Why Are Semiconductor Prices Falling So Fast?

Industry Estimates and Implications for Productivity Measurement

Ana Aizcorbe
Federal Reserve Board
aaizcorbe@frb.gov

March 2002

By any measure, price deflators for semiconductors fell at a staggering pace over much of the last decade. These rapid price declines are typically attributed to technological innovations that lower constant-quality manufacturing costs. But, given Intel's dominance in the microprocessor market, those price declines may also reflect changes in Intel's profit margins. Disaggregate data on Intel's operations are used to explore these issues. There are three basic findings. First, the industry data show that Intel's markups from its microprocessor segment shrank substantially from 1993-99. Second, about 3-1/2 percentage points of the average 24 percent price decline in a price index for Intel's chips can be attributed to declines in these profit margins over this period. And, finally, the data suggest that virtually all of the remaining price declines can be attributed to quality increases associated with product innovation.

* I thank Steve Oliner and Dan Sichel for many useful discussions and extremely helpful comments. I am also grateful to Eric Chiang, Nile Hatch, Bart Hobijn, David Lebow, Kevin Stiroh, Jack Triplett and Philip Webre for useful comments. Kevin Krewell (MicroDesign Resources) kindly provided the data on Intel's operations and Christopher Schildt provided excellent research assistance. The views expressed in this paper are solely mine and should not attributed to the Board of Governors of the Federal Reserve System or its staff.

1. Introduction

By any measure, price deflators for semiconductors fell at a staggering pace over much of the last decade. As shown in table 1, fisher price indexes for integrated circuits—a broad class of semiconductor devices that includes logic and memory chips—fell between 36 and 56 percent each year after 1995, pushed down primarily by declines of over 50 percent per year in the price index for microprocessors—the logic chips that serve as the central processing unit in PCs.

These rapid price declines are typically attributed to technological innovations that lower constant-quality manufacturing costs. Indeed, the industry is credited with one of the fastest rates of product innovation and technical change within manufacturing as chipmakers generate wave after wave of ever-more powerful chips for prices not much higher than those of existing chips. Informal measures of quality change for microprocessors (MPUs) suggest that quality change is the primary driver behind the price declines typically seen for these devices (Aizcorbe, Corrado, and Doms (2000)). At the same time that the quality of chips is increasing, manufacturers are getting better at producing them. However, no one has examined the potential importance of these cost savings for the measured declines in price indexes.

Strictly speaking, several assumptions must hold for these indexes to provide accurate measures of productivity.[1] Among these, the assumption of perfect competition is perhaps most troublesome given Intel's long-standing dominance of the MPU market. This issue is not relevant for the calculation of an *input* price index, because the actual price paid (including any markup) is what matters for the productivity of downstream industries, and that is precisely what the input price index measures. However, use of an *output* price in measuring productivity for the semiconductor industry could be problematic. Under perfect competition, an output "price" index can be used to measure productivity because it tracks changes in (unmeasured) marginal costs. Absent perfect competition, markups form a wedge between prices and marginal costs and the use of a

[1] See Jorgenson and Griliches (1967), Diewert (1983) and Diewert (1999) for the theoretical foundations underlying these productivity measures.

price index might lead one to incorrectly interpret a price decline caused by a shrinking markup as a productivity improvement.[2]

This paper examines data on Intel's prices, costs and shipments over the period 1993-99 to explore these issues. The data show that Intel's markups were, in fact, squeezed over this period, perhaps owing to increased competition from its rivals and from weaker-than-expected demand for personal computers in 1995 and beyond. In either case, these shrinking markups can cause nontrivial distortions in a Tornquist price index that has not been corrected for the presence of markups. The distortions can overstate productivity growth (when markups are falling) or understate productivity growth (when markups are rising). On balance, a typical Tornquist price index for Intel's microprocessors falls an average of 24-1/2 percent per quarter for the period 1993-99 while one that controls for the influence of markups falls about 21 percent per quarter. Although this *average* difference is small, differences for shorter time periods can be large. Over 1998, for example, about 12 percentage points of the 38 percent decline in the usual price index can be accounted for by the sharp decline in markups that occurred that year.

This is an important issue because price indexes like these have been used to make inferences about productivity and technological change in recent studies of multifactor productivity (MFP) trends in the 1990s.[3] In many of these studies, MFP for each sector is measured using the "dual" approach, or, as the difference between an output price index and an input price index. These studies have found that the semiconductor industry played an important role in generating the rapid productivity rates seen in the 1990s and the speedup in productivity that occurred in the mid 1990s. Our findings suggest that these productivity measures for the production of

[2] The importance of market structure for productivity measurement has been studied in many different contexts. In the empirical micro literature, Denny, Fuss and Waverman (1981) and Morrison (1992) econometrically estimate multiproduct cost functions to remove the influence of markups in productivity measures. Elsewhere, Diewert (1983, 1999) suggests that markups be handled in the same way that excise taxes are in productivity measurement. In the macro literature, Hall (1988), Domowitz, Hubbard and Peterson (1988), and Basu and Fernald (1997) expand the Solow growth model to account for the presence of markups. Finally, Anderson, dePalma and Thisse (1992) and Feenstra (1995) examine the effect of markups on price indexes in the context of specific functional forms and Berry, Levinson and Pakes (1995) and Pakes(2001) study the effect of markups on hedonic regressions.

[3] See, for example, Triplett (1998), Jorgenson (2000), Oliner and Sichel (2000), Jorgenson and Stiroh (2000), and Gordon (2001).

microprocessors are probably a bit overstated for the period 1993-99. But, the overstatement is fairly evenly spread over the period and cannot explain any of the speedup in price declines that occurred in the mid-1990s.

