QUALITY-ADJUSTED PRICE MEASUREMENT: A NEW APPROACH WITH EVIDENCE FROM SEMICONDUCTORS

David M. Byrne, Brian K. Kovak, and Ryan Michaels*

Abstract—Many markets exhibit price dispersion across suppliers of observationally identical goods. Statistical agencies typically assume this dispersion reflects unobserved quality, so standard price indexes do not incorporate price declines when buyers substitute toward lower-price suppliers. We show that long-run price differences across suppliers can be used to infer unobserved quality differences and propose an index that accommodates quality-adjusted price dispersion. Using transaction-level data on contract semiconductor manufacturing, we document substantial quality-adjusted price dispersion and confirm that a standard index is biased above our proposed index.

I. Introduction

ACCURATE measures of market prices are important in all branches of applied economic analysis. One of the most persistent challenges in the practice of price measurement is accounting for differences in product quality. Such differences may involve products' characteristics, as well as aspects of the overall transaction, such as customer service or timely delivery (Carlton, 1983). Price indexes seek to measure average price growth, controlling for differences in product quality across goods and over time. This is quite challenging in practice because of difficulties in observing detailed physical product attributes and other less tangible characteristics of transactions.

While these challenges have been known for decades, they have recently taken on particular significance in markets for intermediate inputs.[1] Such markets are characterized by increased internationalization of production chains and shifts toward relatively low-price suppliers in developing countries such as China (Hummels, Ishii, & Yi, 2001). Moreover, an increasing number of "factoryless manufacturers" outsource product fabrication activities altogether (Bayard, Byrne, & Smith, 2015; Bernard & Fort, 2013). These developments have led to greater substitution across suppliers of intermediate inputs, both domestic and international. In this context, failure to accurately estimate differences in quality across suppliers will lead to biased import quantity and productivity measures.

In this paper, we study the problem of price index construction when new suppliers and incumbents may charge different prices for goods of identical quality. Standard price indexes, known as matched-model indexes, typically assume that the law of one price holds, which rules out the possibility of price dispersion for identical goods. In doing so, they omit price declines when buyers shift toward suppliers offering discounts, and hence are biased upward. We propose a simple means to infer differences in unobserved quality based on long-run price differences across suppliers. Early in a product's life cycle, market frictions can impede arbitrage across incumbents and new suppliers, generating price dispersion for goods of similar quality. Yet the influence of these frictions tends to dissipate over time, so that in the long run, several years after entry, price differences largely reflect quality differences. We use this insight to infer unobserved quality and thereby construct a novel price index that both accounts for quality differences across suppliers and allows for deviations from the law of one price.

We apply our method to the contract semiconductor manufacturing industry, using new transaction-level data that include information on prices and all relevant physical characteristics of each product. These data allow us to compare prices for technically identical products across suppliers located in different countries. We find large price differences across suppliers; for example, Chinese producers charged 17% less on average than firms in market leader Taiwan for otherwise identical products. Moreover, the price differences are especially large early in each product's life cycle but partially converge later on, consistent with a setting in which frictions bind less over time. Together, these patterns suggest the presence of cross-supplier price variation that would confound matched-model price indexes. In fact, a standard matched-model index falls almost 1 percentage point per year more slowly than our proposed quality-adjusted index. This substantial upward bias in the standard approach is large enough to meaningfully bias productivity measures and other government statistics.

Although our empirical setting focuses on substitution between suppliers of imported intermediates, the proposed index applies broadly to environments involving the entry of low-price sellers. Thus, our approach also applies to the domestic retail context, where outlet substitution bias resulting from omitting new entrants is a long-standing

Received for publication February 26, 2014. Revision accepted for publication July 13, 2016. Editor: Gordon Hanson.

* Byrne: Federal Reserve Board; Kovak: Carnegie Mellon University and NBER; Michaels: Federal Reserve Bank of Philadelphia.

An earlier version of this paper was prepared for the Measurement Issues Arising from the Growth of Globalization conference sponsored by the W. E. Upjohn Institute for Employment Research and the National Academy of Public Administration and funded by the Alfred P. Sloan Foundation. Thanks to Yong-Gyun Choi, Manuel Gomez, Steven Paschke, and Fabio Rueda for excellent research assistance. We thank Bill Alterman, Wenjie Chen, Ken Flamm, Erica Fuchs, Ross Goodman, Susan Houseman, Ben Keys, Rebecca Lessem, Nina Pavcnik, Marshall Reinsdorf, Jodi Shelton, Lowell Taylor, Sonya Wahi-Miller, Kim Zieschung, and various seminar participants for helpful discussions, and Chelsea Boone, Peter Dziel, Wade Giles, Len Jelinek, and Andrzej Strojwas for help with data and background information on the industry. Remaining errors are our own. This paper reflects our own views and should not be attributed to the Board of Governors of the Federal Reserve System or other members of its staff.

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/REST_a_00640.

[1] It has long been recognized that price differences reflect both quality variation and quality-adjusted price dispersion. The 1961 Price Statistics Review Committee, chaired by George Stigler, described the measurement of quality variation as the most significant challenge in index construction (National Bureau of Economic Research Incorporated, 1961).

The Review of Economics and Statistics, May 2017, 99(2): 330–342
© 2017 by the President and Fellows of Harvard College and the Massachusetts Institute of Technology
doi:10.1162/REST_a_00640

concern.[2] Our proposed index is quite feasible for measurement agencies to implement, as it uses only price and quantity information, which can be collected using existing surveys.

Our empirical findings contribute to a growing literature investigating the measurement implications of the globalization of supply chains, showing that shifts to low-priced offshore suppliers drive systematic bias in standard import price measures.[3] Our quality inference procedure also provides an alternative to a common approach in the international trade literature, which infers relative quality across suppliers from differences in market share conditional on price.[4] While this approach is suitable for measuring long-run trends in quality or differences in quality for aggregate industries across countries, one should be cautious when studying short time spans or narrowly defined product markets. In a setting where market frictions slow the rate of arbitrage across suppliers, new entrants may have low market shares even when offering the same price and quality as an incumbent supplier. In this setting, the share-based approach systematically understates the relative quality of entering suppliers. We empirically confirm this point in the semiconductor market.

The paper proceeds as follows. Section II introduces the measurement problem and proposes a price index using long-run price differences to infer quality differences. Section III provides background on the contract semiconductor manufacturing industry, which differs in important ways from the more commonly studied memory and processor markets. Section IV investigates price differences and price dynamics between suppliers, finding strong evidence for price differences across suppliers of identical quality goods. In section V, we calculate a standard matched model index and show that it is biased upward in comparison to our proposed index. Section VI concludes.

II. Measurement Framework

The U.S. Bureau of Labor Statistics (BLS) and most other government statistical agencies use a matched-model approach to price index construction (Bureau of Labor Statistics, 1997; Triplett, 2006). This approach defines a set of products and collects their prices in two time periods, t and $t-1$.[5] For each product j, the analyst measures the price relative, $p_{j,t}/p_{j,t-1}$. These price relatives are then averaged across products using an index number formula such as a Laspeyres index, which weights products' price relatives using initial-period expenditure shares. The use of price relatives is intended to control for differences in unobserved quality across products. Assume that an observed unit of product j reflects ξ_j quality-adjusted units, so that the price of a quality-adjusted unit is $P_{j,t} = p_{j,t}/\xi_j$. As long as products are sufficiently narrowly defined such that ξ_j is constant across periods, then $p_{j,t}/p_{j,t-1} = P_{j,t}/P_{j,t-1}$. The observed price relative equals the quality-adjusted price relative. The average of observed price relatives is then the average growth in quality-adjusted prices.

However, matched-model indexes do not easily accommodate new products, since one cannot calculate a price relative in the period of entry. While there are a variety of ways to address the introduction of new products, the most common approach is to simply omit them from the index in the period of entry (Triplett, 2006). This choice is justified when the law of one price holds for quality-adjusted units, meaning that $P_{j,t} = P_{k,t}\ \forall j,k$. In this case, including or omitting a newly entering product from the index has no effect on the average quality-adjusted price; any set of continuing products will provide an accurate measure of average price growth.

