Pricing and Product Sourcing in an Offshoring Market: Evidence from Semiconductor Production Services*

David Byrne
Federal Reserve Board

Brian K. Kovak†
Carnegie Mellon University

Ryan Michaels
University of Rochester

Initial Draft: October 30, 2009
This Draft: July 4, 2011

Preliminary and incomplete
Please do not cite without permission.

Abstract

Measuring the breadth of offshore outsourcing, and its implications, remains a focus of academics and practitioners. To this end, the present paper studies the offshoring of production and its effect on tradable intermediates prices in the market for semiconductor wafer manufacturing services. Specifically, the paper makes three contributions. First, it uses a rich, proprietary dataset to document geographic shifts in product sourcing and uses these findings to shed light on how well official import price data capture the rapidly evolving global supply chain. Second, it provides evidence of substantial price dispersion in this offshore market, with new entrants offering deep discounts on observationally identical goods. Measuring prices for this exercise requires controlling for quality differences, and this paper is somewhat unique in analyzing pricing in an offshore contract manufacturing market using standard quality adjustment techniques. Third, the paper provides a formal accounting of the effect of offshoring on the typical wafer price paid by domestic buyers. In particular, it shows that the entrance of these low-price suppliers has accelerated the decline in the average wafer price, pushing it from 10.4 percent year to 11.8 percent year.

* An earlier version of this paper was 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. Thanks to Yong-Gyun Choi, Manuel Gomez, and Steven Paschke for excellent research assistance. We would like to thank Bill Alterman, Ross Goodman, Susan Houseman, Marshall Reinsdorf, Jodi Shelton, Sonya Wahi-Miller, and Kim Zieschung for helpful discussions and Chelsea Boone at GSA and Len Jelinek at iSuppli for help with data and background information on the industry. Remaining errors are our own. This paper reflects the views of the authors and should not be attributed to the Board of Governors of the Federal Reserve System or other members of its staff.

† Corresponding author: Brian K. Kovak, Heinz College, Carnegie Mellon University. Email: bkovak@cmu.edu

1. INTRODUCTION

In the last half century, reductions in transportation and communication costs have driven a fundamental change in the spatial and institutional organization of manufacturing production. It is now commonplace for individual steps in the manufacturing process for a particular good to be implemented in different establishments, which are owned by different firms and located in different countries. This fragmentation of production means that a larger share of international transactions takes the form of processing trade: inputs are assembled or processed abroad and then, in certain cases, shipped back to the firm that designed the good. It is generally thought that the tasks performed overseas are increasingly done by relatively low-price suppliers, as less developed economies such as China capture a greater share of the global market for tradable intermediates.

Measuring the breadth of this trend in offshore outsourcing, and its implications, remains a focus of academics and practitioners. To this end, the present paper studies, for a particular market, the offshoring of production and its quantitative effect on tradable intermediates prices. Specifically, we focus on the market for offshored manufacturing services for semiconductor wafers. By drawing on a variety of rich, proprietary data for this market, we can thoroughly document geographic shifts in product sourcing and measure the effect on international prices of shifting to lower-price suppliers. This is an important first step toward assessing the ultimate impact of offshoring on pricing dynamics at home and abroad, which is a topic of interest to policymakers (see Bernanke, 2007).

The semiconductor industry is frequently a target of academic study. This is, in part, because it has posted stunning rates of price *de*flation and productivity growth over the last 20 years.[1] It is also a fascinating industry for those interested in fragmentation, with relatively low-skilled tasks, such as final assembly, being offshored as early as the late 1960s.[2] More recently, the industry has witnessed the offshoring of relatively more skill-intensive production tasks, such

[1] For estimates of semiconductor price indexes, see Flamm (1993), Grimm (1998), and Aizcorbe (2002). Oliner and Sichel (2000) discuss the contribution of productivity growth in the semiconductor sector to the acceleration in aggregate productivity growth in the 1990s. Oliner, Sichel, and Stiroh (2007) update this work to include a look at the trends in the 2000s.

[2] See Brown and Linden (2005) and GAO (2006) for a detailed discussion of the history of offshoring in this sector.

as the fabrication of computer chips on silicon wafers. In particular, relatively low-price suppliers in China and Singapore have entered this market.

This recent trend in the wafer fabrication industry raises a more notable challenge to measuring prices in an offshore market. The rapid quality growth that has contributed to the stunning rates of price decline also complicates the accurate measurement of wafer prices. Yet as overseas production becomes increasingly sophisticated, researchers and practitioners must grapple with quality adjustment for tradable intermediates. Accordingly, measurement challenges in the wafer fabrication market likely preview similar challenges that will emerge in other industries as further offshoring of high-tech production occurs.

This paper proceeds by first discussing the technology of wafer fabrication, outlining the key physical attributes of wafers relevant for price setting. It then discusses the rise of offshoring, which is exemplified by the emergence of so-called "fabless" firms. These are companies that design computer chips but that outsource all of the manufacturing to contract manufacturers. Section 3 goes on to characterize the extent to which fabrication has recently shifted toward relatively new overseas suppliers, including China and Singapore.

In Section 4, we investigate the differences in prices across different overseas manufacturers. This exercise requires comparing "apples to apples", that is, comparing prices for goods of the same quality. Our discussion of fabrication technology from Section 2 delineates the key indicators of product quality for which we must control. Fortunately, we are able to do this using survey data collected by an industry trade association, the Global Semiconductor Alliance (GSA). Accounting for quality receives substantial attention in domestic price measurement, but this paper is one of few studies to estimate the effect of offshoring using sufficiently detailed data that permits a direct accounting of cross-country dispersion in quality. We find that, even after controlling for quality using transaction-level micro data, prices in this offshore market still differ by almost 30 percent, with relatively new entrants such as China selling wafers at deep discounts.

We pause briefly to note that this degree of price dispersion is potentially an important fact for models of offshoring to consider. In most theoretical work to date, this feature of

offshore markets has been (understandably) omitted in the interest of tractability.[3] But it should be noted that offshoring does *not* imply that the production of a "task" is specialized to a single location. Arguably, the study of offshore markets requires departing from a framework where the law of one price holds, a point to which we return shortly.

Building on our analysis of price dispersion, Section 5 presents a formal accounting of the effect of the relatively new, low-price suppliers on the trend in the average wafer price faced by a domestic buyer. To accomplish this, it uses the GSA data to construct two price indexes, one of which is designed to omit the effect of shifts in product sourcing on the average price and one that captures this impact. We use the difference between the two indexes to infer the proximate effect of offshoring on average prices in this market. The section concludes that the entry of low-price suppliers accelerated the average annual price decline by 1.4 percentage points, pushing it from 10.4 percent per year to 11.8 percent.

One obstacle to interpreting the results of Section 5 is the lingering concern that the price dispersion we document is actually due to *un*observable differences in quality. In that case, accounting for the effect of low-price suppliers is not a well-designed exercise. Therefore, in Section 6, we turn our attention toward understanding the sources of the differences in prices. The paper documents the presence of long-lived relationships between fabless firms and overseas manufacturers. We argue that these persistent arrangements are a response to frictions that raise the cost of switching production across foreign manufacturers. These arrangements underpin the sustained price differences we observe; in other words, at least some of this dispersion is likely to be "true" as opposed to driven by unobservable, or hard-to-measure, quality dimensions. We note that this interpretation is consistent with the steady shift in market share toward lower-price suppliers.

Section 7 concludes the analysis by relating our findings to data on U.S. import prices collected by the Bureau of Labor Statistics (BLS). Unfortunately, published BLS estimates of price inflation are too highly aggregated to directly compare with our results, since the Bureau is inhibited from providing more detailed unpublished price data because of confidentiality

[3] See, among others, Jones and Kierzkowski (1990), Yi (2003), and Grossman and Rossi-Hansberg (2008).

concerns. What we can do instead is gauge the general extent to which the BLS's *methodology* is capable of capturing the effect of the geographic shifting of production on prices.

Drawing on tabulations from BLS microdata provided to us, as well as our data on production locations (Section 3), we argue that the BLS' methods likely lead it to understate the effect of changes in product sourcing. This occurs, at least in part, for the following reason. According to BLS procedures, price indexes are first built at the buyer-product level, and then aggregated to yield the published index. When a buyer reports a different product design, therefore, a new index is initiated. Critically, this happens even if the production services – the input we seek to price – are unchanged. In other words, a change in the *output* (the design) may effectively be treated as a change in the *input*. From the standpoint of measuring price changes due to shifts in product sourcing, this poses a problem because changes in wafer sourcing almost always coincide with the introduction of a new product design. As a result, a shift toward a new wafer producer triggers the calculation of a new index, so any price decline incurred by switching producers is not reflected in the average price of the imported input.

Section 8 relates our work to related areas of research. This paper interacts with three strands of work. One documents the evolution of the global supply chain in order to estimate the volume of trade in intermediate products (see, among others, Hummels, Ishii, and Yi, 2001). Another analyzes the effect of new entrants abroad on average international prices (Kamin, Marazzi, and Schindler, 2006). We diverge from recent papers in this literature, in part, by our focus on a particular market. As we have argued, there is a clear payoff from this strategy: our relatively more disaggregated data enable us to more precisely identify changes in production location and to do a thorough quality adjustment for the purpose of price measurement. Lastly, a recent literature has begun to focus on the capacity of official data collection practices to capture changes in product sourcing. In this section, we also discuss how our more detailed data permit us to contribute to this line of work.

Section 9 concludes.

2. WAFER TECHNOLOGY

To better appreciate the GSA data, we need to briefly review two features of the industry. First, there are at least three distinct, measurable attributes of wafers that significantly affect their price. The GSA data are remarkable in part because they provide information on each of these technological characteristics. Second, it has become more common for designers of final semiconductor chips to outsource the production of the wafers to firms overseas. Hence, there are now many designers that are "fabless": they do not fabricate any of their wafers in-house. The GSA surveys these fabless producers and thus provides a unique insight into a market in offshored production services. This section details the key attributes of wafers. The next section discusses the geography of wafer production.

Semiconductor fabrication involves creating 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 logic into the corresponding network of transistors, determine the physical layout of those transistors, and simulate the behavior of the proposed design for debugging purposes.

Semiconductors are manufactured 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. Design layout software determines the etching pattern for each layer, which is projected onto the wafer through a mask containing 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 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 wafer. Most fabs currently produce 150mm (roughly 6 inches), 200mm (8 inches), or 300mm (12 inches) diameter wafers. Although larger wafers cost more to

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

produce, the move to a larger wafer has generally reduced the cost per die by approximately 30 percent (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 cofounder's famous observation that the number of transistors on a chip doubled every eighteen months (Moore, 1965).[5] 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 25nm. As a rule of thumb, Kumar (2007) estimates that moving a given chip design to a 30 percent smaller line width will result in cost savings of approximately 40 percent, 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 contract semiconductor manufacturer.

There are a number of options regarding the chemical compounds 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 percent of worldwide semiconductor production in 2008.[6] Other transistor

[5] This regularity later slowed to doubling every two years.
[6] Share of wafer starts reported in SICAS Semiconductor International Capacity Statistics.