In terms of explaining the part of these declines that cannot be attributed to changes in markups, virtually all of the remaining price declines can be attributed to increases in quality associated with product innovation; very little of the declines are explained by reductions in the cost per chip that occur over the life of each chip.

1. Measuring Changes in the Average Quality of Intel's MPUs

"Discussions of "quality" in price indexes often place the term in quotation marks and few authors have attempted to provide a rigorous definition."[4]

The difference between a constant-quality index and an average price series is often interpreted as an informal measure of quality both by practitioners in industry and by researchers interested in price measurement.[5] The idea is that if a price index holds quality constant and an average price series does not, then the difference between the two should provide a measure of quality change. The problem is that while theory tells us how to measure *constant-quality* prices, there is nothing to tell us how to measure the *average* prices. Should it be an arithmetic average or a geometric average? Should the weights be fixed or variable? Does it matter?

Given the intuitive appeal of this informal assessment of quality change, the next section looks at some indexes and average prices that yield a sensible measure of quality change.

[4] Greenlees (1999).

[5] In industry, the issue of "quality" usually comes up in trying to explain changes in average sales prices—the data that are typically reported by trade associations. So, for example, changes in average sales prices are often explained as resulting from "mix-shift"—a change in the composition of goods of varying quality. Sometimes—as is the case for the average sales price of automobiles—the gap between average sales prices and a constant-quality price index—like the CPI—is used as a measure of quality improvements. In the academic literature, Hulten (1997) has studied the issue from a theoretical perspective. In the empirical literature, the issue typically comes up in the context of examining biases in the CPI. Reinsdorf (1993) used average sales prices for homogeneous goods as a check on potential biases in the CPI: the check being that if quality is increasing, then average sales prices should rise faster than constant-quality indexes like the CPI. More recently, Bils and Klenow (2001) use this identity in the context of a structural model to identify the degree to which BLS methods adequately control for quality change.

Interpreting Informal Measures of Quality Change

The paradigm that comes to mind when thinking about quality change is the framework implicitly used by BLS to hold quality constant when replacing one good in their basket with another. Chart 1 shows the general idea. The chart shows price profiles for two chips, with chip 2 replacing chip 1 at time t=1. The change in the price per chip from t=0 to t=2 may be stated as the product of the price changes over the life of each chip and the gap in prices of the new and exiting chip. In terms of the diagram, the change in the price per chip is the ratio of the last price for chip 2 ($P_{2,2}$) and the first price for chip 1 ($P_{1,0}$). That ratio may be written as:

$$(1) \qquad P_{2,2} / P_{1,0} = \boldsymbol{(P_{2,2} / P_{2,1})} \, (P_{2,1} / P_{1,1}) \, \boldsymbol{(P_{1,1} / P_{1,0})}.$$

This change in the price per chip could be viewed as a constant-quality price index only in the hypothetical case where the two chips are of equal quality; in that case, price per chip is all that matters. Alternatively, one can allow the chips to be of different quality and assume that any price difference at t=1 is the market's valuation of these quality differences. In this view, one obtains a constant-quality price index by measuring price changes that occur over the life of each chip--shown in bold italics in (1)--and excluding the gap in the two prices at t=1. The gap between the average price measure--the product of all three terms--and the matched-model index--the product of the two terms in italics--is the measure of quality change implicit in a matched-model index calculation—the middle term.

This seems like a sensible way to value quality change. In general, though, there are many chips that coexist in the market, and turnover is characterized by new and existing goods overlapping for some period of time. In that case, if one thinks of the prices in (1) as *average* prices, then the matched-model index still measures price change over the lives of goods existing in both periods, but quality change is measured as a difference of means: *average* prices with entry (the change in an average price series) and *average* prices without entry (the change in the matched-model index).

A geometric mean index provides a simple example to illustrate the point.[6] A matched-model geometric mean of price change over the period t, t-1 (in logged form-- $\ln P^{GEO}_{t,t-1}$) is an arithmetic mean of logged price relatives for the goods that exist in both periods:

$$\ln P^{GEO}_{t,t-1} = \Sigma_{m=match(t)} (\ln P_{m,t} - \ln P_{m,t-1}) / M_t \tag{2}$$

where models that exist in both periods are denoted match(t) and the number of such models at time t is denoted M_t. Suppose that a new good enters at time t. Then, this geometric mean can be restated as a combination of two terms:[7]

$$\ln P^{GEO}_{t,t-1} = [\Sigma_{m=all(t)} (\ln P_{m,t}) / N_t - \Sigma_{m=all(t-1)}(\ln P_{m,t-1}) / N_{t-1}]$$
$$[- (\Sigma_{m=all(t)} (\ln P_{m,t}) / N_t - \Sigma_{m=matched(t)}(\ln P_{m,t}) / M_t)] \tag{3}$$

where the goods that exist at time t are indexed by m=all(t) and the number of such goods is denoted by N_t.

The first term gives the difference of the (geometric) average sales prices over the two periods. The second term compares an average sales price for time t that *includes* the new good to one that *excludes* the new good. This latter term is a measure of quality change in that when the arrival of the new good raises the average sales price, it must be that the new good is viewed superior--or, of higher quality--by the market. This is the same intuition as in the simple case above, except that quality change there was measured as differences in *individual* prices whereas here it is measured as differences in *average* prices. Again, the benchmark for comparison is the hypothetical case where all goods are homogeneous, in which case the price of new goods would be the same as that of existing goods and the second term in (3) equals zero. In that view, any observed difference in

---

[6] Another reason to look at the quality measure implicit in a geometric index of price change is the close link between the geometric mean and a hedonic regression of the type discussed in Aizcorbe, Corrado and Doms (2000). Given their argument, one could say that the quality adjustment implicit in regressions of the type run by ACD is the same as that in the geometric mean discussed above.

