

Online Appendices

(Not for publication)

| | |
|---|-----------|
| A Price Indexes | 42 |
| A.1 BLS Linking | 42 |
| A.2 BLS Quantity Information | 43 |
| A.3 Fisher Index Formula | 43 |
| B Additional Theoretical Results | 44 |
| B.1 Index Comparisons | 44 |
| B.2 Model with Taste for Diversity | 46 |
| B.3 Equilibrium Dynamics | 47 |
| C Industry Background | 55 |
| D Data | 59 |
| D.1 Wafer Prices | 59 |
| D.2 Semiconductor Wafer Shipments | 61 |
| D.3 Weight Construction | 63 |
| D.4 Entry Timing | 63 |
| E Supplementary Empirical Results | 66 |
| E.1 Alternate Hedonic Regression Results | 66 |
| E.2 Singapore vs. Taiwan Price Dynamics | 70 |
| E.3 Within-Technology Price Variation Explained by Supplier | 71 |
| F Yields | 72 |
| F.1 Incidence of Yield Risk | 72 |
| F.2 Scope for Yield Variation | 72 |
| F.3 Yield Data Analysis | 73 |
| G CES Quality Inference | 79 |

A Price Indexes

A.1 BLS Linking

“Linking” refers to the process used to incorporate a new product into an index. In a matched-model framework, the most common approach to dealing with newly entering products is to drop them from the index, since one can not calculate a price relative until the second period in which a good is present. As an example, consider the BLS International Price Program (IPP), which includes the Import Price Index (MPI). The *BLS Handbook of Methods* describes the process when a new product is added to the index Bureau of Labor Statistics (1997). “When a completely new item series is added to a classification grouping, linking is not feasible. Instead, ... the historical movement of the index is used to begin the series for the item.” In other words, since a price relative can not be constructed for the new product, the existing products in the index determine the index growth until the subsequent period, when a price relative for the new product is available. This is equivalent to simply dropping the new product from the index in the entry period. Section 2 describes the pitfalls of this approach when entering and incumbent products have different quality-adjusted prices.

Occasionally, the physical characteristics or other aspects of an existing product may change. In this case, the BLS analyst has two options. The new version of the product may be considered a completely new item, in which case it is omitted from the index, as described in the previous paragraph. Alternatively, the analyst can link the new version into the existing series by estimating “what the ‘old’ item would trade for in the new time period. To do this, reporters are asked for the dollar amount attributed to the change. This value is then subtracted from the ‘new’ item in the current period” Bureau of Labor Statistics (1997). In other words, the pricing survey respondent is asked to estimate how much the old version’s price would have changed had it still been available. This estimate is then used to generate a hypothetical price relative for the product.

This linking procedure involving a subjective quality adjustment is used very rarely Nakamura and Steinsson (2012). The MPI micro data include a “link” field indicating when this procedure has been implemented. Benjamin Mandel, formerly of the Federal Reserve Board of Governors, kindly supplied calculations showing that, from 2004 to 2007, less than 2 percent of items in the full sample indicate using the linking procedure. In no case was linking used in the semiconductor industry (HS 8542) during this time period.

It is possible that subjective quality adjustment is rarely used because it is rarely needed. New products may be sufficiently distinct, or sourcing patterns may be very stable. In practice, however, this does not appear to be the case. William Alterman, then BLS Assistant Commissioner of International Prices, provided us with similar calculations from the MPI micro data for semiconductors. Between May 2009 and May 2011, less than 0.25 percent of items reported a shift in the country from which the item was imported. During the same time period, the Census End Use data indicate very large shifts in market share across countries, with China’s share of HS 8542 increasing by 6 percentage points (authors’ calculations).

Together, these facts imply that much of the substitution across countries occurs via the formation of new buyer-seller relationships. The MPI treats the physical item \times buyer pair as a “product,” and averages across price relatives at this level Bureau of Labor Statistics (1997). Therefore, when a new buyer enters and sources from a lower-price supplier, this transaction is not averaged with price levels of similar items to show an overall price decline. Instead it is omitted until the subsequent period, when a price relative for the buyer-seller pair can be calculated. This is the main source of the outlet substitution bias in the MPI.

A.2 BLS Quantity Information

The BLS calculates Laspeyres indexes using price and revenue share information. Although the Laspeyres formula (see equation (A1)) does not strictly require the collection of quantity information, the BLS already utilizes two sources of quantity information that can be used to calculate our proposed quality-adjusted unit value index in (8). The discussion here focuses on the International Price Program (IPP).³⁹ First, the BLS already collects quantity information from respondents who report that the quantity was important in determining the unit price. Thus, requesting quantity information from all respondents would not require any change to the price collection instruments or the data storage or processing systems already in use. Because the BLS uses sampling probabilities proportional to transaction size to select items for inclusion in the IPP, these quantities would yield an unbiased estimate of the average price we construct for our index. Second, IPP already uses customs data at the 10-digit Harmonized System (HS) level to calculate weights at the “detailed category” level Bureau of Labor Statistics (1997). These data include product \times establishment quantity information necessary to calculate quality-adjusted unit values at the 10-digit HS level.

A.3 Fisher Index Formula

The Fisher index is calculated as follows. First calculate Laspeyres and Paasche indexes, respectively, as

$$P_L^t = \sum_m s_m^{t-1} \frac{p_m^t}{p_m^{t-1}} \quad (\text{A1})$$

and

$$P_P^t = \left[\sum_m s_m^t \left(\frac{p_m^t}{p_m^{t-1}} \right)^{-1} \right]^{-1}, \quad (\text{A2})$$

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 expenditure in time t accounted for by model m . The Fischer index is a geometric mean of the Laspeyres and Paasche indexes.

$$P_F^t = \sqrt{P_L^t P_P^t} \quad (\text{A3})$$

Diewert (1976) notes that this index is superlative, meaning that it is an exact index for a fairly flexible production or utility function.

³⁹Thanks to Robert Sutton at BLS for details on IPP quantity information.

B Additional Theoretical Results

B.1 Index Comparisons

In this section, we compare the standard matched-model index and our proposed quality-adjusted index in (8) to the exact index in (2). We examine cases with constant quality (as assumed by both matched-model and our proposed index) and cases with different relative quality improvements across suppliers.

B.1.1 Matched-Model with Constant Quality

We begin by comparing the standard matched-model index to the exact index when quality is constant over time. Rearranging the exact index in (2) yields

$$I_t = \frac{\sum_{j \in J_t} \theta_{j,t} P_{j,t}}{\sum_{j \in J_{t-1}} \theta_{j,t-1} P_{j,t-1}}, \quad (\text{B1})$$

$$\text{where } P_{j,t} \equiv \frac{p_{j,t}}{\xi_j} \quad \text{and} \quad \theta_{j,t} \equiv \frac{\xi_j y_{j,t}}{\sum_{i \in J_t} \xi_i y_{i,t}}.$$

For simplicity, assume a single entrant in period t and that no products exit. Then the sets of suppliers are $J_{t-1} = \{1, \dots, N_{t-1}\}$ and $J_t = \{1, \dots, N_{t-1}, N_t\}$. Next, to make analytical headway, define $\bar{\theta}_j \equiv \alpha \theta_{j,t} + (1 - \alpha) \theta_{j,t-1}$, where $\alpha \in [0, 1]$, and take a first order approximation of (B1) around $\theta_{j,t} = \bar{\theta}_j$ and $\theta_{j,t-1} = \bar{\theta}_j$, $\forall j \in \{1, \dots, N_{t-1}\}$, and $P_{j,t} = P_{j,t-1}$. This yields

$$I_t \approx \underbrace{\sum_{j=1}^{N_{t-1}} \omega_j (P_{j,t}/P_{j,t-1})}_{(a)} + \underbrace{\sum_{j=1}^{N_{t-1}} \rho_j [\tilde{\theta}_{j,t} - (1 + \Omega_J) \tilde{\theta}_{j,t-1}]}_{(b)}. \quad (\text{B2})$$

$$\text{where } \omega_j \equiv \frac{\bar{\theta}_j P_{j,t-1}}{\sum_{i=1}^{N_{t-1}} \bar{\theta}_i P_{i,t-1}}, \quad \rho_j \equiv \frac{P_{j,t-1}}{\sum_{i=1}^{N_{t-1}} \bar{\theta}_i P_{i,t-1}}, \quad \tilde{\theta}_{j,t} \equiv \theta_{j,t} - \bar{\theta}_j, \quad \text{and} \quad \Omega_J \equiv \frac{\theta_{N_t,t} P_{N_t,t}}{\sum_{i=1}^{N_{t-1}} \bar{\theta}_i P_{i,t-1}}.$$

Term (a) in (B2) is a matched-model index, a weighted average of price relatives for continuing varieties. Since quality is constant over time, the quality-adjusted price relatives in (a) equal observable price relatives. Moreover, it is a Laspeyres index as in (A1) when $\alpha = 0$, since in that case ω_j is simply the expenditure share on product j . Thus, term (a) reflects the standard matched-model index used by the BLS.

To compare the standard matched-model index in (a) to the exact index, we must characterize term (b), which reflects changes in market shares. Dividing term (b) by N_{t-1} , and using the fact that $\tilde{\theta}_{j,t-1} = -\alpha \Delta \theta_{j,t}$, we have

$$(1 + \alpha \Omega_J) \frac{1}{N_{t-1}} \sum_{j=1}^{N_{t-1}} \rho_j \Delta \theta_{j,t} \quad (\text{B3})$$

$$= (1 + \alpha \Omega_J) \left[\text{cov}(\rho_j, \Delta \theta_{j,t}) + \left(\frac{1}{N_{t-1}} \sum_{j=1}^{N_{t-1}} \rho_j \right) \left(\frac{1}{N_{t-1}} \sum_{j=1}^{N_{t-1}} \Delta \theta_{j,t} \right) \right]. \quad (\text{B4})$$

To interpret the term in brackets, first note that ρ_j is the quality-adjusted price for product j compared to the share-weighted average quality-adjusted price across products. Therefore, if market

share weakly shifts away from products with initially high relative quality-adjusted prices, then $\text{cov}(\rho_j, \Delta\theta_{j,t}) \leq 0$. Next, if the entrant acquires positive market share ($\Omega_J > 0$), then incumbents must lose share on average, so $\sum_{j=1}^{N_{t-1}} \Delta\theta_{j,t} < 0$. It follows that term (b) is negative. Thus, we have established our first main result.

Proposition 1. *Assume that market shares weakly increase for suppliers offering low relative quality-adjusted prices. Then, to a first order approximation around initial market shares and prices, the matched-model index overstates true price growth:*

$$I_t < \sum_{j=1}^{N_{t-1}} \omega_j (P_{j,t}/P_{j,t-1}) \quad (\text{B5})$$

B.1.2 Matched-Model with Variable Quality

Now compare the matched-model index to the exact index in (2) when individual products' quality may improve over time, such that $\xi_{j,t}/\xi_{j,t-1} > 1$. In this case observed price relatives no longer reveal quality-adjusted price relatives, $p_{j,t}/p_{j,t-1} \neq P_{j,t}/P_{j,t-1}$. Rewrite (B5) as

$$I_t < \sum_{j=1}^{N_{t-1}} \omega_j (P_{j,t}/P_{j,t-1}) = \sum_{j=1}^{N_{t-1}} \omega_j \frac{p_{j,t}/p_{j,t-1}}{\xi_{j,t}/\xi_{j,t-1}} < \sum_{j=1}^{N_{t-1}} \omega_j (p_{j,t}/p_{j,t-1}). \quad (\text{B6})$$

This expression makes clear that quality improvement exacerbates the upward bias already present in matched-model indexes in the presence of quality-adjusted price variation.

B.1.3 Proposed Index with Variable Quality

Suppose that an entrant's ($j = N_t$) quality steadily improves over time in comparison to a reference supplier ($j = 1$) such that $\xi_{N_t,t}/\xi_{1,t} < \xi_{N_t,T}/\xi_{1,T} \forall t < T$. For simplicity, hold other suppliers' relative qualities constant. Also, as in the main text, assume that the effects of frictions have dissipated by period T . In this case, the long-run price difference overstates the entrant's relative quality in earlier periods, since $p_{N_t,T}/p_{1,T} = \xi_{N_t,T}/\xi_{1,T} > \xi_{N_t,t}/\xi_{1,t} \forall t < T$.

To understand the implications, rewrite the exact index in (2) in the period of entry as

$$I_t = G_t \frac{\sum_{j=1}^{N_{t-1}} \frac{\xi_{j,t-1}}{\xi_{1,t-1}} y_{j,t-1}}{\sum_{j=1}^{N_{t-1}} \frac{\xi_{j,t}}{\xi_{1,t}} y_{j,t} + \frac{\xi_{N_t,t}}{\xi_{1,t}} y_{N_t,t}}, \quad (\text{B7})$$

$$\text{where } G_t \equiv \frac{\sum_{j=1}^{N_t} p_{j,t} y_{j,t}}{\sum_{j=1}^{N_{t-1}} p_{j,t-1} y_{j,t-1}}$$

is the gross change in expenditure. Note that G_t is observable and does not contain any quality terms. Our proposed index in (8) can be written similarly.

$$\hat{I}_t = G_t \frac{\sum_{j=1}^{N_{t-1}} \frac{p_{j,T}}{p_{1,T}} y_{j,t-1}}{\sum_{j=1}^{N_{t-1}} \frac{p_{j,T}}{p_{1,T}} y_{j,t} + \frac{p_{N_t,T}}{p_{1,T}} y_{N_t,t}} \quad (\text{B8})$$

Since $p_{N_t,T}/p_{1,T} > \xi_{N_t,t}/\xi_{1,t}$, the denominator of (B8) is upward biased, and the index is downward biased. Thus, in the period of entry, the proposed index is biased downward, $\hat{I}_t < I_t$.