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. We restrict our analysis to CMOS and refer to each combination of wafer size and line width as a "process technology" (e.g., 200mm wafer, 180nm line width).

The price index calculations below require us to define the set of product characteristics that determine the price of a given wafer. To guide this choice, we have consulted pricing models used by engineers to estimate production costs. Kumar (2008) presents a wafer cost model based on wafer size, line width, and logic family. Thus, we distinguish between models in our price indexes based on process technology (wafer size and line width, restricted to the CMOS logic family).[7]

3. THE GEOGRAPHY OF WAFER FABRICATION

In the early 1970s nearly all semiconductor producers were vertically integrated, with design, wafer fabrication, packaging, testing, and marketing performed within one company. By the mid-1970s, 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 both design and wafer fabrication are referred to as integrated device manufacturers (IDM).

As wafer fabrication technology advanced, however, it became prohibitively costly for younger and smaller semiconductor firms to stay at the frontier of process technology. The cost of building a fabrication facility has increased nearly 18 percent per year since 1970 and now stands at $4.2 billion (IC Knowledge, 2000; Global Foundries, 2009). Consequently, during the middle of the 1980s, younger and smaller firms began contracting with larger U.S. and Japanese IDMs to produce some of their more advanced designs in the latter's existing facilities (Hurtarte

[7] A commercial cost estimation firm, IC Knowledge, distinguishes wafer cost estimates by wafer size, line width, logic family, number of polysilicon layers, and number of metal layers. Although layers are potentially important in measuring the complexity of a given design, they represent an essentially continuous product characteristic. This makes it difficult to include layers in model definitions when constructing matched-model price indexes. We address this issue by additionally calculating a hedonic price index that includes detailed layer information. The hedonic index yields very similar results to the matched-model style indexes presented in Section 5.

and others, 2007). Around the same time, new firms sprung up overseas that were entirely dedicated to manufacturing wafers designed by other parties. These firms, operating principally in Asia, are known as wafer "foundries." Taking advantage of these new overseas facilities, a number of young U.S. semiconductor firms began outsourcing all of their wafer fabrication. These companies, which have little or no 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. It now accounts for about a quarter of total semiconductor industry revenue, as shown in Figure 3.[8] Some of the most prestigious U.S. chip makers, such as Fortune 500 firm Qualcomm, are fabless producers.

Table 1 shows how the share of worldwide foundry capacity has evolved in the last decade. By 2002, the majority of capacity was already in Asia, mainly Taiwan. Since then, the share of capacity in Asia as a whole has only increased moderately. But there has been a notable shift in capacity *within* Asia. In particular, China has more than doubled its share of foundry capacity, largely at the expense of the industry leader, Taiwan. Thus, in Table 1, we get our first look at the emergence of China in the market for wafer fab services. Singapore has also more than doubled its share of capacity since 2002.[9]

In the next section, we study the prices of wafer fab services across countries and show that the shifts in production location shown in Table 1 represent a movement toward relatively low-price suppliers, particularly China.

4. THE DISTRIBUTION OF WAFER PRICES ACROSS COUNTRIES

4.1 International data on wafer prices

Information on wafer prices comes from a proprietary database of semiconductor wafer purchases, collected by the Global Semiconductor Alliance (GSA), a nonprofit industry

[8] Note that the shares in Figure 3 are likely to understate the extent of fabless production activity because it counts only companies that derive 75 percent or more of their semiconductor revenue from fabless production. Many companies not counted as fabless, such as Texas Instruments, nevertheless rely heavily on foundries.

[9] The sharp increase in European capacity from 2008 to 2010 marks the founding of Global Foundries, which was the fab division of integrated device manufacturer, AMD. At the time of the divestiture, AMD had substantial fab capacity located in Germany in particular. These facilities were recorded as foundries in our data beginning in 2008.

organization. The dataset consists of 7,436 individual responses to the *Wafer Fabrication & Back-End Pricing Survey* for 2004-2008. The survey has been conducted quarterly since 2004. The survey responses account for a representative sample of about 20 percent of the wafers processed by the foundry sector worldwide.

The GSA dataset is unique in the amount of detail it provides for contract manufacturing of a high-technology product. For instance, it includes information on the technological attributes that industry analysts and engineers report as being the key price-forming characteristics of wafers. These are process technology, line width, and wafer size, as discussed in Section 2. In addition, GSA reports the location of the foundry for each transaction and the price paid. This information will allow us to investigate the effect of the shift toward lower-price wafer suppliers (see Section 3) on the average price of wafers purchased by a diverse set of foundry customers.

There are two noteworthy limits to the scope of the GSA data. One omission is the lack of firm identifiers. That is, we see information on individual transactions, but not the identity of the producing firm or buyer of wafer fab services. The potential importance of this omission will become clearer below (see Section 6). For our specific purposes in this section, the transaction-level information is sufficient. We also note that the GSA data only includes prices on wafers purchased through arms-length transactions and omits within-firm wafer production by integrated device manufacturers that would be subject to transfer pricing considerations, which are beyond the scope of this study.

We trim the dataset somewhat for the purposes of our analysis. As shown in Table 2, we exclude transactions missing key variables. We also drop observations reporting prices for engineering runs (preliminary fabrication before volume production), products using logic families other than CMOS, and the handful of observations on outdated 100mm wafers and one observation reporting an implausibly large wafer shipment.[10] All told, we use 5,423 observations in our analysis.

Descriptive statistics for key variables in the resulting dataset are shown in Table 3. We observe 275 prices per quarter, on average. 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 3.6 percent of contracts to 20.7 percent of contracts over the survey. Similarly, new

[10] GSA officials confirm that this observation should be dropped from the analysis.

generations of lithography increase in share 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 were just emerging in 2008.[11] Meanwhile, older technologies, with processes larger than 250 nanometers, dwindle in prominence from 45 percent in 2004 to 25 percent in 2008. The number of metal and mask layers per wafer also rose over the period studied, reflecting a trend toward foundries handling increasingly complex designs.

4.2 Wafer price dispersion

The GSA data furnish a unique opportunity to measure international price dispersion within a market for an important intermediate input. It provides information on the location of the foundry and the key product attributes that one needs in order to make sure that the comparison of prices is done on an "apples-to-apples" basis. This adjustment for product quality is particularly important in the market that we study. The average number of chips and transistors per wafer increased substantially due to larger wafers and smaller line widths, so an average wafer in 2008 contained many more transistors and interconnections than an average wafer in 2004.

To measure price dispersion, we estimate a simple hedonic regression. Specifically, we regress the wafer price observed in a transaction on a series of dummy variables. These include indicators for key product attributes (e.g., wafer size, line width), foundry location, and calendar quarter. (The reference category is a 200mm 180nm wafer produced in Taiwan.) We also include the quantity of the wafers purchased as a regressor.[12] All of these variables have a noticeable (and precisely estimated) effect on prices, as shown in Table 4. Together, they account for 88 percent of the variation in log wafer prices.

The signs on individual regressors are intuitive. For instance, more advanced technologies command a higher price. When 300mm wafers came online in 2007 and 2008, the regression results indicate that they commanded a 60 percent premium over the next largest size (200mm).

[11] 2004 and 2007 mark the years when volume production of DRAM began at 90nm and 65nm, respectively (*International Technology Roadmap for Semiconductors* 2007).

[12] Industry analysts have indicated to us that foundries do generally provide discounts for large purchases. The regression results confirm this.

In the same vein, the regression estimates show that as firms substituted from 180nm line widths to 130nm, they paid, on average, 35 percent more per wafer.[13]

The estimated effect of foundry location indicates substantial price differences across countries, even after controlling for detailed technological characteristics. China and Malaysia have markedly lower prices than Taiwan (which is the reference category). Specifically, a Chinese fab charges nearly 27 percent less for an observably identical product than does its counterpart in Taiwan. Japan, Korea, and Singapore's prices are moderately lower, while U.S. and European prices are substantially higher than Taiwan's. Indeed, the average U.S. price is 20 percent more.

That Chinese foundries sell at lower prices is an exceptionally robust result. Digging deeper, we have also looked at price changes in Taiwan and China for each technology, which is defined by wafer size and line width. Significantly lower prices are observed in China relative to Taiwan for 21 separate technologies in the GSA data. Higher prices are not observed for China in any case of substantial production volume.

To the best of our knowledge, these estimates represent one of the first pieces of evidence on price dispersion across source countries in a market for an important intermediate input into domestic manufacturing. In contrast, much of the literature on price dispersion to date has instead focused on consumer goods (see Engel, Rogers, and Wang, 2005 and Bergin and Glick, 2007).

5. THE ENTRY OF LOW-PRICE SUPPLIERS & AVERAGE PRICE INFLATION

The preceding sections presented evidence for cross-country variation in quality-adjusted semiconductor wafer prices and substitution toward the relatively low-price suppliers. As lower-price suppliers entered the market and captured market share, the average price for a given

[13] The reader will notice from the table that another indicator of the wafer's technology included in the regression is the number of polysilicon and metal layers. The number of layers affects the complexity of the fabrication process. Hence, an increase in the number of layers raises the price, as seen in Table 3. Ideally, we would include layers in our matched-model indexes presented in the next section. However, as discussed in Section 5, matched-model indexes rely on discrete indicators of product quality so that "models" may be defined as a collection of particular product characteristics. Because layers represent an essentially continuous product characteristic, with no obvious rules of thumb to discretize it, this variable is somewhat more difficult to include in the index. That said, the quality-adjusted rate of price declines estimated by hedonic regressions are similar to what we obtain from the matched-model index.

product (defined by observable wafer characteristics) presumably fell. In this section, we provide a formal accounting of the effect of geographic substitution on the evolution of average wafer prices.

To proceed, then, we calculate two different wafer price indexes. One index is constructed by first computing price changes *within* country and then averaging these changes across countries. By construction, this index excludes the effect of geographic substitution on average wafer prices. The other index is built by first averaging prices at each point in time *across* countries and then computing the price change. Increases in the share of production at low-price foundries will push this index lower. The difference between these indexes thus represents an effect of substitution. We find that shifts toward low-price suppliers account for an additional average price decline of 1.4% per year.

5.1 Matched-model Index Construction

It is helpful to begin by summarizing a few details concerning general price index construction. In the case of each index we present, we take a matched-model approach by tracking the price path of each product, or "model," over time and then aggregating these individual price indexes across models. The model is defined by a list of relevant product attributes that summarize the "quality" of the good. For instance, recalling that wafer size, line width, and logic family are the key determinants of a wafer's price, one example of a model would be a 200mm wafer with 90nm line width in the CMOS logic family.

Measuring price changes within model is a way of estimating a constant-quality price index. Controlling for quality is particularly important in this industry, as new higher-quality models are introduced frequently and carry steeper prices. Indeed, we have found that, if one simply averaged all wafer prices at each point in time, it would look as if wafers were becoming *more expensive* over time. However, the opposite is true once one controls for quality upgrades. Fortunately, controlling for certain quality-related attributes is also particularly straightforward using our data. Defining the technological parameters of a "model" according to size and line width is feasible because the technologies (size, line width) evolve over time in discrete steps. As the GSA data reports this information for each transaction, we can easily distinguish between models based on process technology.

