Price dispersion and price dynamics when it is costly to switch suppliers: A quantitative example

Brian Kovak, Carnegie Mellon University

Ryan Michaels, University of Rochester

Prepared for "Measuring the Effects of Globalization"

February 2013

1. Introduction

As production has become more fragmented across establishments, firms, and countries, less developed economies such as China are capturing an ever greater share of the global market for tradable intermediates. This trend presents an important challenge for price index measurement (Houseman et al. 2010, 2011). The lower observed prices charged by suppliers in emerging economies may reflect differences in the quality of the products, and/or real discounts for goods of identical quality. Price indexes that assume the former will tend to under-state the degree of price decline faced by domestic firms that source production from these economies. However, it is equally perilous to neglect the scope for unobserved quality variation. The challenge to statistical agencies is that, in practice, it is very difficult to separately identify quality-related versus real price dispersion given data limitations.

In reaction to this, the present paper attempts to provide some guidance for price measurement when real price dispersion may be present. We consider how well imperfect but *feasible* price indexes approximate the true price change in markets characterized by two critical, and realistic, features: i) products are differentiated, but ii) even the same product may be priced differently due to trading frictions that impede arbitrage. In our model, two large suppliers, a leader and a follower, produce an intermediate input for overlapping generations of buyers. Quality variation arises from differences in the quality of the manufacturing services: one supplier (the follower) must be monitored and consulted more closely than the leader when producing complex goods, giving the leader an advantage in producing such goods. The trading friction takes the form of a lump-sum setup cost that is sunk with the original supplier and must be paid once again if a buyer switches to the other supplier during the life of its product. This friction implies that when the follower enters the market, much of the leader's customer base will choose to remain with the leader despite a higher price.

In section 2, we solve the model numerically and investigate its implications for price dispersion and price dynamics. Our results are two-fold. 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 conclude, it will compete more aggressively for new generations of buyers, causing the price dispersion between firms to narrow. This is a very intuitive result, and it is a key distinguishing property of the model – time-invariant quality differences would not induce this pattern in the *dynamics* of price dispersion.

Section 3 presents empirical evidence consistent with these dynamics. We summarize the empirical results of Byrne, Kovak, and Michaels (2012) (hereafter BKM), which provides evidence for real price dispersion in a particular market. The data provide transaction-level prices of semiconductor wafers, along with the key technological attributes of each wafer, facilitating rigorous quality adjustment. The leader in this market is Taiwan, and the follower is China. These data are used to test the prediction that the price difference between these two suppliers narrows after China's entry into a particular wafer market, which typically occurs two years after Taiwan initiates production. On average the price gap 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 5 years. This suggests that the switching-cost mechanism described in the model does contribute to the observed price dispersion across suppliers.

This dispersion across suppliers implies that, if the follower firm captures a share of the product space, its lower price will drive down the aggregate price level. To pursue this point quantitatively, Section 4 uses the model to study equilibrium inflation. This section calculates a benchmark index that is aggregated up from price changes within particular segments of the quality space. In this sense, it gives us a quality-adjusted index, but its construction uses more information than is typically available to the analyst. Therefore, we compare it to two alternative, *feasible* price indexes that can be calculated even when product quality heterogeneity is unobserved. The first feasible index assumes that price dispersion across suppliers results exclusively from quality variation and thus assumes away quality-adjusted price differences across suppliers. The second feasible index does not attempt to adjust for quality but instead compares the average price across suppliers in adjacent periods; this effectively assumes away quality differences across suppliers.

Section 5 concludes.

2. A pricing game with costly switching

This section begins by describing an extension of the simple pricing game in BKM. Our modeling is guided by a large literature that studies price setting in differentiated product markets under costly switching.[1] 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 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 in the sense that it has a disadvantage in producing more complex, higher quality products. Each firm competes to supply an input to overlapping generations of final-goods producers – the consumers, or buyers, in this market.

We use the model to ask two questions. First, in a market with costs of switching suppliers, how does price dispersion across suppliers evolve following the entry of a low-cost supplier? Price dispersion in this model reflects both the cost of switching, which impedes arbitrage, and quality differences across suppliers resulting from the leader's advantage in producing higher quality products. Neither switching costs nor detailed product quality is typically *observable* to the analyst. This means there is no straightforward way in practice to disentangle "real" price dispersion, due to the presence of costly switching, from quality-related price variation. This observation raises our second question: Are there *feasible* price indexes that provide a reasonable approximation to the true quality-adjusted index when product quality heterogeneity is not observed?[2]

[1] See Klemperer (1995) and Farrell and Klemperer (2007) for surveys.

[2] Our approach here has an analogue in the monetary policy literature. A long-studied question in that line of work is, what is the optimal interest rate rule that the Federal Reserve should use to limit fluctuations in output and inflation? But the answer depends unavoidably on one's model of household and firm behavior, to which the Fed's policy must react. Hence, in practice, it is virtually impossible to determine *the* optimal interest rate rule. Moreover, if it were found, it might be too complicated to communicate to private actors. Instead, researchers have looked for simple and *robust* interest rate rules that yield welfare outcomes similar to what is obtained under the optimal rule across a variety of economic models (see Levin, Wieland, and Williams, 2003).

2.1 The model

2.1.1 The basic environment

There are three periods, and 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, and buyers have unit demand and 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, the manufacturing process must be tailored to each buyer's unique design. To preview the example in Section 3, in the market for semiconductor wafers, buyers are designers of integrated circuits. Each buyer purchases a wafer with a particular set of technological characteristics, but each buyer has a unique chip design, and some designs are more difficult to fabricate than others. 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 designs would be unobservable to an econometrician who has data only on broad product characteristics. In this sense, the model nests as a special case the situation where price dispersion merely reflects a constant quality difference across suppliers.

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 that is increasing in design complexity, τy. What enables 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 1 but switches to Firm B in period 2, it must pay s again (independent of its quality). In the semiconductor wafer market, each producer has its own proprietary processes and technologies that are generally incompatible across manufacturers, so investments and calibrations at one producer cannot be reused at another producer. Hence, a buyer that had previously purchased from Firm A pays a price, p_t^A, to continue buying from Firm A in period t, and pays $p_t^B + \tau y + s$ 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 suppliers.

Lastly, to simplify the problem, we rule out hold-up problems by assuming perfect contracting and follow the related literature on switching costs in prohibiting price discrimination. The suitability of these simplifying assumptions will vary from market to market. They are consistent with semiconductor wafer supplier contracts, which i) enumerate measurable requirements for buyers and suppliers, ii) specify sanctions if either party does not meet the specified requirements, and iii) explicitly limit the supplier's freedom in charging appreciably different prices across its customers. Hence, the price p_t^A (p_t^B) applies to all Firm A (B) buyers in period t.

2.1.2 The terminal period problem

The problem is solved by backward induction. To analyze the terminal-period problem, we first conjecture that there is a threshold y_2 such that, in the prior period, Firm B attracts all period-2 entrants with $y \leq y_2 < 1$. In other words, we assume the less efficient 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; 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

$$\begin{aligned}
\sigma_3^{AA}(p_3^A; p_3^B, y_2) &= \Pr[p_3^A \leq p_3^B + \tau y + s \mid y > y_2] \\
\sigma_3^{AB}(p_3^A; p_3^B, y_2) &= \Pr[p_3^A + s \leq p_3^B + \tau y \mid y \leq y_2] \\
\sigma_3^{A0}(p_3^A; p_3^B, y_2) &= \Pr[p_3^A \leq p_3^B + \tau y]
\end{aligned} \tag{1}$$

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

Each of the components of (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 price of its rival, p_3^B, plus the monitoring cost, τy, plus the cost of setting up production with 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 $p_3^B + \tau y$. Lastly, 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 enter in the entrant's decision, since it must pay the setup cost regardless of the supplier from which it sources.

