

Preliminary and incomplete.
Please do not cite without permission of the authors.

The Slow Growth of New Plants: Learning about Demand?*

Lucia Foster
Bureau of the Census
Lucia.S.Foster@census.gov

John Haltiwanger
University of Maryland
and NBER
haltiwanger@econ.umd.edu

Chad Syverson
University of Chicago Booth
School of Business and NBER
chad.syverson@chicagogsb.edu

May 31, 2009

Abstract

Many studies using business-level microdata have documented large size average differences across plant ages. New businesses tend to be much smaller than their established industry competitors. This size gap also closes slowly, taking well over a decade on average. We show that even for producers of commodity-like products, these patterns are not driven by productivity gaps. New plants are just as technically efficient as, if not more than, older plants. They are small in spite of their prices, not because of them. The patterns instead appear to be linked to differences in demand-side fundamentals. New plants start with a considerable demand deficit and only slowly erase it over time—if they survive at all. We document patterns in plants’ idiosyncratic demand levels, and explore the sources of their variance across plants and growth rates within them. We estimate a dynamic model of plant expansion in the presence of a “demand accumulation” process (e.g., building a customer base) that allows both passive accumulation over time and active accumulation related to plants’ past production decisions. We find interesting differences in the levels and growth rates of plants depending on the types of firms that own them.

* Preliminary and incomplete. We thank seminar participants at UC-Berkeley, UCLA, Colorado, Harvard, Penn State, Princeton, SITE, and the CAED and CEPR IO meetings for their comments. Contact information: Foster: Center for Economic Studies, Bureau of the Census, Room 2K130F, Washington, DC 20233; Haltiwanger: Department of Economics, University of Maryland, College Park, MD 20742; Syverson: University of Chicago Booth School of Business, 5807 S. Woodlawn Ave., Chicago, IL 60637. This work is unofficial and thus has not undergone the review accorded to official Census Bureau publications. The views expressed in the paper are those of the authors and not necessarily those of the U.S. Census Bureau.

1. Introduction

The large literature using business-level microdata to study various aspects of production behavior has, without exception, found considerable differences between producers in a given industry. Enormous heterogeneity has been documented along many dimensions. One of the more consistent findings is that entrants are different than incumbents, and in particular with regard to size. New businesses tend to start small (e.g., Dunne, Roberts, and Samuelson (1989), Caves (1998)). These patterns are tied to several facets of industry evolution, from industry lifecycle features to the ways individual producers' growth impact industry aggregates.

In this paper, we look more closely at where the size gaps between young and old plants come from. While earlier work has focused on productivity/cost differences as an explanation (see Bahk and Gort (1993), for example), this does not seem likely to be an explanation. We found in Foster, Haltiwanger, and Syverson (2008) that new plants in our sample of producers of commodity-like product are just as technically efficient as—and often even slightly more efficient than—older plants. That is, entrants are small *in spite of* their prices, not because of them. Their prices in fact actually tend to be lower.

This similarity in supply-side fundamentals suggests that idiosyncratic demand factors might explain the well documented plant size differences. Our earlier work documents some evidence of this. There is a clear dichotomy between the age profiles of plants' physical productivity and demand-side fundamentals. While young plants' technical efficiency levels are similar to established plants' levels, they have much lower idiosyncratic demand measures. Moreover, these demand gaps close very slowly over time. Supply side fundamentals show no such slow convergence.

These patterns can be seen in Table 1, which is analogous to Table 5 in Foster, Haltiwanger, and Syverson (2008). It shows the evolution of physical total factor productivity (TFP)—i.e., units of output per unit input—and idiosyncratic demand across plants of various ages. Plant-level demand can be thought of as the output a plant would sell relative to the average plant in the industry, if all plants charged a common, fixed price. (Further details of the construction of the sample and the variables follow below.) We use four age categories. “Entrants” are plants appearing for the first time in the Census of Manufactures (CM).¹

¹ Because the CM includes all manufacturing plants in the U.S., we observe all entry and exit, though only at five-year intervals.

“Young” establishments are those that first appeared in the census prior to the current time period; that is, they were entrants in the previous census. Establishments first appearing two censuses back are “medium” aged, and establishments that first appeared three or more censuses prior are classified as “old.” Plants that will exit (die) by the next CM are placed in their own category. We separately regress plants’ physical TFP and idiosyncratic demand levels on dummies for each age category (old plants are the excluded category). The specification also includes a full set of industry-year fixed effects, so all comparisons are among plants in the same industry in a given year.

