The Microeconomics of Microprocessor Innovation

Kenneth Flamm

First Draft January 2006
Revised December 2006, February 2007, March 2007, June 2007

VERY ROUGH, PRELIMINARY DRAFT

Comments Only Please

Kenneth Flamm
LBJ School of Public Affairs
University of Texas at Austin
Box Y
Austin, Texas 78713

kflamm@mail.utexas.edu

The late 1990s were a period of unusually rapid technological progress in the semiconductor components responsible for a considerable portion of technological improvement in information technology hardware.[1] Estimates suggest, for example, that from 40 to 60 percent of the decline in quality-adjusted prices for computers around this time was attributable to improvements in price-performance for semiconductors going into computers.[2] Similarly, a rough estimate suggests that from 20 to 30 percent of declines in quality-adjusted communications equipment prices were attributable to improved semiconductors used in building this equipment.[3]

Table 1 lays out the contours of price declines in semiconductors in the late 1990s. The acceleration in quality-adjusted price declines after 1995 is quite large and noticeable. The rate of decline in microprocessor prices picked up by about 50% in the second half of the decade, compared with the early 1990s. Memory chips, particularly the workhorse dynamic random access memories (DRAMs) used in computers, registered even greater acceleration in price declines. Though less impressive than annual declines (exceeding 60 percent) registered in the late 1990s for microprocessors and memory, declines in virtually all other kinds of semiconductor prices accelerated over this period. The faster pace of technical innovation in semiconductors was a major factor behind a broader pickup in the pace of IT innovation over this period, and even more importantly, in productivity growth in the overall economy.[4]

Clearly, common forces seem to have accelerated the rate of technological progress throughout the chip industry in the late 1990s. As some have noted, a faster pace for the

---

[1] I am most grateful to Pablo Cruzat, Anjum Khurshid, Kevin Williams, Erik Schuchmann, Caroline Alexander, Javier Beverinotti, and Angela Newell for their outstanding research assistance on various elements of this project. Without implicating them in my errors, I am also grateful to Vinod Aggarwal , Yaichi Aoshima, Fred Chang, Hiroyuki Chuma, Gary Chapman, Shane Greenstein, Dale Jorgenson, Arati Prabhakar, Bill Raduchel, Marc Snir, Bill Spencer, Jack Triplett, Dave Tuttle, members of the National Research Council's Science, Technology, and Economic Policy Board, and participants in colloquia at the UT Austin LBJ School of Public Affairs, the UC Berkeley Institute of International Studies, and the National Institute of Science and Technology Policy, Japan, for their comments on earlier versions of these ideas, presented in seminars or meetings in January 2005, October 2005, January 2006, February 2006, March 2006, May 2006, October 2006, and February 2007.

[2] See Aizcorbe, Flamm, and Kurshid, "The Role of Semiconductor Inputs in IT Hardware Price Decline: Computers vs. Communications," Federal Reserve Finance and Economics Discussion Paper 2002-37, (Board of Governors, The Federal Reserve Board, Washington) August, 2002; revised, 2004, forthcoming in E. Berndt, Ed., ***Hard to Measure Goods and Services—Essays in Honor of Zvi Griliches***, (Chicago and National Bureau of Economic Research).

[3] Ibid.

[4] See Jorgenson and Stiroh, "Raising the Speed Limit: U.S. Economic Growth in the Information Age", ***Brookings Papers on Economic Activity***, G. Perry and W.C. Brainard, eds., (Washington, DC: Brookings Institution Press), 2000; Jorgenson, "Information Technology and the U.S. Economy," ***American Economic Review***, vol. 91, no. 1, March 2001. This is not to say that forces other than greater use of IT capital did not also play a major role in productivity growth in the broader U.S. economy; see for example Bosworth and Triplett, "The Early 21st Century Productivity Expansion is *Still* in Services," **International Productivity Monitor**, No. 14, Spring 2007.

introduction of new manufacturing technology into the industry after 1995 seems a likely explanation for at least part of the industry-wide acceleration in technical progress.[5]

This acceleration in the speed of technical innovation in semiconductor manufacturing seems to have had an explicit policy component. Greater rates of technical innovation in chip manufacturing were a public goal of newly created and quite unique industrial institutions (SEMATECH, the National Semiconductor Technology Roadmap, and later, the International Technology Roadmap for Semiconductors) coordinating semiconductor R&D across firms, at first within the US, and later, globally.[6] The coordination process implemented by these new, private institutions was sanctioned, to some extent, by laws passed in the 1980s, and the tacit approval of national governments observing this coordination across firms and periodic public announcements of institutional objectives.

This paper explores to what extent the more rapid pace of technical progress in microelectronics in the late 1990s was attributable to technological innovation in manufacturing—coordinated or uncoordinated—and to what extent it was due to other sources of technical progress, for a key semiconductor product—microprocessors. Computer processor units are an intrinsically important product. Microprocessors are the largest single semiconductor input, in terms of value, in personal computers,[7] and are the technological core of all computers, big and small.

Table 1 shows that microprocessors are somewhat unique. They had by far the highest rate of decline in quality-adjusted price for any semiconductor product class in the late 1990s. Microprocessors are also by far the largest value semiconductor product manufactured in the US, accounting for over 46% of US factory shipments of integrated circuits (ICs) in 2004.[8] The disproportionate role of microprocessors (MPUs) in US shipments (contrast this with a share for MPUs of about 17% in global IC sales[9]) almost certainly means that processors (with atypically high margins) account for an even larger share of value added in the US semiconductor industry. Since the semiconductor industry has been the largest manufacturing industry in the US (measured by value added),

---

[5] See Flamm, 2001, 2003a, 2003b, 2004; Aizcorbe, Oliner, and Sichel, 2003.

[6] See Flamm, "Economic Impacts of International R&D Coordination: SEMATECH, the International Technology Roadmap, and Innovation in Microprocessors," presented at the NISTEP/STEP conference on "21st Century Innovation Systems for Japan and the United States: Lessons from a Decade of Change," Tokyo, Japan, January 10-11, 2006.

[7] For example, in 2001, Hennessey and Patterson estimate the processor to be the largest single cost element (22%) of the components used in a $1,000 PC. Memory (DRAM) is 5%, the video card 5%, the hard disk 9%, the monitor 19%, and the operating system 20%. All components together account for 47% of the value of this $1,000 PC. See Hennessy and Patterson (2003), pp. 21-24.

[8] This figure, based on U.S. Census, **Current Industrial Reports: Semiconductors, Printed Circuit Boards, and Other Electronic Components - 2004** (2005) excludes 16-bit microprocessors, where data have been suppressed. The role of MPUs in US semiconductor manufacture also increased greatly in the late 1990s. In 1995, the **Current Industrial Reports** shows all microprocessors accounting for fewer than 30% of US chip shipments.

[9] Calculated by the author from data contained in the World Semiconductor Trade Statistics statistical printout for December, 2004. The author thanks the Semiconductor Industry Association for making this data available to him.

microprocessors are arguably the single most important manufactured product now produced in the US economy.

My analysis is particularly timely because the most recent data, analyzed below, show a sharp decline in the pace of price-performance improvement in microprocessors beginning around 2003. Absent new pathways for innovation in microprocessors, this slowdown in microprocessors may have a broader ripple effect in reducing the pace of progress in computers, the deployment of IT in the US economy, and ultimately slowing gains in productivity across IT-using sectors of the American economy, over the next several years.

In attempting to explain why the rate of quality-adjusted price decline for microprocessors has declined so rapidly, it is helpful to draw an analytical distinction between improvements in quality-adjusted semiconductor price due to changes in the cost of manufacturing an electronic device (I use "transistor" as shorthand to describe all such devices) on an integrated circuit (IC), and improvements in the qualities of the transistors and functionality of the ICs designed to make use of these transistors.

We should recognize that drawing a sharp distinction between improvements in the manufacturing process for the transistors, and improvements in the products designed using these transistors, is to some extent artificial. Much design innovation is induced by the declining cost of the transistors used as building blocks in increasingly complex product designs. And process innovation may to some extent improve the quality of the product even if no improvements are made to a product design—smaller, cheaper transistors may allow an existing product design to run faster, even if no other changes are made to this design. Altering an existing design to optimize use of cheaper, faster transistors may enable still further gains. Nonetheless, we can at least conceptually consider the effect of a "pure" decline in manufacturing cost, where an existing product is simply produced at a lower cost, and additional benefits attainable from tuning designs to make better use of cheaper (and possibly faster or less power-hungry) components simply ignored.

The plan of this paper is to first establish, in an approximate way, what sorts of improvements in quality-adjusted price for leading edge semiconductors might be expected, based solely on the pace of "pure" manufacturing technology innovation observed in the late 1990s. I then measure what the observed rate of quality-adjusted price decline in Intel desktop microprocessors actually was, and how it has varied over time. I argue that the residual—the difference between the actual and the baseline expected effect of "pure" manufacturing innovation—must largely be attributable to design innovations in microprocessors (which may well have been enabled and induced to some extent by innovation in manufacturing). I then sketch the historical origins of major design innovations adopted in Intel desktop processor designs over this period. I argue that most of the design innovations embodied in improved Intel desktop microprocessors since 1995 can be traced back to R&D on high performance computing over the three decades from the 1960s through the 1980s, and that the rapid decline in quality-adjusted price for these processors since 1995 largely reflected this stock of

previously developed ideas, produced by research and experimentation in high performance computer design, being rapidly moved into Intel PC processor designs over this period. Finally, I note that there has been a significant slowdown in the decline of quality adjusted processor price since early 2003, which I argue is consistent with the "pipeline" of proven good ideas for enhancing processor performance through improved design largely being emptied over the previous decade. A final section of this paper considers what trajectory for quality-adjusted price for microprocessors is likely in the near future, and speculates on the possible economic implications of these developments.

Manufacturing Costs for Semiconductors: An Overview

I begin with some general observations on the potential contributions of semiconductor manufacturing innovation to declining semiconductor prices. The most commonly used index of technological advance in semiconductor manufacturing is the size of the smallest features that can be patterned by leading edge lithographic equipment. Prior to 1995, a new generation of lithographic and supporting equipment (commonly referred to within the industry as a "technology node") was introduced at approximately three year intervals. Each new technology node historically reduced minimum feature size by about 30%. The net effect was to reduce the surface area required for a transistor by about 50%.[10] Thus, on any given area of silicon, the introduction of a new technology node could be expected to double the number of transistors etched.

To this I add the observation that, over the long run, manufacturing costs per area of silicon have remained roughly constant (see Figure 1, based on SEMATECH engineering cost estimates).[11] Manufacturing costs per area, on some given diameter silicon wafer, rise over time with each new technology node, but are offset by a sharp downward fall every three or so technology nodes, when a new wafer diameter is introduced. The net effect is a zigzag pattern in leading edge IC manufacturing costs per area of silicon over time, with a trend that is almost flat. If we also factor in a trend to steadily rising yields as successive generations of products reach volume production, it seems reasonable to propose that wafer fabrication costs per area of good silicon have remained roughly constant.[12]

---

[10] Since .7 x .7 = .49, or about a 50% reduction in area.

[11] Over the 1983-1998 period, one estimate is that overall wafer processing cost per square centimeter of silicon increased at a compound annual growth rate of 5.5%. See Carl Cunningham, Denis Fandel, Paul Landler, and Robert Wright, "Silicon Productivity Trends," International SEMATECH Technology Transfer #00013875A-ENG, February 29, 2000, p. 5. Note that this estimate is per total silicon area processed, not cost per good yielded area. Since good yielded area appears to have increased over time as a fraction of total wafer area processed, with improved processing yields, it is not wholly unreasonable to assume that wafer processing cost per good, yielded silicon area was roughly constant over time.

[12] That is, if (c *[1/ fraction yielded good chips]) is substituted for c in the numerator of (1), and has remained roughly constant, i.e., slowly rising increases in area processing costs have just been offset by steadily increasing yields. For solid evidence that dynamic random access memory (DRAM) chip yields have increased steadily over time, for successive generations of DRAMs, see Charles H. Stapper and Raymond J. Rosner, "Integrated Circuit Yield Management and Yield Analysis: Development and Implementation," *IEEE Transactions on Semiconductor Manufacturing*, 8:2, 1995, p. 100; Rainier Cholewa, "16M DRAM Manufacturing Cooperation IBM/SIEMENS in Corbeil Essonnes in France," *Proceedings of the 1996 IEEE/SEMI Advanced Semiconductor Manufacturing Conference*, 1996, p. 222.

