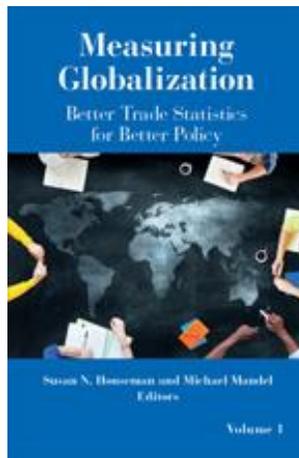

Upjohn Institute Press

Assessing Price Indexes for Markets with Trading Frictions: A Quantitative Illustration

Brian K. Kovak
Carnegie Mellon University

Ryan Michaels
University of Rochester



Chapter 3 (pp. 89-120) in:

**Measuring Globalization: Better Trade Statistics for Better Policy, Volume 1,
Biases to Price, Output, and Productivity Statistics from Trade**

Susan N. Houseman and Michael Mandel, eds.

Kalamazoo, MI: W.E. Upjohn Institute for Employment Research, 2015

DOI: 10.17848/9780880994903.vol1ch3

3

Assessing Price Indexes for Markets with Trading Frictions

A Quantitative Illustration

Brian K. Kovak
*Carnegie Mellon University and
National Bureau of Economic Research*

Ryan Michaels
University of Rochester

In the last half-century, reductions in transportation and communication costs have dramatically reshaped the spatial organization of manufacturing production. It is becoming common, for instance, for an input to be manufactured abroad and then shipped back to the firm that designed it (Hummels, Ishii, and Yi 2001). The physical manufacturing of the good in this case is increasingly concentrated in developing economies such as China, which tend to offer lower prices than incumbent producers.

What is the source of these lower prices? They may represent real discounts on the same physical good. But there is also a possibility the price differentials are spurious. They may reflect, for instance, unobserved differences in the composition of goods. Furthermore, even if the inputs are physically identical, the quality of the production *service* may vary—as judged, for instance, by the timeliness of delivery and the reliability of the finished product. As Carlton (1983) stresses, the service quality factors into the true price of the good to the buyer (and into the real resource cost of the transaction).

The answer to our question is of considerable importance to price index measurement. If price differentials are mistakenly assumed to be spurious, price indexes will be constructed to ignore the true decline that occurs when lower-price suppliers enter an intermediate market.

However, it is equally perilous to neglect the scope for unobserved variation in product and service quality. The challenge to statistical agencies is that, in practice, it is very difficult to isolate real price dispersion given data limitations.

In reaction to this, the present paper attempts to provide some guidance for price measurement. We explore how well imperfect, but *feasible*, price indexes approximate the true price change in markets where quality variation and real dispersion commingle. A price index is feasible if it can be computed from data only on observable outcomes, such as market prices. We apply these feasible indexes to markets characterized by two key features. First, even if physical products are identical, there is scope for variation in service quality that would be unobservable to the analyst. Second, the same product and service can still be priced differently because of a certain trading friction that impedes arbitrage.

We carry out our experiment within a simple duopoly pricing model. The structure of the model is designed to mimic salient features of the market for semiconductor wafers, the subject of our empirical application below. The latter market is an excellent example of the contract manufacturing sector, in which domestic firms design products and offshore all fabrication activities. This sector is expanding at a remarkable rate in the United States (Bayard, Byrne, and Smith 2013).

In the model, two large suppliers—a leader (the founding firm in the market) and a follower (who enters the market last)—produce an input for overlapping generations of buyers. We assume that the physical (observable) dimensions of the input are the same across suppliers. This assumption is relatively safe in our context since, in our empirical application, we have exceptionally detailed data on physical attributes. However, the model allows for variation in the quality of the manufacturing service. We make this notion more precise below; the bottom line is that lower-quality service will raise the effective price of transacting with the *follower*. At the same time, we introduce a trading friction that takes the form of a setup cost that must be paid if a buyer switches suppliers during the life of its product. The setup cost applies regardless of the identity of the supplier to whom the buyer is switching. This friction implies that, when the follower enters the market, the leader's customers may pay its high price even if there is no difference in production service.

In the second section (“A Pricing Game with Costly Switching”), we first solve the model numerically to illustrate its key implications for price dispersion and price dynamics. The presence of the setup cost implies that, when the follower enters, the extent of price dispersion exceeds that which could be attributed to quality variation. However, as the leader’s contracts with its original customers end, it will compete more aggressively for new generations of buyers. This causes price dispersion between firms to narrow. In fact, under certain circumstances, the effect of the setup cost on price dispersion will abate to the point that the price differential at the end of the product life will reflect *only* variation in service quality. This is a key distinguishing property of the model—constant differences in service quality alone do not induce this pattern in the *dynamics* of price dispersion.

Before we apply the model to price index measurement, we first look for evidence consistent with these predictions. To that end, the third section (“An Application to the Semiconductor Industry”) presents results from Byrne, Kovak, and Michaels (2013) that are consistent with these dynamics. The authors have data on transaction-level prices of semiconductor wafers along with the key technological attributes of each wafer. They can therefore control for differences across suppliers in product composition. However, there may still be differences with respect to service quality. Indeed, it is often thought that the leader in this market, Taiwan, has software tools that enable it to provide higher-quality service, for which it presumably charges a higher price than its main competitor, China.¹ The theory in Section Two suggests a way of identifying quality-*adjusted* dispersion in this setting: Byrne, Kovak, and Michaels can test for whether the price difference between these suppliers narrows after China’s entry into a particular wafer market.² They find that, on average, the price differential between Taiwan and China does close substantially over the life of a given semiconductor technology: It falls from 39 percent in the year of Chinese entry to 10 percent after five years. This narrowing is consistent with the presence of real dispersion, although the differential remaining even at the end of the product life is suggestive of quality dispersion.

In light of this evidence, Section Four (“Feasible Price Indexes”) returns to the model and uses it to study the performance of different price indexes when the observed change in average market price reflects both real dispersion and variation in quality. This section first

calculates a benchmark index that assumes the analyst has perfect information regarding quality and is able to directly adjust for the effective cost of transacting with the follower. It then compares the results of this benchmark with feasible price indexes that can be calculated even when the analyst has access only to data on market prices. We consider three examples.

The first feasible index is based on the idea that price dispersion across suppliers derives exclusively from quality variation. In this case, the index can be computed by analogy to a standard superlative index, which treats a supplier's service as a separate "good" and averages price changes across providers. The second feasible index takes the opposite view: All of the observed differential represents a real discount. Accordingly, one can simply average price levels across suppliers and compute the change in the average price across periods. Not surprisingly, this index yields the largest declines in the price level when the lower-priced supplier enters the market. The third index is our preferred index, since it tries to strike a compromise between these two. It relies on a simple implication of the theory, noted above: the effect of the setup cost on price dispersion abates over the course of the product life cycle, leaving a price differential that reflects only variation in service quality. As a result, we can use the *observed* price differentials late in the product life cycle to proxy for unobserved quality variation. This enables a simple correction to market price data while not foreclosing a role for real price dispersion.

Section Four confirms that our preferred index performs best. Yet, as we detail later, the correction here is still somewhat conservative in that it continues to slightly understate the extent of the true price decline that occurs when the follower enters the market. We then illustrate how to adjust our preferred index so that it delivers an *upper* bound on the extent of the price decline. The true price change can then be bracketed.

Section Five offers a conclusion.

A PRICING GAME WITH COSTLY SWITCHING

This section begins by describing an extension of the simple pricing game in Byrne, Kovak, and Michaels (2013). Our modeling is guided

by a large literature that studies price setting in markets under costly switching.³ The model here deviates slightly from this preceding literature, which typically restricted attention to the analytics of games with symmetric players. Reflecting our interest in the quantitative dynamics following the entry of a low-cost supplier, we analyze a calibrated game with asymmetric actors. The leader is the founding firm in the market and enjoys monopoly status for a time. The follower enters the market subsequently and has a lower unit cost of production but inferior production technology. Each firm competes to supply an input to overlapping generations of final-goods producers—the consumers, or buyers, in this market.

The Model

The basic environment

The model is perhaps the minimalist structure needed to consider some of the questions of interest. There are three periods, and there are two types of agents in the market—buyers and manufacturers of an intermediate good. A cohort of buyers enters in each of the three periods. The period-1 cohort is present in periods 1 and 2, the period-2 cohort is present in periods 2 and 3, and the period-3 cohort is present in period 3. Each cohort is of mass 1. Buyers have unit demand, and they purchase from one of the suppliers as long as the price is less than the reservation value, a constraint that we discuss in detail below.