Now consider a period t after entry but prior to T . Assume that the supplier mix is now stable, such that $N_t = N_{t-1} = N$. Following (B7), the exact index can be written as

$$I_t = G_t \frac{\sum_{j=1}^N \xi_{j,t-1} y_{j,t-1}}{\sum_{j=1}^N \xi_{j,t} y_{j,t}}, \quad (\text{B9})$$

where G_t is the same as above. Our quality-inference procedure introduces offsetting errors relative to (B9). Since, by assumption, the quality of the entrant (supplier N) is weakly increasing over time, our estimation strategy overstates $\xi_{N,t-1}$ by more than $\xi_{N,t}$, which implies an upward bias. However, if the entrant gains market share over time, then our overstatement of $\xi_{N,t}$ receives more weight because it is multiplied by a larger $y_{N,t}$, which implies a downward bias. Intuitively, if quality evolves slowly in comparison to market share from one period to the next, then our proposed index will be biased downward on net, as in the period of entry. Put another way, our index is a lower bound if the change in market shares reflect less the changes in quality and more the dynamics that arise because of market frictions.

More formally, take the log of (B9), and a first-order approximation around $\xi_{j,t-1} = \xi_{j,T}$ and $\xi_{j,t} = \xi_{j,T}$, yielding

$$\ln I_t - \ln \hat{I}_t \approx \sigma_{N,t} \frac{\xi_{N,T} - \xi_{N,t}}{\xi_{N,t}} - \sigma_{N,t-1} \frac{\xi_{N,T} - \xi_{N,t-1}}{\xi_{N,t-1}}, \quad (\text{B10})$$

where $\sigma_{j,t} \equiv \xi_{j,t} y_{j,t} / Y_t$, and \hat{I}_t is our proposed index. Note that terms for $j \neq N$ drop out under the assumption that suppliers $j < N$ have constant quality. This expression is positive if and only if

$$\frac{\sigma_{N,t}}{\sigma_{N,t-1}} \geq \frac{\xi_{N,t}}{\xi_{N,t-1}} \left(\frac{\xi_{N,T} - \xi_{N,t-1}}{\xi_{N,T} - \xi_{N,t}} \right). \quad (\text{B11})$$

The right hand side is increasing in $\frac{\xi_{N,t}}{\xi_{N,t-1}}$, so it requires that quality growth for the entrant is not too large relative to the growth in its market share. We have thus established the following result.

Proposition 2. *Assume that a supplier's quality increases after entry but incumbents' qualities are fixed. If the rate of change in the new supplier's quality is not too large relative to its acquisition of market share, as in (B11), then the proposed index, \hat{I}_t , is a lower bound on the true index, I_t .*

Together with the results of the prior section, we can bound the exact price index between a standard matched-model index and our proposed index when i) quality improves over time, ii) the entrant's quality improves relative to that of incumbents, and iii) quality growth for the entrant is not too large compared to the growth in its market share.

B.2 Model with Taste for Diversity

This section presents a generalization of the framework in Section 2 to incorporate a taste for diversity on the part of the buyer. Assume that a buyer orders n_j units from supplier j , but that only $y_j = n_j^a$ units are delivered, with $a \in (0, 1)$. Bernstein et al. (2013) suggest $a < 1$ captures the idea that errors in production are made at a higher rate when a supplier fills a larger order. The resulting diseconomies of scale provide a motive to diversify the supply chain.

The buyer now solves

$$\min_{\{y_{j,t}\}} \sum_j p_{j,t} y_{j,t}^{1/a} \quad \text{s.t. (1)}, \quad (\text{B12})$$

where $p_j y_j^{1/a}$ is the total cost of a delivery y_j from supplier j and p_j is the price per order. The associated first order conditions are

$$\frac{p_{j,t}}{\xi_j} y_{j,t}^\beta = \frac{\mu_t}{1 + \beta} \quad \forall j, \quad (\text{B13})$$

where $1 + \beta \equiv 1/a$. The left side is the quality adjusted unit price of a delivery. Thus, although $a < 1$ breaks the one-to-one mapping from orders to deliveries, the law of one quality-adjusted price holds for deliveries.

This generalization also resolves a counterfactual prediction of the framework in the main text. Substituting the first order condition into the constraint, (1), we can solve for the quality adjusted price of a delivery, which is common across all products.

$$\frac{p_{j,t}}{\xi_j} y_{j,t}^\beta = \left(\frac{\sum_j \xi_j P_{j,t}^{-1/\beta}}{Y_t} \right)^{-\beta} \quad (\text{B14})$$

As $a \rightarrow 1$, as in the main text, then $\beta \rightarrow 0$, and this collapses to $P_{j,t} = 1$. All quality-adjusted prices equal 1, so the overall price index also equals 1, and does not change over time. When $a < 1$, we preserve the law of one price, but allow the overall price level to change over time with total demand, Y_t .

B.3 Equilibrium Dynamics

We introduce a tractable model of industry dynamics where market frictions slow the reallocation of market shares across suppliers. This is a variant on the model, (4), in the main text. To enhance tractability, we deviate from that setup in two respects. First, rather than stipulate an occasionally binding constraint on shifts in market share, we suppose the buyer faces a smooth, convex cost of adjusting its supplier mix. Second, we introduce a convex cost to assemble the inputs into final output, which ensures a unique steady state distribution of market shares.

We consider an environment in which a new supplier has entered, bringing the number of suppliers to N . We assume there is no subsequent supplier entry or exit and characterize the path of relative prices on the transition to the new steady state. The main result is stated in Proposition 3 at the end of this subsection. A supplier who has a market share below its steady state value, i.e. a new entrant, will initially set a low relative price. Over time it will raise its relative price as it (monotonically) acquires market share. To clarify the mechanism, this result is established for a simple case with constant supplier quality and marginal cost.

B.3.1 The buyer's problem

The buyer chooses quantities $\{y_{i,t}\}$ of inputs to minimize

$$\min_{\{y_{i,t}\}} \sum_t \bar{\beta}^t \left\{ \sum_{i=1}^N p_{i,t} y_{i,t} + \text{Cost of assembly}_t + \text{Cost of adjusting suppliers}_t \right\}, \quad (\text{B15})$$

subject to a given amount of final output,

$$Y_t = \sum_i \xi_i y_{i,t}. \quad (\text{B16})$$

There are three pieces to (B15). The first term is expenditure on the intermediates. The second term refers to the cost of assembling intermediates into final output. We assume this has the form,

$$\text{Cost of assembly}_t = Y_t \cdot C\left(\{\sigma_{i,t}\}_{i=1}^N\right) \quad (\text{B17})$$

where $C(\cdot)$ is defined implicitly as the unit cost of assembly, and

$$\sigma_{i,t} \equiv \frac{\xi_i y_{i,t}}{Y_t} \quad (\text{B18})$$

is the share of intermediate i in final output. We assume C is increasing in each argument, that is, the contribution of intermediate i to the unit cost depends positively on its share in final output. Further, we assume C is convex in each argument. Specifically, to facilitate the algebra, we assume it takes a quadratic form,

$$C\left(\{\sigma_{i,t}\}_{i=1}^N\right) = \frac{\delta}{2} \sum_i \sigma_{i,t}^2. \quad (\text{B19})$$

This is a highly tractable way of ensuring that all intermediates are purchased in steady state equilibrium, avoiding the indeterminacy that would result in steady state because expenditure is linear in $y_{i,t}$.

The third term refers to the cost of adjusting the mix of suppliers. We assume this has the form,

$$\text{Cost of adjusting suppliers}_t = \frac{k}{2} \sum_i Y_t [\Delta\sigma_{i,t}]^2, \quad \text{with } k > 0. \quad (\text{B20})$$

This says that the cost of adjusting is proportional to the sum of squared changes in each supplier's share. Note that we also scale the cost by final output, Y . This facilitates some of the algebra, but also captures the reasonable assumption that it is arguably more disruptive to the buyer to alter its supplier mix when output is high.

It is helpful to recast the problem in terms of the choice of market shares, $\{\sigma_{it}\}$. If we substitute from (B18) to replace $y_{i,t}$, we may write the Lagrangian associated with the buyer's problem as

$$\mathcal{L} = \sum_t \bar{\beta}^t \left\{ Y_t \sum_i P_{i,t} \sigma_{i,t} + \frac{\delta}{2} \sum_i Y_t [\sigma_{i,t}]^2 + \frac{k}{2} \sum_i Y_t [\Delta\sigma_{i,t}]^2 + \mu_t \left(1 - \sum_i \sigma_{i,t} \right) \right\}, \quad (\text{B21})$$

where the quality-adjusted price is

$$P_{i,t} \equiv \frac{p_{i,t}}{\xi_i}, \quad (\text{B22})$$

and μ_t is the multiplier on the (B16), re-expressed here as a constraint that market shares must sum to 1.

The FOC is

$$P_{i,t} + \delta\sigma_{i,t} + k\Delta\sigma_{i,t} - \bar{\beta}_{t+1,t} k\Delta\sigma_{i,t+1} = m_t \text{ for all } i \quad (\text{B23})$$

where $\bar{\beta}_{t+1,t} \equiv \bar{\beta} Y_{t+1}/Y_t$. Note that, in any steady state where $\Delta\sigma_{i,t} = \Delta\sigma_{i,t+1} = 0$ for all i , this collapses to $P_{i,t} + \delta\sigma_{i,t} = m_t$. Thus, $\delta > 0$ drives a wedge between quality-adjusted prices even in steady state. But we can interpret δ as being "small" (subject to a weak restriction given below), which is akin to assuming an arbitrarily high elasticity of substitution in a standard constant elasticity of substitution (CES) model (see Appendix G). Quantitatively, this framework captures the idea that long run quality-adjusted prices approximately equalize.

Consider the case where final demand grows at a constant rate, such that Y_{t+1}/Y_t is constant

for all t . Then, $\bar{\beta}_{t+1,t} \equiv B$ is constant, and we assume $B \in (0, 1)$. In the following, we also assume that, though $D \equiv \delta/k$ can be small, it satisfies $D > 1 - B$. Aggregating over all N suppliers, noting that $\frac{1}{N} \sum_i \Delta \sigma_{i,t} = 0$ for any t , and defining the relative (quality-adjusted) price,

$$\hat{P}_{i,t} \equiv P_{i,t} - \frac{1}{N} \sum_i P_{i,t}, \quad (\text{B24})$$

we can rewrite the FOC as

$$B\kappa\sigma_{i,t+1} - \sigma_{i,t} + \kappa\sigma_{i,t-1} = \alpha\hat{P}_{i,t} - \frac{\delta\alpha}{N}, \quad (\text{B25})$$

where $\alpha \equiv \frac{1}{\delta+(1+B)k}$ and $\kappa \equiv k\alpha$. This is a second-order difference equation that can be solved by conventional techniques.

Lemma 1. *The solution to (B25) is*

$$\sigma_{i,t} = \frac{D}{B} \frac{1}{\lambda - 1} + \frac{1}{B\lambda} \sigma_{i,t-1} - \frac{1}{Bk\lambda} \sum_{j=0} \left(\frac{1}{\lambda}\right)^j \hat{P}_{i,t+j}, \quad (\text{B26})$$

where $D \equiv \delta/(Nk)$, $\lambda \equiv \frac{1+\sqrt{1-4B\kappa^2}}{2B\kappa} > 1$, and, if $D > 1 - B$, $B\lambda > 1$.

Proof. Define

$$s_{i,t} \equiv \sigma_{i,t-1} \quad (\text{B27})$$

and rewrite the FOC as a system,

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \sigma_{i,t+1} \\ s_{i,t+1} \end{pmatrix} = \begin{pmatrix} (B\kappa)^{-1} & -B^{-1} \\ 1 & 0 \end{pmatrix} \begin{pmatrix} \sigma_{i,t} \\ s_{i,t} \end{pmatrix} + \begin{pmatrix} \alpha(B\kappa)^{-1} \left(\hat{P}_{i,t} - \frac{\delta}{N}\right) \\ 0 \end{pmatrix} \quad (\text{B28})$$

or,

$$W_{i,t+1} = QW_{i,t} + Z_{i,t}. \quad (\text{B29})$$

There is one real eigenvalue, $\lambda > 1$, associated with Q that lies outside the unit circle. Denote the corresponding eigenvector by v such that

$$v'(Q - \lambda I) = 0. \quad (\text{B30})$$

Now substituting $Q = \lambda I$ into (B29) and premultiplying each side by v' enables us to rewrite (B29) as

$$\omega_{i,t} = -\frac{\alpha}{B\kappa\lambda} \left(\hat{P}_{i,t} - \frac{\delta}{N}\right) + \frac{1}{\lambda} \omega_{i,t+1}, \quad (\text{B31})$$

where

$$\omega_{i,t} \equiv v'W_{i,t} = \sigma_{i,t} - \frac{1}{B\lambda} s_{i,t}, \quad (\text{B32})$$

Since $1/\lambda < 1$, this can be solved forward to yield (B26). ■

B.3.2 The supplier's problem

The supplier in time t chooses its price to maximize the present value of profits,

$$\sum_{j=0} \bar{\beta}^j y_{i,t+j} [p_{i,t+j} - c_{i,t+j}], \quad (\text{B33})$$

subject to (B26). We assume for simplicity that each supplier i faces a flat marginal cost schedule, $c_{i,t+j}$. This can be re-expressed in terms of market shares,

$$Y_t \sum_{j=0} \beta_{t+j,t} \sigma_{i,t+j} [P_{i,t+j} - x_{i,t+j}], \quad (\text{B34})$$

where $x \equiv c/\xi$. For some results below, we will impose that x is fixed for each supplier, but that is not necessary at this point. Again, assume $\beta_{t+j,t} \equiv B^j$ in what follows.

