

**Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.**

**Shifting Trends in Semiconductor Prices and the Pace
of Technological Progress**

Ana Aizcorbe, Stephen D. Oliner, and Daniel E. Sichel

2006-44

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

Shifting Trends in Semiconductor Prices and the Pace of Technological Progress

Ana Aizcorbe, Stephen D. Oliner, and Daniel E. Sichel*

Original version: July 2003
Current version: September 2006

Abstract

This paper examines three questions motivated by previous research on semiconductors and productivity growth: Why did semiconductor prices fall so rapidly in the second half of the 1990s, why has the rate of price decline slowed since 2001, and to what extent are these price swings associated with changes in the rate of advance in semiconductor technology? We show that the price swings are statistically significant and that they reflect changes in both price-cost markups and cost trends. Further analysis indicates that the shift to faster cost declines in the mid-1990s likely corresponded to a speed-up in the pace of advance in semiconductor technology; however, the slower cost declines since 2001 appear not to have been mirrored by a deceleration in technology. Consequently, researchers should be cautious about associating price or cost movements for semiconductors with changes in the pace of underlying technology even over moderately long periods.

*Aizcorbe is affiliated with the Bureau of Economic Analysis; Oliner and Sichel are affiliated with the Federal Reserve Board. Darrell Ashton, Blake Bailey, Ryan Bledsoe, and Tom McAndrew provided excellent research assistance. We thank Ernie Berndt, Iain Cockburn, Mark Doms, John Fernald, Kevin Fox, Dale Jorgenson, Glenn Rudebusch, Eric Swanson, and Dan Wilson for valuable comments on earlier versions of the paper. We also thank David Byrne and Charles Gilbert for providing data on semiconductor prices and shipments. Gene Amromin and Nellie Liang kindly provided estimates of the value of stock-option grants for Intel Corporation and Micron Technology Inc. We are indebted as well to Alan Allan and Mung Chen of Intel Corporation, Robert Doering of Texas Instruments Inc., and Dan Hutcheson of VLSI Research Inc. for sharing their expertise on semiconductor technology trends and to Mike Scherer for a helpful discussion on competition and innovation in the semiconductor industry. The views expressed are those of the authors alone and should not be attributed to the Bureau of Economic Analysis, the Board of Governors of the Federal Reserve System, or other members of the staff of these organizations.

1. INTRODUCTION

The U.S. economy expanded at a rapid pace in the second half of the 1990s, spurred by a resurgence in labor productivity growth. Considerable research has highlighted a central role for information technology (IT) in that resurgence, reflecting the enormous improvements in price-performance ratios for IT capital goods and, more fundamentally, for the semiconductors that power this capital.¹ In recent years, however, semiconductor prices have fallen less rapidly than in the second half of the 1990s, tempering the price declines for IT capital goods and likely contributing to the more restrained spending on these goods. Given the key role of semiconductors in these developments, three questions demand attention: Why did semiconductor prices fall so rapidly in the second half of the 1990s? Why has the rate of price decline slowed more recently? To what extent are these price swings associated with changes in the rate of advance in semiconductor technology?

Several studies have examined the faster rate of price decline in the second half of the 1990s. Jorgenson (2001) linked the steeper price declines to a shift from three-year to two-year technology cycles in the semiconductor industry. Flamm (2004) and Congressional Budget Office (2002) accepted the view that technology cycles had become shorter, but they asserted that technology alone could not fully explain the more rapid price declines in the late 1990s. Similarly, McKinsey Global Institute (2001) and Aizcorbe (2005, 2006) argued that the speed-up in the rate of price decline reflected, at least in part, heightened competition between Intel and its chief rival, AMD.

We build on this research in three ways. First, we examine not only the shift to faster declines in constant-quality semiconductor prices in the mid-1990s but also the reversion to

¹ See, for example, Oliner and Sichel (2000a and 2002), Jorgenson and Stiroh (2000), and Jorgenson, Ho, and Stiroh (2002).

slower declines since 2001. Little research has focused on the latter period.

Second, we formally test for breaks in the rate of price decline using a state-of-the-art framework that allows us to search for multiple unknown breakpoints. Such statistical evidence is important because it determines whether observed swings in price trends represent more than just random variation and, hence, are worthy of further study; as far as we know, our paper is the first to apply any econometric analysis to this question. We focus on prices for a broad aggregate of semiconductors and for two key types of chips: microprocessor units (MPUs) and dynamic random access memory chips (DRAMs). Using this machinery, we find compelling statistical evidence of breaks in MPU price trends in 1994 and 2001 and reasonably strong evidence of similar breaks for DRAM prices. For the series on overall semiconductor prices, the results are weaker but nonetheless point to a break in 1995.

Third, to examine the source of these breaks, we decompose changes in constant-quality DRAM and MPU prices into changes in price per transistor and a residual that captures quality change beyond the number of transistors per chip; we then decompose price per transistor into a price-cost markup and cost per transistor. We implement this decomposition for MPU chips using data on Intel's revenues, costs, and chip output; we implement the same decomposition for DRAM chips using data for Micron Technology, the largest DRAM producer in the United States.

For both DRAM and MPU chips, our framework shows that changes in price-cost markups contributed notably to the observed shifts in the rate of price decline. The role for markups is especially pronounced for DRAM chips, which are subject to wide swings in worldwide supply-demand conditions. The decomposition also shows that cost trends have varied over time. For DRAM chips, the downtrend in cost per transistor sped up in the mid-

1990s and then slowed around 2001. For MPUs, the story concerning cost reductions is more nuanced than for DRAM chips, as it reflects the combined influence of cost per transistor and improvements in chip quality beyond the number of transistors (the residual term in the decomposition). Still, the basic message is similar to that for DRAMs: MPU cost trends, as measured in our decomposition, were unusually favorable from the mid-1990s through 2001, and less so thereafter.

Our decomposition provides the starting point from which to answer a crucial question: To what extent do semiconductor price dynamics reflect fundamental shifts in the pace of advance in semiconductor technology? This issue is important because researchers (see Oliner and Sichel, 2000a and 2002, and Whelan, 2002, for example) have often presumed a tight linkage between prices and technology when measuring productivity growth in the semiconductor industry and other high-tech sectors. However, this may not be an appropriate assumption. Basu, Fernald, Fisher, and Kimball (2005) discuss a number of reasons why relative price trends could be a poor proxy for technological progress, including time-varying markups and increasing returns to scale. In a similar vein, Feenstra, Reinsdorf, Slaughter, and Harper (2005) point to changes in terms of trade as a possible wedge between IT price declines and the pace of technological progress. These studies highlight some important issues, but they do not directly address the questions we are asking about semiconductors.

Zeroing in on semiconductors, our decomposition reveals wide swings in price-cost markups for DRAMs and MPUs, as noted above. These markups have moved with well-documented shifts in the balance of supply and demand that are largely unrelated to the underlying pace of technological progress, undercutting the linkage between prices and technology. After controlling for swings in markups, can one regard the cost shifts that remain

as providing a strong signal about technology trends? This need not be the case because scale effects and changes in product mix could influence measured cost trends even if technology improved at a constant rate. To investigate this issue, the final part of the paper looks directly at technology trends in the industry, drawing on the International Technology Roadmap for Semiconductors (which presents the consensus judgment of industry participants), other indicators of technology trends, and discussions with several industry experts.²

This industry analysis provides strong evidence that the technology cycle for semiconductors became faster during the 1990s. The timing of this shift cannot be determined precisely – a plausible range of dates runs from 1993 to 1998 – but the case for its existence is convincing. Thus, a speed-up in the pace of technological advance appears to have contributed to the more favorable cost trends in the second half of the 1990s.

Regarding the more recent period, the connection between cost and technology trends is far more tenuous. The Roadmap indicates that, as of 2004, technology cycles remained on the faster track that had prevailed since the late 1990s, a conclusion confirmed by a variety of other indicators. Accordingly, the slower decline in cost per transistor over 2001-2004 cannot be ascribed to an adverse shift in technology cycles and, instead, appears to reflect some combination of scale effects and changes in product mix. All in all, our analysis suggests that researchers should be very cautious about associating price or cost movements for semiconductors with changes in the pace of underlying technology even over moderately long periods.

² Our analysis does not emphasize Moore's Law, an often-cited indicator that refers to the period over which the number of transistors doubles in leading-edge chips (Moore, 1965). We focus instead on a more fundamental driver of advances in semiconductor technology, the shrinkage in the size of the smallest component on a chip. This shrinkage is what enables manufacturers to pack more transistors on a chip at an acceptable cost of production. See Flamm (2004) for an interesting discussion of Moore's Law.

The rest of the paper is organized as follows. Section 2 presents the basic facts about semiconductor price trends. Section 3 carries out the econometric tests for structural breaks in these series, and section 4 implements our decomposition of DRAM and MPU prices. Section 5 analyzes the connection between prices and costs on the one hand and technological advance on the other. Section 6 briefly reviews our conclusions.

2. BASIC FACTS ABOUT SEMICONDUCTOR PRICES

We analyze three constant-quality indexes of semiconductor prices: the aggregate index for integrated circuits used by Oliner and Sichel (2000a and 2002) and separate price series for MPUs and DRAMs. The aggregate price index is an annual series extending from 1975 to 2004, while the MPU and DRAM indexes are both quarterly, covering 1987:Q1-2004:Q4 and 1975:Q1-2004:Q4, respectively. For the period from 1992 forward, these series all rely on monthly price data used by the Federal Reserve Board to construct its index of industrial production. Other price series for MPUs and DRAMs are then spliced in to cover the earlier periods.³ The data appendix provides further information about these price series.