The results in the table’s top row indicate that new plants have slightly higher physical TFP levels than established (“old”) incumbents. By the time plants are over five years old, however, this TFP advantage is indistinct from zero. (Incidentally, we also find that exiters of any age are less efficient than incumbents, consistent with the large literature on the subject.)

The patterns are very different for plants’ idiosyncratic demands, shown in the table’s bottom row. The coefficient on the entrant dummy implies that, at the same price, a new plant will sell only 57 percent ($e^{-0.556} = 0.573$)—the demand measure’s units are logged output—of the output of a plant more than 15 years old. This gap is also slow to close. Young plants (five to nine years old) would sell 67 percent of the output of an old plant, and even plants 10-15 years old would only sell 73 percent as much.

We explore the sources of this demand gap and its slow convergence here. Our proposed explanation involves dynamic demand side forces—growth of a customer base or building a reputation, for example—that take considerable time to play out. Caminal and Vives (1999), Radner (2003), and Fishman and Rob (2003) model examples of such processes. These forces lead to gradual growth of an entrant’s “demand stock,” at least among entrants good enough to survive. The uncertainties tied to such processes may also create for the business an option value of waiting to expand until further information about demand is revealed (e.g., Dixit and Pindyck (1994)). It is also likely that the rate of demand stock growth and the level of uncertainty are related to the characteristics of a plant or the firm that owns it.

We are purposefully not too specific about the particular process behind demand-stock growth in our sample. Demand growth may well have multiple sources among our industries. These could include customer learning through “word of mouth” or by the firm’s own advertising efforts, expansion of downstream buyers on either the extensive or intensive margin,

or several other possibilities. We refer to it generically as “learning,” but the building of any sort of relationship capital along buyer-supplier links fits our conceptual framework.² What we seek to do here is characterize the basic mechanics of that generic process and investigate how it interacts with producer behavior.

This paper fits into a new line of research extending the large literature tying productivity to plant and firm survival (Bartelsman and Doms (2000) is an early survey of this literature). This new approach—see Das, Roberts, and Tybout (2007); Eslava et al. (2008); and Foster, Haltiwanger, and Syverson (2008) for other examples—seeks to explicitly account for demand-side effects on plants’ growth and survival. Earlier heterogeneous-productivity industry frameworks (e.g., Jovanovic (1982), Hopenhayn (1992), Melitz (2003), and Asplund and Nocke (2007)) captured differences among industry producers in a single index, often explicitly or implicitly taken to be producer costs/productivity. The new research line expands the sources of heterogeneity to include both technological and demand-based idiosyncratic profitability fundamentals, each following separate (even independent) stochastic processes. The new framework therefore allows an additional and realistic richness in the market forces that determine producers’ fates. Further, this new approach also suggests a reinterpretation of productivity’s effects as inferred from standard measures. This is because typical productivity measures incorporate not just technology but also demand-side shocks through their (often unavoidable because of data limitations) inclusion of producer prices in the output measure.

Our empirical analysis is two-pronged. We first document the evolution of plants’ idiosyncratic demand fundamentals in an atheoretic way. Once these facts are established, we posit a simple dynamic model of producers’ decisions in the face of a dynamic demand process, and estimate the model using producers’ behavior in our data. In the model, producers observe realizations of a stochastic demand process and then choose output levels that in turn feed back into future sales through a demand stock accumulation process. The model generates an Euler equation describing establishments’ incentives for investing in this sort of demand capital. We then use the Euler equation and the demand specification to estimate the model’s parameters.

² Our read of the evidence is that the customer “learning” that drives demand stock growth is much broader than the simple process of buyers finding out about the existence of a producer. While spotty information about mere existence might be consistent with the large gaps in idiosyncratic demand present at plants’ births, it seems unlikely to explain why convergence takes upwards of 15 years. We posit that learning involves much deeper components, like details of producers’ product attributes, the quality and quantity of their bundled services, the consistency of their operations, their expected longevity, and so on. Having to learn about these features can impart considerable inertia into producers’ demand stocks.

The model is informative about the nature and size of the demand growth process and will let us run counterfactual simulations.