[7] To see this, add and subtract two (geometric) means: one for *all* N logged prices at time t and one for all prices at time t-1. Rearranging the expression gives (3).

the price of new and existing goods can be taken to be a measure of their quality differences. A similar expression can be derived for exiting goods.

The particular functional form of the quality measure depends on that of the price measure. In (3), each price gets an equal weight because the functional form for the price measure is a geometric mean. Moreover, note that for this functional form, "quality" only changes when there is turnover. This is because the weights in the index are fixed (at 1/N) and any changes in the relative importance (in terms of sales, say) of one good relative to another are not counted as quality change.

In contrast, superlative indexes do capture changes in the relative importance of goods and place the correct weights on them. For one such superlative index—the Tornquist—one can apply the same logic that was applied above for the geometric mean. Although the algebra is a little more complicated, one can show that the Tornquist price index captures quality changes that result from both turnover--differences in means with and without the new good--and from mix-shift among existing goods--changes in the relative importance of existing goods.

Intel Price Data and Calculations

The quality measures described above are calculated using the dataset used in Aizcorbe, Corrado and Doms (2000). Those data were obtained from MicroDesign Resources (MDR), the industry's primary source for data on Intel's operations. The data are quarterly observations on prices, unit shipments, and revenues for Intel's microprocessors at a high level of product detail. MDR estimates prices by taking Intel's published list prices and making any needed adjustments for volume discounts. They also estimate unit shipments and revenue data using Intel's 10K reports and the World Semiconductor Trade Statistics data published by the Semiconductor Industry Association.

Chart 2 uses price profiles for selected Intel chips to show two features of these profiles that are characteristic of microprocessors and other semiconductors.[8] First, there is a high degree of turnover in this segment as new, faster chips are brought to the market. Second, prices fall steeply over the life of each chip. Prices typically start at

[8] See Flamm(1996) and Irwin and Klenow(1994) for similar profiles for DRAM memory chips.

between \$600 to \$1000 at introduction--substantially higher than the prices of existing chips. By the time the chip exits the market, its price has fallen to under \$100. The steepness of these profiles could reflect demand- or supply-driven forces. On the demand side, the profiles are consistent with the view that users are initially willing to pay high prices for new chips but as the introduction of the new (better) chip nears, they are less willing to do so and prices of the incumbent chips fall. On the supply side, these profiles are also consistent with the view that prices over the life of the chip are pulled down by declining costs as firms find lower-cost ways to produce each chip.

Three aggregate price measures were constructed using these data and are compared in table 2.[9] As shown in the first column, the two chained, matched-model indexes—a geometric mean and a Tornquist index—fall sharply over this period: the geometric mean falls at an average rate of 21.9 percent per quarter over the 1993-99 period while the Tornquist index falls 24.4 percent.[10] In contrast, the (geometric) average price series is essentially flat: that series falls an average of 0.5 percent per quarter. Apparently, the average price series says more about the distribution of prices over time than it says about declines in prices over the life of each chip. Intuitively, it is relatively flat because the declines in prices over the life of each chip are undone when the next chip enters the market at the same high introductory price.

This large gap between the matched-model indexes and the average price series suggests that virtually all of the declines in the price index stem from increases in the quality of chips. Because we think of changes in a matched-model index as changes in an average price series less any changes in quality, finding no change in the average price series means that all the movement in the matched-model index stems from changes in quality. Taking a literal read of the Tornquist index—the better measure of price change—all but 0.5 percentage points of the 24.4 percent average decline in the Tornquist price index are attributable to quality change.

---

[9] A matched-model geometric mean index (in logs) of the aggregate price change from time o to time * is given by $\ln P_{o,*} = \Sigma_{s=o+1,*} \ln P_{s,s-1}$,. For a geometric mean index, $\ln P_{s,s-1}$ is given by equation (2); for a Tornquist index, $\ln P_{s,s-1} = \Sigma_{m=match(s)}\, \omega_{m,s}(\mathit{ln}\, P_{m,s} - \mathit{ln}\, P_{m,s-1}) / M_s$, where $\omega_{m,s} = P_{m,s} X_{m,s} / \Sigma_s(P_{m,s}\, Q_{m,s})$.

[10] The percent changes reported in here do not line up with those reported in ACD(2000). The measures here are calculated as *averages* of the quarter-to-quarter price changes while those in ACD(2000) they are

The implied quality indexes—shown in the first two rows of table 3—are essentially the inverses of the price indexes. The Tornquist index grows faster than the Geometric mean measure because it accounts for two sources of quality differences—changes in the relative importance of existing chips and turnover—while the geometric mean only considers the latter. The last line of the table gives changes in a crude measure of quality that considers only the speed of chips (measured by megahertz). Apparently, speed is not the only characteristic that matters, since the megahertz per chip measure grows at a much slower rate—an average of 9.2 percent per quarter.

2. Measuring Changes in the Costs Per Chip

"In general, Intel's prices are several times the manufacturing cost of the chips, so that cost has little influence on their price."[11]

Quality improvements are only one type of technological change that could drive down the constant-quality price of goods. Improvements that lower the average cost of production by directly shifting down cost curves also represent technological progress. A brief description of the manufacturing process helps explain the two main ways that semiconductor firms lower the cost per chip.