The FOC for the quality-adjusted price, $P_{i,t}$, is straightforward. Recalling that

$$\hat{P}_{i,t+j} \equiv P_{i,t+j} - \frac{1}{N} \sum_{\iota} P_{\iota,t+j} = \frac{N-1}{N} P_{i,t+j} - \frac{1}{N} \sum_{\iota \neq i} P_{\iota,t+j}, \quad (\text{B35})$$

and using (B26) to calculate that $\frac{\partial \sigma_{i,t+j}}{\partial P_{i,t}} = -\left(\frac{1}{B\lambda}\right)^{j+1} \frac{N-1}{Nk}$ for $j \geq 0$, we have

$$P_{i,t} - x_{i,t} = \frac{Nk}{N-1} B\lambda \sigma_{i,t} - \sum_{j=1} \frac{1}{\lambda^j} [P_{i,t+j} - x_{i,t+j}]. \quad (\text{B36})$$

This says that the supplier's markup depends positively on its market share and negatively on future markups. The latter feature reflects the fact that high future markups depress future market share. As a result, if a supplier prices high in the future, it will price low now to "build a customer base."

B.3.3 Equilibrium

We can now solve for equilibrium market shares and prices.

Market shares Defining $\sigma \equiv (\sigma_1 \ \sigma_2 \ \cdots \ \sigma_N)'$ and stacking the demand schedules in (B26) over suppliers, we have

$$\sigma_t = c_\sigma + \frac{1}{B\lambda} \sigma_{t-1} - \frac{1}{B\lambda} \frac{N-1}{Nk} \sum_{j=0} \frac{1}{\lambda^j} \tilde{\mathbf{A}} \mathbf{P}_{t+j}, \quad (\text{B37})$$

where $c_\sigma \equiv \frac{D}{B} \frac{1}{\lambda-1}$, $\tilde{\mathbf{A}} \equiv \mathbf{I} - \frac{1}{N-1} \mathbf{M}$, and \mathbf{M} is a $N \times N$ hollow matrix where each off-diagonal element is one and each diagonal element is zero. Similarly, collecting the suppliers' price-setting rules (B36) into matrix form and using (B26) to substitute out σ yields

$$\mathbf{A} \mathbf{P}_t = c_P + \frac{Nk}{N-1} \sigma_{t-1} + \varkappa_t - \sum_{j=1} \frac{1}{\lambda^j} \mathbf{A} \mathbf{P}_{t+j}, \quad (\text{B38})$$

where $\mathbf{A} \equiv \mathbf{I} + \tilde{\mathbf{A}}$, $\mathbf{P} \equiv (P_1 \ P_2 \ \cdots \ P_N)'$, $c_P \equiv D \frac{Nk}{N-1} \frac{\lambda}{\lambda-1}$,

$$\varkappa_t \equiv (\varkappa_{1,t} \ \cdots \ \varkappa_{N,t})', \text{ and } \varkappa_{i,t} \equiv \sum_{j=0} \frac{1}{\lambda^j} x_{i,t+j}. \quad (\text{B39})$$

We are now prepared to characterize the dynamics of market share.

Lemma 2. *Assume there is a finite φ such that, for all $j \geq \varphi$, $x_{i,t+j} = x_i$ for each i . Then there is a unique, stable steady state solution to (B37), σ^* .*

Proof. Substituting for \mathbf{P}_t in (B37) using (B38) collapses the former to

$$\sigma_t = \mathbf{C}_0 + \mathbf{C}_\sigma \sigma_{t-1} - \mathbf{C}_\varkappa \varkappa_t, \quad (\text{B40})$$

where

$$\mathbf{C}_0 \equiv c_\sigma - \frac{1}{B\lambda} \frac{N-1}{Nk} \tilde{\mathbf{A}} \mathbf{A}^{-1} c_P, \quad (\text{B41})$$

$$\mathbf{C}_\sigma \equiv \frac{1}{B\lambda} \left[\mathbf{I} - \tilde{\mathbf{A}} \mathbf{A}^{-1} \right], \text{ and } \mathbf{C}_\varkappa \equiv \frac{1}{B\lambda} \frac{N-1}{Nk} \tilde{\mathbf{A}} \mathbf{A}^{-1}. \quad (\text{B42})$$

\mathbf{A} is indeed invertible by the Sherman-Morrison theorem. Defining \mathbf{J} as the unit matrix (ie, the matrix of ones) and noting that $\mathbf{M} = \mathbf{J} - \mathbf{I}$, we have

$$\mathbf{A} = \left(2 + \frac{1}{N-1} \right) \mathbf{I} - \frac{1}{N-1} \mathbf{J}. \quad (\text{B43})$$

This has the form $\mathbf{A} = \eta \mathbf{I} - \mathbf{u} \mathbf{v}^T$, where $\eta \equiv 2 + \frac{1}{N-1}$; \mathbf{u} is a $N \times 1$ vector with each element equal to $\frac{1}{N-1}$; and \mathbf{v} is a $N \times 1$ vector of ones. Since $\eta \mathbf{I}$ is invertible, Sherman-Morrison implies $\mathbf{A}^{-1} = (\eta \mathbf{I} - \mathbf{u} \mathbf{v}^T)^{-1} = \eta^{-1} \left(\mathbf{I} - \frac{\mathbf{u} \mathbf{v}^T}{\eta + \mathbf{v}^T \mathbf{u}} \right)$.

The steady state solution to (B40) is

$$\sigma^* = [\mathbf{I} - \mathbf{C}_\sigma]^{-1} [\mathbf{C}_0 - \mathbf{C}_\varkappa \varkappa], \quad (\text{B44})$$

where $\varkappa \equiv (\varkappa_1 \ \dots \ \varkappa_N)'$ and $\varkappa_i = \frac{\lambda}{\lambda-1} x_i$ in a steady state. To confirm that $\mathbf{I} - \mathbf{C}_\sigma$ is invertible, one can first prove by induction that

$$\mathbf{I} - \tilde{\mathbf{A}} \mathbf{A}^{-1} = \frac{1}{2N-1} [N\mathbf{I} + \mathbf{M}]. \quad (\text{B45})$$

Again using $\mathbf{M} = \mathbf{J} - \mathbf{I}$ and the definition of \mathbf{C}_σ , we now have that

$$\mathbf{I} - \mathbf{C}_\sigma = \left(1 - \frac{1}{B\lambda} \frac{N-1}{2N-1} \right) \mathbf{I} - \frac{1}{B\lambda} \frac{1}{2N-1} \mathbf{J}, \quad (\text{B46})$$

which has the same form as (B43) and is thus invertible.

To establish stability, first rearrange (B40) to read

$$\hat{\sigma}_t = \mathbf{C}_\sigma \hat{\sigma}_{t-1} - \mathbf{C}_\varkappa \hat{\varkappa}_t, \quad (\text{B47})$$

where $\hat{\sigma}_t \equiv \sigma_t - \sigma^*$ and $\hat{\varkappa}_t \equiv \varkappa_t - \varkappa^*$. Iterating backward L periods yields

$$\hat{\sigma}_t = \mathbf{C}_\sigma^L \hat{\sigma}_{t-L} - \sum_{l=0}^{L-1} \mathbf{C}_\sigma^l \mathbf{C}_\varkappa \hat{\varkappa}_{t-l}. \quad (\text{B48})$$

If $\mathbf{C}_\sigma^L \rightarrow 0$ as $L \rightarrow \infty$, the system is stable. This can be confirmed by inspection: from (B45), each element of $\mathbf{I} - \tilde{\mathbf{A}} \mathbf{A}^{-1}$ is positive and less than or equal to $N/(2N-1)$, which is in turn less than 1 for integer-valued $N > 1$. Since $B\lambda < 1$, it follows that each element of \mathbf{C}_σ lies on the unit line. ■

An instructive special case of (B40) permits an even tighter characterization of market shares along the transition to steady state.

Corollary. *Assume $x_{i,t+j} = x_i$ for all j and each i . Then suppliers that start with market shares below (above) steady state will acquire (lose) market share monotonically until steady state*

is reached.

Proof. Suppose that, in period $t - 1$, supplier i has market share 1 percentage point below its steady state and, therefore, all suppliers $j \neq i$ have, in the aggregate, market share 1 percentage point above steady state. Since $\hat{\varkappa}_t \equiv 0$, we can calculate $\hat{\sigma}_{it}$ using the i th row of \mathbf{C}_σ and the fact that $\sum_{j \neq i} \hat{\sigma}_{jt} = -1$. We have

$$\hat{\sigma}_{i,t} = \frac{1}{B\lambda} \frac{1}{2N-1} [N - (N-1)] = \frac{1}{B\lambda} \frac{1}{2N-1} > 0. \quad (\text{B49})$$

Thus, supplier i acquires market share, and continues to do so until steady state is reached. ■

Prices We can next solve for equilibrium relative prices, $\{\hat{P}_{i,t}\}$. Leading (B38) one period and using this to substitute for \mathbf{P}_{t+1} enables us to collapse the price setting rule to

$$\mathbf{P}_t = \frac{\lambda-1}{\lambda} \mathbf{A}^{-1} c_P + \mathbf{A}^{-1} \left[\varkappa_t - \frac{1}{\lambda} \varkappa_{t+1} \right] + \frac{Nk}{N-1} \mathbf{A}^{-1} \left[\sigma_{t-1} - \frac{1}{\lambda} \sigma_t \right]. \quad (\text{B50})$$

Now substituting for σ_t using (B40) and premultiplying each side by $\frac{N-1}{N} \tilde{\mathbf{A}}$ shows that the relative prices are

$$\hat{\mathbf{P}}_t = \frac{N-1}{N} \tilde{\mathbf{A}} \left\{ \mathbf{F}_0 + (\lambda \mathbf{A})^{-1} [\mathbf{F}_\varkappa \varkappa_t - \varkappa_{t+1}] \right\} + \frac{1}{\lambda} \tilde{\mathbf{A}} \mathbf{A}^{-1} [\lambda \mathbf{I} - \mathbf{C}_\sigma] k \sigma_{t-1}, \quad (\text{B51})$$

where

$$\mathbf{F}_0 \equiv (\lambda \mathbf{A})^{-1} \left[(\lambda - 1) c_P - \frac{Nk}{N-1} \mathbf{C}_0 \right], \quad \text{and} \quad \mathbf{F}_\varkappa \equiv \lambda \mathbf{I} + \frac{Nk}{N-1} \mathbf{C}_\varkappa. \quad (\text{B52})$$

Again consider the case $\varkappa_t = \varkappa_{t+1} \equiv \varkappa$. Therefore, the evolution of market share is the exclusive driver of movements in relative prices, and the pass through depends on $\tilde{\mathbf{A}} \mathbf{A}^{-1} [\lambda \mathbf{I} - \mathbf{C}_\sigma] k / \lambda$. Using (B45), we can see that

$$\begin{aligned} \tilde{\mathbf{A}} \mathbf{A}^{-1} [\lambda \mathbf{I} - \mathbf{C}_\sigma] &= \left\{ \frac{N-1}{2N-1} \left(\lambda - \frac{1}{B\lambda} \frac{N}{2N-1} \right) + 2 \frac{1}{B\lambda} \left(\frac{1}{2N-1} \right)^2 \right\} \mathbf{I} \\ &\quad - \left\{ \lambda - \frac{1}{B\lambda} \frac{2}{2N-1} \right\} \frac{1}{2N-1} \mathbf{M}, \end{aligned} \quad (\text{B53})$$

where we have used that $\mathbf{M}^2 = 2\mathbf{I} + \mathbf{M}$. The effect on supplier i 's relative price of a 1 percentage point increase in own market share, matched by a 1 percentage point decline in market share among all suppliers $j \neq i$, can now be calculated using the i th row of this matrix. We find that the relative price shifts by

$$\frac{2k}{\lambda} \left\{ \frac{N-1}{2N-1} \left(\lambda - \frac{1}{B\lambda} \frac{N+2}{4N-2} \right) + \frac{1}{B\lambda} \left(\frac{1}{2N-1} \right)^2 \right\} > 0. \quad (\text{B54})$$

The following proposition collects the implications of the preceding results.

Proposition 3. *Assume $x_{i,t+j} = x_i$ for all j and each i . If a supplier's market share starts below steady state, its relative price and market share will start low but increase monotonically on the transition to the new steady state.*

B.3.4 Quantitative Illustration

We now briefly illustrate the model’s quantitative implications. We consider a scenario in which the initial steady state includes two suppliers. One represents Taiwan and the other refers to the rest of the world, excluding China. Then a new supplier, China, enters. We refer to Taiwan as T , the rest of the world as R , and China as C . We solve for market shares and relative prices on the transition path to the new steady state with $N = 3$.

