Offshoring and Price Measurement in the Semiconductor Industry*

David Byrne
Federal Reserve Board

Brian K. Kovak†
University of Michigan

Ryan Michaels
University of Michigan

October 30, 2009

Abstract

The recent growth in offshore outsourcing of intermediate input production makes it especially critical that statistical agencies are able to accurately measure quality-adjusted trade flows. This paper focuses in particular on the implications of global production sharing for measuring the price of semiconductors, a critical input to high-end domestic manufacturing and U.S. productivity growth. We analyze new transaction-level data on semiconductor wafer fabrication around the world, including prices and detailed information on key physical attributes of semiconductor wafers. We estimate that, after adjusting for changes in product characteristics, the average annual price decline in processed wafers was roughly 12.5 percent during the last five years. We also find that shifts in the location of production to lower-cost countries can contribute an additional price decline of up to 0.8 percent per year.

* Paper prepared for the conference "Measurement Issues Arising from the Growth of Globalization," sponsored by the W.E. Upjohn Institute for Employment Research and the National Academy of Public Administration. We would like to thank Susan Houseman for helpful discussions and Chelsea Boone at GSA for all of her help with data and background information on the industry. Remaining errors are our own. The views expressed do not necessarily represent the views of the Federal Reserve Board.
† Brian K. Kovak, Ph.D. Candidate, Department of Economics, University of Michigan. Email: bkovak@umich.edu

1 Introduction

The recent growth in offshore outsourcing of intermediate input production has generated concern that standard government data collection methods are ill suited to an increasingly international productive structure (Houseman 2007). This paper focuses on the semiconductor industry to directly measure the effects of offshore outsourcing on input price measurement. We find that offshoring in this industry necessitates the collection of very detailed product data to adequately adjust prices for input quality, and that shifting sourcing patterns may cause standard price measures to understate price declines for processed semiconductor wafer inputs by as much as 0.8% per year.[1]

We choose to examine wafer fabrication, an intermediate stage in semiconductor production, for a number of reasons. First, semiconductor production has moved offshore to a dramatic degree in the last forty years, with continual shifts in the geographic distribution of semiconductor manufacturing capacity. Second, China's entrance in the semiconductor manufacturing market in 2001 was much heralded in the media, and provides an interesting case study on the effects of growing Chinese economic strength on an important industry. Third, the discrete nature of technological progress in semiconductor fabrication techniques makes careful quality adjustment feasible, as we describe in detail below. Finally, we have obtained a new dataset of semiconductor input prices with information on country of origin, making possible an empirical investigation of the effects of offshoring on input price measurement.

Offshoring poses a number of challenges for price measurement in the semiconductor manufacturing sector in particular. First, suppose a U.S.-based manufacturer contracts out all production to a firm overseas and that, prior to its decision to offshore, it had purchased final goods from an independent supplier here in the U.S. or had made the good itself. The one-time decline in the price level associated with the decision to offshore is not captured by current data-collection procedures. The Producer Price Index's universe does not include imports, so it does not reflect the price reduction. The Bureau of Labor Statistics (BLS) International Price Program (IPP) will measure price changes beginning in the *second* month in which the imported good is

[1] Semiconductor wafers are described in detail in Section 2.

observed, as it is not designed to measure the initial price decline that occurs when a domestic producer first off-shores a segment of production. A similar problem can arise if the firm has already contracted out production overseas but now sources from a low-cost supplier in China rather than from a producer in Taiwan.[2]

The problem posed by shifting sourcing arrangements is essentially equivalent to the problem of outlet substitution bias in the CPI, described in detail by the Boskin Commission Report (1996) and Diewert (1998). While these studies were concerned with consumers shifting their consumption toward low-cost retail outlets, this paper confronts the problem of semiconductor producers shifting their intermediate input sourcing toward low-cost suppliers located internationally. The bias is most acute whenever the inputs, as in our case, are approximately identical, which implies that the unmeasured price change when production is shifted to a new location does in fact represent a genuine price decline for the same good.

The final significant challenge is quality adjustment. As a greater share of production is shifted abroad, the composition of imports becomes increasingly sophisticated. This is particularly true within the semiconductor industry, which imports many complex intermediate inputs at various stages in the production process. This process places much greater demands on quality adjustment procedures for import prices, as semiconductor technology changes so quickly. The challenge of quality adjustment in the semiconductor industry is well known and has been demonstrated in many previous studies.[3]

We address these concerns using new transaction-level data on semiconductor wafer purchases, collected by the Global Semiconductor Alliance (GSA). These data contain fine detail on product characteristics, allowing us to generate constant-quality price indexes. They also report the source country for each transaction, making it possible to examine the effects of shifting geographic production on price measurement. Our results demonstrate the importance of having

[2] In principle, the IPP would measure this change if the manufacturer imports the good itself or if it continues to work through the same intermediary that is surveyed by IPP. If, on the other hand, the manufacturer contracts with a different intermediary in order to access a new market overseas, the IPP will miss the price decline since it surveys the importer, which in this case was the original intermediary. Unfortunately, to the best of our knowledge, there is little information on the relative importance of intermediaries in the IPP.

[3] See, among others, Flamm (1993), Grimm (1998), Doms, Aizcorbe, and Corrado (2003)

such detailed data when constructing price indexes in industries with large amounts of offshoring. This necessity is likely to increase as more countries move up the technical ladder and begin exporting ever more complex products.

The paper proceeds as follows. Section 2 describes aspects of the semiconductor manufacturing process that are relevant to price measurement and the dramatic shift to offshore production. Section 3 describes the data we utilize to build input price measures. Descriptive results demonstrate the importance of controlling for process technology and reveal substantial price differences across countries and shifting production toward lower cost locations. Section 4 presents our price index calculations. We begin with a standard matched model index as a baseline and then follow Reinsdorf (1993) to bound the potential effect of outlet substitution bias due to shifting input sourcing across countries. This section concludes with comparisons to a hedonic index and a publicly available official semiconductor price index. Section 5 concludes.

2 Semiconductor Production

This section describes the semiconductor manufacturing process and recent changes in the business models employed by semiconductor firms, highlighting characteristics of the industry that are important for price measurement. Semiconductor production technology progresses in distinct measurable steps, allowing us to account for technological improvement when constructing price indexes in spite of rapid changes over time. The continuing movement to outsource semiconductor production to offshore firms raises the possibility of outlet-substitution bias in standard price indexes and motivates our choice to focus on foundry wafer fabrication.

2.1 Semiconductor Production Technology

Semiconductor fabrication involves creating interconnected networks of transistors on the surface of a thin piece of semiconducting material.[4] The process begins with the design and layout of a new chip. Semiconductor designers use suites of complex software to specify the functionality of the chip, convert that functionality into the corresponding network of transistors,

[4] Turley (2003) provides an accessible overview of semiconductor technology, manufacturing, and business.

determine the physical layout of those transistors, and simulate the behavior of the proposed design for debugging purposes.

Semiconductors are generally manufactured on a thin wafer of pure silicon in a facility called a fab. Transistors are created on the surface of the wafer through a photolithography process, in which successive layers of conducting and insulating materials are deposited on the surface of the wafer and chemically etched away in the appropriate places to form the desired pattern of transistors and necessary interconnections. The design layout software determines the etching pattern for each layer, which is projected onto the wafer through a mask containing the negative of the desired pattern, in a process similar to developing a photograph by projecting light through a negative. Each step of the etching process is repeated multiple times across the wafer, resulting in a grid pattern of many identical copies of the chip. Once all transistors and connection layers are complete, the chips are tested in a process called wafer probe, and any faulty chips are marked to be discarded. The wafer is then cut up, leaving individual chips, called die. The die are then placed inside protective packages and connected to metal leads that allow the chip to be connected to other components.