Our preliminary results show that while almost all entrants have lower idiosyncratic demand levels than incumbents, the gap is especially large for those owned by small and to some extent less experienced firms. There are some differences in demand convergence rates across plants in different types of firms, but all exhibit limited speed of catching up to their more established competitors. These patterns might in part reflect large or established firms' pre-existing "brand capital" imparting higher initial demand levels and perhaps faster customer base growth on their new plants. The estimates from our dynamic model indicate that the rate of demand-stock building is tied to plants' past activity (sales) levels. Selling more output today serves to shift out demand tomorrow. Interestingly, conditional on plants' sales histories, age itself is actually negatively related demand. That is, there is no passive learning; in the absence of production, a plant's demand will continually shift in over time. This "gravitational" force could be responsible for much of the exit observed among younger businesses.

The paper proceeds as follows. The next section describes data and measurement issues. Section 3 documents basic empirical facts about the evolution of producers' idiosyncratic demands in our sample. Section 4 describes the empirical model that we estimate using plants' dynamic choices. The results are discussed in Section 5, and Section 6 concludes.

2. Data and Measurement Issues

This paper uses the same data set of homogenous goods producers we used in Foster, Haltiwanger, and Syverson (2008). Details on the selection of our sample and construction of the variables we use are in that paper as well as the Appendix, so we only highlight key points here.

The data is an extract of the U.S. Census of Manufactures (CM). The CM covers the universe of manufacturing plants and is conducted quinquennially in years ending in "2" and "7". We use the 1977, 1982, 1987, 1992, and 1997 CMs in our sample based upon the availability and quality of physical output data. Information on plants' production in physical units is important because we must be able to observe plants' output quantities and prices, not just total revenue (which is often the only output measure available in producer microdata). The CM collects information on plants' shipments in dollar value and physical units by seven-digit SIC product

category.³

The 17,669 plant-year observations in the sample include producers of one of eleven products: corrugated and solid fiber boxes (which we will refer to as “boxes” from now on), white pan bread (bread), carbon black, roasted coffee beans (coffee), ready-mixed concrete (concrete), oak flooring (flooring), gasoline, block ice, processed ice, hardwood plywood (plywood), and raw cane sugar (sugar).⁴ These products were chosen because their physical homogeneity. This allows plants’ output quantities and unit prices to be more meaningfully compared.

Note that physical homogeneity does not necessarily imply that producers operate in an undifferentiated product market. Prices vary within industries because, for instance, geographic demand variations or webs of history-laden relationships between particular consumers and producers create producer-specific demand shifts. Further, as we have already shown, quantities sold differ tremendously even holding price fixed. Trying to explain *why* they differ is the very point of our analysis. Our quantity data are meaningful not due to the complete absence of differentiation, but rather because there is no differentiation along the dimension in which we measure output—the physical unit. The notion behind the selection of our sample products is that a consumer should be roughly indifferent between unlabeled units of the industry output. But that does not rule out consumers view as equivalent other products or services (real or perceived) that are tied to those units of output. Much of such differentiation, we argue in our earlier work, is horizontal rather than vertical in nature.

³ A problem with CMs prior to our sample is that it is more difficult to identify balancing product codes (these are used to make sure the sum of the plant’s product-specific shipment values equals the plant’s separately reported total value of shipments). Having reliable product codes is necessary to obtain accurate information on plants’ separate quantities and prices, important inputs into our empirical work below. A related problem is that there are erratic time series patterns in the number of establishments reporting physical quantities, especially in early CMs. We thus choose to focus on the data in 1977 and beyond. However, we do use revenue data from prior censuses as far back as 1963 when constructing plants’ ages and demand stocks Z_{it} .

⁴ Our product definitions are built up from the seven-digit SIC product classification system. Some of our eleven products are the only seven-digit product in their respective four-digit SIC industry, and thus the product defines the industry. This is true of, for example, ready-mixed concrete. Others are single seven-digit products that are parts of industries that make multiple products. Raw cane sugar, for instance, is one seven-digit product produced by the four-digit sugar and confectionary products industry. Finally, some of our eleven products are combinations of seven-digit products within the same four-digit industry. For example, the product we call boxes is actually comprised of roughly ten seven-digit products. In cases where we combine products, we base the decision on our impression of the available physical quantity metric’s ability to capture output variations across the seven-digit products without introducing serious measurement problems due to product differentiation. The exact definition of the eleven products can be found in the Appendix.