Absent from (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,* the buyer's choice of supplier depends only 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 bind in periods 2 and 3.

Firm A's terminal-period problem may now be stated as follows. From (1), we have that total sales by Firm A in period 3 are given by

$$Y_3^A = \sigma_3^{AA}(p_3^A; p_3^B, y_2)(1 - y_2) + \sigma_3^{A0}(p_3^A; p_3^B, y_2) + \sigma_3^{AB}(p_3^A; p_3^B, y_2)y_2. \tag{2}$$

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 gives $p_3^B(p_3^A, y_2)$. 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. Over a range of low Firm B prices, Firm A will concede all new entrants to its rival. The reason is that it can earn greater profits by setting a higher price and selling exclusively to its partially “locked-in” buyers who have already paid the setup cost with Firm A. As Firm B raises its price, Firm A will at some point be able to profitably compete for new entrants. Although it will have to reduce its price to do so (relative to the reservation price of its customer base), it can attract a sufficiently large share of new entrants so as to make up for the lost profit. Even higher Firm B prices will enable Firm A to poach from Firm B’s customer base, as switching will be cost-minimizing for Firm B’s buyers at this point.

2.1.3 The period-2 period 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 discounted sum of period-2 and 3 prices is less than that the entrant would face if it purchased from Firm B. This implies that y_2 satisfies the indifference relation,

$$\begin{aligned} &p_2^A + \beta \min[P_3^A(y_2), P_3^B(y_2) + \tau y_2 + s] \\ &= p_2^B + \tau y_2 + \beta \min[P_3^A(y_2) + s, P_3^B(y_2) + \tau y_2], \end{aligned} \tag{3}$$

where $\beta < 1$ is the discount factor. This 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 $1 - y_2(p_2^A, p_2^B) \equiv \sigma_2^{A0}(p_2^A, p_2^B)$.[3]

In addition, Firm A begins the period with a customer base. Let y_1 denote the threshold level of quality such 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 in period 2 and then exit. Hence, their problem in period 2 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 captures a market share of the new entrants equal to

[3] 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 (3) cutting the left side from below.

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

Firm A now solves

$$\pi_2^A(p_2^B, y_1) = \max_{p_2^A}\left(p_2^A - c^A\right) \cdot Y_2^A\left(p_2^A, p_2^B;\ y_1\right) + \beta\left(P_3^A\left(y_2\left(p_2^A, p_2^B\right)\right) - c^A\right) \cdot Y_3^A\left(y_2\left(p_2^A, p_2^B\right)\right),$$

subject to period-2 sales

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

and period-3 sales $Y_3^A\left(y_2\left(p_2^A, p_2^B\right)\right)$, given in (2). Firm B solves the analogous problem. We denote the equilibrium prices in the period by $P_2^A(y_1)$ and $P_2^B(y_1)$.

Note that a firm's individual pricing policies in periods 2 and 3 generally exhibit discontinuities. To see why, consider Firm A's problem in the terminal period and suppose that Firm A's strategy were continuous. This would mean its share of period-3 entrants varied continuously as a function of p_3^B. There must then be a Firm B price such that Firm A chooses to attract a share of new (period 3) entrants that only marginally exceeds zero. But this cannot be optimal if $s > 0$. If Firm A were in this position, it could raise its price, discretely increase profits from its old buyers (because they face a discrete cost to switch), and lose only an infinitesimal profit by foregoing sales to entrants. Therefore, profit from "gouging" its customer base must be higher at this point. Accordingly, Firm A will maintain a high price and wait for Firm B's price to climb further. At some sufficiently high Firm B price, it will be optimal for Firm A to reduce its price *discretely* and attract a sufficiently large share of new buyers so as to make up for the lost profits from its locked-in base. This accounts for the "jump" in the pricing rule in Figure 1.[4]

When there are discontinuities in the pricing policies of the two actors, there may exist no pure-strategy equilibrium. In short, one firm's best response can pass through the

[4] Although not shown in Figure 1, there is also a jump at the point where Firm A poaches Firm B buyers. The idea is similar. If the cost to switch is sufficiently large, there will be range of Firm B prices over which Firm A attracts all new entrants but is unable to poach. In this range, Firm A sets a relatively high price to exploit its customers. But when Firm B's price rises sufficiently, it will be optimal for Firm A to reduce its price discretely and acquire such a large share of Firm B buyers as to make up for the lost profits from its other buyers.

"hole" in the other's. Later, we discuss the selection of a calibration that does in fact support an equilibrium.[5]

2.1.4 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 to not participate at all. To solve this cohort's problem, 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 3 entrants into the market. To this end, we assume the payoff, F, to the buyer from its (unit) sale of the final good has the form,

$$F(y) = R + ry, \tag{4}$$

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

Given (4), the buyer's problem in period 1 can be made straightforward. The key simplifying assumption is that ideas are not storable – the buyer must act in period 1 when it obtains the idea for the final good, or else not participate in the market at all. In that case, since it is costless for the buyer to exit in the *next* period, it will participate in period 1 as long as $F(y) > p_1^A$.[6]

This buyer's decision leads to a threshold $y_1(p_1^A)$ such that all entrants with $y \geq y_1(p_1^A)$ participate. Given (4), this threshold is simply $y_1(p_1^A) = (p_1^A - R)/r$, and so the monopolist 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, $(p_1^A - c^A)\left(1 - y_1(p_1^A)\right) + \beta\pi_2^A\left(y_1(p_1^A)\right)$, where $\pi_2^A(y_1) \equiv \pi_2^A(P_2^B(y_1), y_1)$ is the discounted present value of profits as of the start of period 2 conditional on the equilibrium plays in periods 2 and 3.

---

[5] 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 do.

[6] That is, if $F(y) < \min\{p_2^A, p_2^B + \tau y\}$, the buyer can costlessly exit the market altogether. Hence, its problem in period 1 is essentially a static one. The assumption of costless exit is a certainly a strong one, but it keeps this portion of the problem quite tractable.

To be sure, the assumption that the idea is perishable is a stark one. We use it here principally because it delivers a simple period-1 problem. One could appeal to the notion that ideas tend to "leak", that is, another firm is likely to get word of it by the start of period 2 and put it into place. (Input prices will be lower then, due to entry, so production will be profitable.) Nonetheless, a more thorough investigation of the period-1 problem warrants future work.

We now calibrate (4). To ensure participation in periods 2 and 3, R must be sufficiently high. This guarantees that even the lowest-quality buyer ($y = 0$) makes a purchase. In particular, it is sufficient that R exceed $\hat{R} \equiv \max_t\{p_t^A, p_t^B\}$.[7] For our purposes, then, we do not have to point-identify R – it is indeterminate conditional on $R > \hat{R}$. But we note that the monopolist's profits in period 1 are a monotone function of R. Since profits are likely limited in real-life markets by the *threat* of entry, from which we abstract, we set R to be as small as possible in order to contain the profit rate implied by the model.[8]

We then 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 semiconductor 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. We then set r so that y_1 is close to zero (0.05, to be precise); hence, nearly a measure one of buyers participate in period 1. Since an additional measure 1 of buyers enters in period 2, this corresponds to about one-half of the size of the market in the mature phase of the product life cycle.