The manufacture of semiconductor chips is extremely complex.[12] The process involves taking a silicon wafer of fixed size, etching chips on this wafer, and eventually separating out the individual chips and packaging them for sale. The manufacturing cost per wafer is constant, and anything that increases the number of usable chips on a wafer reduces the average cost per usable chip. The smaller the size of the chip, the more chips one can fit on a wafer and the lower the average cost of production. This *"chip shrinking"* is one important source of cost reductions for semiconductor producers. Firms have also reduced the cost per chip by increasing the size of the wafer upon which the chips are etched.

---

reported as *compound annual growth rates*. While both measures give similar qualitative results, the former is more intuitive.

[11] Gwennap and Thomsen (1998), P. 67

[12] See Hatch and Mowery (1998) for a full description of the manufacturing process in general, and the learning curve in particular.

A final way that firms lower the average cost per usable chip is by increasing the *yield* of production. The complexity of the manufacturing process is such that the early months of production of a new chip are marked by high defect rates that hold down yields—defined as the ratio of usable chips to all chips. As the wrinkles in the new process are ironed out, yields rise and the average cost of production falls. However, it's not clear that all learning economies should be viewed as "technological progress." Lessons learned over a long span of time--like how to make faster chips--are clearly technical change. But, the increase in yields that occurs every time a new chip is introduced may best be viewed as a form of increasing returns or an adjustment cost like the kind faced by automakers when a changes at a new model year require a ramp-up to full production volumes.

These sources of cost reductions have received a lot of attention in the literature and the conventional wisdom is that the cost savings from the learning curve are significant.[13] Below, data on Intel's costs are examined to see whether these cost savings contribute much to the steep price declines seen in semiconductor price indexes.

Data on Intel's Costs

Average cost data for each Intel chip were obtained from MDR—the same source as the price data. Their cost estimates include labor and material costs plus depreciation of the equipment and part of the building,[14] but do not include an adjustment for the design of the chip or other R&D costs. Thus, the cost concept is closer to variable cost than total cost.[15]

[13] There is a large literature aimed at estimating the importance of the learning curve in the semiconductor industry. See Hatch and Mowery (1998) for a recent review of the literature. Because the needed cost data are not readily available, empirical studies of the learning curves typically use prices as a proxy for cost. The two exceptions are Irwin and Klenow (1994)—where a structural model that specified the relationship between price and marginal cost was used to obtain learning curve estimates—and Hatch and Mowery(1998)—where a unique survey was used to obtain the needed data.

[14] MDR uses a four-year straight-line depreciation for the cost of equipment and clean room. MDR P. 68.

[15] Use of variable costs–rather than total costs—is consistent with a short-run view of production, where once the firm incurs these set-up costs (R&D and plant and equipment investment), these costs are sunk and the relevant cost concepts (marginal and average) are based on variable costs. Flamm (1996) uses a similar concept of marginal cost in his model of semiconductor production; Danzon (2000) also takes this view when discussing the cost structure for pharmaceuticals—another industry characterized by large setup costs.

Often, the way that data are collected by industry experts is as telling as the data themselves. Here, the data collection process shows two features of the cost structure that highlight a disconnect between price and costs. First, although some characteristics are very important determinants of price (like speed), they do not seem to matter much for average costs: within a chip designation (like the Pentium II with 256K cache), MDR does not collect data for chips of different speeds. When asked, an analyst at MDR stated that they do not consider speed to be an important determinant of cost--the direct quote is "speed is free."

Second, MDR collects costs for each chip only once during the life of a chip because they think costs bottom out after the initial ramp-up. Specifically, MDR collects the data somewhere between the "sixth and twelfth month after the release of a new processor, when defect rates are approaching or have reached maturity. Costs will be higher than that during the first few months of production."[16]

Another way to say that costs are constant over the life of a chip in our data is to say that MDR's definition of a "chip" is at a sufficiently detailed level so that features of the production process (and hence costs) don't change over the life of the chip. Given MDR's definition, the cost-reducing changes in the production process that were discussed above coincide with the introduction of "new" chips in this data set.[17]

Turning to the data, manufacturing costs for the typical chip are very low relative to price. The dark line in Chart 3 shows a typical price profile for one of the Pentium I chips. A horizontal dotted line is also drawn to represent MDRs cost estimate. The chip was introduced in the first quarter of 1994 at $1000. By the fourth quarter, price had fallen to $800 and costs, at that point, were estimated at $53 per chip. When the chip was removed from the market in 1997, the selling price was still well above MDR's average cost. Clearly, there is a wide gulf between price and average variable cost. The implication for markups over the life of the chip is that they are large and decline as the chip ages.

[16] MDR (1998), P. 74.
[17] Specifically, the following attributes of a chip that determine cost define a "chip" in the MDR data: gate length (minimum width of the polysilicon layer used to form transistors), process type (BiCMOS or CMOS), die size, number of metal layers used in the chip and the size of the wafer.

To get some rough idea of how markups change, on average, across chips, an estimate of the average markup is calculated under the assumption that the average variable cost of a chip is always equal to the MDR estimate. An estimate for the average markup is then calculated as revenues less variable cost divided by revenues.

As seen in table 4, the average markup is large--ranging between 70 to 90 percent from 1993 to 1999. Importantly, the average markup declined over this period from nearly 90 percent in 1993 to 73 percent by 1999. The largest declines occurred in 1995-96--when Intel was reportedly under intense competition from its rivals--and again in 1998--when the recession in Asia began to affect world demand for electronic goods. [18]

The Influence of Margins on the Price Index

In industries where firms have market power, price indexes used to measure MFP should, in principle, be calculated using marginal cost rather than price. So, one way to assess the importance of markups on these price indexes is to use the cost data described above to recalculate the price indexes using marginal costs rather than price. In principle, differences in this index and the usual price index should provide some information on the potential distortions to productivity inferences from having used the standard price index.