2.1. Idiosyncratic Demand: Concept and Measurement

The plant-level idiosyncratic demand measures that we used in Table 1 above and that we will use in our descriptive analysis in the next section are obtained by estimating demand for each of the eleven products in our sample. We describe this process briefly here; again, details can be found in Foster, Haltiwanger, and Syverson (2008).

We begin by estimating the following demand system separately for each of our eleven products:

$$(1) \quad \ln q_{it} = \alpha_0 + \alpha_1 \ln p_{it} + \sum \alpha_t YEAR_t + \alpha_2 \ln(INCOME_m) + \eta_{it},$$

where q_{it} is the physical output of plant i in year t , p_{it} is the plant's price, and η_{it} is a plant-year specific disturbance term. We also control for a set of demand shifters, including a set of year dummies ($YEAR_t$), which adjust for any economy-wide variation in the demand for the product, as well as the average income in the plant's local market m . We define local markets using the Bureau of Economic Analysis' Economic Areas (EAs).⁵

Plant quantities are simply their reported output in physical units. We calculate unit prices for each producer using their reported revenue and physical output.⁶ These prices are then adjusted to a common 1987 basis using the revenue-weighted geometric mean of the product price across all of the plants producing the product in our sample.

Of course, estimating the above equation using ordinary least squares (OLS) methods could lead to positively biased estimates of the price elasticity α_1 . Producers may optimally respond to demand shocks in η_{it} by raising prices, creating a positive correlation between the error term and p_{it} . A solution to this is to instrument for p_{it} using supply-side (cost) influences on prices. While such instruments can sometimes be hard to come by in practice, we believe we have very suitable instruments at hand: namely, plants' physical TFP levels. These embody producers' idiosyncratic technical efficiency levels—their physical production costs. As such, they should have explanatory power over prices. They do. The correlation between plants' physical TFP and prices in our sample is -0.54. Further, it is unlikely they will be correlated

⁵ EAs are collections of counties usually, but not always, centered on Metropolitan Statistical Areas. The 172 EAs that are mutually exclusive and exhaustive of the land area of the United States. See U.S. Bureau of Economic Analysis (1995) for detailed information.

⁶ The reported revenues and physical quantities are annual aggregates, so the unit price is an annual average. This is equivalent to a quantity-weighted average of all transaction prices charged by the plant during the year.

with any short-run plant-specific demand shocks embodied in η_{it} . Hence they appear quite suitable as instruments for plant prices.⁷

The price and income elasticity estimates from the above demand equation are not reported here for space reasons, but are available in Foster, Haltiwanger, and Syverson (2008). The estimates are reassuring about our estimation strategy. All estimated price elasticities are negative, and for all but carbon black, they exceed one in absolute value. This is what one should expect; price-setting producers should be operating in the elastic portion of their demand curves. (Carbon black's inelastic point estimate may be due to the small number of producers of that product in our sample; we cannot in fact reject that carbon black producers face elastic demand.) Further, all products, again except for carbon black, have more elastic IV demand estimates than in the OLS estimations. This is consistent with the theorized simultaneity bias present in the OLS results as well as the ability of TFPQ to instrument for endogenous prices.

The idiosyncratic demand estimates for our sample plants are simply the residual from this IV demand estimation, along with the estimated contribution of local income added back in. Thus the measure essentially captures across-plant output variation that reflects shifts in the demand curve rather than movements along the demand curve.

The dispersion of our producer-specific demand measure is huge. Its within-product-year standard deviation is 1.16 (recall the measure's units are logged output). This implies that a plant sells 3.2 times as much output at a given price as another in its industry that is one standard deviation lower in the idiosyncratic demand distribution. By way of comparison, the comparable standard deviations of logged physical TFP and logged prices are 0.26 and 0.18, respectively.

3. Facts about Plants' Idiosyncratic Demands

In this section, we expand on the exercise done in Table 1 to explore how the relative levels and convergence of idiosyncratic demand levels change with plants' attributes. As briefly

⁷ There are two potential problems with using physical TFP as an instrument. The first is that selection on profitability can lead to a correlation between TFP and demand at the plant level, even if the innovations to both series are orthogonal as assumed. Producers with a higher TFP draws can tolerate lower demand draws (and vice versa) while still remaining profitable. The second potential problem is measurement error. We compute prices by dividing reported revenue by quantity and any measurement error in physical quantities will overstate the negative correlation between prices and physical TFP, potentially contaminating the first stage of the IV estimation. We describe in Foster, Haltiwanger, and Syverson (2008) how we deal with these issues. We found the patterns of demand estimates to be quite robust, reducing concerns about either measurement issue. In Tables 1-3 in the next section, we use the innovation to physical TFP as the instrument since this approach is more consistent with the estimation approach for demand and Euler equations used later in the paper.

mentioned above, a possible source of differences in idiosyncratic demand patterns are the types of firms that own the plants. Dynamic demand effects from customer learning or other similar processes might be impacted by the type and form firms to which plants are tied.