The index construction problem becomes more difficult when the law of one price does not hold. For example, imagine that newly entering products typically have lower quality-adjusted prices than continuing products, so that buyers shift toward the entrants. In this case, average price growth for continuing products will overstate overall price growth. More generally, when new entrants exhibit systematic quality-adjusted price differences from continuing varieties, omitting new entrants from the matched-model index leads to systematic upward bias.[6]

This potential problem with matched-model indexes is well known to the BLS and other measurement agencies. It is referred to as new product bias in the case of fundamentally new goods, and new outlet bias or outlet substitution bias when new suppliers of existing goods enter the market.[7] The BLS attempts to mitigate these biases using hedonic regressions or subjective quality estimates provided by survey respondents, but these adjustments are implemented for

[2] See, for example, Reinsdorf (1993), the Boskin Commission Report (Boskin, Dulberger, & Gordon, 1998), Greenlees and McClelland (2011), and Hausman and Leibtag (2009).

[3] Houseman et al. (2011) argue that substitution toward offshore suppliers in developing countries substantially biased measures of productivity and real value-added growth in the United States. Feenstra et al. (2013) study three issues that bias the import price index upward: using a nonsuperlative index number formula, omitting the effect of tariffs on purchase prices, and omitting the effects of increased variety. Houseman and Ryder (2010) provide a summary of other papers discussing measurement problems arising from globalization.

[4] Prominent examples include Feenstra and Romalis (2014), Hallak and Schott (2011), Hottman, Redding, and Weinstein (2016), Hummels and Klenow (2005), Khandelwal (2010), Khandelwal, Schott, and Wei (2013), and Kugler and Verhoogen (2011).

[5] A "product" can be interpreted broadly as a bundle including a physical item, its supplier, and/or its buyer. We will specify the relevant product dimensions in our application.

[6] Appendix B.1.1 formalizes this statement using the measurement framework we introduce in the paper.

[7] The landmark Boskin Commission Report (Boskin et al., 1998) addressed both new product and new outlet bias in the U.S. Consumer Price Index (CPI). BLS economists Greenlees and McClelland (2011) empirically examine outlet substitution bias in the CPI and cite BLS studies on the topic as early as 1964. Former BLS assistant commissioner of international prices William Alterman (2015) recently proposed the creation of a new Input Price Index program to address outlet substitution bias in the BLS Producer Price Index (PPI) and Import Price Index (MPI) due to offshoring.

very few products.[8] In the overwhelming majority of cases, new products and new suppliers are simply omitted from the index when they first appear.

In this paper, we focus on new outlet bias and propose an alternative strategy for incorporating entering suppliers that allows for quality-adjusted price variation. We examine a framework that exhibits the law of one price as a special case by assuming that different firms supply the same physical good, and thus quality-adjusted units are perfect substitutes. The representative buyer combines output, $y_{j,t}$, from each supplier, j, using a linear aggregator over quality-adjusted units,

$$Y_t = \sum_{j \in J_t} \xi_j y_{j,t}, \tag{1}$$

where J_t is the set of available products. Balk (1998) shows that with a perfect substitutes aggregator like equation (1), the exact price index measures the growth in quality-adjusted unit value:

$$I_t \equiv \frac{\sum_{j \in J_t} p_{j,t} y_{j,t} \Big/ \sum_{j \in J_t} \xi_j y_{j,t}}{\sum_{j \in J_{t-1}} p_{j,t-1} y_{j,t-1} \Big/ \sum_{j \in J_{t-1}} \xi_j y_{j,t-1}}. \tag{2}$$

Because goods are perfect substitutes, cost minimization implies that the law of one price holds—that $p_{j,t}/\xi_j = P_t \; \forall j$.[9] Substituting this into equation (2) yields $I_t = P_t/P_{t-1} = p_{j,t}/p_{j,t-1} \; \forall j \in (J_t \cap J_{t-1})$; average quality-adjusted price growth equals price growth for each continuing product under the law of one quality-adjusted price. The matched-model index remains accurate even when dropping new entrants.

Now consider a more general setting in which product market frictions impede buyers from freely adjusting their mix of suppliers. These impediments might include search costs (Sorensen, 2000), slow diffusion of information about new suppliers (Foster, Haltiwanger, & Syverson, 2015; Perla, 2016), relationship-specific investments between buyers and suppliers (Tirole, 1988; Klemperer, 1995), investments in systems of compatible components (Greenstein, 1995), and others.[10] We capture these frictions by assuming the buyer faces a constraint on the rate at which it can alter its mix of suppliers,

$$\sum_j \mathcal{A}\left(\sigma_{j,t}, \sigma_{j,t-1}\right) \leq k \quad \text{for } k > 0, \tag{3}$$

where $\sigma_{j,t} \equiv \xi_j y_{j,t}/Y_t$ is supplier j's share of quality-adjusted purchases, k represents the buyer's capacity to adjust its supplier mix in any one period, and $\mathcal{A}\left(\sigma_{j,t}, \sigma_{j,t-1}\right) > 0$ measures the amount of this capacity that is exhausted by adjusting the share of purchases from a single supplier from $\sigma_{j,t-1}$ to $\sigma_{j,t}$. Assume that $\mathcal{A}(\cdot)$ is a smooth, convex function of both arguments.[11] Equation (3) parsimoniously captures the ideas that frictions slow the reallocation of market shares and that they relax as a product market matures. Early in a product's life cycle, buyers may be unaware of new suppliers or may be locked in by specific investments with the founding suppliers in the market. These frictions initially impede the reallocation of market share, since equation (3) binds. But over time, information diffuses, specific investments depreciate, and new buyers enter the market without existing supplier commitments. When the market reaches a point where large reallocations of market shares are no longer sought, equation (3) no longer binds.[12]

Frictions make the problem dynamic, so the buyer solves

$$\min_{\{y_{j,t}\}} \sum_t \delta^t \sum_j p_{j,t} y_{j,t}, \qquad \delta \in (0,1), \tag{4}$$

subject to equations (1) and (3). Using the definition of $\sigma_{j,t}$, the first-order conditions are

$$\frac{p_{j,t}}{\xi_j} + \lambda_t \mathcal{A}_1\left(\sigma_{j,t}, \sigma_{j,t-1}\right) + \delta\lambda_{t+1}\mathcal{A}_2\left(\sigma_{j,t+1}, \sigma_{j,t}\right) = \mu_t \; \forall j, \tag{5}$$

where μ_t is the multiplier on equation (1) and λ_t is the multiplier on equation (3). If the set of suppliers has stabilized and desired market share adjustments have largely been realized, then equation (3) will not bind, and equation (4) collapses to the static problem above. In this long-run setting, which we refer to as $t = T$, equation (5) collapses to $p_{j,T}/\xi_j = \mu_T$, that is, quality-adjusted prices equalize. Now consider the prior period, $T-1$, in which the adjustment constraint still binds. Since the constraint does not bind in T, $\lambda_T = 0$, and equation (5) becomes

$$\frac{p_{j,T-1}}{\xi_{j,T-1}} + \lambda_{T-1}\mathcal{A}_1\left(\sigma_{j,T-1}, \sigma_{j,T-2}\right) = \mu_{T-1}. \tag{6}$$

In the second term, $\mathcal{A}_1\left(\sigma_{j,T-1}, \sigma_{j,T-2}\right)$ reflects the marginal cost of increased market share in terms of its depletion of the buyer's limited adjustment capacity. This term is positive when $\sigma_{j,T-1} > \sigma_{j,T-2}$, which implies that a supplier must offer a lower quality-adjusted price to increase its market share. This intuition applies more generally: new entrants

[8] For the CPI, the BLS uses hedonic quality adjustments for certain products in apparel and electronics and for textbooks and rental units. The PPI and MPI use hedonic adjustment for computers. Subjective adjustments are common for computers and vehicles but rare otherwise; less than 2% of observations in the MPI are quality-adjusted in this manner. See online appendix A.1 for details, including a discussion of the particular sources of new outlet bias in import price measurement.

[9] The law of one price may still hold when buyers have an explicit taste for diversity. Online appendix B.2 provides an example.

[10] Such frictions appear to apply quite broadly. Monarch (2016) finds that 45% of U.S. arm's-length importers maintain their primary Chinese supplier across years, inferring the presence of substantial costs of switching international suppliers.

[11] A few other natural restrictions on $\mathcal{A}$ are (a) $\mathcal{A}(\sigma,\sigma) = 0$, (b) $\mathcal{A}_1(\sigma,\sigma) = \mathcal{A}_2(\sigma,\sigma) = 0$, (c) $\mathcal{A}_1(\sigma_t,\sigma_{t-1}) > (<) \, 0$ as $\sigma_t > (<) \, \sigma_{t-1}$; and (d) $\mathcal{A}_2(\sigma_t,\sigma_{t-1}) < (>) \, 0$ as $\sigma_t > (<) \, \sigma_{t-1}$.