We use the aggregate Oliner-Sichel series because it is known in the literature and is similar to series used by other researchers, including Jorgenson and Stiroh (2000), Jorgenson (2001), and Jorgenson, Ho, and Stiroh (2002). This aggregate series includes many types of integrated circuits. We focus on DRAM and MPU chips for the following reasons. MPUs represented nearly half of the dollar value of integrated circuits shipped by U.S. producers during 1992-2004, and they accounted for about three-quarters of the decline in the overall semiconductor price index over this period. Thus, MPUs are central to understanding the dynamics of semiconductor prices. DRAMs, in contrast, account for a much smaller share of the

³ For MPUs and DRAMs, the series used for the earlier periods are available for several years after 1992. We confirmed that these series moved in sync with the Federal Reserve series during the overlap period.

total semiconductor market.⁴ However, DRAMs historically have been the pace-setting chip for key technological advances in semiconductors. In addition, the data on DRAM prices extend back to the mid-1970s; having such a long series is useful in a search for structural breaks.

Figure 1 displays the aggregate index of semiconductor prices and the separate series for DRAM and MPU prices. As can be seen in the upper panel, the aggregate price index has fallen dramatically over the past three decades, with an especially rapid decline during the second half of the 1990s. Table 1 shows that this price index dropped at an average rate of 22.5 percent per year over 1975-94 and at roughly double that pace over 1994-2001, before reverting to an average rate of about 28 percent over 2001-04.

The lower panel of figure 1 displays the separate time series on MPU and DRAM prices, with the average rates of change again presented in table 1. Consistent with the aggregate price index, the drop in MPU prices accelerated around 1994, with the rate of decline averaging 63 percent annually over 1994-2001, a marked speed-up from the average 30 percent rate over 1988-1994.⁵ Since 2001, however, the rate of decline has slowed from the extraordinary 1994-2001 pace. The DRAM series is choppier than that for MPUs, which makes it more difficult to discern changes in trends. Nonetheless, the average rate of decline over 1994-2001, at about 48 percent annually, was substantially faster than the roughly 28 percent pace over the preceding two decades. The rate of decline, however, has slowed noticeably since 2001. Thus, all three price series recorded especially fast declines over the second half of the 1990s, with signs of moderation more recently.

⁴ From 1992 to 2004, DRAMs made up 10.6 percent of nominal shipments of integrated circuits from U.S. producers and accounted for about 6 percent of the decline in overall semiconductor prices. These figures and those cited above for MPUs were calculated by Federal Reserve staff based on source data from the Semiconductor Industry Association and the Census Bureau.

⁵ Although the MPU price data are available back to 1987:Q1, the table shows the percent change over 1988-1994 to match the time period used for our price decomposition in section 4.

These swings in semiconductor prices have had a noticeable effect on the prices of computing equipment, which use semiconductors as a key input. Figure 2 displays the price index for computers and peripheral equipment in the National Income and Product Accounts along with the aggregate index of semiconductor prices. Both series are plotted as rolling percent changes over three-year periods to make the underlying trends more apparent. As shown, both price series fell especially rapidly in the late 1990s, and both have since reverted to a pace of decline more characteristic of the period before the mid-1990s. The swings are more pronounced for semiconductor prices than for computer prices, as would be expected given that semiconductors represent only a portion of the production cost for computers. Nonetheless, semiconductor prices clearly influence the prices of computing equipment and thus indirectly affect the pace of business investment in IT capital and the growth of productivity throughout the economy.

3. IDENTIFYING STRUCTURAL BREAKS IN SEMICONDUCTOR PRICES

To the best of our knowledge, previous discussions of changing trends in semiconductor prices have been based on casual observation. Although pictures such as figure 1 can be very instructive, statistical tests allow one to determine whether the variation is more than random.⁶

Tests for Structural Change

For the case of a single breakpoint of unknown timing, many tests – beginning with Quandt (1960) – have been proposed for identifying the most likely breakpoint. However, figure 1 suggests that the semiconductor price series might contain more than one breakpoint. To account for the possibility of multiple breaks, we use tests proposed by Bai and Perron (1998, 2003, and 2006) to search for multiple breakpoints in a unified statistical framework. (From this point on, we will refer to Bai and Perron as BP.) BP recommend a multi-stage test procedure.

⁶ See Hansen (2001) for an overview of the literature on tests for structural change.

The first stage of the procedure tests the null of no breaks versus the alternative of an unknown number of breaks given an upper bound on the number of possible breakpoints. If the first stage suggests the existence of a break, the subsequent stages identify the date of the first break and test for two breaks versus one break, three breaks versus two breaks, and so on.

Because semiconductor prices have a strong downward trend, we pre-tested (the log of) these series for unit roots using Dickey-Fuller tests and tests proposed by Banerjee, Lumsdaine, and Stock (1992). These tests uniformly failed to reject the null hypothesis of a unit root in the semiconductor price series at the 5 percent level. Therefore, we will conduct our break tests on the log differences of prices, and the starting point for the tests is a regression of the form:

$$\Delta y_t = \alpha_0 + \alpha_1 I_{t,k} + \alpha_2 \Delta y_{t-1} + \varepsilon_t \quad (1)$$

$$I_{t,k} = 0 \text{ if } t \leq k$$

$$I_{t,k} = 1 \text{ if } t > k,$$

where Δy represents the log difference of y . This regression can be run for every possible breakpoint k , with the indicator variable equal to zero for all periods prior to the breakpoint and unity for all periods after the breakpoint. For each breakpoint, the coefficient on the indicator variable, α_1 , measures the amount by which the average growth rate of y differs between the first and second subsamples.

To conduct the first stage of BP's procedure, we use the unweighted version of the double maximum test, the so-called UDMAX test. For this test, we first estimate equation 1 at each possible breakpoint and calculate the chi-squared statistic described by BP. The date associated with the maximum value of this chi-squared statistic across all possible breaks is used to divide the sample into two subsamples. Given this proposed break, we then roll through the first subsample to identify the date associated with the maximum value of the chi-squared

statistic in that subsample and repeat this procedure in the second subsample. The maximum value of the chi-squared statistic from the first subsample is compared to that from the second subsample. The date associated with the larger of the two is used to divide the subsample that contains that date into two further subsamples. This procedure is repeated until one reaches the previously selected upper bound on the number of possible breakpoints. For our implementation, we assumed that the maximum number of breakpoints is two. Figure 1 strongly suggests that the price series have no more than two breaks, and in any case, a test for three breaks would have had limited power given the relatively small number of observations in the price series.

The largest chi-squared statistic from all of the steps is the UDMAX test statistic. This statistic is compared to critical values in BP (2003) to evaluate the null of no breaks against the alternative of at least one break. If the UDMAX test fails to reject the null of structural stability, the test procedure is finished and there is no significant evidence of a break in the series. Conversely, if the UDMAX test rejects the null, the next step of the multi-stage procedure is implemented using the $supF_T(l+1/l)$ test.

The $supF_T(l+1/l)$ test evaluates the null of l breaks versus the alternative of $l+1$ breaks and identifies the date of potential breaks.⁷ To identify the date of the first breakpoint, we estimate equation 1 at each possible single breakpoint and select the date that minimizes the residual sum of squares. Using this date to split the sample, we then estimate equation 1 over the first subsample with a break at each possible date and do the same for the second subsample. The date from the two subsamples that minimizes the residual sum of squares identifies the most likely date of a second break. The $supF_T(2/1)$ test statistic is then calculated and compared to the

⁷ In principle, the $supF_T$ test could have been used in the first stage to test for the existence of any breaks. However, BP (2006) indicate that, for power reasons, it is preferable to use the UDMAX test in the first stage.

appropriate critical value from BP (2003) to determine whether the second breakpoint is significant.

Results

Table 2 summarizes the results for the aggregate index of semiconductor prices, DRAM prices, and MPU prices. For the index of aggregate semiconductor prices, the UDMAX test in the first stage is significant at the 10 percent level. The most likely date of a first break in this annual series is 1995, and the $supF_T(2/1)$ test provides no evidence of a second break. For the quarterly series on DRAM prices, the evidence of breaks is stronger than for the aggregate price series. The UDMAX test is significant at the 5 percent level, and the most likely date of a first break is 1995:Q4. There also is evidence of a second break in 2001:Q4 based on the $supF_T(2/1)$ test, but that evidence is significant only at the 10 percent level. For the quarterly series on MPU prices, the evidence of two breaks is even stronger. The UDMAX test provides evidence of a significant break at the 5 percent level. The most likely date of the first break is 1994:Q2, and the $supF_T(2/1)$ provides evidence of a second break at the 5 percent level in 2001:Q4.

In sum, there is compelling evidence of two breaks in MPU prices, one in 1994 and one in 2001. For DRAM prices, the tests indicate a break in 1995 and provide some evidence of a second break in 2001. For overall semiconductor prices, the test results are weaker and point to only a single break in 1995.⁸

What accounts for this variation? The stronger evidence of breaks in MPU prices than in DRAM prices may reflect, at least in part, the greater volatility of DRAM prices. As we showed in figure 1, DRAM prices have not fallen in a smooth fashion but rather have oscillated in

⁸ To assess the robustness of these tests, we considered variations in the sample periods, the minimum number of data points required for subsamples, and the lag length on the Newey-West covariance matrices. Although the test statistics were somewhat sensitive to these variations, the results remained significant at the levels reported in table 2.

periodic cycles around a declining trend. These cycles line up closely with industry accounts of global supply and demand imbalances in the DRAM market.⁹ For example, DRAM prices fell rapidly from late 1995 through mid-1998, which followed a period of large increases in DRAM production capacity. Then, from mid-1998 through late 2000, DRAM prices held nearly steady, supported by a consolidation in the industry and strong demand for computing equipment. These market dynamics may well have made it more difficult for the statistical tests to identify breaks in the underlying downtrend in prices.

