

Preliminary
Comments Welcome

Product Introductions and Productivity at Intel in the 1990s

Ana Aizcorbe*
Ana.Aizcorbe@bea.gov

*Bureau of Economic Analysis and The Brookings Institution. I wish to thank Marshall Reinsdorf for many useful conversations on this issue both informally and in his role as discussant at the ASSA meetings in San Diego and seminar participants at BEA and at Brookings for their comments.

Introduction

The semiconductor 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 at prices not much higher than those of existing chips. This industry has undoubtedly been an important driver of productivity growth as advances in semiconductors paved the way for co-invention in downstream industries that, taken together, provide firms with more efficient ways to do business and the ability to provide new goods and services that ultimately increase consumer welfare.

In the mid-1990s, measured productivity growth for the industry shows a pickup that coincided with an economy-wide pickup in labor productivity growth. The acceleration in the semiconductor market stems from an increase in the growth of real output that was, in turn, generated by what Jorgenson (2001) calls an “inflection point” in price indexes for the semiconductor industry. Within semiconductors, microprocessors (MPUs) produced by Intel were the primary contributor to both the trend and inflection point in this price index (figure 1).

This paper explores movements in the price index for MPU chips over the 1990s to better understand sources of the pickup in measured productivity growth. Section 1 reviews three major developments in MPU markets that roughly coincided with the measured increase in productivity: 1) the introduction of more sophisticated lithography equipment that could have allowed Intel to increase its rate of product innovation; 2) an increase in competitive pressure from AMD; and 3) a pickup in the rate of product introductions at Intel. Section 2 provides a stylized model of decision-making at Intel and uses it to show that the increase in the rate of product introductions at Intel could have been a profit-maximizing response to increased competition from AMD. Section 3 considers the implications of that increase in competition and the rate of product introductions for the price index for MPUs.

This paper identifies conditions under which changes in competition and introduction rates may have distorted measures of quality change in this market. For example, in matched-model indexes, the value of quality improvement in a new good is measured as the gap between the introduction price for a new good and the prices of existing goods. Because the quality of AMD chips tended to be a few quarters behind

that of Intel's chips, increased competition from AMD might have lowered the prices of (older) existing chips more than those of newly-introduced goods. If so, this could widen the gap between prices of new and existing chips and, hence, increase measured quality change for the new chip.

Although the potential problems discussed here are theoretically possible, empirical work is needed to ascertain whether the types of parameters necessary to give rise to these potential distortions actually hold in this market. Moreover, the numerical magnitude of these effects needs to be explored.

Developments in MPU Markets in the 1990s

The 1990s witnessed several innovations that had implications for the microprocessor market. The explosion of the internet that began in the early 1990s and the introduction of sophisticated software applications ranging from new operating systems to complex graphics for computer games all raised the bar on system requirements for the microprocessor chips that power the personal computer (PC). Demand for higher-quality microprocessors was also bolstered by innovations in the devices that serve as complements to microprocessors: disk drives, memory chips, and CD and DVD capabilities. The microprocessor chips themselves also underwent significant transformation over the decade; the exponential growth in the number of transistors in MPUs raised clock speeds over the decade and, beginning with the Pentium chips, the inclusion of embedded features—like cache—in the MPU also enhanced their performance.

This section reviews three developments in MPU markets that occurred in the mid-1990s that may have had implications for measured productivity.

Shortened product cycle

The development of equipment capable of etching finer circuitry is often referred to as “process innovation” and the length of time between the introduction of the new

equipment as the “product cycle.” Beginning in 1995, the industry moved from a three-year to a two-year product cycle.¹

Understanding the link between the arrival of new equipment and productivity requires a basic understanding of the production process. Semiconductor chips are produced using lithography equipment that etches circuitry on silicon wafers. The etched circuitry is then cut into individual die that, once tested, become the MPU chip. Because lithography equipment determines the size of features on chips—e.g., transistors—equipment that can etch narrower circuitry can include more features on each chip and, hence, can produce a higher quality chip. On the other hand, higher-quality chips are typically larger, *ceteris paribus*, and larger chips typically cost more to produce. This is because, for a given generation of equipment, the cost of etching a wafer is fixed; smaller chips allow more chips to fit on the wafer and reduces average costs.

A shortening of the product cycle opens up possibilities to increase the rate of product innovation. However, the link between the two is neither immediate nor direct. The translation of shortened product cycles to faster product innovation depends on 1) the rate of diffusion of the new equipment (how many plants switch to the cutting-edge process when it arrives), 2) whether the equipment is used to improve product quality or reduce average costs, and 3) the presence of any learning curves required to fully implement the new process.

It is difficult to assess whether the rate of quality improvement increased given the existing data. Data on the features of Intel’s chips that we normally associate with the quality of the chip do not exhibit a compelling inflection point to match that in the MPU price index. For example, data on the number of transistors and the speed of Intel’s cutting-edge chips appear to have remained on the same trajectory over the 1990s (Aizcorbe, Oliner and Sichel (2003)). On the other hand, it is possible that the quality improvement in these devices occurred in ways that are not immediately apparent in the data. For example, MPUs introduced in the mid 1990s included a memory chip (cache) intended to increase the performance of the chip. Similarly, the external bus of an MPU

¹ For additional details about this shift to a shorter technology cycle, see the International Technology Roadmap for Semiconductors (2001 and 2002 update)

increased over the 1990s; external bus affects the speed with which the processor and a PC's memory can communicate and, thus, affects the system's functionality.

Increased Competition from AMD

Another development in MPU markets over the 1990s was increased competition that intensified, in part, because AMD won rights to produce chips under its own brand name in a lawsuit that was settled in 1995; before then, AMD produced chips under contract for Intel. Shortly after the case was settled, AMD purchased another semiconductor firm—NextGen—and launched an attempt to establish itself as a credible alternative to Intel.