Formally, a matched-model index is calculated as a Fisher index of quarterly price relatives for each model. First we calculate Laspeyres and Paasche indexes, respectively, as

$$P_L^t = \sum_m s_m^{t-1} \cdot \frac{p_m^t}{p_m^{t-1}} \quad (1)$$

$$P_P^t = \left[\sum_m s_m^t \cdot \left(\frac{p_m^t}{p_m^{t-1}}\right)^{-1}\right]^{-1}. \quad (2)$$

where m represents each model, t is time (quarter), p is the average price for a given model, and s_m is the share of total output in time *t* accounted for by model m. As it is well-known that the Laspeyres index overstates the welfare-theoretic "true" price change and the Paasche understates it, we construct the superlative Fischer index, which is a geometric mean of the Laspeyres and Paasche indexes:

$$P_F^t = \sqrt{P_L^t P_P^t}. \quad (3)$$

We normalize all indexes to 100 in the first quarter of 2004.

To calculate the price index (3), we need information on quantities to calculate s_m. Before we proceed to the results, we briefly describe here the source of these data. Although the GSA survey includes information on the size of each order, some gaps in reporting remain. This can make weights based on the GSA data unstable at quarterly frequencies. As an alternative, we construct weights based on global foundry capacity. For capacity data, we turn to the Pure Play Foundry Market Tracker database managed by market research firm iSuppli. This provides us with quarterly data from 2004 to 2008 for 78 specific fabs operated by 20 companies. This information, coupled with additional iSuppli data, enables us to calculate capacity in each quarter by country and process technology.

Although capacity may not perfectly reflect production or purchases due to changes in utilization rates, we must choose between somewhat erratic sales measures and highly credible capacity estimates.[14] Our calculations use the more stable capacity weights.[15]

[14] Since capacity utilization is higher for leading edge geometries, the application of capacity shares underweights these geometries. For example, utilization on fab lines using 90nm and smaller geometries was 94 percent in 2007, noticeably higher than the 86 percent utilization for larger geometries (SICAS 2008).

5.2 Within Technology and Country Index

Our first matched-model index is explicitly designed to *ignore* the effect of substitution toward lower-price suppliers on the average global wafer price. To accomplish this, we define a model to include a chip's technology (size, line width, and logic family) *and* the country where the foundry is located. For instance, extending the example from earlier, a 200mm-90nm CMOS wafer produced in China would be a model. We refer to the resulting index as the "within technology and country" index.

To see more clearly why this index omits the effect of substitution, consider the following simple scenario. Suppose two countries produce the same wafer type and both charge a constant price over two years. However, one of these countries has a consistently lower price level. Because a model includes the foundry country, we first calculate price growth within each country – zero in this case. Then averaging price changes across the two countries (with weights that reflect each country's market share) yields zero. Even if market share shifted sharply toward the low-cost country, this shift would have no effect on the price index.

Table 5 presents our price index calculations. Column (1) contains the within technology and country index just described. We present the quarterly index, yearly averages, and the average yearly growth rate between 2004 and 2008. The index falls by 10.4 percent per year, which implies that the price of a given quality wafer falls by half roughly every 6 years.[16]

This is consistent with estimates of finished semiconductors' prices. For instance, the Federal Reserve Board's semiconductors price index fell 28 percent per year during the period covered by our GSA data.[17] The difference between our result and the price of finished ICs –

---

[15] We also note a second, relatively more minor challenge to price index construction. The GSA survey contains sporadic information on prices for some technologies in certain geographic regions, despite independent evidence that such production took place. For such cells where we believe there was production (based on our iSuppli capacity database) we interpolate prices using values from surrounding periods.

[16] We also calculated a hedonic index that parallels the within country and technology matched-model index. This index was calculated based on the wafer price regressions reported in Table 4 and falls by 10.5% per year. This indicates that the additional wafer characteristics (i.e., number of layers and contract size) included in the hedonic regressions but not used in the matched-model index do not have that large of an influence on the overall rate of price decline.

[17] Semiconductor production is one component of the Board's industrial production index. Real semiconductor output is computed by deflating nominal output using the price index referenced in the main text, reweighted to reflect the product composition of U.S. production. Several important papers have investigated the secular decline in prices in finished semiconductors over earlier periods. These include Flamm (1993), Grimm (1998), Aizcorbe (2002), and Aizcorbe, Oliner, and Sichel (2006).

recall that the wafer is one input into the final chip -- suggests potentially large declines elsewhere in the supply chain. This is an important avenue for future work.[18]

We should note that the headline 10.4 percent result does mask some heterogeneity in price declines across regions. The dispersion in price inflation is perhaps accentuated by the unusually large price *in*creases observed in some countries in 2008. Through 2007, though, average annual price inflation across all regions ranged narrowly between -5 percent (Europe) and -11 percent (Taiwan). Perhaps not surprisingly, the technological leader among overseas foundries, Taiwan, also displays the fastest price declines.

5.3 Within Technology Index

Our second index defines a model based on process technology alone, so we refer to it as the "within technology" index. This explicitly omits country from the model definition and is designed to capture the effect of shifts in wafer sourcing on the global average wafer price.

To see more precisely how the index accomplishes this, return to the example of two countries that produce the same type of wafer. The within technology index treats wafers from both countries as the same "good". Accordingly, the index calculation procedure first calculates the market share-weighted average price *across* the two countries and then calculates the growth in that average price. Since both countries' prices are constant, the price index would exhibit no change in the absence of shifts in sourcing patterns across countries. However, if production shifts toward the low-price country, this procedure will show a decline in the average price.[19]

The within technology index is shown in column (7) of Table 5. It declines at a rate of 11.8% per year. Comparing this to the within-technology, within-country index, we conclude that substitution toward lower-price locations contributed an additional price decline of 1.4% per year.

In order to visualize the source of the difference between the two price indexes, consider Figure 4. This shows the price data for a particular process technology -- 300mm wafers with 130nm line widths -- produced in Taiwan, Singapore, and China. Taiwan was the sole producer of this process technology for three years, until Singapore entered in Q4 2006, followed by

[18] In preliminary work we have found that price declines in other material inputs, such as the assembly and testing of chips, have been modest. This suggests that the implicit price of the research and design performed by fabless firms has fallen substantially.

[19] Except for the difference in the model definition, the index calculation procedure is identical to that described in the previous section.

China a year later. Note that prices decline at similar rates across countries, but the price levels differ substantially. The within technology index reflects the share-weighted average price across countries and hence falls with the arrival of Singapore. In contrast, the within technology and country index continues to track the Taiwan price, since it averages the rate of price decline, which is similar across countries, and omits variation in the price level across countries.

Although the measured effect of cross-country substitution is substantial, our measure likely understates its effect for two reasons. First, we consider the effect of substitution only in an accounting sense, and do not model the equilibrium impact of the entry of lower-price suppliers on the decisions of the incumbent firms. We would expect entry to reduce the prices of the incumbents relative to what the latter would otherwise charge. The dynamics of price adjustment in general equilibrium remains a topic for further research. Second, we also only account for substitution across countries and not between producers within a given country.

5.4 Interpretation

We have found that shifts in product sourcing had a substantial effect on the typical wafer price. But the interpretation of this finding is not entirely straightforward. This is because one may question whether it is appropriate to regard a wafer made in one country as the same good as a wafer made elsewhere. We briefly elaborate on this concern here.

Arguably, the within technology, cross-country price dispersion documented above may be due to unobservable dimensions of quality. In the case of China, for instance, it is well known that the government, in general, aggressively pushes foreign firms to share their intellectual property with Chinese manufacturers.[20] If that is done in the wafer fabrication industry, the difference in price between China and Taiwan may simply be the differential needed to compensate foreign firms for this transfer. As a result, it is perhaps hard to argue that Chinese producers charge a lower price for the same "product", as IP protection associated with a product may be interpreted as an attribute of the good itself. This argument would suggest that true inflation is better measured by the within-technology, *within-country* price index given in Section 5.2, since this essentially removes the country-specific quality effect associated with the price.

[20] Dean, Fung, and Wang (2011) find evidence that highly R&D-intensive industries are less likely to establish supply chains involving processing in China.

We agree that it is hard to know how to decompose the price differential across true price dispersion and unobservable quality. But we provide two pieces of evidence to suggest that at least some of this dispersion is "genuine", and thus the true measure of price inflation lies somewhere between the two indexes given above. First, the interpretation of the differential as due entirely to unobservable quality is inconsistent with the sustained increase in market share obtained by lower-price foundries. If the law of one price held, then ostensibly lower-price suppliers, such as China and Singapore, would not have roughly doubled their respective market shares. A similar observation has been made, in other contexts, by Reinsdorf (1993) and Bils (2009). Second, it is well understood that price dispersion can emerge in markets characterized by frictions that impede the instantaneous arbitraging of price differentials. We find compelling evidence for these types of frictions in this market. In the next section, we provide a more extended discussion of this point. [21]

6. FOUNDRY RELATIONSHIPS

Thus far, we have documented the existence of substantial price dispersion in the offshore market for processed semiconductor wafers. This dispersion appears to be underpinned by the presence of "sticky", or long-term, relationships in this product market: Even after a lower price is available, fabless design firms continue to purchase from higher-priced suppliers such as TSMC in Taiwan. In this section, we document these long-term arrangements between foundries and fabless firms. We also discuss reasons why such arrangements persist. Specifically, we argue that the long-lived "matches" between fabless firms and suppliers reflect the presence of substantial frictions in this product market.

To begin, we turn to Figure 5, which shows the evolution of market share for a given technology. At any given point in time, Taiwan's TSMC and UMC are the only suppliers of the

[21] We would argue that our suspicion that there is true, and potentially substantial, price dispersion is consistent with the general thinking of the economic measurement community. For instance, many observers have argued that the Bureau of Labor Statistics ought to track the intermediate-good purchases of firms and use these data to build an *input* price index. The advantage of the input price index is precisely that it picks up price movements due to shifts in suppliers (Diewert and Nakamura, 2009). The BLS has said it is interested in such an index (Alterman 2009). Thus, given the widespread interest in this sort of index, there must be some openness to the notion that a portion of price differentials reflect "true" dispersion.

frontier technology among major foreign foundries -- but this national monopoly is temporary.[22] One to three years after TSMC and UMC introduce a new process technology, SMIC (China) and Chartered (Singapore) begin to ramp up production. These Chinese and Singaporean foundries charge less -- in the case of China, substantially less -- than Taiwanese foundries for an observationally equivalent wafer. Thus, it is notable that TSMC remains the leader in terms of volume well after lower-price suppliers enter the market.