Consider the following example. Two new plants are built in an industry. One is a *de novo* entry by a firm with no prior experience; the other is opened by a large firm with considerable history in the industry (though perhaps in a different local market). We might expect that the latter will enter with a higher idiosyncratic demand, because customers may already be familiar with the plant's product (or at least its firm). This might also impact the speed at which demand convergence occurs.

To begin exploring these possibilities, we again project plants' idiosyncratic demand measures on plant age indicators, but this time interact those indicators with variables tied to characteristics of the firms that own the plants. In the first specification, we will simply allow the age dummies to differ for plants that are part of a multi-plant firm. (The firm's other plants need not make the same product, or even be manufacturers for that matter.) This is a crude proxy for firm size. The second specification will use a series of dummies for the age of *the firm*, defined as the age of the firm's oldest plant. These will be interacted with the plant age dummies. The notion is that plants of older, more established firms may start larger and grow faster than those of newer firms. (Again, the remainder of the firm's activities need not be in the same industry. We can look at product-specific experience in future drafts.)

The results looking at the impact of multi-unit firm status are shown in Table 2. The upper row shows the coefficients on the age categories, the lower those for the age categories interacted with the multi-plant firm indicator. Hence the upper row reflects the evolution of idiosyncratic demand for single-unit plant/firms, while the column-wise sum of the two rows' values shows the same evolution for plants in multi-plant firms (multi-plant firms account for 59 percent of the observations in our sample). Note here that the excluded group is different from that in Table 1. There, it was all old plants—those having first appeared three or more CMs prior, and are therefore at least 15 years old. Here, it is only old plants in single unit firms. Hence the age coefficients show groups' average idiosyncratic demands relative to this group rather than all old plants. Since, as we will see, old plants in multi-plant firms are the largest plants in our sample, their separation from the excluded group will be noticeable.

The results are interesting. Single-unit plants exhibit similar patterns to those seen before. Entrants have considerably smaller idiosyncratic demand levels than do established incumbents; they sell 27 percent less output at a given price than do old single-unit plants, and undersell old multi-unit plants by 58 percent. There is some convergence between entry and being 5-9 years old (“young”), where single-unit plants have demand levels 16 percent below old single-unit plants. But then convergence largely stalls; medium-aged plants still have 14 percent demand deficits.

For plants in multi-plant firms, similar qualitative relationships are present, but their demand levels are higher than single-unit plants at every age. That said, they’re still considerably smaller than old plants in multi-unit firms, with average demand levels that are only two-thirds that of their older counterparts. Convergence is also slow among multi-unit plants. (Interestingly exiting plants in multi-unit firms have lower average demand levels than single-unit exiters. We will see this inversion again below in the interaction with firm age.)

It therefore appears that new plants in small firms (by our crude size measure) face significantly lower idiosyncratic demand levels than do their new competitors in multi-plant firms. Nevertheless, both types of plants see the inertial convergence patterns observed in the broader sample, suggesting demand dynamics are at work in both cases.⁸

Estimating the interactions between firm and plant age yields the results in Table 3. A fully interacted model with four plant and firm age categories each, for both single- and multi-unit firms, would unfortunately create some subsample cells that are too small to be useful for identification and would possibly violate data confidentiality standards. So we pool some categories together. First, we only break out firm age effects for plants in multi-unit firms. Further, we pool young- or medium-aged firms (i.e., whose first plant was observed either one or two CMs prior). Note also that some plant-firm-age categories cannot exist by definition, and are such missing from the estimation. There cannot be a medium-aged plant in an entering or young firm, for example. Old plants in single-unit firms are again the excluded group.

Focusing on the multi-unit plant results in the bottom three rows, we see that among firms that are at least 15 years old, the basic convergence patterns seen before hold here as well.