---

[7] Since that the lowest-quality buyer in Firm A's cohort has payoff $R + ry_t$ in period $t = \{2,3\}$, it follows that Firm A buyers will in fact participate if R exceeds $\max_t\{p_t^A\}$. Furthermore, if $r > \tau$ (as it is, in our calibration), then the lowest-quality customer of Firm B will participate if R is greater than $\max_t\{p_t^B\}$.

[8] In reality, a monopolist would likely price low in period 1 to deter or delay the follower's entry. This is omitted from the present model, which assumes exogenous entry. This omission is perhaps the most important reason why we do not emphasize the *absolute* magnitude of inflation; the absolute drop in prices upon entry is counterfactually large, arguably because this notion of deterrence is absent. Instead, we use the model to compare measured inflation using different feasible measures, with the idea that these results may be robust to the introduction of a richer pre-entry problem.

2.2 Quantitative analysis

We now calibrate and solve the model numerically. There are five parameters, $(\beta, c^A, c^B, s, \tau)$, that have to be chosen. Of these, only the discount factor β can be set without reference to a particular input market. Assuming the period is one year, we set $\beta = 0.95$, which implies an annual real interest rate slightly higher than 5 percent. The remaining parameters will vary across markets. We proceed by calibrating the model to the market for which we have data, namely, the semiconductor wafer fabrication market.

Regarding the costs of production, only the relative price affects the allocation of buyers across suppliers. Therefore, we can normalize $c^B = 1$. We then set $c^A = 1.2$, so the unit cost of production is 20 percent higher at the leader. This choice results in a close match with the empirical leader-follower price gap in the period of entry. Note that we do *not* target the *change* in the price gap across periods; we let the model speak in this regard.

Next, as equation (1) suggests, the extent of switching hinges on the value of s relative to τ. For example, a Firm B customer weighs the cost of switching to Firm A against the premium, τy, that it saves by ending its contract with Firm B. We thus normalize $\tau = 1$ and select s. Two factors guide our choice of s. First, since switching appears to be relatively rare in our application market, s must be sufficiently large. On the other hand, if s is too large, there may be no equilibrium. Intuitively, as s increases, the "jumps" in the policy rule grow larger: at the point where Firm A elects to withdraw from the market for new entrants, it will raise its price discretely to exploit its locked-in buyers. But as these jumps grow larger, it becomes increasingly difficult to find a point of intersection between the two firms' best responses. Our strategy, then, is to set s to be near the highest possible value consistent with the existence of pure-strategy equilibrium, in this case $s = 0.5\tau$.

Table 1 reports the model's predictions regarding the dynamics of price dispersion. There are a few 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 40 log points in the period in which the lagging supplier enters, and a gap of around 20 log points in the next period. These price gaps are quite close to the empirical estimates reported in the next section. For this reason, we believe that our calibrated model, although quite simple, provides a useful perspective on the dynamics of price dispersion in this market.

The source of these findings 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 stronger incentive to compete aggressively for new entrants. The difference in prices between the leader and follower therefore narrows. This result is a simple but important property of the model, as it yields a testable prediction about price dispersion that is *not* replicated by any time-invariant sources of quality (technology or customer service) difference across suppliers.

Second, the average log gap is 16 percent higher than it would be in the frictionless model. To see this, note that the average of the logarithmic price gaps under $s > 0$ is 29. If we set $s = 0$ in our model, one may show that the optimal prices are given by $p^A = \frac{2}{3}(\tau + c^A) + \frac{1}{3}c^B$ and $p^B = \frac{1}{3}(\tau + c^A) + \frac{2}{3}c^B$. These results imply a constant log difference in prices of 25 points. Hence, the gap implied by $s > 0$ is $\frac{4}{25} = 16$ percent higher. Lastly, the calibration of s implies minimal switching. 14 percent of Firm A's customer base (carried over from period 1) switches to Firm B in period 2. No buyer switches away from either supplier in period 3.

3. An Application to the Semiconductor Industry

We believe that the model presented in the previous section captures features common to many intermediate input markets. In prior work (BKM), we have used transaction-level data from the contract semiconductor manufacturing industry to test various predictions of the model. There we show that the model conforms closely with salient features of price behavior in this market that are not predicted by alternative models, suggesting that switching costs are a relevant feature of this intermediate input market. In the remainder of this section, we briefly review these findings.

Semiconductor production involves a number of discrete steps.[9] 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.

We 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 characteristics 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.

Our empirical results use data on arms-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 U.S. 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% of foundry output. In the model, the leader, Firm A, represents Taiwanese producers, while the follower, Firm B, represents Chinese producers. Taiwanese foundries enter a given product market, defined by a wafer-size and line-width combination, at least 8 quarters ahead of Chinese foundries. The dominant Taiwanese foundry, Taiwan Semiconductor Manufacturing Company (TSMC), is the overall market leader, providing the most advanced design integration tools and engineering support, reflected in the model through Firm A's advantage in producing highly complex designs. Semiconductor fabrication also involves many equipment investments and calibrations that

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

are specific to each combination of chip and producer, implying very large costs of moving a given chip from one producer to another, consistent with the model's assumption of large switching costs.

Our data come from a proprietary database collected by the Global Semiconductor Alliance (GSA), a nonprofit industry organization. Our 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 quality-adjusted prices across suppliers located in different countries.

We implement such a comparison in a hedonic regression framework relating wafer prices to observable technological characteristics, quarter indicators, and indicators for supplier's location.[10] By controlling for product characteristics, the hedonic regression allows us to compare average prices across suppliers located in different countries while holding product characteristics constant. Table 2 shows the results of the hedonic regression estimation, reproducing Table 3 of BKM. We find 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.[11] This finding is consistent with the model, given that Chinese producers enter the market well after Taiwanese producers and represent the following Firm B.

Along with substantial average price discounts at the following supplier, the model also predicts closing price gaps between the leading and following suppliers. We examine this process by restricting attention to Chinese and Taiwanese producers and use an event-study style framework to show that Chinese foundries charge less than Taiwanese foundries upon Chinese entry, and that this gap closes as more time elapses following Chinese entry.[12] Figure 2 reproduces Figure 5 of BKM (2012) showing the results of this analysis. In the four quarters following Chinese entry, Chinese producers offer a 39

[10] The GSA data do not provide firm identifiers. They only contain information on country in which the supplier is located.
[11] Since the dependent variable is the log price, we interpret the coefficient as exp(-0.186) – 1 = 17%.
[12] See BKM (2012) Section 5 for details.

percent discount on average compared to Taiwanese producers.[13] As more time elapses, this discount erodes. By the fifth year since Chinese entry, the gap has closed to only 10 percent.

The presence of closing price gaps after follower entry is central to our argument that the switching cost model just described captures the salient features of the semiconductor foundry market. In the absence of this result, one could argue that some unobserved aspect of product quality or customer service was driving the observed price differences across suppliers. Such an omitted factor would likely lead to a) constant price gaps across suppliers over time or b) price gaps that evolve steadily over calendar time but without a systematic pattern following Chinese entry for each technology. Our event-study estimation approach controls for both of these patterns, ensuring that our results are not driven by mechanisms predicting these alternative price dynamics.

The semiconductor wafer fabrication industry closely reflects the central assumptions of the model: sequential entry of suppliers into a given product market with large costs of switching suppliers. The empirical results strongly confirm the model's central predictions of quality-adjusted price discounts at the following supplier and closing price gaps between leader and follower as time elapses after the follower's entry. Although we do not have access to similarly detailed data for other industries, we suspect that these features are present in many other intermediate input markets as well.