Semiconductor fabrication technology has advanced over time in discrete steps, defined by wafer size and line width (also called feature size). Increases in wafer size allow larger numbers of chips to be produced on a given wafer. Most fabs currently produce 150mm, 200mm, or 300mm diameter wafers. Although larger wafers cost more to produce, they can fit many more die of a given size. For representative line widths and die size, the move to a larger wafer has generally resulted in overall savings of approximately 30% per die (Kumar, 2007).

Line width is the size of the smallest feature that can be reliably created on the wafer. Decreased line width means that individual transistors are smaller, and more functionality can be integrated into a given area of silicon. This makes chips of a given functionality smaller, lighter, and faster, and also makes it feasible to include more functions on a single chip. The number of transistors that can be produced on a chip has grown exponentially over time, following Moore's Law, the Intel co-founder's famous prediction that the number of transistors on a chip would double every two years (Moore 1965). Figure 1 shows the maximum number of transistors per chip and the

minimum line width used to produce Intel processors over the last 40 years (both plotted on logarithmic scales).

Current line widths are measured in microns (μm) or nanometers (nm). The smallest line width currently being produced in volume is 32nm. As a rule of thumb, Kumar (2007) estimates that moving a given chip design to a 30% smaller line width will result in cost savings of approximately 40%, assuming the same number of defects in both processes. The primary drawback of smaller line widths is increased cost per wafer, particularly early in the technology's life span. Masks are much harder to produce when creating smaller features, and new process technologies often result in higher defect rates and lower yields, the fraction of chips on a wafer that function correctly. In spite of these challenges, the benefits of increased die per wafer and better performance outweigh the problems of decreased yields, particularly as the fabrication technology matures and yields increase. Given the benefits of smaller line widths, semiconductor manufacturers have steadily moved toward newer technology. This is apparent in Figure 1 for Intel processors and can be seen even more clearly in Figure 2, which plots the technology composition of sales at Taiwan Semiconductor Manufacturing Company (TSMC), the largest semiconductor foundry.

There are a number of options regarding the chemicals used to create the transistors themselves and how the transistors are arranged to implement logical functions. The most common technology, called complementary metal-oxide semiconductor (CMOS), accounted for 97% of worldwide semiconductor production in 2008.[5] Other transistor arrangements, such as bipolar logic, and other chemical processes, such as Gallium Arsenide (GaAs) or Silicon Germanium (SiGe), generally focus on niche markets for high-frequency, high power, or aerospace devices, rather than the storage and computational logic products comprising the majority of the CMOS market. In the following analysis, we will refer to each combination of wafer size, line width, and logic family as a "process technology" (e.g. 200mm, 180nm, CMOS constitutes one process technology).

2.2 Changing Semiconductor Business Models

[5] Share of actual wafer starts reported in SICAS *Semiconductor International Capacity Statistics*.

In the early 1970's nearly all semiconductor producers were vertically integrated, with design, wafer fabrication, packaging, testing, and marketing performed within one company. By the mid '70's, firms began moving packaging and test operations to East Asia to take advantage of lower input costs (Scott and Angel 1988, Brown and Linden 2005). In spite of outsourcing these relatively simple steps in the production process, firms maintained their complex wafer fabrication operations in house. Firms that perform design and wafer fabrication are referred to as Integrated Device Manufacturers (IDM). As wafer fabrication technology advanced, the cost of production facilities increased dramatically; the cost of a fabrication facility has risen from $6 million in 1970 (IC Knowledge 2004) to $4.2 billion in 2009 (GlobalFoundries 2009). This sharp increase in cost has made it ever more difficult to stay at the leading edge of process technology. In the mid 1980's, small semiconductor firms began producing some of their more advanced designs on the manufacturing lines of larger, more established semiconductor manufacturers that were better able to bear the capital costs of maintaining a state of the art fab facility. Many Japanese semiconductor firms had substantial excess manufacturing capacity during this time period, making such production partnerships particularly attractive (Hurtarte et al. 2007).

These production sharing arrangements led to the creation of a new business model through the emergence of wafer foundries that manufacture semiconductors designed by other firms. At first, foundries were used by IDM's as an alternative source of capacity for older process technologies (Kumar 2008). By the late 1980's a number of new semiconductor firms avoided wafer fabrication by doing all of their manufacturing through foundries. Semiconductor companies without any in-house wafer manufacturing capability are called "fabless" firms. In general, fabless firms perform chip design and layout, and use foundries and other contractors for mask production, wafer fabrication, packaging, and testing. The fabless business model has grown quickly over the last 30 years, accounting for more than 20% of total semiconductor industry revenue in 2008, as shown in Figure 3. Since the largest foundries are located in Asia, and the largest fabless semiconductor producers are located in North America and Europe, the

growth of the fabless model has increased the internationalization of semiconductor production.[6] Although the fabless share of the global semiconductor industry only edged up from 2006 to 2008, as new process technologies continue to raise the costs of fab facilities, the prominence of the fabless model may well increase even more. Indeed, even very large IDM's such as AMD and Texas Instruments, have markedly increased their reliance on foundries (EE Times, Mar 11, 2002).[7]

2.3 Implications for Price Measurement

The extremely fast pace of technological change in semiconductor manufacturing poses a large challenge to quality-adjusted price measurement. Aizcorbe (2002) demonstrates the difficulty government price indexes have had in tracking rapid price declines in finished semiconductors. However, as just described, technological advance in semiconductor production proceeds in discrete, measurable steps, in contrast to continuous and difficult to measure quality improvements seen in other industries (Flamm 1993). This discrete nature of technological advance in the semiconductor industry makes it possible to control for quality changes, given detailed enough data on product characteristics. In this study we construct constant-quality price indexes for wafer production using quarterly pricing data that includes the relevant aspects of process technology: wafer size, line width, and logic family.

This section has also documented the increasing internationalization of the semiconductor supply chain coinciding with offshoring various steps in the production process and the growth of the fabless model of semiconductor production. Houseman (2007) describes the challenges faced by statistical agencies attempting to measure price changes when producers switch suppliers, particularly when the suppliers are located abroad. In particular, substitution toward low-cost suppliers is likely to be missed in standard price index calculations (see below for a more detailed discussion), understating the rate of input price decline. As semiconductor production

[6] In 2008, the 5 largest foundries (accounting for 84% of foundry revenue) were all located in Asia. Of the 25 largest fabless semiconductor companies (accounting for 75% of fabless revenue), 19 were located in North America or Europe. These figures were calculated from proprietary reports from iSuppli and GSA, respectively.
[7] A recent report (IC Insights) predicts that between 2008 and 2013, total foundry sales will grow at double the rate of the overall semiconductor industry.

technology advances and the fabless business model becomes more prominent, it is likely that these price measurement challenges will remain relevant in the foreseeable future.

In the remainder of this paper, we focus on foundry wafer production, leaving analysis of IDM production for future work. We make this choice for practical reasons. Our pricing data include only wafer purchases from foundries, though those purchases could have been made by fabless firms or IDM's choosing to utilize foundry suppliers. Also, the issue of within-firm transfer pricing raises a number of complications that are beyond the scope of this study and makes data collection essentially impossible.

3 Data Sources and Descriptive Results

To construct the price indexes used in our analysis, we require information on prices paid and quantities purchased for foundry services, specified by the characteristics relevant for pricing. We obtain prices from a survey conducted by the Global Semiconductor Alliance and we calculate quantities by merging several different sources. Observations are quarterly, and our data span the period from 2004 to 2008. Descriptive results demonstrate the importance of controlling for process technology. They also reveal substantial shifting of production toward lower cost countries.