Over the second half of the decade, the rate of product innovation at AMD was sufficiently fast to narrow the gap between the quality of its chips and that of Intel's. An analysis by McKinsey and Company illustrates this point by plotting the length of time between Intel's introduction of a cutting-edge chip and AMD's introduction of a comparable chip (figure 2). As may be seen, before 1995, the lag between Intel's and AMD introduction of comparable chips was over six quarters. But, by 1998, the gap had become negligible.

Anecdotal reports of price wars and chip shortages suggest that AMD's presence affected Intel's pricing. This change in competitive conditions does appear to have had some effect on Intel's margins but that reduction in margins but does not appear to have *directly* generated the inflection point (Aizcorbe, Oliner and Sichel(2003)). Nonetheless, changes in the competitive environment could have prompted Intel to change strategies in ways that might have implications for productivity measurement and those possibilities are examined below.

Increase in the Rate of Introductions

The final development that merits analysis is a striking change in the rate of product introductions at Intel. The horizontal lines in figure 3 depict production runs for Intel's desktop chips that were introduced between 1993 and 2000. The vertical axis notes the clock speed of each chip—a rough indicator of quality.

As may be seen, chips were introduced every 6-8 quarters early in the 1990s but, by 1999 were introduced nearly every quarter. This shortening in the lag between introductions coincided with a shortening of production runs that shifted production from older chips towards younger, higher-quality, chips.

It is tempting to view this increase in the rate of introductions simply as a by-product of possible increases in the rate of innovation—if manufacturers had been able to increase the quality of chips faster, they might have brought them to market faster. However, as is argued below, this need not have been the case. The model below demonstrates that even if Intel's quality trajectory had remained unchanged over the 1990s, increased competition could have prompted it to bring new chips out sooner in order to shift production away from older chips—that are more susceptible to competition from AMD—and toward newer chips.

A Stylized Model

This section lays out a model for decision-making at Intel to explore how Intel might have responded to these developments and the effect that those actions could have had on Intel's price and productivity measures. The model is necessarily stylized. The production and sales of Intel's MPUs are governed by several key features that, even when taken one at a time, generate extremely complex models: Intel is a durable goods monopolist that faces various degrees of competition, and chooses an optimal rate of innovation to generate new (higher quality) goods. Production involves learning along several important dimensions and these learning curves act as a binding supply constraint in the short run. The complexity of these models precludes the possibility of combining existing approaches to specify a general model that allows for all of these features. Instead, this paper develops a highly stylized model that captures the essence of the landscape within which Intel operated in the 1990s.

Demand

Demand for Intel chips is ultimately determined by end users' demand for the PCs that contain the chips and the applications they power. Although the buyers of Intel chips

are literally PC makers like Dell and Compaq, demand here is specified with an eye to the end user while keeping in mind that this demand is actually revealed to Intel through PC makers' derived demand.

A specification for the demand side of the model must be consistent with the downward-sloping price contours typically exhibited by these devices.² As seen in figure 4, introduction prices for MPUs begin around \$600-1000 and fall almost immediately. One way to explain the curvature of these price profiles is that Intel maximizes profits through intertemporal price discrimination.³ That is, Intel introduces chips at a high price and sells to those buyers most willing to pay the high price. Once Intel has sated that segment of the market, it lowers the price and sells to buyers willing to pay the next-highest price, etc.

With this type of price-skimming behavior, recently-produced chips would be purchased by high tech consumers that need to use cutting-edge applications and are willing to pay top dollar for the latest devices. As a chip ages, purchases are increasingly made by low-tech consumers who are content to be a few generations behind because they use mostly older applications. The important distinction between these high- and low-tech consumers is that low tech buyers are assumed more flexible and more susceptible to both inter-temporal and contemporaneous substitution possibilities. (Parker (1992)).

Each segment has its own implicit demand function that captures factors affecting those consumers' willingness to pay for chips. The implicit demands are potentially functions of the same arguments—high- and low-tech buyers view the same goods as substitutes—but their elasticities differ. Each implicit demand for chip c of age a is expressed as a function of number of chips sold ($Y_{c,a}$), the quality of the chip (q_c), and the quality and price of all available substitutes. The contemporaneous substitutes are yesterday's cutting-edge chip and chips produced by the firm's competitors. The model

² See Flamm(1996) and Irwin and Klenow(1994) for evidence on DRAMs and Aizcorbe(2002) for evidence on MPUs.

³ There is a theoretical argument, known as the Coase conjecture, which states conditions under which a durable-goods monopolist will tend to expand production until price equals marginal cost. As discussed below, an important feature of semiconductor production is that output is constrained by steep learning curves. This type of supply constraints violates one of the conditions required for the Coase conjecture to hold and makes intertemporal price discrimination a theoretical possibility for firms like Intel.

assumes that Intel's only competitor (AMD) produces one chip with quality q_{AMD} and price P_{AMD} and that Intel takes the AMD prices and quantities as given. There is also an inter-temporal dimension to demand in that buyers' willingness to pay for today's chip also depends on the expected quality and introductory price of tomorrow's cutting-edge chip (q_{c+1} and P_{c+1}) and, importantly, on the amount of waiting time until that chip is introduced (W). That waiting time is a function of the age of the chip ($W(a)$), where age is defined as the time that has elapsed since the chip was introduced. Finally, demand can also be affected by exogenous factors (X) like the price and quality of complements (like software) or network effects that affect the price buyers are willing to pay for chips (like the evolution of the internet).