However, applying matched-model methods to these cost data points to a potential problem with these methods. A literal read of the MDR cost data is that the average cost of each chip (AC) is fixed over the life of the chip and, therefore, equal to marginal cost (MC). Under the assumption that MC is flat over the life of each chip, a constant-quality cost index obtained from the matched-model method would show no change over the entire period. Numerically, this happens because the matched-model method only measures cost change over the life of a chip. Conceptually, the problem seems to be that, unlike in perfect competition, the cost side does not provide any information on users' valuations of the chips and, hence, cannot provide a measure of

[18] Using firm-level data from Compustat, Oliner and Sichel (2000) checked the movement of markups over time and found that Intel's aggregate markup over all lines of business actually rose slightly from 1990-95 to 1996-99. The differences in their results and those reported here stem from the differences in cost concept—they use a broad cost concept that includes design and R&D costs to define markups—and differences in the coverage—they calculated markups for all of Intel's operations.

quality.[19] This is clearly wrong because even a very crude constant-quality cost index (like cost/mhz) falls over time to reflect the faster speeds (higher quality) of newer generations of chips.

Absent a direct solution to the problem, an indirect approach is taken to assess the importance of markups. As was done earlier, a standard price index ($P^{PRICE\text{-}BASED}$) may be viewed as an average price measure that has been purged of changes in quality:

(4) $$d\left(P^{PRICE\text{-}BASED}\right) = d\,(\text{average price}) - d(\text{quality})$$

Suppose that the presence of markups does not affect the measure of quality. Then, removing markups from this price index boils down to replacing the average price measure implicit in the usual price-based index with an average cost measure. That is, rather than calculating a price index as a change in average *prices* that is purged of any changes in quality, calculate it as a change in average *costs* that is purged of quality changes:

(5) $$d\left(P^{COST\text{-}BASED}\right) = d\,(\text{average cost}) - d(\text{quality})$$

Numerically, the expression for the quality measure implicit in a Tornquist price ($Qual^{TORN}$) index is:

(6) $$\ln Qual^{TORN}_{t,t-1} = \ln P^{PRICE\text{-}BASED}_{t,t-1} - \left[\Sigma_{m=all(t)}(\mathit{ln}\, P_{m,t})/N_t - \Sigma_{m=all(t-1)}(\mathit{ln}\, P_{m,t-1})/N_{t-1}\right],$$

where $\ln P^{PRICE\text{-}BASED}_{t,t-1}$ is a standard tornquist price index and the term in the brackets is the change in average prices (logged geometric averages).

The cost-based price index is then this quality measure plus the change in average costs:

[19] This problem would also arise if one tried to obtain estimates of quality change from a hedonic regression that uses average or marginal cost as the dependent variable. For example, it can be shown that a regression of the sort estimated in Aizcorbe, Corrado and Doms (2001) would also generate a price measure that showed no price change.

(7) $$\ln P^{COST\text{-}BASED}_{t,t-1} = \ln Qual^{TORN}_{t,t-1} + [\, \Sigma_{m=all(t)} (\, \mathit{ln}\, AC_{m,t}\,) / N_t \;-\; \Sigma_{m=all(t-1)} (\, \mathit{ln}\, AC_{m,t-1}) / N_{t-1}\,],$$

where AC denotes average cost. For these functional forms, the cost-based measure can be restated as the price-based measure less (geometric) average markups:

(8) $$\ln P^{COST\text{-}BASED}_{t,t-1} = \ln P^{PRICE\text{-}BASED}_{t,t-1} - [\Sigma_{m=all(t)}\, \mathit{ln}\, (P_{m,t}/MC_{m,t})/N_t - \Sigma_{m=all(t-1)} \mathit{ln}\, (P_{m,t-1}/MC_{m,t-1}) / N_{t-1}\,].$$

Calculations for the two Tornquist indexes are shown in table 5.[20] As may be seen, the price-based index (column 1) falls an average of about 3-1/2 percentage points per quarter faster that the cost-based index. Note, also, that the distortion is about the same in 1993-95 as it is in the latter part of the sample: a distortion of about 3 percent in 1993-95 versus 4 percent in 1996-99. Correcting these price measures for shrinking margins does not return price declines after 1995 to rates closer to those observed in the earlier part of the decade.

The distortions in the year-to-year measures can be quite large. In 1998, for example, the price-based index shows an average rate of decline of about 38 percent per quarter while the cost-based index falls an average of 26 percent: about 30 percent of the price declines in the price-based index can be attributed to falling margins. In 1995—another year where Intel's margins were squeezed—about 20 percent of the price decline there could be attributed to falling margins. Note, also, that though the distortion is most often positive, it can also be negative, as was the case in 1996.

3. Conclusions and Future Research

This paper has assessed the relative importance of technological progress and markups in generating price declines in indexes for microprocessors over the 1993-99 period. Disaggregate data on Intel's price, cost, and shipments of microprocessor chips

[20] Calculations using the geometric mean give similar qualitative results.

were explored and established that technological progress is the primary driver of the steep price declines seen in price indexes for Intel's chips.

Shrinking markups over the 1990s account, on average, for about a tenth of the measured rate of price declines in these data. Over the entire period, the rate of price decline for the price-based index is 3.6 percent faster than the adjusted index: a 24.4 average percent decline for the price-based index versus a 20.8 percent average decline for the cost-based index. Distortions in the annual measures can be quite large.