TSMC's sales do not fall off because it apparently has "locked in" purchases from its customers. Unfortunately, the GSA data do not permit us to explore these long-term supplier-buyer arrangements directly, because it does not include firm identifiers for producers or buyers. However, we do have two pieces of supplementary evidence in this regard. One comes from specific contracts between foundries and fabless firms. The other comes from supplier relationships mentioned in fabless firms' annual reports.

First, we have been able to inspect a handful of contracts between fabless firms and foundries.[23] These provide information on periodic price reviews, production schedules, and the duration of the agreement. Contracts generally run for a term of at least three years. By itself, the duration of the agreement suggests a degree of "stickiness" in product-market relationships. In addition, the contracts indicate that production at any point in time is predetermined by prior agreement. For instance, in each month, the contract requires the buyer to provide a forecast of its production needs to the foundry over the next 12 months. The first three months of each production forecast are binding on each party and thus compensation must be made if the buyer were to suspend production. Clauses such as these discourage the movement of production from one foundry to another.

Second, Figure 6 displays the relationships between major fabless companies and foundries. This information is culled from the fabless firms' annual (10K) reports.[24] Although

[22] Korea's Samsung and IBM offer leading-edge foundry services, as has Global Foundries, since entering the market in 2009. But as TSMC is the world's largest foundry during the period covered by our price data, we will, for the sake of simplicity and concreteness, use TSMC throughout this section as our example of a leading-edge manufacturer.

[23] A publicly traded fabless firm must provide the wafer supply agreement to the Securities and Exchange Commission if the "registrant's business is substantially dependent" on a particular supplier. The contracts are filed as Exhibit 10.7. The agreements we have obtained are those for Marvell Technology and TSMC; Altera Corporation and TSMC; Conexant and Jazz Semiconductor; and Advanced Power Technology and CSMC Manufacturing.

[24] Publicly traded companies provide a description of important risk factors to their business in their annual 10-K filing to the Securities and Exchange Commission. Among these disclosed risk factors are significant reliance on particular suppliers of intermediate goods and services.

all foundry relationships are not necessarily reported in 10K reports, the most important suppliers will be listed. We observe the formation of new relationships with low-price suppliers, such as SMIC. But note that fabless firms that initiate a relationship with SMIC generally do not terminate their arrangements with other providers. For instance, when Qualcomm signed an agreement with SMIC in 2006, their relationships with TSMC and UMC remained. Thus, once established, foundry-fabless relationships are highly persistent; the hazard rate at which they break up is only 6 percent per year.[25]

The pattern of long-lived product market relationships evident in Figure 6 also clarifies that the notion of substitution toward a lower-price supplier is more subtle than a fabless firm "firing" its foundry and replacing it with another. This is not what appears to happen. Among mature fabless firms, what instead happens may be better illustrated by the Qualcomm-SMIC agreement. Qualcomm signed a deal with SMIC in July 2006 to produce wafers with 130nm line width for use as power management chips in cellular phones. (These are devices that, among other functions, perform battery charging). At the same time, Qualcomm was under contract with TSMC to ramp up production of a 65nm processor design for cellular handsets. Thus, it seems that Qualcomm turned to China for production of trailing-edge devices but ramped up production of the more sophisticated 65nm wafers by sourcing to the industry leader in Taiwan.[26]

Qualcomm's decision represents substitution along the *intensive* margin: less services are sourced to TSMC than would perhaps otherwise be the case, but the relationship itself is not severed. Rather, TSMC remains the supplier of choice for leading-edge products, while SMIC serves Qualcomm's needs for trailing-edge products that could have been produced at TSMC, but for a higher price. In contrast, if an IDM decides to change business models by sourcing from a foundry and chooses SMIC *rather than* TSMC, this represents movement along the *extensive* margin. This is, of course, another avenue through which the market as a whole substitutes toward lower-price suppliers. Unfortunately, as noted above, the GSA data do not allow us to track individual firms across time, so we cannot observe these changes directly.

[25] Figure 6 is consistent with what industry participants have repeatedly stressed to us. For instance, a design manager at LSI Corp, a major U.S. fabless firm, said he was aware of only one instance in his career when his firm shifted a product between foundries. Executives at a large foreign fabless firm also told us that they only switched foundries when their original partner has reached full capacity and was not able to fill their orders quickly enough.

[26] This example is reconstructed based on several articles in the *Electronic Engineering Times*, the electronics industry newspaper of record (cf.LaPedus, 2006).

We now ask what enforces these long-lived relationships. In short, there are substantial costs, in terms of real resources and reputation, of interrupting a wafer supply contract and re-sourcing the product to lower-price supplier.

First, it is very costly to reallocate the production of a particular design from one foundry to another once it has begun. The reason for this begins with the fact that chip designs are highly differentiated and must be tailored to specific equipment to be manufactured. A good example of this is the mask set – the key guide used by lithography equipment in etching the pattern of transistors onto a semiconductor wafer. The mask set is built for a particular design, but it also must be specifically fabricated to be compatible with the machines in a particular foundry. Critically, foundries have their own propriety processes and technologies that are generally incompatible across manufacturers. This implies that a mask often cannot be reused at a foundry other than the one where the wafer is initially sourced. In the case of the mask set in particular, that implies an exceptionally high cost of switching: a typical mask set designed for fabricating a 90nm wafer is priced at around $1.2 million. As a result, as one industry association noted, "*The time and cost associated with [switching] tend to lock customers into a particular foundry*." [27]

There are also other, though perhaps less visible, resource costs of beginning a new foundry relationship. In brief, negotiating a new supplier's agreement is not trivial. A vast amount of information has to be exchanged about the specific details of the wafer design and the technological requirements, on the foundry's part, needed to manufacture the chips. These talks can absorb much attention by senior management.[28]

Second, if a fabless firm reneged on an arrangement with a foundry by initiating production with a lower-price supplier, we conjecture that it would likely suffer extensive reputational costs. As seen in Figure 5, among overseas foundries, there is *consistently* a dominant producer of the frontier technology -- TSMC. TSMC stands to gain substantially over time from this leading

[27] This quotation is taken from the Common Platform alliance. This industry group consists of a few large chip manufacturers, such as IBM and Samsung. In an effort to reduce these costs, the alliance has begun to advocate for a "common platform" that would standardize production technology for certain high-volume wafers (such as 90nm chips). However, this alliance has not yet had a material impact on, for instance, standardizing mask sets (McGregor, 2007). We are also grateful to Ross Goodman for insightful discussions on this topic.

[28] A summer 2002 article in the magazine, *Electronics Design Chain*, gives a good example of this. It discusses how a new fabless firm began a relationship with TSMC. The executives at the fabless firm stressed the importance of establishing clear channels of communication so that the foundry management knew precisely who to contact regarding details of the wafer design and production needs. Otherwise, if there were a single liaison in the fabless firm fielding all questions, that person "will be getting 250 communications per day" from the foundry.

position, so it has a strong incentive to "punish" fabless designers who interrupt a contract and re-source their product. This would make it very difficult for the fabless firms to purchase frontier technologies in the future.

To summarize, there are monetary and reputational considerations that appear to discourage the movement of a particular design from one foundry to another, even in the presence of substantial differences in price for the relevant process technology at a given moment in time. If a fabless firm's product requires leading-edge process technology to be competitive, it will contract with a leading-edge foundry such as TSMC in Taiwan. Even after lower-price suppliers begin manufacturing that process technology more cheaply, the large fixed costs of moving production to a new foundry make switching almost prohibitively expensive. Instead, movement toward lower-price suppliers occurs at the introduction of new products. This close relationship between substitution across suppliers and the introduction of new products will be important for our discussion of official price measures in the following section.

7. IMPLICATIONS FOR OFFICIAL MEASURMENT PRACTICES

In this paper, we have presented new data on prices for processed semiconductor wafers, a key component in the semiconductor supply chain. At this point, it would be natural to compare our estimates to the official source of international price data. However, there is no analogue to our indexes in the Bureau of Labor Statistics' (BLS) International Price Program (IPP). The closest product category for which the IPP produces an index differs from our indexes in three respects: i) our data shows prices for wafer fabrication services purchased by chip makers around the world, while the BLS samples only U.S. importers, ii) Our data include only arms-length transactions while the IPP includes these and intra-company transfers, and iii) the IPP index is much more aggregated than our indexes, as the IPP includes finished semiconductor chips and microassemblies in addition to processed wafers and die.[29]

This does not mean, however, that our data are uninformative with regard to official indexes. Consider, for instance, the question of how much fabless firms shift toward lower-price suppliers. Our data can shed light on whether this phenomenon is adequately reflected in the IPP. The GSA and other proprietary data (namely, iSuppli's foundry database) clearly reveal

[29] The IPP index to which we refer is Harmonized System Code (HSC) 8542, electronic integrated circuits and microassemblies.

shifts toward lower-price wafer suppliers. Tabulations from the IPP indicate that domestic buyers virtually never switch providers for the same physical product. Specifically, between 2004 and 2007, no such switches were observed in the IPP within the product category most related to wafers. Yet we know that fabless firms continued to substitute toward low-price suppliers throughout this period (see Table 4).[30]

To understand how this difference might arise, it is necessary to discuss how the IPP index is constructed. We focus specifically on two features of IPP's methodology.

First, there is a crucial difference between how price inflation is measured in this paper and in the IPP. To see this, suppose there are *N* respondents in the IPP survey and assume each respondent reports prices for *M* products. The IPP rolls a separate price index for each physical product imported by each individual respondent – a total of *N*x*M* indices, one for each unique *buyer-product* pair. An appropriately weighted average of these indices is then calculated to deliver the overall price index. In our case, the order is reversed: for a given physical product, we first calculate the price *level* averaged across the universe of importers of fabrication services. This would give us *M* observations in each period. We then roll a price index from these data on average price levels and then average across products. Our unit of analysis is, in other words, the *physical* product; the buyer is not treated as an attribute of the good.

This difference in procedures may help account for the difference in results. The result from the IPP may reveal that there is quite a lot of substitution along the extensive margin, to use the language from Section 6. In other words, more fabless firms choose to source new products from SMIC rather than from TSMC.[31] In our data, these new fabless buyers would show up as recipients of relatively low prices and thus drive the average wafer price down. This follows directly from the fact that the unit of analysis in our data is the wafer, and prices for that wafer are averaged across buyers. But in the IPP, the unit of analysis is the buyer-wafer pair, so a new

[30] We thank Ben Mandel for these calculations. The IPP staff also kindly provided us with similar calculations over the last two years. They report that less than 0.25 percent of the items in the IPP's HSC 8542 sample reported a shift in source country.

[31] This contrasts with the example given in Section 6 regarding Qualcomm, where the firm continued to source leading-edge products from TSMC but purchased trailing-edge products from SMIC. In principle, this shift for a given product should be captured in the IPP.

index would be initiated for a new pair. By construction, this approach captures price changes *after* the buyer-product pair forms; it does not measure the initial drop in price.[32]

The second noteworthy methodological issue is related to the definition of a "product". Whereas we take an explicit stand on the key price-forming attributes of wafers, the IPP largely leaves it to the respondents to define the parameters of the good. We claim that our approach accurately categorizes the good from the perspective of the producer. That is, if the fabless firm alters a chip design but the underlying wafer production technology remains fixed, the *input* (namely, the wafer fabrication service) is the same even if the *output* (the final semiconductor) is somewhat different. This is the most appropriate approach, we argue, if the imported product is treated as a material input that is, in turn, marketed and sold as a finished good by the fabless firm. In this case, the objective is to price the input, wafer fabrication, not quality-adjust the finished good, which includes the value of the design performed by the U.S. fabless firm.