A limited amount of actual data on IC manufacturing costs is available by studying the financials of so-called "pure play" foundry semiconductor manufacturers in Taiwan. These foundries manufacture semiconductors designed by others in their fabrication facilities, as a form of contract manufacturing. Their manufacturing costs are entirely related to the cost of fabricating wafers. Since these firms also typically report the number of 8" equivalent wafers shipped as well, it is possible to calculate their average manufacturing costs per area of silicon, on a quarterly basis. Figure 2 shows COGS (Cost of Goods Sold, which I interpret as roughly equal to variable costs plus depreciation on plant and equipment) per area of silicon shipped for the two largest foundry semiconductor manufacturers in the world, Taiwan Semiconductor Manufacturing Corporation (TSMC) and United Microelectronics Corporation (UMC), which together account for about 75% of the "pure play" foundry capacity in the world. As can be seen, actual manufacturing costs per area of silicon seem to have remained roughly constant over the last decade, with no obvious trend, despite frequent introduction of new technology nodes, and somewhat less frequent increases in wafer sizes. The spikes in manufacturing cost per area of silicon seem to be correlated with decreased capacity utilization, as might be expected, as well as the introduction of new technology nodes; the declines in fabrication cost per wafer with increased capacity utilization and transitions to larger size wafers.[13]

Combining these two observations, we can approximate manufacturing cost per transistor, m, as follows:

(1) $m = c / t$

where m is manufacturing cost per device (i.e., per transistor);
c is processing cost per area silicon wafer;
t is device density, number transistors per area silicon wafer.

If the exact same chip design is produced using the new and smaller transistor manufacturing process every three years,[14] then, manufacturing cost drops by one half every three years. This works out to a decline rate in manufacturing cost of -21% per year, purely as the result of manufacturing cost improvements associated with the shift to newer manufacturing technologies. If we are willing to assume that price-cost margins have been roughly constant over time, we would see prices falling at the same rate.[15]

---

[13] Since the cost of goods sold (COGS) concept used in company accounts includes an allocation for depreciation costs, the higher capacity utilization, the lower the fixed cost of depreciation per wafer actually fabricated. These data are taken from public company quarterly and annual reports, and earnings presentations for TSMC and UMC.

[14] For a much more detailed analysis relating the analysis of this section to "Moore's Law" (which holds that transistors per IC—not transistor density t—doubles every 18 months), see Flamm, 2001, 2003a, 2003b, 2004.

[15] For a discussion of changes in price-cost margins in microprocessors and memory chips, see Aizcorbe, Oliner, and Sichel (2003). There is little or no evidence of a secular trend in price-cost markup evident in the data they examine; see their figure 5. My conversations with Intel executives suggest that at Intel, the

In reality, data for the period 1975-1995 show prices for leading edge ICs[16], which presumably show the effect of new manufacturing technology introduction most directly, falling at a rate about 50% faster! What was going on?

The answer is that other improvements in manufacturing technology, in addition to lithography, have made further improvements in device density possible. The lithographic improvements allow more devices to be crammed into a two-dimensional plane, but other advances have permitted manufacturers to cram more devices on a chip in three dimensions. Better etching and deposition equipment makes it easier to create precise vertical features, like trenches, and has reduced the two-dimensional footprint required for many silicon devices. The complex interconnections among devices, rather than using up precious two-dimensional real estate on the surface of the silicon wafer, have been moved upward, in three dimensions, in additional layers making use of new technologies for producing acceptably flat surfaces on which additional layers of silicon insulation and metal interconnections are deposited.[17] In DRAMs (memory), historically, t, rather than merely doubling every three years, actually increased 2.9 times. Equivalently, t increased at a rate of 43% annually, rather than 26%.[18] (See the first three rows of Table 2 for the arithmetic of this linkage.)

If, in equation (1), t is increased by 2.9X (an annual increase of 43%), rather than merely doubling every three years, the decline rate for transistor cost increases to 30 percent per year. This is pretty much the average historical decline rate observed for both DRAM and microprocessor prices over the years 1975-1995.

In undertaking the approximation in equation (1), we have ignored a number of factors. These include yield improvements (reductions in defective products produced) at all stages of semiconductor production, due to learning economies in semiconductor production, intra-node reductions in minimum feature size ("die shrinks", also associated with learning economies), the costs of assembly, packaging, and testing of finished products, scale economies (since prices charged for products must cover substantial, relatively fixed costs of R&D, and rising minimum efficient scales of output for production facilities), and changes in price-cost margins that might be associated with, for example, increased or decreased levels of competition within different product segments. That the long-run rate of price decline for leading edge products—DRAMs and microprocessors—was so close to this "fundamental" rate of cost decline associated with the pace of manufacturing innovation over this two decade period suggests that, on balance, the role of changes in these other factors was relatively modest compared with the introduction of new technology nodes in manufacturing.

---

perceived behavior of microprocessor profit margins is one of up and down fluctuation, but no secular trend.

[16] In the 1970s and 1980s, DRAMs were the product most likely to utilize the newest manufacturing technology. In the 1990s, microprocessors too joined DRAMs at the leading edge, and most recently, flash memory has joined microprocessors and DRAMs at the manufacturing technology frontier.

[17] This was enabled by development of chemical-mechanical planarization techniques in the 1990s.

[18] Flamm, 2001, 2003a, 2003b.

Manufacturing Technology Acceleration in the Late 1990s

In the mid-1990s, US semiconductor R&D consortium SEMATECH adopted a strategic plan intended to focus its efforts on working with suppliers of equipment and materials to accelerate the introduction of new technology nodes. SEMATECH's intention was to reduce the time between introductions of new technology nodes from 3 years to 2 years. At about the same time and in collaboration with SEMATECH, the US National Advisory Commission on Semiconductors (a government entity) kicked off a series of workshops and reports that culminated in the first National Technology Roadmap, in 1994. A 1997 update to this roadmap codified the goal of a 2-year cycle of new technology nodes for the entire US semiconductor industry. Effectively the entire US semiconductor industry, along with academia and interested agencies in the US government, participated in these efforts to set technical goals, and coordinate research efforts across firms in order to accomplish those goals.[19]

In the late 1990s, this roadmap evolved into a global roadmapping effort (the International Technology Roadmap for Semiconductors, ITRS), involving the semiconductor industries of the US, Japan, Europe, Korea, and Taiwan. While the extent to which closer coordination among specialized suppliers of equipment and materials, and their users in the semiconductor industry, has effectively worked to accelerate innovation in the industry may be debated, one thing that is clear is that the industry did shift to a new accelerated timetable, approaching two years per node. While a fascinating subject for further research, the precise role of the roadmap, versus other factors, in accelerating semiconductor manufacturing innovation is not central to the analysis within this paper. Despite periodic attempts to slow down this accelerated schedule for technical innovation, the industry has apparently continued field new technology nodes roughly every two years, since the mid-1990s.

If we posit that density, t, continues to increase by 2.9X between nodes—now a 2-year time span, then using equation (1), we would now predict a 41% annual decline in the cost of manufacturing a transistor. (See the last row of Table 2 for the arithmetic.) Looking back at Table 1, we can see that the decline in prices for microprocessors in the late 1990s greatly exceeded even this much higher rate. Understanding what accounted for the difference between predicted declines in manufacturing costs for microprocessors, and microprocessor prices, since 1995, is the focus of the balance of this paper.

Manufacturing Innovation in Microprocessors: A Baseline

While there is no reliable data publicly available on actual manufacturing costs, or wafer processing costs, for microprocessors, there is some public data available on transistor density. Figure 3 shows transistors per area of silicon for newly introduced Intel desktop microprocessor ICs over the period 1993-2004. A trend line fitted to these data shows a 53% annual rate of increase for transistor density in Intel desktop microprocessors.

---

[19] A more fully developed and documented version of this history may be found in Flamm (2006).

Table 2 calibrates equation (1) to different scenarios for microprocessors. The third line of this table shows that if, as in DRAMs, there were a 2.9 X increase in transistor density every three years, we would expect manufacturing cost to drop by 66% every three years (row 3). This scenario, discussed above, corresponds to an annual growth in transistor density of 43%, and an annual drop in transistor cost of 30%. With a two year technology node cycle, transistor density would instead increase by 70% annually, and transistor cost would drop by 41% annually (line 9 in Table 2). This is a significantly higher rate of increase in transistor density than actually observed in Intel desktop processor introductions. As the sixth line in this table, shows, if we instead posit a 2.5 year period between introductions of new technology nodes, we get exactly the 53% rate of increase in transistor density observed in the actual Intel data. Equivalently, if we assumed a lesser degree of "ingenuity" than in DRAMs (t increasing by only 2.34 with every new technology node, rather than 2.9), with a 2 year cycle, we would again exactly predict the observed historical trend in Intel transistor density.

I conclude that either the technology acceleration in the late 1990s was not quite as aggressive for the "average" Intel desktop processor fabrication facility (2.5 rather than 2 years, between nodes), or the pace of improvement due to "ingenuity" in squeezing out more transistors per area, above and beyond the gains from "pure" lithographic improvement at a new technology node, slowed (from 2.9X to 2.34X).[20]

As Table 2 also shows, either of these last scenarios produces a 53% annual increase in transistor density, and an expected annual decline in transistor cost of 35%. Therefore, regardless of what produced it—a 2 year cycle and a slowdown in "ingenuity," or a 2.5 year cycle with no slowdown, or some other combination of manufacturing parameters that yielded a 53% annual increase in transistor density—this actual historical transistor density improvement observed with Intel desktop microprocessors since 1995 would be expected to produce a 35% annual decline in transistor cost, based solely on continuing manufacturing innovation. Having established a conceptual baseline for the expected decline in manufacturing cost due solely to introduction of new semiconductor manufacturing technology nodes, let us next examine the actual behavior of Intel desktop microprocessor prices.

Measuring Quality-Adjusted Prices for Microprocessors

[20] The slowdown in "application of ingenuity" in this latter scenario represents a genuine slowing, and goes beyond what might be attributed to the shorter time period between nodes (i.e., 2 vs. 3 years). The compound average growth rate in transistor density attributable to "ingenuity" historically was about 13% (that is, a thirteen percent annual improvement applied to the 26% CAGR associated with a doubling in transistor density, compounded over three years, yields an increase from 2X to 2.9X in density). If we apply that same additional annual increment in density (13%) to the underlying annual improvement in density associated with a two year doubling in transistor density (41.4%), we get a rate of improvement in transistor density of 60%, which exceeds the 53% Intel rate actually observed. Equivalently, a 13% annual increase in density from ingenuity applied over 2 years, increases the underlying two year doubling in transistor density to 2.56X, not the 2.34X needed to produce the observed 53% annual growth in transistor density.

I begin by tracing the contours of change over time in microprocessor prices using a unique, highly detailed data set. Since the mid-1990s, Intel has periodically published, or posted on the Internet, current list prices for its microprocessor product line, in 1000-unit trays. These list prices are available at a very disaggregated level of detail, distinguishing between similar models manufactured with different packaging, for example, and are typically updated every 4 to 8 weeks—though price updates have sometimes come at much shorter or longer intervals.[21] By combining these detailed prices with detailed published attributes of different processor models, it is possible to construct a very rich data set relating processor prices to processor characteristics.

This permits one to construct so-called "hedonic" price indexes, which relate processor prices to processor characteristics. There are three virtues of these hedonic price indexes, compared with other methodologies (so-called "matched model" indexes, in particular). First, the data are disaggregated at a level greater than is typically available in the data sources used with matched model indexes[22]. Most "matched model" indexes for microprocessors are actually based on industry consultant estimates of average selling prices and revenues across a set of close, but not precisely identical products. My data set permits measuring differences in processor characteristics down to individual models of processors, controlling for such things as processor speed, clock multiplier, bus speed, differing amounts of L1, L2, and L3 cache memory, architectural changes, voltage levels, package types, and particular new processor features and instructions. The latter have become particularly important recently—in mid-2004, Intel dropped processor clock speed as the principle characteristic used to differentiate processors, and introduced a complex "processor model number" system that distinguished between very small and arguably minor differences between processors that have proliferated in its recent product introductions.[23]

Second, it is typically difficult to deal satisfactorily with new products entering matched model indexes in their period of first sale, and obsolete products exiting matched model indexes in their last period of sale. Because very rapidly falling prices are typical when new semiconductors are first offered for sale, there is reason to believe that matched model indexes may somewhat underestimate quality-adjusted price declines in periods

---

[21] My data initially (over the 1995-1998 period) made use of compilations of this data collected by others and posted on the web; since 1998-99, most of this data was collected and archived directly off the Intel web site. Additional relevant details of data set construction go here in the next draft.

[22] Need to discuss data sources used by US statistical agencies here.

[23] This is not to suggest that Intel's processor number system was unique in spawning confusion about the processors it was describing. Back when Intel was using clock speed as the principal descriptor for different models of its Pentium processors, Intel's main competitor, AMD, devised a processor rating system based on supposed equivalence (or better) in performance to Intel processors with given clock speeds. So, for example, an "Athlon 64 3400+" (supposedly with performance exceeding that of an Intel Pentium 4 with a clock speed of 3.4 Gigahertz) is actually any of three quite different CPUs: an Athlon CPU running at 2.2 Ghz, with 1 MB of L2 cache, and a 64-bit wide memory bus; a 2.4 Ghz Athlon with 512KB L2 cache and a 64-bit wide memory bus, and a 2.2Ghz Athlon with 512KB L2 cache and a 128-bit wide memory bus. See "The AMD K8 Architecture," accessed at http://www.cpuid.com/reviews/K8/index.php, December 28, 2005.

when new processors are introduced.[24] Further, the rate of introduction of new models of Intel desktop processors has considerably increased in recent years.