3.1 Wafer Pricing Survey

Our primary dataset consists of 7,455 individual responses to GSA's *Wafer Fabrication & Back-End Pricing Survey*, provided to us for 2004 to 2008.[8] The survey has been conducted quarterly since 2004 and provides extensive detail on contracts for foundry services, including key technological features, foundry location, price paid, and volume for a diverse set of foundry customers. The survey responses account for a representative sample of about 20 percent of the wafers processed by the foundry sector.

[8] Individual respondents are not identified in our data.

As shown in Table 1, we drop observations missing key variables. We also drop observations reporting prices for engineering runs, preliminary fabrication before volume production. To focus on substitution between onshore and offshore production, and between offshore locations, we retain only contracts for production at the major offshore locations (Taiwan, Singapore, and China), U.S. foundry contracts, and European contracts for comparison.[9] A small number of observations with internally inconsistent responses are dropped, as are the handful of observations on 100mm wafers – a very dated technology. All told, we use 5,464 observations for index construction.

3.2 Descriptive Price Results

Descriptive statistics for key variables in the resulting dataset are shown in Table 2. We observe 273 prices per quarter, on average. Wafer prices average $1,575 over the period covered. Interestingly, no substantial time trend is evident before adjusting for composition. The average contract was for 2,307 wafers, and the average contract size climbs over time. The number of layers per wafer also rose significantly over the period studied, from 23 in 2004 to 28 in 2008, reflecting a trend toward foundries handling increasingly complex products.

The changing technological characteristics of the fabrication process are evident in the statistics for wafer diameter and geometry. Pilot lines for 300 mm wafers were first introduced in 2000 and the share for this emerging technology rises from 3.5 percent of contracts to 20 percent of contracts over the survey. Similarly, new generations of lithography increase in penetration over time: 90 nanometer technology reached volume production in the overall semiconductor industry in 2004 and slowly gained share in the foundry market, ending at 7 percent in 2008; 65 nanometer contracts are still emerging in 2008.[10] Meanwhile, older technologies, with processes above 250 nanometers, dwindle in prominence from 45 percent in 2004 to 28 percent in 2008.

[9] Significant omissions from the global foundry industry are Japan and Korea. Our approach to estimating capacity, described below, does not allow us to assign reasonable weights on technologies in Korea. Our preliminary price index for Japan behaved erratically, and suggested that the product composition was changing in a way not captured by our data. We have obtained more detailed data extracts that may assist in alleviating this problem in subsequent versions.

[10] 2004 and 2007 mark the years when volume production of DRAM began at 90nm and 65nm, respectively (International Semiconductor Technology Roadmap, 2007).

92 percent of contracts reported in the survey are for CMOS technology, but prices are available for other processes as well.

A challenge with the GSA pricing survey is sporadic reporting for some technologies in certain geographic regions, despite independent evidence that such production existed. For such cells where we believe there was production (based on our capacity database described in the next subsection) we linearly interpolate prices using values from surrounding periods or extrapolated based on higher-level prices.[11]

3.3 Quantities and the Shifting Geography of Production

To ensure that our price indexes are representative regarding technology, we collected data on global foundry capacity. Although the GSA survey includes information on the size of each order, some gaps in reporting remain. This makes weights based on the GSA data unstable at quarterly frequencies. As an alternative, we turned to a number of different sources on semiconductor fab capacity.

The Gartner *Semiconductor Fab Database* provided us with quarterly capacity data from 2004 to 2007. For specific fabs, key features are reported, including planned wafer start capacity, minimum line width, operating status, and whether the fab was operating as a foundry. We extended these data with GSA's *IC Foundry Almanac (2009)* which provided as snapshot of capacity and technology by fab as of 2009.

Merging these data sets gives us a preliminary set of weights, but we address several remaining shortcomings. First, Gartner only reports *planned* capacity by fab and ramp-up status, leaving the contours of the ramp-up process unknown. Fortunately, many major foundries provide quarterly information on actual operational capacity, showing the actual path of capacity as equipment is added incrementally. We employ these directly reported capacities, when available, and add a comparable ramp-up period to fabs for companies without direct reporting.

[11] Note that dropping these periods for lack of directly observed prices is not neutral, since it amounts to (1) assuming the product mix within the industry is different than we know it is, and (2) throwing out price information from this period for cells with similar with regard to technology or geography. See discussion in (Gordon, 2006).

Second, the data do not distinguish CMOS production quantitatively, though GSA does indicate whether a fab uses CMOS and other processes. Since CMOS prices behave rather differently than non-CMOS prices, we assigned a weighted average of the CMOS and non-CMOS prices to each fab for the technology in operation, using overall industry weights from the GSA. Third, in the Gartner fab database, we only observe the minimum line width in use at a fab, but we know that fabs often operate multiple geometries at any point in time, raising the possibility that we overweight leading edge technologies. On the other hand, we only observe capacity, though we would prefer to construct actual production weights. Since utilization is significantly higher for leading edge geometries, this raises the possibility that we underweight these geometries.[12]

Table 3 compares two aggregate measures of foundry capacity, constructed as just described, to industry estimates from other sources. First, wafer fab capacity as reported to the SICAS survey suggests our wafer fab measure is not fully capturing the overall size of the sector. However, the growth rate from 2004 to 2008 for the measure constructed from our bottom-up approach is very close to the SICAS measure, suggesting we are catching major trends in the industry. Our measure of revenue is also somewhat lower than the measure of foundry company revenue published by the consultancy iSuppli. This may simply reflect that not all foundry revenues are for the services we are studying. Table 4 shows shifting revenue weights among the largest offshore foundry suppliers. While Taiwan's share falls somewhat, China and Singapore both gain revenue share, representing movement toward lower cost foundry locations.

4 Price Index Results

This section presents our price index calculations using the database just described. We generate indexes under a variety of assumptions in an effort to quantify the effects of substitution across foundries in different countries. The results imply that, under an assumption that price differences across countries are not based on quality differences, such substitution contributes 0.8 percentage points per year to the price index for processed foundry wafers. Our findings also

[12] Utilization on fab lines using 90nm and smaller geometries was 94% in 2007, noticeably higher than the 86% utilization for larger geometries. (SICAS, 2007)

support the established importance of careful quality adjustment to capture the effects of rapid technological change on semiconductor prices.

4.1 Fisher Matched Model Index

Our data set includes price information by detailed semiconductor wafer type and source country at the quarterly frequency. As discussed in Section 2, a wafer's process technology (defined by wafer size, line width, and logic family) determines its performance, along with circuit design. Process technologies proceed in discrete steps, so our detailed data on prices by process technology yields a time series of price observations for each wafer type, with attributes held constant over time. This high level of detail allows us to construct a matched model price index tracking quarterly price changes for each wafer type.

In constructing our price indexes we need to determine what characteristics determine the performance, and hence the price, of a given wafer. To guide this choice, we have consulted pricing models used by engineers at fabless firms to estimate production costs when they are developing business plans. Kumar (2008) presents a wafer cost model based on wafer size, line width, and logic family. A commercial cost estimation firm, IC Knowledge, distinguishes wafer cost estimates by wafer size, line width, logic family, number of polysilicon layers, and number metal layers. Given this potential importance of the number of layers in a given design, indicating the design's complexity, we calculate price per layer rather than price per wafer (although results for wafer prices are qualitatively similar to those presented here). These pricing models support the use of process technology (wafer size, line width, and logic family) to distinguish between goods in our price indexes.