Even though buyers purchase the same physical input from both suppliers, we assume there are details of the production process that have to be tailored to the buyer's order. To preview the example in Section Three, consider the market for semiconductor wafers, where buyers are designers of integrated circuits. Suppliers are firms that fabricate silicon wafers on which the design is implanted. Each buyer purchases a wafer with the same size and density of transistors, but there are details of the design—the precise manner in which transistors are arrayed on the wafer—that require some specialization of the production process. Formally, we follow in the spirit of Klemperer (1995) and assume that design complexity, y , is distributed uniformly from 0 (lowest quality) to 1 (highest quality). This heterogeneity across buyers would be unobservable to an econometrician who has data only on the physical wafer

size and line width. In this sense, the model allows for price dispersion that reflects unobserved heterogeneity.

Turning to the manufacturers, Firm A is the leader and is present in the market from period 1 onward. Firm B is the follower; it joins the market in period 2. We assume Firm A is at the technology frontier. To again borrow an example from the semiconductor wafer market, it is thought that Taiwan's fabrication firms have intellectual property that enables them to more efficiently produce a highly complex design. This means that, although Firm B (China, in the case of the wafer market) can fabricate any chip, the consumer must pay a cost to monitor and consult with this supplier. We assume that buyers who purchase from Firm B pay a per-period monitoring cost, τy (with $\tau > 0$), that is increasing in design complexity. What helps Firm B to compete in the face of this disadvantage is that it enjoys a lower unit cost of production, which we denote by c^B . Specifically, we assume that both firms have constant unit costs, and that $c^A > c^B$.

When a buyer initiates production with a supplier, it must pay a startup cost, s . This cost has to be paid again if the buyer switches suppliers. Thus, if a buyer purchased from Firm A in period $t - 1$ but switches to Firm B in period t , it must pay s again (independent of its quality). Hence, this buyer would pay a price, p_t^A , to remain with Firm A in period t , and would pay $p_t^B + \tau y + s$ to switch to Firm B, where p_t^B is Firm B's posted price in period t , τy is the monitoring cost, and the startup cost s acts as a cost of switching.

There are very clear examples of switching costs in the wafer market. To illustrate, certain equipment has to be supplied by the customer and calibrated to the processes and technologies of each supplier. For instance, the customer supplies the mask, through which its design of transistors is projected onto a wafer. The mask must be specified to sync with the supplier's proprietary technologies, which are generally incompatible across manufacturers. This makes it difficult to re-source a product once wafer production begins. In the case of a mask, the price of a new one is high, at over \$1 million. As a result, notes one industry association, "The time and cost associated with [switching] tend to lock customers into a particular [supplier]" (Gabriel Consulting Group 2006, p. 1).⁴

Last, following much of the literature on costly switching, the model prohibits price discrimination. This restriction is roughly con-

sistent with wafer supplier contracts, which limit a supplier's freedom in charging appreciably different prices across its customers.⁵ Thus, we assume the price p_t^A (p_t^B) applies to all Firm A (B) buyers in period t .

The terminal period problem

The problem is solved by backward induction. To analyze the period-3 problem, we first conjecture that there is a threshold y_2 so that Firm B attracts all period-2 entrants with designs y that satisfy $y \leq y_2$. In other words, we assume the least "advanced" producer attracts buyers with the least complex designs. This conjecture will be confirmed in equilibrium. In what follows, since y is uniformly distributed, we refer to the mass of buyers y_2 as Firm B's customer base at the start of period 3. The mass of higher-quality buyers $1 - y_2$ makes up Firm A's customer base.

There are three groups of buyers to whom Firm A may sell: members of its own customer base, members of Firm B's customer base, and buyers who enter in period 3 (period-3 entrants). The demand schedules for each of these cohorts are given below. Throughout, we let σ_t^{jj} represent the share of Firm j 's customer base that it retains in period t , σ_t^{j0} the share of period- t entrants that it attracts, and σ_t^{ji} the share of Firm i 's customer base acquired by Firm j . Hence, for Firm A, we have

$$(3.1) \quad \begin{aligned} \sigma_3^{AA}(p_3^A; p_3^B, y_2) &= \Pr[p_3^A \leq p_3^B + \tau y + s \mid y > y_2] \\ \sigma_3^{AB}(p_3^A; p_3^B, y_2) &= \Pr[p_3^A + s \leq p_3^B + \tau y \mid y \leq y_2] \\ \sigma_3^{A0}(p_3^A; p_3^B, y_2) &= \Pr[p_3^A \leq p_3^B + \tau y] \end{aligned}$$

where p_t^j denotes the price of Firm j in period t .

Each of the components of Equation (3.1) is straightforward. Firm A retains a member $y > y_2$ of its customer base if its price, p_3^A , is less than the quality-adjusted price of its rival, $p_3^B + \tau y$, plus the cost of switching production to a new supplier, s . It poaches a buyer $y \leq y_2$ in Firm B's customer base if its price, plus the cost of switching, is less than $p_3^B + \tau y$. Last, Firm A attracts a new (period-3) entrant if its price is less than the quality-adjusted price of Firm B. Observe that s does not appear in the entrant's decision, since it must pay the cost of setting up regardless of the supplier from which it sources.

Absent from Equation (3.1) is any mention of the buyer's (gross) payoff from the sale of its final good. This is because the gross payoff is independent of the identity of the supplier. Thus, conditional on participation in the market, the buyer's choice of supplier depends only on the relative (quality-adjusted) prices and setup costs. We only assume at this stage that the gross payoff exceeds the minimum cost to the final-goods maker. Later, we will specify the payoff and calibrate it so that the participation constraint does not in fact bind in periods 2 and 3.

Firm A's terminal-period problem may now be stated as follows. From Equation (3.1), we have that total sales by Firm A in period 3 are given by

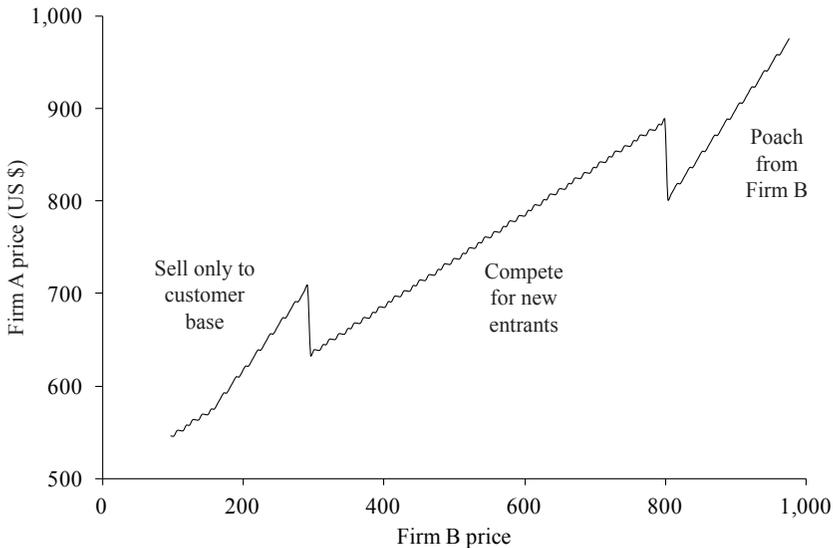
(3.2)

$$Y_3^A = \sigma_3^{AA}(p_3^A, p_3^B, y_2)(1 - y_2) + \sigma_3^{A0}(p_3^A; p_3^B, y_2) + \sigma_3^{AB}(p_3^A; p_3^B, y_2)y_2.$$

The leader then sets its price to maximize profits, $(p_3^A - c^A)Y_3^A$, which yields an optimal price of the form $p_3^A(p_3^B, y_2)$. Firm B faces the analogous problem, the solution of which is represented by $p_3^B(p_3^A, y_2)$. The intersection of the two best responses yields the terminal-period equilibrium, conditional on y_2 . We denote the equilibrium prices by $P_3^A(y_2)$ and $P_3^B(y_2)$.

The (pure-strategy) pricing policy of a firm can typically be partitioned into three regions. Consider, for instance, the behavior of Firm A, whose optimal price is shown in Figure 3.1 as a function of Firm B's price. Over a range of low Firm B prices, Firm A will concede all new entrants to its rival. The reason for doing so is that it can earn greater profits by setting a higher price and selling exclusively to its partially "locked-in" buyers. As Firm B raises its price, it becomes profitable for Firm A to compete for new entrants. Thus, there is an intermediate range of Firm B prices over which Firm A both retains its own customer base and captures a share of new entrants. Lastly, still higher Firm B prices enable Firm A to poach from Firm B's customer base.