[12] In an earlier draft, we explicitly modeled a lump-sum cost of switching suppliers (Byrne, Kovak, & Michaels, 2013). Though this cost endows sellers with some market power that elevates prices across the board, we find that the relative prices of suppliers largely reflect quality differences late in the product life cycle, as the inflow of new buyers helps to equalize market power for incumbent sellers and new suppliers.

offer quality-adjusted discounts to accumulate market share, but their relative prices rise over time until the law of one quality-adjusted price is restored.[13] This is the pattern we will observe in the semiconductor market in section IV.

Our proposed measurement strategy takes advantage of the fact that in the long run (period T), the effects of frictions dissipate, such that differences in observed prices fully reflect quality differences. This makes it straightforward to infer quality differences, since

$$p_{j,T}/p_{k,T} = \xi_j/\xi_k. \tag{7}$$

We can therefore use the ratio of observed prices in the long run, $p_{j,T}/p_{k,T}$, to correct for unobserved quality differences in earlier periods. Substituting equation (7) into equation (2) yields

$$\hat{I}_t \equiv \frac{\sum_j p_{j,t}y_{j,t} \Big/ \sum_j \frac{p_{j,T}}{p_{1,T}} y_{j,t}}{\sum_j p_{j,t-1}y_{j,t-1} \Big/ \sum_j \frac{p_{j,T}}{p_{1,T}} y_{j,t-1}}, \tag{8}$$

where the baseline supplier ($k = 1$) can be any firm that participates in both periods $t-1$ and t. Unlike equation (2), this index can be calculated using observable data. However, it cannot be implemented in real time, since it uses long-run price differences to infer earlier quality differences. We propose that price measurement agencies continue to produce existing indexes, but later issue revised series that update earlier estimates once the the set of suppliers and their market shares has stabilized, and the long-run condition, equation (7), holds. Section V discusses these implementation issues in more detail.

Our proposed index accurately measures the exact index in equation (2) when relative quality is constant over time, because long-run price relatives reflect earlier quality differences.[14] In contrast, even with constant relative quality, matched-model indexes are biased upward because they omit price declines occurring when entrants offering low quality-adjusted prices gain market share.[15] We will argue that constant relative quality is likely in the market we study, and if products are sufficiently narrowly defined (e.g., with barcode data), it will hold in many other markets as well. However, in many contexts, the constant quality assumption will be less plausible.[16] In appendix B.1, we show that in situations where an entrant's quality catches up with that of incumbents, under reasonable assumptions one can bound the exact index between the standard matched-model index (an upper bound) and our proposed index (a lower bound). Thus, even with changing quality, our proposed index may be used to bound the true price path.

III. Contract Semiconductor Manufacturing

We study these price measurement issues in the context of the contract semiconductor manufacturing industry.[17] Semiconductor fabrication involves creating networks of transistors on the surface of a thin piece of semiconducting material called a "wafer." Transistors and the connections between them are created by successively depositing and etching away layers of conducting and insulating materials on the wafer's surface, with more complex designs requiring more layers. The etching pattern for each layer is projected onto the wafer through a "mask" containing the desired pattern.

Semiconductor fabrication technology has advanced over time in discrete steps, defined by wafer size and line width (also called feature size). Larger wafers accommodate a larger numbers of chips, reducing the cost per chip. During our sample period, producers made 150 mm, 200 mm, or 300 mm diameter wafers. Line width is the size of the smallest feature that can be fabricated reliably. Decreased line width means that individual transistors are smaller, making chips of a given functionality smaller, lighter, faster, and more energy efficient. Smaller line widths are more difficult to produce, initially resulting in lower yields, the fraction of chips on a wafer that function correctly, when a new line width is introduced. In our sample, line widths range from 45 nm to more than 500 nm. We refer to each combination of wafer size and line width as a process technology (e.g., 200 mm wafer, 180 nm line width).

A primary benefit of studying the contract semiconductor manufacturing industry is that technological change proceeds discretely, with wafer sizes and line widths taking on a small set of particular values. The number of layers also determines the manufacturing complexity of a particular design. These discrete observable features allow us to flexibly control for product characteristics when examining prices.

We focus on transactions between semiconductor design firms and contract semiconductor manufacturers called foundries. These firms, operating principally in Asia, exclusively produce chips designed by other firms. This vertically disintegrated business model has grown substantially over time as the costs of fabrication facilities have become prohibitive for all but the largest semiconductor companies. Importantly, using arm's-length transactions avoids price measurement difficulties potentially arising from transfer pricing between related parties.

The foundry market differs in important ways from the microprocessor and memory markets that have been the

[13] See appendix B.3 for a full characterization of equilibrium dynamics when the reallocation of market shares is impeded by frictions.

[14] Our index requires no assumptions about costs or markups, which do not appear in equation (2) or equation (8).

[15] See appendix B.1.1.

[16] Bilir (2014) provides an example in which intellectual property rights concerns wane as a product ages, so the difference in quality between suppliers with strong versus weak intellectual property protections will decline over time.

[17] See appendix C for additional detail on contract semiconductor manufacturing. Turley (2003) provides an accessible overview of semiconductor technology, manufacturing, and business.

focus of most prior economic studies of semiconductors.[18] Foundries instead specialize in custom chips for specific models of electronic devices such as cellular phones, hard drives, automobiles, and digital cameras, with chips produced in smaller batches and requiring substantial investment in design. They also use mature fabrication technologies that are one generation or more behind the leading-edge processes used by processor and memory manufacturers. Thus, the arrival of new process technologies in the foundry market is largely determined by external forces, and there is less scope for yield improvement in these mature technologies.[19]

Contract semiconductor fabrication involves large costs of switching suppliers, resulting from the custom aspect of each product and substantial investments required to have a particular foundry produce a particular design. For example, a set of masks costs more than $1 million and cannot be transferred across suppliers. This led one industry association to state, "The time and cost associated with [switching] tend to lock customers into a particular foundry."[20] Other examples of start-up costs include the many chemical and mechanical adjustments and calibrations to manufacturing equipment that are implemented during the engineering phase of production for a particular chip design. These adjustments must be redone when moving to a new production line, even within the same firm. An additional example is negotiating a new supplier agreement, which requires a vast amount of technical information to be exchanged and can absorb much attention of senior management (Allan, 2002). Together these supplier-specific investments for each design create frictions that impede buyers from switching an existing design to a new supplier. As we showed in section II, frictions of this nature tend to generate quality-adjusted price dispersion that biases traditional matched-model price indexes.

IV. Price Dispersion across Countries

Information on semiconductor wafer prices comes from a proprietary database of purchases from foundries, collected by the Global Semiconductor Alliance (GSA), a nonprofit industry organization. The data set consists of 6,916 individual responses to the Wafer Fabrication & Back-End Pricing Survey for 2004–2010, providing a representative sample of transactions that accounts for about 20% of the semiconductor wafers produced by the foundry sector worldwide. The data report the price paid, location (country) of the foundry, and the line width, wafer size, and layers for each transaction. This information allows us to examine average prices by foundry location, controlling for all relevant physical characteristics.

Although this level of product detail is remarkable, the GSA data have two important limitations. First, we only observe the country in which the foundry is located and cannot identify the producing firm. However, a single firm accounts for the vast majority of output in each of the three largest sources of contract semiconductor manufacturing services: Taiwan Semiconductor Manufacturing Corportation (TSMC) in Taiwan, Semiconductor Manufacturing International Corporation (SMIC) in China, and Chartered Semiconductor in Singapore.[21] As shown in table 1, these countries account for about 80% of worldwide foundry output. Thus, for the three largest supplying countries, country identifiers closely approximate producing-firm identifiers. The second main limitation is that we have no information on buyers, so we are unable to control for any buyer-specific information.[22]

Descriptive statistics for key variables in the GSA database are shown in table 2.[23] The rapid changes in semiconductor technology are immediately apparent: 300 mm wafers' market share rises from 3% to 20% over the sample period, while older technologies' shares fall. Newer line widths also generally increase in share over time: 65 nm technology slowly gained share in the foundry market, accounting for 8% by 2010, while older technologies, with line widths larger than 250 nm, dwindle from 40% in 2004 to 33% in 2010. The number of metal and mask layers per wafer also rose somewhat over time, reflecting a trend toward more complex designs. These steady improvements in wafer technology demonstrate the importance of controlling for product characteristics when examining price differences across suppliers and price changes over time.