The implicit demand functions are written:

$$(1) \quad P_{c,a} = P(q_{c-1}, P_{c-1,0}; q_c, Y_{c,a}; q_{c+1}, P_{c+1,0}; q_{AMD}, P_{AMD}; W(a), X)$$

Within the each segment, two factors drive down prices between chip introductions. First, waiting times decline as the new introduction nears; those reductions in waiting time increase the substitutability of today's chip with tomorrow's chip and push down current prices (i.e., $dP_{c,a}/dW < 0$). Similarly, as a chip ages, the AMD chip becomes a closer substitute—although the quality of the Intel chip is constant over the life of the chip, the quality of AMD chips continue to climb as the Intel chip ages. As the Intel chip ages, this gap narrows and that puts downward pressure on the prices of today's chips.

Evolution of Quality

The quality of a microprocessor is a function of the design of the chip, the sophistication of the equipment that produces it and the firm's familiarity with the equipment. At a given point in time, the first two factors are the result of past investment in R&D and equipment. The third—increases in the quality of chips over the life of the equipment—is governed by a learning curve.

The model developed below is a model for short-run decisions—those that do not involve investment—and focuses on production and pricing decisions that are made while

using a particular generation of equipment to produce a particular chip (like the Pentium III). The significance of the equipment for the quality of chips is that it places a ceiling on the quality of chips that it can produce. The firm inches towards that cap as it gains familiarity with the equipment. Because the manufacturing process is extremely complex—production of a chip involves hundreds of steps—it necessarily involves defects. The learning curve involves engineering and managerial activities that discover and remove defects and these activities are the main force behind improvements in the quality of chips produced under a given process (Hatch and Mowery(1998)).

Formally, the highest quality level that can be achieved at a point in time (q_t) increases with learning and is modeled as a linear quality trajectory that lies everywhere below the maximum quality the existing process can deliver:

$$(2) \quad q_t = q_{t-1} + \alpha \quad \text{if} \quad q_c < q^{\max}$$

Although this level of quality increases continuously, firms typically increase quality in discrete steps with the introduction of new goods. In what follows, it will be useful to keep track of time in terms of the number of periods between chip introductions (the introduction lag). The quality trajectory is thus rewritten as:

$$(2') \quad q_c = q_{c-1} + \alpha I$$

where the quality of chip c is equal to that of its predecessor (chip $c-1$) plus some function of the number of periods between their introduction (written here as a scalar multiple of I). Given α , the introduction lag determines the quality level assigned to each introduced chip.

It is assumed that AMD has its own quality trajectory that, for simplicity, lags behind Intel's trajectory by L periods: $q_{\text{AMD}} = q_c - L$ and that AMD is able to narrow the gap by reducing L .

Output

With regard to production levels, the maximum number of chips that can be produced on a given process, Y^{MAX} , is assumed exogenous; it is a function of the sizes of silicon wafers and the chips etched on them, both of which are fixed in the short run (since changing either of these requires investment). At any point in time, some of the manufactured chips will turn out to be defective and will be discarded. The fraction of chips that are usable (called the yield) increases over the life of the process as defects that render chips unusable are removed. This is an important issue for chipmakers; during the years that this type of learning curve takes place, yields (the ratio of usable chips to total) can increase from as low as [10] percent when the new process is launched to [60] percent. In this model, the yield is modeled as a function of time at the introduction of the chip $\lambda(cI)$ and the level of usable output is written $\lambda(cI)Y^{\text{MAX}}$.

The quality of a chip is determined in one of the last steps of the production process, where chips are tested under a variety of conditions (temperature, air pressure, etc.) to determine their functionality. The presence of defects is such that the quality of chips produced by a given process at a point in time is actually a distribution rather than a specific level. For example, the Pentium III 750, 866 and 1000 MHz chips were generated from the same wafers, using the same process but the chips tested at different levels (a process called 'binning').

Thus, firms must choose how many grades of chips to bring to the market—how many bins to sort chips into—and how long to market each chip. To keep the model simple, it is assumed that the firm chooses to market only two types of chips at any given period. Those chips that test at the maximum speed (q_c) are sold to high-tech buyers as "cutting-edge" chips; it is assumed that all others successfully test at the quality of yesterday's cutting edge chip and are sold to low-tech buyers as such. As time passes, the quality distribution edges up and the number of cutting-edge chips that can be produced increases.

Output levels for the two types of chips are modeled as shares of the number of usable chips. It is assumed that the share of chips that test as "cutting-edge" chips, η , is a function of the number of periods since introduction, or the age of the chip (a). The

number of cutting edge chips of age a that are available for sale may then be written as a product of this share and the number of usable chips:

$$(3) \quad Y_{c,a} = \eta(a) \lambda(cI) Y^{\text{MAX}}$$

Output of the low-tech chip is written as $Y_{c-1,a+I}$, where $c-1$ denotes that these chips are sold as yesterday's cutting edge chip and $a+I$ tracks the number of periods since those chips were introduced. The number of these chips brought to market is equal to the number of usable chips that did not test as a cutting-edge chip:

$$(4) \quad Y_{c-1,a+I} = \lambda(cI) Y^{\text{MAX}} - Y_{c,a} = [1 - \eta(a)] \lambda(cI) Y^{\text{MAX}}$$

This lines up with what one sees in the data. Output is typically low at introduction, rises over the life of the chip and then tails off as new chips are introduced.

This implicitly assumes that the firm produces at capacity, an assumption that is consistent with the anecdotal evidence. The enormous setup costs for production, the rapid obsolescence of the devices, and the learning curves that constrain production all place incentives for the plant to run at maximum capacity once the equipment is in place (Flamm(1996)).

Over the M periods that Intel uses a particular process, the introduction lag (I) implies the number of chip introductions (M/I). Given the assumption that Intel produces at capacity, the introduction lag also implies how long chips are marketed. Since we've assumed that Intel serves two segments (markets two types of chips each period), chips live $2I$ periods: a chip is initially sold as a high tech chips for I periods before it is eclipsed by the next cutting-edge chip; once that occurs, the chip is sold as a low tech chip for another I periods.