The matched model index is calculated as a Fisher index of price relatives for each process technology and country pair. First we calculate Laspeyres and Paasche indexes, respectively, as

$$P_L^t = \sum_i \sum_j s_{ij}^{t-1} \frac{p_{ij}^t}{p_{ij}^{t-1}},$$

$$P_P^t = \left[\sum_i \sum_j s_{ij}^t \left(\frac{p_{ij}^t}{p_{ij}^{t-1}} \right)^{-1} \right]^{-1},$$

where *i* represents process technology, *j* represents source country, *t* is time (quarter), and *p* is the average price for a given process technology, country, and quarter, from the GSA survey.[13] *s* is the share of total output value in time *t* accounted for by wafers in the relevant process technology and country cell, calculated using our capacity database. The Fisher index is calculated as a geometric mean of the Laspeyres and Paasche indexes.

$$P_F^t = \sqrt{P_L^t \cdot P_P^t}$$

We normalize the index to 100 in the first quarter of 2004.

The procedure just described treats observations from different source countries as separate "models" by calculating separate price relatives by country. This parallels similar treatment of prices across outlets in the U.S. CPI, and is subject to similar assumptions (Reinsdorf 1993). When a new process technology and country combination appears, it is assumed that any difference in the price *level* across countries for that process technology entirely reflects quality differences. This is the "link-to-show-no-price-change" method in Tripplett's (2004) classification of linking methods for matched model indexes. This linking strategy is based upon the assumption that the law-of-one-price holds when considering quality adjusted units across outlets. As we argue below, there is reason to believe that this assumption does not hold in the semiconductor wafer fabrication industry, potentially leading the standard matched model index to understate the true rate of price decline.

As expected, entry and exit are a prominent feature of the data. As shown in Table 5, 27 cells are new entrants in the 2004-2008 period, and 23 cells are exits. This raises the challenge of estimating price changes for the first and last periods in the series for a large share of the data. However, because our data is high frequency (quarterly), the number of entrants or exits in any given quarter is small, at 2.5 on average. In addition, the weights on these periods are small as new technologies ramp up gradually.

[13] Note that we use price per layer for the results presented here to account for the increased cost of producing more complex wafers containing more layers. As we expect, an index based on price per wafer falls somewhat more slowly, but the qualitative conclusions using price per wafer are the same as those presented here.

Table 6 presents our price index calculations. Column (1) contains the Fisher matched model index just described. We present the quarterly index, yearly averages, and the average yearly change between 2004 and 2008. The index falls by 12.6% per year. As has been known since at least Flamm (1993), Grimm (1998) and more recently Aizcorbe (2002), quality adjustment of prices for semiconductors, and indeed for all high-tech products is critical. The substantial differences across countries points to the necessity of accurate weights by country.

4.2 Relaxing the Location-as-Quality Assumption

Our previous index maintained the assumption that price differences across countries for otherwise identical goods reflect unspecified differences in quality. We now construct an index assuming that these price differences reflect price dispersion for goods of identical quality. That is, we calculate unit values by technology, averaging across observations without regard to country.

Incorrectly specifying the features that matter for quality may lead to bias in the price index, as follows. Assume two countries exhibit similar price trends for a given wafer type, but one has a consistently lower price level. Any shifts toward the lower cost country's foundries will have no effect on the price index, since the prices are assumed to decline at the same rate in both countries. The linking procedure implicitly assumes that the savings accrued in shifting suppliers are offset by lower quality of the goods being purchased. If, however, the goods are actually identical, then the shift to the lower cost country represents a genuine price drop for the relevant customer. The matched model presented in the previous section would miss this price drop achieved in switching suppliers, and thus understate the true rate of price decline. This is the so-called "outlet substitution bias" discussed in the Boskin Commission (1996) report.

Ideally one would be able to directly observe buyers substituting between different outlets. Since our data do not include purchaser identifiers, directly observing substitution is impossible. Instead, we follow Reinsdorf (1993) and calculate an average price index across outlets. This index is motivated by the opposite quality assumption of the index presented above. If models

are very well defined, one can assume that quality for a given model is identical across outlets. In our context, this amounts to assuming that a given process technology is identical across foundries in different countries. If this assumption is correct, then there is no reason to distinguish price relatives by country. Instead, we calculate average prices across countries for each process technology:

$$\overline{P_i^t} = \sum_j w_{ij}^t p_{ij}^t ,$$

where *w* is country *j*'s fraction of the total number of units of process technology *i* produced at time *t*. We then generate price relatives of these average prices for each process technology and use them to generate a Fisher price index as described above. This approach is able to capture the effect of substitution toward low cost countries as the weights on the lower prices increase with substitution.

If demand for wafers is shifting toward low cost suppliers, and the matched model is missing this substitution effect, we expect to find that the average price index declines more quickly than the matched model index. The results are presented in Column (7) of Table 6. The index falls by 13.4% per year, which is 0.8 percentage points faster than the matched model index in Column (1). This result supports the notion that outlet substitution bias causes the standard measure to understate the price declines for wafer fabrication, suggesting an outlet substitution problem no bigger than 0.8 percentage points per year. Note, however, that the scale of quality change over time is much larger, as indicated by the sharp overall price declines.

This result should be interpreted with a number of caveats in mind. Both the law-of-one-price assumption and the alternative assumption of uniform quality across countries are extreme. The data likely reflect both quality differences across countries and some persistent quality-adjusted price differences. Thus, the two approaches bound the true quality-adjusted price change, and the difference between them is an upper bound on the effect of outlet substitution.[14] This discussion raises the question of why quality-adjusted price differences should be able to occur in equilibrium. In this semiconductor fabrication market, a number of observations support the idea that quality-adjusted price differences can persist over time. There have been substantial

[14] There may be additional substitution to lower price producers *within* countries, which would be reflected in our country-specific price index. To the extent this is true, our bounds would be accordingly larger.

shifts toward low cost countries. This behavior suggests the presence of quality-adjusted discounts at the low cost countries. Why might that be? Although Reinsdorf (1993) discusses the role of costly information gathering in generating real price dispersion, we think that this explanation is unlikely to hold in a market as concentrated as this one. Rather, we propose an alternative reason for price dispersion based on the particular characteristics of the wafer fabrication industry.

Very large fixed costs are incurred when getting a production line up to capacity with a given design. Discussions with engineers at a large U.S. fabless firm indicate that it takes a large number of sensitive calibrations to fabricate a particular design on a particular production line. This creates substantial startup cost, such that semiconductor firms are very reluctant even to switch production lines within the same foundry, much less to move a product to a different foundry. This fact, coupled with the nature of new product introduction across countries leads us to a potential explanation for equilibrium price dispersion.

Consider the price plots presented in Figure 4. The top panel plots prices by country for a leading edge technology. Taiwan entered the market first, with a high price. Singapore and China each entered later, each at a lower price level. In spite of the increased competition from competitors entering the market, the Taiwanese price continued to decline at a steady rate, maintaining a roughly constant price differential relative to the others. A similar pattern for a more mature process technology is apparent in the bottom panel of Figure 4, in which a roughly constant price differential is maintained between the U.S. and Taiwan relative to Singapore and China.

To understand the implications of these observations, consider only Taiwanese and Chinese foundries for simplicity. If a given design requires the newest technology, it will have to be produced in Taiwan. In two years' time, when the Chinese foundry brings the same process technology on line, they charge a lower price in order to win market share away from their Taiwanese competitors. However, the lower wafer price in China does not outweigh the fixed cost of moving the existing products from Taiwan. The Taiwanese foundry can maintain a discretely higher price without losing its existing business, and only new products using the now

year-old technology will go to the lower priced Chinese foundry. The Chinese foundry may adopt the new technology more slowly due to a relative lack of technical expertise or due to U.S. export license restraints on advanced semiconductor fabrication equipment going to China (EE Times, Apr 27, 1998). In any case, the presence of large fixed costs of switching foundries coupled with staggered entry into a given technology makes persistent quality-adjusted price differences across countries possible.