In table 3, we use a hedonic regression framework to investigate cross-country variation in wafer prices. First, we regress the log wafer price on quarter indicators and indicators for the location of production, with Taiwan as the omitted category. In all specifications, we cluster standard errors by quarter.[24] Column 1 finds very large unconditional price differences across suppliers, with wafers produced in China costing 25.7% less than those produced in Taiwan on average.[25] These large unconditional price differences can partly be explained by differences in product attributes. Column 2 adds indicators for each wafer size and each

[18] Microprocessor and memory products account for a minimal share of foundry output, according to the IHS iSuppli Pure Play Foundry Market Tracker. Unfortunately, our data agreement prohibits us from citing a precise figure for this share.

[19] See appendix F for detailed analysis of yields.

[20] This quotation is from the Common Platform Alliance, an industry group advocating a common platform that would standardize aspects of semiconductor production technology.

[21] These statements are based on IHS iSuppli data, but our confidentiality agreement prohibits us from quoting precise figures at the firm level. Note also that although TSMC operated one small production facility in China during our sample period, it accounted for a very small share of Chinese output.

[22] For more information on the data sources, sample, and cleaning, see appendix D.

[23] Observations are weighted by combining transaction-level quantities from GSA and quarter × country × technology quantity information from IHS iSuppli. See appendix D.2 for details.

[24] Results without clustering the standard errors are presented in appendix E.1. For the vast majority of coefficients in table 3, the clustered standard errors are larger than those assuming independent errors.

[25] $\exp(-0.297) - 1 = -0.257$.

TABLE 1.—FOUNDRY CAPACITY BY COUNTRY

	Global Capacity									
	Foundry	Total Industry	Taiwan	China	Singapore	Europe	United States	Japan	South Korea	Malaysia
2000	875	9,462	63.2%	7.1%	7.5%	6.8%	4.4%	7.8%	3.2%	0.0%
2001	972	8,286	60.5	8.9	8.4	6.3	4.2	7.1	2.9	1.7
2002	1,011	8,646	56.0	10.7	10.0	6.2	5.3	7.0	2.9	1.8
2003	1,150	9,018	51.1	15.5	10.1	6.5	5.1	6.4	3.2	2.2
2004	1,429	10,000	50.3	19.0	9.6	5.6	4.1	5.3	3.1	3.0
2005	1,739	11,073	48.5	23.2	9.2	4.9	3.6	4.5	3.0	3.1
2006	1,951	12,320	48.0	24.1	9.9	4.4	3.3	4.1	3.2	3.0
2007	2,157	13,588	48.9	23.4	10.0	4.6	3.1	3.7	3.5	3.0
2008	2,401	14,297	49.9	22.1	11.2	5.1	2.9	2.9	3.3	2.7
2009	2,546	14,058	49.2	22.0	10.9	6.5	2.8	2.7	3.5	2.5
2010	2,812	14,230	49.4	21.5	11.3	7.4	2.6	2.5	3.3	2.1
2011	3,177	14,923	50.2	21.8	10.7	7.9	2.4	2.3	3.0	1.8

Authors' calculations from IHS iSuppli data. See appendix D.2 for more detail on this data source. Sample includes contract manufacturers (pure-play foundries) only. Capacity measured in thousand 8-inch equivalent wafers per month.

TABLE 2.—WAFER PRICE DESCRIPTIVE STATISTICS

			Yearly Means						
	Mean	SD	2004	2005	2006	2007	2008	2009	2010
Price per wafer ($)	1,204.22	949.28	1,321.01	1,293.51	1,298.54	1,189.68	1,120.34	1,080.39	1,126.09
Number of wafers contracted	11,865.00	22,302.74	8,344.61	14,912.17	10,544.97	16,767.34	9,260.24	9,692.42	13,533.24
Layers									
Metal layers	4.60	1.77	4.22	4.56	4.85	4.76	4.62	4.56	4.64
Mask layers	25.72	7.31	23.69	24.52	26.37	26.74	25.86	26.27	26.58
Polysilicon layers	1.23	0.46	1.36	1.21	1.20	1.19	1.30	1.14	1.20
Wafer size									
150 mm	0.16	0.37	0.18	0.17	0.15	0.10	0.17	0.16	0.17
200 mm	0.73	0.44	0.79	0.76	0.76	0.81	0.68	0.68	0.63
300 mm	0.11	0.32	0.03	0.07	0.09	0.09	0.15	0.16	0.20
Line width									
45 nm	0.00	0.06	0.00	0.00	0.00	0.00	0.00	0.00	0.03
65 nm	0.03	0.16	0.00	0.00	0.00	0.01	0.04	0.07	0.08
90 nm	0.04	0.21	0.00	0.02	0.04	0.05	0.08	0.06	0.06
130 nm	0.15	0.36	0.12	0.16	0.18	0.17	0.14	0.14	0.15
150 nm	0.06	0.23	0.07	0.08	0.10	0.09	0.03	0.01	0.01
180 nm	0.26	0.44	0.25	0.23	0.24	0.28	0.26	0.31	0.26
250 nm	0.12	0.33	0.17	0.17	0.13	0.15	0.10	0.07	0.08
Older vintage	0.33	0.47	0.40	0.35	0.30	0.25	0.34	0.35	0.33

Authors' calculations based on GSA Wafer Fabrication and Back-End Pricing Survey. Summary statistics based on sample of 6,253 transaction-level observations. Transactions weighted by combining technology-level (wafer × line width) shipments information from IHS iSuppli and within-technology order size information from the GSA transaction data.

line width, layer controls, and the size of the transaction.[26] The omitted category is for 200 mm wafers with 180 nm line width, produced in Taiwan. The signs of all regression coefficients are intuitive. More advanced production technologies, with larger wafers, smaller line widths, or more layers, command higher prices. When controlling for this exhaustive list of technological attributes, the estimated price differences across suppliers change. For example, the measured China discount falls in magnitude from 25.7% to 17.0%, indicating that China produced more trailing-edge products than Taiwan. Column 3 uses more flexible technology controls, including indicators for each wafer size × line width combination. The substantial price differences persist, indicating that Chinese suppliers provide substantial average discounts, even for technologically identical products.

Even if physical product attributes are identical across suppliers, there may still be unobserved differences in the quality of the overall production service that explain the price differences in table 3. For example, a supplier with systematically lower yields would likely provide a discount to compensate the customer for the fact that each wafer represents fewer usable chips as compared to a supplier with higher yields.[27] Foundries may differ in the quality of engineering support they provide customers to implement their designs in the foundry's manufacturing process. Other aspects of customer service, reliability, or brand reputation may also differ across foundries in different countries.

While hedonic methods are unable to account for these potential sources of unobserved heterogeneity across suppliers, the measurement strategy proposed in section II shows how to use long-run prices to infer unobserved quality

[26] Note that hedonic regression estimates reflect a joint envelope of demand and supply functions (Rosen, 1974). Including an order size regressor to control for nonlinear pricing is reasonable in this context, unlike when trying to separately identify a demand or supply curve. That said, all results are similar when omitting this control.

[27] Appendix F provides a detailed analysis of yields, finding that they are very unlikely to account for the price differences across countries in table 3.