First Order Conditions

Intel times its introductions so as to maximize profits over the M periods it uses a given set of equipment. Specifically, Intel chooses I to maximize expected revenues from

each segment less a per-period manufacturing cost, θ , and an introduction cost for each introduced chip, ι :

$$(5) \quad \text{Max}_I \sum_{c=1, M/I} \sum_{a=1, 2I} \text{REV}_{c,a} - \theta - \iota (M/I)$$

where $\text{REV}_{c,a}$ is the revenue generated by each chip: $\text{REV}_{c,a} = P_{c,a} Y_{c,a}$. The distinction between high- and low-tech consumers is deferred until the effect of competition on price indexes is examined (at which point it becomes important).

To motivate the first order conditions, figure 5 illustrates the tradeoffs involved in changing the introduction lag. The chart compares two scenarios, both assume that Intel is using a particular process for 8 periods ($M=8$) and serves two segments. The solid lines depict the first scenario, where Intel introduces a new chip every period. This allows it to make 8 introductions and produce each chip for two periods before switching to a new chip. The dashed lines depict an alternative scenario where there is a lag between introductions of 2 periods. This allows Intel to introduce only 4 chips and produce each one for 4 periods.

There are two types of effects from increasing the introduction lag. For what follows, it is useful to think of the segments that both scenarios have in common separate from the others; in this simple example, both scenarios include at least 4 chip introductions and chip lives of at least 2 periods. Over these common segments, buyers' willingness to pay for all chips increases both because increases in I increase the waiting time (which increases the amount they're willing to pay for today's chips) and also because increases in I increases the quality assigned to each chip. The second type of effect trades periods where newer chips are marketed for those that involve older chips: increasing the introduction lag from $I=1$ to $I=2$, Intel loses the revenue from the 4 chips it no longer introduces but gains revenues from extending the production run for each introduced chip to 4 periods (from 2). Thus, increasing the introduction lag shifts the distribution of production towards older chips.

Formally, the first order conditions say that Intel should increase the introduction lag until these effects just offset one another. Let $C^*=M/I^*$ and $A^*=NI^*$ be the number of

chip introductions and length of a chip's market life at the optimal introduction lag and C^0 and A^0 be the counterparts for an introduction lag one period less than the optimal. The first order conditions may be stated as:

$$(6) \quad \sum_{c=1, C^*} \sum_{a=1, A^0} \delta \text{REV}_{c,a}(I^0) / \delta I \\ + \sum_{c=1, C^*} \sum_{a=A^0+1, A^*} \text{REV}_{c,a}(I^*) \\ - \sum_{c=C^*+1, C^0} \sum_{a=1, A^0} \text{REV}_{c,a}(I^0) - \tau (C^0 - C^*) = 0$$

The first term gives the marginal revenue from the segments that are common to the two scenarios: both introduce at least C^* chips and market each chip for at least A^0 periods. The revenue from these common segments increases when I is raised from I^0 to I^* as the result of three effects. First, the increase in the introduction lag pushes introductions out in time when the quality of chips is higher and the price that buyers are willing to pay is also higher.⁴ Second, output levels are also higher when I increases because chips are introduced later when yields are higher. Finally, increases in I increases the waiting time until the next cutting-edge chip becomes available. Because this reduces intertemporal substitution possibilities, the price buyers are willing to pay for today's chips increases.

The last two summations in (10) capture the effect on revenues from the reshuffling towards older chips that occurs when I increases to I^* . The second term considers the additional revenue generated by extending each chips life by $A^* - A^0$ periods while the third term measures the loss in revenues that occurs when the number of chip introductions are cut back by $C^0 - C^*$. The last term measures the change in introduction costs.

⁴ Although the quality of the future chips is also higher, those effects are discounted and, thus, lower than the own-quality effect,

Competition from AMD⁵

A first look at how the optimal introduction lag changes in response to changes in competitive conditions assumes that output is constant both across chip introductions and as a chip ages: $Y_{ca} = \Psi$ for all c and a . Under that assumption, an increase in q_{AMD} prompts Intel to shorten the product cycle. Intuitively, this occurs because increased competition from AMD reduces the revenue Intel can earn from the older chips (they are closer substitutes to AMD's chips) and increases the incentive to produce newer chips. To see this formally, consider a version of the first order condition in (6) that holds output constant across chips:

$$(6') \quad \sum_{c=1, C^*} \sum_{a=1, A^0} (\delta P_{c,a}(\Gamma^0) / \delta I) \Psi \\ + \sum_{c=1, C^*} \sum_{a=A^0+1, A^*} P_{c,a}(\Gamma^0) \Psi \\ - \sum_{c=C^*+1, C^0} \sum_{a=1, A^0} P_{c,a}(\Gamma^0) \Psi = 0$$

The effect of an increase in competition on the timing of introductions will depend on the sign of the derivative of this first order condition with respect to q_{AMD} . Assuming that the sensitivity of demand to changes in the timing of introductions is not affected by competitive conditions (i.e., $\delta (\delta P_{c,a} / \delta I) / \delta q_{AMD} = 0$). The tradeoff boils down to whether the increase in revenues from extending chips' lives is affected more than the loss in revenues from fewer chip introductions:

$$(7) \quad \sum_{c=1, C^*} \sum_{a=A^0+1, A^*} \delta P_{c,a}(\Gamma^*) / \delta q_{AMD} \Psi \\ - \sum_{c=C^*+1, C^0} \sum_{a=1, A^0} \delta P_{c,a}(\Gamma^0) / \delta q_{AMD} \Psi$$

The terms in the summations are negative because an increase in the quality of AMD's chips reduces the prices that buyers are willing to pay for Intel's chips. But, this price

⁵ A complete analysis would include an investigation of other factors that might have prompted Intel to increase introduction rates. Among those possibilities are other exogenous shocks in demand (such as the introduction of complementary software), increases in production capacity, changes in the parameters of the yield curve, and the like. The effect of competition is examined first because the change in market structure over the 1990s is widely thought to be important whereas the importance of changes in the other factors are not immediately apparent.

effect is assumed greater on older chips (the first term) than that on relatively newer chips. Therefore, the negative effect on older chips outweighs the reduced losses on the newer chips and the expression in (7) is negative. The firm returns to equilibrium by shortening the market lives of chips (marketing newer chips) and it does that by reducing the period of time between introductions.