4.3 Hedonic Price Index

To check the robustness of our results, we next generate a hedonic price index. Table 7 presents some information on the importance of the characteristics we observe. We regress log price per wafer on indicators for foundry location, technological characteristics, contract size, and quarter indicators using the 5,000 observations on contracts for CMOS technology. [15] All of these variables have a noticeable effect on prices and are estimated precisely. Collectively, they account for 88 percent of the variation in wafer prices. Controlling for technology, China has markedly lower prices than Taiwan, which serves as the baseline case in the regression. Singapore prices are moderately lower than Taiwan's, while U.S. and European prices are substantially higher. Production using more advanced technologies clearly commands a higher price. Compared to the baseline case of production on 200 mm wafers with 180 nm geometry, production on larger (300 mm) wafers and production with narrower line widths is significantly more expensive. More overall layers per chip, and more metal layers in particular, both proxies for the complexity of the circuitry, also drive up the price. Finally, contracts involving a greater scale of production do appear to draw a volume discount; other things equal, doubling contract size would be expected to reduce wafer costs by 5.5 percent.

Like the matched-model index, the hedonic index also falls rapidly, though the 11 percent average yearly rate of decline is 2 percentage points short of the rate for the matched model.[16] From this we conclude that our baseline results are fairly robust to choice of price index

[15] As mentioned above, non-CMOS technology is generally used in specialized niche markets. Although we do use non-CMOS prices when calculating industry price indexes, we omit them here for simplicity of exposition. Results for non-CMOS prices, not shown, indicate that location explains little of the variation in pricing, but technological characteristics do play a role.

[16] Aizcorbe, Corrado, and Doms (2003) find a similar result for microprocessors.

construction methodology. The hedonic specification also controls for characteristics not addressed in the matched model index, which suggest that contract size and the composition of layers contracted does affect pricing. The regression statistics indicate that these features explain over 80 percent of the variation in prices.

4.4. Official Indexes

For completeness, this section compares our results to the Bureau of Labor Statistics' price series for imported semiconductors. The BLS' International Price Program (IPP) publishes a price index for Harmonized System code 8542, Electronic integrated circuits. These include microprocessors and memory, the final products of the semiconductor production chain.

IPP draws its sample from Customs lists at the more detailed 10-digit Harmonized System level.[17] For instance, until recently, IPP would draw a sample of establishments whose product(s) are recorded under the just phased-out HS classification 8542.21.80.05 for "unmounted chips, die, and wafers." Price indexes are calculated at this more disaggregated level and IPP then aggregates across the price relatives to produce the published index. Unfortunately, this more detailed data is sealed to outside researchers for confidentiality reasons.

Perhaps *the* measurement challenge for IPP is to control for quality improvements in ICs. We do this via a matched model price index that controls for several important performance-related characteristics of wafers. IPP does not necessarily observe as many characteristics of each IC, but it does have a potentially promising way to identify quality improvements. At least some respondents provide BLS staff with their own internal product code assigned to the surveyed item. It is likely that new, higher quality products would receive a new product code. If IPP observes that the product code attached to the surveyed item changes, it will follow up with the respondent to ask what the price of the new product *would have been* last month so that it can record the true price change for the quality-enhanced good. These follow-ups based on observed

[17] This discussion draws on a number of conversations with Sonya Wahi-Miller of the IPP. We are very grateful for the time she spent educating us on the IPP's procedures. Any errors in our characterization of the IPP, however, are our own.

changes in firm product codes appear to be one of the principal ways by which IPP adjusts goods, at least in HS 8542, for quality improvements.[18]

The ICs observed by IPP are not directly comparable to the wafers studied in this paper. To see this more clearly, it is useful to recall that we can break up the production of ICs into four stages – design, wafer fabrication, test, and assembly. Our data pertain to the input produced in stage two whereas IPP measures the price of final output shipped at the conclusion of stage four. Nonetheless, it is instructive to ask how average price per wafer compares to the IPP estimate of the price of the finished product.

Table 6 Column (9) presents the IPP index by quarter over the period 2004-2008. Over this time period, the index falls on average 2.9% per year. Even though this is not directly comparable to our indexes, the discrepancy is quite large. It would imply that the prices in the remainder of the production chain (development, wafer test, and assembly) fall implausibly slowly. Consider, for instance, that recent research has found price declines that approach 40-50 percent per year for finished semiconductors sold in the U.S. (see, among others, Aizcorbe (2002), Table 1). This work suggests that prices at other stages of the production chain, such as test and assembly, fall faster than the price of wafer fabrication. This contrasts starkly with the message sent by the IPP series. A critical task for future work is to dig deeper into the sources of these discrepancies. In particular, it seems worthwhile to investigate whether the IPP's follow-up procedure for product code changes does in fact effectively capture key quality improvements.

5 Conclusion

Our analysis exploits a rich new data set to calculate constant quality price indexes for processed semiconductor wafers. We calculate matched a matched model price index, finding that wafer prices fall on average by 12.6% per year. Given that average prices, unadjusted for quality, remain fairly constant over the time period, the sharp yearly price decline demonstrates the importance of careful quality adjustment in this industry. Our results support the conclusion of

[18] Thus far, we have been unable to obtain information on how often this procedure is generally used in generating the HS 8542 index.

numerous previous studies that official statistics substantially understate the rate of semiconductor price decline.

Since our data set includes information on the source country for wafer purchases, we can also measure how geographic changes in sourcing patterns affect price measurement. Our approach is analogous to Reinsdorf's (1993) measurement of retail outlet substitution bias in the CPI. We calculate an average price index that captures the effects of shifting sourcing patterns toward wafer foundries in low cost countries. Our results imply that the baseline matched model approach understates the yearly price decline by at most 0.8 percentage points.

Although this problem is not overwhelming, particularly in comparison to the much larger issue of quality adjustment in the semiconductor industry, it is suggestive that continued shifts in international sourcing patterns will cause the problem to persist and potentially grow. Our findings here should motivate research into other industries that have seen large shifts in sourcing patterns across countries. Since there are large fixed costs of shifting suppliers in semiconductor production, the finding here may be smaller than the bias in more footloose industries that can substitute quickly in response to smaller price differences. Note however, that future analyses will need to motivate the assumption of persistent quality adjusted price differences across suppliers, as we do here.

References

Aizcorbe, Ana. 2002. "Price Measures for Semiconductor Devices". FEDS Working Paper 2002-13.

Brown, Clair and Linden, Greg. 2005. "Offshoring in the Semiconductor Industry: Historical Perspectives" *Institute for Research on Labor and Employment Working Paper Series*.

Clendenin, Mike and Junko Yoshida. 2002. "Game, set, match for foundries?: Even top-tier IC makers partnering as CMOS process rigors intensify." Electronic Engineering Times,(1209).

Diewert, W. E. 1998. "Index Number Issues in the Consumer Price Index" *Journal of Economic Perspectives*, 12(1): 47-58.

Doms, Mark, Ana Aizcorbe, and Carol Corrado. 2003. "When do Matched-Model and Hedonic Techniques Yield Similar Measures?" .

Flamm, Kenneth. 1993. "Measurement of DRAM Prices: Technology and Market Structure." In *Price Measurements and their Uses*, ed. Murray F. Foss, Marilyn E. Manser, and Allan H. Young, 157-197: National Bureau of Economic Research Studies in Income and Wealth, vol. 57; Chicago and London:; University of Chicago Press.

Gordon, Robert J. 2006. "The Boskin Commission Report: A Retrospective One Decade Later" *International Productivity Monitor*(12): 7-22.

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

Houseman, Susan. 2007. "Outsourcing, Offshoring and Productivity Measurement in United States Manufacturing" *International Labour Review*, 146(1-2): 61-80.