Third, and most importantly, my detailed processor, price, and characteristics data set permits us to connect different models of Intel processors to the manufacturing technology used in their production. This allows us to evaluate the direct and indirect impacts on processor prices of changes in characteristics linked directly or indirectly to the introduction of new manufacturing technology, and assess what portion of the improvement in processor price-performance has been linked to this and other factors. This permits us to ask "why" a quality-adjusted price index changes over time, a question that cannot be answered using a matched-model price index.

My basic methodology in constructing a hedonic price index was to estimate a regression equation over one year periods, typically running from May of one year, through May of the next. The regression equation was of the form

$$(2) \qquad \ln p_{it} = b_0 + b_1 D_t + b_S \ln proc_{it} + b_A A_{it} + u_{it}$$

where p_{it} is processor model i's price at time t, $proc_{it}$ is processor clock rate at time t, A_{it} is one of a series of dummy variables for other measured processor characteristics at time t, D_t is a monthly time dummy representing an offset for a general price level from a base period incorporated into intercept term b_0, and u_{it} is an unexplained residual error term. We interpret this equation as a "reduced form," reflecting both consumer valuations of characteristics, and supplier costs and product introduction decisions over the period in question.[25] By allowing all coefficients to vary from one year to the next, we allow for structural change in technology, costs, market structure, and consumer preferences over time.[26]

---

[24] It has been argued that one should estimate a reservation price (a price where demand would just equal zero had the product been available, presumably using a hedonic methodology) for the period preceding the period of introduction, and use this as the "matched model" price for the missing observation. See for example, W.E. Diewert, "The Early History of Price Index Research", pp. 33-65 in W. E. Diewert and A. O. Nakamura (eds.) **Essays in Index Number Theory**, Volume 1, (Amsterdam: North Holland) 1993; Ana Aizcorbe, Carol Corrado, and Mark Doms, "When Do Matched-Model and Hedonic Techniques Yield Similar Price Measures?," Federal Reserve Board of San Francisco Working Paper 2003-14, 2003.

[25] A useful derivation of how a hedonic price surface can be derived as a reduced form from a model of monopolistically competitive producers and consumers with heterogeneous tastes may be found in D. Prentice and X. Yin, "Measuring Quality-Adjusted Inflation Rates for a Heterogeneous Oligopoly," Discussion Paper 00.06, La Trobe University School of Business, June 2000. An influential derivation of a hedonic function as the reduced form in a model of Bertrand competition among single product firms may be found in A. Pakes, "A Reconsideration of Hedonic Price Indices With An Application to PC's," **American Economic Review**, December, 2003. Pakes adopts the hypothesis of coefficient stability over one year time periods in the empirical portion of his article.

[26] A future version of this paper may break my sample into six-month subperiods, and test for coefficient stability over these six-month subperiods where feasible. One limit to coefficient stability tests is that adjusting the estimated variance-covariance matrix for possible heteroskedasticity produces tests on coefficients that are only asymptotically valid. As the number of observations producing any given set of coefficients shrinks with length of time period, the validity of large sample statistical test results becomes increasingly questionable. Note that, as previously mentioned, there seems to be no secular trend in

Data on Intel desktop processors (i.e., Pentiums and Celerons) was used. Price data was quite scanty in 1995 and 1996, when relatively few models of Pentiums were available and the pace of new product introduction relatively slow compared to later periods. Architectural family dummy (0-1) variables were coded for each new design family for these processors.[27]

Typically, there was exact collinearity between some of the characteristics collected in the data set, since, for example, all models having one design architecture may have the same amount of L1 or L2 cache, and the same bus speed. Differing cache sizes and bus speeds were measured as dummy variables within processor architecture families, and were only identifiable when there was variation in these variables within a processor architectural family. A subset of characteristics that were not exactly collinear was used as regressors in the hedonic equation above. In cases where a variable was exactly collinear with other characteristics, however, we must interpret the coefficient of the variable as reflecting the joint impact of all the exactly collinear variables, not just the one actually included in the regression. Table 3 provides an overview of design architectures making up my sample.

The one-year periods over which this regression was run were chosen to overlap in a single month. The coefficients of time dummy variables in the above regression were exponentiated and used to construct an index of monthly price levels within any given one year period. (If s is the base time period, then price in time period t relative to base period s is just $\exp(D_t)$.) The overlapping month in successive one year periods was used to chain these monthly indexes into a longer price index for multiple years. Because of gaps in months in which Intel published data, prices—and price indexes—were not available for all months. The actual one-year regressions are reproduced in a supplementary appendix to this paper.[28]

Table 4 shows the price indexes produced using the above methods. The price decline is impressive, falling from an index of 735,000 in June 1995, to an index of 47 in December 2006. Figure 4 shows these data in graphical form, with price on a logarithmic scale.

---

producer markups over most of this period. Similarly, available data on Intel and AMD market shares show only small changes over the entire period.

[27] Design "families" are commonly denominated as codenames by Intel—e.g., Katmai, Coppermine, Northwood, Conroe, etc. Within each of these design families, there are minor revisions and corrections to the original design, but these very small changes in processor design—known in the industry as "steppings"—within a family are viewed as quite minor, and generally not known to the consumer when a microprocessor (by itself, or within a computer) is purchased. Since all members of a processor design family are manufactured with the same generation of production technology (the same technology node), dummy variables for design families/chip architecture capture both supply side (manufacturing cost) and demand side (processor feature) characteristics. While we do have data on technology node used to manufacture processors, it is exactly collinear with (and therefore indistinguishable from) the design family/architecture dummies.

Similarly, our monthly dummy variables capture all demand and supply (cost) factors which vary from month to month that are not controlled for explicitly by including other variables.

[28] Because heterogeneity in consumers underlies a hedonic demand model, in which different groups of consumers choose different bundles of characteristics, robust, heteroscedasticity-consistent estimators for standard errors are also shown, in addition to uncorrected standard errors.

Figure 5 charts compound monthly decline rates, calculated back to the last previous month in which I was able to estimate the price index. The marked slowing of declines in quality-adjusted price after 2003 is quite apparent. A similar post-2003 slowdown seems to have occurred in absolute processor performance (see Figure 6, taken from a presentation by computer scientist David Patterson, of UC Berkeley.)

Table 5 shows that my estimated price indexes are quite consistent with other available price indexes for microprocessors. My estimates over comparable time periods are quite similar to the matched model index results of Aizcorbe, Corrado, and Doms, and to the producer price index for microprocessors, microcontrollers, and other non-memory semiconductors (which currently is a fixed weight price index making use of hybrid price adjustments).[29]

Detailed results for the May 2004-May 2005 period are shown in Table 6, and discussion of these results is useful in understanding results for other periods, where a similar methodology was used. For reasons of time and space, this discussion is foregone in this draft of this paper.

The hedonic results show that processor speed had a very large impact on price. The estimated elasticity of price with respect to clock speed, holding all else constant, was 3.25. That is, a ten percent increase in clock speed was associated with a 30 percent increase in price, cet. par. This, in fact, was roughly true for all years in which this hedonic equation was estimated. This coefficient (which we can interpret as an elasticity with the functional form used) was generally in the range of 2 to 3.5, and rising, in successive one-year estimation periods.

Interpreting the Hedonic Results

Figure 7 reviews different ways in which quality changes can affect an index of quality adjusted-price. For simplicity, assume a single continuous characteristic, S (think "speed"), affecting processor quality. In some base period, the upper line (containing point O) represents the reduced form relationship between price and quality combinations produced by firms and sold to consumers. A decline in quality-adjusted price, shown on the graph, is represented as a downward shift in this surface, i.e., the same quality S is now available at a lower price.

This downward shift in quality-adjusted price could be generated in a number of distinct ways. First, identical products could be produced, but sold at a lower price. This case is labeled "A" in Figure 6, a "pure" nominal price decline. Secondly, the quality of all existing products could be improved, but their price left unchanged. This is labeled "B" in Figure 6, a "pure" quality improvement. In estimating how changes like either A or B lower quality-adjusted price, we have made the assumption that the slope (i.e., the

[29] Since microprocessors alone probably account for greater than 50% of all US IC production, they must account for a much higher percentage of non-memory US IC production. I need a discussion of the PPI methodology here.

coefficient of log S) in the price-quality relationship remains constant within a one year time frame.

Note that not all product improvements necessarily affect quality-adjusted price. For example, from one month to the next, the product with characteristics O could be replaced by a higher quality product like C, with a commensurately higher price (this might occur because of shifts in demand; for example, if new software required a higher quality processor in order to run in a fashion acceptable to consumers). This improvement in quality might represent a movement along the price-quality line, not a shift of the line, and would not affect our index of quality-adjusted price.

Sources of Quality Improvement

We can actually use the estimated hedonic equation, equation (2), to decompose the sources of quality improvement over time. Taking means on both sides of (2) over all observations within some month t during our one year estimation period, and then differencing both sides of equation (2) between month t and our base period, and rearranging, we have

$$(3) \qquad \Delta\text{mean}(\ln I) = b_1 = \Delta\text{mean}(\ln P) - b_S\, \Delta\text{mean}(\ln S) - b_A\, \Delta\text{mean}(\ln A).$$

That is, the change in log of quality-adjusted price from the base period to month t is just the change in mean of log nominal price in the two periods, less the coefficient of each measured determinant of price times the change in the mean of the log of that determinant between the two periods. Figure 8 provides a simple geometric interpretation of this decomposition of quality-adjusted price change in the special case of a single determinant of price.[30]

Differencing between the first and last month during every one-year period over which regression (2) was run, we can use (3) to decompose changes in log of quality adjusted price index I into a contribution of nominal price change, and a net contribution of other factors to improved quality. We can also break the improvement of improved quality into the contribution of all distinct factors measured over the period, after calculating geometric means of these factors at the beginning and end of each period. In practice, since the measured determinants of quality change relatively rapidly from one year to the next, with technical change, I shall partition changes in log quality into the effect of processor clock rate, and the summed effect of all other measured determinants of quality.

Table 7 shows this decomposition of quality-adjusted price change into contributions of nominal price change, and contributions of quality improvement. Over three periods (1999-2000, 2001-2002, and 2004-2005, the geometric mean of the price of desktop microprocessor products Intel was offering actually increased, but was more than offset

[30] Note that the sum of residuals in every time period in equation (2) is zero, by construction, since there is a dummy variable taking on one for every time period, 0 elsewhere, and the least squares estimator produces residuals orthogonal to such dummy variables.

by quality improvements going into those higher-priced processors, and produced net declines in quality-adjusted price. In all periods other than 1995-96 (where the data is quite limited) and 2004-2005, the vast bulk of decline in quality-adjusted price came from quality improvements.

Table 8 shows the decomposition of quality change into the contribution of processor speed, and all other factors, over time. Clock speed dominates in all periods, and actually offsets negative contributions of other quality factors in 2000-2003. This makes some intuitive sense, since one characteristic of the Pentium 4 architecture was that it was optimized to provide greater performance through increased processor frequency. Intel computer architects opted to trade other aspects of performance for higher frequency clock. For given clock rate, the earlier Pentium III architecture was actually significantly faster on many tasks[31], and the later models of the P4—particularly the Prescott family architecture—were notorious for actually reducing performance compared even with the earlier Northwood P4 family, at equivalent processor clock rates.[32]

During 2003, the P4 architecture hit a "brick wall," as it became increasingly difficult to deal with increased power and heat dissipation requirements for this architecture. A new Intel emphasis on other determinants of processor performance over clock frequency became apparent with the introduction of multicore processors in 2005, and the Core 2 Duo processor architecture in 2006. Geometric mean clock rate actually fell in the second half of 2006, as the Core 2 Duo was introduced, but was more than offset by other benefits of the new architecture.

Recent discussions by Intel frankly acknowledge the shift from a strategy centered on continuing increases in processor clock speed, to an architectural approach more oriented toward increasing throughput of instructions processed at lower clock rates. (See Figure 9.) The power and heat "brick wall" shows up very clearly in Figures 10 and 11, as clock speeds cease to increase significantly after 2003. Referring back to Table 5 and Figure 4, we can see that this coincides with a rapid decline in the rate of price-performance improvement, which bottoms out in 2004-2005. Since then, with the introduction of the Core 2 Duo processor architecture, price-performance bounces back with an annual quality-adjusted price decrease of about 38%. This is just a little greater than the price-

---

[31] See Hennessy and Patterson, 2003, for a discussion of this point.

