yes
2007	Core™2 Quad Q9300 (6M Cache, 2.50 GHz, 1333 MHz FSB)	Yes
2007	Core™2 Quad Q9450 (12M Cache, 2.66 GHz, 1333 MHz FSB)	Yes
2007	Core™2 Quad Q9550 (12M Cache, 2.83 GHz, 1333 MHz FSB)	Yes
2007	Pentium® D 935 (4M Cache, 3.20 GHz, 800 MHz FSB)	No
2007	Pentium® E2140 (1M Cache, 1.60 GHz, 800 MHz FSB)	Yes
2007	Pentium® E2160 (1M Cache, 1.80 GHz, 800 MHz FSB)	Yes
2007	Pentium® E2180 (1M Cache, 2.00 GHz, 800 MHz FSB)	No
2007	Pentium® E2200 (1M Cache, 2.20 GHz, 800 MHz FSB)	No
2008	Atom™ 230 (512K Cache, 1.60 GHz, 533 MHz FSB)	No
2008	Celeron® E1200 (512K Cache, 1.60 GHz, 800 MHz FSB)	No
2008	Celeron® E1400 (512K Cache, 2.00 GHz, 800 MHz FSB)	No
2008	Celeron® E1500 (512K Cache, 2.20 GHz, 800 MHz FSB)	No
2008	Core™2 Duo E4700 (2M Cache, 2.60 GHz, 800 MHz FSB)	No
2008	Core™2 Duo E7200 (3M Cache, 2.53 GHz, 1066 MHz FSB)	Yes
2008	Core™2 Duo E7300 (3M Cache, 2.66 GHz, 1066 MHz FSB)	Yes
2008	Core™2 Duo E7400 (3M Cache, 2.80 GHz, 1066 MHz FSB)	Yes
2008	Core™2 Duo E8300 (6M Cache, 2.83 GHz, 1333 MHz FSB)	No
2008	Core™2 Duo E8600 (6M Cache, 3.33 GHz, 1333 MHz FSB)	Yes
2008	Core™2 Extreme QX9770 (12M Cache, 3.20 GHz, 1600 MHz FSB)	Yes
2008	Core™2 Extreme QX9775 (12M Cache, 3.20 GHz, 1600 MHz FSB)	No
2008	Core™2 Quad Q8200 (4M Cache, 2.33 GHz, 1333 MHz FSB)	Yes
2008	Core™2 Quad Q8300 (4M Cache, 2.50 GHz, 1333 MHz FSB)	Yes
2008	Core™2 Quad Q9400 (6M Cache, 2.66 GHz, 1333 MHz FSB)	Yes
2008	Core™2 Quad Q9650 (12M Cache, 3.00 GHz, 1333 MHz FSB)	No
2008	Pentium® E2220 (1M Cache, 2.40 GHz, 800 MHz FSB)	Yes
2008	Pentium® E5200 (2M Cache, 2.50 GHz, 800 MHz FSB)	Yes
2008	Pentium® E5300 (2M Cache, 2.60 GHz, 800 MHz FSB)	Yes
2009	Celeron® E1600 (512K Cache, 2.40 GHz, 800 MHz FSB)	No
2009	Celeron® E3200 (1M Cache, 2.40 GHz, 800 MHz FSB)	No
2009	Celeron® E3300 (1M Cache, 2.50 GHz, 800 MHz FSB)	No
2009	Core™ i3-530 (4M Cache, 2.93 GHz)	Yes
2009	Core™ i3-540 (4M Cache, 3.06 GHz)	Yes
2009	Core™ i5-650 (4M Cache, 3.20 GHz)	Yes
2009	Core™ i5-660 (4M Cache, 3.33 GHz)	Yes
2009	Core™ i5-661 (4M Cache, 3.33 GHz)	Yes
2009	Core™ i5-670 (4M Cache, 3.46 GHz)	Yes
2009	Core™ i5-750 (8M Cache, 2.66 GHz)	Yes
2009	Core™ i7-860 (8M Cache, 2.80 GHz)	Yes
2009	Core™ i7-870 (8M Cache, 2.93 GHz)	Yes
2009	Core™ i7-920 (8M Cache, 2.66 GHz, 4.80 GT/s Intel® QPI)	Yes
2009	Core™ i7-940 (8M Cache, 2.93 GHz, 4.80 GT/s Intel® QPI)	Yes
2009	Core™ i7-950 (8M Cache, 3.06 GHz, 4.80 GT/s Intel® QPI)	Yes
2009	Core™ i7-960 (8M Cache, 3.20 GHz, 4.80 GT/s Intel® QPI)	Yes
2009	Core™ i7-965 Extreme Ed. (8M Cache, 3.20 GHz, 6.40 GT/s Intel® QPI)	Yes
2009	Core™ i7-975 Extreme Ed. (8M Cache, 3.33 GHz, 6.40 GT/s Intel® QPI)	Yes