TABLE 3.—HEDONIC WAFER PRICE REGRESSIONS
Dependent Variable: Log of Price per Wafer

	(1) No Attribute Controls		(2) Linear Attribute Controls		(3) Flexible Attribute Controls		(4) China and Taiwan Only	
Variable	Coefficient	SE	Coefficient	SE	Coefficient	SE	Coefficient	SE
Foundry location								
China	−0.297	(0.042)***	−0.186	(0.027)***	−0.188	(0.027)***	−0.196	(0.028)***
Malaysia	−0.274	(0.065)***	−0.278	(0.042)***	−0.286	(0.041)***		
Singapore	−0.046	(0.026)*	−0.061	(0.016)***	−0.068	(0.016)***		
United States	−0.093	(0.021)***	0.068	(0.030)**	0.064	(0.031)**		
Wafer size								
150 mm			−0.467	(0.032)***				
300 mm			0.671	(0.021)***				
Line width								
≥ 500 nm			−0.245	(0.053)***				
350 nm			−0.167	(0.033)***				
250 nm			−0.061	(0.026)**				
150 nm			0.169	(0.027)***				
130 nm			0.356	(0.018)***				
90 nm			0.479	(0.032)***				
65 nm			0.676	(0.030)***				
45 nm			0.962	(0.062)***				
Wafer Size × Line Width Indicators					X		X	
Number of metal layers			0.076	(0.007)***	0.076	(0.007)***	0.081	(0.007)***
Number of polysilicon layers			0.027	(0.024)	0.028	(0.024)	0.025	(0.024)
Number of mask layers			0.005	(0.002)***	0.005	(0.002)***	0.004	(0.002)**
Epitaxial layer indicator			0.064	(0.037)*	0.067	(0.037)*	0.066	(0.038)*
log number of wafers Contracted			−0.056	(0.004)***	−0.057	(0.004)***	−0.058	(0.005)***
Quarter indicators	X		X		X		X	
R^2	0.046		0.909		0.913		0.922	
Observations	6,253		6,253		6,253		5,378	

Observations represent individual semiconductor wafer transactions from GSA data. The omitted category is a 200 mm 180 nm wafer produced in Taiwan. Transactions weighted by combining technology-level (wafer × line width) shipments information from IHS iSuppli and within-technology order size information from the GSA transaction data. Standard errors clustered by 28 quarter clusters. Significant at ***1%, **5%, *10%. See appendix E.1 for alternative weighting assumptions and results without clustering standard errors.

differences. To visualize these long-run price differences, figure 1 shows the evolution of prices for four important technologies produced by both Chinese and Taiwanese foundries. Each point represents the average price per wafer for transactions in the same quarter × country × technology cell. In general, Taiwan begins production at least eight quarters before Chinese foundries.[28] When Chinese firms do enter, their prices are below their Taiwanese competitors, but the gap generally narrows over time. This pattern is consistent with newly entering Chinese producers initially offering quality-adjusted discounts to attract customers in the presence of frictions, followed by declining price gaps that eventually reflect the underlying average quality difference as frictions dissipate. The increase in Chinese market share shown in table 1 corroborates the interpretation that Chinese producers initially offer quality-adjusted discounts.

Figure 2 examines these technology-specific price gaps directly. The y-axis shows the China/Taiwan price ratio, allowing for comparisons across technologies with different average prices. The x-axis measures the number of quarters since the last supplier entered the relevant technology.[29] The dashed gray line plots the price ratio for 200 mm 180 nm wafers, the technology with the most sales in our sample period. Because our data begin in Q1 2004 and China entered this market in Q3 2002, we first observe the price gap six quarters following Chinese entry. Although the price ratio is volatile over time, it exhibits a clear upward trend, reflecting the closing price gap visible in the upper right panel of figure 1. The black line in figure 2 averages the technology-specific price ratios over technologies within each entry quarter, showing that the pattern of closing price gaps applies on average across technologies. To ensure that this pattern is not driven by changes in the technology mix or changes in design complexity over time, table 4 implements a within-technology analysis, regressing quarterly price ratios on technology fixed effects, layer controls, and the time since Chinese entry. All specifications find closing price gaps, and quadratic specifications in columns 3 and 4 suggest that the price gap closes faster in the years just after entry and stabilizes later. The dashed black line in figure 2 plots the predicted values from column 4 of table 4, confirming that the closing price gaps are a within-technology phenomenon and that the average China/Taiwan price ratio levels out at approximately 0.9. Appendix E.2 documents a

[28] Authors' calculations using IHS iSuppli data. See appendix D.4 for details on measurement and the reasons for delayed Chinese entry.

[29] We restrict attention to Taiwan, China, and Singapore, omitting very small suppliers, and record the quarter in which the last of these three suppliers entered the relevant market. China was the last entrant for all but one technology, for which Singapore entered last. See appendix D.4 for details.

FIGURE 1.—AVERAGE WAFER PRICES BY QUARTER AND TECHNOLOGY, TAIWAN AND CHINA

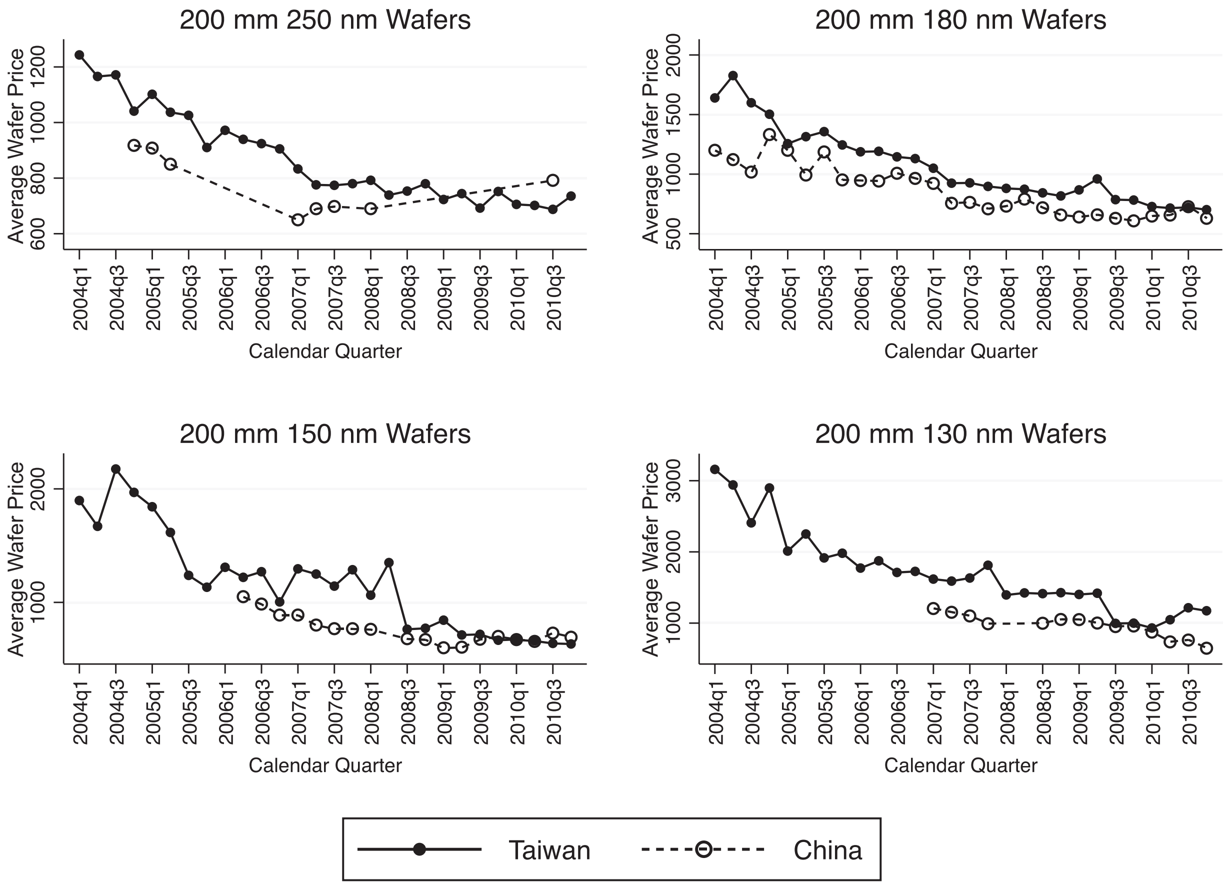

Each point represents the average price per wafer for transactions in the relevant quarter, produced by the relevant country, and with the relevant technology (wafer size × line width). When multiple transactions appear in a given quarter, country, technology cell, prices are averaged using quantity weights to yield the average price per wafer across the transactions. The figure shows price trajectories for four large technologies with substantial output from both Taiwan and China.

similar pattern of closing price gaps between Singapore and Taiwan.[30]

The robust pattern of large initial price gaps that converge over time suggests the presence of frictions that initially drive large quality-adjusted price differences across suppliers and then dissipate over time. Fixed quality differences or quality differences that evolve over calendar time for all technologies cannot explain this within-technology pattern.[31] Only if quality differences evolve systematically within each technology can they explain the observed price dynamics. As an example, if Chinese yields start well below those in Taiwan for each new technology and then catch up over time, this could explain the convergence in relative prices. However, appendix F provides extensive evidence against this possibility. We use detailed foundry yield data for 65 nm technology, GSA yield data, industry publications, and interviews to show that it is very unlikely that yield changes account for the observed gradual convergence in prices. Moreover, even if relative quality does change over time, our measurement strategy can still be used to bound the true index, as noted in section II. We return to this possibility in the next section.