[32] The newer architecture, optimized to scale up to higher clock rates, was never actually used at these higher clock rates because, apparently, of heat problems. The deeper "pipeline" used in executing instructions in this processor (31 stages, vs. 20 in the previous model of Pentium 4, Northwood), designed to enable performance at much higher clock rates, actually slowed down performance at current clock rates for many applications. See, for example, W. Fink, "Intel 3.2E vs. 3.2EE vs. 3.2C: Comparing Baseline Performance," available at http://www.anandtech.com/cpuchipsets/showdoc.aspx?i=1965,; "Intel Pentium 4 Prescott CPU Processor," http://www.a1-electronics.net/Intel_Section/CPUs/Pentium4_Prescott_Feb04.shtml,; P. Schmid, A Roos, B. Topelt, "Intel's New Weapon: Pentium 4 Prescott," http://www.tomshardware.com/2004/02/01/intel/; V. Freeman, "Intel Pentium 4-3.4E GHz Processor Review," http://www.sharkyextreme.com/hardware/cpu/article.php/3329681; K. Schmerer, "Intel Prescott: the benchmarks," http://reviews.zdnet.co.uk/hardware/processorsmemory/0,39024015,39145079-2,00.htm; D. Mepham, "The next Pentium 4 processor, Prescott arrives," http://www.hardwareanalysis.com/content/article/1686/; all articles accessed on 1/15/2006.

performance baseline improvement I suggested was attributable to "pure" manufacturing innovation alone, however.

The Bottom Line

Our earlier analysis of manufacturing innovation suggested that, with a speedup in node introduction after 1995, with manufacturing innovation alone simply reducing transistor cost, 35% annual decline in cost/transistor would have been predicted. The balance of rates of decline in quality-adjusted price for microprocessors after 1995, which at times exceeded 70% annually, appears to be attributable to other factors, particularly architectural innovation and higher clock speeds (though the latter was in part enabled by manufacturing innovation). Due to power and heat considerations, clock speeds are no longer increasing at pre-2003 rates. This would seem to suggest that continuing architectural innovation will be needed to fuel continued declines in quality-adjusted price, if they are to significantly exceed the roughly 35% annual declines predicted with continuation of the 2-year cycle for the introduction of new lithographic technology.

To date, the industry response to these developments has been to rapidly move toward using ever-cheapening transistors to produce processors with multiple cores, operating at lower clock rates, a shift in focus toward architectural changes enabling greater numbers of instructions per clock (IPC) to be processed, and development of designs with lower power consumption. There are, however, problems visible with these new directions.

Dual and multi-core processors, unlike faster uniprocessors (employing higher clock rates), do nothing to improve performance of applications written as single threads. (They do well in improving performance when running multiple applications, running so-called "embarrassingly parallel" applications—like rendering a graphic image or processing multimedia, or running multiple instances of a single application on a server.) Rewriting other existing applications, to "parallelize," and divide work into parallel threads is difficult and expensive—this is a lesson learned repeatedly in the supercomputer industry.

The quest for for new architectural innovations is also unlikely to be an easy one. Many of the architectural innovations applied to Intel desktop processors in the 1990s actually can be traced back to ideas that were first explored in supercomputers or in research projects of the 1950s, 1960s, and 1970s. (See Figure 12, with my historical annotations paired to Intel computer architects' own evaluation of their major architectural innovations over this period.[33]) Intel started out as a semiconductor manufacturer, not a computer producer, and with the integration of more and more of a computer's design into its processors, effectively became the computer architect for the PC industry.

---

[33] Note that I have excluded several Intel innovations related to power-saving or energy efficiency—as opposed to computing performance—on my version of the Intel list, since these were first introduced in mobile or laptop processors, not desktop chips. My estimated hedonic functions actually attribute a statistically significant, positive value to one such set of innovations, measurable when EIST (enhanced Intel Speedstep) was later added to an Intel desktop processor's feature set in order to improve energy efficiency and reduce heat.

Much of the existing stock of innovation in computer design over the previous 30 years was imported into the PC world by Intel in the 1990s. The first instance in which one of the concepts in Figure 11 was *first* commercialized in an Intel microprocessor appears to have been with Simultaneous Multithreading (SMT, branded as Hyperthreading by Intel). In this case, the idea was explored in the 1990s, the first designs surfaced at the end of that decade, and Intel purchased the intellectual property[34] shortly before utilizing the concept commercially in its microprocessor products, in 2003. The incident dramatized the extent to which the existing stock of computer architectural knowledge, the fruit of research going back forty years, had been entirely imported and absorbed into the desktop PC domain during the decade of the 1990s. The question now is, once the old ideas have all been assimilated, how long will it take to explore truly new ideas, and where are these ideas going to be coming from?

Implications

All the analysis developed in this paper suggests that price/performance in desktop microprocessors is likely to be mainly driven by continued improvement in semiconductor manufacturing technology in the near future. While falling costs for manufacturing transistors on the one hand, and improvements in clock speed and architecture on the other, seem to have played roughly equal roles in driving microprocessor prices down in the period from 1995 to 2003, this no longer seems to be the case.

The pace of improvement in microprocessor price/performance currently can be almost entirely explained by continuing manufacturing innovation. Indeed, the main current trend in microprocessor design—the integration of multiple computing cores onto a single chip—is a strategy that is simply utilizing ever cheaper transistors in the most direct possible way, by replicating a processor design as many times as are possible, and economic, on a single chip. Unless this approach can be scaled up indefinitely, which knowledgeable experts do not believe to be the case,[35] even this formula for utilizing cheaper transistors would seem to be headed for difficulties in the medium term.

This also means that sustaining continued innovation in semiconductor manufacturing is even more important to the health of the semiconductor industry, and its downstream users in the wider economy, than was the case in the previous decade. Current efforts to coordinate R&D through the ITRS, and the variety of public and private consortia that are attempting to maintain the pace of technical progress in semiconductor manufacturing, must be relatively more important than they were prior to 2003.

---

[34] When it bought out now-defunct Digital Equipment's semiconductor operations from Compaq Computer.

[35] Because of bottlenecks in transferring data in and out of the processor, communication bottlenecks and coordination overheads among multiple computing cores, and the difficulties of programming multithreaded applications that utilize large numbers of processors effectively. For a thorough and technically competent discussion, see National Research Council, Computer Science and Technology Board, **Getting Up to Speed: The Future of Supercomputing**, (Washington: The National Academies), 2005.

Today, the problems of design innovation in desktop microprocessors seem to be converging with the problems of supercomputing. The problem of writing software that utilizes large numbers of processors or cores efficiently in working on a single problem, an issue on the frontiers of cutting edge R&D in supercomputers, has now become the new challenge for desktop computers. Truly new ideas for computer and software design are the domain of leading edge R&D in computer systems. Increased investments in high performance computing R&D are ultimately likely to be needed to fuel the continuing "spillover" of economic benefits to IT users—and the broader economy.

Recent research has shown that rapid improvement in price performance for processors and memory are a major source for high rates of improvement in PC price-performance. Rapid improvement in PC price-performance, in turn, has stimulated widespread take-up of IT across many sectors of the economy, which in turn has stimulated productivity improvements in these sectors.[36]

Thus, slower improvement in PC price-performance is likely to reduce incentives to purchase new computers. A slowdown in purchases of PCs, and the application of IT, could well have significant ripple effects on the global economy.

Finally, it is possible that my reading of the historical record of quality-adjusted price decline for Intel desktop processors is too narrow a slice of a bigger landscape, and biases me toward excessive pessimism. It is possible that data for AMD processors may tell a different story than Intel data, and may even provide a persuasive explanation for recent gains in market share by AMD. Also, the PC market has been shifting toward laptops and mobile computing, rather than desktops, and it is possible that laptop processors may tell an entirely different story than desktops.[37] Lastly, it is entirely plausible that PCs are no longer at the cutting edge—that today it is the Internet and communications to which the economic center of gravity for IT innovation has shifted.

[36] The most influential portrait of this nexus is Jorgenson, 2001.

[37] Future results from this research effort will compare results for Intel desktop processors with those manufactured by AMD, as well as desktop with mobile microprocessors. It is interesting to speculate that differences in quality-adjusted price declines may explain some of the recent shifts from desktop toward mobile computing, and possibly, small gains in market share by AMD in the 2004-05 period.

Table 1
Rates of Decline in Quality-adjusted Price for Semiconductors, 1991-1999
Compound Annual Decline Rates (%)

	CAGR 91-95	CAGR 95-99	CAGR 91-99
MOS MPU	-40.36	-61.89	-52.3
MOS Memory	-8.02	-47.87	-30.8
of which, DRAM	-7.76	-53.46	-34.5
MOS MPR	-3.89	-23.01	-14.0
Other MOS Logic	-6.76	-19.13	-13.2
Thyristors & Rectifiers	-0.84	-12.94	-7.1
MOS MCU	0.36	-13.87	-7.0
Power Transistors	-0.78	-10.27	-5.6
Small Signal Transistors	0.26	-10.50	-5.3
Optoelectronics	3.25	-10.04	-3.6
Diode & All Other Discrete	4.28	-9.03	-2.6
Digital Bipolar	5.37	-4.01	0.6

Source: Author's calculations based on data in Aizcorbe, Flamm, and Khurshid (2006).

Table 2

Density Scenarios

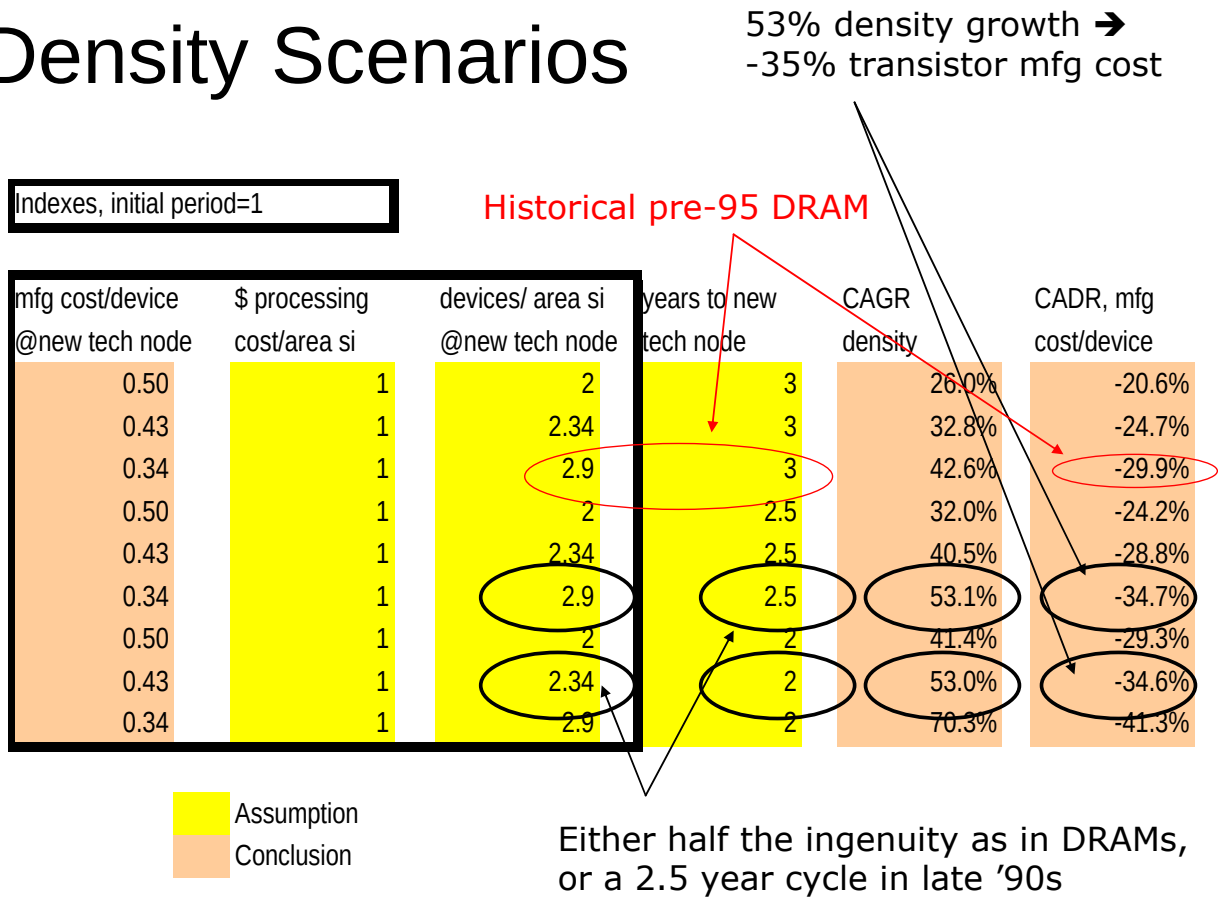

Table 3

Summary of Intel Desktop Processor Sample

Number of Observations, by Design Architecture Family and Selected Characteristics