Interestingly, the pricing rule in Figure 3.1 is not necessarily continuous across these regions. As a result, one firm's best response can pass through the "gap" in the other's, yielding *no* equilibrium (in pure strategies).⁶ The reason for these discontinuities can be traced to the fact that, given $s > 0$, no firm wishes to charge a price so as to acquire only

Figure 3.1 Firm A Best Response

NOTE: This shows Firm A's optimal price, given the Firm B price shown along the horizontal axis. The three regions of the graph are discussed in the main text.

SOURCE: Authors' calculations of simulation results from the model in Section Two.

a marginal share of new entrants. If Firm A does this, for instance, it renders the $y = 1$ entrant (the most complex design) indifferent between suppliers. But in that case, A's incumbents will be strictly inframarginal because they face $s > 0$. As a result, the firm can increase profit by discretely raising its price: It makes a higher profit from incumbents while sacrificing an infinitesimal share of new entrants. Accordingly, Firm A delays reducing its price to compete for incoming buyers. Then, when p^B is sufficiently high, Firm A can increase profits by reducing its price discontinuously and capturing a *discrete* share of new entrants, even while still charging a reasonably high price level to its incumbents.

Despite these discontinuities in the best responses, we identify a realistic calibration of the model under which there does in fact exist a Nash equilibrium in pure strategies, in which both suppliers sell to new entrants in each period (a “no-sale” equilibrium, to borrow from Farrell and Klemperer's [2007] language). We discuss this calibration in greater detail below.

The period-2 problem

We now turn to the period-2 problem. There are two types of buyers: new entrants and members of Firm A's customer base. We begin with the former. A period-2 entrant with design y will purchase from Firm A only if the presented discounted sum of period-2 and period-3 prices is less than what the entrant would face if it purchased from Firm B. This implies that the buyer at the threshold $y = y_2$ must be indifferent across suppliers. Accordingly, y_2 solves

$$(3.3) \quad p_2^A + \beta \min[p_3^A(y_2), p_3^B(y_2) + \tau y_2 + s] \\ = p_2^B + \tau y_2 + \beta \min[p_3^A(y_2) + s, p_3^B(y_2) + \tau y_2],$$

where $\beta < 1$ is the discount factor. Equation (3.3) implicitly defines the threshold, y_2 , as a function of period-2 prices, $y_2(p_2^A, p_2^B)$. Thus, Firm A's demand schedule among period-2 entrants is $1 - y_2(p_2^A, p_2^B) \equiv \sigma_2^{A0}(p_2^A, p_2^B)$.⁷

In addition, Firm A begins the period with a customer base. Let y_1 denote the threshold level of quality so that all entrants in period 1 (the initial period of the market) with $y \geq y_1$ participate and so purchase from Firm A. Thus, Firm A's customer base is $1 - y_1$. These buyers remain in the market for period 2 and then exit. Hence, their problem is a static one: they remain with Firm A if $p_2^A \leq p_2^B + \tau y + s$. Since y is uniformly distributed (conditional on $y \geq y_1$), Firm A retains a measure of its old customers equal to

$$\sigma_2^{AA}(p_2^A, p_2^B; y_1) = \Pr[p_2^A \leq p_2^B + \tau y + s \mid y \geq y_1] = \frac{\tau - p_2^A + p_2^B + s}{(1 - y_1)\tau}.$$

Firm A now solves

$$\pi_2^A(p_2^B, y_1) = \max_{p_2^A} \left\{ (p_2^A - c^A) \cdot Y_2^A(p_2^A, p_2^B; y_1) \right. \\ \left. + \beta \left(p_3^A(y_2(p_2^A, p_2^B)) - c^A \right) \cdot Y_3^A(y_2(p_2^A, p_2^B)) \right\},$$

subject to period-2 sales

$$Y_2^A(p_2^A, p_2^B; y_1) = \sigma_2^{A0}(p_2^A, p_2^B) + \sigma_2^{AA}(p_2^A, p_2^B; y_1)(1 - y_1)$$

and period-3 sales $Y_3^A(y_2(p_2^A, p_2^B))$, given in Equation (3.2). Firm B solves the analogous problem. We denote the equilibrium prices in the period by $P_2^A(y_1)$ and $P_2^B(y_1)$.

The initial period problem

The period-1 problem is a monopoly problem, as Firm A is the only supplier. The period-1 cohort's problem is to source its input from Firm A or not participate in the market at all. To solve this cohort's problem, then, we must make more explicit the demand side of the market. Our goal here is modest: We wish to introduce a reduced-form demand schedule that enables us to pose a simple monopoly problem for Firm A in period 1 and is consistent with the full participation of all period-2 and period-3 entrants. To this end, we assume that the payoff, F , to the buyer from its (unit) sale of the final good has the form

$$(3.4) \quad F(y) = R + ry,$$

where $R, r > 0$. This assumes, reasonably in our view, that higher-quality final goods “fetch” a higher price, so the payoff is increasing in y .

Given Equation (3.4), the buyer's problem in period 1 can be made straightforward, if we make three assumptions. First, if the buyer chooses to leave the market altogether in period 2, *exit is costless*. This means that a *sufficient* condition for participation in period 1 is $F(y) > p_1^A$. To see this, note that a buyer who enters in period 1 and remains in the market through period 2 has a present discounted payoff from participation equal to

$$F(y) - p_1^A + \beta \max \{F(y) - p_2^A, F(y) - p_2^B - \tau y - s, 0\}.$$

Since exit is costless, the buyer will leave the market if the maximum payoff across the two suppliers is negative. Furthermore, it pays no cost to leave. In this case, could a lower- y buyer ever be better off if it waited and signed up with the lower-price supplier, Firm B, in period 2? If it did, its discounted payoff would be $\beta \max \{F(y) - p_2^B - \tau y - s, 0\}$. Notice the presence of s , since the setup cost has to be paid upon entry. Comparing the two payoffs, it is clear that, as long as $F(y) - p_1^A > 0$, the buyer is always better off participating in period 1 than waiting until period 2.

However, the latter is not, in general, a necessary condition. Even if $F(y) < p_1^A$, a buyer may stand to make a profit in period 2. Thus, it may still enter in period 1 if that is the only opportunity for it to enter. There are a number of ways to make $F(y) > p_1^A$ a necessary condition. We choose to do this by assuming that the firm has no access to external finance. This implies the firm cannot borrow to cover losses during period 1, which in turn implies a nonnegativity constraint on dividends: $F(y) - p_1^A > 0$. Thus, $F(y) > p_1^A$ is a necessary and sufficient condition for participation in period 1.

Last, what happens to a firm if it declines to participate on account of $F(y) < p_1^A$? We assume that ideas are *not* storable. This means that firms for which $F(y) < p_1^A$ do not retain the option to return to the market in period 2.⁸ Therefore, we do not have to keep track of buyers that decline to enter in period 1.

These assumptions achieve a substantial simplification. In particular, if $F(y) > p_1^A$ is necessary and sufficient, then the choice of participation collapses to a static problem. From Equation (3.4), we see that the buyer, y , participates if $y \geq y(p_1^A) \equiv (p_1^A - R)/r$. It follows that the monopolist supplier faces a linear demand schedule $1 - y_1(p_1^A) = (R + r - p_1^A)/r$. The monopolist then selects its price p_1^A to maximize present discounted profits,

$$\max_{p_1^A} \{ (p_1^A - c^A)(1 - y_1(p_1^A)) + \beta \pi_2^A(y_1(p_1^A)) \},$$

where $\pi_2^A(y_1)$ is the discounted present value of profits as of the start of period 2 conditional on the equilibrium plays in periods 2 and 3.

Quantitative Analysis

Calibration

We now illustrate the model's mechanics. To that end, we calibrate and solve it numerically. There are seven parameters ($\beta, c^A, c^B, s, \tau, R, r$) that have to be chosen. Of these, only the discount factor β can be set without reference to a particular input market. We assume the period is one year and set $\beta = 0.95$, which implies an annual real interest rate slightly higher than 5 percent (Table 3.1). The remaining parameters will vary across markets. We calibrate the model to the offshore semiconductor wafer market.⁹

Table 3.1 Calibration

Parameter	Interpretation	Value	Target moment/reason
β	Discount factor	0.95	Real interest rate
c^A	Unit cost, Firm A	400	Long-run Taiwan price
c^B	Unit cost, Firm B	334	Long-run China price
τ	Monitoring cost	395	Long-run Taiwan mkt. share
s	Switching cost	197.5	Probability of switching
r	Buyer's payoff	432	Period-1 mkt. size
R	Buyer's payoff	688	Firm A profits

NOTE: This presents the calibration of the pricing game. The parameters are chosen so that the model induces the moments on the far right side of the table.

SOURCE: Authors' calculations of the values for the parameters used in the model simulation.