In this simple version of the model, then, an increase in competition prompts the firm to increase the rate of product introductions. Under more realistic depictions of the production process, this result becomes ambiguous. To see this, consider a version of (6) that defines output as in (3) and (4): that is, $Y_{c,a} = \eta(a) \lambda(cI) Y^{\text{MAX}}$ and $Y_{c-1,a+1} = Y^{\text{MAX}} - Y_{c,a}$. This creates two types of complications. First, the effect on revenue from changing the introduction lag now involves an output effect in addition to the price effects considered above: because yields increase over time, an increase in the introduction lag means that, ceteris paribus, each chip is produced later in time and, hence, at a higher level of output. The second type of complication deals with how output varies over the life of a chip. For a given cap on production (Y^{MAX}), the output of high-tech chips ($Y_{c,a}$) increases as the chip ages while the output of low-tech chips decreases. This makes it difficult to compare output levels across chips of different ages and, hence, to assess the effects of changes in the exogenous variables.

The effect of increased competition on introduction rates will, in general, depend on how fast output increases across chip introductions versus over the lives of chips. To see this, consider the following version of the first order condition:

$$(6') \quad \sum_{c=1,C^*} \sum_{a=1,A^0} [(\delta P_{c,a}(I^0)/\delta I) Y_{c,a}^0 + P_{c,a}(Y_{c,a}(I^*) - Y_{c,a}(I^0))] \\ + \sum_{c=1,C^*} \sum_{a=A^0+1,A^*} P_{c,a}(I^*) Y_{c,a}(I^*) \\ - \sum_{c=C^*+1,C^0} \sum_{a=1,A^0} P_{c,a}(I^0) Y_{c,a}(I^0) = 0$$

As before, the terms in the first summation represent the change in revenue over the common segments, the terms in the second summation represent the increase in revenue from extending chips' market lives and the last terms are the lost revenue from introducing fewer chips.

The effect of an increase in competition on the timing of introductions depends on the sign of the derivative of the first order condition with respect to q_{amd} . Again assuming that the sensitivity of demand to changes in the timing of introductions is not affected by competitive conditions (i.e., $\delta (\delta P_{c,a}/\delta I) / \delta q_{amd} = 0$), that expression is:

$$(7') \quad \sum_{c=1,C^*} \sum_{a=1,A^0} [(\delta P_{c,a}(I^0)/\delta q_{AMD}) (Y_{c,a}(I^*) - Y_{c,a}(I^0))] \\ + \sum_{c=1,C^*} \sum_{a=A^0+1,A^*} [(\delta P_{c,a}(I^*)/\delta q_{AMD}) Y_{c,a}(I^*)] \\ - \sum_{c=C^*+1,C} \sum_{a=1,A^0} [(\delta P_{c,a}(I^0)/\delta q_{AMD}) Y_{c,a}(I^0)]$$

As before, the three terms in brackets are negative because of the negative impact on prices from increased competition. In this case, however, an increase in competition prompts a shortening of introduction lags when:

$$(8) \quad \sum_{c=1,C^*} \sum_{a=1,A^0} [(\delta P_{c,a}(I^0)/\delta q_{AMD}) (Y_{c,a}(I^*) - Y_{c,a}(I^0))] \\ + \sum_{c=1,C^*} \sum_{a=A^0+1,A^*} [(\delta P_{c,a}(I^*)/\delta q_{AMD}) Y_{c,a}(I^*)] \\ < \sum_{c=C^*+1,C} \sum_{a=1,A^0} [(\delta P_{c,a}(I^0)/\delta q_{AMD}) Y_{c,a}(I^0)]$$

and this will depend not just on the responsiveness of prices to competition, but also on the magnitude of output.

Assuming that the responsiveness of prices to increased competition does not depend on the length of the introduction lag (i.e., $\delta P_{c,a}(I^0)/\delta q_{AMD} = \delta P_{c,a}(I^*)/\delta q_{AMD}$ for all c,a), this expression may be simplified. In particular, a firm will reduce the introduction lag in response to competition when the drop in revenue from increased competition is greater (a bigger negative) under introduction lags of I^* than under a shorter introduction lag (like I^0):

$$(9) \quad \sum_{c=1,C^*} \sum_{a=1,A^*} (\delta P_{c,a}(I^*)/\delta q_{AMD}) Y_{c,a}(I^*) \\ < \sum_{c=1,C} \sum_{a=1,A^0} (\delta P_{c,a}(I^0)/\delta q_{AMD}) Y_{c,a}(I^0)$$

The firm will choose the introduction lag that best protects its sales from competition. If Intel's output was distributed such that the longer introduction lags left it more vulnerable to competition, the profit-maximizing response would have been to shorten the introduction lags.

Implications for Price Indexes

The primary function of price indexes is to separate out changes in average prices that are due to "pure price change" from those that are due to changes in goods' quality. For MPUs, the average price over all chips was fairly flat over the 1990s so that the rapid price declines seen in price indexes were mostly generated by measured quality increases.

In what follows, the focus is on matched-model techniques only because the price index in figure 1 is a matched-model index.^{6,7} Matched-model indexes that are constructed using panel data measure quality change by assuming that price differences at a point in time reflect the market's valuation of differences in quality across the goods. In this method, models are assumed homogeneous over time so that quality change occurs only when new goods are introduced. Conceptually, the gap between the introductory price and existing prices is the method's estimate of the value of quality improvement in the new chip. Mechanically, these indexes boil down to a weighted average of the price changes for models that exist in the two periods. As such, the steepness of the index depends on the steepness of the price contours.