		L2 Cache								Bus Speed								Architecture Family							date in sample	
Architecture Number:	Name	0	128	256	512	1000	2000	4000	8000	60	66	100	133	400	533	800	1066	cel	Pent	P2	P3	P4	P4D	Core2	beg	end
1	Cedar Mill	0	0	0	17	0	28	0	0	0	0	0	0	0	17	28	0	17	0	0	0	28	0	0	200601	200612
2	Conroe	0	0	0	0	0	8	12	0	0	0	0	0	0	0	0	20	0	0	0	0	0	0	20	200607	200612
3	Coppermine	0	0	579	0	0	0	0	0	0	0	205	374	0	0	0	0	0	0	0	579	0	0	0	199910	200209
4	Coppermine-128	0	207	0	0	0	0	0	0	0	102	105	0	0	0	0	0	207	0	0	0	0	0	0	200003	200202
5	Covington	17	0	0	0	0	0	0	0	0	17	0	0	0	0	0	0	17	0	0	0	0	0	0	199804	199903
6	Deschutes	0	0	0	157	0	0	0	0	0	35	122	0	0	0	0	0	0	0	157	0	0	0	0	199705	200001
7	Gallatin	0	0	0	47	0	0	0	0	0	0	0	0	0	0	41	6	0	0	0	0	47	0	0	200402	200506
8	Katmai	0	0	0	62	0	0	0	0	0	0	56	6	0	0	0	0	0	0	0	62	0	0	0	199902	200005
9	Kentsfield	0	0	0	0	0	0	0	2	0	0	0	0	0	0	0	2	0	0	0	0	0	0	2	200611	200612
10	Klamath	0	0	0	8	0	0	0	0	0	8	0	0	0	0	0	0	0	0	8	0	0	0	0	199705	199808
11	Mendocino	0	99	0	0	0	0	0	0	0	99	0	0	0	0	0	0	99	0	0	0	0	0	0	199808	200010
12	Northwood	0	0	0	207	0	0	0	0	0	0	0	0	38	94	75	0	0	0	0	0	207	0	0	200201	200503
13	Northwood-128	0	142	0	0	0	0	0	0	0	0	0	0	142	0	0	0	142	0	0	0	0	0	0	200211	200506
14	P54CS	39	0	0	0	0	0	0	0	10	29	0	0	0	0	0	0	0	39	0	0	0	0	0	199506	199802
15	P55C	26	0	0	0	0	0	0	0	0	26	0	0	0	0	0	0	0	26	0	0	0	0	0	199611	199808
16	Prescott	0	0	252	0	338	0	0	0	0	0	0	0	0	267	323	0	252	0	0	0	338	0	0	200402	200612
17	Prescott 2M	0	0	0	0	0	86	0	0	0	0	0	0	0	0	77	9	0	0	0	0	86	0	0	200503	200612
18	Presler	0	0	0	0	0	0	44	0	0	0	0	0	0	0	36	8	0	0	0	0	0	44	0	200601	200612
19	Smithfield	0	0	0	0	0	45	0	0	0	0	0	0	0	2	43	0	0	0	0	0	0	45	0	200505	200612
20	Tualatin	0	0	75	0	0	0	0	0	0	0	41	34	0	0	0	0	41	0	0	34	0	0	0	200108	200209
21	Willamette	0	11	256	0	0	0	0	0	0	0	0	0	267	0	0	0	11	0	0	0	256	0	0	200011	200302
	Total	82	459	1,162	498	338	167	56	2	10	316	529	414	447	380	623	45	786	65	165	675	962	89	22		

Table 4
A Hedonic Price Index for Desktop Intel Processors

	Index,2/2004=100		
1995 Jun	734,807.39	2001 Jan	1,175.22
1996 Jan	345,947.39	2001 Mar	1,056.31
1996 Feb	270,857.70	2001 Apr	897.52
1996 May	210,467.95	2001 May	786.81
1996 Jun	180,776.20	2001 Jun	778.97
1996 Aug	169,797.14	2001 Jul	794.18
1996 Nov	159,196.50	2001 Aug	662.77
1997 Feb	124,980.15	2001 Sep	613.56
1997 May	98,236.21	2001 Oct	583.15
1997 Jun	107,342.75	2001 Dec	538.55
1997 Aug	59,413.15	2002 Jan	540.24
1997 Nov	49,176.86	2002 Feb	540.24
1998 Feb	38,487.46	2002 Mar	452.77
1998 Mar	37,347.53	2002 Apr	453.60
1998 Apr	28,997.51	2002 May	423.86
1998 May	28,997.51	2002 Jun	371.65
1998 Jun	23,906.10	2002 Sep	371.65
1998 Jul	19,048.24	2002 Nov	192.45
1998 Aug	18,877.98	2003 Jan	191.44
1998 Sep	16,280.53	2003 Feb	170.32
1998 Oct	13,904.33	2003 Apr	156.14
1998 Nov	13,500.49	2003 Jul	134.74
1998 Dec	12,124.24	2003 Aug	131.27
1999 Jan	12,907.72	2003 Oct	119.63
1999 Feb	10,279.51	2004 Feb	100.00
1999 Mar	10,446.84	2004 Apr	99.50
1999 Apr	8,781.54	2004 May	99.50
1999 May	7,622.31	2004 Jun	99.93
1999 Jun	7,277.27	2004 Aug	87.68
1999 Jul	7,069.10	2004 Oct	86.63
1999 Aug	5,652.06	2004 Dec	86.23
1999 Sep	5,109.80	2005 Jan	86.23
1999 Oct	4,911.82	2005 Feb	82.67
1999 Dec	4,857.91	2005 Mar	82.67
2000 Jan	4,612.63	2005 May	82.62
2000 Feb	3,252.52	2005 Jun	82.47
2000 Mar	3,188.11	2005 Jul	81.25
2000 Apr	2,724.91	2005 Aug	77.27
2000 May	2,280.57	2005 Sep	77.27
2000 Jun	2,221.07	2005 Dec	76.56
2000 Jul	2,049.60	2006 Jan	76.12
2000 Aug	1,744.91	2006 Apr	61.13
2000 Oct	1,541.80	2006 Jun	60.22
2000 Nov	1,374.24	2006 Jul	50.00
2000 Dec	1,338.44	2006 Oct	48.10
		2006 Nov	47.38
		2006 Dec	47.38

Source: Author's calculations, described in text.

Table 5
Comparison of Price Indexes

Compound Annual Decline Rates, Intel Desktop Microprocessors

Aizcorbe, Corrado, Dom Fisher Ideal Matched Model Price Index				
Q293-Q294	-28.27%			BLS
Q294-Q295	-57.39%		Flamm Hedonic Index	Microprocessors, microcontrollers, other non-memory ICs
Q295-Q296	-66.22%			
Q296-Q297	-48.54%	May96-May97	-53.32%	"Sliding substitutes"
Q297-Q298	-71.82%	May97-May98	-70.48%	Hybrid Index
Q298-Q299	-68.06%	May98-May99	-73.71%	
		May99-May00	-70.08%	-65.60%
		May00-May01	-65.50%	-55.91%
		May01-May02	-46.13%	-43.05%
		May02-Apr03	-66.36%	-43.91%
		Apr03-May04	-34.03%	-46.91%
		May04-May05	-16.96%	-20.62%
		May05-Jun06	-25.32%	-30.11%
		Jun06-Dec06	-38.10%	-33.79%

Table 6

Hedonic Regression Equation, 2004-05

Source	SS	df	MS		
Model	251.944728	23	10.9541186	Number of obs	428
Residual	6.06193693	404	.015004794	F(23, 404)	730.04
Total	258.006665	427	.604231066	Prob > F	0.0000
				R-squared	0.9765
				Adj R-squared	0.9752
				Root MSE	.12249

lp	Coef.	Std. Err.	t	P>|t|	[95% Conf.	Interval]
lproc	3.14283	.0812105	38.70	0.000	2.983182	3.302478
a7b1066	-.0378972	.0667829	-0.57	0.571	-.1691826	.0933883
a12b800	-.0096687	.0365394	-0.26	0.791	-.0814998	.0621624
a16b800	.1951877	.0379609	5.14	0.000	.1205622	.2698131
a17b1066	.6594852	.0973689	6.77	0.000	.4680723	.8508982
arch7	1.367857	.0481989	28.38	0.000	1.273105	1.462609
arch13	-.1968557	.03806	-5.17	0.000	-.2716762	-.1220353
arch16	-.8698693	.1211131	-7.18	0.000	-1.10796	-.6317786
arch17	-.3842368	.1403317	-2.74	0.006	-.6601084	-.1083652
arch19	.7571755	.1813871	4.17	0.000	.4005951	1.113756
hvolt	-5.575525	1.071103	-5.21	0.000	-7.681156	-3.469894
lvolt	1.455206	.6286528	2.31	0.021	.2193663	2.691045
ht	.2820691	.0358078	7.88	0.000	.2116761	.352462
lga775	-.003009	.0232709	-0.13	0.897	-.0487561	.0427381
exdisab	.0115304	.0236996	0.49	0.627	-.0350595	.0581203
d200406	.0043476	.0248819	0.17	0.861	-.0445666	.0532618
d200408	-.1264262	.0282396	-4.48	0.000	-.1819412	-.0709112
d200410	-.1384226	.0242471	-5.71	0.000	-.1860887	-.0907564
d200412	-.1431436	.027764	-5.16	0.000	-.1977234	-.0885637
d200501	-.1431436	.027764	-5.16	0.000	-.1977234	-.0885637
d200502	-.1852775	.0281581	-6.58	0.000	-.2406322	-.1299228
d200503	-.1853107	.0278123	-6.66	0.000	-.2399856	-.1306357
d200505	-.1858412	.0292746	-6.35	0.000	-.2433907	-.1282917
_cons	-13.55221	1.251253	-10.83	0.000	-16.01199	-11.09242

Table 7

Share of Change In Quality-Adjusted Price Due to

Period	Change in mean log of Price	Change in mean log of Quality
6/1995-5/1996	66.6%	33.4%
5/1996-5/1997	6.2%	93.8%
5/1997-5/1998	22.4%	77.6%
5/1998-5/1999	33.4%	66.6%
5/1999-5/2000	-19.0%	119.0%
5/2000-5/2001	48.1%	51.9%
5/2001-5/2002	-37.6%	137.6%
5/2002-4/2003	-3.5%	103.5%
4/2003-5/2004	8.7%	91.3%
5/2004-5/2005	-90.0%	190.0%
5/2005-6/2006	75.2%	24.8%
6/2006-12/2006	9.7%	90.3%

Note: Negative numbers correspond to increase in geometric mean of price for available models.

Table 8

	Share of Change In Quality Due to	
	Change in mean log of chip clock freq	Change due to all other sources
Period		
6/1995-5/1996	101.5%	-1.5%
5/1996-5/1997	70.9%	29.1%
5/1997-5/1998	54.9%	45.1%
5/1998-5/1999	129.4%	-29.4%
5/1999-5/2000	100.4%	-0.4%
5/2000-5/2001	116.7%	-16.7%
5/2001-5/2002	170.0%	-70.0%
5/2002-4/2003	139.0%	-39.0%
4/2003-5/2004	65.8%	34.2%
5/2004-5/2005	71.1%	28.9%
5/2005-6/2006	112.9%	-12.9%
6/2006-12/2006	-49.9%	149.9%

Note: Negative numbers correspond to decline in geometric mean of clock frequency for available models.