4. Measuring Price Inflation with Switching Costs

We have found evidence consistent with real price dispersion in the wafer market. This dispersion across suppliers implies that, if Firm B captures a share of the product space, its lower price will drive down the aggregate price level. To pursue this point quantitatively, we now use the model to study equilibrium inflation.

We present three measures of inflation. The first is an index that assumes that the analyst is able to observe the thresholds, $\{y_t\}$ with $t \in \{1,2,3\}$, that partition the product space. We refer to this as the *benchmark* index. They key to this index is that, since it is endowed with knowledge of the thresholds, it compares prices over time within the same segment of the quality space. The other two indexes may be thought of as approximations to

[13] The differences in Figure 2 are in log points, so exp(-0.494) -1 = 39%.

the benchmark that could be implemented with only information that is typically available to analysts.

We first describe the benchmark index. As noted, this index compares prices over time within quality segment. This means that, to compute the index, we must observe *repeated* purchases within a given quality segment. To illustrate in a simple case, consider the change between periods 1 and 2 and assume $y_2 < y_1$. Then the only qualities for which we see repeated purchases are those sold by Firm A in *both* periods 1 and 2, namely $y \geq y_1$. Therefore, the benchmark price index in this example just corresponds to the change in Firm A's price.[14]

If $y_2 > y_1$, on the other hand, the problem requires more attention. First imagine that Firm A retains all of its period-1 customers, but only sells to new entrants with $y \geq y_2$. In this case, there are, in effect, two goods that comprise the index. There is a "high" good, which consists of qualities $y > y_2$ and which is exclusively supplied by Firm A in each period. And there is a "low" good, which consists of qualities $y \in [y_1, y_2]$. This quality segment is supplied by *both* firms in period 2, as Firm A sells its customer base and Firm B sells to new entrants.

Now (still assuming $y_2 > y_1$) consider the case in which Firm B does poach some of Firm A's customer base in period 2. This is in fact what happens in the calibrated model. Firm B poaches from Firm A for all qualities less than a threshold, y_2^{p}, where the threshold satisfies $y_1 < y_2^{\mathrm{p}} < y_2$. Hence, Firm A supplies to all period-2 entrants $y \geq y_2$ and to all incumbents $y \geq y_2^{\mathrm{p}}$, as illustrated in Figure 3.

In this case, the product space can be partitioned into four regions (or, goods). The first (i) is the quality segment $y \geq y_2$, which is supplied in each period by Firm A. The second (ii) consists of the buyers with $y \in [y_2^{\mathrm{p}}, y_2]$. These buyers purchase from Firm A in period 1. But in period 2, a measure $(y_2 - y_2^{\mathrm{p}})$ of buyers continues to purchase from Firm A (its incumbents), but there is an equal measure of new entrants that purchases from Firm B. Therefore, the *average* price paid by buyers in this segment is the geometric mean,

[14] Throughout, we will compute the baseline price change as the Törnqvist index. But in this example, Firm A is the only supplier to qualities $y \geq y_1$ in each period. Hence, the period-1 and 2 expenditure weights are each just 1.

$p_2 = \sqrt{p_2^A p_2^B}$. The third (iii) is comprised of qualities $y \in [y_1, y_2^{\mathrm{p}}]$, which are supplied exclusively by Firm A in period 1. But in period 2, both new entrants *and* Firm A incumbents in this quality segment chose Firm B. Lastly, there is (iv) the segment, $y \leq y_1$. However, since these qualities did not participate in period 1, this segment is excluded from the index.

We can now calculate the benchmark index as a Törnqvist index,

$$\frac{1}{2}\left[\omega_1^i + \omega_2^i\right] \cdot \ln\left(p_2^A / p_1^A\right) + \frac{1}{2}\left[\omega_1^{ii} + \omega_2^{ii}\right] \cdot \ln\left(p_2 / p_1^A\right) + \frac{1}{2}\left[\omega_1^{iii} + \omega_2^{iii}\right] \cdot \ln\left(p_2^B / p_1^A\right),$$

where (as noted) p_2 is the average price in segment (ii) and the ωs denote the expenditure weights. If we let

$$W_1 = (1 - y_1) p_1^A$$

$$W_2 = \left[2(1 - y_2) + \left(y_2 - y_2^{\mathrm{p}}\right)\right] p_2^A + \left[\left(y_2 - y_2^{\mathrm{p}}\right) + 2\left(y_2^{\mathrm{p}} - y_1\right)\right] p_2^B$$

denote, respectively, total period-1 and 2 expenditure on all qualities $y \geq y_1$, then it follows that the period-1 and 2 weights are

$$\omega_1^i = (1 - y_2) p_1^A \,/\, W_1$$

$$\omega_1^{ii} = \left(y_2 - y_2^{\mathrm{p}}\right) p_1^A \,/\, W_1$$

$$\omega_1^{iii} = \left(y_2^{\mathrm{p}} - y_1\right) p_1^A \,/\, W_1$$

and

$$\omega_2^i = 2(1 - y_2) p_2^A \,/\, W_2$$

$$\omega_2^{ii} = \left(y_2 - y_2^{\mathrm{p}}\right)\left(p_2^A + p_2^B\right) / \, W_2$$

$$\omega_2^{iii} = 2\left(y_2^{\mathrm{p}} - y_1\right) p_2^B \,/\, W_2.$$

We note that the benchmark index uses only *market* prices and revenues as inputs – it does *not* incorporate the monitoring premium paid by Firm B customers. This might be suitable if we need a price index in order to deflate revenue data and recover the quantity of *physical* goods. But such an index is not a comprehensive measure of the price change of the production *service*.[15] In the appendix, we present results for the latter. The distinction is not necessarily minor: the monitoring premium elevates Firm B prices and so mitigates the measured price decline if a quality segment switches from Firm A to B.

Two other inflation measures are worth consideration, since analysts do not typically observe all sources of quality differences. In the first, we follow our understanding of BLS-IPP practice, which is to typically treat the *identity* of the seller as a price-forming characteristic (Diewert and Nakamura, 2009). In this case, there are, in effect, two goods: those sold by Firm A, and those sold by Firm B. Price changes are computed within supplier and then these changes are averaged. We refer to this as the *within* index. This contrasts with the benchmark index, which computes price changes within quality segment regardless of the seller's identity.

Applied to the period-2 data, the within index is very simple. Since Firm B does not participate in period 1, it follows that the within price change is computed only for qualities sold by Firm A in each period. Note that this coincides with our benchmark index only if $y_2 < y_1$, that is, only if Firm B does not sell to buyers in period 2 that would have been supplied by Firm A in period 1.

The second measure takes the opposite approach to the problem of unobserved quality. The strategy here is to average the period-2 prices across both suppliers and compare it to the period-1 price. We refer to this as the *average* index. The concern that guides this approach is that the new supplier is likely to sell the same qualities at lower prices, which does in fact happen when $y_1 < y_2$. In these instances, the within index misses these quality-adjusted price declines occurring due to the entry of a lower price supplier.[16]

---

[15] Note that the use of this broader measure to deflate revenue data allows one to account for inputs *other than* the physical good, namely, the real resources invested in a product-market relationship. Thus, one could argue that this deflator is more appropriate for Solow-type decompositions, as it controls, to some degree, for unobservable inputs.

[16] In period 2, Firm A sells to all new entrants with $y \geq y_2$ and retains incumbents $y_2 \geq y_2^{\mathrm{p}}$, whereas Firm B sells to a measure y_2 of new entrants and poaches Firm A incumbents $y \leq y_2^{\mathrm{p}}$. Hence, the average period-2 price is the weighted geometric mean $[p_2^A]^w [p_2^B]^{1-w}$, where $w \equiv \frac{(1-y_2)+(1-y_2^{\mathrm{p}})}{2-y_1}$.