Entry Year	MPU Description ¹	SPEC Score?
2009	Core™2 Duo E7500 (3M Cache, 2.93 GHz, 1066 MHz FSB)	Yes
2009	Core™2 Duo E7600 (3M Cache, 3.06 GHz, 1066 MHz FSB)	Yes
2009	Core™2 Quad Q8200S (4M Cache, 2.33 GHz, 1333 MHz FSB)	No
2009	Core™2 Quad Q8400 (4M Cache, 2.66 GHz, 1333 MHz FSB)	Yes
2009	Core™2 Quad Q8400S (4M Cache, 2.66 GHz, 1333 MHz FSB)	No
2009	Core™2 Quad Q9400S (6M Cache, 2.66 GHz, 1333 MHz FSB)	No
2009	Core™2 Quad Q9505 (6M Cache, 2.83 GHz, 1333 MHz FSB)	No
2009	Core™2 Quad Q9505S (6M Cache, 2.83 GHz, 1333 MHz FSB)	No
2009	Core™2 Quad Q9550S (12M Cache, 2.83 GHz, 1333 MHz FSB)	No
2009	Pentium® E5400 (2M Cache, 2.70 GHz, 800 MHz FSB)	Yes
2009	Pentium® E6500 (2M Cache, 2.93 GHz, 1066 FSB)	No
2010	Atom™ D410 (512K Cache, 1.66 GHz)	No
2010	Atom™ D425 (512K Cache, 1.80 GHz)	No
2010	Atom™ D510 (1M Cache, 1.66 GHz)	No
2010	Atom™ D525 (1M Cache, 1.80 GHz)	No
2010	Celeron® E3400 (1M Cache, 2.60 GHz, 800 MHz FSB)	No
2010	Celeron® E3500 (1M Cache, 2.70 GHz, 800 MHz FSB)	No
2010	Core™ i3-550 (4M Cache, 3.20 GHz)	Yes
2010	Core™ i3-560 (4M Cache, 3.33 GHz)	Yes
2010	Core™ i5-655K (4M Cache, 3.20 GHz)	No
2010	Core™ i5-680 (4M Cache, 3.60 GHz)	Yes
2010	Core™ i5-750S (8M Cache, 2.40 GHz)	No
2010	Core™ i5-760 (8M Cache, 2.80 GHz)	No
2010	Core™ i7-860S (8M Cache, 2.53 GHz)	No
2010	Core™ i7-870S (8M Cache, 2.66 GHz)	Yes
2010	Core™ i7-875K (8M Cache, 2.93 GHz)	No
2010	Core™ i7-880 (8M Cache, 3.06 GHz)	Yes
2010	Core™ i7-930 (8M Cache, 2.80 GHz, 4.80 GT/s Intel® QPI)	Yes
2010	Core™ i7-970 (12M Cache, 3.20 GHz, 4.80 GT/s Intel® QPI)	No
2010	Core™ i7-980X Extreme Ed. (12M Cache, 3.33 GHz, 6.40 GT/s Intel® QPI)	Yes
2010	Core™2 Quad Q9500 (6M Cache, 2.83 GHz, 1333 MHz FSB)	No
2010	Pentium® E5500 (2M Cache, 2.80 GHz, 800 MHz FSB)	No
2010	Pentium® E5700 (2M Cache, 3.00 GHz, 800 MHz FSB)	No
2010	Pentium® E5800 (2M Cache, 3.20 GHz, 800 MHz FSB)	No
2010	Pentium® E6800 (2M Cache, 3.33 GHz, 1066 FSB)	No
2010	Pentium® G6950 (3M Cache, 2.80 GHz)	Yes
2011	Atom™ D2500 (1M Cache, 1.86 GHz)	No
2011	Atom™ D2700 (1M Cache, 2.13 GHz)	No
2011	Celeron® G440 (1M Cache, 1.60 GHz)	No
2011	Celeron® G460 (1.5M Cache, 1.80 GHz)	No
2011	Celeron® G530 (2M Cache, 2.40 GHz)	No
2011	Celeron® G530T (2M Cache, 2.00 GHz)	No
2011	Core™ i3-2100 (3M Cache, 3.10 GHz)	Yes
2011	Core™ i3-2100T (3M Cache, 2.50 GHz)	Yes
2011	Core™ i3-2105 (3M Cache, 3.10 GHz)	No
2011	Core™ i3-2120 (3M Cache, 3.30 GHz)	Yes
2011	Core™ i3-2120T (3M Cache, 2.60 GHz)	Yes
2011	Core™ i3-2125 (3M Cache, 3.30 GHz)	No
2011	Core™ i3-2130 (3M Cache, 3.40 GHz)	Yes
2011	Core™ i5-2300 (6M Cache, up to 3.10 GHz)	Yes
2011	Core™ i5-2310 (6M Cache, up to 3.20 GHz)	Yes
2011	Core™ i5-2320 (6M Cache, up to 3.30 GHz)	Yes
2011	Core™ i5-2390T (3M Cache, up to 3.50 GHz)	Yes
2011	Core™ i5-2400 (6M Cache, up to 3.40 GHz)	Yes
2011	Core™ i5-2400S (6M Cache, up to 3.30 GHz)	Yes
2011	Core™ i5-2405S (6M Cache, up to 3.30 GHz)	No
2011	Core™ i5-2500 (6M Cache, up to 3.70 GHz)	Yes
2011	Core™ i5-2500K (6M Cache, up to 3.70 GHz)	Yes