As for overall semiconductor prices, this series contains a wide variety of chips other than DRAMs and MPUs that are subject to quite different market and technological forces. Evidently, the price behavior of these chips differs enough from that of DRAMs and MPUs to partially obscure the structural breaks that are evident in the DRAM and MPU price series.

4. DECOMPOSITION OF PRICE CHANGE

Given the evidence of breaks in DRAM and MPU price trends, we now develop a decomposition – in the spirit of Flamm (2004) – to explore the sources of these changing trends.

Framework

The number of transistors on a chip is a key determinant of its quality, and we build up our decomposition of constant-quality prices from the following expression for price per transistor:

$$price/transistor = (price/cost)(cost/transistor) \quad (2)$$

In words, price per transistor can be decomposed into a price-cost markup and the average cost per transistor. To convert equation 2 into an expression for the rate of change in price per transistor, let p denote price per transistor, m denote the price-cost markup, and c denote the cost

⁹ See, for instance, Standard & Poor's (2005). The annual reports of the major semiconductor producers also provide useful information about industry developments.

per transistor. In addition, let x^* denote the value of x in a subsequent time period, where $x = p, m, \text{ or } c$. With this notation, equation 2 implies

$$p^* / p = (m^* / m)(c^* / c)$$

or

$$1 + \dot{p} = (1 + \dot{m})(1 + \dot{c}), \quad (3)$$

where the dot above a variable signifies the percent change over a given period.¹⁰

As indicated, our ultimate interest is in decomposing the changes in constant-quality prices rather than the price per transistor. Constant-quality prices, denoted by p_{cq} , can be expressed as the price per transistor, divided by an index of quality improvements not captured by the number of transistors per chip (q): $p_{cq} = p/q$, or $p = p_{cq}q$, which implies that

$$1 + \dot{p} = (1 + \dot{p}_{cq})(1 + \dot{q}) = 1 + \dot{p}_{cq} + \dot{q}(1 + \dot{p}_{cq}). \quad (4)$$

Now, combine equations 3 and 4 and rearrange terms to yield

$$\dot{p}_{cq} = \left[(1 + \dot{m})(1 + \dot{c}) - 1 \right] + R, \quad (5)$$

where $R = -\dot{q}(1 + \dot{p}_{cq})$.

Equation 5 is our decomposition of the percent change in constant-quality price indexes for DRAM and MPU chips. The terms in brackets explain the percent change in price per transistor, based on the contributions from changes in the price-cost markup and cost per transistor. This expression accounts for the cross products between the individual terms, which would be excluded (incorrectly) if we were to simply sum up the percent changes in m and c .

¹⁰ Although using log differences would have simplified the arithmetic of the decomposition, we chose to use percent changes because log differences are a poor approximation for rates of change in rapidly-changing series, like semiconductor prices.

The remaining term, R , represents the difference between the percent change in the constant-quality price index and the percent change in the price per transistor. Any improvement in chip quality over and above increases in the number of transistors per chip is unobserved and enters our decomposition as a residual. The residual term captures both the amount of unobserved quality improvement and the value that purchasers place on that functionality. The residual also will impound any measurement error in the other terms in the equation.

Empirical Implementation of the Decomposition

We use equation 5 to shed light on why semiconductor prices fell so rapidly from the mid-1990s through 2001 and why these price declines have moderated since 2001. The decomposition begins in the earliest year for which all necessary data are available (1988 for MPUs and 1990 for DRAMs). We then select the break year in the mid-1990s based on the results of the structural break tests reported in section 3. Recall that these tests identified a break in 1994 for MPU prices and one in 1995 for DRAM prices.

The decomposition of MPU prices relies on data for Intel, the dominant producer of MPU chips, while the DRAM decomposition employs data for Micron Technology, the only major DRAM producer in the United States and one of the largest firms in this market worldwide.

Terms in the Decomposition

Price-Cost Markup. The first term on the right side of equation 5 is the rate of change in the price-cost markup, which captures the cyclical swings in market conditions and longer-term changes in market structure. To measure the markup, we start with data from Intel's and Micron's annual financial statements on the ratio of profits to sales revenue. We then translate this profit margin, denoted by π , into the implied markup of price over average cost, $m \equiv p/c$. Letting Q denote the total number of transistors in the chips sold by either Intel or Micron,

$$\pi \equiv (p - c)Q/pQ = (p - c)/p = 1 - (1/m), \text{ which implies that } m = 1/(1 - \pi).^{11}$$

Our figures for π are based on operating income, rather than the bottom-line measure of profits reported on financial statements, net income after tax. We prefer operating income for two reasons. First, reported net income includes the effects of infrequent charges (“special items” in accounting parlance) that can distort the underlying pattern of earnings over time. Second, net income includes the earnings generated by activities outside the firm’s core line of business. For example, Intel maintains an active program of equity investments in other companies, with the aim of nurturing ventures that have the potential to spur demand for its products. The gains from sales of these securities generated more than one-third of Intel’s reported net income in 2000. Operating income excludes both special items and the gains (or losses) from financial activities and thus provides a cleaner measure of trends in earnings related to the production of semiconductors.

We make one adjustment to operating income. During our sample period, neither Intel nor Micron Technology recorded the value of stock options granted to employees as a labor cost on their income statements. The accounting rules in place through 2004 only required firms to disclose information about stock options in the footnotes to their financial reports. As a result, the companies’ reported operating income overstates their true earnings, sometimes by a wide

¹¹ Note that we calculate the markup over average cost rather than over marginal cost. Although standard models of firm behavior (Hall, 1988, for example) generate markups defined in terms of marginal cost, such cost data are difficult to obtain. Moreover, the optimization problem for semiconductor firms is complex, entailing decisions along several margins that range from the short-term choice of variable inputs given existing plants and technology to the long-horizon allocation of resources to research and development activities. There is no fully articulated model of the choices along the various margins, though Aizcorbe and Kortum (2005) is a step in this direction. In the absence of specific guidance from theory, we use the markup over average cost as a rough indicator of market conditions. As a check on these results, section 5 presents a cost series for Intel that is closer to short-run marginal cost and discusses the behavior of the markup implied by this series.

margin. To compute a more accurate measure of earnings, we subtract an estimate of the value of stock option grants from the firms' reported operating income.¹²

Figure 3 shows the annual price-cost markups for Intel and Micron. As indicated by the solid line, Intel's markup has experienced two cyclical declines over the past twenty years. The first decline was in 1985-86, when Japanese producers flooded the market with DRAM chips, which caused prices to dive and prompted Intel (and most other U.S. firms) to exit the DRAM market. The second downturn reflects the post-2000 meltdown in the technology sector. Focusing on the period between these two episodes, Intel's markup trended up from the late 1980s through the mid-1990s (with some year-to-year variation) and then was relatively stable during the second half of the 1990s.

The dashed line in figure 3 presents the analogous markup for Micron Technology. Micron's markup exhibits much greater variation than Intel's and is somewhat lower on average.¹³ These differences reflect the fact that DRAM production is a highly competitive business subject to periodic imbalances in global supply and demand.

Cost per Transistor. The second term on the right side of equation 5 is the rate of change in the average cost per transistor. The numerator of this series is the measure of operating cost for Intel and Micron that we used to calculate the companies' markups. We then divide this measure of operating costs by the total number of transistors in the MPU chips shipped by Intel and the DRAM chips shipped by Micron.¹⁴

¹² In December 2004, the Financial Accounting Standards Board issued a new rule (Statement of Financial Accounting Standards No. 123(R), "Share-Based Payment") that requires public companies to include the value of stock-option grants as a labor cost on their income statements for fiscal years that begin after June 15, 2005. Thus, starting in fiscal 2006, it will no longer be necessary to adjust reported income for the value of option grants.

¹³ Between 1995 and 2001, Micron derived a substantial part of its total revenue from the production of personal computers. For these years, we strip out its PC subsidiary when calculating the profit margins and the markup.

¹⁴ The data appendix provides additional detail on the definition and source for each series used in the decomposition.

Decomposition of DRAM Prices

Table 3 shows our decomposition of DRAM prices based on equation 5. Line 1 of the table displays the percent changes in constant-quality DRAM prices that we seek to explain with the decomposition. Line 2 shows the percent change in price per transistor – the term in brackets in equation 5 and the core of the decomposition. Lines 3, 4, and 5 show the contributions to the percent change in price per transistor from the price-cost markup, cost per transistor, and the cross-product between these factors, respectively. Finally, line 6 shows the residual term R , which captures both quality change beyond increases in the number of transistors per chip and measurement error in the other terms of the decomposition.

As shown in the table, constant-quality DRAM prices fell at an average annual pace of about 15 percent from 1990 to 1995. Prices were driven down by the technological improvements that enabled Micron to reduce its cost per transistor at an average rate of nearly 25 percent per year. This cost reduction was partially offset by an increase in Micron's markup. On balance, Micron's price per transistor fell at about a 13½ percent annual pace over 1990-95. This drop in price per transistor was quite close to the rate of decline in constant-quality DRAM prices, leaving only a small residual in the decomposition.

The constant-quality DRAM price index and the price per transistor both fell much more rapidly during 1995-2001 than they had in the earlier period. These faster price declines reflect two factors. First, cost per transistor fell at nearly double its rate during 1990-95. In addition, the markup moved down, reversing the rise in the earlier period and illustrating the boom-bust cycles in the DRAM market. This decline in the markup, along with a sharp drop in Micron's cost per transistor, fully accounts for the nearly 54 percent annual rate of decline in constant-quality DRAM prices over 1995-2001.