There are four parameters to pin down: δ , the marginal cost of integrating an individual supplier’s delivery into final output; k , the marginal cost of adjusting purchases from an individual supplier; and x_T and x_C , the ratio of, respectively, Taiwan’s and China’s marginal cost of production to their quality. Note that x_R is normalized to 1, so x_T and x_C are expressed relative to the rest of the world.

We use four moments of the data to calibrate these parameters. The parameters x_T and x_C should communicate considerable information about the supplier’s market shares in the new steady state. Consistent with Table 1, we target Taiwan’s market share of 50 percent and China’s share of 22.5 percent. Next, we have shown that, since the cost of adjusting purchases across suppliers gives incumbents an advantage, entrants will initially set low relative prices. Hence, the adjustment cost parameter, k , determines the size of the initial quality-adjusted discount offered by the entrant. Consistent with Figure 2, we target a quality-adjusted discount of 30 percent one year since entry.

Lastly, we turn to δ . As noted above, $\delta > 0$ breaks the assumption of perfect substitutes and implies a degree of steady state quality-adjusted price dispersion. This conflicts with the presumption in the main text that the law of one price obtains in the long run. However, the notion of a long-run law of one price is still attractive for a few reasons. First, in other reasonable variants of the model, a law of one price may obtain in the long run, even if the assumption of perfect substitutes is relaxed (see Appendix B.2). The specification with $\delta > 0$ is used here only because it preserves the highly tractable linear-quadratic structure of the model. Second, a long run law of one price is highly practical, since it yields measurement strategies that are simpler for measurement agencies to implement. For these reasons, we select δ to minimize deviations from the law of one price in the long run, i.e. $P_C = P_T$.

The results of the calibration are as follows. First, Taiwan’s relatively high quality implies a relatively low value of x_T . Specifically, we find $x_T = 0.8175$ and $x_C = 1.0425$. Next, the cost of adjusting supplier mix is rather modest. We find that $k = 0.36$, which implies that total adjustment costs (i.e., $0.5k \times \sum_i [\Delta\sigma_{i,t}]^2$) amount to just 0.7 percent of final output (Y) in the year of entry. Thus, adjustment frictions do not have to be particularly large to drive the price dynamics we observe. Lastly, $\delta = 0.03$. The model reproduces nearly exactly the targeted long-run market shares and the size of the initial discount. The choice of $\delta = 0.03$ restricts the long-run difference between TSMC and SMIC to just under 5 percent.⁴⁰


Figure B.1 shows the main result of this quantification exercise. It reports China’s (the entrant) price relative to Taiwan’s (the market leader).⁴¹ China’s quality-adjusted relative price is 30 percent lower in the first year following entry and increases monotonically thereafter. Relative to Figure 2,

⁴⁰If China’s quality-adjusted price is indeed lower than Taiwan’s, even in steady state, then our quality-adjustment over-corrects for quality by attributing too much long-run price variation to quality differences.

⁴¹Proposition 3 characterizes the relative price, defined as the difference between the price level of the entrant and the arithmetic average of prices. This relative price follows a similar

the model tends to over-predict the rate at which China catches up in the second year after entry. Nonetheless, the relative price does not fully settle into steady state for 5 years, which is consistent with the data. Note that these dynamics in relative prices occur in spite of the fact that we have restricted quality differences across suppliers to be constant over time.

Figure B.1: Quality-Adjusted Price of Entrant Relative to Market Leader



Results of model calibration, reporting the transition path for China's quality-adjusted price relative to Taiwan in each year following China's entry. See text for details of the model and calibration.

path to that in Figure B.1. For comparability to our empirical work, the figure plots the ratio of the entrant's price to the price of the leading-edge supplier.

C Industry Background

This section provides additional detail on the contract semiconductor manufacturing industry as a supplement to Section 3.

In semiconductor fabrication, transistors are created on the surface of a semiconducting wafer through a photolithography process. 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 cut up, leaving individual chips, called “die,” that are placed inside protective packages and connected to metal leads that allow the chip to be connected to other components.

As discussed in Section 3, semiconductor wafer technology progresses in discrete steps, either involving a larger wafer or a smaller line width. Because larger wafers contain more chips of a given size, and many steps in the fabrication process are implemented for the entire wafer at once, there are economies of scale in wafer size. The move to a larger wafer has generally reduced the cost per die by approximately 30 percent Kumar (2007). Similarly, moving a given chip design to a 30 percent smaller line width results in cost savings of approximately 40 percent, assuming the same number of defects in both processes Kumar (2007).

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. In spite of these challenges, the benefits of increased die per wafer and better performance have outweighed the costs of yield reductions, which improve as the fabrication technology matures. Given the benefits of smaller line widths, semiconductor manufacturers have steadily moved toward newer technology. This is apparent in Figure C.1, which plots the line-width 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 transistors and for how the transistors are arranged to implement logical functions. The most common technology, called complementary metal-oxide semiconductor (CMOS), a silicon-based chemical process, accounted for 97 percent of worldwide semiconductor production in 2008 (SICAS Semiconductor International Capacity Statistics). Other transistor arrangements, such as bipolar logic, and other chemical processes, such as gallium arsenide (GaAs) or silicon germanium (SiGe), generally focus on niche markets for high-frequency, high power, or aerospace devices, rather than the storage and computational logic products comprising the majority of the CMOS market. We therefore restrict our analysis to CMOS.

In the early 1970s, nearly all semiconductor producers were vertically integrated, with design, wafer fabrication, packaging, testing, and marketing performed within one company. Firms that perform both design and wafer fabrication are referred to as integrated device manufacturers (IDM). By the mid-1970s, IDMs began moving packaging and test operations to East Asia to take advantage of lower input costs Scott and Angel (1988); Brown and Linden (2006). In spite of offshoring these relatively simple steps in the production process, firms maintained the more complex wafer fabrication operations in the home country. 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 (2001); 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 et al. (2007). Around the same time, new contract manufacturing firms sprang 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 factoryless goods producers, which have little or no in-house manufacturing capability, are called "fabless" firms. See Bayard et al. (2015) and Bernard and Fort (2013) for other studies of factoryless goods producers. 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 C.2. Note that the shares in Figure C.2 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. Some of the most prestigious U.S. chip makers, such as Fortune 500 firms Broadcom and AMD, are fabless firms. The foundry and fabless firm structure creates arms-length transactions for semiconductor wafer fabrication services that avoid price measurement difficulties arising from transfer pricing between related parties. We focus on these arms-length transactions in our empirical analysis.


Along with the shift from an integrated manufacturer to a foundry business model came a shift in production capacity toward Asia, where most large foundries are located. Table 1 shows how the share of worldwide foundry capacity has evolved in the last decade. In 2000, the majority of foundry capacity was already in Asia, mainly Taiwan. Since then, the share of capacity in Asia as a whole has only increased modestly, but there has been a notable shift in capacity within Asia. In particular, China has more than tripled its share of foundry capacity, largely at the expense of the industry leader, Taiwan. Note that the sharp increase in European capacity from 2008 to 2010 marks the founding of Global Foundries, which was the fab division of former integrated device manufacturer AMD.

Economists have devoted substantial attention to studying semiconductor production in part because of semiconductors' outsized contribution to productivity growth Byrne et al. (2013), and partly in an effort to uncover the sources of rapid constant-quality price declines observed for high-tech products such as computers Berndt and Rappaport (2001) and communications equipment Doms (2005); Byrne and Corrado (2015). Attention has focused on the most important semiconductor components of computers, namely microprocessors Dulberger (1993); Grimm (1998); Doms et al. (2003); Holdway (2001); Flamm (2007); Byrne et al. (2015) and memory chips Flamm (1993); Grimm (1998); Aizcorbe (2002). There is little price variation across suppliers in these markets, so related studies instead focus on learning-by-doing to explain stunning rates of average price decline across suppliers. For models of learning by doing in semiconductor production, see Baldwin and Krugman (1988), Irwin and Klenow (1994), and Flamm (1996), all of which assume the law-of-one-price holds.

Foundries instead specialize in custom products called Application-Specific Integrated Circuits (ASICs). These products have been the subject of limited previous research and differ from memory and processors in important ways. ASICs are produced in smaller batches, and each model


requires a substantial investment in design. Foundries also use process technologies that are one generation or more behind the leading edge. New process technologies are first used outside the contract semiconductor manufacturing industry by large integrated device manufacturers that produce primarily memory chips and processors. Only when these technologies are relatively mature do contract manufacturers begin using them to produce ASICs. Hence, the arrival of new process technologies in the foundry market is largely determined by forces external to the contract manufacturing industry.

Figure C.1: TSMC Sales by Line Width



Authors' calculations based on quarterly reports from the largest semiconductor contract manufacturing firm, Taiwan Semiconductor Manufacturing Corporation (TSMC). Each line represents the share of TSMC sales accounted for by wafers with the stated line width(s). While the TSMC reports group together certain sets of adjacent line widths (e.g. 250/350nm), our main analysis uses data reporting each line width individually.

Figure C.2: Fabless Firms' Share of Global Semiconductor Revenue



Global fabless firm revenue from the Global Semiconductor Alliance (GSA) divided by global semiconductor industry revenue from the Semiconductor Industry Association's (SIA), Semiconductor Industry Bluebook.

D Data

In this appendix, we describe the data sources and steps taken in the construction of the dataset used in our analysis.

D.1 Wafer Prices

As discussed in Section 4, our wafer price data come from a proprietary database of semiconductor wafer purchases from foundries, collected by the Global Semiconductor Alliance (GSA). The survey began in 2001, but was substantially revised in 2004, so the sample we received from GSA begins in the first quarter of that year and ends in the last quarter of 2010. We implement a number of data cleaning procedures before the main analysis.

First, we impose a variety of sample restrictions. We drop observations corresponding to engineering runs that occur during the design stage prior to volume production. We omit a few observations reporting the obsolete 100mm wafer diameter. We keep only observations corresponding to CMOS process technology, which dominates the foundry market, and omit other processes that are quite distinct and serve niche markets for high-power, defense, or aerospace applications. We also keep only observations for wafers produced by so-called “pure-play” foundries, omitting transactions involving integrated device manufacturers (IDMs) that do both design and fabrication. Finally, we omit output of countries that focus heavily on non-ASIC products, leading us to drop Europe, Japan, and Korea, although results including these producers are nearly identical. Compared to the countries in our sample, foundries producing primarily in the dropped countries derived a much larger share of their revenue from analog devices (41.5% vs. 11.0%), a larger share from discrete devices and memory products (19.0% vs. 7.7%), and a much smaller share from computational logic devices characterizing the ASIC market (35.3% vs. 70.3%) (calculations from IHS iSuppli). The countries in our sample accounted for 86.9% of foundry revenue in 2010.

Table D.1 shows that 6,916 observations satisfy these sample restrictions. We drop an additional 663 observations for a variety of reasons. 576 observations do not report the location of production, and 1 observation reports an implausibly large order, which distorts the price per wafer. We confirmed with GSA that this order size must reflect misreporting. 86 observations report technologies (wafer-size, line-width combinations) for which data from IHS iSuppli (see below) report no production in the reported country and quarter. Because the IHS iSuppli database is based on in-person visits to each factory, we conclude these few inconsistent price reports are erroneous, reflecting reporting error for one or more of the technological characteristics.

Finally, we combine a few closely related line widths. Any line width greater than or equal to 500nm is combined into the 500nm code. 140 and 150nm widths are combined into the 150nm code. 80 and 90nm widths are combined into the 90nm code. 60 and 65nm widths are combined into the 65nm code. 40 and 45nm widths are combined into the 45nm code. In all of these cases, one of the combined widths is vastly dominant in the market, and it would be difficult to separately identify prices for the less prevalent line width with very few observations.

Table D.1: Dropped Observations

| | |
|-----------------------------------|-------|
| Total observations | 6,916 |
| Used in analysis | 6,253 |
| Dropped | 663 |
| Missing foundry location | 576 |
| Implausibly large order reported* | 1 |
| Inconsistent | 86 |

There may be multiple reasons to drop an observation

* Confirmed with GSA

Authors' calculations based on GSA Wafer Fabrication and Back-End Pricing Survey.

D.2 Semiconductor Wafer Shipments

As discussed in Section 4, we combine the transaction-level quantity information from GSA with quantity information at the quarter \times country \times technology cell level from the IHS iSuppli Pure Play Foundry Market Tracker database. This report is a census of all semiconductor foundries worldwide, including 91 fabs belonging to 20 companies. Annual and quarterly frequency data begin in 2000 and 2002, respectively. Characteristics provided for each fab include company of ownership, location, wafers shipped per month at full capacity, and diameter of wafers shipped. In addition, the report provides information on wafer shipments for specific technologies (line widths) by *company*, but not by *plant*. For 11 of the companies covered, this information is sufficient to construct wafer size \times line width \times country weights for our analysis without further assumptions. In two cases, only one fab is active in the period we cover, and in nine cases, all the company's fabs employing the same wafer diameter and are located in the same country.

The remaining companies either have fabs in multiple regions, fabricate wafers of multiple diameters, or both. In these cases, we first calculate estimates of wafer shipments by technology for each fab, which allows us to divide production between countries for each technology. To do this, we employ three additional resources: the partial information on the timing of technology introduction by plant from the IHS iSuppli report; company information provided in public statements; and extensive discussions with iSuppli personnel. Specifically, the Market Tracker lists the technologies employed (without output shares) in each fab at the time of publication, which we have for previous vintages of the database beginning in Q1 2010. This information allows us to construct technology-by-country-by-quarter weights for an additional 2 companies. In these cases, each with two fabs in operation, we were able to identify one company fab exclusively employing a single category of legacy technology (500nm and above), leaving the remaining fab to account for the residual company production.