Entry Year	MPU Description ¹	SPEC Score?
2011	Core™ i5-2500S (6M Cache, up to 3.70 GHz)	Yes
2011	Core™ i5-2500T (6M Cache, up to 3.30 GHz)	Yes
2011	Core™ i7-2600 (8M Cache, up to 3.80 GHz)	Yes
2011	Core™ i7-2600K (8M Cache, up to 3.80 GHz)	Yes
2011	Core™ i7-2600S (8M Cache, up to 3.80 GHz)	Yes
2011	Core™ i7-2700K (8M Cache, up to 3.90 GHz)	Yes
2011	Core™ i7-3930K (12M Cache, up to 3.80 GHz)	No
2011	Core™ i7-3960X Extreme Ed. (15M Cache, up to 3.90 GHz)	Yes
2011	Core™ i7-980 (12M Cache, 3.33 GHz, 4.8 GT/s Intel® QPI)	No
2011	Core™ i7-990X Extreme Ed. (12M Cache, 3.46 GHz, 6.40 GT/s Intel® QPI)	Yes
2011	Pentium® G620 (3M Cache, 2.60 GHz)	Yes
2011	Pentium® G620T (3M Cache, 2.20 GHz)	Yes
2011	Pentium® G630 (3M Cache, 2.70 GHz)	Yes
2011	Pentium® G630T (3M Cache, 2.30 GHz)	No
2011	Pentium® G6960 (3M Cache, 2.93 GHz)	No
2011	Pentium® G840 (3M Cache, 2.80 GHz)	Yes
2011	Pentium® G850 (3M Cache, 2.90 GHz)	Yes
2011	Pentium® G860 (3M Cache, 3.00 GHz)	Yes
2012	Atom™ D2550 (1M Cache, 1.86 GHz)	No
2012	Celeron® G465 (1.5M Cache, 1.90 GHz)	No
2012	Celeron® G540T (2M Cache, 2.10 GHz)	Yes
2012	Celeron® G550T (2M Cache, 2.20 GHz)	No
2012	Celeron® G555 (2M Cache, 2.70 GHz)	No
2012	Core™ i3-3220 (3M Cache, 3.30 GHz)	Yes
2012	Core™ i3-3220T (3M Cache, 2.80 GHz)	Yes
2012	Core™ i3-3225 (3M Cache, 3.30 GHz)	No
2012	Core™ i3-3240 (3M Cache, 3.40 GHz)	Yes
2012	Core™ i3-3240T (3M Cache, 2.90 GHz)	Yes
2012	Core™ i5-2380P (6M Cache, up to 3.40 GHz)	No
2012	Core™ i5-2450P (6M Cache, up to 3.50 GHz)	No
2012	Core™ i5-2550K (6M Cache, up to 3.80 GHz)	No
2012	Core™ i5-3330 (6M Cache, up to 3.20 GHz)	Yes
2012	Core™ i5-3350P (6M Cache, up to 3.30 GHz)	No
2012	Core™ i5-3450 (6M Cache, up to 3.50 GHz)	Yes
2012	Core™ i5-3450S (6M Cache, up to 3.50 GHz)	No
2012	Core™ i5-3470 (6M Cache, up to 3.60 GHz)	Yes
2012	Core™ i5-3470S (6M Cache, up to 3.60 GHz)	Yes
2012	Core™ i5-3470T (3M Cache, up to 3.60 GHz)	Yes
2012	Core™ i5-3475S (6M Cache, up to 3.60 GHz)	No
2012	Core™ i5-3550 (6M Cache, up to 3.70 GHz)	No
2012	Core™ i5-3550S (6M Cache, up to 3.70 GHz)	No
2012	Core™ i5-3570 (6M Cache, up to 3.80 GHz)	Yes
2012	Core™ i5-3570K (6M Cache, up to 3.80 GHz)	Yes
2012	Core™ i5-3570S (6M Cache, up to 3.80 GHz)	Yes
2012	Core™ i5-3570T (6M Cache, up to 3.30 GHz)	Yes
2012	Core™ i7-3770 (8M Cache, up to 3.90 GHz)	Yes
2012	Core™ i7-3770K (8M Cache, up to 3.90 GHz)	Yes
2012	Core™ i7-3770T (8M Cache, up to 3.70 GHz)	Yes
2012	Core™ i7-3820 (10M Cache, up to 3.80 GHz)	No
2012	Pentium® G2100T (3M Cache, 2.60 GHz)	Yes
2012	Pentium® G2120 (3M Cache, 3.10 GHz)	Yes
2012	Pentium® G640 (3M Cache, 2.80 GHz)	Yes
2012	Pentium® G640T (3M Cache, 2.40 GHz)	Yes
2012	Pentium® G645 (3M Cache, 2.90 GHz)	No
2012	Pentium® G645T (3M Cache, 2.50 GHz)	No
2012	Pentium® G860T (3M Cache, 2.60 GHz)	Yes
2012	Pentium® G870 (3M Cache, 3.10 GHz)	Yes
2013	Celeron® G1620 (2M Cache, 2.70 GHz)	No