During 2001-2004, these prices fell much more slowly than they had over 1995-2001, returning to roughly the pace of decline observed before 1995. In contrast to the two earlier periods, the decline in price per transistor over 2001-2004 was more rapid than the drop in constant-quality prices, with the difference showing up as the residual on line 6. If taken as a measure of chip quality, this difference suggests that, apart from the continued increase in transistors per chip, the quality of DRAM chips *worsened* in the recent period. This seems implausible because the number of transistors per DRAM chip is proportional to the key performance characteristic of these chips – their storage capacity – and thus serves as a rough summary statistic for DRAM chip quality. More likely, the residual is capturing measurement error of some sort in other parts of the decomposition.

Columns 4 and 5 show that the decomposition accounts reasonably well for the swings in DRAM price changes across the three periods. As shown in column 4, the decomposition of price per transistor fully explains the shift to faster declines in constant-quality prices during 1995-2001. A sharp negative swing in the price-cost markup accounts for the bulk of this shift. In addition, cost per transistor fell more rapidly during 1995-2001 than before, consistent with a move to shorter technology cycles in the second half of the 1990s.¹⁵ With regard to the more recent period, column 5 shows that our decomposition accounts for nearly all of the reversion to slower declines in constant-quality prices. A big part of the story is that the markup recovered from its depressed level in 2001. But, over and above this cyclical effect, prices fell less rapidly during 2001-2004 because cost trends were not as favorable as they had been during 1995-2001.

¹⁵ Our results support Flamm's (2004) conclusion that the move from three-year to two-year technology cycles, by itself, cannot fully explain the steep declines in DRAM prices in the second half of the 1990s.

Decomposition of MPU Prices

Table 4 presents our decomposition of MPU prices. As shown on lines 1 and 2, both constant-quality MPU prices and price per transistor fell at an average annual pace of 30 percent during 1988-1994. Our decomposition attributes the entire decline in these price measures to a steep downtrend in cost per transistor. During 1994-2001, the decline in constant-quality MPU prices sped up dramatically, similar to the pattern for DRAMs. However, as can be seen in column 4, our decomposition of price per transistor captures only about a quarter of this pickup; the rest of the swing shows up in the residual.

Although the bulge in the residual during 1994-2001 may owe partly to measurement error, we would argue that it likely also reflects a combination of improvements in the quality of MPU chips beyond the number of transistors per chip and an increase in the willingness of purchasers to pay for that additional functionality. The performance of an MPU chip depends heavily on its architecture and, unlike DRAM chips, need not be well summarized by the number of transistors on the chip. During the second half of the 1990s, an important set of innovations improved the architecture of MPU chips. These innovations included the introduction of cache memory (a temporary storage area for frequently accessed data) and the adoption of new technologies that allowed faster bus speeds (the rate at which the MPU communicates with other key components). Regarding purchasers' willingness to pay for this enhanced quality, the rapid spread of the Internet and other graphics-intensive applications, along with the introduction of more complex operating systems, likely boosted the demand for faster processors in the second half of the 1990s. Quantifying the contribution of these factors is difficult and beyond the scope

of this paper, but Chwelos (2003) provides some support for the interpretation of the residual described here.¹⁶

Turning to the more recent period, constant-quality prices of MPUs declined about 40 percent per year during 2001-2004, well below the extraordinary 63 percent annual rate observed during 1994-2001. As can be seen in column 5, our decomposition of price per transistor can account for virtually the entire reversal: Intel's price-cost markup moved up substantially from its low in 2001, and its cost per transistor declined less rapidly over 2001-2004 than it had over 1994-2001. Although the markup is the more significant of the two factors over this set of years, the relative importance of the two factors is sensitive to the year chosen to separate the period of rapid price declines in the 1990s from the more recent period. In particular, if we split the two periods at 2000 rather than 2001, the positive swing in the markup shrinks to 5 percentage points, and the unfavorable shift in cost per transistor becomes the main factor accounting for the slowdown in MPU price declines. This slowdown reflects some combination of a rising markup and less rapid cost reduction, but our analysis cannot assign precise weights to each factor.

Figure 4 displays the decomposition graphically. For DRAMs in the upper panel and MPUs in the lower panel, the solid black lines show the change in constant-quality semiconductor prices, and the shaded bars show the pieces of the decomposition of these price changes. The upper panel highlights the important role of swings in markups for DRAM prices. The lower panel shows the significant role of the "other" term for MPU prices, which includes

¹⁶ As for the improvement in quality beyond the number of transistors per chip, Chwelos shows that an MPU performance measure based on benchmark tests for laptops increased much more rapidly during 1994-98 than it did during 1990-94. Although Chwelos does not report analogous results for transistors per chip, our data (which cover a broader set of MPU chips than Chwelos' sample) show no acceleration in transistors per chip between 1990-94 and 1994-98. A change in purchasers' willingness to pay for functionality would appear in Chwelos' hedonic regression as a structural break in the coefficients on performance characteristics. Chwelos does not formally test for such breaks, but the reported standard errors suggest that the coefficient on the benchmark performance measure was significantly higher in 1997-98 than in earlier years.

both the residual term in our decomposition and the cross product term (which is quite small). For the reasons discussed above, we believe that the residual term in the MPU decomposition largely reflects improvements in quality over and above those captured by the number of transistors per chip.

5. HOW TIGHT IS THE LINK BETWEEN SEMICONDUCTOR PRICES AND TECHNOLOGICAL PROGRESS?

As indicated above, some analysts have used constant-quality prices to gauge the pace (and changes in the pace) of technological progress in the semiconductor industry and other high-tech sectors. However, our decomposition of constant-quality prices highlights that swings in margins are quite significant; in addition, these swings can be linked to well-documented shifts in the balance of supply and demand that are largely unrelated to the underlying pace of technical progress. Thus, one would want to strip out margins before drawing any inferences from prices about technology.¹⁷ Moreover, our measure of cost per transistor may be affected by factors that are not directly tied to technology. In particular, shifts in scale economies over time and swings in product mix toward and away from more costly chips could affect our measure of cost per transistor.

In a perfect world, we would combine detailed data on prices and quantities of individual chips with a comprehensive economic model of the semiconductor industry to account for these factors. Then, having done so, we would apply our structural break machinery to the appropriate measures to search for breaks in the pace of technical progress. Unfortunately, we do not have sufficiently detailed data to undertake this exercise, and developing a comprehensive model of the industry is beyond the scope of this paper. Thus, we turn to other evidence – more in the

¹⁷ To be clear, we are not criticizing the use of changes in relative prices to infer trends in multifactor productivity via the “dual” approach to production analysis. Rather, we are arguing that multifactor productivity in the semiconductor industry is not a reliable gauge of technological progress over periods during which margins vary or during which other factors largely unrelated to technology induce price swings.

nature of industry analysis or fieldwork – to assess whether there were breaks in the pace of technical progress in the mid-1990s and around 2001. This evidence includes information on technology cycles in the semiconductor industry as well as data from Intel on their cost per transistor for MPUs that control for shifts in product mix and, to a degree, for scale economies. Although this approach cannot answer all of the relevant questions, we believe that it can enhance our understanding of an industry that contributed importantly to U.S. productivity growth in the past decade.

Technology Cycles

The International Technology Roadmap for Semiconductors (2001, 2003, 2005) provides an important source of information on technological progress in the semiconductor industry. The Roadmap is produced every two years (with off-year updates) under the auspices of the semiconductor industry associations in the United States, Europe, Japan, Korea, and Taiwan. It contains a consensus view of technology trends among experts from private industry, government, and academia. The Roadmap identifies the most significant technical and production challenges facing the industry and describes potential solutions to these challenges, with the aim of coordinating research activities among different segments of the industry.

The Roadmap's key summary measure of the pace of innovation is the length of a technology cycle, defined as the amount of time needed to achieve a 30 percent reduction in the width of the smallest feature on a chip.¹⁸ This simple concept summarizes a wealth of technical material and makes the bottom-line assessment in the Roadmap accessible to a broad audience.

¹⁸ Prior to the 2005 edition, the Roadmap also characterized the minimal feature size associated with each technology cycle. This size measure was known as the "technology node." Because DRAMs historically were the pace-setting chip for scale reduction, technology nodes referred to the minimal width for circuitry in these chips. The 2005 Roadmap, however, stopped using the DRAM-based technology node as a general summary measure of minimal feature size because advances specific to MPU and flash memory chips have been driving the scale reductions for those chips in recent years. The Roadmap now refers to separate scaling benchmarks for DRAM, MPU, and flash memory chips. In addition, the Roadmap's projections for the length of technology cycles now allow for differences across the three chip types.

Jorgenson (2001) was among the first to bring the Roadmap to the attention of the economics profession. The Roadmap available at that time documented a shift from three-year to two-year technology cycles starting in the mid-1990s.¹⁹ Relying on this assessment, Jorgenson argued that the especially rapid drop in semiconductor prices in the second half of the 1990s owed to the shift to shorter technology cycles.

However, the linkage between technology cycles and price declines in the 1990s has been muddied somewhat by a subsequent revision to the Roadmap. Based on new information from chip producers, the 2005 edition of the Roadmap re-dated the speed-up in technology cycles from the mid-1990s to 1998 for both DRAM and MPU production. Thus, the mid-1990s break in semiconductor prices identified by our structural break tests now precedes the shift to shorter technology cycles, as measured in the Roadmap, by about three years.

Moreover, the Roadmap cannot explain the emergence of slower price declines since 2001. According to the 2005 edition of the Roadmap, the two-year technology cycle was still in effect for DRAMs as of 2004. For MPU chips, the chronology in the 2005 Roadmap is somewhat ambiguous; however, industry experts indicated to us that the Roadmap should be interpreted as saying that technology cycles for MPUs also had not lengthened as of 2004. All told, the characterization of technology cycles in the Roadmap tells a different story about shifts in technology trends than does our econometric evidence for price declines.