Figure 1
Wafer Processing Costs for Leading Edge Logic

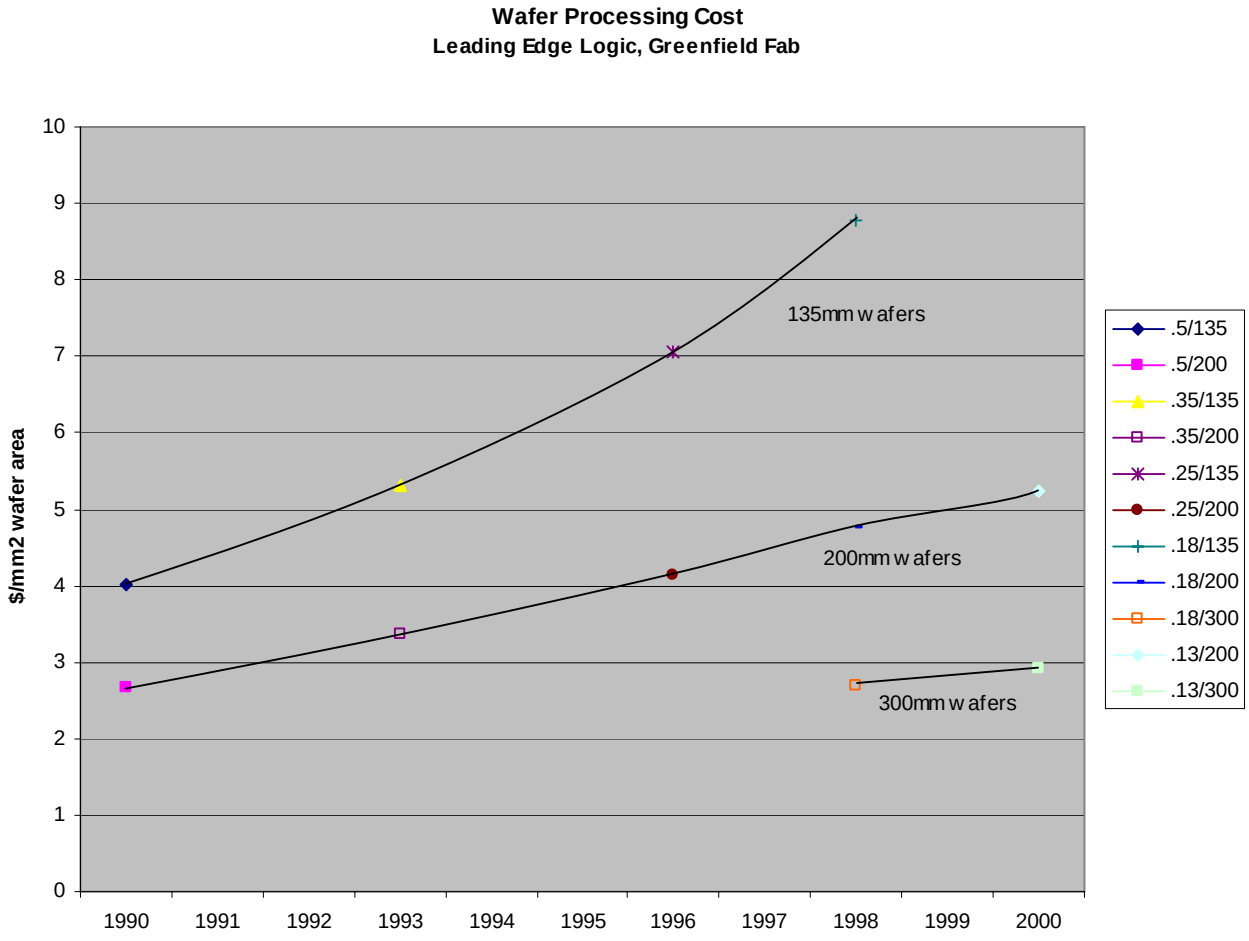

Figure 2

Costs per Area of Wafers Fabricated at TSMC and UMC

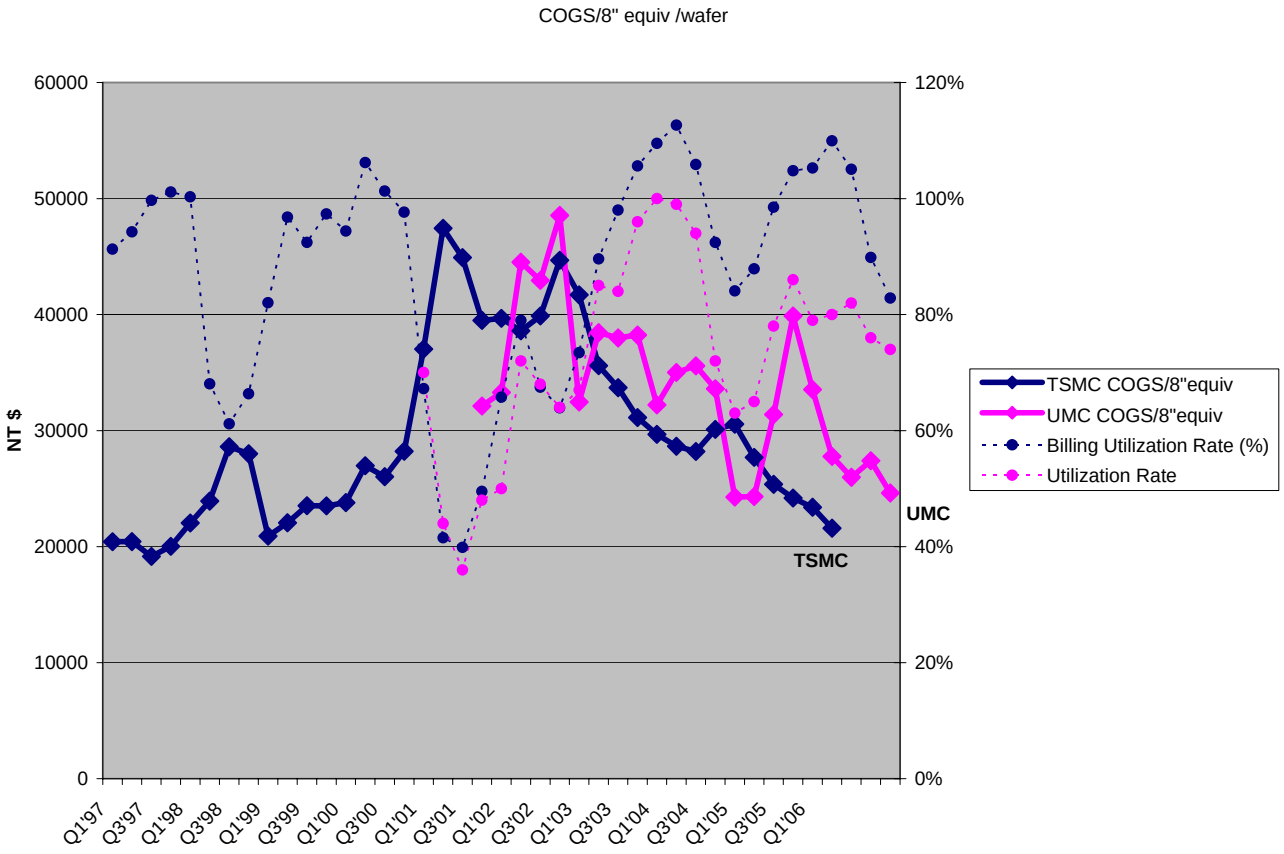

Figure 3
Pentium Microprocessor Transistor Density

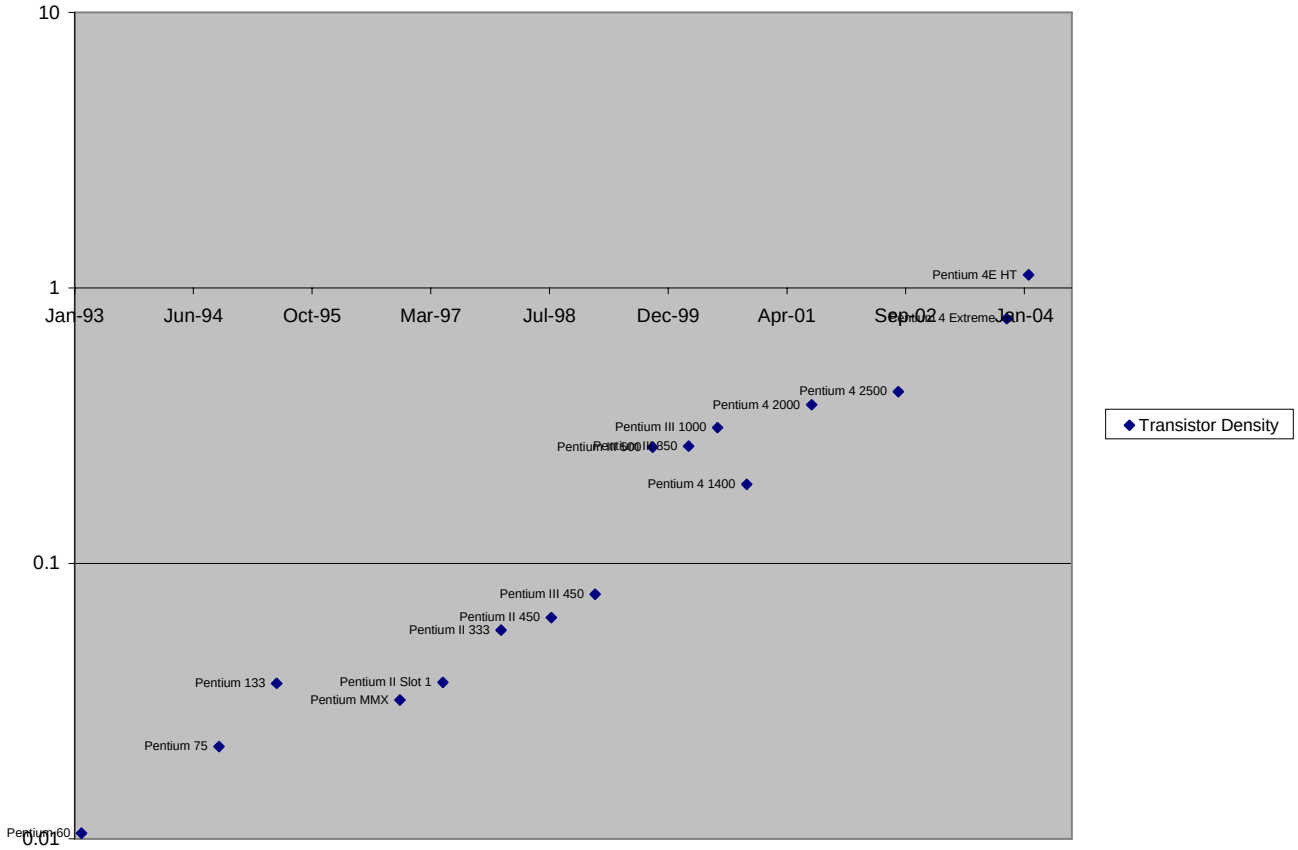

Figure 4
Hedonic Price Index for Intel Desktop Processors

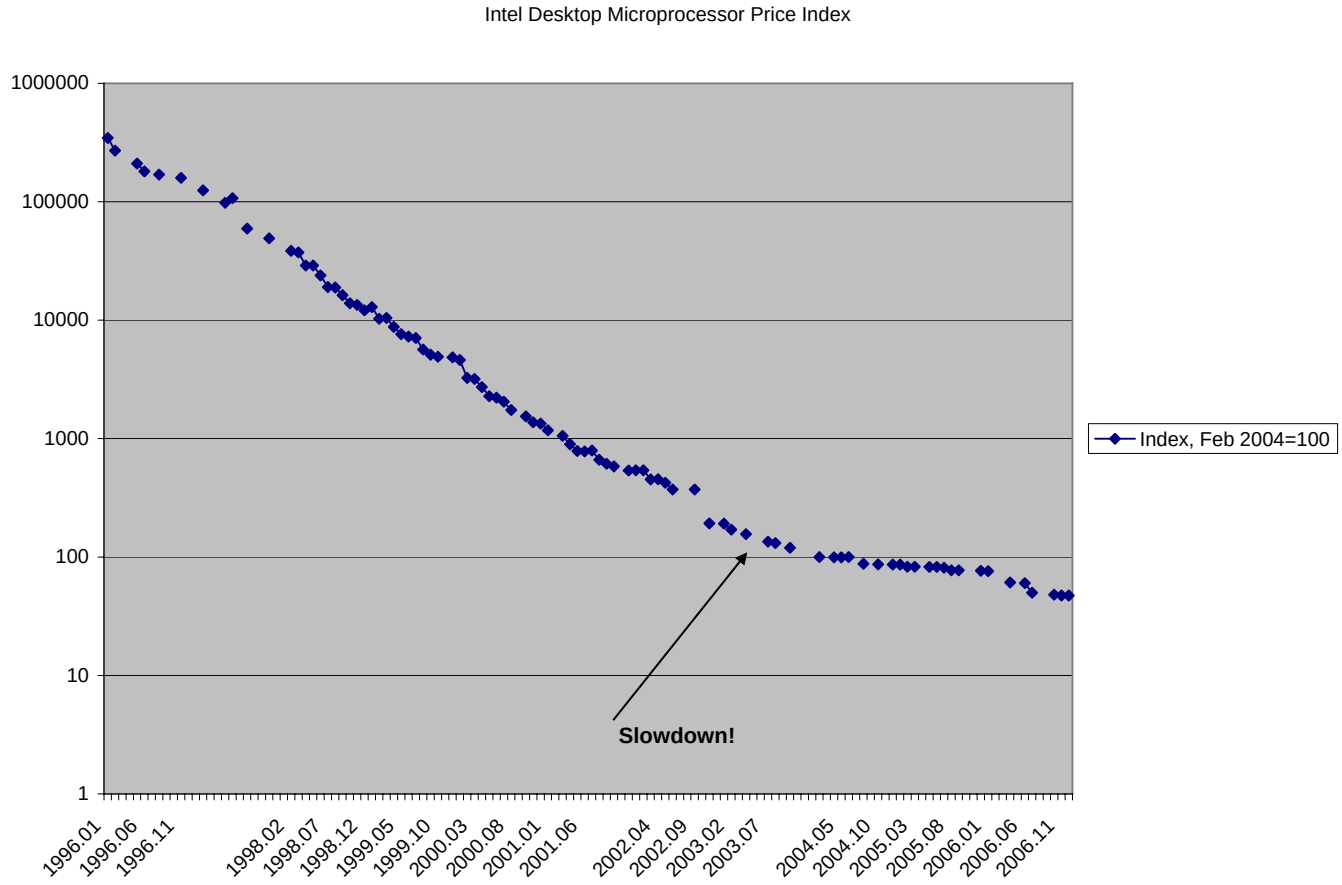

Figure 5
Monthly Price Decline Rate