V. Price Index Results

In this section, we implement the measurement strategy proposed in section II and compare the results to a standard matched-model approach. The first step is to choose a reference supplier ($k = 1$ in equation [8]) that is present in both periods. We choose Taiwan, since it is the first entrant for all technologies and remains present throughout our sample period.[32] Next, we must define the long run, which is reached when the distribution of market shares has stabilized and frictions no longer distort relative prices. We suggest that analysts use a combination of institutional knowledge and observed price and market share dynamics to determine the relevant long run for a given market. For semiconductors, large fixed costs of starting production of a particular design with a given supplier deter the movement of a design

[30] Appendix E.3 uses an alternative approach to document price convergence across all suppliers.

[31] This rules out explanations based on steady improvements in the quality or reliability of China's production service or changes in brand recognition, customer service, or tax policy.

[32] Note that the reference supplier can change when considering different pairs of time periods, so it is not necessary that a particular supplier be present throughout the life of an index.

FIGURE 2.—CLOSING CHINA/TAIWAN PRICE RATIO

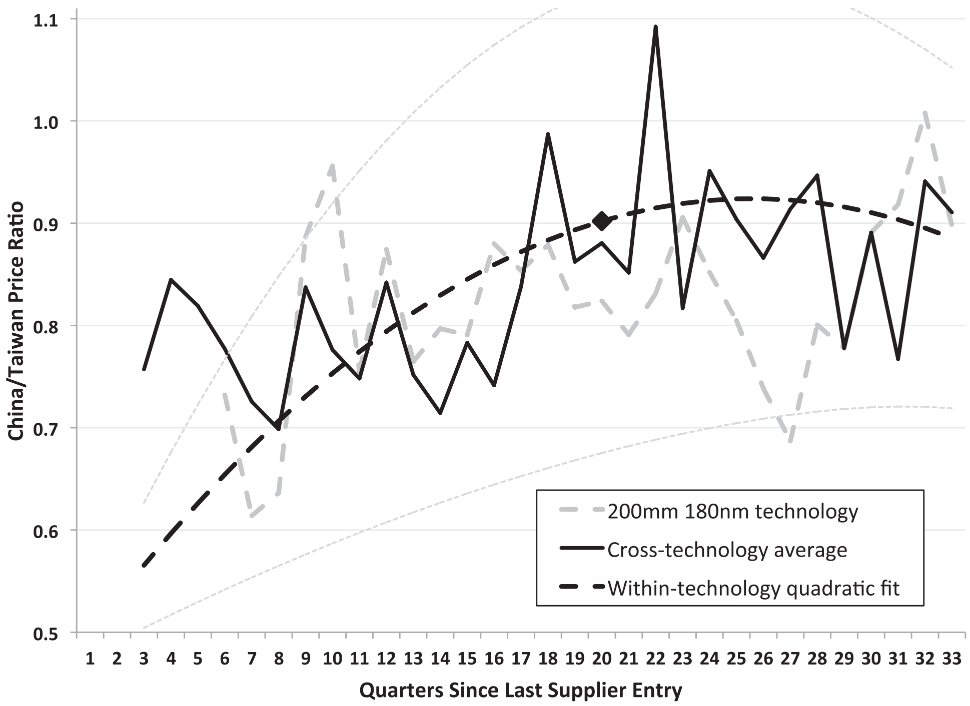

The x-axis measures the number of quarters since the last supplier began producing the relevant technology. The y-axis measures the China/Taiwan price ratio. The gray dashed line shows the raw quarterly price ratio for 200 mm 180 nm technology, which had the largest market share during our sample period. We construct similar series for all technologies and then average the price ratios across technologies in each quarter to generate the solid black line, which also exhibits a closing price gap over time. The dashed black line shows predicted values from a within-technology quadratic trend estimated using technology fixed effects. The associated regression estimates are reported in column 4 of table 4. The light dotted lines show a 90% confidence interval around the quadratic prediction. The black diamond shows the point on the quadratic trend that we use to measure the long-run price ratio between China and Taiwan for the purposes of quality adjustment.

TABLE 4.—CLOSING PRICE GAPS WITHIN TECHNOLOGY

Dependent Variable: China-Taiwan Price Ratio

	(1)	(2)	(3)	(4)
Time since last supplier entry	0.007**	0.009***	0.027*	0.036***
	(0.003)	(0.002)	(0.015)	(0.013)
(Time since last supplier entry)^2			−0.0005	−0.0007**
			(0.0003)	(0.0003)
Difference in average:				
Number of metal layers		0.310***		0.323***
		(0.090)		(0.087)
Number of polysilicon layers		0.079		0.081
		(0.075)		(0.072)
Number of mask layers		0.148		0.208**
		(0.102)		(0.101)
Epitaxial layer		−0.011**		−0.010***
		(0.004)		(0.004)
Log number of wafers contracted		−0.006		−0.007
		(0.009)		(0.007)
Technology (wafer size × line width) indicators	X	X	X	X
R^2	0.166	0.433	0.190	0.472
Observations	91	91	91	91

Observations represent technology (wafer size × line width) and quarter combinations. The dependent variable is the China/Taiwan price ratio. "Time since last supplier entry" reflects the elapsed number of quarters since the last supplier first began producing the relevant technology. All specifications include technology fixed effects, so trends reflect within-technology changes in the price ratio as time elapsed following the last supplier's entry. All specifications exhibit closing price gaps within technology over time. Heteroskedasticity-robust standard errors in parentheses. Significant at ***1%, **5%, *10%.

from one supplier to another. However, designs have finite life spans and new designs continually enter the market. After all of the initial designs phase out, the newer remaining designs face a stable set of suppliers during their entire period of production, and there is little subsequent reallocation of market shares. The vast majority of designs produced at foundries last for five years or less, so nearly all of the locked-in designs leave within five years following the last supplier's entry.[33] We therefore assume that the long run is reached after five years, which is corroborated by figure 2,

[33] This conclusion is based on helpful discussions with Len Jelinek, chief analyst at IHS, and Falan Yinug of the Semiconductor Industry Association, who solicited related information from the association's members.

FIGURE 3.—PRICE INDEXES

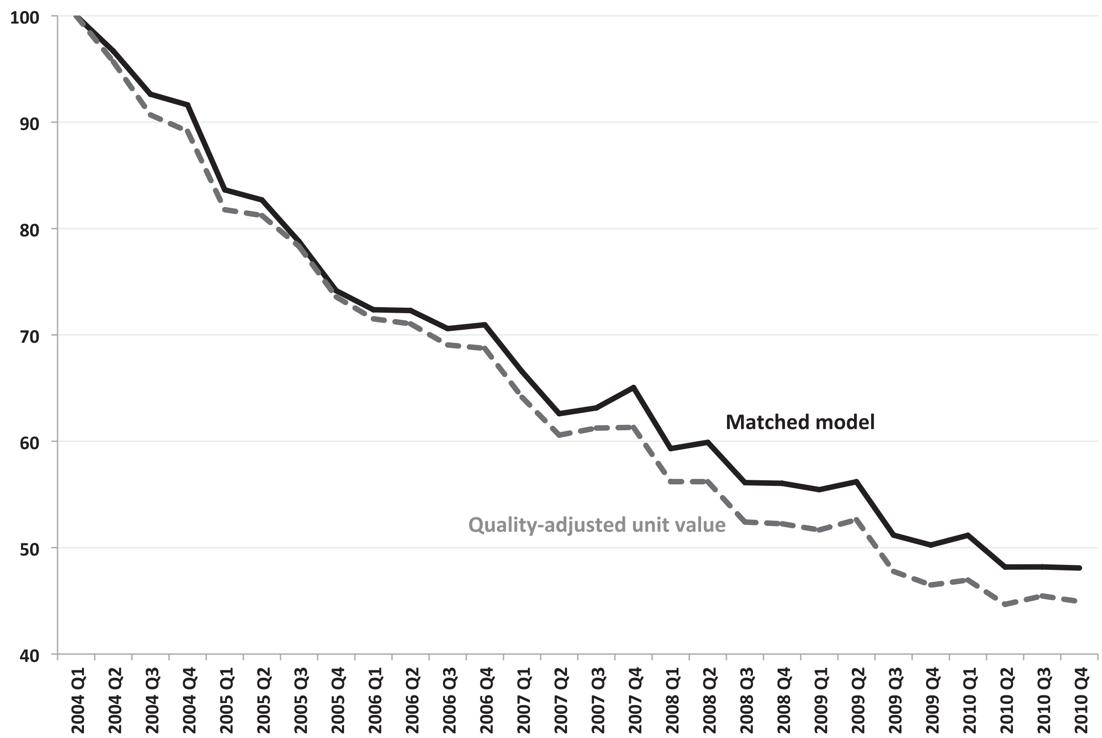

Authors' calculations using GSA wafer price data and IHS iSuppli quantity data. Sample restricted to wafers produced in Taiwan, China, and Singapore. The solid black line matched-model index is a Fisher index of price relatives for each country × technology, following the standard matched-model index construction method. The dashed gray line, quality-adjusted unit value index, first generates quality-adjusted unit values across countries within technology, following equation (8), and then forms technology-level price relatives from these unit values. It then aggregates these price relatives across technologies using a Fisher index. The quality-adjusted unit value falls more quickly than the standard matched-model index, reflecting the fact that the standard approach omits price declines when suppliers with lower quality-adjusted prices enter the market. See the text for details and table 5 for specific index values.

in which relative prices stabilize around five years after the last supplier's entry.