Although the Roadmap is a rich source of information about semiconductor technology, the process by which the Roadmap committee reaches its judgments about technology cycles is

¹⁹ The length of a technology cycle, as defined in the Roadmap, is distinct from Moore's Law, which refers to the historical tendency for the number of transistors on leading-edge chips to double roughly every eighteen months. The Roadmap (2005, page *i*) notes that this rate of increase in transistors per chip, combined with a three-year technology cycle for shrinking feature width, resulted in a continuous increase in chip size; by the mid-1990s, the larger chips had begun to substantially boost production costs. The shift to a two-year technology cycle allowed Moore's Law to remain on track while relieving the pressure on the chipmakers' cost structure.

quite opaque. To avoid exclusive reliance on this “black box” approach, we obtained two other indicators of semiconductor technology cycles. The first indicator is the sequence of introduction dates for advances in lithography techniques. Lithography refers to the process by which semiconductor producers imprint a chip’s circuitry on the silicon base material. Accordingly, technological advances in lithography govern the rate of shrinkage in chip components. Our second indicator consists of the introduction dates for the first Intel MPU chip produced with the most advanced lithographic process.

Table 5 presents these two indicators of technology cycles. The date shown for each lithography process represents the first year in which at least one semiconductor manufacturer produced chips in volume with that process.²⁰ As shown, new lithography processes have been introduced every two or three years since the late 1960s. The latest generation allows manufacturers to produce chips with individual features as small as 65 nanometers (a nanometer is one billionth of a meter). According to the dates in the table, Intel has been very close to the frontier of lithography processes, if not at the frontier. Since the late 1980s, Intel has introduced MPU chips made with the most advanced lithography process either in the first year that process was used by any semiconductor firm or in the following year.²¹

We used the information in table 5 to calculate the length of technology cycles as defined in the Roadmap – i.e., the number of years needed to achieve a 30 percent reduction in scaling. The results are shown below. Based on the dates for the adoption of new lithography processes, the technology cycle averaged three years over 1969-1993, before speeding up to an average of

²⁰ We thank Dan Hutcheson of VLSI Research Inc. for providing these data.

²¹ For the 1500 nanometer process introduced in the early 1980s, our data indicate that Intel sold chips based on this technology two years *before* the process was used anywhere in the industry. Although we were unable to resolve this inconsistency, it has no effect on the average length of technology cycles over the periods that we examine. Those results depend only on the beginning-of-period and end-of-period values for the MPU and lithography series and not on the values for intermediate years.

about two-year cycles over 1993-2005. Within this later period, the especially rapid scaling reduction from 1993 to 1995 was followed by cycles that averaged roughly two years over both 1995-2001 and 2001-2005. Importantly, these data show no reversion to longer cycles after 2001. The results based on the introduction of Intel MPU chips are very similar. Overall, these results accord reasonably well with the description of technology cycles in the Roadmap.

**Technology Cycles
(Years needed for 30 percent reduction in scaling)**

Lithography process	
1969-1993	3.0
1993-2005	1.9
1993-1995	1.3
1995-2001	2.2
2001-2005	2.1
Intel MPU chips	
1971-1994	2.9
1994-2005	1.8
1994-1995	.7
1995-2001	2.2
2001-2005	2.1

The only notable difference vis-à-vis the Roadmap concerns the timing of the shift to shorter cycles in the 1990s. The Roadmap focuses on 1998, while our data series identify 1993 or 1994. It may not be possible to resolve this discrepancy given the opaque nature of the Roadmap methodology. That said, both the Roadmap and other sources point to a speed-up in the pace of technological progress during the 1990s (with some uncertainty about the precise timing) and no slowdown around 2001.

Additional Data on Cost per Transistor

The evidence about technology cycles is all well and good, but those data say nothing about the cost of achieving successive advances in technology and so are a step removed from a measure of cost per transistor, a key ingredient in our decomposition of constant-quality prices of semiconductors. However, we obtained cost data from Intel that help bridge the gap. These data show Intel's cost per transistor from 1995 to 2005 for leading-edge, high-volume MPU chips.²² By focusing on chips at the technological frontier, these data abstract from shifts in product mix. In addition, the Intel data exclude research and development (R&D) and other non-production costs, thereby controlling in part for shifts in scale economies. As can be seen in figure 5, these data show a fairly steady downtrend of about 35 percent per year in cost per transistor from 1995 to 2005. Thus, controlling for mix shifts and partly for scale economies appears to eliminate any evidence of a slowdown in the pace of cost reduction in 2001 for MPUs.

By excluding R&D and other non-production costs, the Intel data are closer in concept to short-run marginal cost than are the primary cost data employed in the paper. If we were to use the Intel cost series in our price decomposition, it would show a nearly constant rate of decline over 1995-2001 and over subsequent years, and hence would explain little of the shift to slower MPU price declines in 2001. Instead, the associated markup would have to do more of the work in the decomposition. This result provides a useful robustness check on our conclusion about the importance of time-varying markups for understanding the dynamics of MPU prices.

6. CONCLUSIONS

Some analysts (including two of the authors of this paper) have used changes in semiconductor prices to infer the pace of technological progress in this industry. The formal

²² These data are reproduced by permission of Intel Corporation, which retains the copyright. We are grateful to Mung Chen of Intel for providing the data.

econometric tests in this paper document a statistically significant shift to faster price declines for DRAM and MPU chips in the mid-1990s, followed by a significant reversion to slower price declines in 2001. Taken at face value, this evidence would suggest that the pace of technical progress in semiconductors sped up in the mid-1990s and then slowed around 2001.

However, the analysis presented here casts doubt on some aspects of that conclusion. First, our decompositions of DRAM and MPU prices indicate that swings in price-cost markups account for a considerable part of the price dynamics over the past fifteen years. After controlling for these movements in markups, the implied cost trends point to notably smaller swings in the pace of technical progress than do the original price series. Second, the implied cost trends themselves may be affected by factors (such as scale economies or shifts in product mix) that are largely unrelated to the pace of technical progress. To assess this possibility, we examined several indicators of the length of technology cycles, obtained a cleaner cost measure directly from Intel, and consulted with industry experts. On balance, this assessment points to a speed-up in the pace of technical progress in the mid-1990s, but it indicates no slowdown around 2001.

Where does this leave us? As noted above, the econometric evidence points clearly to breaks in semiconductor price trends in the mid-1990s and in 2001. We believe that the first of these breaks corresponds to an important speed-up in the pace of technical advance in the semiconductor industry during the mid-1990s. That said, pinning down the precise timing of this break is difficult, and the magnitude is smaller than one would think from looking at price data alone. As for the break around 2001, the totality of the evidence leads us to conclude that there was no slowdown in the pace of technical advance at that time, despite the more modest rate of decline in constant-quality semiconductor prices since then. Given the apparent lack of a

technology slowdown in 2001, we are fairly optimistic about the contribution of the semiconductor industry to economic growth going forward.

REFERENCES

- Aizcorbe, Ana (2005). "Moore's Law, Competition, and Intel's Productivity in the Mid-1990s," *American Economic Review* 95(2), May, 305-8.
- _____ (2006). "Why Did Semiconductor Price Indexes Fall So Fast in the 1990s? A Decomposition," *Economic Inquiry* 44(3), July 485-96.
- Aizcorbe, Ana and Samuel Kortum (2005). "Moore's Law and the Semiconductor Industry: A Vintage Model," *Scandinavian Journal of Economics* 107(4), December, 603-30.
- Bai, Jushan and Pierre Perron (1998). "Estimating and Testing Linear Models with Multiple Structural Changes," *Econometrica* 66(1), January, 47-78.
- _____ (2003). "Critical Values for Multiple Structural Change Tests," *The Econometrics Journal* 6(1), 72-8.
- _____ (2006). "Multiple Structural Change Models: A Simulation Analysis," in *Econometric Theory and Practice: Frontiers of Analysis and Applied Research*, edited by Dean Corbae, Steven N. Durlauf, and Bruce E. Hansen. New York: Cambridge University Press, 212-40.
- Banerjee, Anindya, Robin L. Lumsdaine and James H. Stock (1992). "Recursive and Sequential Tests of the Unit-Root and Trend-Break Hypotheses: Theory and International Evidence," *Journal of Business and Economic Statistics* 10(3), July, 271-87.
- Basu, Susanto, John Fernald, Jonas Fisher, and Miles Kimball (2005). "Sector-Specific Technical Change," paper presented at the NBER Summer Institute, July.
- Chwelos, Paul (2003). "Approaches to Performance Measurement in Hedonic Analysis: Price Indexes for Laptop Computers in the 1990's," *Economics of Innovation and New Technology* 12(3), June, 199-224.
- Congressional Budget Office (2002), "The Role of Computer Technology in the Growth of Productivity," CBO Paper, May.
- Feenstra, Robert C., Marshall B. Reinsdorf, Matthew J. Slaughter, and Michael Harper (2005). "Terms of Trade Gains and U.S. Productivity Growth," paper presented at the NBER Summer Institute, July.
- Flamm, Kenneth (2004). "Moore's Law and the Economics of Semiconductor Price Trends," in *Productivity and Cyclicalities in Semiconductors: Trends, Implications, and Questions – Report of a Symposium*, edited by Dale W. Jorgenson and Charles W. Wessner. Washington, D.C.: National Academies Press, 151-70.