Entry Year	MPU Description ¹	SPEC Score?
2013	Celeron® G1620T (2M Cache, 2.40 GHz)	No
2013	Celeron® G1630 (2M Cache, 2.80 GHz)	No
2013	Celeron® G470 (1.5M Cache, 2.00 GHz)	No
2013	Celeron® J1750 (1M Cache, 2.41 GHz)	No
2013	Celeron® J1850 (2M Cache, 2.00 GHz)	No
2013	Core™ i3-3210 (3M Cache, 3.20 GHz)	No
2013	Core™ i3-3245 (3M Cache, 3.40 GHz)	No
2013	Core™ i3-3250 (3M Cache, 3.50 GHz)	No
2013	Core™ i3-3250T (3M Cache, 3.00 GHz)	No
2013	Core™ i3-4130 (3M Cache, 3.40 GHz)	No
2013	Core™ i3-4130T (3M Cache, 2.90 GHz)	No
2013	Core™ i3-4330 (4M Cache, 3.50 GHz)	Yes
2013	Core™ i3-4330T (4M Cache, 3.00 GHz)	No
2013	Core™ i3-4340 (4M Cache, 3.60 GHz)	No
2013	Core™ i5-3340 (6M Cache, up to 3.30 GHz)	No
2013	Core™ i5-3340S (6M Cache, up to 3.30 GHz)	No
2013	Core™ i5-4430 (6M Cache, up to 3.20 GHz)	Yes
2013	Core™ i5-4430S (6M Cache, up to 3.20 GHz)	No
2013	Core™ i5-4440 (6M Cache, up to 3.30 GHz)	No
2013	Core™ i5-4440S (6M Cache, up to 3.30 GHz)	No
2013	Core™ i5-4570 (6M Cache, up to 3.60 GHz)	Yes
2013	Core™ i5-4570S (6M Cache, up to 3.60 GHz)	No
2013	Core™ i5-4570T (4M Cache, up to 3.60 GHz)	No
2013	Core™ i5-4670 (6M Cache, up to 3.80 GHz)	No
2013	Core™ i5-4670K (6M Cache, up to 3.80 GHz)	Yes
2013	Core™ i5-4670S (6M Cache, up to 3.80 GHz)	No
2013	Core™ i5-4670T (6M Cache, up to 3.30 GHz)	No
2013	Core™ i7-3970X Extreme Ed. (15M Cache, up to 4.00 GHz)	No
2013	Core™ i7-4765T (8M Cache, up to 3.00 GHz)	No
2013	Core™ i7-4770 (8M Cache, up to 3.90 GHz)	Yes
2013	Core™ i7-4770K (8M Cache, up to 3.90 GHz)	No
2013	Core™ i7-4770R (6M Cache, up to 3.90 GHz)	No
2013	Core™ i7-4770S (8M Cache, up to 3.90 GHz)	No
2013	Core™ i7-4770T (8M Cache, up to 3.70 GHz)	No
2013	Core™ i7-4771 (8M Cache, up to 3.90 GHz)	No
2013	Pentium® G2010 (3M Cache, 2.80 GHz)	No
2013	Pentium® G2020 (3M Cache, 2.90 GHz)	Yes
2013	Pentium® G2020T (3M Cache, 2.50 GHz)	No
2013	Pentium® G2030 (3M Cache, 3.00 GHz)	No
2013	Pentium® G2030T (3M Cache, 2.60 GHz)	No
2013	Pentium® G2120T (3M Cache, 2.70 GHz)	No
2013	Pentium® G2130 (3M Cache, 3.20 GHz)	Yes
2013	Pentium® G2140 (3M Cache, 3.30 GHz)	No
2013	Pentium® G3220 (3M Cache, 3.00 GHz)	Yes
2013	Pentium® G3220T (3M Cache, 2.60 GHz)	No
2013	Pentium® G3240T (3M Cache, 2.70 GHz)	No
2013	Pentium® G3420 (3M Cache, 3.20 GHz)	Yes
2013	Pentium® G3430 (3M Cache, 3.30 GHz)	No
2013	Pentium® J2850 (2M Cache, 2.41 GHz)	No