It is worth pausing to note that this line of reasoning suggests a more general point relevant for designing input price surveys.[33] In a world where outsourcing is becoming increasingly common, a survey that gathers data on input prices from buyers must exercise care to ensure that the reported price applies to a given quality input and excludes the value-added associated with the original buyer's design.

The IPP's approach with regard to product definition may explain the lack of observed production switching. Without a clear parameterization of the product space, any slight alterations in the respondent's report that coincide with a change in source country may lead the IPP staff to determine that a new buyer-product index is needed, even if the underlying wafer fabrication technology is fixed. In other words, the IPP may, in practice, also treat the *location of the seller* as a product attribute, since in this market shifts across sellers almost always occur with the introduction of a new chip design.[34] Nakamura and Steinsson's (2011) analysis of the IPP micro data persuaded them that a product in the IPP is essentially a "contract between a particular buyer and seller." In that case, the IPP will, by construction, not capture a shift in market share towards lower-price suppliers.

---

[32] Ideally, we would replicate the IPP's procedure and compute buyer-seller price indexes in the GSA data as a means of investigating how the results from this approach contrast with our baseline estimates. However, this is not feasible because the GSA data do not include firm identifiers.

[33] As noted in footnote 22, the economic measurement community is actively debating the merits and mechanics of producing an input price index.

[34] Gopinath and Rigobon (2008) note that the BLS' *intent* is not to treat the identity of the seller as a "price-forming" characteristic.

8. RELATED RESEARCH

At this point, it is instructive to place our work in the context of recent literature. The present paper touches on a few related lines of research. First, there is considerable interest in developing more systematic accounts of the overseas supply chain. Several studies have recently begun to address this.[35] These efforts often rely on relatively aggregated data on trade flows and input-output tables to construct measures of a country's exports of intermediates (see, for instance, Hummels, Ishii, and Yi, 2001).[36] Our paper has taken a complementary approach by focusing on a particular market and using transaction-level data to document product sourcing.

Second, some researchers have explored the effects of offshoring on final product and factor prices. One issue relevant to inflation-targeting central banks, for instance, is the effect of the entry of overseas intermediate-goods suppliers on the price inflation faced by domestic final-goods producers (see Bernanke, 2007).[37] Work along these lines (Kamin, Marazzi, and Schindler, 2006) has used published data on import prices and trade flows to estimate the association between the Chinese share of U.S. imports in a particular sector and import price inflation in that sector. We have, again, taken a complementary approach. The present paper trades off comprehensiveness for more detailed transaction-level price data available for a particular market. As emphasized, these data permit a rigorous adjustment for dispersion in product quality across supplying countries. Our data has also enabled us to carry out an alternative, and perhaps more direct, accounting of the effect of product sourcing on average prices.

Lastly, there is increasing interest in the question of whether official published data adequately capture the effect of offshore outsourcing on the prices faced by domestic buyers. In an important recent paper, Houseman et al. (2011) use the micro data underlying the U.S. import

[35] See the papers and panel discussions in several recent conferences, including "Measurement Issues Arising from the Growth of Globalization," organized by the W.E. Upjohn Institute for Employment Research and the National Academy of Public Administration, November 2009; "Measurement Issues in Trade" session of the Federal Economic Statistics Advisory Committee, December 2010; and "The Fragmentation of Global Production and Trade in Value-Added - Developing New Measures of Cross Border Trade," World Bank, June 2011.

[36] The input-output tables are needed in order to assign an end use to the imports; the trade data generally do not report whether a product is a final or intermediate good.

[37] Presumably, the introduction of lower-price intermediates dampens international price changes and, ultimately, leads to lower domestic price inflation, which is, of course, of great interest to central bankers.

price index to estimate the impact of shifting sourcing patterns on the prices of inputs into the U.S. manufacturing sector. To control for quality differences, the paper isolates the price changes that occur when the U.S. importer reports a different source country for the same product, defined as a 10-digit Harmonized system (HS) code.

Although the microdata afford considerable product-level information, the proprietary data we use in this paper permit a much more detailed analysis that is arguably necessary in markets with highly differentiated goods. For example, all of the processed semiconductor wafers in our sample fall under the same 10-digit HS code. Thus, Houseman et al. implicitly treat all wafers as the same product in their analysis. As shown in Table 3, when controlling for technological characteristics, Chinese foundries charge approximately 27% less than Taiwanese foundries for identical wafers. Without technological controls, the Chinese wafer price is 66% lower than Taiwan's on average, reflecting Chinese foundries' focus on trailing-edge products. Thus, even when using microdata at the 10-digit HS level, substantial quality differences remain across countries that may confound the analysis conducted by Houseman et al.[38] This issue is likely to be less severe in sectors with more detailed HS codes and slower technological change, but the findings just discussed demonstrate the value of using highly detailed transaction-level data.

9. CONCLUSION

In this paper, we have investigated pricing and product sourcing in the offshore market for wafer fabrication services. Using detailed, transaction-level data, we detect significant shifts in production toward relatively new suppliers, such as China. We also found notable variation in prices across suppliers for observationally equivalent goods, with China specifically offering deep discounts. A formal accounting of the effect of these new suppliers on average wafer prices indicates that they accelerated the rate of price decline by over 10 percent, pushing the rate of deflation from 10.4% to 11.8%. We argue that the large price differentials we observe are at least partly due to significant frictions in this product market that raise the cost of switching across suppliers and thereby enforce long-lived buyer-supplier relationships.

[38] Klier and Rubenstein (2009) document shifts in auto parts production toward Mexico and China, and rely on estimates from industry publications to demonstrate lower production costs for aluminum wheels in Mexico relative to the U.S.

Going forward, there are at least two important avenues for related research. First, we conjecture that frictions are not unique to the market for wafer fabrication services. Although the difficulties of re-sourcing a product may be particularly large in this sector, the costs involved in initiating a new overseas relationship are likely to be appreciable in most cases. To this end, analyses of price dispersion in other offshoring markets would be revealing. Systematic variation in the prices of observationally equivalent goods argues strongly, in our view, for producing an official input price index along the lines of the "within-technology" index we introduced in Section 5. As a result, additional studies would be a valuable input into the deliberations of official statistical agencies as they consider the best way of measuring the effect of the evolving global supply chain on international and domestic prices.

Second, though we have focused on one part of the overseas supply chain, it may be possible to also use our data to examine other stages in the semiconductor production process. By doing this we may be able to infer the implicit price of product design, which is primarily performed domestically in the United States and Europe. In preliminary work, we have used GSA data to first price the other stages of overseas production, including assembly and testing. Combining these estimates of materials inputs prices with additional information on the price of *finished* semiconductors, it is, in principle, possible to "back out" the price associated with chip design. Pricing research and development is a top priority for U.S. statistical agencies (Copeland, Medeiros, and Robbins, 2007).

REFERENCES

Aizcorbe, Ana (2002), "Price Measures for Semiconductor Devices," Finance and Economics Discussion Series 2002-13. Washington: Board of Governors of the Federal Reserve System (U.S.), March.

Aizcorbe, Ana, Steven Oliner, and Daniel Sichel (2006), "Shifting Trends in Semiconductor Prices and the Pace of Technological Progress," Finance and Economics Discussion Series 2006-44. Washington: Board of Governors of the Federal Reserve System (U.S.), December.

Bils, Mark (2009). "Do Higher Prices for New Goods Reflect Quality Growth or Inflation?", *Quarterly Journal of Economics*, vol. 124(2), pp. 637-675, May.

Boskin, Michael J. _et al (1998). "Toward a More Accurate Measure of the Cost of Living: Final Report to the Senate Finance Committee from the Advisory Commission to Study the Consumer Price Index," in Dean Baker, ed., *Getting Prices Right: The Debate over the Consumer Price Index.* Armonk, N.Y. and London: Sharpe, pp. 5-77.

Brown, Clair and Linden, Greg (2006). "Offshoring in the Semiconductor Industry: Historical Perspectives," in Lael Brainard, Susan M. Collins, eds., *Brookings Trade Forum: 2005.* Washington, DC: Brookings Institution Press, .

Campa, Jose, and Linda S. Goldberg (1997). "The Evolving External Orientation of Manufacturing: A Profile of Four Countries," *Federal Reserve Bank of New York Economic Policy Review*, vol. 3 (2), pp. 53-81.

Clenndenin, Mike, and Junko Yoshida (2002). "Game, Set, Match for Foundries? Even Top-Tier IC Makers Partnering as CMOS Process Rigors Intensify". *Electronic Engineering Times.* March 11.

Copeland, Adam, Gabriel W. Medeiros, and Carol A. Robbins (2007). "Estimating Prices for R&D Investment in the 2007 R&D Satellite Account," Bureau of Economic Analysis working paper.

Dean, Judith, K. C. Fung, and Zhi Wang (2011). "Measuring Vertical Specialization: the Case of China," *Review of International Economics,* .

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

Doms, Mark, Ana Aizcorbe, and Carol Corrado (2003), "When Do Matched-Model and Hedonic Techniques Yield Similar Measures?" Working Papers in Applied Economic Theory 2003-14. San Francisco: Federal Reserve Bank of San Francisco, June.

Electronic Engineering Times (1998). "U.S., China at Odds on Fab-Gear Export". *Electronic Engineering Times.* April, 1998.

Feenstra, Robert C., Benjamin R. Mandel, Marshall B. Reinsdorf, and Matthew J. Slaughter (2009), "Effects of Terms of Trade Gains and Tariff Changes on the Measurement of U.S. Productivity Growth," NBER Working Papers 15592. Cambridge, Mass.: National Bureau of Economic Research, Inc., December.

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

Global Foundries Fact sheet: Fab 2 module 1. Global FoundriesAvailable from http://globalfoundries.com/newsroom/2009/Fab_2_fact_sheet.pdf [cited March 2009].

Golub, Stephen S., Ronald W. Jones, and Henryk Kierzkowski (2007). "Globalization and Country-Specific Service Links," *Journal of Economic Policy Reform*, vol. 10 (2), pp. 63-88.

Gopinath, Gita, and Roberto Rigobon (2008). "Sticky Borders," *Quarterly Journal of Economics*, vol. 123 (2), pp. 531-75.

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

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

Grossman, Gene M., and Esteban Rossi-Hansberg (2008). "Trading Tasks: A Simple Theory of Offshoring," *American Economic Review*, vol. 98 (5), pp. 1978-97.

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

Hummels, David, Jun Ishii, and Kei-Mu Yi (2001). "The Nature and Growth of Vertical Specialization in World Trade," *Journal of International Economics*, vol. 54 (1), pp. 75-96.