- Grimm, Bruce T. (1998). "Price Indexes for Selected Semiconductors, 1974-96," *Survey of Current Business* 78(2), February, 8-24.
- Hall, Robert E. (1988). "The Relation between Price and Marginal Cost in U.S. Industry," *Journal of Political Economy* 96(5), October, 921-47.
- Hansen, Bruce E. (2001). "The New Econometrics of Structural Change: Dating Breaks in U.S. Labor Productivity," *Journal of Economic Perspectives* 15(4), Fall, 117-28.
- International Technology Roadmap for Semiconductors* (2001, 2003, and 2005 editions). Jointly sponsored by the Semiconductor Industry Association in the United States and similar associations in Europe, Japan, Korea, and Taiwan. Available at <http://www.itrs.net/reports.html>
- Jorgenson, Dale W. (2001). "Information Technology and the U.S. Economy," *American Economic Review* 91(1), March, 1-32.
- Jorgenson, Dale W. and Kevin J. Stiroh (2000). "Raising the Speed Limit: U.S. Economic Growth in the Information Age," *Brookings Papers on Economic Activity* 1, 125-211.
- Jorgenson, Dale W., Mun S. Ho, and Kevin J. Stiroh (2002). "Projecting Productivity Growth: Lessons from the U.S. Growth Resurgence," *Economic Review*, Federal Reserve Bank of Atlanta, third quarter, 1-13.
- McKinsey Global Institute (2001). "US Productivity Growth 1995-2000: Understanding the Contribution of Information Technology Relative to Other Factors," Washington, DC: McKinsey & Company.
- Moore, Gordon E. (1965). "Cramming More Components onto Integrated Circuits," *Electronics*, 38(8), April 19.
- Oliner, Stephen D. and Daniel E. Sichel (2000a). "The Resurgence of Growth in the Late 1990s: Is Information Technology the Story?" *Journal of Economic Perspectives* 14(4), Fall: 3-22.
- _____ (2000b). "The Resurgence of Growth in the Late 1990s: Is Information Technology the Story?" Federal Reserve Board, Finance and Economics Discussion Series, no. 2000-20, March. Available at www.federalreserve.gov/pubs/feds/2000/200020/200020pap.pdf
- _____ (2002). "Information Technology and Productivity: Where Are We Now and Where Are We Going?" *Economic Review*, Federal Reserve Bank of Atlanta, third quarter, 15-44.
- Quandt, Richard E. (1960). "Tests of the Hypothesis that a Linear Regression Obeys Two Separate Regimes," *Journal of the American Statistical Association* 55, June, 324-30.
- Standard & Poor's (2005). *Standard & Poor's Industry Surveys: Semiconductors*. February 24. New York: The McGraw-Hill Companies.

VLSI Research Inc. (2006). "Did Acceleration from a Three to Two Year Node Life Really Occur?" *The Chip Insider*, April 6.

Whelan, Karl (2002). "Computers, Obsolescence, and Productivity," *The Review of Economics and Statistics* 84(3), August, 445-61.

DATA APPENDIX

Semiconductor Price Indexes

We relied heavily on internal Federal Reserve price indexes for semiconductor products, which are constructed as part of the Fed's program to publish estimates of industrial production and capacity utilization. For the period since 1992, the Federal Reserve calculates chained, matched-model price indexes for different types of semiconductors (e.g. MPUs) from large datasets of prices and shipments of individual chips. We extrapolated the price indexes back from 1992 using price measures from a variety of different sources. Details are provided below.

Aggregate price index, 1975-2004 (annual)

For the period 1992-2004, we used the internal Federal Reserve price index for shipments for NAICS product class 3344131 (integrated circuits). We extrapolated this series back to 1977 using an internal Federal Reserve price index for SIC 36741 (integrated circuits). We then extrapolated back to 1975 using a price index for memory chips constructed by Grimm (1998). For details on this extrapolation backwards from 1977, see the data appendix in Oliner and Sichel (2000b).

MPU price index, 1987:Q1-2004:Q4 (quarterly)

For the period 1992:Q1-2004:Q4, we used the internal Federal Reserve price index for microprocessors. We extrapolated this series back to 1987:Q1 using a matched-model geometric-means index that we created from quarterly price data for individual Intel MPU chips from Dataquest, Inc.

DRAM price index, 1975:Q1-2004:Q4 (quarterly)

For the period from 1992:Q1-2004:Q4, we used the internal Federal Reserve price index for DRAMs. We extrapolated this series backward to 1975:Q1 using a series on price per megabit from Dataquest, Inc.

Series Used in the Decomposition of Price per Transistor

Price-cost markup

As described in the text, we calculated the price-cost markup (m) for Intel and Micron Technology from the equation $m = 1/(1 - \pi)$, where π represents the ratio of pre-tax operating income (adjusted for the value of stock-option grants) to net sales. The data sources for the components of π are as follows.

Net sales and operating income. These annual series were obtained from Compustat: data item A12 for net sales and data item A178 for pre-tax operating income after depreciation. Because Micron Technology derived a substantial part of its total revenue during 1995-2001 from selling personal computers, we excluded this line of business from the net sales and

operating income reported in Compustat for these years. This adjustment was based on data in the company's 10-K filings with the Securities and Exchange Commission.

Value of stock-option grants. For 1994-2004, we estimated the value of option grants using the Black-Scholes formula and data on the companies' option programs reported in their 10-K filings. Prior to 1994, accounting rules did not require firms to report detailed information about their option programs. We extrapolated the estimated 1994 grant values back to earlier years by assuming that the value of grants was a constant share of the value of the firm's common stock.

Operating cost

We calculated operating cost as net sales minus operating income (adjusted for the value of stock-option grants). The series on net sales, operating income, and option grants are the same as those used to calculate the price-cost markup.

Number of transistors shipped

For DRAM chips, we used data from Gartner, Inc. on annual unit shipments of such chips by Micron Technology for each memory class (i.e., 128 megabit chips, 256 megabit chips, and so on). We multiplied the unit shipments in each memory class by the number of transistors per chip and then aggregated the result across the memory classes.

For MPU chips, we matched data on Intel's annual unit shipments by chip (from Dataquest for 1988-92 and from Instat MDR for 1993-2004) with data on the number of transistors per chip. We multiplied the unit shipments by the number of transistors on each chip and then aggregated the result across all chips. For all MPU chips, except Celeron chips produced after 1999, the number of transistors per chip was taken directly from Intel's "Microprocessor Quick Reference Guide" found at <http://www.intel.com/pressroom/kits/quickreffam.htm>. For the more-recent Celeron chips, we assumed that the number of transistors contained in each chip was the same as that contained in its Pentium counterpart (that is, the Pentium chip produced using the same die and process).

Table 1
Semiconductor Prices
 (Percent change over periods shown, annual rate)

	Aggregate index	MPUs	DRAMs
1975 – 1994	-22.5	n.a.	-28.4
1975 – 1988	-24.3	n.a.	-31.0
1988 – 1994	-18.6	-30.0	-22.5
1994 – 2001	-47.1	-63.1	-48.3
2001 – 2004	-28.2	-40.5	-11.1

Note: The aggregate index is an annual series; the MPU and DRAM series are quarterly and were converted to annual averages before calculating the numbers in the table.

n.a.: Not available.

Table 2
Parameter Constancy Tests Allowing for Two Possible Breakpoints
in Semiconductor Price Trends¹

	Aggregate Index	DRAMs	MPUs
<i>First stage: UDMAX test²</i>			
Test statistic	7.4	8.6	40.3
Critical value, 5%	8.0	8.0	8.0
Critical value, 10%	6.6	6.6	6.6
Evidence of a break at 5%	no	yes	yes
Evidence of a break at 10%	yes	yes	yes
<i>Second stage: $supF_T[l+1/l]$ test³</i>			
Date of first break	1995	1995:Q4	1994:Q2
Test statistic	7.3	9.4	15.8
Critical value, 5%	10.1	10.1	10.1
Critical value, 10%	8.5	8.5	8.5
Evidence of second break at 5%	no	no	yes
Evidence of second break at 10%	no	yes	yes
Date of second break	--	2001:Q4	2001:Q4

¹We follow the multi-stage procedure described in Bai and Perron (2006). The first stage in this procedure uses a double maximum test (the UDMAX test) to test for the presence of a break. If the UDMAX test indicates a break, the second stage uses the $supF_T(l+1/l)$ test to test for the presence of a second break conditional on the first break. The data used for the aggregate semiconductor price index are annual from 1975 to 2004, the data for DRAM prices are quarterly from 1975:Q1 to 2004:Q4, and the data for MPU prices are quarterly from 1987:Q1 to 2004:Q4.

²The UDMAX test is described in Bai and Perron (1998). We use the unweighted version of the test with a maximum number of two breakpoints. The test statistic is calculated with Newey-West standard errors, which are computed with a lag length of four quarters for MPU and DRAM prices and one year for overall semiconductor prices. The sample trim is 25 percent for each price series. Critical values are from the unpublished tables associated with Bai and Perron (2003).

³The $supF_T(l+1/l)$ test is described in Bai and Perron (1998). If the UDMAX test indicates the presence of a break, we then implement the $supF_T(l+1/l)$ test to consider the null of one break versus the alternative of two breaks. The sample trim is 15 percent for each price series. Critical values are from the unpublished tables associated with Bai and Perron (2003).

Table 3
Sources of Price Change for DRAMs
(Average annual rates; price change in percent, contributions in percentage points)

	1990-1995 (1)	1995-2001 (2)	2001-2004 (3)	(2)-(1) (4)	(3)-(2) (5)
1. Constant-quality price	-14.9	-53.6	-11.1	-38.7	42.5
<i>Contributions from:</i>					
2. Price/transistor (\dot{p})	-13.6	-53.1	-16.6	-39.5	36.5
3. Price/cost markup (\dot{m})	14.5	-14.7	10.9	-29.2	25.6
4. Cost/transistor (\dot{c})	-24.6	-45.0	-24.8	-20.5	20.2
5. Cross product ($\dot{c} \times \dot{m}$)	-3.6	6.6	-2.7	10.2	-9.3
6. Other quality change/ residual (R)	-1.2	-.5	5.5	.7	6.0

Note: Rounding may cause slight differences between column 4 and column 2 minus column 1, between column 5 and column 3 minus column 2, between line 1 and the sum of lines 2 and 6, and between line 2 and the sum of lines 3, 4, and 5.