One problem with this method is that differences in prices at a point in time can reflect the presence of different consumers that sort into purchasing different goods.⁸ The issue is that the matched-model method assigns the gap between introduction prices and

⁶ The index is actually the result of splicing two indexes. The first index used data from DATAQUEST from 1974-1996 to construct matched-model Fisher indexes that used hedonic regressions to impute missing prices (Grimm(1998)). The second index used data from MicroDesign Resources from 1994-2002 to construct matched-model indexes (Aizcorbe(2002)). Using available data from DATAQUEST, an extension of Grimm's index yields very similar measures that also show an inflection point.

⁷ In this context, "matched-model techniques" refers to the construction of indexes using near-universe data sets. The only similarity between this and the techniques applied by statistical agencies like the BLS is that both assume that differences in prices across goods at a point in time may be viewed as the market's valuation of differences in the characteristics of those goods.

⁸ The possibility that the gap between introductory and existing prices may not yield an accurate reflection of quality change has been recently discussed in Kokoski, et. al.(2000). In their case, producers used model introductions as an opportunity to introduce pure price change, hence potentially distorting the quality

existing prices to quality differences when, in fact, buyers that pay the introductory price and those that pay the existing prices may be different people.

This section considers conditions under which the increase in competition that Intel faced over the 1990s might have distorted the measure of quality change implicit in matched-model indexes. The increase in competition has both a direct effect on Intel's prices and indirect effects through potential increases in the rate of introductions.

Direct Effects

The direct effect of an increase in competition is a fall in the prices of Intel's chips. If competition from AMD affected the high-tech and low-tech segments equally, then the increase in competition would push down the price contours without altering the steepness of the contours. In that case, there would only be a transitory effect on the price indexes, as the price contours move from the pre-competition to post-competition levels.

However, if buyers in the two segments respond differently to increases in competition, then the estimate for quality change might be affected.⁹ For example, suppose that only the low-tech buyers view AMD chips as substitutes.¹⁰ In that case, increased competition would push down prices of the low-tech chips without affecting those of high-tech chips and would widen the gap between the (unchanged) introductory prices and those of existing chips. In this case, increased competition would generate a larger estimate for quality change and steeper declines in a constant-quality price index.

Indirect Effects

The indirect effect of increased competition arises if Intel responds to changes in competition by increasing the rate of product introductions. The presence of shorter introduction lags implies shorter waiting times until new chips arrive. To the extent that buyers view this waiting time (W) as a negative attribute of soon-to-arrive chips, declines

change measure.

⁹ It would be interesting to explore how this argument plays out in the context of specific functional forms for demand that allow for heterogeneity but are consistent with indexes based on representative consumers (Feenstra(1995); Hobijn(2000)).

¹⁰ This might be a reasonable assumption for those years where the quality of AMD's chips lagged well behind that of Intel's cutting-edge chip.

in W would have the same effect as an increase in quality: prices of today's chips will be driven down as the future chip looks more attractive.

That effect could be amplified if low-tech users react more to changes in W than high-tech users. One could argue that low-tech buyers are more susceptible to substitution possibilities so that changes in the wait time would generate a bigger price decline for low-tech chips than for high-tech chips. If so, this could impart a distortion on matched model indexes.

Summary and Future Work

This paper explored the possibility increased competition in MPU markets over the latter half of the 1990s may have contributed to the observed inflection point in price indexes for those devices. Two types of effects were considered. First, the direct effect of increased competition is to lower prices. However, if markets are segmented in that heterogeneous consumers respond differently to demand shocks, it is possible that changes in competitive conditions could affect measures of quality change. One possibility occurs when the market is made up of high-tech and low-tech consumers and the competition is felt mostly by the low-tech consumers. In that case, prices of existing (older) chips fall but introductory prices do not so that the index will attribute a greater amount of quality improvement to the new good.

A potential solution to this problem would be to construct separate indexes for the different segments in the market. The main problem in implementing this solution is in defining the segments. This is a difficult problem that economists in the antitrust area have grappled with for years. There, antitrust cases often hinges on measures of market power that, in turn, depend on how broadly the market is defined. An alternative to defining market segments would be to estimate price change using hedonic regressions that control not just for the different attributes for the goods but also for the different attributes of consumers. This framework would allow one to construct price indexes that control for consumer heterogeneity.

The second potential effect of increased competition that was considered arises if competition prompts the firm to increase the rate of product introductions. An increase in

the rate of product introductions reduces the amount of time consumers have to wait for the arrival of tomorrow's cutting-edge chip. If consumers view the waiting time until the next cutting-edge chip arrives as a negative attribute of the future chip, then the shorter wait times would increase intertemporal substitution possibilities and could put downward pressure on the prices of today's chips. Another way to describe this effect is that if waiting time is a negative attribute for tomorrow's chip, then a decrease in the waiting time has the same effect as an increase in quality.

These issues raise questions about the ability of standard price measures to properly account for quality change. Theoretical work is needed to develop methods that allow one to control for the presence of heterogeneous consumers. Empirical work is also needed to ascertain whether the types of parameters necessary to give rise to these potential distortions actually hold in this market. Moreover, the numerical magnitude of these effects needs to be explored.