We can now use long-run relative prices to calculate the quality adjustment factors $p_{j,T}/p_{1,T}$ in equation (8). The preferred approach given sufficient data would be to calculate these adjustment factors separately for each technology. However, in our data, this approach would result in many missing values for smaller technologies that appear sporadically in the GSA transaction sample (e.g., 200 mm 250 nm wafers for China in figure 1). Instead, we pool across technologies to estimate average within-technology price ratios. In particular, we evaluate the quadratic fit in figure 2 at five years (twenty quarters) following the last supplier's entry, yielding a quality adjustment of 0.915 (marked with a black diamond on the figure). This implies that Chinese products provide 8.5% fewer effective units than otherwise equivalent Taiwanese products. A similar process for Singapore yields an adjustment of 0.956.

In appendix G, we compare these quality estimates to those based on a commonly used method in the international trade literature, which relies on the structure of CES demand to infer relative quality based on relative sales conditional on price.[34] We implement a version of this approach following Khandelwal et al. (2013) and find far lower relative quality estimates for China and Singapore than those just reported. This difference reflects the fact that new entrants slowly accumulate market share in a setting with frictions. Quality inference procedures assuming a frictionless market observe the low quantities for entrants and systematically understate their relative quality when frictions are present.

Given our relative quality estimates, it is straightforward to calculate the proposed quality-adjusted unit-value index. Let t index quarters, j supplying country, and h technology. We plug the average price in each quarter × country × technology cell ($p_{j,h,t}$), the quantities for each cell ($y_{j,h,t}$), and the quality adjustment factors just listed ($p_{j,T}/p_{1,T}$) into equation (8), yielding unit value relatives for each quarter × technology, $P_{h,t}/P_{h,t-1}$. We then average these relatives across technologies using a Fisher ideal index, which flexibly captures potential substitution across technologies.[35] We compare our proposed index to a standard matched-model index, which takes price relatives for each quarter × country × technology, $p_{j,h,t}/p_{j,h,t-1}$, and averages them across across country × technology pairs using a Fisher index. Note that this price relative is undefined in the quarter of supplier j's entry into the market for technology h, so the new supplier is omitted from the standard matched-model index in its quarter of entry.

The index results appear in figure 3 and table 5. Both indexes fall over time, consistent with rapid improvement

[34] See note 4 for citations.

[35] See appendix A.3 for the Fisher index formula.

TABLE 5.—PRICE INDEXES

Quarterly			Yearly		
Index:	Matched Model	Quality-Adjusted Unit Value		Matched Model	Quality-Adjusted Unit Value
2004 Q1	100.0	100.0	2004	95.2	93.9
2004 Q2	96.7	95.7	2005	79.8	78.7
2004 Q3	92.6	90.7	2006	71.6	70.1
2004 Q4	91.6	89.1	2007	64.3	61.8
2005 Q1	83.6	81.8	2008	57.8	54.3
2005 Q2	82.7	81.2	2009	53.3	49.6
2005 Q3	78.8	78.3	2010	48.9	45.5
2005 Q4	74.2	73.6			
2006 Q1	72.4	71.5	Average yearly		
2006 Q2	72.3	71.0	change 2004–2010	−0.105	−0.114
2006 Q3	70.6	69.1			
2006 Q4	71.0	68.7			
2007 Q1	66.6	64.1			
2007 Q2	62.6	60.6			
2007 Q3	63.1	61.2			
2007 Q4	65.0	61.3			
2008 Q1	59.3	56.2			
2008 Q2	59.9	56.2			
2008 Q3	56.1	52.4			
2008 Q4	56.1	52.2			
2009 Q1	55.5	51.7			
2009 Q2	56.2	52.6			
2009 Q3	51.2	47.8			
2009 Q4	50.3	46.5			
2010 Q1	51.2	47.0			
2010 Q2	48.2	44.7			
2010 Q3	48.2	45.5			
2010 Q4	48.1	44.9			

Authors' calculations using GSA wafer price data and IHS iSuppli quantity data. Sample restricted to wafers produced in Taiwan, China, and Singapore. The matched-model index is a Fisher index of price relatives for each country × technology, following the standard matched-model index construction method. The quality-adjusted unit value index first generates quality-adjusted unit values across countries within technology, following equation (8), and then forms technology-level price relatives from these unit values. It next aggregates these price relatives across technologies using a Fisher index. The quality-adjusted unit value falls more quickly than the standard matched-model index, reflecting the fact that the standard approach omits price declines when suppliers with lower quality-adjusted prices enter the market. See the text for details and figure 3 for a graphical representation.

in wafer characteristics alongside relatively stable average prices seen in table 2. As expected, the quality-adjusted unit value index falls more quickly than the standard matched-model index, with a difference of 0.9 percentage points per year.[36] The difference reflects the fact that the matched-model index omits price declines when new suppliers offering quality-adjusted discounts enter the market. To understand the practical scale of this difference, consider a hypothetical situation in which the prices of all intermediate inputs imported into the United States experienced an upward bias of 0.9 percentage points per year. This upward bias would result in understated input quantities and overstated productivity growth. We accordingly adjust the materials input price index from the BEA National Income and Product Accounts and recalculate multifactor productivity.[37] In this situation, U.S. productivity growth would be biased upward by 0.14 percentage points per year over 2004 to 2010, implying that the official productivity growth measure of 1.5% per year during the time period is biased upward by 9.5%. Thus, pervasive price index bias of the scale identified in the contract semiconductor manufacturing industry would have an important effect on productivity measurement.

If producers in China and Singapore improve their relative product quality over time for each technology, then our method overstates the entrants' relative quality when they first appear in the market. Although the scope for such within-product quality improvement is limited (see section III), it remains possible. However, under reasonable assumptions, the exact index is bounded between the matched-model and quality-adjusted unit value indexes (see section II). Thus, even if the constant-quality assumption is violated, one can plausibly bound the path for the exact index using easily calculated indexes.

A few implementation issues are worth emphasizing. The index requires only data on prices and quantities, which are generally available to measurement agencies.[38] Analysts must identify a period of supplier stability, which is a common stage of industry evolution (Klepper, 1996), but

[36] The scale of bias is similar to findings in prior work on other globalization-related measurement issues. Houseman et al. (2011) examine offshoring to developing countries and find upward bias for manufacturing import prices of 0.8 percentage points per year. Feenstra et al. (2013) document several problems biasing the import price index upward by 1.5% per year. Note that these various upward biases are likely additive across the three papers.

[37] To implement this calculation, we invoke the commonly used import comparability assumption that imports are split between intermediate and final use in the same proportion as domestic production (Feenstra & Jensen, 2012). This allows us to estimate the share of materials inputs accounted for by imported intermediates.

[38] Appendix A.2 describes the sources of quantity information already collected by BLS.

is not universal. One must also determine the length of time required before the effects of frictions dissipate. As in our analysis, this choice should be informed by a combination of institutional knowledge regarding the sources of frictions and observed behavior of prices or market shares. Finally, since the quality-adjustment procedure is retroactive, measurement agencies can continue to produce standard matched-model indexes in a timely fashion, later issuing revised or supplementary research series that use long-run prices to infer earlier quality differences. Researchers and other government agencies could then use these adjusted indexes.