It would be interesting to explore whether movements in markups can explain the apparent break in the trend that occurred in 1995. This paper shows that profit margins decreased after 1993. If one could show that margins increased over the 1980s and early in the 1990s, then, removing the influence of markups from the price index might smooth the apparent break in the trend. An exploration of this possibility is left for future work.

References

Aizcorbe, A. M., C. Corrado and M. Doms (2000) Constructing Price and Quantity Indexes for High Technology Goods, presented at the CRIW workshop on Price Measurement at the NBER Summer Institute, July. Available at www.nber.org

Anderson, S.P., A. de Palma and J. Thisse (1992) *Discrete Choice Theory of Product Differentiation*. Cambridge, MA: MIT Press.

Basu, S. and J. Fernald (1997) "Returns to Scale in U.S. Production: Estimates and Implications," *Journal of Political Economy,* 105:249-283, April.

Berndt, E. R. (1991) "The Measurement of Quality Change: Constructing an Hedonic Price Index for Computers Using Multiple Regression Methods" in *The Practice of Econometrics: Classic and Contemporary*, Addison-Wesley Publishing Co., Inc.

Berry, S., J.A. Levinson, and A. Pakes (1995) "Automobile Prices in Market Equilibrium," *Econometrica* 63:841-890.

Bils, M. and P.J. Klenow (2001) "Quantifying Quality Growth," mimeo, May.

Danzon, P. (2000) Testimony to the U.S. Senate Committee on Health, Education, Labor and Pensions, June 13. Available at: http://www.senate.gov/~labor/hearings/june00hrg/061300wt/061300jmj/061300emk/gorton613/dorgan613/rhodes/danzon/danzon.htm

Denny, M., M. Fuss and L. Waverman (1981) "The Measurement and Interpretation of Total Factor Productivity in Regulated Industries, with an Application to Canadian Telecommunications", Pp. 179-218 in T. Cowing and Stevenson, eds., *Productivity Measurement in Regulated Industries*, New York: Academic Press.

Diewert, W.E. (1999) "Appendix A: A Survey of Productivity Measurement," in *Measuring New Zealand's Productivity*, Draft Report, January

Diewert, W.E. (1983) "The Theory of the Output Price Index and the Measurement of Real Output Change," in *Price Level Measurement*, editors W.E. Diewert and C. Montnmarquette, Statistics Canada, Ottawa, Ontario (December 1983), pp. 1049-1113.

Domowitz, I.R., G. Hubbard and B.C. Petersen (1988) "Market Structure and Cyclical Fluctuations in United States Manufacturing," *Review of Economics and Statistics* 70:55-66.

Feenstra, R.C. (1995) "Exact Hedonic Price Indexes," *Review of Economics and Statistics*, Pp. 634-653.

Fisher, F.M. and K. Shell (1998) *Economic Analysis of Production Price Indexes*, Cambridge, U.K.: Cambridge University Press,

Flamm, K. (1996) *Mismanaged Trade? Strategic Policy and the Semiconductor Industry*,Washington, D.C.: Brookings Institution

Gordon, R. (2001) "Does the New Economy Measure Up to the Great Inventions of the Past?" in *Journal of Economic Perspectives*.

Greenlees, J.S. (1999) "Consumer Price Indexes: Methods for Quality and Variety Change," Paper presented at the Joint ECE/ILO meeting on Consumer Price Indices, Geneva, 3-5 November 1999.

Gwennap, L and M. Thomsen (1998) "Intel Microprocessor Forecast, 4th ed." Sebastopol, CA: MicroDesign Resources, Inc.

Hall, R. E., (1988) "The Relation Between Price and Marginal Cost in United States Industry," *Journal of Political Economy*, 96:921-947.

Hatch, N. and D.C. Mowery (1998) "Process Innovation and Learning by Doing in Semiconductor Manufacturing," *Management Science*, 44:1461-1477.

Hobijn, Bart (2001) "Is Equipment Price Deflation a Statistical Artifact?" mimeo.

Hulten, Charles R. (1997) "Quality Change in the CPI," *Federal Reserve Bank of St. Louis Review*, May/June, Pp. 87-106.

Irwin, D. A. and P. Klenow (1994) "Learning by Doing Spillovers in the Semiconductor Industry," *Journal of Political Economy*, 102(6):1200-1227, December

Jorgenson, D.W. (2000) "Information Technology and the U.S. Economy," *Presidential Address to the American Economic Association*, New Orleans, Louisiana, January 6.

Jorgenson, D.W. and K.J. Stiroh (2000) "Raising the Speed Limit: U.S. Economic Growth in the Information Age," mimeo.

Jorgenson, D.W. and Z. Griliches (1967) "The Explanation of Productivity Change," in *Review of Economic Studies*, 34, 249-283.

McKinsey Global Institute (2001) "Semiconductor Manufacturing," in "Productivity in the United States," report available at http://www.mckinsey.com/knowledge/mgi/reports/productivity.asp

Morrison, C. J. (1992) “Unraveling the Productivity Growth Slowdown in the United States, Canada and Japan: The Effects of Subequilibrium, Scale Economies and Markups,” in *The Review of Economics and Statistics* 74(3):381-393.

Oliner, S. and D. Sichel (2000) “The Resurgence of Growth in the Late 1990s: Is Information Technology the Story?” *Journal of Economic Perspectives*, 14(4):3-22.

Pakes, A. (2001) “New Goods, Hedonics, and Price Indices; With an Application to PC’s,” mimeo, May 5.