We also use each of these indexes to compute inflation between periods 2 and 3. To compute the benchmark index, we follow the results from our calibration in which no buyers switch suppliers in the terminal period, and poaching is limited in period 2 in that $y_2^{\mathrm{p}} < y_3$. In particular,

$$y_1 < y_2^{\mathrm{p}} < y_3 < y_2.$$

Figure 4 catalogues the pattern of sales in the two periods. There are, in effect, four goods that comprise the benchmark index: (i) $y > y_2$; (ii) $y \in (y_3, y_2]$; (iii) $y \in \left(y_2^{\mathrm{p}}, y_3\right]$; and (iv) $y \leq y_2^{\mathrm{p}}$. Firm A always sells to the very high-quality buyers ($y \geq y_2$) and Firm B always sells to the very low-quality buyers ($y \leq y_2^{\mathrm{p}}$). The other quality segments deserve a little more attention. Consider (iii) First. In period 2, Firm A retains all incumbents $y \in \left(y_2^{\mathrm{p}}, y_3\right]$, and Firm B attracts entrants in this segment. Hence, we see two prices in this segment in period 2. In the terminal period, Firm B retains all of its customer base, while Firm A only attracts new entrants $y \geq y_3 > y_2^{\mathrm{p}}$. Therefore, in the terminal period, Firm B is the only supplier to this quality segment. As for customers in (ii), they purchase from both firms in each period. This is because Firm B retains its customer base in the terminal period while Firm A attracts entrants in this range.

The benchmark price index may now be computed as

$$\frac{1}{2}\left[\overline{\omega}_2^{i} + \overline{\omega}_3^{i}\right] \cdot \ln\left(p_3^A / p_2^A\right) + \frac{1}{2}\left[\overline{\omega}_2^{ii} + \overline{\omega}_3^{ii}\right] \cdot \ln\left(p_3^{ii} / p_2^{ii}\right)$$

$$+\frac{1}{2}\left[\overline{\omega}_2^{iii} + \overline{\omega}_3^{iii}\right] \cdot \ln\left(p_3^B / p_2^{iii}\right) + \frac{1}{2}\left[\overline{\omega}_2^{iv} + \overline{\omega}_3^{iv}\right] \cdot \ln(p_3^B / p_2^B),$$

where p_2^{ii} is the (geometric) average period-2 price within segment $y \in (y_3, y_2]$; p_3^{ii} is the (geometric) average period-3 price within the same segment; and p_2^{iii} is the (geometric) average period-2 price within segment $\left(y_2^{\mathrm{p}}, y_3\right]$. These averages have to be computed because both firms sell to buyers in these segments in these periods. As for the expenditure weights $\overline{\omega}$s, let

$$W_2 = \left(2(1 - y_2) + \left(y_2 - y_2^{\mathrm{p}}\right)\right) p_2^A + \left(y_2 + y_2^{\mathrm{p}}\right) p_2^B$$

$$W_3 = \left(2(1 - y_2) + (y_2 - y_3)\right) p_3^A + (y_2 + y_3) p_3^B$$

denote total expenditure in each period. Then the period-2 weights are

$$\overline{\omega}_2^{i} = 2(1 - y_2)p_2^A \;/\; W_2$$

$$\overline{\omega}_2^{ii} = (y_2 - \ y_3)\big(p_2^A + p_2^B\big) \;/\; W_2$$

$$\overline{\omega}_2^{iii} = \big(y_3 - y_2^{\mathrm{p}}\big)\big(p_2^A + p_2^B\big) \;/\; W_2$$

$$\overline{\omega}_2^{iv} = y_2^{\mathrm{p}} p_2^B \;/\; W_2$$

and the period-3 weights are

$$\overline{\omega}_3^{i} = 2(1 - y_2)p_3^A \;/\; W_3$$

$$\overline{\omega}_3^{ii} = (y_2 - \ y_3)\big(p_3^A + p_3^B\big) \;/\; W_3$$

$$\overline{\omega}_3^{iii} = \big(y_3 - y_2^{\mathrm{p}}\big)p_3^B \;/\; W_3$$

$$\overline{\omega}_3^{iv} = y_2^{\mathrm{p}} p_3^B \;/\; W_3.$$

To conclude the period-3 calculations, we present the within and average indexes. Recall that the within index calculates prices changes within each supplier and then averages. In other words, from this perspective there are two goods: those sold by Firm A and those sold by Firm B. Letting $Y_2^A = \Big(2(1 - y_2) + \big(y_2 - \ y_2^{\mathrm{p}}\big)\Big)\,p_2^A$ and $Y_3^A = \big(2(1 - y_2) + (y_2 - \ y_3)\big)p_3^A$ denote, respectively, total spending on Firm A goods in periods 2 and 3, the appropriate Törnqvist index is

$$\frac{1}{2}\left[\frac{Y_2^A}{Y_2} + \frac{Y_3^A}{Y_3}\right] \cdot \ln\big(p_3^A / p_2^A\big) + \frac{1}{2}\left[\left(1 - \frac{Y_2^A}{Y_2}\right) + \left(1 - \frac{Y_3^A}{Y_3}\right)\right] \cdot \ln(p_3^B / p_2^B).$$

The average index is given by $\ln\big(\overline{p}_3/\overline{p}_2\big)$ where the average period-t price $\overline{p}_t$ is simply the weighted geometric mean

$$\overline{p}_t = [p_t^A]^{Y_t^A/Y_t} \cdot [p_t^B]^{1-(Y_t^A/Y_t)}.$$

The price index calculations are shown in Table 3, showing the ratio of the two feasible price change measures (the within and average) to the benchmark price change, which requires complete information on the unobserved quality of each product. As one might have anticipated, the within index understates the price decline in the period in which Firm B enters; the within index price decline is about 3/4 of the baseline price change. Firm B's lower price enables it to capture a portion of the product space that had been exclusively supplied by Firm A in period 1. Hence, the average price paid by customers in this quality segment falls notably, but this is not picked up by the within index.

The average index is designed to include this decline, since the average includes Firm B's lower price in period 2. Indeed, the average index *over*-states the decline in price relative to the baseline index. To see why, note that, if we use $p_2 = \sqrt{p_2^A p_2^B}$ and rearrange slightly, the benchmark and average indexes imply period-2 price changes, respectively, of

benchmark: $\frac{1}{2}\left[\omega_1^i + \omega_2^i + \frac{1}{2}\left(\omega_1^{ii} + \omega_2^{ii}\right)\right] \cdot \ln\left(\frac{p_2^A}{p_1^A}\right) + \frac{1}{2}\left[\omega_1^{iii} + \omega_2^{iii} + \frac{1}{2}\left(\omega_1^{ii} + \omega_2^{ii}\right)\right] \cdot \ln\left(\frac{p_2^B}{p_1^A}\right)$

average: $\left(\frac{Y_2^A}{Y_2}\right) \cdot \ln\left(\frac{p_2^A}{p_1^A}\right) + \left(1 - \frac{Y_2^A}{Y_2}\right) \cdot \ln\left(\frac{p_2^B}{p_1^A}\right)$.

Clearly, the difference lies in the weighting. The average index applies a weight to $\ln\left(p_2^A / p_1^A\right)$ equal to the *current* (period-2) market share acquired by Firm A. In contrast, the benchmark index weights according to a combination of the period-1 and 2 market shares in the quality segments that were exposed to the change. The latter is higher, since Firm A was the sole supplier in period 1. This elevates the weight on $\ln\left(p_2^A/p_1^A\right)$ in the benchmark index relative to that in the average index.