Table 4
Sources of Price Change for MPUs
(Average annual rates; price change in percent, contributions in percentage points)

	1988-1994 (1)	1994-2001 (2)	2001-2004 (3)	(2)-(1) (4)	(3)-(2) (5)
1. Constant-quality price	-30.0	-63.1	-40.5	-33.1	22.6
<i>Contributions from:</i>					
2. Price/transistor (\dot{p})	-30.2	-38.1	-17.1	-7.9	21.0
3. Price/cost markup (\dot{m})	2.8	-5.3	11.2	-8.1	16.5
4. Cost/transistor (\dot{c})	-32.1	-34.7	-25.5	-2.6	9.2
5. Cross product ($\dot{c} \times \dot{m}$)	-.9	1.8	-2.9	2.7	-4.7
6. Other quality change/ residual (R)	.2	-25.0	-23.4	-25.2	1.6

Note: Rounding may cause slight differences between column 4 and column 2 minus column 1, between column 5 and column 3 minus column 2, between line 1 and the sum of lines 2 and 6, and between line 2 and the sum of lines 3, 4, and 5.

Table 5
Year of Introduction for New Semiconductor Technology

Process (nanometers)	Lithography frontier	Intel MPU chips
10,000	1969	1971
8000	1972	n.a.
6000	n.a.	1974
5000	1974	n.a.
4000	1976	n.a.
3000	1979	1979 ¹
2000	1982	n.a.
1500	1984	1982
1250	1986	n.a.
1000	1988	1989
800	1990	1991
600	1993	1994
350	1995	1995
250	1997	1997
180	1999	1999
130	2001	2001
90	2003	2004
65	2005	2005

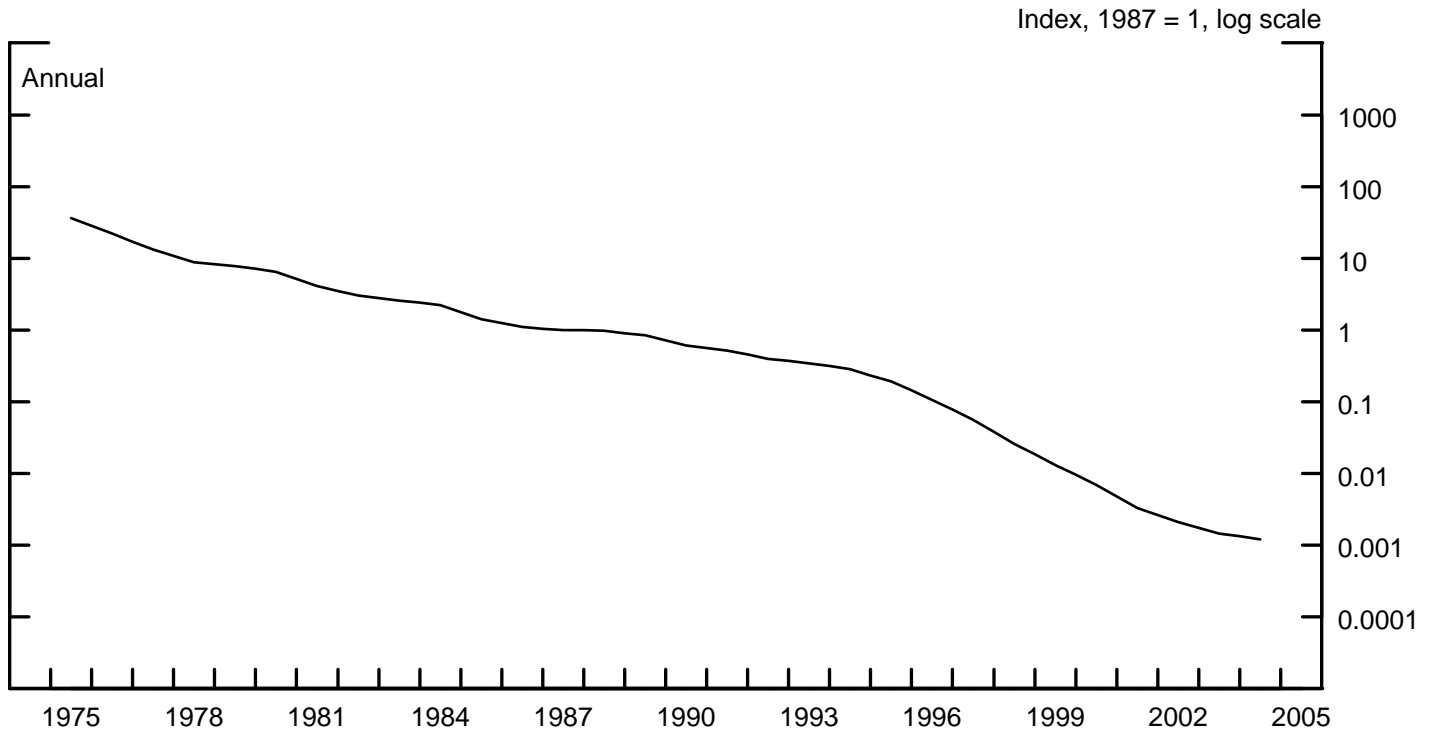
¹ Intel began making MPU chips with this process in 1979. We omitted Intel's earlier use of the 3000 nanometer process (starting in 1976) to produce less complex devices, such as scales.

n.a.: Not available.

Source. VLSI Research Inc. (2006) for the introduction dates for frontier lithography processes. Intel's introduction date for the 65-nanometer process was obtained from its 2005 annual report; all other Intel introduction dates were obtained from Intel's website (<http://www.intel.com/pressroom/kits/quickreffam.htm>).

Figure 1
Semiconductor Prices

Aggregate Index



MPUs and DRAMs

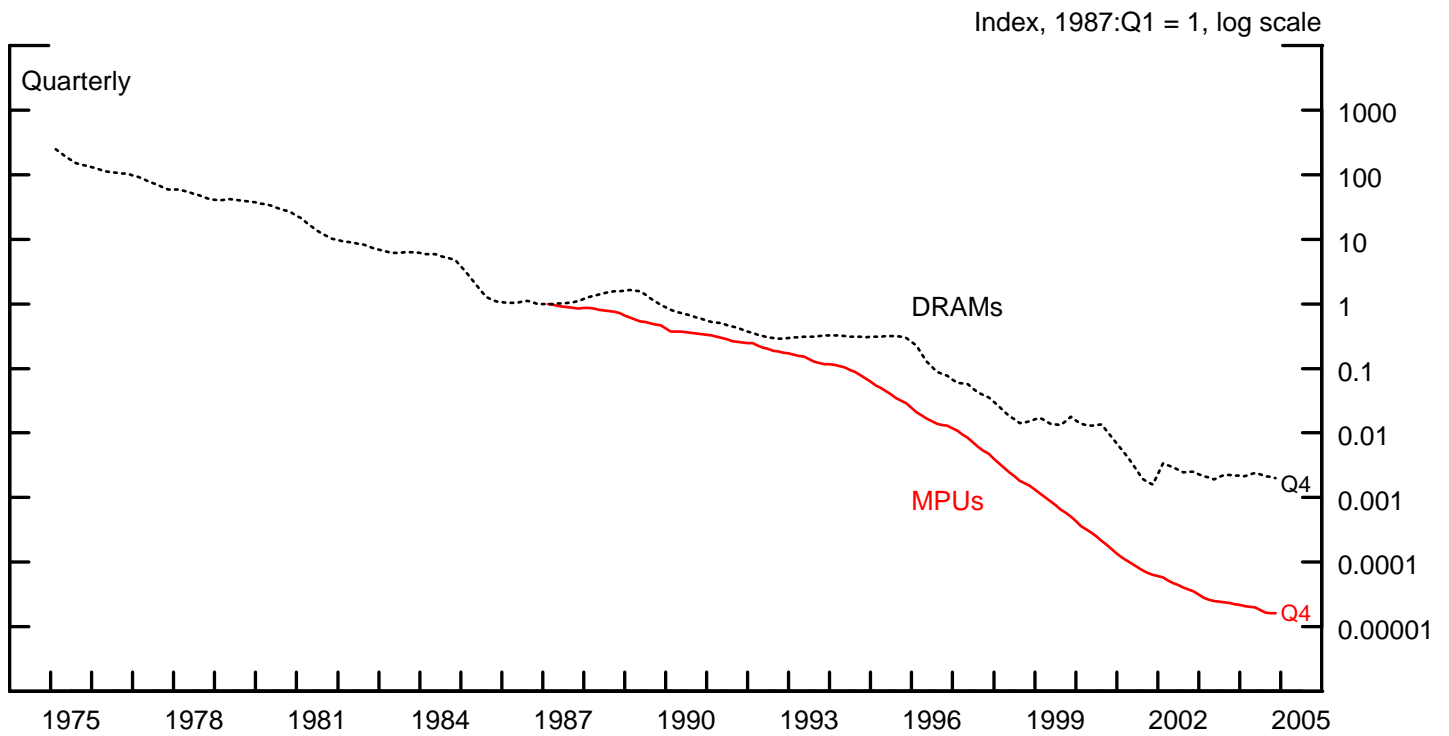
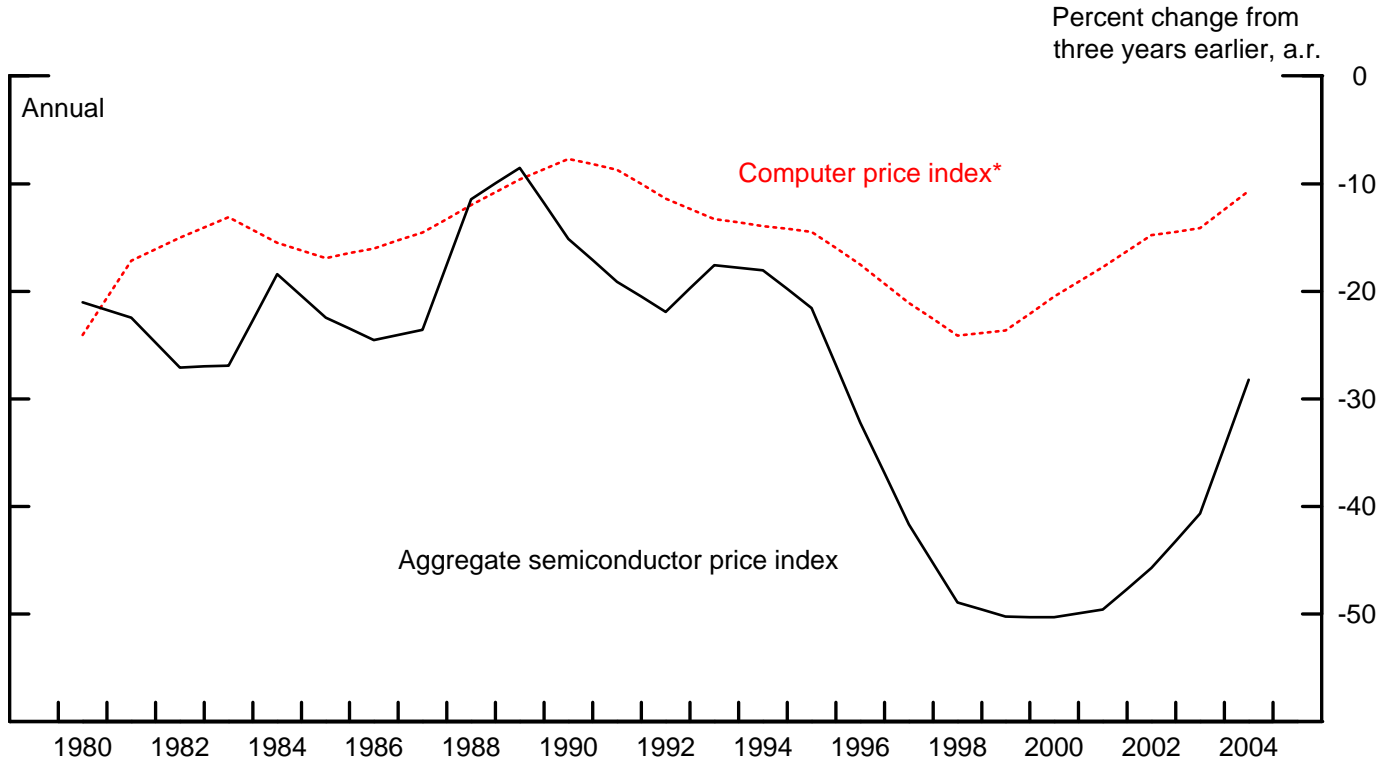