Reinsdorf, M. (1993) “The Effect of Outlet Price Differentials on the U.S. Consumer Price Index,” in Foss, Murray F., et.al., eds., *Price Measurements and Their Uses*. Chicago, Ill: University of Chicago.

Triplett, J. (1998) “The Solow Productivity Paradox: What Do Computers Do to Productivity?”   *Canadian Journal of Economics*, Volume 32, No. 2, April 1999. p. 319. Available at http://www.csls.ca/jrn/v32n2_04.pdf.

Table 1. Chained Fisher Price Indexes for Integrated Circuits, 1993-2000
Annual Percent Changes

	1993	1994	1995	1996	1997	1998	1999
ICs	-9.34	-14.33	-36.3	-45.54	-44.27	-55.29	-49.83
Memory chips	-4.57	0.7	-9.62	-38.04	-43.7	-49.05	-17.58
DRAM	2.64	7.56	0.59	-47.16	-58.72	-61.87	-16.5
Other	-8.99	-4.78	-22.12	-23.28	-26.19	-37.26	-22.04
Logic chips	-18.79	-25.81	-53.82	-59.16	-51.42	-64.34	-61.98
MPU	**-26.07**	**-32.94**	**-63.51**	**-66.98**	**-53.6**	**-70.53**	**-69.12**
Other	-4.1	-2.36	-6.43	-35.26	-42.17	-28.33	-23.96
Other	7.86	5.62	1.9	-4.26	-11.67	-6.41	1.97
Contributions:							
Memory chips							
DRAM	0.35	1.58	0.14	-5.71	-5.35	-4.91	-1.98
Other	-1.84	-0.64	-3.98	-2.73	-2.71	-2.98	-2.12
Logic chips							
MPU	-17.55	-20.61	-43.23	-42.7	-33.69	-45.49	-47.99
Other	-1.13	-0.3	-1.34	-5.56	-6.11	-3.87	-4.12
Other	2.13	0.92	0.42	-0.8	-1.94	-0.81	0.27

Source: Author's Calculations

Table 2. Price Measures for Intel's Microprocessors, 1993-99
(average quarterly percent change)

	1993-99	1993-95	1996-99
Matched-model Indexes			
Geometric Mean	-21.9	-16.2	-25.9
Tornquist	-24.4	-17.1	-29.4
Change in Average Prices	- 0.5	2.0	- 2.3

Source: Author's calculations based on proprietary data from MDR.

Table 3. Quality Measures for Intel's Microprocessors, 1993-99
(average quarterly percent change)

	1993-99	1993-95	1996-99
Matched-model Indexes			
Geometric Mean	21.4	18.2	23.6
Tornquist	23.9	19.1	27.1
Ave. megahertz per chip	9.2	8.1	9.9

Source: Author's calculations based on proprietary data from MDR.

Table 4. Revenue, Manufacturing Costs and Implied Margin for Intel's Microprocessors.

	1993	1994	1995	1996	1997	1998	1999
Revenue	6.8	8.8	12.0	14.9	19.9	22.4	25.0
Manufacturing Cost	0.8	1.2	2.2	3.5	4.8	6.2	6.8
Implied Margin	6.0	7.6	9.8	11.4	15.1	16.2	18.2
Margin/Revenue	88.2	86.4	81.7	76.5	75.9	72.3	72.8

Source: MicroDesign Resources

Table 5. Tornquist Price Measures for Intel's Microprocessors, 1993-99
(average quarterly percent change)

	Price-based (1)	Cost-based (2)	Difference (2)-(1)
1993-99	-24.4	-20.8	-3.6
1993-95	-17.1	-14.1	-2.9
1996-99	-29.4	-25.3	-4.1
1993	- 7.4	- 0.5	- 6.9
1994	-14.4	-16.8	2.4
1995	-26.9	-21.7	- 5.3
1996	-22.8	-24.4	1.6
1997	-27.1	-26.0	-1.1
1998	-37.7	-25.7	-11.9
1999	-30.2	-25.2	-5.0

Source: Author's calculations based on proprietary data from MDR.

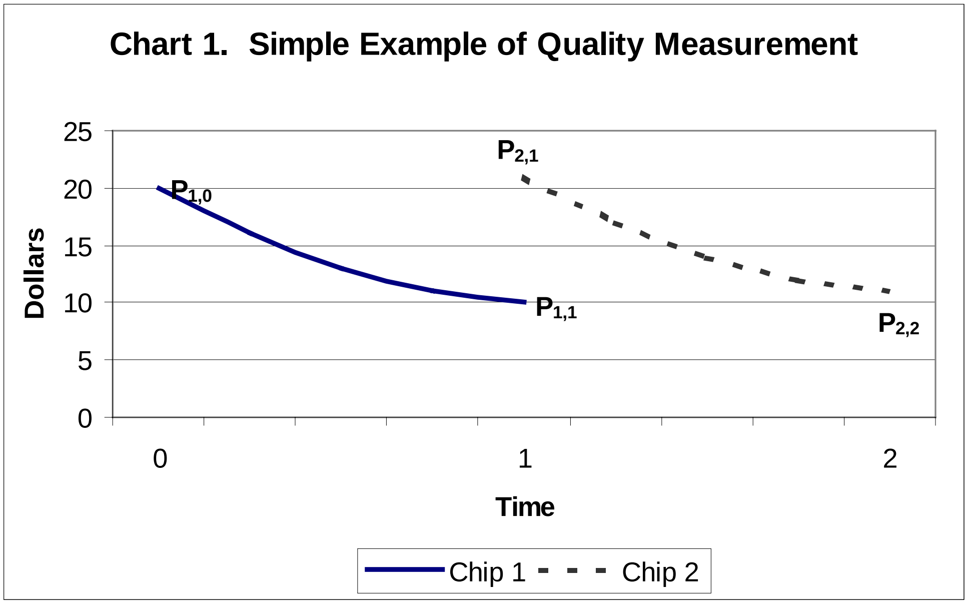
Chart 1. Simple Example of Quality Measurement
25
20
15
10
5
0
Dollars
$P_{2,1}$
$P_{1,0}$
$P_{1,1}$
$P_{2,2}$
0
1
2
Time
Chip 1
Chip 2