References

- Aizcorbe, Ana (2002). "Why Are Semiconductor Prices Falling So Fast? Industry Estimates and Implications for Productivity Measurement," Finance and Economics Discussion paper 2002-20, Federal Reserve Board.
- Coase, R. ()
- Feenstra, R. C. (1995) "Exact Hedonic Price Indexes," *Review of Economics and Statistics*, 77(4):634-53, November.
- Flamm, Kenneth (1996). *Mismanaged Trade? Strategic Policy and the Semiconductor Industry*, Washington D.C.: Brookings Institution.
- Flamm, Kenneth (2003). "Moore's Law and the Economics of Semiconductor Price Trends," unpublished paper, University of Texas at Austin.
- Gwennap, L and M. Thomsen (1998) "Intel Microprocessor Forecast, 4th ed." Sebastopol, CA: MicroDesign Resources, Inc.
- Hatch, N. and D. C. Mowery (1998). "Process Innovation and Learning by Doing in Semiconductor Manufacturing," *Management Science*, 44, 1461-77.
- Hobijn, B. (2000) "On both sides of the quality bias in price indexes," No 157 in Staff Reports from Federal Reserve Bank of New York
- "International Technology Roadmap for Semiconductors," (2001 Edition and 2002 Update), jointly sponsored by the Semiconductor Industry Association and similar association from other producer countries.
- 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, Dale W. (2001). "Information Technology and the U.S. Economy," *American Economic Review* 91, March: 1-32.
- Kokoski, M., K. Waehrer and P. Rozaklis (2000) "Using Hedonic Methods for Quality Adjustment in the CPI: The Consumer Audio Products Component" BLS Working Paper No. 344, Office of Prices and Living Conditions, March.
- McKinsey Global Institute (2001). "US Productivity Growth 1995-2000: Understanding the Contribution of Information Technology Relative to Other Factors," Washington, DC: McKinsey & Company.

Parker, P. (1992) "Price Elasticity Dynamics Over the Adoption Life Cycle," *Journal of Marketing Research*, August, Pp. 358-367

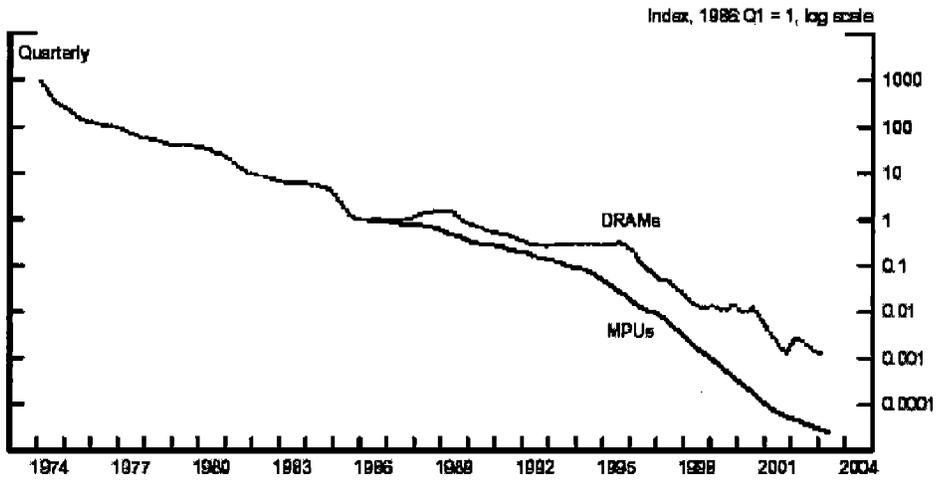
Silver, M, Ioannidis, C and Webb, B - Using Scanner Data to Estimate Quality-Adjusted Price Changes: An Application to VCR's, *Scottish Journal of Political Economy*, 48(1) (2001) 48-68, ISSN 0036-9292.

Figure 1
Semiconductor Prices

Aggregate Index



MPUs and DRAMs



Source: Aizcorbe, Oliner and Sichel (2003)

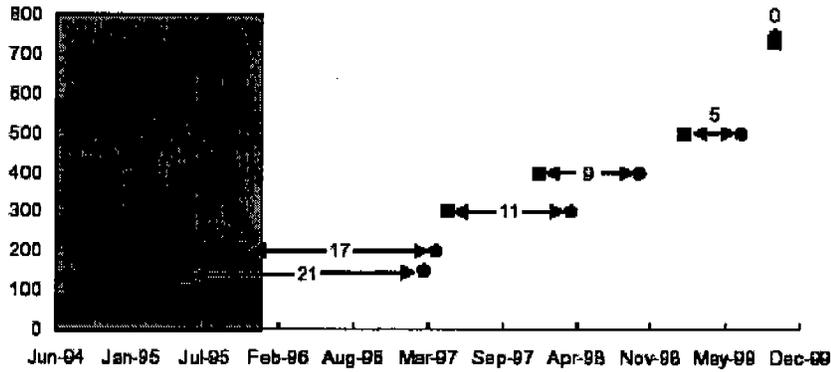
Figure 2
 Competition from AMD, 1994-1999