Hurtarte, Jeorge, Evert Wolsheimer, and Lisa Tafoya. 2007. *Understanding Fabless IC Technology*. Newnes.

"International Technology Roadmap for Semiconductors" 2007.

Kumar, Rakesh. 2007. "The Business of Scaling." *IEEE Solid-State Circuit Society News*, 12(1): 22-27

Kumar, Rakesh. 2008. *Fabless Semiconductor Implementation*, illustrated ed: McGraw-Hill Professional.

Moore, Gordon E. 1965. "Cramming More Components Onto Integrated Circuits." *Electronics*.

Reinsdorf, Marshall. 1993. "The Effect of Outlet Price Differentials on the U.S. Consumer Price Index." In *Price Measurements and their Uses*, ed. Murray F. Foss, Marilyn E. Manser, and Allan H. Young, 227-254: National Bureau of Economic Research Studies in Income and Wealth, vol. 57; Chicago and London:; University of Chicago Press.

Scott, A. J. and D. P. Angel. 1988. "The Global Assembly-Operations of U.S. Semiconductor Firms: A Geographical Analysis" *Environment and Planning A*, 20(8): 1047-1067.

Triplett, Jack. 2006. *Handbook on Hedonic Indexes and Quality Adjustments in Price Indexes: Special Application to Information Technology Products*: Paris and Washington, D.C.:; Organisation for Economic Co-operation and Development.

Turley, James L. 2003. *The Essential Guide to Semiconductors*, illustrated ed: Prentice Hall PTR.

"U.S., China at odds on fab-gear export." 1998, April. Electronic Engineering Times,(1004)

Figure 1: Moore's Law – Intel Processors

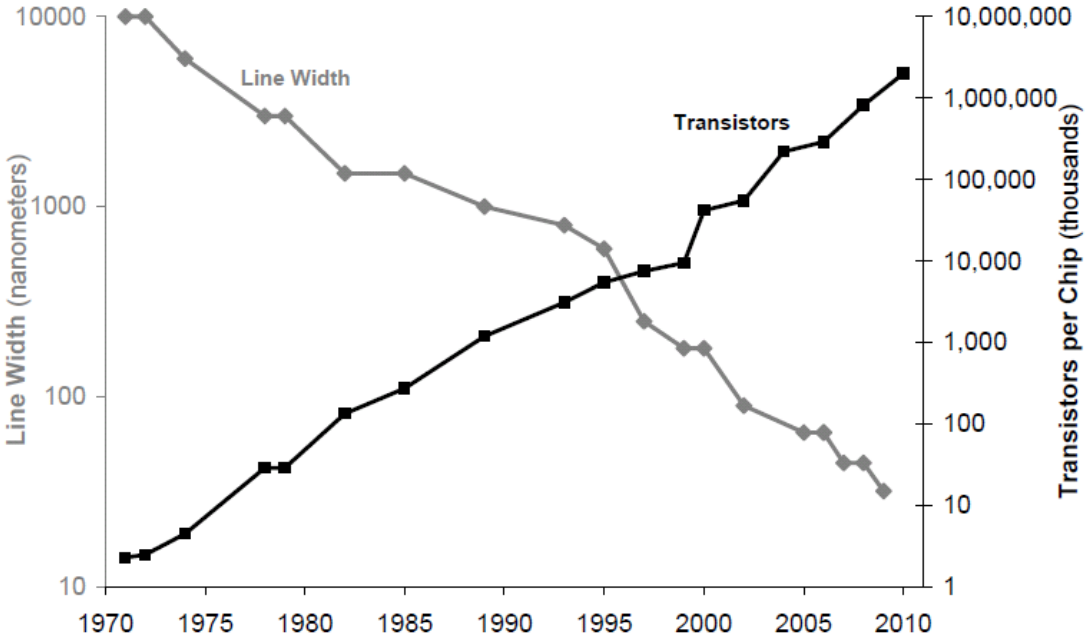

Sources: http://www.intel.com/technology/timeline.pdf
http://www.intel.com/pressroom/kits/quickreffam.htm