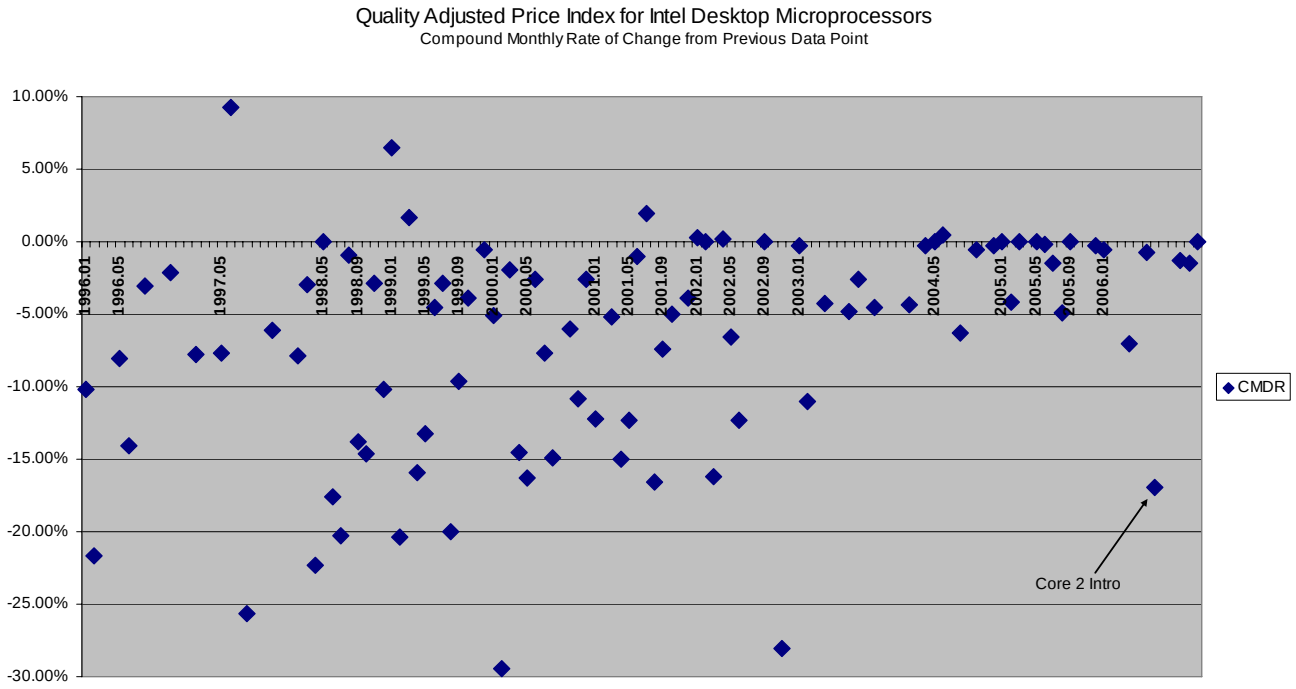

Figure 6
Slowdown in Computer Performance

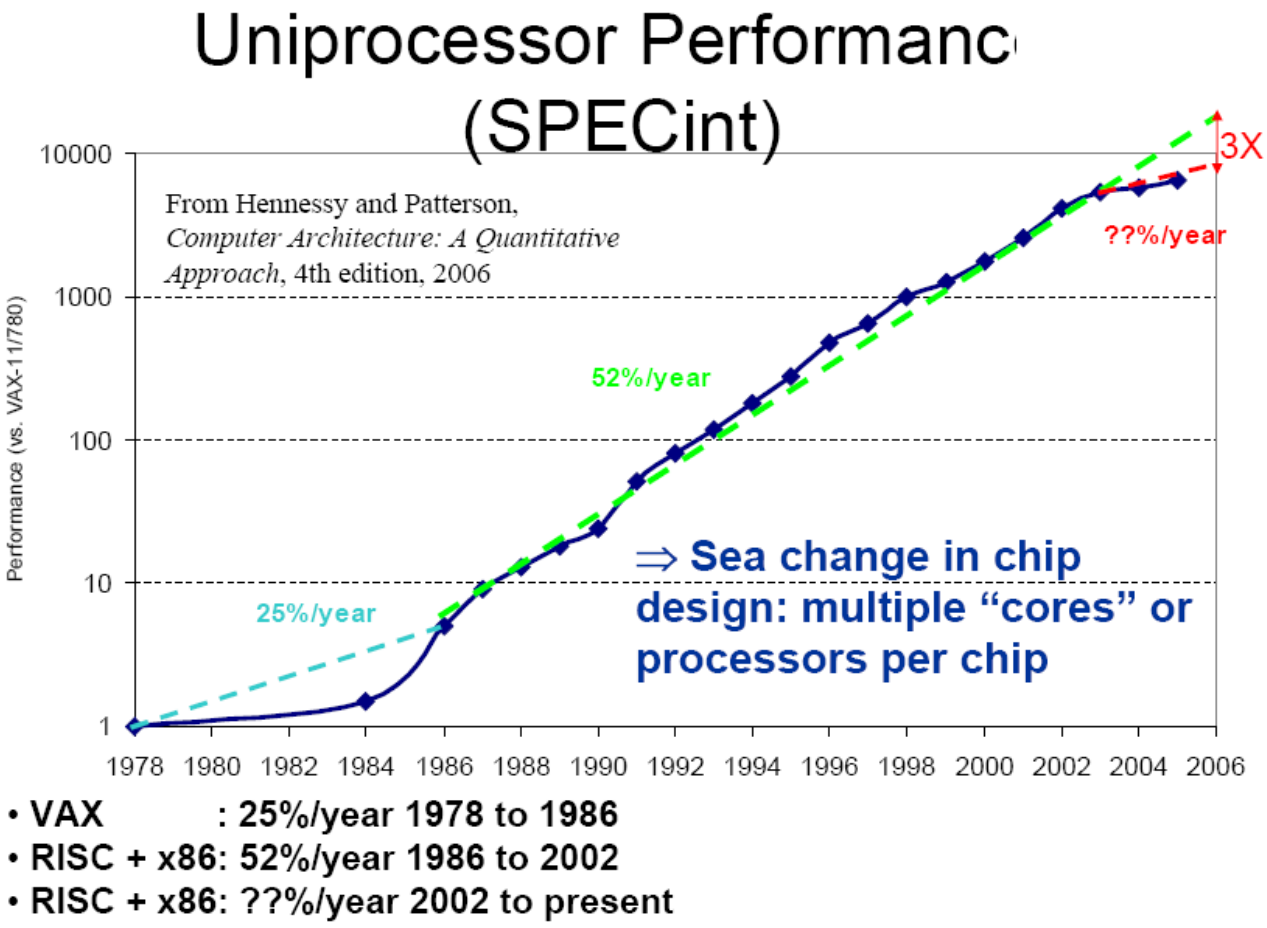

(Slide from presentation by D. Patterson, Feb. 2006)

Figure 7
Hedonic Price Index Concept

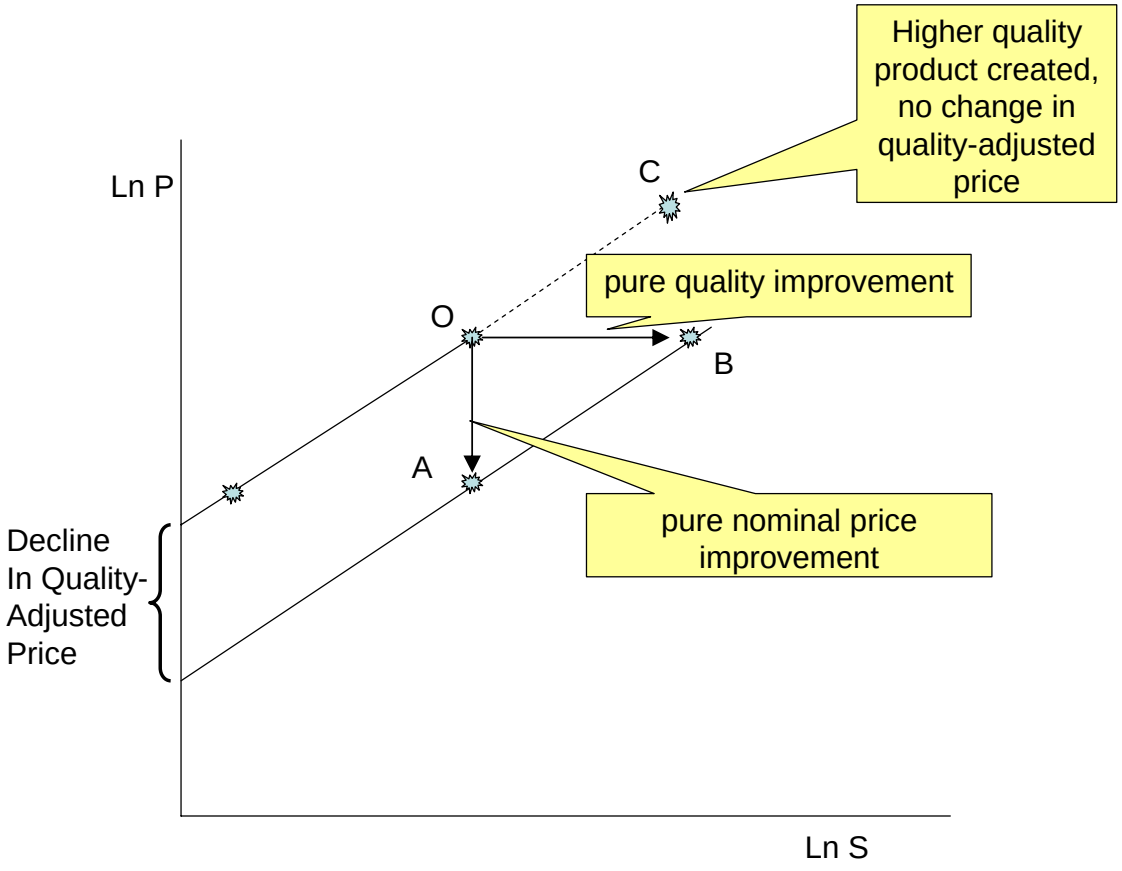

Figure 8

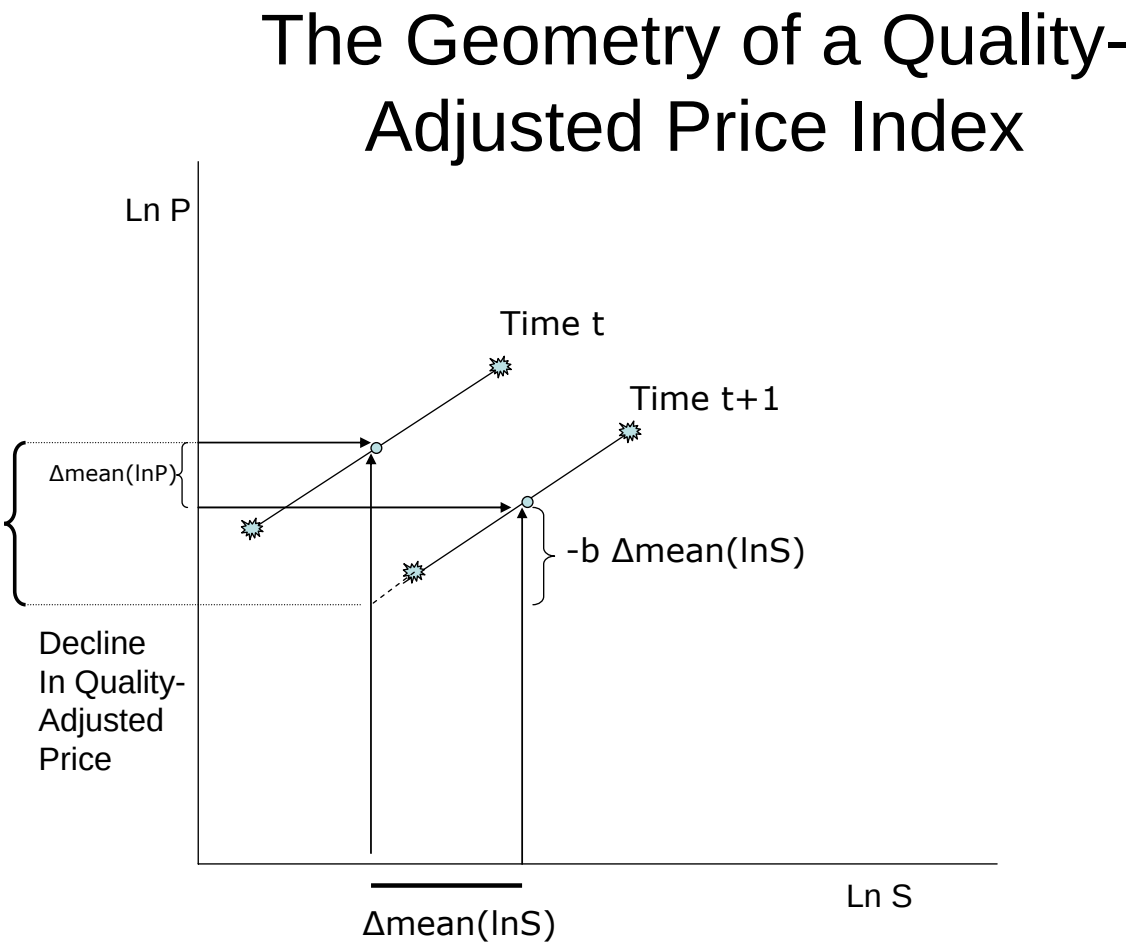
The Geometry of a Quality-Adjusted Price Index
Ln P
Time t
Time t+1
Δmean(lnP)
-b Δmean(lnS)
Decline
In Quality-
Adjusted
Price
Ln S
Δmean(lnS)