1. MPU description is identical to that on Intel's ARK database except for the following changes to save space: "Intel®" at beginning of description has been omitted, "Processor" has been omitted, and "Edition" and "Technology" have been abbreviated to "Ed." and "Tech." respectively.

APPENDIX C: REGRESSION RESULTS USING MPU CHARACTERISTICS

Table C.1. Introduction period only, 2000-06

	2000-01	2001-02	2002-03	2003-04	2004-05	2005-06
<i>Year dummy</i>	-1.227** (.242)	-.465 (.304)	-.561** (.168)	-.093 (.283)	-.092 (.143)	-.155 (.351)
<i>ln Clock speed</i>	5.28* (1.88)	3.05 (2.17)	5.28* (1.99)	3.53 (1.74)	4.35** (1.06)	2.73 (1.80)
<i>ln Power</i>	-3.19* (1.24)	-1.10 (1.61)	-2.50 (1.95)	-.29 (1.62)	.69 (.66)	-.19 (1.18)
<i>Number of cores</i>					.31 (.28)	.50 (.55)
<i>Number of threads</i>						
<i>ln Lithography</i>	2.45* (1.02)	.54 (1.32)		2.51* (.68)	1.22** (.36)	.29 (1.16)
<i>ln Cache</i>	-.70 (.64)	-.73 (.59)	.28* (.09)	.19* (.07)	.39* (.17)	.17 (.58)
<i>GPU dummy</i>						
Number of Obs.	20	18	14	11	24	31
Adjusted R ²	.64	.56	.81	.91	.88	.26