The costs of production and the quality premium, τ , are chosen to target the two suppliers' long-run price levels and the leader's (Firm A's) market share. To be more precise, we seek to have the model's terminal-period outcomes match observed outcomes "late" in a product's life cycle. As for how "late" ought to be measured, the model suggests that we would like to observe market outcomes after the initial cohort of Taiwan's customers conclude their production runs. The evidence available from supplier agreements indicates that customers arrange for at least three-year production runs, but with an option to renew.¹⁰ To allow for some "slippage" around the three-year mark, we focus on market outcomes after the first five years of a product's life.

Next, we select s . We have no direct estimates of this, but the testimony of industry experts (see footnote 4) suggests that switches are very rare. We also observe that customers remain in very long-term arrangements with suppliers. Fabless firms' annual reports to shareholders, for instance, show that fabless firms maintain relationships with Taiwan's TSMC and China's SMIC for at least four to five years at a time. Therefore, we choose s to imply a "low" switch rate, which we take to be on the order of 10–15 percent of Firm A's period-2 customer base.

Last, we now calibrate Equation (3.4). For given R , the slope, r , in Equation (3.4) pins down the firm's incentive to target high- y buyers: if r is high, the supplier charges a high price to take advantage of these buyers' willingness to pay. But as a result, many low- y buyers elect not to participate. We therefore set r to target a size for the period-1 market

relative to the size of the period-2 market. The idea here is that, in many markets, there is a ramp-up in terms of the volume of business after the introduction of a new product. Our data from the wafer market suggest that the size of the market at the time of product introduction is around one-half of its size in the mature phase of the product's life. Since there will be a measure 2 of period-2 buyers, we must then set r so that nearly a measure 1 of buyers elect to participate. This means that y_1 is not far from zero. Since there is a measure 2 of buyers in period 2, this corresponds to about one-half of the size of the market in the mature phase of the product life cycle.

As for R , this is chosen to ensure that all period-2 and period-3 entrants wish to participate. If R is sufficiently high, then Equation (3.4) indicates that even the lowest-quality buyer ($y = 0$) will make a purchase. In particular, one can easily show that, in order to guarantee full participation in periods 2 and 3, it is sufficient that $R > \hat{R} \equiv \max_{t=2,3} \{p_t^A, p_t^B\}$.¹¹ Of course, this provides only a lower bound; it does not point-identify R . To do the latter, we note that our model very likely understates the degree of competition in this market. Though Taiwan and China are the most significant producers, there are others. Therefore, we choose R in order to contain the rather outsized profits implied by the model. This means that R is set to roughly target \hat{R} .

Results: Price dispersion

We focus here on the model's predictions regarding the dynamics of price dispersion. We delay a discussion of aggregate price changes until later. Table 3.2 reports the results. There are two we wish to highlight.

First, the model implies that the degree of price dispersion declines over the product life cycle. The model implies a gap of roughly \$250 in the period in which the lagging supplier enters, and a gap of around \$150 in the next period. The source of these dynamics is very intuitive. In period 2, the leader charges a relatively high price to its customers, who are partially locked in because of the cost to switch. However, as the leader's original customers exit the market, it has a stronger incentive to compete aggressively for new entrants. The difference in prices between the leader and the follower therefore narrows. This result is a simple but important property of the model, and it is one that is absent if the only source of price dispersion is time-invariant unobserved het-

Table 3.2 Equilibrium Prices and Market Shares

Firm B period-2 price, p_2^B	456.37
Firm B period-3 price, p_3^B	686.80
Period-2 price differential, Δ_2	252.00
Period-3 price differential, Δ_3	150.38
Frictionless ($s = 0$) price differential	153.67
Measure of participants in period 1	0.99
Firm A period-2 market share	0.67
Firm A period-3 market share	0.55
Measure of switchers in period 2	0.13
Measure of switchers in period 3	0.00

NOTE: This presents the equilibrium of the pricing game discussed in the main text.

The calibration underlying this solution is shown in Table 3.1.

SOURCE: Authors' calculations of simulation results from the model in Section Two.

erogeneity. Thus, it gives us a testable prediction to take to data, which we do in the next section.

We also note here that the magnitude of dispersion in period 2, and the extent of its decline in period 3, line up reasonably well with the estimates from the semiconductor wafer market discussed in Section 3. For this reason, we believe that our calibrated model, although quite simple, provides some insight into price determination in this market. As such, it should serve as a useful laboratory in which to study the properties of various price indexes, a topic to which we return in Section Four, "Feasible Price Indexes."

Second, the price differential not only narrows in period 3, but it very nearly approaches the differential in the *frictionless* model where $s = 0$. To see more clearly how this comes about, return to the period-3 problem for a moment. We impose the restriction that each supplier retains its customer base from the period-2 cohort, as occurs in equilibrium for our calibration (see the final row in Table 3.2). Under these conditions, a little bit of algebra reveals that the difference, $\Delta_3 \equiv p_3^A - p_3^B$, in period-3 prices is given by the expression

$$(3.5) \quad \Delta_3 = \Delta^* + (2(1 - y_2) - 1) \frac{\tau}{3},$$

where Δ^* represents the difference between Firm A and B prices in the frictionless ($s = 0$) equilibrium.

We unpack Equation (3.5) in two steps. First, it is straightforward to show that Δ^* is the difference in market prices that makes the marginal buyer with design y^* indifferent across suppliers. This means that Δ^* compensates for the transaction cost, so that $\Delta^* = y^* \tau$. We therefore interpret Δ^* as the difference in market prices that could be accounted for by (unobserved) heterogeneity in quality. Second, the source of the wedge between Δ_3 and Δ^* is intuitive. To see this, note that the wedge vanishes if $y_2 = 1/2$. Each supplier in this case charges a higher price level than in the frictionless model, but the two suppliers' incentive to exploit their customer bases is the same. Hence, the price difference reflects entirely the difference in the quality of the production service. If $y_2 < 1/2$, in contrast, then Firm A's market share is relatively large. As a result, it charges a relatively high price to "milk" its customer base, and $\Delta_3 > \Delta^*$.

For our calibration, it is true that $y_2 \approx 1/2$, which implies $\Delta_3 \approx \Delta^*$. This suggests that one might use the *observed* difference in market prices late in the product cycle to proxy for the contribution of *unobserved* heterogeneity to the price differentials. In particular, one can subtract Δ_3 from price differentials earlier in the product life cycle, such as Δ_2 , and thereby adjust prices all along the product life cycle for unobserved heterogeneity. This yields an estimate of the share of the period-2 differential that is due to frictional dispersion. More exactly, we have that

$$(3.6) \quad \frac{\Delta_2 - \Delta_3}{\Delta_2} \approx \frac{250 - 150}{250} = 40\% ,$$

which is the percentage of the observed differential that is "real."¹² As we discuss in greater detail below, this simple correction will significantly aid our measurement strategy in Section Four.

Before we turn to the model's implications for aggregate price changes, we digress slightly in the next section to consider some recent evidence for the model's key prediction regarding the dynamics of price dispersion.

AN APPLICATION TO THE SEMICONDUCTOR INDUSTRY

We believe that the model presented in the previous section captures features common to many intermediate input markets. Indeed, it

is, in principle, relevant to any market with entry and clearly defined product turnover. But our empirical exploration of the model's implications is, of course, limited by available data. In a prior work (Byrne, Kovak, and Michaels 2013) we focused on the contract semiconductor manufacturing industry, for which we have detailed, transaction-level data. In the remainder of this section, we briefly review the structure of this market and the findings reported in Byrne, Kovak, and Michaels.

Semiconductor production involves a number of discrete steps.¹³ A chip is first designed using computer-aided tools that convert the desired functionality into a network of transistors and interconnections. The chip is then fabricated by depositing and etching away conducting and insulating materials to create a three-dimensional pattern of transistors and connections on the surface of a silicon wafer. Each step in the process is repeated for each of many chips, called "die," resulting in a grid of identical, completed die on the surface of the wafer. The die are then tested, sliced up, and placed in protective packages with leads allowing the chips to be connected to circuit boards in a final product.

Byrne, Kovak, and Michaels (2013) focus on the second step in this production process—the fabrication of semiconductor chips based on a particular design. Semiconductor fabrication technology has evolved steadily over time and can be characterized by a few observable technological traits such as the size of the wafer and the size of the smallest feature that can be produced on the surface of the wafer, called the "line width." The number of physical layers needed to create the chips has also increased over time, reflecting increased design complexity and leading to increased fabrication cost. Semiconductor technology evolves discretely over time, with only a few specific wafer sizes and line widths present in the market at any moment in time, making it possible to control for technological differences across products very flexibly.