Figure 2
Declines in Semiconductor and Computer Prices



* NIPA chain-weighted price index for computers and peripheral equipment.

Figure 3
Price-Cost Markup*



* The ratio of price to average operating cost (adjusted to account for the value of stock option grants).
Source. Compustat and company 10-K statements.

Figure 4
Decomposition of Constant-Quality Semiconductor Price Changes

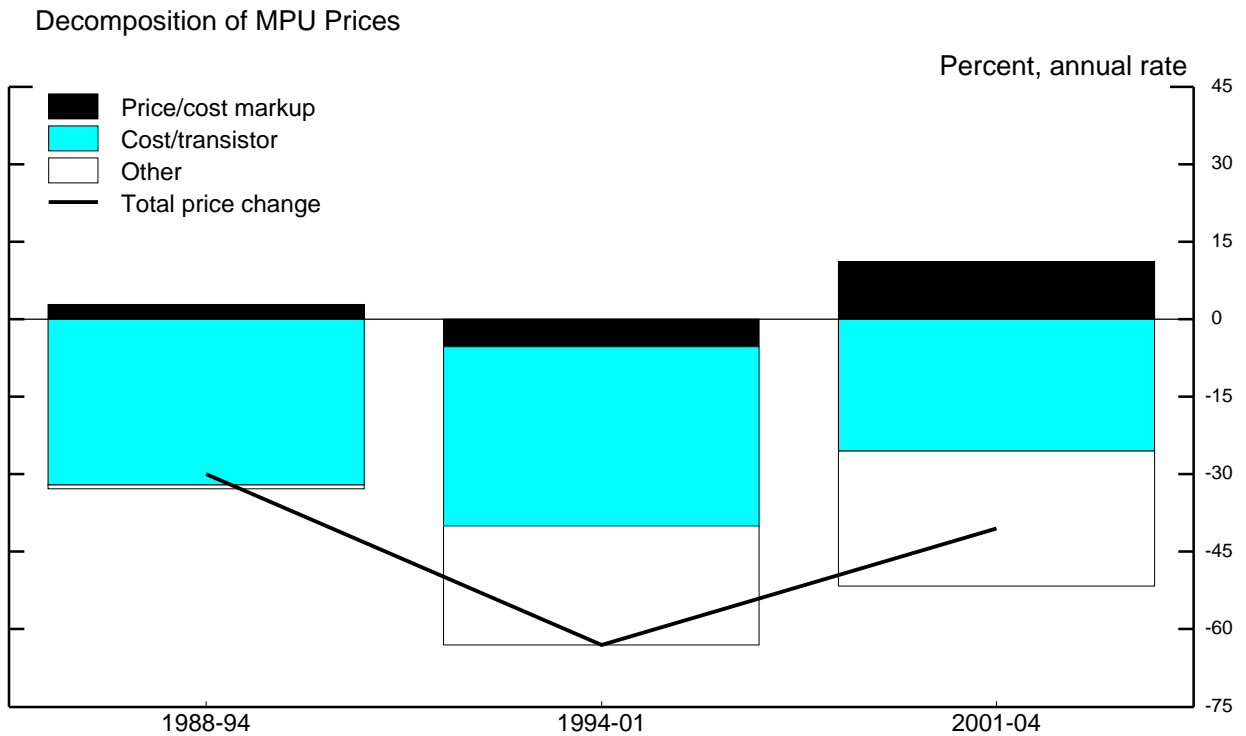
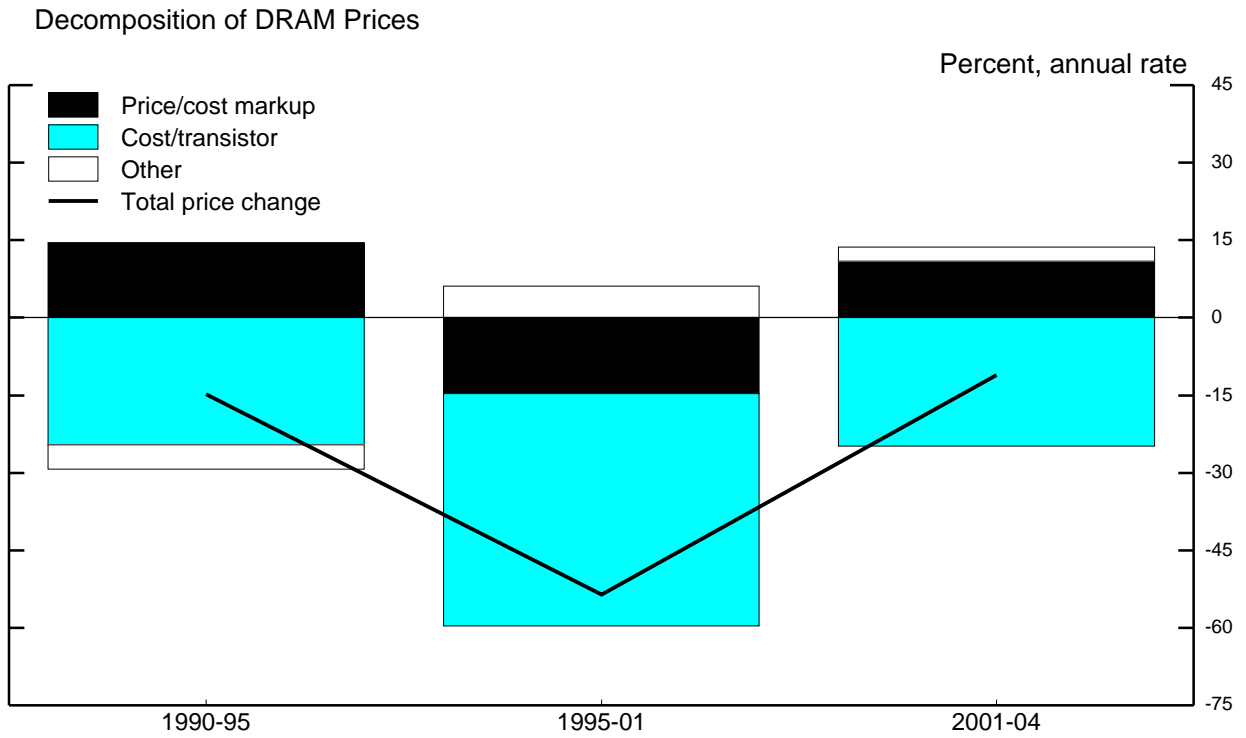



Figure 5
Intel's Production Cost per Transistor for Frontier MPU Chips



*At an annual rate of 34.7 percent, the average annual decline in the Intel series over 1995-2005.

Source. Intel Corporation. These data are reproduced by permission of Intel Corporation, which retains the copyright.