In period 3, the within index again commits the larger of the two errors. We suspect the key to this result is that, in the terminal period, Firm B overtakes Firm A as the sole supplier to quality segment $y \in \left[y_2^{\mathrm{p}}, y_3\right]$. As a result, these qualities experience a price decline that is missed by the within index, which captures price declines that occur only within supplier. This leads the within index to over-state inflation more than the average index, since the latter picks up these price changes across suppliers (and within quality).

In interpreting Table 3, we caution that our model may overstate the market share captured by Firm B, relative to the data in the semiconductor wafer market. Our calibration implies that Firm B, by poaching and capturing new entrants, acquires around

40 percent of the market. Although there are a few quarters in which China accounts for 50 percent or more of a given technology's output from China and Taiwan combined, it typically captures no more than 1/3 of a particular wafer market. This likely causes the simulations to exaggerate the inaccuracy of the within index. In the event that Firm B captures little of the market, the within index has to be a good approximation.[17]

5. Conclusion

This paper has investigated a simple model of costly switching in product markets. It stresses two applications. First, models of costly switching with entry can leave a clear imprint on the *dynamics* of price dispersion: the switching cost allows the leader to keep price elevated at first, before lowering it as the firm's original customers exit the market. This simple property of the model provides a means of diagnosing the presence of real price dispersion in markets. Second, if there is a real dispersion, then the entry of the follower firm will lead to quality-adjusted price declines. However, given practical limitations on data, it is often hard to *quantitatively* disentangle the effect of real dispersion from changes in the composition of quality (unobserved heterogeneity). One way to tackle this is to investigate, as we did in Section 4 (and in the appendix) how well alternative feasible price indexes approximate the true price change. Here, we have only illustrated this approach in a model calibrated to a single market, but we believe this is promising avenue of future research across other industries.

[17] We suspect the model overstates the Firm B share because it fails to account for the fact that buyers may wish to source from the leader *now* if that guarantees it access in the *future* to new wafer technologies. Our model misses this because it is confined to a single product and thus omits investment by the leader in new technologies.

6. Appendix: Pricing the manufacturing service

In the main text, we present a benchmark price index that seems suitable for measurement of the price change of the *physical* input. But if we wish to measure the change in the production *service*, we ought to account for the monitoring cost, τy, that a supplier pays to source from Firm B. This will tend to attenuate the within-segment price changes calculated in the main text, since it elevates the price of the service offered by Firm B relative to the price charged by Firm A. As a result, we anticipate this version of the index will diverge more significantly from the average price index.[18]

In regards to the period-2 price change, we calculate this alternative index as

$$\frac{1}{2}\left[\omega_1^i + \omega_2^i\right]\cdot\Pi^i + \frac{1}{2}\left[\omega_1^{ii} + \omega_2^{ii}\right]\cdot\Pi^{ii} + \frac{1}{2}\left[\omega_1^{iii} + \omega_2^{iii}\right]\cdot\Pi^{iii},$$

where the ωs denote the expenditure weights (these are identical to those in the main text) and Π^j is the average of the log price relative in segment j. In segment i, Π^i is the log change in Firm A's price, $\ln\left(p_2^A/p_1^A\right)$, since the latter firm is the exclusive supplier to this segment. The other two segments require a little more attention.

To proceed, it is convenient to begin with segment (iii). Here, Firm A is the period-1 supplier, and Firm B is the period-2 supplier. We wish to evaluate the average log *quality-adjusted* price relative in this segment in that the Firm B price at each quality $y \in \left[y_1, y_2^{\mathrm{p}}\right]$ is taken to be $p_2^B + \tau y$. Hence, over the segment $\left[y_1, y_2^{\mathrm{p}}\right]$, the average log change is

$$\Pi^{iii} = \int_{y_1}^{y_2^{\mathrm{p}}} \ln\left(\frac{p_2^B + \tau y}{p_1^A}\right)\frac{\mathrm{d}y}{y_2^{\mathrm{p}} - y_1} = \frac{\rho_2\left(y_2^{\mathrm{p}}\right)\left(\ln\rho_2\left(y_2^{\mathrm{p}}\right) - 1\right) - \rho_2(y_1)(\ln\rho_2(y_1) - 1)}{\rho_2\left(y_2^{\mathrm{p}}\right) - \rho_2(y_1)} - \ln p_1^A$$

where $\rho_t(y) \equiv p_t^B + \tau y$. The second equality follows from the assumption that y is uniform.

The average log price change in segment (ii) is handled analogously. The only innovation here is that, whereas Firm A was the exclusive supplier in period 1, an equal measure of buyers purchase from Firm A and B in period 2. To accommodate this, we interpret the gross price change at quality $y \in \left[y_2^{\mathrm{p}}, y_2\right]$ as the ratio between the straight

[18] Note that the average and within indexes are unaffected by the considerations in this appendix, as the quality distribution is (deliberately) not used in their construction.

geometric mean $\sqrt{(p_2^B + \tau y) \cdot p_2^A}$ and the Firm A period-1 price. Hence, the average log change within the segment is

$$\Pi^{ii} = \int_{y_2^{\mathrm{p}}}^{y_2} \ln\left(\frac{\sqrt{(p_2^B + \tau y) \cdot p_2^A}}{p_1^A}\right) \frac{\mathrm{d}y}{y_2 - y_2^{\mathrm{p}}}$$

$$= \frac{1}{2}\left\{\frac{\rho_2(y_2)(\ln \rho_2(y_2) - 1) - \rho_2(y_2^{\mathrm{p}})(\ln \rho_2(y_2^{\mathrm{p}}) - 1)}{\rho_2(y_2) - \rho_2(y_2^{\mathrm{p}})} + \ln p_2^A\right\} - \ln p_1^A.$$

Moving to period 3, this alternative index is computed as

$$\frac{1}{2}\left[\overline{\omega}_2^i + \overline{\omega}_3^i\right] \cdot \Pi^i + \frac{1}{2}\left[\overline{\omega}_2^{ii} + \overline{\omega}_3^{ii}\right] \cdot \Pi^{ii}$$

$$+\frac{1}{2}\left[\overline{\omega}_2^{iii} + \overline{\omega}_3^{iii}\right] \cdot \Pi^{iii} + \frac{1}{2}\left[\overline{\omega}_2^{iv} + \overline{\omega}_3^{iv}\right] \cdot \Pi^{iv},$$

where Π^j is the average log price relative in segment j. In segment (i), of course, Π^i is just $\ln(p_3^A/p_2^A)$. In the other segments, we must account for the heterogeneity in monitoring costs. Following the same steps we took to calculate the period-2 price change, the average log price relatives Π^{ii}, Π^{iii}, and Π^{iv} are given by:

$$\Pi^{ii} = \frac{1}{2}\frac{1}{\tau y_2 - \tau y_3}\left\{\begin{array}{c} [\rho_3(y_2)(\ln \rho_3(y_2) - 1) - \rho_3(y_3)(\ln \rho_3(y_3) - 1)] \\ -[\rho_2(y_2)(\ln \rho_2(y_2) - 1) - \rho_2(y_3)(\ln \rho_2(y_3) - 1)] \end{array}\right\} + \frac{1}{2}\ln(p_3^A/p_2^A)$$

$$\Pi^{iii} = \frac{\rho_3(y_3)(\ln \rho_3(y_3) - 1) - \rho_3(y_2^{\mathrm{p}})(\ln \rho_3(y_2^{\mathrm{p}}) - 1)}{\tau y_3 - \tau y_2^{\mathrm{p}}}$$

$$-\frac{1}{2}\left\{\frac{\rho_2(y_3)(\ln \rho_2(y_3) - 1) - \rho_2(y_2^{\mathrm{p}})(\ln \rho_2(y_2^{\mathrm{p}}) - 1)}{\tau y_3 - \tau y_2^{\mathrm{p}}} + \ln p_2^A\right\}$$

$$\Pi^{iv} = \frac{1}{\tau y_2^{\mathrm{p}}}\left\{\begin{array}{c} [\rho_3(y_2^{\mathrm{p}})(\ln \rho_3(y_2^{\mathrm{p}}) - 1) - \rho_3(0)(\ln \rho_3(0) - 1)] \\ -[\rho_2(y_2^{\mathrm{p}})(\ln \rho_2(y_2^{\mathrm{p}}) - 1) - \rho_2(0)(\ln \rho_2(0) - 1)] \end{array}\right\}$$