Table C.2. Full sample, 2000-06

	2000-01	2001-02	2002-03	2003-04	2004-05	2005-06
<i>Year dummy</i>	-.928** (.097)	-.427** (.070)	-.422** (.082)	-.378** (.056)	-.140* (.062)	-.275** (.087)
<i>ln Clock speed</i>	3.14** (.76)	2.71** (.55)	2.08** (.63)	2.12** (.70)	4.16** (.49)	3.41** (.58)
<i>ln Power</i>	-1.55** (.54)	-1.47** (.40)	-1.08* (.54)	.77 (.60)	.39 (.33)	-.24 (.37)
<i>Number of cores</i>					.55** (.12)	.57** (.14)
<i>Number of threads</i>						
<i>ln Lithography</i>	.94 (.75)	.69 (.46)	1.33 (.71)	2.65** (.23)	1.99** (.16)	1.09** (.30)
<i>ln Cache</i>	-.27 (.50)	-.36 (.24)	.34** (.07)	.27** (.03)	.26** (.02)	.26** (.07)
<i>GPU dummy</i>						
Number of Obs.	100	111	94	78	110	168
Adjusted R ²	.51	.54	.56	.91	.88	.49

Table C.3. Introduction period only, 2006-13

	2006-07	2007-08	2008-09	2009-10	2010-11	2011-12	2012-13
<i>Year dummy</i>	-.286 (.302)	.183 (.114)	-.460** (.134)	-.115 (.091)	-.228* (.083)	-.353** (.063)	-.060 (.073)
<i>ln Clock speed</i>	.02 (1.45)	.02 (.42)	4.34** (.75)	5.05** (.59)	2.86** (.39)	2.04** (.29)	2.02** (.51)
<i>ln Power</i>	.54 (.98)	2.53** (.39)	.93* (.44)	-.24 (.34)	-1.22** (.17)	-.89** (.13)	-.80** (.25)
<i>Number of cores</i>	.11 (.43)	-.38** (.11)	.29 (.15)	.61** (.18)	.37** (.11)	.38** (.09)	.39** (.14)
<i>Number of threads</i>			.03 (.03)	.11* (.04)	.02 (.03)	.08** (.02)	.09** (.03)
<i>ln Lithography</i>	1.08 (.91)	1.37** (.39)	.56 (.79)	3.73** (1.08)	.41 (.50)	-.33 (.19)	-.15 (.30)
<i>ln Cache</i>	.87 (.56)	.85** (.15)	.10 (.22)	-.42 (.39)	.44 (.36)	.15 (.25)	.15 (.35)
<i>GPU dummy</i>			-.03 (.30)	1.03* (.46)	-.41 (.22)	-.24 (.15)	
Observations	29	29	30	27	35	50	32
Adjusted R ²	.60	.96	.93	.94	.97	.97	.93

Table C.4. Full sample, 2006-13

	2006-07	2007-08	2008-09	2009-10	2010-11	2011-12	2012-13
<i>Year dummy</i>	-.414** (.132)	-.271** (.086)	-.348** (.060)	-.148** (.038)	-.161** (.033)	-.188** (.026)	-.028 (.022)
<i>ln Clock speed</i>	1.59* (.62)	.47 (.42)	1.14** (.31)	2.08** (.26)	2.10** (.21)	2.36** (.17)	2.21** (.17)
<i>ln Power</i>	-.43 (.40)	-.49 (.34)	.09 (.27)	-.53* (.19)	-.63** (.09)	-.90** (.08)	-.96** (.08)
<i>Number of cores</i>	.55* (.18)	.43** (.12)	.25** (.08)	.33** (.06)	.22** (.04)	.28** (.03)	.47** (.05)
<i>Number of threads</i>			.07** (.02)	.09** (.01)	.07** (.01)	.07** (.01)	.10** (.01)
<i>ln Lithography</i>	.86 (.44)	.71* (.28)	.38 (.20)	.37 (.20)	-.43* (.17)	.13 (.10)	.41** (.08)
<i>ln Cache</i>	.35 (.24)	.64** (.13)	.47** (.08)	.36** (.06)	.44** (.07)	.44** (.07)	-.03 (.15)
<i>GPU dummy</i>			.13 (.18)	.06 (.10)	-.36** (.07)	-.23** (.04)	-.13 (.08)
Observations	162	170	209	236	291	296	217
Adjusted R ²	.44	.67	.76	.85	.87	.91	.95

Note: The dependent variable is ln(MPU price); the regressions include a constant, not shown above. Standard errors are in parentheses. * and ** indicate significance at the 5% and 1% levels, respectively. An omitted coefficient indicates there was no variation in that variable across models.