Figure 9

Intel Computer Architect View of Intel Processor Evolution, 2006

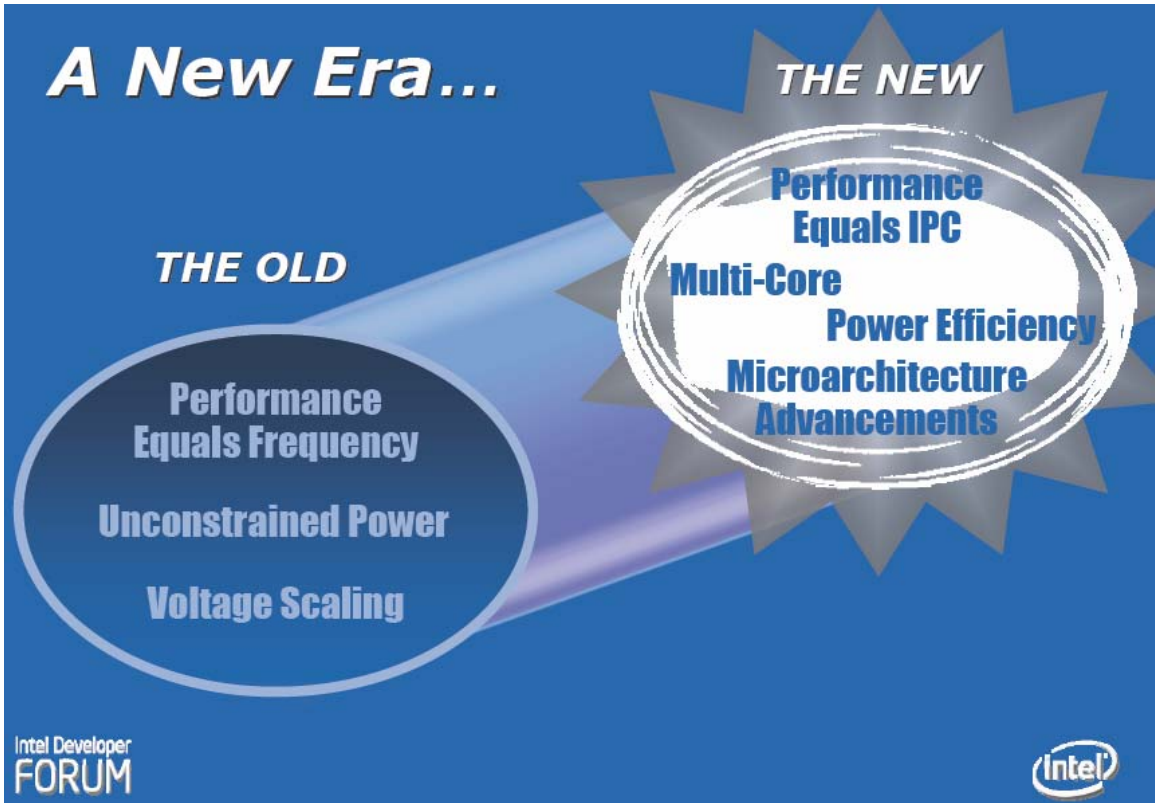

Slide from presentation by S. Pawlowski and O. Wechsler, Intel Developers Forum, 2/2006.

Figure 10
Geometric Mean of Intel Desktop Processor Speed, by Technology Node

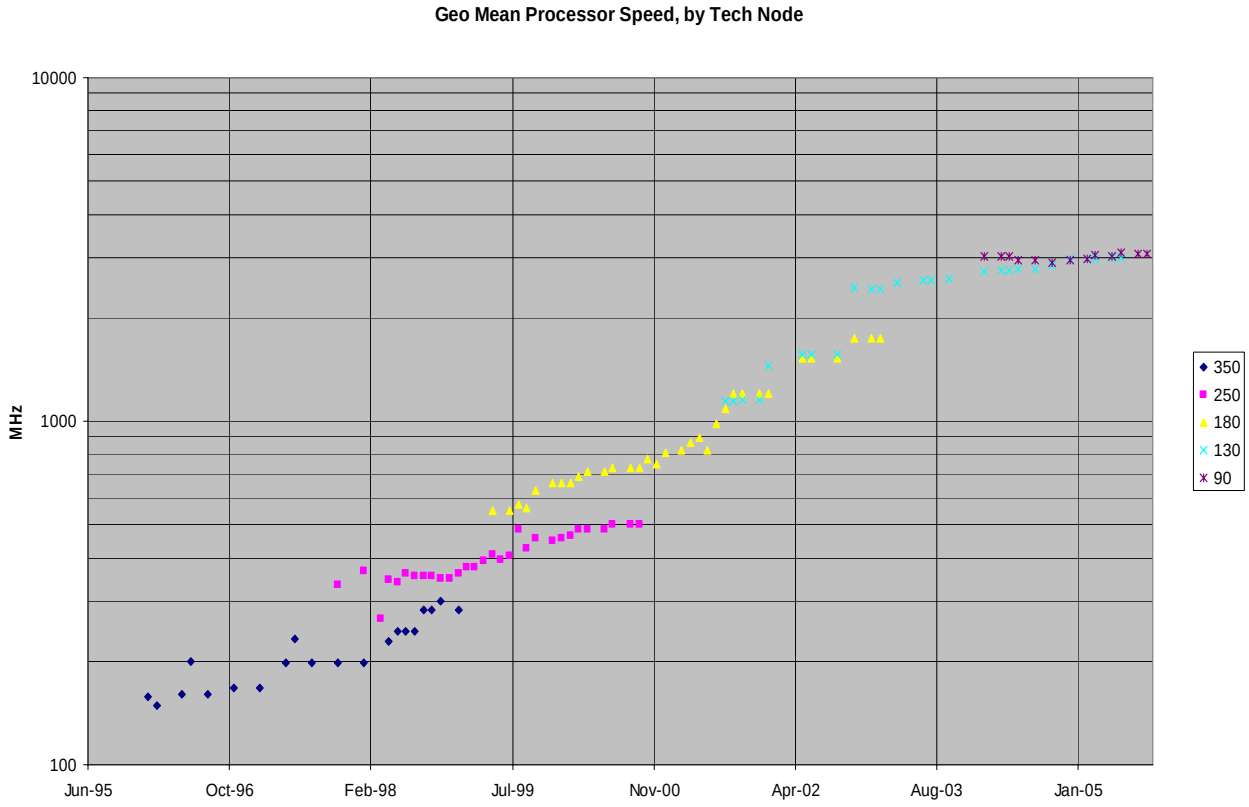

Figure 11
Geometric Mean of Intel Desktop Processor Speed, All Technology Nodes

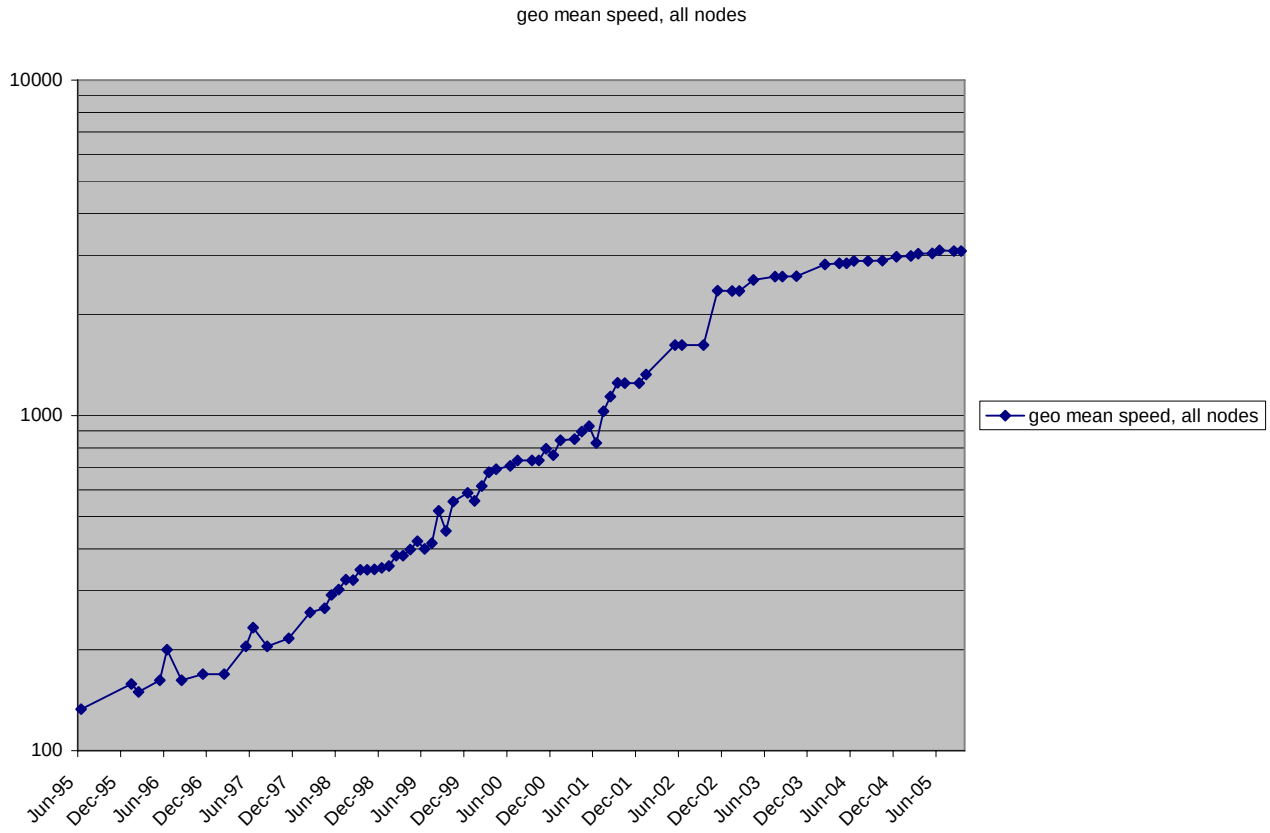

Figure 12

First Origins of Architectural Innovations in Intel Desktop Processors

- i486 (1989)
  - Integrated Level 1 cache — IBM 360/85 (1969)
  - Integrated FPU — EDVAC (1951); IBM 7030 (Stretch, 1961) multiple FPUs
  - Pipelining — IBM 7030 (1961)
- Pentium (1993)
  - Branch prediction — IBM 7030 (1961); IBM 360/91 (1969)
  - Superscalar — IBM ACS (concept, late '60's); Astromatics ZS-1(1988); IBM RS/6000 (1990)
- Pentium MMX (1995)
  - Integer vector processor — TI ASC(1972);CDC Star-100 (1972); Cray 1 (vector registers, 1976)
- Pentium II (1997)
  - Register renaming — IBM 360/91(1969)
  - Out of order execution — CDC 6600 (1964);IBM 360/91 (1969)
  - Speculative execution — IBM 360/91(1969)
- Pentium III (1999)
  - SSE (floating point vector processor) — TI ASC(1972);CDC Star-100 (1972); Cray 1 (vector registers, 1976)
- Pentium 4
  - SSE2 (2000) — TI ASC(1972);CDC Star-100 (1972); Cray 1 (vector registers, 1976)
  - Hyperthreading (Simultaneous Multithreading, SMT) (2003) — Hardware multithreading: MIT TX-2 (1959); CDC 6600 PPU (1964); HEP (1982); SMT: DEC Alpha EV-21464 (2003, design, never produced)

Sources: S. Pawlowski and O. Wechsler, "Intel Core Microarchitecture," presented at Intel Developer's Forum, Spring, 2006, available at ww.intel.com/pressroom/kits/core2duo/pdf/icm_tech_overview.pdf; Flamm, 1988; Hennessey and Patterson, 2003; email communications with Prof. Bill Dally, 2004; Prof. C. Kozvrakis, "Lecture 16: SMT & CMP," EE382A Handout, Winter 2006, Stanford University; "Simultaneous multithreading resources," available at www.princeton.edu/~jdonal/research/hyperthreading/. Andrew Orlowski, "Project Jackson- why SMT is the joker in the chip pack," **The Register**, Feb. 25, 2001; Mark Hachman, "Update: Compaq Licenses Alpha to Intel," **ExtremeTech**, June 2001.