In the remaining 7 cases, to calculate shipments by technology, we need further information about the technology employed at specific plants. A search of company press releases and industry press reports yielded information on the timing of introduction of particular line widths at specific fabs in several cases, but information on the relative importance of each technology *by fab* is not available. To fill this gap we appeal to information on industry norms gathered from iSuppli and other reports and discussions with industry analysts. We assume that several rules hold in general: (1) fabs add new technologies progressively, adding a line width more advanced than all technologies used in previous periods; (2) fabs ramp up output of new technologies linearly over a four-quarter period; (3) companies introduced new technologies first at the company's most advanced fab; (4) when not ramping a new technology, fabs split production evenly among the 2-4 technologies in production. Individual companies often required judgmental deviations from these rules based on information regarding specific fabs to match the overall company technology mix.

Table D.2 shows the specific assumptions we made for each foundry in generating the weights. We have made available the resulting quantities by country, wafer size, line width, and quarter at http://www.andrew.cmu.edu/user/bkovak/bkm_semi_con_data/index.html

Table D.2: Technology Assumptions by Foundry

| Company | # of Fabs | # of Countries | # of Wafer Diameters | Notes on Assumptions |
|----------------------------------|-----------|----------------|----------------------|--|
| Altis Semiconductor | 1 | 1 | 1 | |
| ASMC | 3 | 1 | 2 | Small diameter fab produces legacy technology. Large diameter fab produces remaining, more advanced technology. |
| CRMC | 4 | 1 | 2 | All fabs are same diameter until 2009. No assumptions needed. 2009-2011: small diameter fabs produce all reported legacy technology and amount of 350nm indicated by historical pattern. Remaining shipments from large diameter fabs. |
| Dongbu Hi Tek | 2 | 1 | 1 | |
| Episil | 4 | 1 | 1 | |
| Globalfoundries/ Chartered | 11 | 2 | 3 | 2004-2008: Single country Small diameter fab capacity split evenly between 350nm and legacy technology. Large diameter fabs produce all shipments using 130nm and more advanced technology. Medium diameter fabs produce remaining, more advanced technology. 2009-2011: Begin with 2008 mix from existing locations. Ramp up 65nm and 45nm with timing indicated in reports. Residual is production in single remaining location. |
| Grace | 2 | 1 | 1 | |
| He Jian | 2 | 1 | 1 | |
| HHNEC | 3 | 1 | 1 | |
| HuaLi | 1 | 1 | 1 | |
| Landshunt Silicon Foundry GmbH | 2 | 1 | 1 | |
| Phenitec | 3 | 1 | 1 | |
| Silterra | 3 | 1 | 1 | |
| SMIC | 10 | 1 | 2 | All 150nm and less advanced technology produced at medium-diameter fabs. All 65nm and 90nm is produced at large-diameter fabs. Assign 130nm production based on guidance from industry analysts. |
| SSMC | 1 | 1 | 1 | |
| TowerJazz | 5 | 3 | 2 | Small diameter fab produces 350nm and legacy technology. Remainder split among other locations proportional to capacity. |
| TSMC | 13 | 4 | 3 | Production at one location known with certainty. Second set of fabs known to be divided between 250nm and 350nm, assumed to be split evenly. Third set of fabs assumed to be split evenly among 5 technologies (130nm, 150nm, 250nm, 350nm, 500nm) Small diameter fab accounts for nearly all legacy technology and small share of 350nm & 250nm. Small amount of legacy production in second location indicated by data. Large diameter fabs account for all production using 90nm and more advanced technologies. Residual capacity at these fabs split evenly between 130nm, 150nm, and 180nm until drawing down to minimal to offset ramp-up of 65nm technology. Residual goes to medium diameter production at remaining location. |
| UMC | 12 | 3 | 3 | Production at one location known with certainty. Small diameter fab split evenly between 350nm and legacy technology. Remaining 350nm and legacy technology and all 150nm, 180nm, and 250nm production from medium-diameter fabs. Large diameter fabs account for all production using 65nm and more advanced technologies. assumptions required to split 90nm-130nm between 200mm and 300mm. 2008-2011: Remaining medium-diameter capacity split evenly between 90nm and 130nm. Residual production using these technologies at large-diameter fabs. |
| Vanguard | 2 | 1 | 1 | |
| X-Fab Semiconductor Foundries AG | 6 | 4 | 2 | Split legacy technology production between small-diameter fabs in two locations proportional to capacity. 350nm at known location, medium diameter. Residual is implied at third location. |

Notes: All companies require an estimate of the share of production using CMOS process.

D.3 Weight Construction

While the transactions in the GSA database are representative of the global industry, the limited sample size leads to volatile quantity information in each quarter \times country \times technology cell, when calculated using the GSA transaction sizes alone. To avoid this sampling volatility, we supplement the GSA quantity information with data from IHS iSuppli, described in the previous section. Because this data source covers all shipments of all global foundries, it yields quantities in each cell that are much more stable over time than those from the GSA sample. We construct quantity weights for each transaction by first applying the IHS iSuppli shipments to each quarter \times country \times technology cell and then apportioning that quantity to each transaction in the cell proportionally to the transaction quantity reported in the GSA data. Thus, the IHS iSuppli census determines the relative weight across cells, while the GSA data determine the relative weight across transactions within each cell. Appendix E.1 shows the main results with alternative weighting assumptions, yielding similar results.

D.4 Entry Timing

The time of entry for the last supplier among Taiwan, China, and Singapore was assessed using data from the IHS iSuppli Pure Play Foundry Market Tracker. The iSuppli data are provided three ways, including annual (quarterly) company wafer shipments by feature size beginning in 2000 (2003), annual (quarterly) capacity by fab beginning in 2000 (2002), and annual (quarterly) company capacity by line width beginning in 2000 (2002). Because each fab has a unique wafer size, and the only important Chinese foundry at the time (Semiconductor Manufacturing International Corporation, SMIC) employed 200 mm wafers exclusively in the early 2000s, combining these data unambiguously determines the year of Chinese for each technology introduced from 2001 to 2003. Identifying the quarter of introduction within 2002 required an assumption that SMIC began producing technologies in decreasing order of line width during the year. The oldest technologies were omitted from the entry analysis, as these were present years before Chinese firms entered the foundry market. In all technologies except one, China was the last country to enter the market. In the case of 300mm 130nm wafers, Singapore entered last, coinciding with the opening of Chartered Semiconductor’s first 300mm fab. The resulting entry timing is shown in Table D.3.

China’s delayed entry into the market for each technology results from a number of factors. First, Chinese foundries have less expertise introducing new technologies to the foundry market. New technologies are first implemented by vertically integrated processor and memory manufacturers. Later, these technologies are adapted to the contract manufacturing industry, generally by leading-edge producers in Taiwan. 8 to 12 quarters later, trailing-edge producers such as those in China adopt the technology, and by this point it is a few generations behind the leading-edge technologies in the processor and memory markets.

There are also explicit policies restricting the transfer of leading-edge technologies to producers in China. The *Wassenaar Arrangement on Export Controls for Conventional Arms and Dual-Use Goods and Technologies*, is a multilateral export control agreement that allows participating countries to restrict trade in so-called “dual-use” technologies, which can be used in both conventional goods and in weapons systems. Integrated circuits fall under Category 3 of the dual-use list covered by the arrangement, allowing participating countries to restrict the sale of the latest semiconductor manufacturing equipment. Based on 2006 Gartner reports, essentially all such equipment is produced in Europe, Japan, or the U.S, all of which are Wassenaar Arrangement participants. While it is plausible that these export restrictions slow the diffusion of new technologies to Chinese producers, the fact that individual countries have discretion over what products to restrict may limit the

Arrangement's practical effect on the transfer of semiconductor fabrication technology Brown and Linden (2006). The Taiwanese government imposes strictly binding export control restrictions for a variety of semiconductor related technologies. For example, TSMC has consistently faced limits on the technologies it may produce at their fab in Shanghai LaPedus (2010). Many current and former managers at foundries and fabless firms with whom we spoke cited these Taiwanese government restrictions as the most important barriers to adoption of the most advanced semiconductor fabrication technologies.

Endogenous technology adoption on the part of Chinese foundries could pose a problem for our measurement technique if it precedes relative quality improvement between China and Taiwan within an individual technology. This could occur if Chinese foundries improve their relative quality over time, enter the market once relative quality exceeds some threshold, and continue to improve their quality relative to Taiwan. There are two reasons to discount this possibility. First, the technology transfer restrictions just described limit the ability of Chinese foundries to adopt the newest technologies. Second, as discussed in detail in Appendix F, the primary potential source of quality differences across foundries is the yield, the share of chips on a wafer that function correctly, and there is minimal scope for or evidence of relative yield improvement between Chinese and Taiwanese foundries once volume production begins. Endogenous technology adoption related to costs rather than quality will have no impact on our measurement technique, as costs do not appear in the target exact index, (2), or our proposed index, (8).

Table D.3: Last Supplier Entry Timing

| Wafer Diameter | Line Width | Entry Period | Last Entrant |
|----------------|------------|--------------|--------------|
| 150 mm | 500 nm | . | |
| 150 mm | 350 nm | . | |
| 200 mm | 500 nm | . | |
| 200 mm | 350 nm | 2002 Q1 | China |
| 200 mm | 250 nm | 2002 Q2 | China |
| 200 mm | 180 nm | 2002 Q3 | China |
| 200 mm | 150 nm | 2003 Q4 | China |
| 200 mm | 130 nm | 2005 Q1 | China |
| 300 mm | 130 nm | 2005 Q2 | Singapore |
| 300 mm | 90 nm | 2007 Q3 | China |
| 300 mm | 65 nm | 2009 Q1 | China |

Quarter in which the last supplier among Taiwan, China, and Singapore started production for each process technology. Authors' calculations based on IHS iSuppli data.

E Supplementary Empirical Results

E.1 Alternate Hedonic Regression Results

Table E.1 presents a version of Table 3 with an alternative weighting scheme that weights all transactions equally. The pattern of coefficients across countries is quite similar to the main results, but the magnitudes are quite a bit larger, particularly in column (1), which does not include controls for product attributes. When technology controls are included, the magnitudes are closer to those in Table 3 and maintain the same ranking across countries. Since the magnitudes are still larger in columns (2) - (4) of Table E.1, the main results are conservative.

Table E.2 weights the transactions using only quantity information from the GSA data (omitting the IHS iSuppli information). Again, the unconditional price gaps in column (1) are much larger than those in the main results, but once we include technology controls in columns (2) - (4), the results are very similar to those in Table 3.

Table E.3 presents heteroskedasticity robust standard errors, without clustering. We present this alternative to confirm that the main results in Table 3, which are clustered by quarter, yield larger standard errors in the vast majority of cases and hence are conservative.

Table E.1: Hedonic Wafer Price Regressions - Equal Weights

| Variable | <i>dependent variable: log of price per wafer</i> | | | |
|------------------------------------|--|--|--|--|
| | (1) No Attribute Controls Coefficient Std. Err. | (2) Linear Attribute Controls Coefficient Std. Err. | (3) Flexible Attribute Controls Coefficient Std. Err. | (4) China and Taiwan Only Coefficient Std. Err. |
| Foundry Location | | | | |
| China | -0.620*** (0.030) | -0.252*** (0.019) | -0.249*** (0.019) | -0.247*** (0.018) |
| Malaysia | -0.748*** (0.059) | -0.291*** (0.026) | -0.297*** (0.026) | |
| Singapore | -0.345*** (0.017) | -0.069*** (0.010) | -0.080*** (0.010) | |
| United States | -0.283*** (0.039) | 0.203*** (0.050) | 0.203*** (0.051) | |
| Wafer Size | | | | |
| 150 mm | | -0.434*** (0.027) | | |
| 300 mm | | 0.671*** (0.011) | | |
| Line Width | | | | |
| ≥ 500 nm | | -0.228*** (0.052) | | |
| 350 nm | | -0.146*** (0.029) | | |
| 250 nm | | -0.048*** (0.016) | | |
| 150 nm | | 0.120*** (0.013) | | |
| 130 nm | | 0.304*** (0.014) | | |
| 90 nm | | 0.428*** (0.019) | | |
| 65 nm | | 0.637*** (0.015) | | |
| 45 nm | | 0.827*** (0.022) | | |
| Wafer Size x Line Width Indicators | | | | |
| Number of Metal Layers | | 0.080*** (0.006) | X | X |
| Number of Polysilicon Layers | | 0.053*** (0.013) | 0.082*** (0.006) | 0.081*** (0.006) |
| Number of Mask Layers | | 0.009*** (0.001) | 0.058*** (0.014) | 0.048*** (0.010) |
| Epitaxial Layer Indicator | | 0.114*** (0.011) | 0.008*** (0.001) | 0.010*** (0.001) |
| log Number of Wafers Contracted | | -0.048*** (0.002) | 0.117*** (0.011) | 0.104*** (0.008) |
| Quarter Indicators | X | X | -0.047*** (0.002) | -0.042*** (0.002) |
| R-squared | 0.071 | 0.896 | X | X |
| Observations | 6253 | 6253 | 0.900 | 0.916 |
| | | | 6253 | 5378 |

Observations represent individual semiconductor wafer transactions from GSA data. The omitted category is a 200mm 180nm wafer produced in Taiwan. Transactions equally weighted. Standard errors clustered by 28 quarter clusters, * significant at 10, ** 5, and *** 1 percent level. See Table 3 for the main results.