Hurtarte, Jorge S., Evert A. Wolsheimer, and Lisa M. Tafoya (2007). *Understanding Fabless IC Technology*. Oxford: Elsevier.

IC Knowledge (2001). *"Can the semiconductor industry afford the cost of new fabs?"*. IC Knowledge, www.icknowledge.com/economics/fab_costs.html (accessed March 2010).

Kamin, Steven B., Mario Marazzi, and John W. Schindler (2006). "The Impact of Chinese Exports on Global Import Prices," *Review of International Economics,* vol. 14 (2), pp. 179-201.

Klier, Thomas and James Rubenstein (2009). “Imports of Intermediate Parts in the Auto Industry – A Case Study.” Prepared for the conference, “Measurement Issues Arising from the Growth of Globalization,” Washington, D.C., November.

Kumar, Rakesh (2008). *Fabless Semiconductor Implementation.* New York: McGraw-Hill Professional.

———. (2007). "The Business of Scaling," *IEEE Solid-State Circuit Society News,* vol. 12 (1), pp. 22-27.

LaPedus, Mark (2006). "Qualcomm's Foundry Push Closes Gap with TI". *Electronic Engineering Times.* www.eetimes.com/electronics-news/4064790/Qualcomm-s-foundry-push-closes-gap-with-TI,September 8.

———. (2006). "SMIC to Manufacture Qualcomm Chips in BiCMOS". *Electronic Engineering Times.* www.eetasia.com/ART_8800427539_480200_NT_4bd09293.HTM,July 31.

McGregor, Jim (2007). *"The common platform technology: A new model for semiconductor manufacturing"*. Scottsdale, Ariz.: In-Stat, .

Moore, Gordon E. (1965). "Cramming More Components onto Integrated Circuits". *Electronics.*

Nakamura, Emi, and Jon Steinsson (2011), "Lost in Transit: Product Replacement Bias and Pricing to Market," Columbia University, June.

Oliner, Stephen D., and Daniel E. Sichel (2000). "The Resurgence of Growth in the Late 1990s: Is Information Technology the Story?" *Journal of Economic Perspectives,* vol. 14 (4), pp. 3-22.

Oliner, Stephen D., Daniel E. Sichel, and Kevin J. Stiroh (2007). "Explaining a Productive Decade," *Brookings Papers on Economic Activity,* (1), pp. 81-137.

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

Scott, A. J., and D. P. Angel (1988). "The Global Assembly-Operations of U.S. Semiconductor Firms: A Geographical Analysis," *Environment and Planning A,* vol. 20 (8), pp. 1047-67.

Semiconductor Industry Association (2007). *International Technology Roadmap for Semiconductors.* Austin, Texas: International SEMATECH.

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.* Upper Saddle River, New Jersey: Prentice Hall PTR.

Yeats, Alexander J. (2001). "Just How Big Is Global Production Sharing?" in Sven W. Arndt, Henryk Kierzkowski, eds., *Fragmentation: New production patterns in the world economy.* Oxford and New York: Oxford University Press, pp. 108-143.

Yi, Kei-Mu (2003). "Can Vertical Specialization Explain the Growth of World Trade?" *Journal of Political Economy,* vol. 111 (1), pp. 52-102.

Figure 1: Moore's Law – Intel Processors

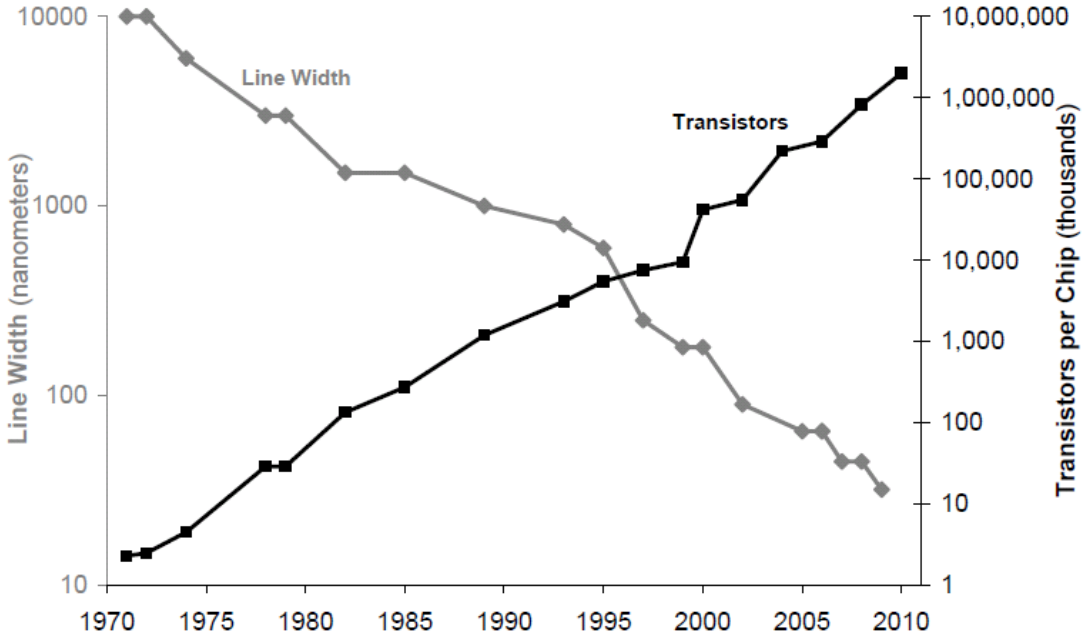

Sources: http://www.intel.com/technology/timeline.pdf
http://www.intel.com/pressroom/kits/quickreffam.htm

Figure 2: Technology Cycle – TSMC Sales by line width

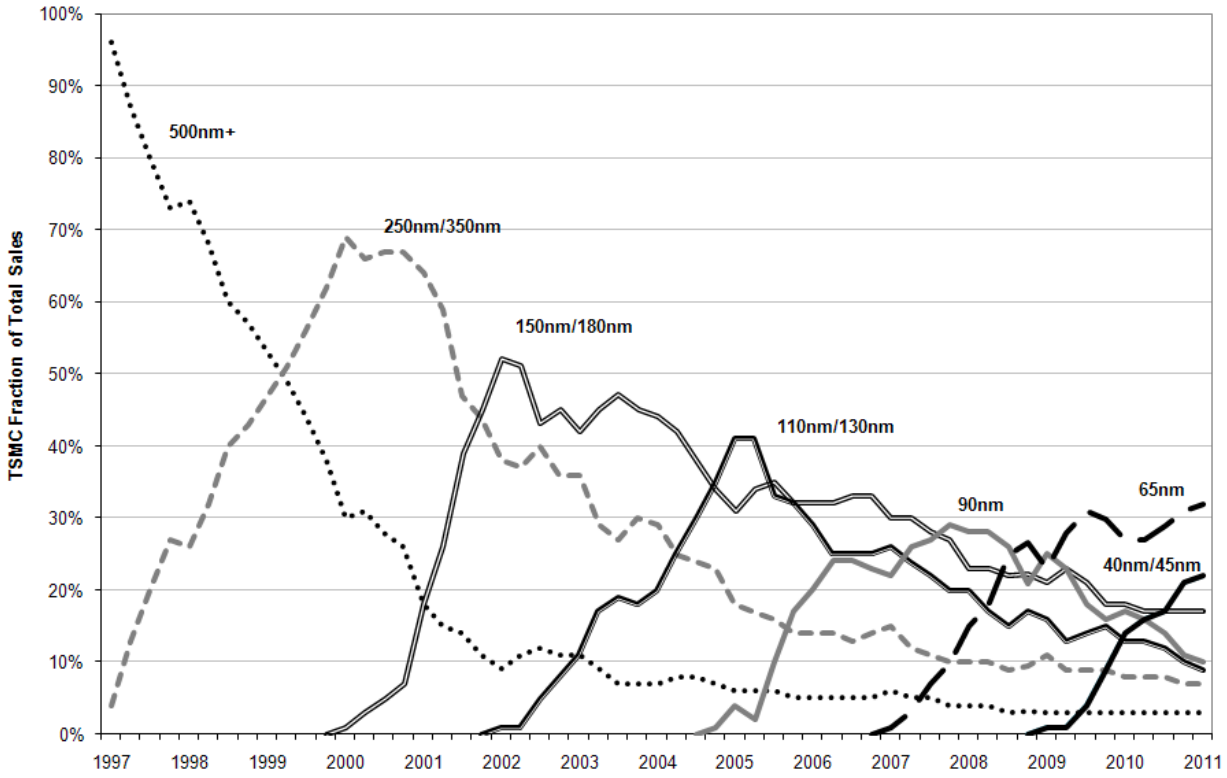

Sources: TSMC quarterly reports

Figure 3: Growth of the Fabless Business Model

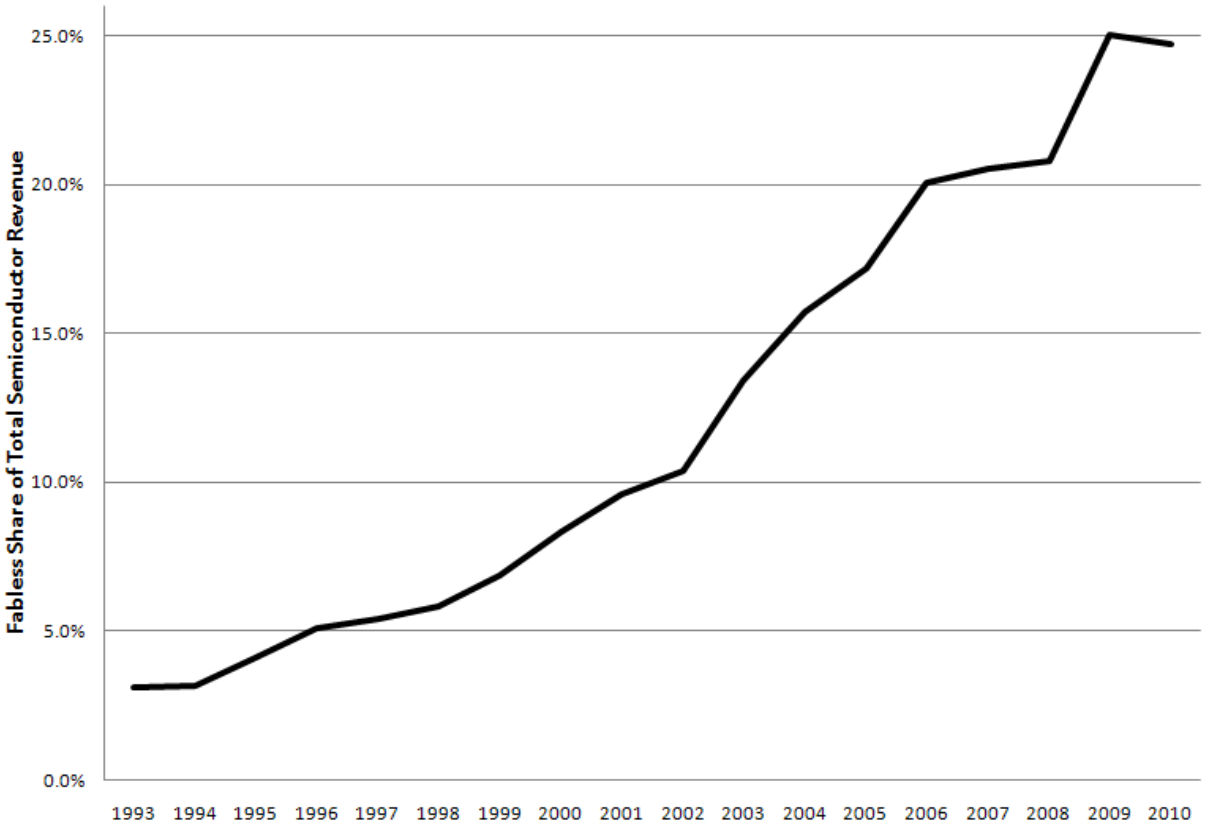

Sources: Global Semiconductor Association (GSA) and Semiconductor Industry Association (SIA)