Figure 2: Technology Cycle – TSMC Sales by line width

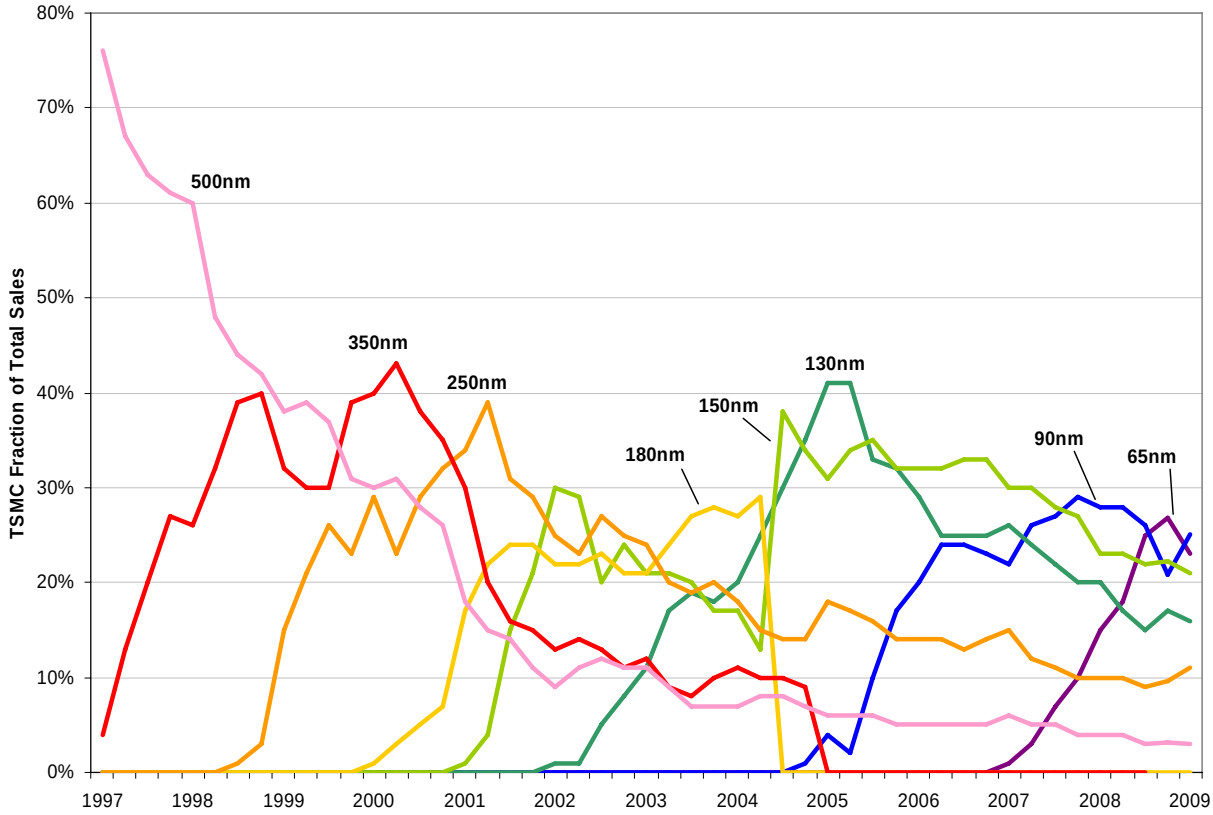

Source: TSMC quarterly reports

Figure 3: Growth of the Fabless Business Model

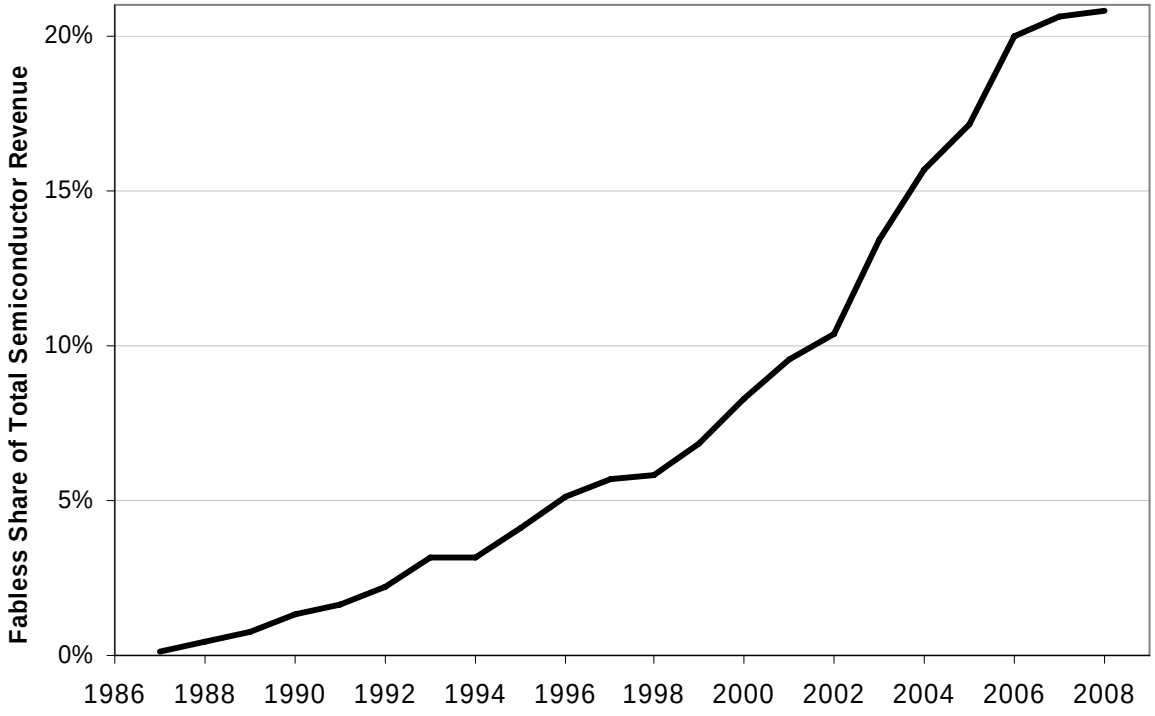

Sources: Global Semiconductor Association (GSA) and Semiconductor Industry Association (SIA)

Figure 4: Price Differences Across Locations

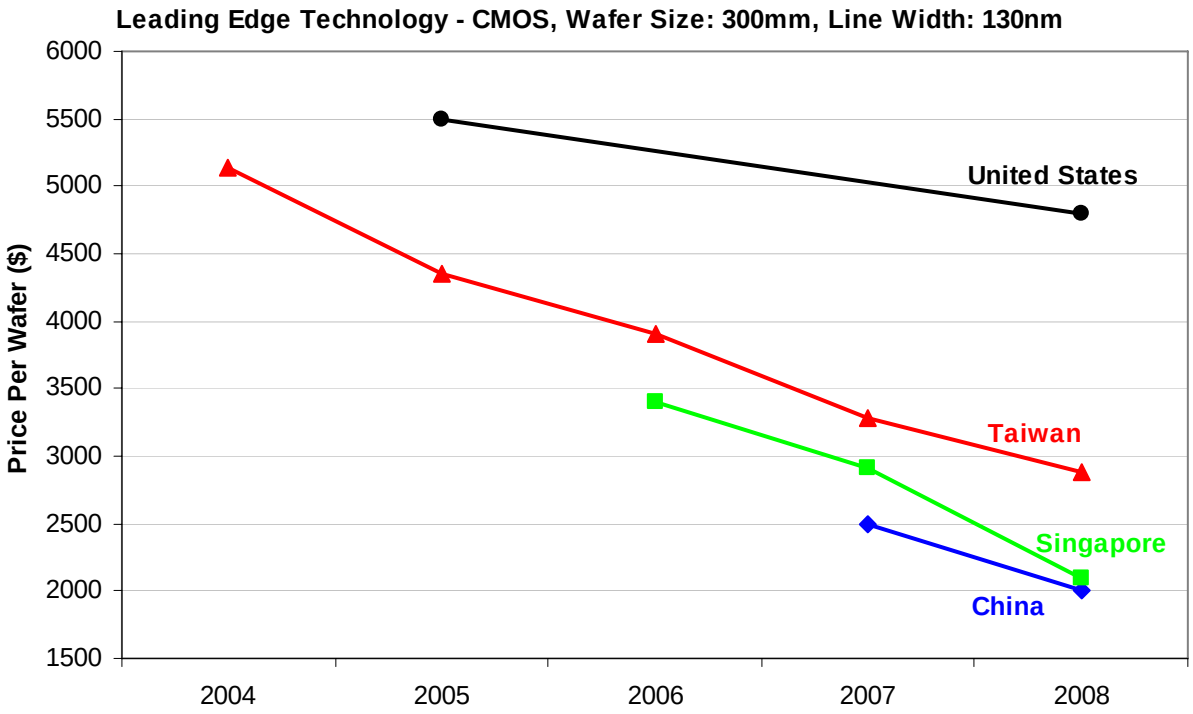

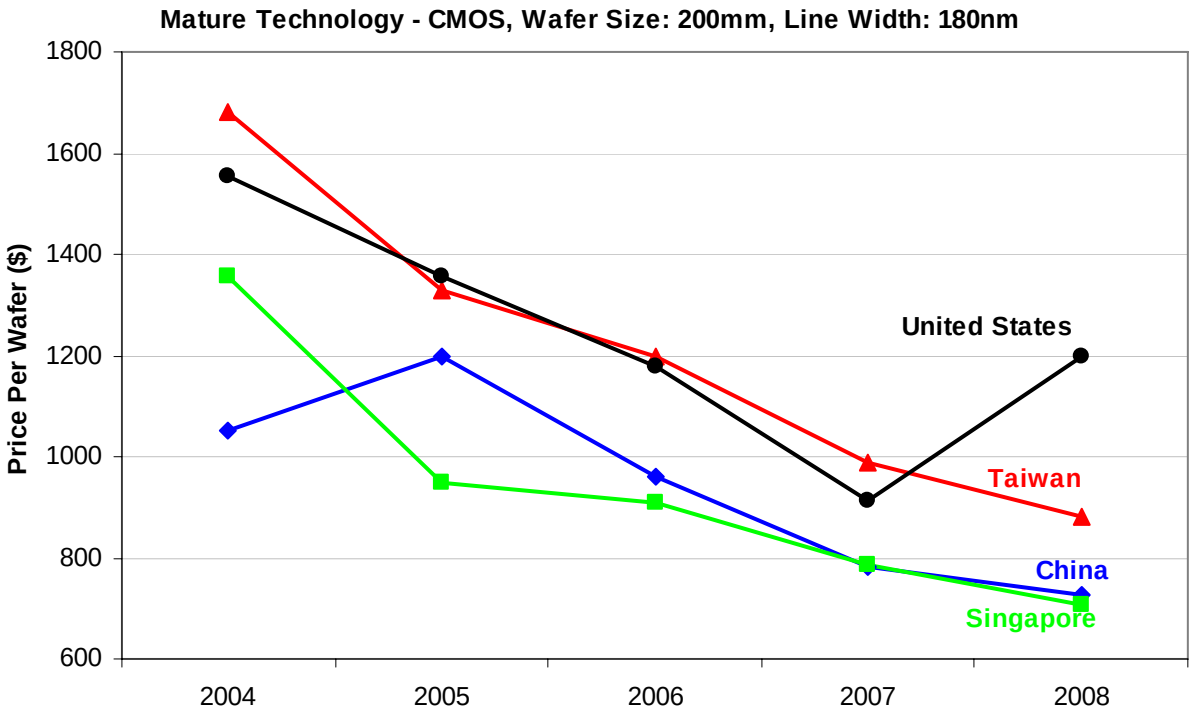

Table 1: Dropped Observations

Total observations	7455
Used in analysis	5464
Dropped	1991
Missing:	
foundry location	813
wafers purchased	19
price	19
Other reason:	
engineering run	778
location	499
100mm wafer	3
inconsistent	3