Byrne, Kovak, and Michaels' (2013) empirical results use data on arm's-length transactions between firms specializing in chip design and marketing, called "fabless firms" since they have no fabrication facilities, and firms called "foundries" that specialize in fabricating other firms' chips. Most fabless firms are located in the United States and Europe, and they correspond to the buyers in the model just described. The largest foundries are located in Taiwan and China, which together account for 74 percent of foundry output. Taiwanese foundries enter a

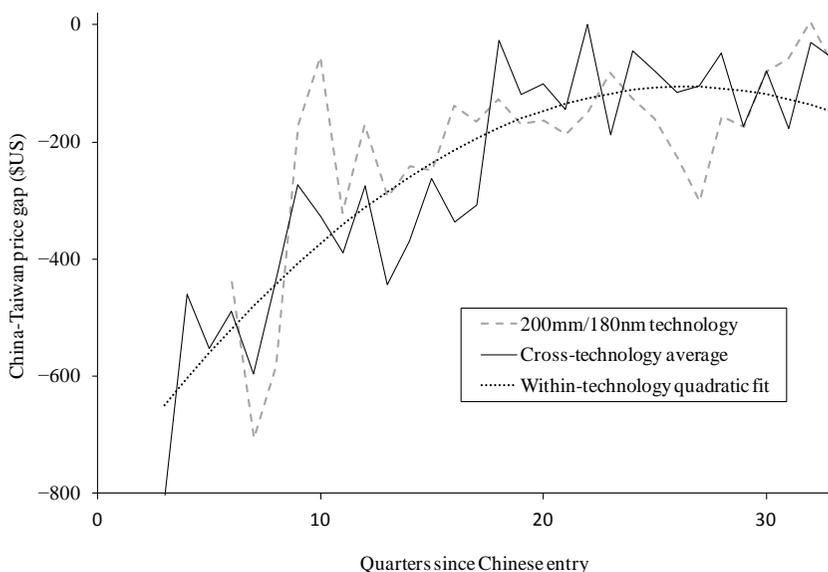
product market, defined by a wafer-size and line-width combination, at least eight quarters ahead of Chinese foundries. The dominant Taiwanese foundry, Taiwan Semiconductor Manufacturing Company (TSMC), is the overall market leader—it is Firm A in the model. TSMC is widely considered as possessing the most advanced design integration tools and engineering support.

The data come from a proprietary database collected by the Global Semiconductor Alliance (GSA), a nonprofit industry organization. Byrne, Kovak, and Michaels' (2013) extract spans 2004–2010 and covers a representative sample of about 20 percent of the wafers produced by the worldwide foundry sector. The GSA data are unique in providing details on transaction prices, along with all technological characteristics of finished semiconductor wafers that are relevant for pricing, including wafer size, line width, and numbers of various types of layers. This detailed product-characteristic information makes it possible to compare average prices for physically identical inputs across suppliers located in different countries.

Formally, Byrne, Kovak, and Michaels (2013) implement such a comparison in a hedonic regression framework that relates wafer prices to observable technological characteristics, quarter indicators, and indicators for supplier's location.¹⁴ The controls for product characteristics enable them to estimate the effect of location on price, holding fixed the composition of goods. The data reveal substantial price differences across suppliers. Comparing the two largest suppliers, a Chinese wafer sells at a 17 percent discount compared to an otherwise identical Taiwanese wafer.¹⁵

This average difference masks, however, interesting *dynamics* in price dispersion. Byrne, Kovak, and Michaels (2013) go on to estimate how the price differential evolves following Chinese entry. The key result is replicated in Figure 3.2 (which is Figure 4 in Byrne, Kovak, and Michaels). The dashed gray line plots raw quarterly price differences for the process technology with the largest sales during our time period, 200mm wafers with 180nm line width. Despite the noise in the series, it is clear that the average price differential closes considerably over the life of this technology. It falls from around \$600 to around \$150 more than five years after Chinese entry. This pattern applies to other technologies with smaller sales as well. The black solid line in the figure plots the difference in price averaged across all technologies in each

Figure 3.2 The Closing China-Taiwan Price Gap Following China's Entry



NOTE: This presents the difference between China's and Taiwan's price for certain categories of wafers. The 200mm/180nm is one of the most popular wafers in the sample. The cross-technology average measures the mean of the price differentials across wafers. The within-technology fit is derived from a regression model with wafer fixed effects and thus uses only within-technology variation in the price differential. See Byrne, Kovak, and Michaels (2013) for more.

SOURCE: Based on regression results in Byrne, Kovak, and Michaels (2013).

quarter following Chinese entry and exhibits quite consistent declines in the gap between Chinese and Taiwanese prices.¹⁶ Last, whereas the solid line pools *across* technologies, the dotted line in the figure estimates the price differential based exclusively on the typical variation *within* the life of a technology. It reveals a very similar pattern.

As Byrne, Kovak, and Michaels (2013) stress, this dynamic pattern is unlikely to be driven by unobserved differences in products or services across Chinese and Taiwanese suppliers. The price differences start out large and then converge for each new process technology, so constant differences across suppliers or differences that evolve over calendar time for all technologies are unlikely to explain the observed

pattern. In other words, steady improvements in the quality or reliability of China's production service may explain price differentials *across* technologies, but they are unlikely to account for the sharp, *within*-technology dynamics we observe. This also rules out explanations related to brand recognition, customer service, intellectual property-rights protection, tax policy, and other factors that might make Chinese producers more attractive over time.

Accordingly, Byrne, Kovak, and Michaels (2013) argue that the dynamics reflect the presence of real, frictional price dispersion. The pattern of narrowing differentials is clearly consistent with that predicted by the theory of switching costs sketched in Section Two, "A Pricing Game with Costly Switching." This finding motivates our work in the next section, as we consider developing a price index that admits roles for both frictional dispersion and constant, unobserved heterogeneity.¹⁷

FEASIBLE PRICE INDEXES

Our goal is to measure the change in price of a production service in an environment in which the quality of service varies across suppliers. The *source* of the difference in service quality is not especially critical, though, for the purpose of this exercise. In our model, the quality of service varies inversely with the complexity of the design, y . But this is merely one way to operationalize the idea; heterogeneity across designs is *not*, per se, significant.

It is worth taking a moment at the outset to elaborate on this point. The following discussion should help reveal the generality of the problem confronting price index construction. In so doing, it also points the way toward developing the "ideal" price index in this setting, which will serve as the benchmark against which all feasible alternatives are judged.

A Benchmark

There is a way to reinterpret the model that is particularly helpful. Imagine that suppliers produce *and ship* the input to customers. Assume, moreover, that Firm B has an inferior transport technology. In

this context, it is natural to reinterpret y as distance from Firm B. The cost, τy , is then read as a transport cost, so that a Firm B customer with unit demand who is y units away must purchase $1 + \tau y$ units because τy are “lost in transit.” In this interpretation, it is the *customer* who implants its own design on the chip after receipt of the product.

This problem is formally identical to our own. But what design the retailer implants on the wafer *after* it is shipped by Firm B is clearly orthogonal to how we measure the price of the production (and transportation) service. In other words, in this (re)interpretation, the measurement of the input’s price is unrelated to, and unaffected by, the presence of heterogeneous designs. All that matters, in terms of the real resource cost to customers, is the transaction cost, τy .

It follows that, in our preferred interpretation of the problem, design heterogeneity matters for price measurement only insofar as it implies a particular transaction cost. If we could just observe these costs (τy), we would fold them into a comprehensive, or quality-adjusted, measurement of the price paid by Firm B buyers, $\hat{p}^B \equiv p^B + \tau y$, for the production service. After adjusting for τy , this is the same service provided by Firm A. Hence, at that point, we simply aggregate across the p^A ’s and \hat{p}^B ’s in a particular period and compare the result with the average price in the prior period. This is, in fact, how we will build our benchmark price index, to which we now turn.

Our benchmark index requires the most information on the part of the analyst. In particular, the analyst is assumed to observe the transaction cost, τy , paid by a Firm B customer with design y . Hence, the analyst measures prices, p^A and \hat{p}^B .

Since each price in the model (inclusive of τy) pertains to the same service, it is not hard to aggregate across observed prices. In any period t , there is a set I_t of buyers with measure μ_t . The period- t (with $t = 2$ or 3) price is then given simply by

$$\frac{1}{\mu_t} \int_{I_t} p_t(i) di ,$$

where $p_t(i)$ is the price paid by buyer $i \in I_t$ (and equal to either p_t^A or \hat{p}_t^B). This aggregates prices paid across the measure μ_t of buyers, weighting each equally, since all participants purchase one unit of the input.¹⁸ This

price is compared to the average price in the prior period to derive the price change.