Figure 4: Index Calculation for Selected Technology

300mm wafer diameter, 130nm line width

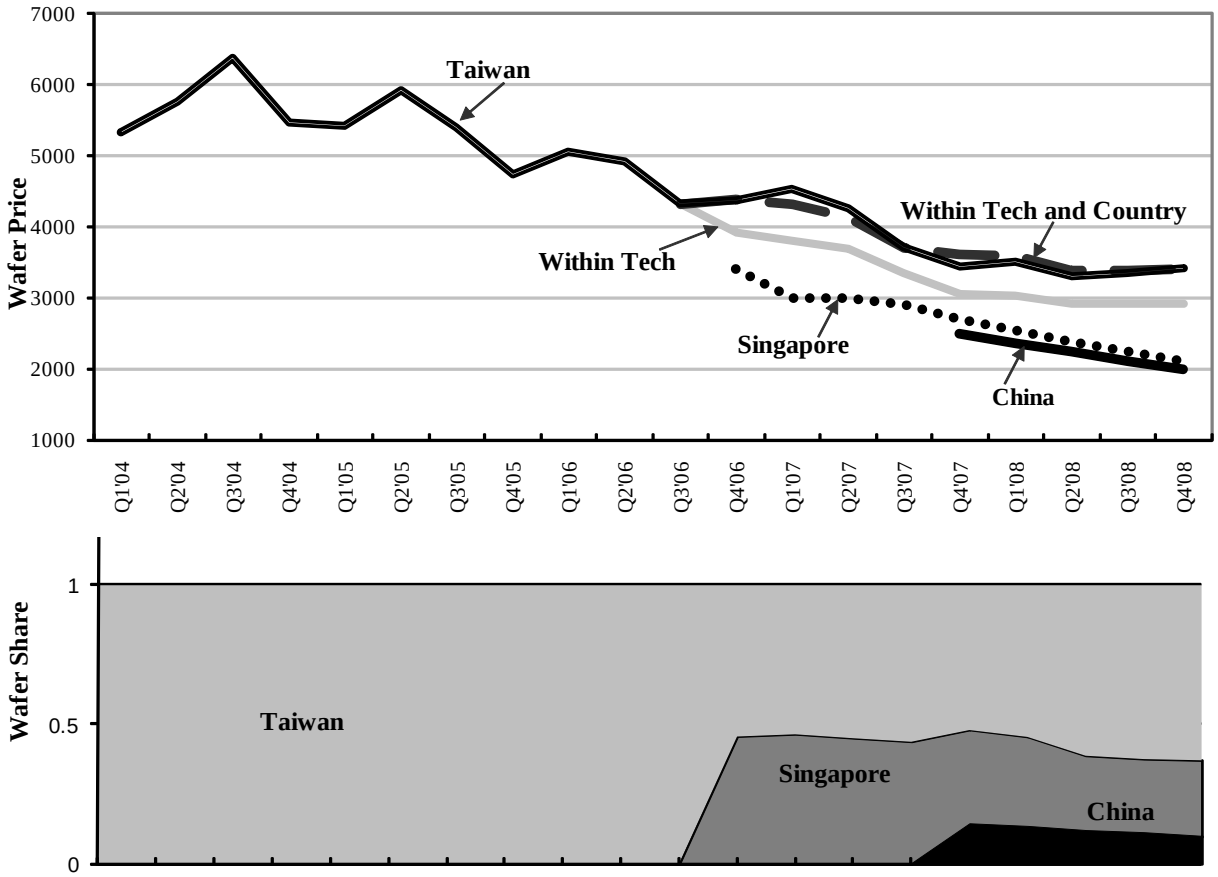

Source: Authors' calculations using Global Semiconductor Alliance (GSA) Wafer Fabrication Pricing Report and iSuppli Pure Play Foundry Market Tracker

Figure 5: Foundry Capacity by Country and Technology

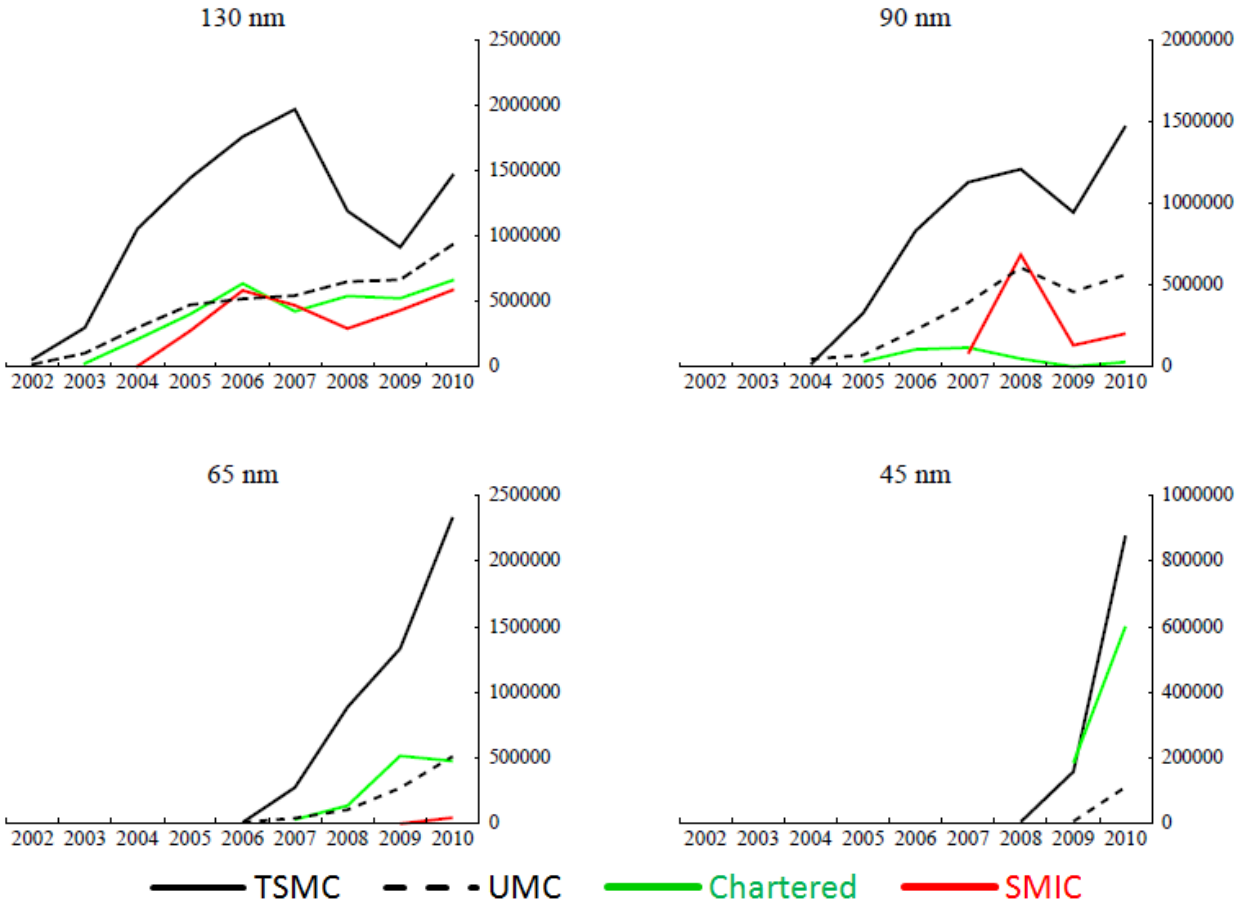

Note: Capacity in 8-inch equivalent starts per month
Source: iSuppli Pure Play Foundry Market Tracker