Table A.1 reports results. We show the ratio of the within and average indexes to this alternative benchmark As we anticipated, the average index now more dramatically overstates the price decline relative to the alternative benchmark. Intuitively, accounting for the monitoring cost attenuates the price change associated with Firm B taking over a quality segment from Firm A. However, it is interesting to note that the average index continues to out-perform in period 3. We are pursuing the source of this latter result; it is not immediately clear to us.

References

Byrne, David, Brian Kovak, and Ryan Michaels (2012). “Price and quality dispersion in an offshoring market: Evidence from Semiconductor Production Services,” mimeo.

Diewert, Erwin and Alice Nakamura (2010). “Bias due to input source substitutions: Can it be measured?”, in *Measurement issues arising from the growth of globalization*.

Farrell, Joseph and Paul Klemperer (2007). “Co-ordination and Lock-in: Competition with Switching Costs and Network Effects”, *Handbook of Industrial Organization*, M. Armstrong and R. Porter (eds) pp. 1967-2072.

Klemperer, Paul (1995). “Competition when Consumers have Switching Costs”, *Review of Economic Studies*, Vol. 62, pp. 515-39.

Levin, Andrew, Volker Wieland, and John C. Williams (2003). "The Performance of Forecast-Based Monetary Policy Rules Under Model Uncertainty," *American Economic Review,* vol. 93, no. 3, pp. 622-645.

Nishimura, Kazuo and James Friedman (1981). “Existence of Nash equilibrium in n person games without quasiconcavity,” *International Economic Review*, Vol. 22 No. 3, pp. 637–648.

Houseman, Susan, Christopher Kurz, Paul Lengemann, and Benjamin Mandel. (2011). “Offshoring Bias in U.S. Manufacturing,” *Journal of Economic Perspectives*, Vol. 25 No. 2, 111-132.

Houseman, Susan, Christopher Kurz, Paul Lengermann, and Mandel Benjamin (2010). “Offshoring Bias in U.S. Manufacturing: Implications for Productivity and Value Added,” Board of Governors of the Federal Reserve System, International Finance Discussion Papers.

Turley, James L. (2003). *The Essential Guide to Semiconductors*, Prentice Hall PTR.

Table 1 – Simulation Results
(Table 4 from BKM)

	$\ln(p^A / p^B)$		Pr(switching)	Firm A market share
	s > 0	s = 0		
Period 2	0.394	0.25	0.14	0.67
Period 3	0.186	0.25	0.00	0.55

Note: The column "s > 0" gives the log difference in prices in the model with a switching cost. The column, "s = 0", relates to the model with no switching cost, in which price dispersion purely reflects differences in product design. The column, "Pr(switching)" gives the share of Firm A's customer base that switches to Firm B. The column "Firm A market share" reports the Firm A's share of the total market (which consists of a continuum of buyers of mass two).

Table 2 – Hedonic Wafer Price Regression
(Table 3 of BKM)

dependent variable: log of price per wafer

Variable	Coefficient	Std. Err.	t-Stat
Foundry Location			
China	-0.186	(0.027)***	-6.87
Malaysia	-0.278	(0.042)***	-6.64
Singapore	-0.061	(0.016)***	-3.77
United States	0.068	(0.030)**	2.23
Wafer Size			
150 mm	-0.467	(0.032)***	-14.54
300 mm	0.671	(0.021)***	31.60
Line Width			
≥ 500 nm	-0.245	(0.053)***	-4.61
350 nm	-0.167	(0.033)***	-4.99
250 nm	-0.061	(0.026)**	-2.34
150 nm	0.169	(0.027)***	6.21
130 nm	0.356	(0.018)***	19.80
90 nm	0.479	(0.032)***	14.77
65 nm	0.676	(0.030)***	22.49
45 nm	0.962	(0.062)***	15.62
Number of Metal Layers	0.076	(0.007)***	10.56
Number of Polysilicon Layers	0.027	(0.024)	1.10
Number of Mask Layers	0.005	(0.002)***	2.98
Epitaxial Layer Indicator	0.064	(0.037)*	1.72
log Number of Wafers Contracted	-0.056	(0.004)***	-13.18
R-squared	0.909		
Observations	6253		

Specification also includes quarterly indicator variables
non-CMOS production not included
Baseline case (omitted category) is Taiwan, 200mm, 180nm
Weighted using iSuppli shipment weights
Standard errors adjusted for 28 quarter clusters
* significant at 10% level, ** 5%, *** 1%

Table 3 - Comparison of Measures of Annual Inflation

Period	Ratio of Within to Benchmark	Ratio of Average to Benchmark
2	0.743	1.125
3	1.350	1.083

Figure 1 – Firm A Best Response in Period 3
(from Figure 4 of BKM)

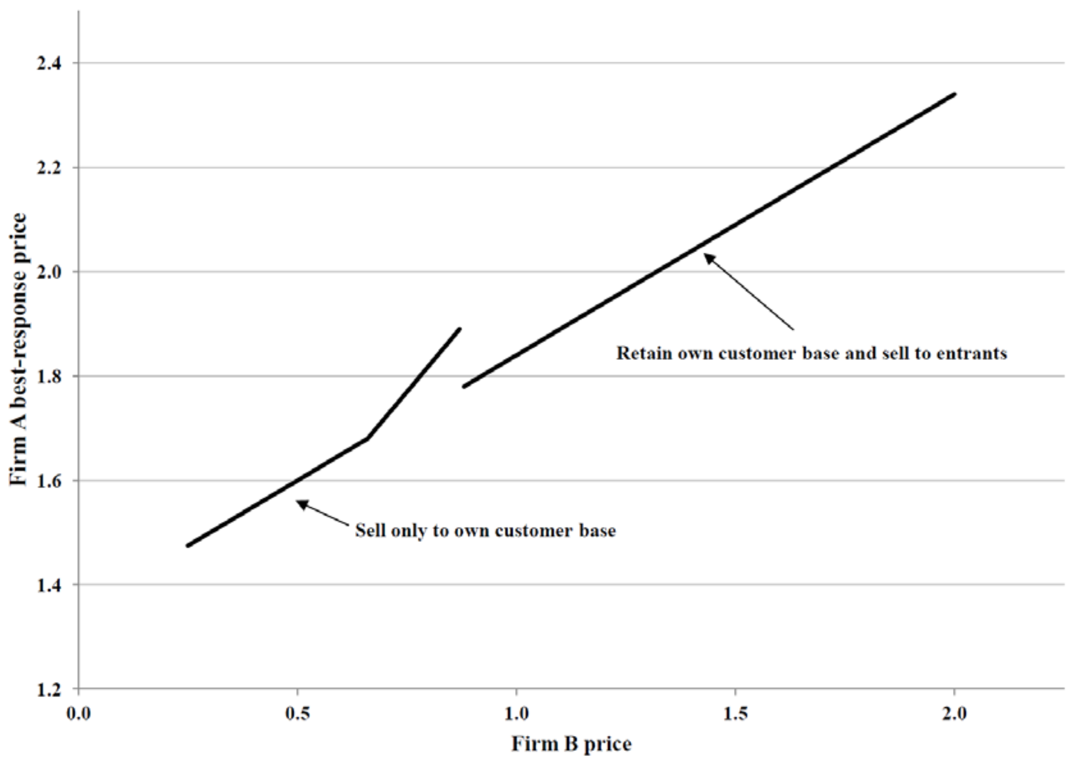

Figure 2 – Closing Price Gaps Within Technology
(Figure 5 of BKM)