To illustrate the calculations, consider the period-2 problem. It is helpful to first take the case where Firm A retains the full measure $1 - y_1$ of its period-1 customers but only sells to new entrants with $y \geq y_2$. It follows that a measure $2 - y_1 - y_2$ pays p_2^A for the input, whereas each new entrant $y \in [0, y_2]$ pays $p_2^B + \tau y$. Therefore, the average period-2 price, P_2^* , is

$$P_2^* \equiv \frac{1}{2 - y_1} \int_{[y_1, 1] \cup [0, 1]} p_2(i) di = \frac{2 - y_1 - y_2}{2 - y_1} p_2^A + \frac{1}{2 - y_1} \int_0^{y_2} [p_2^B + \tau y] dy.$$

Since the only price in period 1 is p_1^A , the price index in period 2 would be P_2^*/p_1^A .

The problem is slightly more cumbersome if Firm B poaches in period 2, as occurs in the model. In this case, Firm B poaches from Firm A for all qualities less than a threshold, y_2^p , where the threshold satisfies $y_1 < y_2^p < y_2$. Therefore, Firm A supplies a measure $1 - y_2$ of entrants and $1 - y_2^p$ of incumbents. Firm B supplies, in turn, a measure of y_2 of entrants and $y_2^p - y_1$ of incumbents. Given this distribution of buyers across designs, the average price becomes

$$(3.7) \quad P_2^* \equiv \frac{2 - y_2 - y_2^p}{2 - y_1} p_2^A + \frac{1}{2 - y_1} \left\{ \int_0^{y_2} [p_2^B + \tau y] dy + \int_{y_1}^{y_2^p} [p_2^B + \tau y] dy \right\}.$$

In what follows, it will be instructive to integrate the terms enclosed in braces on the right half of Equation (3.7) and rewrite this as a weighted average of the suppliers' quality-adjusted prices,¹⁹

$$(3.8) \quad P_2^* \equiv \frac{2 - y_2 - y_2^p}{2 - y_1} p_2^A + \frac{y_2 + y_2^p - y_1}{2 - y_1} (p_2^B + \theta \tau),$$

where

$$(3.9) \quad \theta \equiv \frac{y_2}{y_2 + y_2^p - y_1} \cdot \frac{y_2}{2} + \frac{y_2^p - y_1}{y_2 + y_2^p - y_1} \cdot \frac{y_2^p - y_1}{2}$$

is the average design supplied by Firm B. Hence, the average price (Equation [3.8]) is a weighted average of the market prices plus a measure of the average transaction cost paid by Firm B buyers. Again, the price index is simply P_2^*/p_1^A .

Feasible Alternatives

In practice, the BLS does not observe the level of detail—namely, τ —needed to calculate the benchmark index. We now consider several indexes with less demanding data requirements. We refer to these as *feasible price indexes*.

Within index

The first of these is consistent with our understanding of BLS-IPP practice, which typically treats the identity of the seller as a price-forming characteristic (Nakamura and Steinsson 2012). In this case, there are, in effect, two types of goods: those sold by Firm A and those sold by Firm B. Under these circumstances, standard practice is to compute the price index by first calculating price changes within each supplier and then averaging these changes across suppliers. We refer to this as the *within* index. This contrasts with the benchmark index, which first averages prices across suppliers and then takes the difference.

Applied to the period-2 data, the within index is very simple. Since Firm B does not participate in period 1, the within index is computed just by taking the ratio of Firm A's market, p_2^A/p_1^A .

Average index

The second measure takes the opposite approach to the problem of unobserved quality. The strategy here is to take the average period-2 *posted market price* across suppliers for all qualities greater than y_1 and compare it to the period-1 price. We refer to this as the *average* index.²⁰

This index is distinguished by the fact that it takes no account of the transaction costs—it makes no quality adjustment. Accordingly, the average index is calculated by simply excluding τy from the price paid by each of Firm B's customers in the benchmark index (Equation [3.7]). The average period-2 price is then

$$\begin{aligned}
 (3.10) \quad P_2^\alpha &\equiv \frac{2 - y_2 - y_2^p}{2 - y_1} p_2^A + \frac{1}{2 - y_1} \left\{ \int_0^{y_2} p_2^B dy + \int_{y_1}^{y_2^p} p_2^B dy \right\} \\
 &= \frac{2 - y_2 - y_2^p}{2 - y_1} p_2^A + \frac{y_2 + y_2^p - y_1}{2 - y_1} p_2^B,
 \end{aligned}$$

and the index is calculated according to P_2^a/p_1^A . This approach is more concerned that the new supplier is likely to sell the same qualities at lower prices, which does in fact happen if $y_1 < y_2$. In these instances, quality-adjusted price declines faced by buyers are not recorded by the within index.

Diff-in-diff index

Last, we present an index that attempts to strike a compromise between the within and average indexes. The index confronts the challenge of unobserved heterogeneity but does not abandon the idea that there may be quality-adjusted price dispersion in equilibrium. At the same time, it does not place the same data requirements on the analysts as the benchmark index does.

The construction of what we will call the diff-in-diff index is guided by the model in Section Two (“A Pricing Game with Costly Switching”). One of the key points of the section was the idea that one can use the *observed* price differentials late in the product life cycle, Δ_3 , to proxy for the contribution of *unobserved* heterogeneity, denoted by Δ^* . This implies that the quality-adjusted period-2 differential, Δ_2 , can be estimated by netting off Δ_3 , which is the most that could be attributed to quality. This boils down to doing quality adjustment by simply inflating Firm B’s price by Δ_3 . To see this, let \check{p}_2^B denote our estimate of Firm B’s period-2 price adjusted for service quality. We define \check{p}_2^B according to the quality-adjusted period-2 differential, $p_2^A - \check{p}_2^B = p_2^A - p_2^B - (p_3^A - p_3^B)$. Canceling the p_2^A ’s, we have that $\check{p}_2^B = p_2^B + \Delta_3$.

This corrected price differential is the key input into the diff-in-diff index. The index itself is now easy to construct. We add Δ_3 to p_2^B in the average index (Equation [3.10]), use the proxy $\Delta_3 = \Delta^* = \tau y^*$, and integrate. (Recall that y^* is Firm B’s market share in the frictionless equilibrium.) The result can then be written as

$$(3.11) \quad P_2^\delta \equiv \frac{2 - y_2 - y_2^p}{2 - y_1} p_2^A + \frac{y_2 + y_2^p - y_1}{2 - y_1} (p_2^B + y^* \tau).$$

Comparing Equation (3.11) to the benchmark (Equation [3.8]), we see our adjustment is exact only if $y^* = \theta$. In fact, a discrepancy between y^* and θ will likely arise. To see why, recall from Section Two that what drives a wedge between market prices in the frictionless equilibrium

is the transaction cost faced by the *marginal* buyer (who is indifferent across the two suppliers). That is, the difference, Δ^* , in market prices is given by $y^*\tau$. However, a quality-adjusted price index like those in Equations (3.7) and (3.8) requires calculation of the *average* price inclusive of transaction costs among *all* Firm B buyers, as represented by θ . Since the transaction cost increases with y , the marginal cost exceeds the average. It is very likely, then, that $\theta < y^*$. In that case, P_2^δ would overestimate P_2^* .

The comparison between y^* and θ becomes clearer if we consider a special, but informative, case. Suppose that $y_1 = 0 = y_2^p$. This is an instructive case because, in our calibrated model, most buyers do in fact participate in period 1 (so that $y_1 \approx 0$), and switching is minimal in period 2. It follows that Equation (3.9) collapses to $\theta = y_2/2$, where y_2 is Firm B's share of period-2 entrants under $s > 0$. Hence, in this case, $\theta = y^*$ only if $y_2 = 2y^*$: Firm B claims *twice* as many period-2 entrants under $s > 0$ as it does in the frictionless equilibrium. The intuition behind this is that the average transaction cost, θ , approaches y^* only if Firm B supplies very complex designs when $s > 0$. This is unlikely, and our calibrated model doesn't bear this out. It follows that the quality adjustment is too large in Equation (3.11), so that $P_2^\delta > P_2^*$.

Reacting to this, we wish to make two observations. First, if we make no changes to Equation (3.11), it could still be used productively by agencies with the understanding that it provides an upper bound on the true quality-adjusted price. Second, we can complement this upper bound by considering a lower bound; a comparison of the two will help better identify the true change. To see this, suppose the switching cost did not distort the distribution of market shares, so that $y^* = y_2$. Accordingly, in the special case where $\theta = y_2/2$, one can align P_2^δ with P_2^* by just dividing $\Delta^* = y^* \tau$ in Equation (3.11) by 2. This yields an alternative to Equation (3.11),

$$\hat{P}_2^\delta \equiv \frac{2 - y_2 - y_2^p}{2 - y_1} p_2^A + \frac{y_2 + y_2^p - y_1}{2 - y_1} \left(p_2^B + \frac{y^*}{2} \tau \right).$$

In all likelihood, the switching cost *will* affect the distribution of market shares—in particular, the entering firm will compete relatively aggressively to attract customers, since the buyers will be subsequently locked in. This suggests that $y^* < y_2 = 2\theta$. As a result, $y^*/2 < \theta$, which

means that $\hat{P}_2^\delta < P_2^*$: we obtain an estimate of P_2^* that is *downwardly* biased. Comparing \hat{P}_2^δ to the baseline diff-in-diff index (Equation [3.11]), one can better gauge the true price change. We implement this procedure below.