Table E.2: Hedonic Wafer Price Regressions - GSA Weights

| Variable | (1) | | (2) | | (3) | | (4) | |
|------------------------------------|--------------------------------------|-----------|--|-----------|--|-----------|--------------------------------------|-----------|
| | No Attribute Controls Coefficient | Std. Err. | Linear Attribute Controls Coefficient | Std. Err. | Flexible Attribute Controls Coefficient | Std. Err. | China and Taiwan Only Coefficient | Std. Err. |
| Foundry Location | | | | | | | | |
| China | -0.609*** | (0.064) | -0.164*** | (0.014) | -0.169*** | (0.015) | -0.175*** | (0.016) |
| Malaysia | -0.848*** | (0.035) | -0.274*** | (0.020) | -0.279*** | (0.021) | | |
| Singapore | -0.065 | (0.056) | -0.085*** | (0.016) | -0.087*** | (0.016) | | |
| United States | -0.586*** | (0.063) | 0.129*** | (0.031) | 0.125*** | (0.031) | | |
| Wafer Size | | | | | | | | |
| 150 mm | | | -0.494*** | (0.073) | | | | |
| 300 mm | | | 0.678*** | (0.029) | | | | |
| Line Width | | | | | | | | |
| ≥ 500 nm | | | -0.434*** | (0.084) | | | | |
| 350 nm | | | -0.228*** | (0.030) | | | | |
| 250 nm | | | -0.133*** | (0.026) | | | | |
| 150 nm | | | 0.0882** | (0.035) | | | | |
| 130 nm | | | 0.367*** | (0.025) | | | | |
| 90 nm | | | 0.575*** | (0.043) | | | | |
| 65 nm | | | 0.703*** | (0.040) | | | | |
| 45 nm | | | 1.026*** | (0.065) | | | | |
| Wafer Size x Line Width Indicators | | | | | | | | |
| Number of Metal Layers | | | 0.064*** | (0.007) | X | 0.063*** | (0.007) | X |
| Number of Polysilicon Layers | | | -0.003 | (0.033) | | -0.006 | (0.034) | 0.064*** |
| Number of Mask Layers | | | 0.005* | (0.003) | | 0.005 | (0.003) | 0.001 |
| Epitaxial Layer Indicator | | | -0.014 | (0.020) | | -0.018 | (0.020) | 0.005* |
| log Number of Wafers Contracted | | | -0.048*** | (0.004) | | -0.048*** | (0.004) | -0.018 |
| Quarter Indicators | X | | X | | X | | | -0.049*** |
| R-squared | 0.129 | | 0.940 | | 0.942 | | 0.947 | |
| Observations | 6253 | | 6253 | | 6253 | | 5378 | |

Observations represent individual semiconductor wafer transactions from GSA data. The omitted category is a 200mm 180nm wafer produced in Taiwan. Transactions weighted using within-technology order size information from the GSA transaction data. Standard errors clustered by 28 quarter clusters, * significant at 10, ** 5, and *** 1 percent level. See Table 3 for the main results.

Table E.3: Hedonic Wafer Price Regressions - Robust (not Clustered) Standard Errors

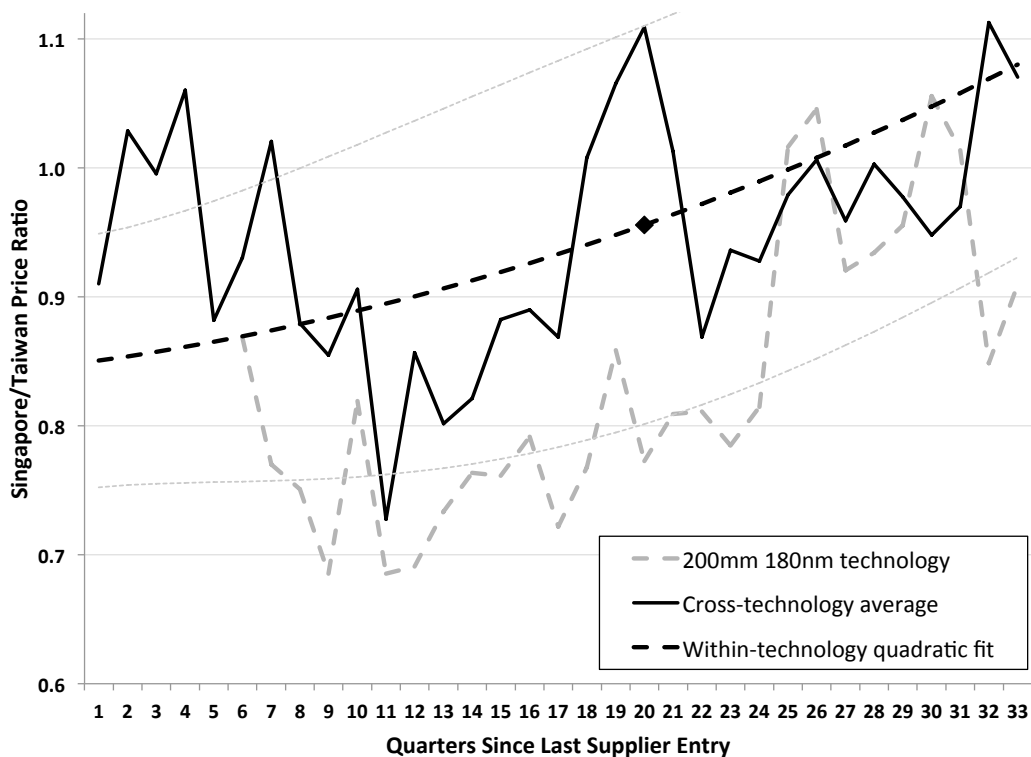
| Variable | <i>dependent variable: log of price per wafer</i> | | | |
|------------------------------------|---|---------------------------|-----------------------------|-----------------------|
| | (1) | (2) | (3) | (4) |
| | No Attribute Controls | Linear Attribute Controls | Flexible Attribute Controls | China and Taiwan Only |
| | Coefficient | Coefficient | Coefficient | Coefficient |
| | Std. Err. | Std. Err. | Std. Err. | Std. Err. |
| Foundry Location | | | | |
| China | -0.297 | -0.186 | -0.188 | -0.196 |
| Malaysia | -0.274 | -0.278 | -0.286 | |
| Singapore | -0.046 | -0.061 | -0.068 | |
| United States | -0.093 | 0.068 | 0.064 | |
| Wafer Size | | | | |
| 150 mm | | -0.467 | | |
| 300 mm | | 0.671 | | |
| Line Width | | | | |
| ≥ 500 nm | | -0.245 | | |
| 350 nm | | -0.167 | | |
| 250 nm | | -0.061 | | |
| 150 nm | | 0.169 | | |
| 130 nm | | 0.356 | | |
| 90 nm | | 0.479 | | |
| 65 nm | | 0.676 | | |
| 45 nm | | 0.962 | | |
| Wafer Size x Line Width Indicators | | | | |
| Number of Metal Layers | | 0.076 | 0.076 | 0.081 |
| Number of Polysilicon Layers | | 0.027 | 0.028 | 0.025 |
| Number of Mask Layers | | 0.005 | 0.005 | 0.004 |
| Epitaxial Layer Indicator | | 0.064 | 0.067 | 0.066 |
| log Number of Wafers Contracted | | -0.056 | -0.057 | -0.058 |
| Quarter Indicators | X | X | X | X |
| R-squared | 0.046 | 0.909 | 0.913 | 0.922 |
| Observations | 6253 | 6253 | 6253 | 5378 |

Observations represent individual semiconductor wafer transactions from GSA data. The omitted category is a 200mm 180nm 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. Heteroskedasticity robust standard errors, * significant at 10, ** 5, and *** 1 percent level. See Table 3 for the main results.

E.2 Singapore vs. Taiwan Price Dynamics

Although Singapore constitutes a much smaller share of the overall market than Taiwan and China (see Table 1), we include it in our index calculations to demonstrate how to construct the quality-adjusted unit value index with multiple suppliers. Figure E.1 shows the price gap analysis for Singapore / Taiwan price ratios, paralleling the results for China in Figure 2. As with China, Singapore exhibits prices below those in Taiwan early in the life of each technology, and a convergence in the relative prices over time, suggesting long run price differences more closely reflect quality differences than price differences shortly after the last supplier enters the market.

Figure E.1: Closing Singapore / Taiwan Price Gap



The x-axis measures the number of quarters since the last supplier began producing the relevant technology. The y-axis measures the Singapore/Taiwan price ratio. The gray dashed line shows the raw quarterly price ratio data for the 200mm 180nm 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 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 light dotted lines show a 90 percent 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 Singapore and Taiwan for the purposes of quality adjustment.

E.3 Within-Technology Price Variation Explained by Supplier

As frictions dissipate following the last supplier’s entry, we expect relative prices across suppliers to converge. Figure 2 shows converging prices for China and Taiwan, and the previous section showed the same for Singapore and Taiwan. Here, we present an alternative analysis that shows converging prices across all suppliers, as opposed to individual pairs of suppliers.⁴²

Frictions drive price variation across suppliers for otherwise identical goods. As the effects of frictions dissipate, we should see that supplier identity accounts for a declining share of within-technology price variation as time elapses following the last supplier’s entry. We first partial out variation related to observable characteristics by regressing log prices and supplier indicators on technology indicators (wafer size line width), layer and contract size controls, and calendar quarter indicators (results are very similar when omitting calendar quarter indicators). The residuals from these regressions retain variation that cannot be explained by observable characteristics or calendar time. We then regress these within-technology log price residuals on supplier residuals. The R^2 in this regression reflects the share of within-technology price variation explained by supplier information. In Table E.4, we estimate this regression separately for observations in each year following the last supplier’s entry and report the R^2 values. As expected, they decline monotonically over time, demonstrating that supplier country accounts for less and less of the within-technology price variation as time elapses following the last supplier’s entry.

Table E.4: Share of Within-Technology Price Variation Explained by Supplier

| Quarters since Chinese entry: | 1-4 | 5-8 | 9-12 | 13-16 | 17-20 |
|--|-------|-------|-------|-------|-------|
| Share of within-technology price variation explained by supplier (R^2) | 0.657 | 0.506 | 0.227 | 0.197 | 0.170 |

Data and sample identical to that in Table 3. R^2 values report the share of within-technology price variation accounted for by supplier country in each year following the entry of the last supplier into the market for the transaction’s technology. The dependent variable is the log wafer price residual after controlling for wafer size, line width, layer, contract size, and calendar quarter controls, as in Table 3. Independent variables are country indicator residuals for China, Malaysia, Singapore, and the USA (Taiwan is the omitted category), using the same controls.

⁴²Thanks to an anonymous referee for suggesting this analysis.

F Yields

F.1 Incidence of Yield Risk

It is common for contract manufacturing service providers to bear the contractual risk of defective products.⁴³ If a fraction of the products produced are defective, the supplier must simply produce more to fill an order for a certain number of functional products. In this case, changes in yields have no effect on the service quality perceived by buyers. Instead, they just raise the supplier’s cost of fulfilling a given order. However, if yields fall low enough, the producer may not have sufficient capacity to fill the buyer’s order in a timely fashion. If low yields lead to missed delivery dates or unfilled orders, then quality would still be affected irrespective of the contractual arrangement.

In the contract semiconductor manufacturing industry, buyers and producers explicitly share the risk if yields fall below expectations. Although foundry contracts are often closely guarded, we have examined a number of contracts appearing as exhibits in semiconductor companies’ annual 10-K filings with the U.S. Securities and Exchange Commission. Most define a minimum acceptable yield, above which the buyer agrees to purchase the wafer and below which the supplier must replace the wafer free of charge. As an example, a contract between communication device maker Integrated Device Technology, Inc. and the foundry Taiwan Semiconductor Manufacturing Corporation states that “If the Minimum Yield is not set forth in the applicable Specifications, the Minimum Yield shall be set to sixty-five percent (65%) of the average Yield of the first three hundred (300) completed Wafers for such Product” Integrated Device Technology, Inc. (2012).

Given that buyers explicitly face the yield risk above the minimum yield threshold and may also face risk below that threshold if very low yields lead to supply disruptions, low yields are a plausible source of quality variation across suppliers. However, the following sections show that in practice there is little scope for yield variation across suppliers.


F.2 Scope for Yield Variation

When producing interchangeable commodity products such as memory or general-purpose processors, yields are a primary concern. Producers in these markets use only the newest technologies (largest wafer sizes and smallest line widths) in an effort to produce more powerful processors or higher capacity memory chips at lower cost than their competitors. Higher yields lower the cost per functional chip and allow a supplier to more effectively compete. In contrast, foundries use older technologies in which yield rates are already relatively high. Foundries mainly produce custom products according to the buyer’s specifications rather than commodity products in which small cost differences can determine success or failure.

Because foundries use relatively mature technologies, there is less scope for yield improvements in their production processes than for firms using the leading-edge production technology. Figure F.1 replicates Figure 1 from Bordelon and Maniar (2006), showing a hypothetical yield curve broken into three stages of production. First, a new production technology is introduced with low yields. During this “process integration” stage, the foundry works to identify and address the sources of yield problems. During the “pilot production” stage, the foundry implements potential customers’ designs with the intention of improving yields to a point where the production process becomes economically feasible. Finally, once yields are quite high, “volume production” begins. Breaux and Collins’s (2007) chapter in the *Handbook of Semiconductor Manufacturing Technology* provides a very similar treatment of yield evolution (see their Figure 27.4).