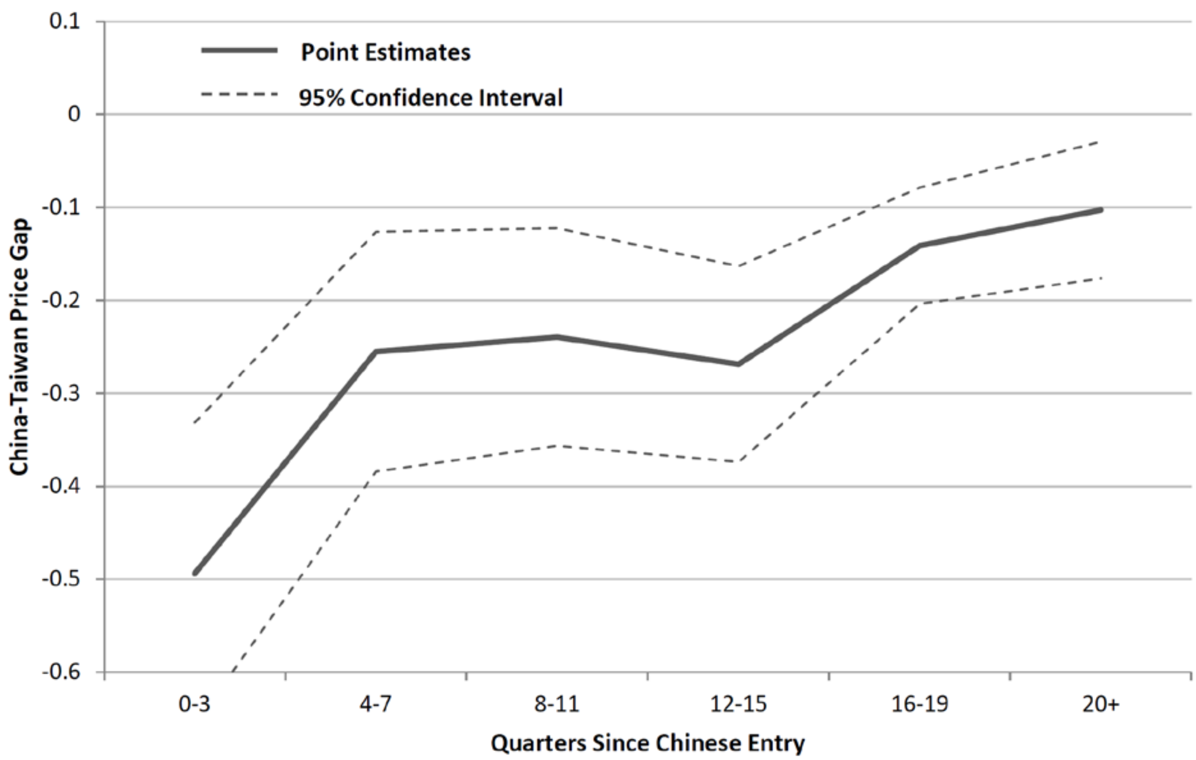

Figure 3 – Period 1-2 Sales Pattern

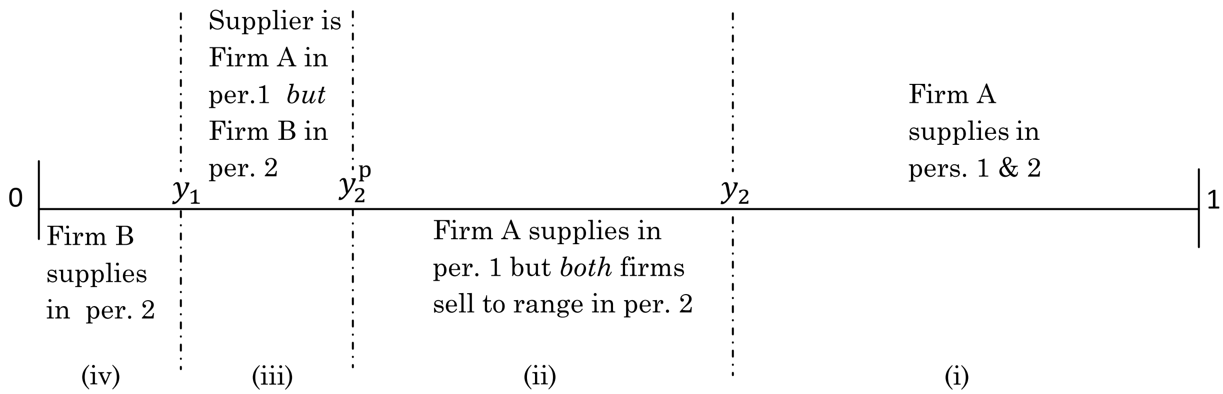

Figure 4 - Period 2-3 Sales Pattern

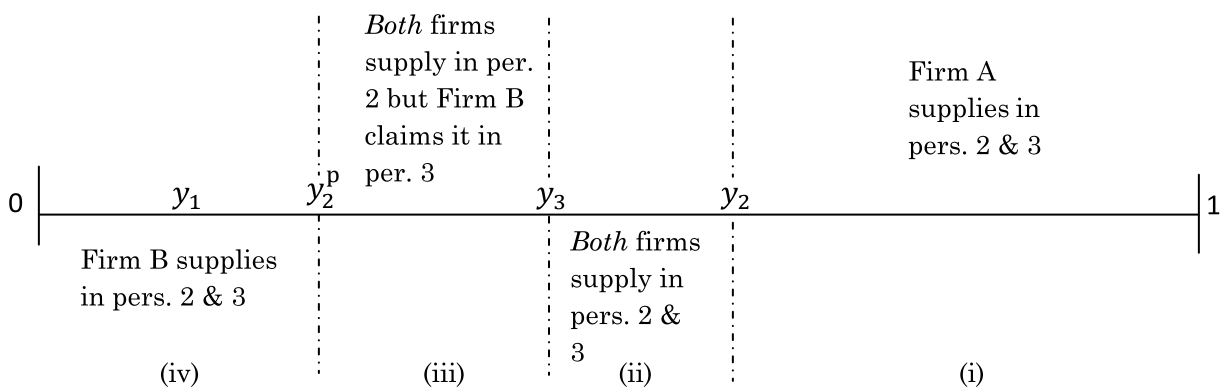

Table A.1 - Comparison of Measures of Annual Inflation

Period	Ratio of Within to Benchmark	Ratio of Average to Benchmark
2	0.89	1.35
3	1.36	1.09