Chart 2. Price Profiles for Selected Intel Chips, by Speed

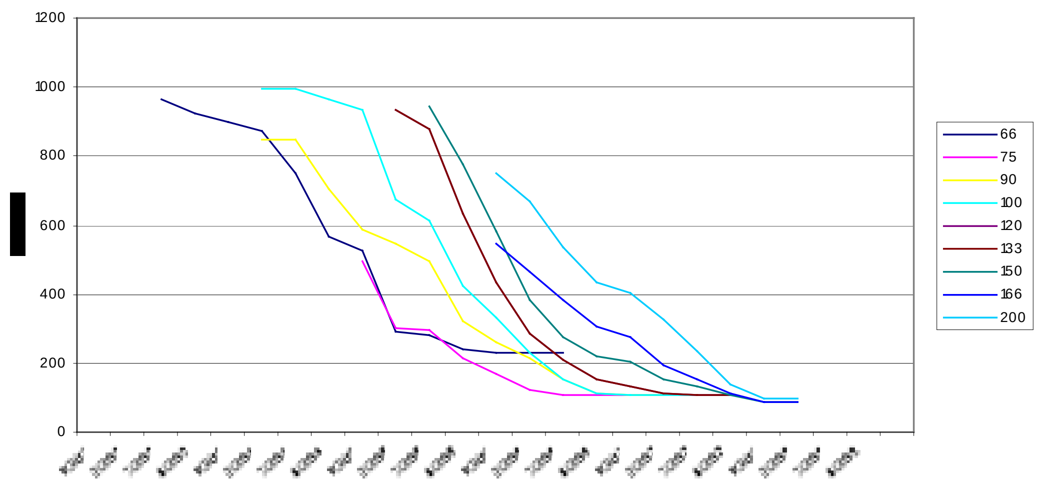

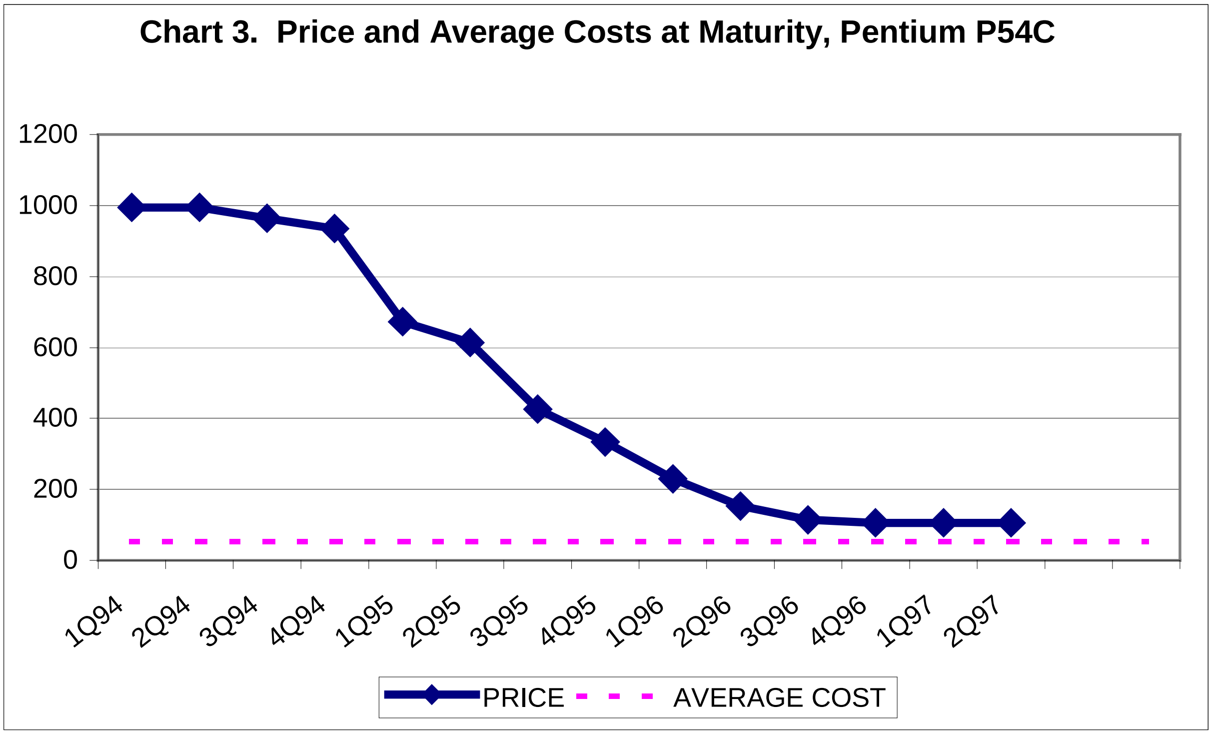
Chart 3. Price and Average Costs at Maturity, Pentium P54C
1200
1000
800
600
400
200
0
1Q94
2Q94
3Q94
4Q94
1Q95
2Q95
3Q95
4Q95
1Q96
2Q96
3Q96
4Q96
1Q97
2Q97
PRICE
AVERAGE COST