⁴³Thanks to Teresa Fort for helpful discussions on this topic.

Figure F.1: Yield Evolution



Replicates Figure 1 from Bordelon and Maniar (2006), showing that the vast majority of yield increases occur prior to volume production.

Although Figure F.1 is stylized, it captures the fact that the vast majority of yield improvements occur before the foundry begins volume production of its customers' designs, and remaining yield improvements are minimal once volume production has begun. In our analysis, we restrict attention to this volume production stage because this is the phase in which price measurement agencies would observe transactions between foundries and their customers. Doing so means that we should expect little variation in yields over time for each technology, because yields are already close to their long-run level by the time volume production begins at a given foundry.

The empirical analysis in the following section confirms this expectation as did our many conversations with managers at foundries and fables firms. We had extensive conversations on this issue with a design manager at LSI Corp., a major U.S. fables firm; the CEO and Managing Director of a large Korean fables firm; and the GSA Supply Chain Performance Working Group, an advisory panel consisting of executives and managers from various fables firms. These foundry customers universally reported receiving similar service from the major Taiwanese and Chinese foundries in terms of both yields and timeliness.

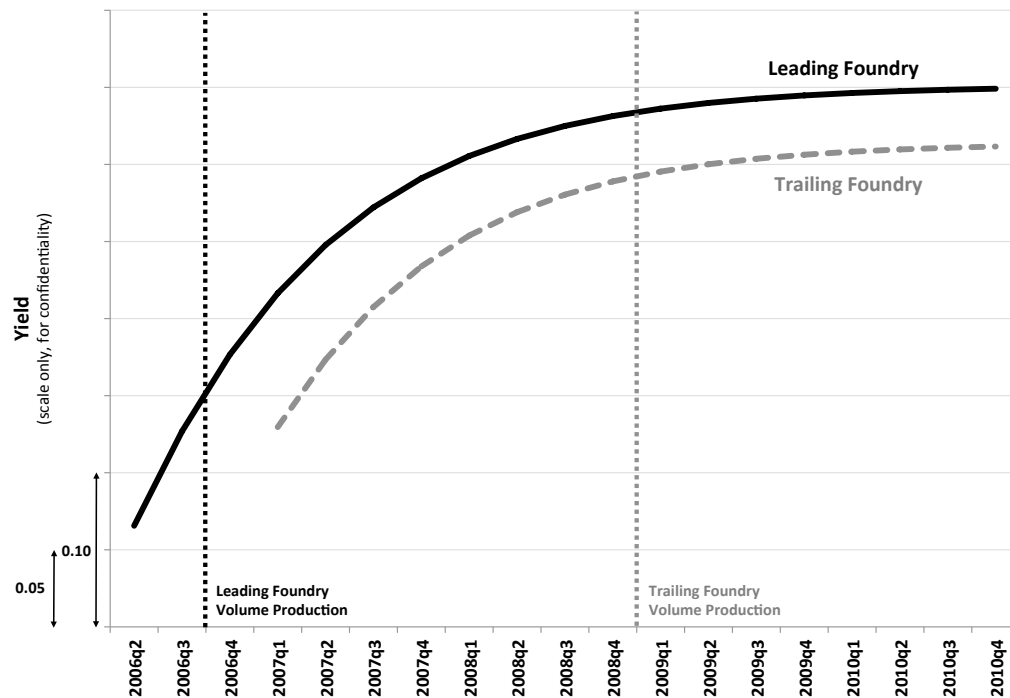
F.3 Yield Data Analysis

The most detailed yield information we were able to obtain comes from PDF Solutions, Inc., which provides yield improvement systems and services to foundries worldwide.⁴⁴ They provided defect density data for two foundries, the best-in-class leading foundry and a trailing foundry, for 65nm technology. To maintain firm confidentiality, we will not identify the firm names. We calculated predicted yields from the defect-density information using the standard Poisson yield model for a 50mm² chip, which is a representative size for 65nm technology. Figure F.2 shows the results. While we are not permitted to report specific yield values on the y-axis, we have labeled its scale, allowing us to measure yield growth over time and to compare yields between the two foundries. The leading foundry began working with this technology at the end of 2004, but did not begin

⁴⁴Thanks to Andrzej Strojwas for providing this information.

volume production runs until Q4 2006. The trailing foundry began pilot production in Q1 2007 and started volume production in Q1 2009.

Figure F.2: 65nm Yields for Leading and Trailing Foundries



Authors' calculations based on defect density information from PDF Solutions, Inc. For confidentiality reasons, we cannot report the location of the Y-axis, but only the scale; the gap between tick marks represents a yield difference of 5 percentage points.

This figure highlights three important points. First, the vast majority of yield improvements for the trailing foundry occur during pilot production, with yields increasing by less than 2 percentage points after beginning volume production. In contrast, the leading foundry does begin volume production at relatively low yields, which seems intuitive; it faces relatively little competition at this point in the product life cycle, so buyers will accept lower yields. But our identification of long-run quality differences is not sensitive to leading-edge yields prior to the entry of new suppliers. Second, although the trailing foundry does not converge to the same yield level as the leading foundry, the yield difference between leading and trailing foundries is quite constant over time, declining by only 0.4 percentage points after the trailing foundry begins volume production. Third, more than 90 percent of this minor yield convergence is realized within 6 quarters. These observations make clear that yield differences and dynamics do not drive the price dynamics seen in Figures 1 and 2. Relative yields barely change following trailing supplier entry, in contrast with clear price convergence. What little yield convergence that does occur happens very quickly; the process of price convergence is far slower, taking about 5 years.

While yield differences are unable to explain the observed price dynamics, yields may explain a portion of the long-run quality difference that we infer in Section 5. In Figure F.2, the two foundries exhibit long-run yield differences of 3.7 percentage points, which could account for about

40 percent of the 8.5 percent quality difference we infer between China and Taiwan, for example.

Expanding our analysis to additional technologies (other than 65nm) involves a tradeoff in terms of the detail of the yield data. From Q4 2004 to Q3 2008 the GSA survey included a question reporting yields for all technologies, but only in the following four ranges: 0-25, 26-50, 51-75, and 76-100 percent. Table F.1 tabulates these yield ranges for each supplier country. No observations report a yield below 76% for a transaction produced by a Chinese supplier, so there is no evidence of far lower yields for Chinese foundries than for Taiwanese foundries. The lack of variation also means that controlling for yield-range indicators in Table 3 would have no effect on measured price dispersion across suppliers.

Table F.1: Yield Ranges

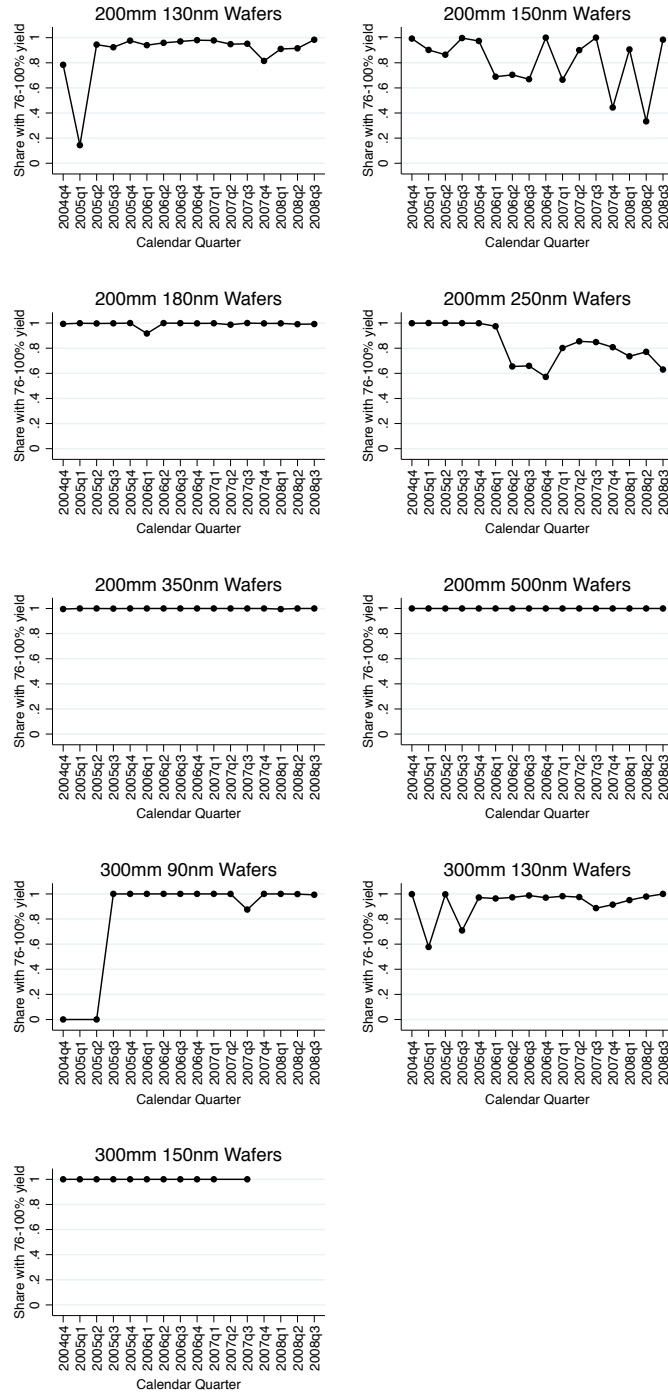
| Yield Range: | 0-25% | 26-50% | 51-75% | 76-100% |
|---------------|-------|--------|--------|---------|
| China | 0.0 | 0.0 | 0.0 | 100.0 |
| Malaysia | 0.0 | 0.0 | 0.0 | 100.0 |
| Singapore | 0.0 | 0.0 | 6.3 | 93.7 |
| Taiwan | 0.0 | 1.8 | 4.6 | 93.7 |
| United States | 0.0 | 0.6 | 8.1 | 91.3 |

Share of volume production transactions from each country reporting yields in the relevant range. Observations represent individual semiconductor wafer transactions from GSA data, weighted as in Table 3. Yield range measure available from Q4 2004 to Q3 2008. 3,208 nonmissing observations.

Because there is variation in the yield range measure for Taiwanese products, we can use it to look for evidence of yield improvements over time within technology. Figure F.3 shows the share of Taiwanese wafers reporting 76-100 percent yields in each quarter for each technology with nontrivial Taiwanese output during our sample period.⁴⁵ Note that zero for this measure does not imply zero yield, but simply that none of the wafers fall in the highest yield range in Table F.1. If Taiwanese producers exhibit substantial yield improvements over time, we expect to see systematically upward sloping lines for each technology, at least in earlier time periods. We do not see this pattern. There are a handful of quarters in which a very small share of Taiwanese wafers have high yields, but the profiles are as likely to slope downward as they are to slope upward. Although the yield range measure is rough, these patterns are consistent with the expectation that yields are relatively stable over time once volume production is reached.

⁴⁵Thanks to an anonymous referee for suggesting this analysis.

Figure F.3: Share of Taiwanese Wafers with 76-100% Yield



Authors' calculations based on GSA data. Each dot records the share of wafers produced in Taiwan that report yields in the 76-100% range in the quarter listed on the x-axis.

At our request, GSA added a continuous yield question to the Q3 2011 survey, in which respondents reported the yield as a share between 0% and 100%. Because this question is available only in one quarter, it does not allow us to track technologies over time. There was also a very high (60%) rate of non-response, which sharply reduces the power of any tests of yield differences across suppliers. That said, the available information supports the industry perception that Chinese yields are not substantially lower than those in Taiwan. Column (1) of Table F.2 examines yield differences across suppliers using a specification similar to that in Table 3. Although the supplier location coefficients are imprecisely estimated, there is no evidence for lower yields at Chinese producers (note that the Malaysia coefficient cannot be estimated due to a lack of valid yield data). Moreover, column (2) shows that, if anything, observations corresponding to Chinese transactions were *more* likely to have a valid value for the continuous yield measure, which is inconsistent with the notion that customers were simply less likely to report their lower yields from Chinese foundries.

While the power of these analyses using GSA data is curtailed by a lack of detail in the yield measure or by substantial non-response, the results all point to similar yields across suppliers, consistent with our discussions with industry insiders. As a whole, the empirical evidence strongly supports our assumption of constant relative quality of goods produced by Chinese and Taiwanese foundries.