VI. Conclusion

Accounting for differences in product quality across suppliers is one of the central challenges in price measurement, and this challenge is amplified as production chains continue to fragment internationally and market shares shift toward low-price suppliers in developing countries. Standard matched-model index methods assume that entrants and incumbents charge identical prices for quality-adjusted units of output, reflecting the law of one price in frictionless markets. Using a novel database of contract semiconductor manufacturing transactions, we find strong evidence for quality-adjusted price differences across suppliers. These results coincide with prior work on cross-supplier price dispersion, which finds copious evidence for frictions in various markets, calling into question the law-of-one-price assumption (Stigler & Kindahl, 1970; Abbott, 1992; Sorensen, 2000; Syverson, 2007).

We demonstrate how violations of the law of one price drive upward bias in standard matched-model indexes and propose a novel index using long-run price differences to infer quality differences across suppliers. We implement this measurement strategy in the context of the contract semiconductor manufacturing industry, finding that a standard matched model index falls almost 1 percentage point per year more slowly than the quality-adjusted unit value index. Given that many markets involve supplier-specific investments, network effects, information frictions, and search costs, we anticipate that the measurement tools we developed here will be applicable to a wide range of other industries. Because our method uses long-run price differences to infer earlier quality differences, we propose that measurement agencies continue to produce standard matched-model indexes in a timely manner and issue retrospective revisions once our quality adjustment procedure can be implemented.

REFERENCES

Abbott, T. A., "Price Dispersion in U.S. Manufacturing: Implications for the Aggregation of Products and Firms," CES discussion papers CES 92-3 (1992).

Allan, G., "A Marriage of Unequals," *Electronics Design Chain* 1 (2002), 17–20.

Alterman, W., "Producing an Input Price Index" (pp. 331–357), in S. N. Houseman and M. Mandel, eds., *Measuring Globalization: Better Trade Statistics for Better Policy* (Kalamazoo, MI: W. E. Upjohn Institute for Employment Research, 2015).

Balk, B. M., *On the Use of Unit Value Indices as Consumer Price Subindices* (Amsterdam: Division of Research and Development, Statistics Netherlands, 1998).

Bayard, K., D. Byrne, and D. Smith, "The Scope of U.S. Factoryless Manufacturing," S. N. Houseman, ed., *Measuring Globalization: Better Trade Statistics for Better Policy*, vol. 1 (Kalamazoo, MI: W. E. Upjohn Institute for Employment Research, 2015).

Bernard, A. B., and T. C. Fort, "Factoryless Goods Producers in the US," NBER Working paper 19396 (2013).

Bilir, L. K., "Patent Laws, Produce Life-Cycle Lengths, and Multinational Activity," *American Economic Review* 104 (2014), 1979–2013.

Boskin, M. J., E. R. Dulberger, and R. J. Gordon, "Toward a More Accurate Measure of the Cost of Living: Final Report to the Senate Finance Committee from the Advisory Committee to Study the Consumer Price Index," in D. Baker, ed., *Getting Prices Right: The Debate over the Consumer Price Index* (Armonk, NY: M. E. Sharpe, 1998).

Bureau of Labor Statistics, *BLS Handbook of Methods* (Washington, DC: Government Printing Office, 1997).

Byrne, D. M., B. K. Kovak, and R. Michaels, "Price and Quality Dispersion in an Offshoring Market: Evidence from Semiconductor Production Services," NBER working paper 19637 (2013).

Carlton, D. W., "Equilibrium Fluctuations When Price and Delivery Lag Clear the Market," *Bell Journal of Economics* 14 (1983), 562–572.

Feenstra, R. C., and J. B. Jensen, "Evaluating Estimates of Materials Offshoring from US Manufacturing," *Economics Letters* 117 (2012), 170–173.

Feenstra, R. C., B. R. Mandel, M. Reinsdorf, and M. Slaughter, "Effects of Terms of Trade Gains and Tariff Changes on the Measurement of U.S. Productivity Growth," *American Economic Journal: Economic Policy* 5 (2013), 1–38.

Feenstra, R. C., and J. Romalis, "International Prices and Endogenous Quality," *Quarterly Journal of Economics* 129 (2014), 477–527.

Foster, L., J. Haltiwanger, and C. Syverson, "The Slow Growth of New Plants: Learning About Demand?" *Economica* 83 (2015), 91–129.

Greenlees, J. S., and R. McClelland, "New Evidence on Outlet Substitution Effects in Consumer Price Index Data," this REVIEW 93 (2011), 632–646.

Greenstein, S. M., "Lock-In and the Costs of Switching Mainframe Computer Vendors in the US Federal Government in the 1970s," *IEEE Annals of the History of Computing* 17 (1995), 58–66.

Hallak, J. C., and P. Schott, "Estimating Cross-Country Differences in Product Quality," *Quarterly Journal of Economics* 126 (2011), 417–474.

Hausman, J., and E. Leibtag, "CPI Bias from Supercenters: Does the BLs Know That Wal-Mart Exists?" In E. W. Diewert, J. S. Greenlees, and C. R. Hulten, eds., *Price Index Concepts and Measurement* (Cambridge, MA: National Bureau of Economic Research, 2009).

Hottman, C., S. J. Redding, and D. E. Weinstein, "Quantifying the Sources of Firm Heterogeneity," *Quarterly Journal of Economics* 131 (2016), 1291–1364.

Houseman, S., C. Kurz, P. Lengermann, and B. Mandel, "Offshoring Bias in U.S. Manufacturing," *Journal of Economic Perspectives* 25 (2011), 111–132.

Houseman, S. N., and Kenneth F. Ryder, *Measurement Issues Arising from the Growth of Globalization: Conference Summary* (Kalamazoo, MI: W. E. Upjohn Institute for Employment Research, 2010).

Hummels, D., J. Ishii, and K.-M. Yi, "The Nature and Growth of Vertical Specialization in World Trade," *Journal of International Economics* 54 (2001), 75–96.

Hummels, D., and P. J. Klenow, "The Variety and Quality of a Nation's Exports," *American Economic Review* 95 (2005), 704–723.

Khandelwal, A., "The Long and Short (of) Quality Ladders," *Review of Economic Studies* 77 (2010), 1450–1476.

Khandelwal, A. K., P. K. Schott, and S.-J. Wei, "Trade Liberalization and Embedded Institutional Reform: Evidence from Chinese Exporters," *American Economic Review* 103 (2013), 2169–2195.

Klemperer, P., "Competition When Consumers Have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade," *Review of Economic Studies* 62 (1995), 515–539.

Klepper, S., "Entry, Exit, Growth, and Innovation over the Product Life Cycle," *American Economic Review* 86 (1996), 562–583.

Kugler, M., and E. Verhoogen, "Prices, Plant Size, and Product Quality," *Review of Economic Studies* 79 (2011), 307–339.

Monarch, R., "It's Not You, It's Me: Breakups in U.S.-China Trade Relationships," International Finance Discussion Paper 1165. Board of Governors of the Federal Reserve System (2016).

National Bureau of Economic Research. Price Statistics Review Committee, and George Joseph Stigler. *The Price Statistics of the Federal Government: Review, Appraisal, and Recommendations.* National Bureau of Economic Research Incorporated (1961).

Perla, J., "Product Awareness, Industry Life Cycles, and Aggregate Profits," University of British Columbia mimeograph (2015).

Reinsdorf, M., "The Effect of Outlet Price Differentials on the U.S. Consumer Price Index," (pp. 227–254), in Murry Foss, Marklin Manser, and Allan Young, eds., *Price Measurements and their Uses* (Chicago: University of Chicago Press, 1993).

Rosen, S., "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," *Journal of Political Economy* 82 (1974), 34–55.

Sorensen, A. T., "Equilibrium Price Dispersion in Retail Markets for Prescription Drugs," *Journal of Political Economy* 108 (2000), 833–850.

Stigler, G. J., and J. K. Kindahl, *The Behavior of Industrial Prices* (Cambridge, MA: NBER, 1970).

Syverson, C., "Prices, Spatial Competition and Heterogeneous Producers: An Empirical Test," *Journal of Industrial Economics* 55 (2007), 197–222.

Tirole, J., *The Theory of Industrial Organization* (Cambridge, MA: MIT Press, 1988).

Triplett, J., *Handbook on Hedonic Indexes and Quality Adjustments in Price Indexes* (Paris: OECD, 2006).

Turley, J. L., *The Essential Guide to Semiconductors* (Upper Saddle River, NJ: Prentice Hall PTR, 2003).