We complete this discussion by quickly mentioning how to apply these indexes to period 3. In correspondence with the calibrated model, we assume both firms retain all buyers in period 3. Hence, Firm A sells to a measure $1 - y_2 + 1 - y_3$, and Firm B sells to a measure $y_2 + y_3$. The calculations of each index then follow by analogy to their period-2 counterparts. For instance, the benchmark index is P_3^*/P_2^* , where P_3^* is given by

$$P_3^* \equiv \frac{2 - y_2 - y_3}{2} p_3^A + \frac{y_2 + y_3}{2} (p_3^B + \Theta\tau),$$

$$\text{and } \Theta \equiv \frac{y_2}{y_2 + y_3} \frac{y_2}{2} + \frac{y_3}{y_2 + y_3} \frac{y_3}{2}$$

is the average design supplied by Firm B in period 3. Next, the within index is obtained by first computing the change in each supplier's *market* price and then aggregating these price changes across suppliers. We use Tornqvist weights in the latter step, which yields

$$\omega \frac{p_3^A}{p_2^A} + (1 - \omega) \frac{p_3^B}{p_2^B},$$

where

$$\omega = \frac{1}{2} \left(\frac{2 - y_2 - y_3}{2} + \frac{2 - y_2^p - y_2}{2 - y_1} \right)$$

is the average Firm A market share across periods 2 and 3. Finally, using the appropriate period-3 market shares, the average and diff-in-diff indexes can be computed according to the expressions contained in Equation (3.10) and Equation (3.11).

Results

Table 3.3 uses our calibrated model to assess the accuracy of the feasible indexes. Each column corresponds to a period. Each row reports the gross price change implied by the index *relative* to the gross price

Table 3.3 Feasible Indexes Relative to Benchmark

	Period 2	Period 3
Within	1.082	1.047
Average	0.958	0.990
Diff-in-diff, baseline	1.032	1.001
Diff-in-diff, alternative	0.995	0.997

NOTE: This table uses the solution to the calibrated model to calculate the gross price changes implied by a variety of price indexes. The results are expressed here relative to the true gross price change. See main text for a discussion of the indexes and Table 3.1 for the calibration.

SOURCE: Authors' calculations of simulation results from the model in Section Two.

change implied by the benchmark. For instance, with respect to period 2, the row that reports the average index presents

$$\frac{P_2^\alpha / p_1^A}{P_2^* / p_1^A}.$$

Thus, if the estimate in the row is less than 1, the feasible index *understates* the true change. Equivalently, the feasible index *overstates* the decline in the price level between periods 1 and 2.

The within and average indexes yield estimates in line with our expectations. As we showed in Section Two, roughly 40 percent of the period-2 price differential cannot be attributed to quality dispersion; Firm B does provide a real discount. As a result, the within index fails to capture the full extent of the decline in the average price level driven by the entry of Firm B. The table reveals that it overstates the price change (understates the price decline) by 8 percentage points. At the same time, as our discussion in Section Two noted, there is a quantitatively significant component of price dispersion owing to difference in service quality. The average index fails to account for this and so understates the true price change. Equivalently, it overstates the price decline—in this case by about 4 percentage points.

We now turn to the performance of the diff-in-diff index. The table reports results for both the baseline index derived from P_2^δ and the alternative based on \hat{P}_2^δ . As we anticipated, the baseline index (Equation [3.11]) outperforms the within index, since it treats a portion of Firm B's price as a real, quality-adjusted discount relative to Firm A's price. Accordingly, it better captures the decline in the average price level

when Firm B enters. Yet it still understates the extent of the price decline by 3 percentage points. Interestingly, the alternative index, based on \hat{P}_2^δ , performs noticeably better. As we noted earlier, it should be the case that $\hat{P}_2^\delta > P_2^*$, but this discrepancy depends on the distance between y_2 and y^* . This distance turns out to be rather limited in this calibration, but we stress that it is hard to judge the robustness of this result.

As for period 3, the diff-in-diff indexes perform very well. Mechanically, the reason is that P_2^δ and P_3^δ both overstate the corresponding true prices. These errors appear to cancel each other out, so the gross change, P_3^δ/P_2^δ , turns out to very nearly equal P_3^*/P_2^* . The same idea applies to the average index. However, with respect to the within index, errors do not cancel each other out so fortuitously; this continues to overstate the true price. Again, it is difficult to know if these results hint at a more general lesson. We hope continued work in this area will help elucidate this.

CONCLUSION

This paper has studied the problem of price index construction for intermediate inputs when observed price differentials are combinations of unobserved heterogeneity and real, or frictional, price dispersion. In particular, it assesses several price indexes that can be *feasibly* constructed. Our results provide some guidance for how to adjust price indexes when a new low-price supplier, such as China, joins a market. In our application to the semiconductor market, we find that if frictional dispersion is ignored, the price index overstates the true price decline due to entry by five to eight percentage points. Ignoring unobserved heterogeneity, in contrast, means that lower-quality service by the entrant is not accounted for, and thus the price decline is overstated. We then try to provide a pathway between these extrema. Our diff-in-diff index exploits a simple insight: the cost of switching to a new supplier in this market sustains frictional dispersion during the early life of a product, but this influence abates as the market matures and the market leader's original customers exit. Thus, late in the product life cycle, the price difference largely reflects time-invariant quality differences. Accordingly, one can use these observed late-in-life price differences to correct

for unobserved heterogeneity and thereby isolate the extent of real price dispersion. For this reason, our diff-in-diff index performs quite well as an approximation to the true price change.

Our assessment, of course, is confined to a particular market in the semiconductor sector. Yet we believe our approach provides a fruitful way forward in this literature. That approach, in sum, consists of a few components: gather detailed data for a particular industry; develop a quantitative model of industry dynamics that can be fitted to these data; and assess alternative, feasible price indexes within the context of the parameterized model. If applied to several industries, we believe this approach has the promise of revealing more general lessons for price index measurement.

Notes

1. By “Taiwan,” we mean Taiwan Semiconductor Manufacturing Corporation, or TSMC, the largest wafer fabrication firm in Asia. Most of its properties are in Taiwan, though it has one plant in Shanghai. However, the vast majority of production in China is due to Semiconductor Manufacturing International Corporation, or SMIC. When we refer to “China,” then, we mean SMIC.
2. China typically enters two years after Taiwan initiates production.
3. See Farrell and Klemperer (2007) and Klemperer (1995) for surveys.
4. This quote is from a report describing the Common Platform technology alliance. This is an industry group consisting of a few large chip manufacturers—IBM, Chartered Semiconductor Manufacturing, and Samsung. The group advocates for a “common platform” that would standardize aspects of semiconductor production technology. However, this alliance has not yet had a material impact on standardizing mask sets (McGregor 2007).
5. We have obtained a handful of these contracts. A representative agreement in terms of how price discrimination is handled is one between Altera Corp. and TSMC. It states that “TSMC shall calculate an average price for such Process in use at all of TSMC’s . . . plants,” and if the buyer’s price “deviates, up or down, by more than three percent (3%) from the [average price],” the buyer’s price will be adjusted in the direction of the average price. Note that this agreement does not commit TSMC to a particular price path over time. The contract merely restricts price discrimination in a given period, consistent with the model’s assumptions. This agreement is found at <http://corporate.findlaw.com/contracts/operations/purchase-agreement-taiwan-semiconductor-manufacturing-co-ltd.html> (accessed April 22, 2014).
6. See Nishimura and Friedman (1981) for an analysis of this class of games. They provide sufficient conditions to ensure a pure-strategy equilibrium, but these conditions can only be confirmed ex post. This is, in effect, what we aim to do.