Table F.2: Continuous Yield Analysis, Q3 2011

| Dependent Variable: | (1) | | (2) | |
|------------------------------------|------------------------------------|-----------|------------------------------------|-----------|
| | <u>yield (in percent, 0...100)</u> | | <u>indicator for missing yield</u> | |
| | Coefficient | Std. Err. | Coefficient | Std. Err. |
| Foundry Location | | | | |
| China | 0.485 | (1.419) | -0.308 | (0.133)** |
| Malaysia | . | . | 0.189 | (0.132) |
| Singapore | -5.954 | (6.471) | 0.141 | (0.109) |
| United States | -0.470 | (1.945) | -0.002 | (0.106) |
| Wafer Size x Line Width Indicators | X | | X | |
| Number of Metal Layers | 0.754 | (0.467) | 0.040 | (0.033) |
| Number of Polysilicon Layers | -0.368 | (1.637) | 0.139 | (0.069)** |
| Number of Mask Layers | -0.246 | (0.122)** | -0.010 | (0.008) |
| Epitaxial Layer Indicator | 0.650 | (1.313) | -0.199 | (0.084)** |
| log Number of Wafers Contracted | 0.352 | (0.350) | 0.004 | (0.017) |
| R-squared | 0.254 | | 0.113 | |
| Observations | 106 | | 265 | |

Column (1) examines cross country differences in yields, as reported in the continuous yield question in Q3 2011, while column (2) examines differences in the probability of reporting a valid (nonzero) yield, using a linear probability model. Observations represent individual semiconductor wafer transactions from GSA data. Sample restricted to Q3 2011 for availability of continuous yield measure. Baseline case (omitted category) is a 200mm 180nm wafer produced in Taiwan. Robust standard errors in parentheses, * significant at 10%, ** 5, *** 1 percent level

G CES Quality Inference

In this section, we implement a standard quality inference procedure following the international trade literature (see footnote 4 for references). We assume that buyers maximize a nested CES objective function over technologies k and suppliers j in each quarter t .

$$U_t = \prod_k u_{kt}^{b_k} \quad \text{where} \quad \sum_k b_k = 1, b_k \geq 0 \quad \forall k, \text{ and} \quad (\text{G1})$$

$$u_{kt} = \left[\sum_j (\xi_{jkt} y_{jkt})^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \quad (\text{G2})$$

where y is the quantity purchased and ξ is the variety's quality. With this objective function, demand is given by

$$\ln y_{jkt} + \sigma \ln p_{jkt} = \ln(E_{kt} P_{kt}^{\sigma-1}) + \ln(\xi_{jkt}^{\sigma-1}), \quad (\text{G3})$$

where p_{jkt} is the price, E_{kt} is expenditure, and $P_{kt} = (\sum_k \sum_j p_{jkt}^{1-\sigma} \xi_{jkt}^{\sigma-1})^{\frac{1}{1-\sigma}}$ is the CES price index for technology k in year t .

We follow the estimation strategy in Khandelwal et al. (2013). Given an estimate of σ , the left hand side of (G3) is observable. The first term on the right hand side can be captured by a technology \times year fixed effect, α_{kt} , so we estimate the following specification

$$\ln y_{jkt} + \sigma \ln p_{jkt} = \alpha_{kt} + \varepsilon_{jkt}. \quad (\text{G4})$$

Each product's log quality can be inferred from the estimated residual as

$$\ln \hat{\xi}_{jkt} = \frac{\hat{\varepsilon}_{jkt}}{\sigma - 1}. \quad (\text{G5})$$

We then calculate quality relative to Taiwan ($j = 1$) for each technology as $\exp(\ln \hat{\xi}_{jkt} - \ln \hat{\xi}_{1kt})$ and average across technologies and quarters to yield a summary measure of relative quality, for comparison to the estimates in Section 5 (results are very similar using the median rather than the mean).

To implement this process, we need an estimate of σ . We follow Khandelwal et al. (2013) in drawing elasticity estimates from Broda et al. (2006). We use estimates for the semiconductor industry, HS 854, and consider a very wide range of values including the 5th percentile (1.34), median (3.26), 95th percentile (25.03), and the maximum value (125.24). The results appear in columns (1) and (2) of Table G.1. The estimate 0.446 in column (1), corresponding to the median substitution elasticity, implies that Chinese wafers provide 55.4 percent fewer effective units than otherwise equivalent Taiwanese wafers. Even the highest estimate in column (1) suggests that China's quality is 18.5 percent lower than Taiwan's, which is more than twice the size of the quality gap we infer using long-run prices (8.5 percent) in Section 5. Thus, irrespective of the substitution elasticity, all estimates in the table are far below the relative quality estimates we find when using long-run price differences.

Why do the two approaches yield such different answers? The main text assumes that identical technologies produced in different countries are perfect substitutes. As the elasticity of substitution becomes arbitrarily large in the CES context, the model infers that price differences can only persist if they exactly reflect quality differences. Hence, for the highest value of sigma, the model infers that Chinese quality is 18.5 percent lower than Taiwanese quality, which is comparable to the

estimated China fixed effect in the hedonic regressions (Table 3, column (2)). This is about twice as large as our estimate of the quality gap in Section 5. Our framework allows for the possibility that a substantial share of China’s discount reflects its attempt to attract market share in the face of large market frictions. In contrast, this price discount pulls down the CES-based estimate of relative Chinese quality.

Now consider the influence of lowering the elasticity of substitution. This places more weight on quantity (as opposed to price) in the quality inference equation (G4). In this case, a strong taste for variety tends to compress the distribution of market shares, all else equal. Therefore, if a supplier nonetheless captures a relatively high market share, it must have much higher quality. As a result, the model infers from Taiwanese market share that Chinese quality is dramatically lower—just 44.6 percent of Taiwan’s when using the median substitution elasticity. Our framework instead recognizes that China’s share reflects, in part, how market frictions slow the reallocation of market share to new entrants.

This exercise shows how standard quality-inference procedures assuming frictionless product markets systematically understate the relative quality of new entrants in the presence of frictions.

Thus far, we have applied (G4) at the individual technology level, but in practice much research in trade and quality does not have access to such detailed data. All of the products in our data fall within a single 10-digit HS code, 8542.21.80.05. Thus, in standard trade data, the researcher would not be able to distinguish between the various technologies (wafer size and line width) indexed by k above. To see the implications of this lack of product detail, we implement an alternative version of the quality inference procedure that ignores differences in wafer technology and only considers differences in supplying country j . In this case, the demand equation is

$$\ln y_{jt} + \sigma \ln p_{jt} = \ln(E_t P_t^{\sigma-1}) + \ln(\xi_{jt}^{\sigma-1}), \quad (\text{G6})$$

where the CES price index in quarter t is $P_t = (\sum_j p_{jt}^{1-\sigma} \xi_{jt}^{\sigma-1})^{\frac{1}{1-\sigma}}$. We replace the first term on the right side with a time fixed effect, and estimate

$$\ln y_{jt} + \sigma \ln p_{jt} = \alpha_t + \varepsilon_{jt}, \quad (\text{G7})$$

inferring quality as above.

The results appear in columns (3) and (4) of Table G.1. China’s quality relative to Taiwan is inferred to be even lower than before, which is to be expected. By omitting product technology information, we are comparing the mix of Chinese products to the mix of Taiwanese products in a given quarter. In Table 3 we showed that the unconditional relative price for Chinese products (column (1)) is much lower than the relative price controlling for product mix (column (2)), because China produces more trailing-edge products than Taiwan. By omitting technology information, it becomes part of the unobserved quality attributes reflected in the estimates in Table G.1, amplifying the inferred differences between the trailing-edge entrants and the more advanced incumbent producers in Taiwan.

Table G.1: Quality Relative to Taiwan, Inferred From CES Demand Without Frictions

| Product definition: | (1) country × technology | (2) | (3) | (4) country |
|-----------------------------------|-----------------------------|------------------|--------------|------------------|
| <u>Elasticity of substitution</u> | <u>China</u> | <u>Singapore</u> | <u>China</u> | <u>Singapore</u> |
| 1.34 (5th percentile) | 0.019 | 0.011 | 0.001 | 0.002 |
| 3.26 (median) | 0.446 | 0.454 | 0.228 | 0.351 |
| 25.03 (95th percentile) | 0.769 | 0.864 | 0.564 | 0.777 |
| 125.24 (max) | 0.815 | 0.919 | 0.604 | 0.831 |

Product quality estimates for China and Singapore relative to Taiwan, inferred from standard CES demand, following Khandelwal et al. (2013). The numbers in the table correspond to ξ_j/ξ_{Taiwan} where $j \in \{\text{China, Singapore}\}$; 0.446 in column (1) implies that Chinese products provide 55.4 percent fewer effective units than otherwise equivalent Taiwanese products. Elasticities of substitution come from Broda et al. (2006) for HS 854, which contains integrated circuits. Results are shown using the 5th percentile, median, 95th percentile, and maximum value estimated by Broda et al. (2006) for this industry. Columns (1) and (2) define varieties based on supplier country and technology (wafer size and line width), while columns (3) and (4) define varieties based on country only. Since all technologies fall within a single 10-digit HS code, columns (3) and (4) correspond more closely to the level of detail available in standard trade datasets.

References

- Aizcorbe, A. (2002). Price Measures for Semiconductor Devices. *Board of Governors of the Federal Reserve System (U.S.), Finance and Economics Discussion Series*, 2002-13.
- Baldwin, R. E. and Krugman, P. R. (1988). Market access and international competition: A simulation study of random access memories. In Feenstra, R. C., editor, *Empirical Methods for International Trade*, pages 171–202. MIT Press.
- Bayard, K., Byrne, D., and Smith, D. (2015). The scope of u.s. factoryless manufacturing. In Houseman, S. N., editor, *Measuring Globalization: Better Trade Statistics for Better Policy*, volume 1. W.E. Upjohn Institute for Employment Research.
- Bernard, A. B. and Fort, T. C. (2013). Factoryless goods producers in the us. *NBER Working Paper*, (19396).
- Berndt, E. R. and Rappaport, N. J. (2001). Price and Quality of Desktop and Mobile Personal Computers: A Quarter-Century Historical Overview. *American Economic Review*, 91:268–273.
- Bernstein, F., Deng, C., and Pekeč, S. (2013). Single versus multiple sourcing: Drivers of supplier diversification. *mimeo*.
- Bordelon, J. and Maniar, P. (2006). The sub-100-nm imperative: parametric yield ramp. *EE Times*.
- Breaux, L. and Collins, S. (2007). Yield management. In Doering, R. and Nishi, Y., editors, *Handbook of Semiconductor Manufacturing Technology*, chapter 27, pages 27-1–27-29. CRC Press, Boca Raton, 2 edition.
- Broda, C., Greenfield, J., and Weinstein, D. (2006). From groundnuts to globalization: A structural estimate of trade and growth. *NBER Working Paper*, (12512).
- Brown, C. and Linden, G. (2006). Offshoring in the Semiconductor Industry: Historical Perspectives. In *Brookings Trade Forum: 2005*. Washington, DC: Brookings Institution Press.
- Bureau of Labor Statistics (1997). *BLS Handbook of Methods, Chapter 15*.
- Byrne, D., Oliner, S. D., and Sichel, D. E. (2013). Is the information technology revolution over? *International Productivity Monitor*, (25):20–36.
- Byrne, D. M. and Corrado, C. A. (2015). Prices for communications equipment: Rewriting the record. *Finance and Economics Discussion Series*, (2015-069).
- Byrne, D. M., Oliner, S. D., and Sichel, D. E. (2015). How fast are semiconductor prices falling? *NBER Working Paper*, (21074).
- Diewert, E. W. (1976). Exact and superlative index numbers. *Journal of Econometrics*, 4(2):115–145.
- Doms, M. (2005). Communication Equipment: What has Happened to Prices? In *Measuring Capital in the New Economy*, volume 65 of *NBER Studies in Income and Wealth*, pages 323–362. University of Chicago Press.

- Doms, M., Aizcorbe, A., and Corrado, C. (2003). When Do Matched-Model and Hedonic Techniques Yield Similar Measures? *Federal Reserve Bank of San Francisco, Working Papers in Applied Economic Theory*, 2003-14.
- Dulberger, E. R. (1993). Sources of Price Decline in Computer Processors: Selected Electronic Components. In *Price Measurements and their Uses*, volume 57 of *National Bureau of Economic Research Studies in Income and Wealth*, pages 103–124. University of Chicago Press.
- Flamm, K. (1993). Measurement of DRAM Prices: Technology and Market Structure. In *Price Measurements and their Uses*, volume 57 of *National Bureau of Economic Research Studies in Income and Wealth*, pages 157–197. University of Chicago Press.
- Flamm, K. (1996). *Mismanaged Trade: Strategic Policy and the Semiconductor Industry*. Brookings.
- Flamm, K. (2007). The Microeconomics of Microprocessor Innovation. *mimeo*.
- Global Foundries (2009). Fact sheet: Fab 2 module 1.
- Grimm, B. T. (1998). Price Indexes for Selected Semiconductors, 1974-96. *Survey of Current Business*, 78 (2):8–24.
- Holdway, M. (2001). Quality-adjusting computer prices in the producer price index: An overview. Accessed January 26, 2011.
- Hurtarte, J. S., Wolsheimer, E. A., and Tafoya, L. M. (2007). *Understanding Fabless IC Technology*.
- IC Knowledge (2001). Can the semiconductor industry afford the cost of new fabs?
- Integrated Device Technology, Inc. (2012). Form 10-k.
- Irwin, D. A. and Klenow, P. J. (1994). Learning-by-doing spillovers in the semiconductor industry. *Journal of Political Economy*, 102(6):1200–1227.
- Khandelwal, A. K., Schott, P. K., and Wei, S.-J. (2013). Trade liberalization and embedded institutional reform: Evidence from chinese exporters. *American Economic Review*, 103(6):2169–2195.
- Kumar, R. (2007). The Business of Scaling. *IEEE Solid-State Circuit Society News*, 12(1):22–27.
- LaPedus, M. (2010). Tsmc gets nod for 130-nm production in china. *EE Times*.
- Nakamura, E. and Steinsson, J. (2012). Lost in Transit: Product Replacement Bias and Pricing to Market. *American Economic Review*, 102(7):3277–3316.
- Scott, A. and Angel, D. (1988). The Global Assembly-Operations of U.S. Semiconductor Firms: A Geographical Analysis. *Environment and Planning A*, 20(8):1047–67.