⁸ Of course, single-unit plants are not restricted to remaining in single-unit firms their entire life, nor for that matter are multi-unit plants restricted to that type of firms. The more common transformation between these is for a plant in a single-unit firm to become part of a multi-unit firm, either through acquisition by another firm or through its own firm acquiring additional plants. From this perspective, the low demand levels and slow convergence of single-unit entrants becomes even starker vis-à-vis their demand levels relative to old plants in multi-unit firms.

Entering plants of old firms have demand levels that are 63 percent of old plants in this type of firm. Growth is slow for the first five years: old firms' young plants have 65 percent of the demand level. Demand growth accelerates after this somewhat, but medium-aged plants still have notably (24 percent) lower demand levels.

For young- and medium-aged firms, we also observe that entrants are smaller than longer-lived plants in such firms (though there can be no old plants in these firms). Notice, too, that plants in young- and medium-aged firms have lower demands than plants of the same age in older firms. The only result that is not in accordance with these general patterns across firm and plant ages involves new plants in new multi-unit firms. While as might be expected their demand levels are smaller than that of old plants in old firms (on average 68 percent of the level), their idiosyncratic demands are higher than new plants in older firms. Another interesting result is that exiting plants in old firms tend to have exceptionally low demand levels—lower, in fact, than new single-unit plants.

The results in Table 3 show there are nontrivial distinctions in the levels and growth of plant demand in firms of different ages. The broadest pattern is one of older firms being tied to higher demand levels at any plant age, just as with firm size again. But also as with the firm-size results above, the demand gaps are still large within any firm type, and these diffuse demands close only slowly over time.

While these results themselves do not uniquely identify an explanation for the patterns seen in the evolution of plants' idiosyncratic demand levels, they are consistent with dynamic demand explanations. We can explore more specific possibilities further in future drafts, looking across a greater array of firm attributes. There are several possible avenues. These include, as mentioned above, whether demand varies not just across firms of different ages, but also with how long the firm has made the plant's particular product. (We observe the product-by-product production history of every plant, making this measurement possible.) When a firm already makes the same product at other locations, consumers' familiarity with the firm's products should make demand spillovers to a new plant more likely, and perhaps lead to faster demand growth as well. This product-level effect may be distinct from the broader firm-level mechanisms hinted at in our Table 3 results above.

Another approach is to explicitly account for geography. Several of our sample products have localized markets (e.g., concrete, ice, and boxes). We can measure whether new producers

are in marginal locations by comparing their sales to more established producers in the same area within the market. Further, if other evidence suggests demand spillovers in customer bases are important, we will be able to see the role that physical proximity to the firms' other plants plays in determining their magnitudes.

4. Model

The analysis above shows various relationships between the attributes of plants and firms and the evolution of producers' idiosyncratic demand levels. The patterns suggest dynamic demand factors are at play—perhaps involved with producers having to build a customer base or reputation, for instance. To address the inherent dynamics more directly, we now pose a model that explicitly builds in a dynamic demand process with both exogenous and endogenous growth components. We will estimate this model using producers' expansion patterns in our dataset to obtain guidance as to the nature of the processes driving demand growth.

We assume the plant faces an isoelastic contemporaneous demand curve:

$$(2) \quad q_t = \theta_t Age_t^\phi Z_t^\gamma p_t^{-\eta},$$

where p_t is the current price charged by the plant. Several factors shift the demand curve. θ_t is an exogenous demand shock that we assume follows an AR(1) process. Age_t is the plant's age. Along with parameter ϕ , this accounts for deterministic changes in plants' demand as they age. Finally, Z_t is a demand shifter that with parameter γ links a plant's current activity to its future expected demand level. Specifically, we assume that Z_t evolves according to the following process:

$$(3) \quad Z_t = (1 - \delta)Z_{t-1} + (1 - \delta)R_t.$$

Thus, Z_t is a sort of operating history of the plant. It grows with past plant sales R_t (defined as $p_t q_t$), subject to depreciation at a rate δ . This process captures dynamic demand processes where a plant's *potential* customer base is related to its past sales activity. For instance, the process embodies many types of “word of mouth” effects consumers are more likely to have heard about a producer or its product if it has operated more in the past. This nests the demand-side analog to the specification common in the supply-side learning-by-doing literature, where learning depends only on cumulative output; i.e., $\delta = 0$. We consider both this and the more general specification in our estimation.