Exhibit 13

**INTEL FACED AN INCREASING
 COMPETITIVE THREAT FROM AMD**

Time between comparable Intel and AMD chip introductions*
 Months

MHz



Source: McKinsey and Company

Figure 3
Product Introductions of Intel Desktop MPUs, 1993-2002

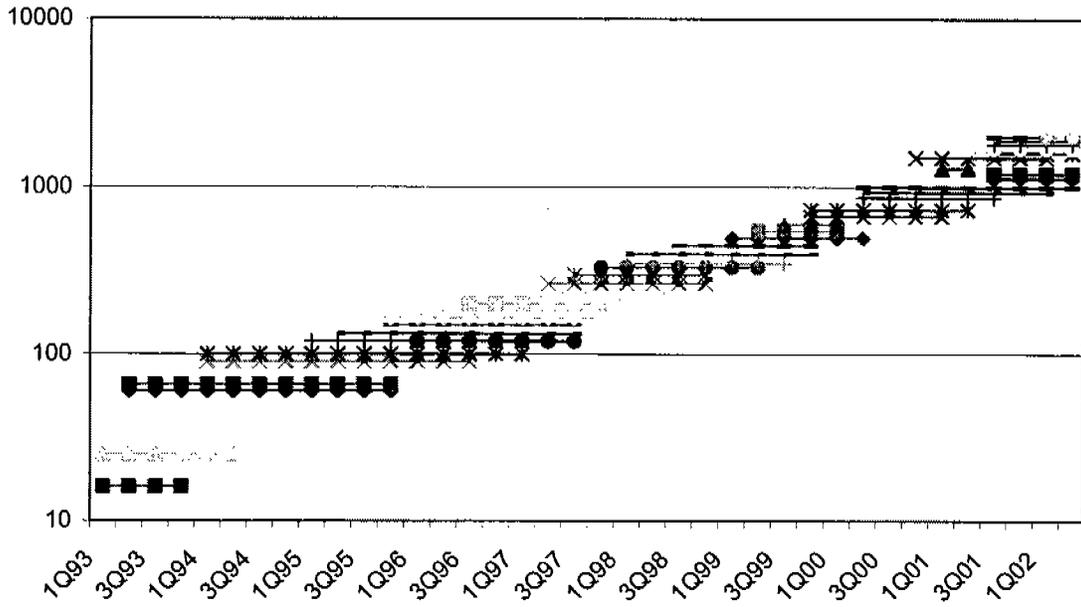


Figure 4
Price Contours for Intel Desktop MPUs, 1993-1998

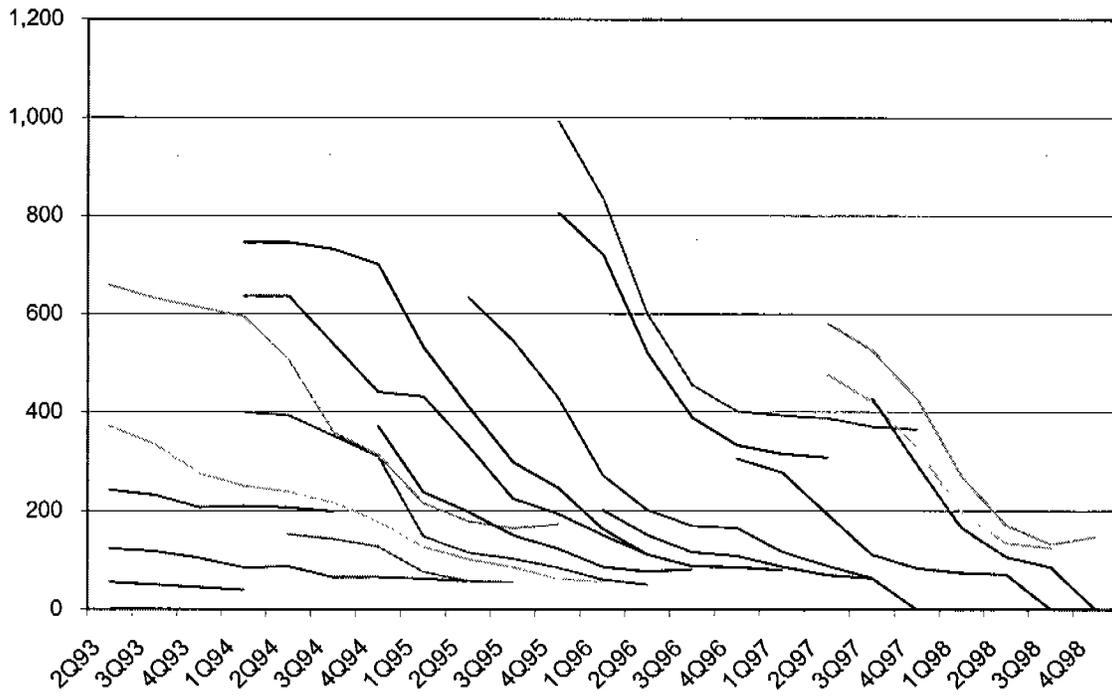
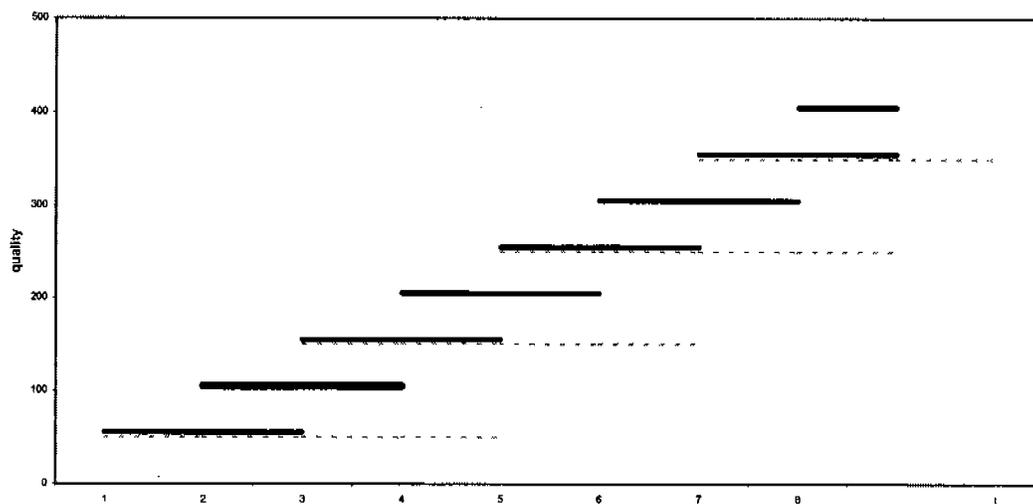


Figure 5

Product Introduction Scenarios: $I=1$ vs $I=2$
($M=8, N=2$)



Legend _____ market lives when $I=1$
 - - - - - market lives when $I=2$