7. It is not immediate that there is a unique solution for y_2 , but we have always located one in practice. The intuition for this is as follows. The discounted sum of Firm B prices is relatively low when y_2 is low (i.e., when Firm B's customer base is small, it sets lower prices in period 3 to attract new entrants). But the discounted sum of Firm B prices is also increasing at a relatively fast rate as y_2 rises. This reflects the quality premium, as captured by τy_2 . Together, these features imply a single crossing, with the right side of Equation (3.3) cutting the left side from below.
8. Interestingly, another way to make $F(y) > p_i^d$ a necessary condition is to drop the nonnegativity constraint and to assume that ideas are instead *perfectly storable*. This would imply that a firm would never produce in period 1 if its instantaneous profits were negative; it would just store the idea and join the market in period 2. Note that, since production runs for two periods, these late entrants would presumably live through period 3. This points to the downside of this approach: delayed entry reverberates through the model's periods 2 and 3 and creates a more complicated dynamic problem. Moreover, the payoff from this added complication is rather small. As we discuss, the model will be calibrated in such a way that the measure of firms that delay entry is very small, so its quantitative implications cannot be too great. For this reason, we choose the simpler approach in the main text.
9. The subsequent two paragraphs are taken from Byrne, Kovak, and Michaels (2013) (see Appendix 6D).
10. For example, a contract between Quicklogic and TSMC states, "The term of this Agreement shall . . . continue for a period of three (3) years, renewable annually as a rolling three (3) year Agreement."
11. Since the lowest-quality buyer in Firm A's cohort has payoff $R + \tau y_i$ in period $t = \{2,3\}$, it follows that Firm A buyers will in fact participate if R exceeds $\max_i \{p_i^d\}$. Furthermore, if $r > \tau$ (as it does, in our calibration), then the highest-quality customer of Firm B will participate if R is greater than $\max_i \{p_i^b\}$.
12. We also stress that, although $y_2 \approx 1/2$ in our model, the approach suggested by Equation (3.4) can be applied robustly to real-world data even if there are certain deviations from this. For instance, if y_2 is smaller than $1/2$, then Δ_3 overestimates its frictionless counterpart. As a result, if we used Δ_3 as a proxy for the contribution of unobserved heterogeneity to the price differential, we would obtain a *lower* bound on the degree of pure (frictional) price dispersion. This property can be desirable in certain circumstances. For instance, though we assume the suppliers provide the same physical input here, data limitations may make it impossible for a statistical agency to do any direct hedonic-style quality adjustments for product composition. In that case, it may want to err on the side of unobserved heterogeneity.
13. Turley (2003) provides an accessible overview of semiconductor technology, manufacturing, and business.
14. The GSA data do not provide firm identifiers, only the country in which the supplier is located.
15. See Table 3 in Byrne, Kovak, and Michaels (2013) for the full list of regression coefficients from the hedonic model.

16. To be more precise, Byrne, Kovak, and Michaels (2013) regress the price differential on, among other controls, a quadratic time trend and product fixed effects; the latter control for changes in the composition of technologies. The dotted line, referred to in Figure 3.2 as the “within-technology fit,” is the estimated time trend.
17. Byrne, Kovak, and Michaels make the correction embodied in Equation (3.6) to the raw wafer price differentials in order to isolate the component that is due to real dispersion. In their application, they interpret “late in the product life cycle” to be roughly five years after Chinese entry, based on the length of typical semiconductor fabrication contracts (see the subsection titled “Quantitative Analysis,” beginning on p. 100). Then, drawing from Figure 3.2, the authors interpret $\Delta_3 = \Delta^* \approx \150 . Netting this off of the observed period-2 differential, Δ_2 , yields the quality-adjusted component. For instance, the average differential 10 quarters after Chinese entry is about \$375, so the authors estimate that 60 percent $\left(\frac{375 - 150}{375}\right)$ reflects real price dispersion. Hence, the wafer data indicate more frictional dispersion than implied by the model.
18. Simple averaging across buyers is appropriate within our theoretical model because the production service, modulo τy , is in fact identical. The BLS does not follow this approach when aggregating across outlets’ prices at the most detailed level (the item-area stratum) of the CPI. This is because BLS staff worry that different outlets’ items are not in fact the same. See Hausman and Leibtag (2009) for more on this practice. See Shapiro and Wilcox (1996) for a general discussion of aggregation within the CPI.
19. The weights are formed from the quantity of units sold by each supplier. Throughout, we do assume that the statistical agency has access to price *and* revenue data from the supplier, so that quantities may be inferred. The BLS International Price Program does request data on the dollar value of trade for each good when a firm is initiated into the survey.
20. The average index does embody a slight recognition of quality differences, in that it excludes never-before-priced designs in period 2. In this sense, the index acknowledges that some period-2 goods are “too different” from the basket of goods in period 1 to be included in the index. We take this approach to try to capture the fact that the statistical agency does observe repeated sales of the same product, even if it does not observe quality precisely. Still, this inclusion only of designs $y > y_1$, as opposed to all designs, makes little quantitative difference to our results, since y_1 is so close to zero.

References

- Bayard, Kimberly, David Byrne, and Dominic Smith. 2013. "The Scope of U.S. Factoryless Manufacturing." Unpublished paper, Board of Governors of the Federal Reserve System, Washington, DC.
- Byrne, David, Brian K. Kovak, and Ryan Michaels. 2013. "Price and Quality Dispersion in an Offshoring Market: Evidence from Semiconductor Production Services." NBER Working Paper No. 19637. Cambridge, MA: National Bureau of Economic Research.
- Carlton, Dennis W. 1983. "Equilibrium Fluctuations When Price and Delivery Lag Clear the Market." *Bell Journal of Economics* 14(2): 562–572.
- Farrell, Joseph, and Paul Klemperer. 2007. "Coordination and Lock-In: Competition with Switching Costs and Network Effects." In *Handbook of Industrial Organization*, Vol. 3, Mark Armstrong and Robert H. Porter, eds. Handbooks in Economics 10. Amsterdam: North-Holland, pp. 1967–2072.
- Gabriel Consulting Group. 2006. GCG Abstract: Common Platform Technology Research Report. Beaverton, OR: Gabriel Consulting Group.
- Hausman, Jerry, and Ephraim Leibtag. 2009. "CPI Bias from Supercenters: Does the BLS Know That Wal-Mart Exists?" In *Price Index Concepts and Measurement*, W. Erwin Diewert, John S. Greenlees, and Charles R. Hulten, eds. Studies in Income and Wealth 70. Cambridge, MA: National Bureau of Economic Research, pp. 203–238.
- Hummels, David, Jun Ishii, and Kei-Mu Yi. 2001. "The Nature and Growth of Vertical Specialization in World Trade." *Journal of International Economics* 54(1): 75–96.
- Klemperer, Paul. 1995. "Competition When Consumers Have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade." *Review of Economic Studies* 62(4): 515–539.
- McGregor, Jim. 2007. *The Common Platform Technology: A New Model for Semiconductor Manufacturing*. Scottsdale, AZ: In-Stat.
- Nakamura, Emi, and Jón Steinsson. 2012. "Lost in Transit: Product Replacement Bias and Pricing to Market." *American Economic Review* 102(7): 3277–3316.
- Nishimura, Kazuo, and James Friedman. 1981. "Existence of Nash Equilibrium in nPerson Games without Quasi-Concavity." *International Economic Review* 22(3): 637–648.
- Shapiro, Matthew D., and David W. Wilcox. 1996. "Mismeasurement in the Consumer Price Index: An Evaluation." In *NBER Macroeconomics Annual 1996*, Vol. 11, Ben S. Bernanke and Julio J. Rotemberg, eds. Cambridge and London: MIT Press, pp. 93–142.
- Turley, Jim. 2003. *The Essential Guide to Semiconductors*. Upper Saddle River, NJ: Prentice Hall Professional Technical Reference.

Measuring Globalization

Better Trade Statistics for Better Policy

Volume 1

Biases to Price, Output, and Productivity Statistics from Trade

Susan N. Houseman
and
Michael Mandel
Editors

2015

W.E. Upjohn Institute for Employment Research
Kalamazoo, Michigan

Library of Congress Cataloging-in-Publication Data

Measuring globalization : better trade statistics for better policy / Susan N. Houseman and Michael Mandel, editors.

volumes cm

Includes bibliographical references and indexes.

ISBN 978-0-88099-488-0 (v. 1 : pbk. : alk. paper) — ISBN 0-88099-488-6 (v. 1 : pbk. : alk. paper) — ISBN 978-0-88099-489-7 (v. 1 : hardcover : alk. paper) — ISBN 0-88099-489-4 (v. 1 : hardcover : alk. paper)

1. Commercial statistics. I. Houseman, Susan N., 1956- II. Mandel, Michael J. HF1016.M44 2015
382.01'5195—dc23

2014047579

© 2015

W.E. Upjohn Institute for Employment Research
300 S. Westnedge Avenue
Kalamazoo, Michigan 49007-4686

The facts presented in this study and the observations and viewpoints expressed are the sole responsibility of the authors. They do not necessarily represent positions of the W.E. Upjohn Institute for Employment Research.

Cover design by Alcorn Publication Design.
Index prepared by Diane Worden.
Printed in the United States of America.
Printed on recycled paper.