Note: there may be multiple reasons to drop a particular observation

Table 2: Descriptive Statistics

			Yearly Means				
	Mean	Std. Dev	2004	2005	2006	2007	2008
Price Per Wafer ($)	1575.40	1145.54	1,576.58	1,609.53	1,502.86	1,545.03	1,655.18
Number of Wafers Contracted	2307	7514	1924	2357	1941	2710	2627
Number of Layers Per Wafer	25.74	7.57	23.25	24.64	25.79	26.64	27.93
Metal Layers	4.77	1.81	4.23	4.55	4.75	4.97	5.27
Wafer Size							
150 mm or less	0.14	0.35	0.17	0.17	0.15	0.12	0.10
200 mm	0.76	0.42	0.80	0.77	0.79	0.76	0.70
300 mm	0.10	0.30	0.03	0.06	0.06	0.12	0.20
Line Width							
65 nm	0.00	0.06	0.00	0.00	0.00	0.00	0.01
90 nm	0.03	0.16	0.00	0.08	0.01	0.03	0.07
130 nm	0.23	0.42	0.14	0.18	0.22	0.27	0.32
180 nm	0.25	0.43	0.26	0.27	0.26	0.26	0.22
250 nm	0.13	0.34	0.13	0.16	0.12	0.12	0.09
older vintage	0.36	0.48	0.45	0.38	0.38	0.31	0.28
CMOS process	0.92	0.28	0.92	0.92	0.92	0.91	0.91

5464 Observations

Source: Authors' calculations based on GSA Wafer Fabrication & Back-End Pricing Survey

Table 3: Coverage of Constructed Capacity and Revenue

	Wafer Start Capacity (1,000 Wafers per Week)		Revenue (US $ Billion)	
	SICAS	Constructed	iSuppli	Constructed
2004	194	123	16.6	9.1
2005	252	139	16.3	9.0
2006	285	151	19.5	9.6
2007	288	172	19.7	9.8
2008	297	188	20.1	9.9

Source: SICAS, iSuppli, and author's calculations from sources described in text

Table 4: Foundry Revenue and Share for Major Offshore Locations

	Revenue ($million)	Taiwan	China	Singapore
2004	7232	66.0%	19.7%	14.3%
2005	8517	61.7%	20.4%	17.8%
2006	8549	62.0%	20.1%	17.9%
2007	8668	60.3%	21.6%	18.1%
2008	8432	59.8%	21.7%	18.5%

Note: Includes pure-play foundries only.

Source: Authors' calculations based on data from GSA, Gartner, and company reports

Table 5: Entry and Exit Statistics, CMOS Process

country technology cells with data	74
ave. no. quarterly prices per cell	10.18
new entrants	27
exits	23
cells with entry or exit	38
ave. quarters with missing prices	5.375

Table 6: Price Index Results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Fisher Matched-Model Indexes						Average	Hedonic	BLS IPP
Quarter	Overall	Taiwan	China	Singapore	USA	Europe	Price Index	Index	HS 8542
2004Q1	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
2004Q2	101.5	101.7	99.9	102.7	97.5	108.5	100.7	99.5	98.3
2004Q3	99.6	103.2	90.1	101.7	97.2	94.6	98.4	97.7	97.1
2004Q4	93.5	95.1	84.3	102.1	95.5	89.1	89.5	91.3	95.9
2005Q1	91.3	86.9	100.5	101.5	93.0	89.5	87.7	87.4	95.5
2005Q2	83.5	76.9	95.2	94.1	95.2	85.1	79.1	87.3	95.1
2005Q3	81.7	79.5	85.5	88.7	80.0	86.8	79.5	83.8	93.9
2005Q4	82.0	79.2	90.6	85.8	79.3	92.1	77.8	82.7	93.5
2006Q1	76.4	73.6	83.5	82.0	69.4	87.2	72.8	78.4	94.0
2006Q2	74.4	71.6	77.8	84.2	70.8	82.0	70.4	74.1	93.8
2006Q3	72.4	69.4	78.0	80.8	66.4	82.0	68.7	73.6	94.6
2006Q4	69.6	65.9	78.6	76.4	64.1	82.1	66.3	71.0	95.3
2007Q1	70.3	67.1	77.7	75.7	65.6	87.6	66.9	69.2	93.3
2007Q2	67.9	63.3	77.4	77.1	59.1	90.0	64.8	67.6	88.8
2007Q3	62.8	58.7	67.2	74.8	56.0	88.4	59.7	65.3	90.0
2007Q4	59.6	55.5	65.3	70.1	52.2	84.1	56.5	64.5	90.3
2008Q1	60.4	55.7	66.3	71.7	58.2	83.7	57.3	65.0	88.5
2008Q2	57.1	51.9	63.9	68.2	57.1	83.3	54.3	61.9	87.5
2008Q3	58.2	52.5	68.2	65.0	69.2	85.3	55.3	59.6	85.8
2008Q4	54.2	49.2	63.1	59.7	63.4	82.9	51.4	59.7	85.6
Year									
2004	98.6	100.0	93.6	101.6	97.6	98.1	97.2	97.1	97.8
2005	84.6	80.6	93.0	92.5	86.9	88.4	81.0	85.3	94.5
2006	73.2	70.1	79.5	80.9	67.7	83.3	69.5	74.3	94.4
2007	65.2	61.2	71.9	74.5	58.2	87.5	62.0	66.6	90.6
2008	57.5	52.3	65.4	66.2	62.0	83.8	54.6	61.6	86.9
Avg. Yearly Change '04-'08	-12.6%	-14.9%	-8.6%	-10.2%	-10.7%	-3.9%	-13.4%	-10.8%	-2.9%

Table 7: Descriptive Wafer Price Regression Results

dependent variable: log of price per wafer

Variable	Coefficient	Std. Err.	t-Stat
Foundry Location			
China	-0.272	0.019	-14.59
United States	0.218	0.014	15.58
Europe	0.119	0.018	6.54
Singapore	-0.062	0.012	-5.30
Wafer Size			
150 mm	-0.344	0.015	-22.20
300 mm	0.645	0.014	47.68
Line Width			
≥ 1000 nm	-0.696	0.038	-18.24
800 nm	-0.353	0.027	-13.31
600 nm	-0.358	0.022	-16.12
450 nm	-0.355	0.019	-18.35
350 nm	-0.194	0.013	-14.59
250 nm	-0.092	0.012	-7.74
130 nm	0.306	0.012	26.31
90 nm	0.511	0.025	20.19
65 nm	0.737	0.050	14.63
layers per wafer	0.012	0.001	13.83
no. metal layers	0.057	0.004	14.30
log wafers contracted	-0.055	0.002	-32.92
constant	6.743	0.030	223.47
R-squared	0.8773		
Observations	5000		

Specification also includes quarterly indicator variables

non-CMOS production not included

Baseline case (omitted category) is Taiwan, 200mm, 180nm