Figure 6: Significant Fabless-Foundry Relationships

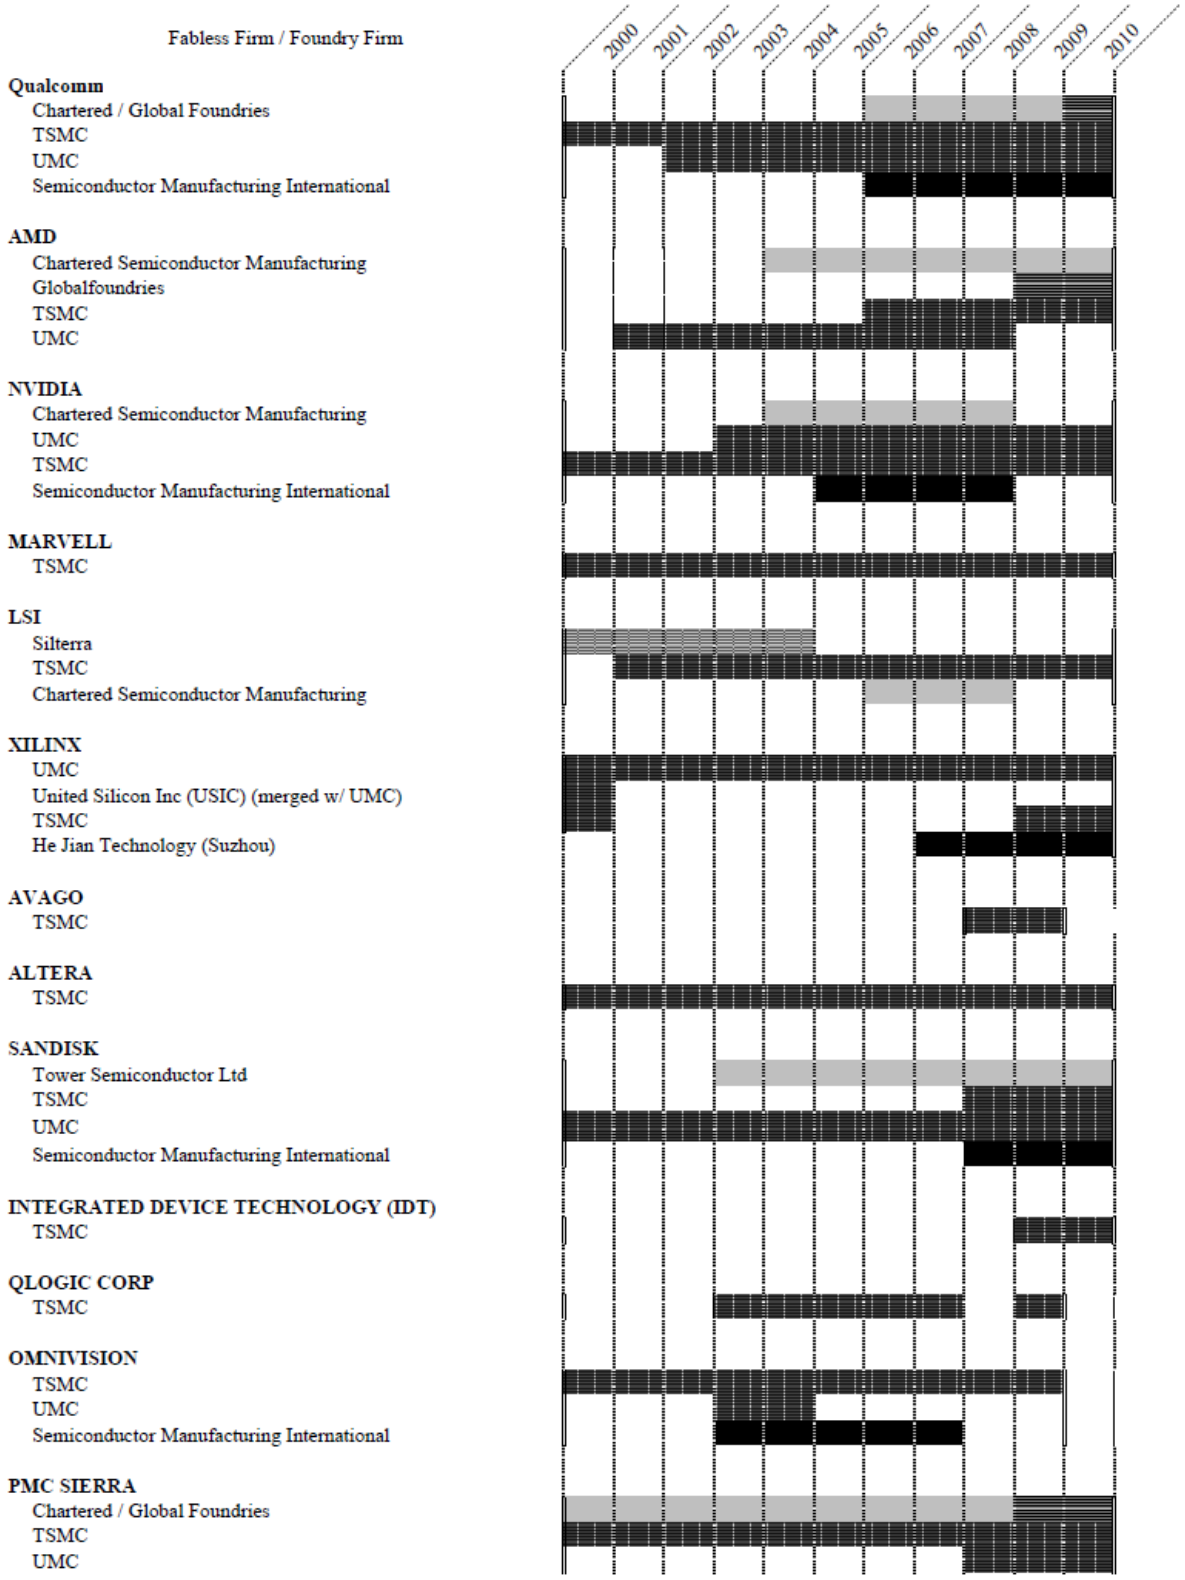

Source. SEC financial filings.
Location of foundry firm headquarters: grid = Taiwan, black = China, grey = Singapore, horizontal = United States, diagonal = Malaysia

Table 1: Foundry Capacity by Country

	Global Capacity	Taiwan	China	Singapore	Europe	USA	Japan	S. Korea	Malaysia
2002	946	60.3%	7.8%	5.8%	7.7%	5.8%	7.5%	3.2%	2.0%
2003	1072	55.4%	12.5%	6.1%	7.8%	5.5%	6.9%	3.4%	2.4%
2004	1344	53.4%	16.5%	6.9%	6.7%	4.4%	5.7%	3.3%	3.2%
2005	1629	51.0%	19.8%	8.1%	5.8%	3.9%	4.8%	3.2%	3.4%
2006	1821	49.9%	21.1%	9.0%	5.2%	3.6%	4.4%	3.5%	3.3%
2007	2025	50.2%	20.6%	9.8%	5.0%	3.3%	3.9%	3.8%	3.4%
2008	2267	50.3%	19.7%	11.4%	5.5%	2.9%	3.1%	3.8%	3.3%
2009	2396	49.5%	19.7%	11.6%	6.9%	2.8%	2.8%	3.6%	3.1%
2010	2630	47.3%	19.1%	13.1%	8.9%	2.7%	2.7%	3.4%	2.9%

Note: Includes pure-play foundries only. Global capacity measured as 8-inch equivalent wafer starts per week
Source: iSuppli Pure Play Foundry Market Tracker

Table 2: Dropped Observations

Total observations	7436
Used in analysis	5423
Dropped	2013
Missing:	
foundry location	794
wafers purchased	34
Other reason:	
engineering run	778
not CMOS	651
100mm wafer	3
inconsistent	1

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

Table 3: Descriptive Statistics

			Yearly Means				
	Mean	Std. Dev	2004	2005	2006	2007	2008
Price Per Wafer ($)	1528.76	1150.92	1,516.34	1,552.43	1,453.03	1,519.16	1,614.90
Number of Wafers Contracted	2408	7793	2037	2422	2081	2696	2833
Layers							
Metal Layers	4.81	1.84	4.22	4.56	4.78	5.07	5.33
Mask Layers	25.31	7.65	22.71	24.15	25.37	26.36	27.62
Polysilicon Layers	1.21	0.43	1.26	1.21	1.21	1.19	1.17
Wafer Size							
150 mm or less	0.15	0.35	0.18	0.18	0.15	0.12	0.11
200 mm	0.75	0.43	0.79	0.77	0.78	0.75	0.68
300 mm	0.10	0.30	0.04	0.05	0.07	0.13	0.21
Line Width							
65 nm	0.00	0.07	0.00	0.00	0.00	0.00	0.02
90 nm	0.03	0.16	0.00	0.01	0.01	0.04	0.07
130 nm	0.20	0.40	0.10	0.14	0.21	0.25	0.31
150 nm	0.03	0.17	0.04	0.04	0.02	0.03	0.02
180 nm	0.27	0.44	0.26	0.28	0.28	0.27	0.24
250 nm	0.13	0.34	0.15	0.16	0.13	0.13	0.09
older vintage	0.34	0.44	0.45	0.37	0.35	0.28	0.25

5423 Observations

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

Table 4: Descriptive Wafer Price Regression Results

dependent variable: log of price per wafer

Variable	Coefficient	Std. Err.	t-Stat
Foundry Location			
China	-0.267	0.019	-14.37
Europe	0.084	0.018	4.68
Japan	-0.104	0.022	-4.72
Korea	-0.077	0.018	-4.20
Malaysia	-0.318	0.054	-5.84
Singapore	-0.061	0.012	-5.23
United States	0.201	0.014	14.38
Other	-0.095	0.024	-3.96
Wafer Size			
150 mm	-0.363	0.015	-24.87
300 mm	0.603	0.014	44.62
Line Width			
≥ 1000 nm	-0.655	0.032	-20.58
800 nm	-0.404	0.026	-15.56
600 nm	-0.401	0.022	-18.49
500 nm	-0.391	0.019	-20.97
350 nm	-0.241	0.013	-17.93
250 nm	-0.101	0.011	-8.68
150 nm	0.165	0.020	8.44
130 nm	0.347	0.012	28.12
90 nm	0.568	0.025	22.85
65 nm	0.791	0.047	16.79
Number of Metal Layers	0.055	0.004	14.180
Number of Polysilicon Layers	0.044	0.008	5.150
Number of Mask Layers	0.010	0.001	11.630
Epitaxial Layer Indicator	0.105	0.009	11.660
log Number of Wafers Contracted	-0.055	0.002	-33.830
Constant	6.738	0.030	223.130
R-squared	0.884		
Observations	5423		

Specification also includes quarterly indicator variables
non-CMOS production not included
Baseline case (omitted category) is Taiwan, 200mm, 180nm

Table 5: Price Index Results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		Within Technology and Country					Within Technology
Quarter	Overall	Taiwan	China	Singapore	USA	Europe	
2004Q1	100.0	100.0	100.0	100.0	100.0	100.0	100.0
2004Q2	104.0	100.8	98.5	127.9	108.9	134.7	102.6
2004Q3	102.4	103.0	90.6	106.5	108.6	114.4	100.7
2004Q4	94.3	95.0	81.0	97.1	98.4	94.4	92.5
2005Q1	93.8	91.7	99.8	106.1	93.5	99.1	91.0
2005Q2	92.5	90.9	93.9	95.0	97.1	105.8	89.5
2005Q3	91.1	88.1	87.6	97.3	142.7	114.2	88.1
2005Q4	90.9	89.3	87.7	106.3	129.6	88.7	87.7
2006Q1	82.9	81.1	84.6	90.6	79.1	94.0	80.0
2006Q2	82.8	81.6	77.4	71.4	93.7	102.8	79.9
2006Q3	77.9	76.6	83.0	71.5	86.1	88.4	75.2
2006Q4	76.4	73.8	92.1	81.7	81.6	77.8	72.9
2007Q1	76.6	76.1	77.5	73.0	91.6	81.7	73.5
2007Q2	72.7	71.8	74.5	81.4	82.0	88.3	67.9
2007Q3	69.7	68.5	71.2	74.0	76.5	98.2	65.1
2007Q4	67.2	64.3	80.3	79.8	78.9	109.9	62.7
2008Q1	64.0	60.7	71.2	85.1	87.2	129.4	59.7
2008Q2	62.3	56.3	69.9	149.2	139.6	120.4	58.3
2008Q3	69.2	66.2	71.3	149.2	122.3	112.8	64.0
2008Q4	62.9	59.4	64.9	122.1	103.0	119.5	57.1
Year							
2004	100.2	99.7	92.5	107.9	104.0	110.9	98.9
2005	92.1	90.0	92.3	101.2	115.7	102.0	89.1
2006	80.0	78.3	84.3	78.8	85.1	90.8	77.0
2007	71.5	70.2	75.9	77.1	82.3	94.5	67.3
2008	64.6	60.7	69.3	126.4	113.0	120.5	59.8
Avg. Yearly Growth '04-'08	-10.4%	-11.7%	-7.0%	4.0%	2.1%	2.1%	-11.8%