The plant's production function is given by

$$(4) \quad q_t = A_t x_t,$$

where q_t is the plant's output, A_t is its TFP level, and x_t is its input choice. This input can be thought of as a composite of labor, capital, energy, and materials inputs, weighted appropriately. (For example, if the technology is Cobb-Douglas and there are constant returns to scale, the composite would be the plant's inputs raised to their respective input elasticities.)

The plant faces two costs: a factor cost of c_t per unit of x_t and a fixed operating cost of f per period. This, along with the production function, implies the plant's periodic profit function is

$$(5) \quad \pi_t = p_t A_t x_t - c_t x_t - f.$$

Using the demand curve to substitute in for price and simplifying, we have

$$(5a) \quad \pi_t = \theta_t^{\frac{1}{\eta}} Age_t^{\frac{\phi}{\eta}} Z_t^{\frac{\gamma}{\eta}} (A_t x_t)^{1-\frac{1}{\eta}} - c_t x_t - f.$$

The plant manager maximizes the present value of the plant's operating profits.⁹ This problem can be expressed recursively as follows:

$$(6) V(Z_t, A_t, Age_t, \theta_t) = \max \left\{ 0, \sup_{x_t} \theta_t^{\frac{1}{\eta}} Age_t^{\frac{\phi}{\eta}} Z_t^{\frac{\gamma}{\eta}} A_t^{1-\frac{1}{\eta}} x_t^{1-\frac{1}{\eta}} - c_t x_t - f + \beta EV(Z_{t+1}, A_{t+1}, Age_{t+1}, \theta_{t+1}) \right\},$$

where $V(\cdot)$ is the plant's value given state variables. Z is endogenously affected by the plant's input choices; the plant's age, TFP, and demand shock θ_t evolve exogenously. The future is discounted by a factor of $\beta < 1$.

The plant's continuation decision is made explicit in (6): it can operate and earn the profits this entails (the second item in the braces), or it can exit and earn the outside option (normalized to zero here). If it chooses to operate, it takes as given its past operating history as summarized in Z_t and chooses current inputs x_t to maximize its present value. Because of the form of the production function and the demand curve, this choice of x_t simultaneously pins down the plant's output and its price as well.

The dynamics inherent in the plant's choice problem are apparent: by producing more today, the plant can shift out its demand curve tomorrow. The optimal production level

⁹ We abstract from any agency issues that may arise between plants' managers and the owners of these establishments (if they are different people).

(equivalently: the optimal price) in this case will be higher (lower) than that implied by a purely static problem where current price is not tied to future demand. This is consistent with what we found in Foster, Haltiwanger, and Syverson (2008): young plants had lower average prices than older plants in the same industry.

It's important to note that the only sources of dynamics in this model come through the demand process. That means that if other dynamic forces affect plant behavior, it will be interpreted through the lens of our model as demand. It's therefore important that we consider any other such forces and how they might impact the interpretation of our results. We will do this in detail in Section 6 below.

Optimal dynamic behavior (the plant's x_t trajectory) conditional on survival is given by the Euler equation implied by the supremum in (6):

$$(7) \quad \begin{aligned} & \frac{c_t}{(1-\delta)p_t A_t} - \frac{1}{1-\delta} \left(1 - \frac{1}{\eta}\right) \theta_t^\eta \text{Age}_t^\eta Z_t^\eta (A_t x_t)^{-\frac{1}{\eta}} p_t^{-1} \\ &= \beta E \left\{ \theta_{t+1}^\eta \text{Age}_{t+1}^\eta Z_{t+1}^\eta (A_{t+1} x_{t+1})^{-\frac{1}{\eta}} p_{t+1}^{-1} \left[\frac{\gamma}{\eta} \frac{p_{t+1} A_{t+1} x_{t+1}}{Z_{t+1}} - \left(1 - \frac{1}{\eta}\right) \right] + \frac{c_{t+1}}{p_{t+1} A_{t+1}} \right\}. \end{aligned}$$

This expression is slightly unwieldy. Moreover, it includes a state variable θ_t that is observable to the plant manager but unobserved by us.¹⁰ While there are techniques for estimating Euler equations with unobserved state variables, it is preferable to work only with observables.

Fortunately, we can use the demand curve to substitute for the unobservable. We solve (2) for θ_t and substitute the result into (7). This yields, after some algebra,

$$(7a) \quad \frac{c_t}{p_t A_t} - \left(1 - \frac{1}{\eta}\right) = \frac{\beta(1-\delta)\gamma}{\eta} \frac{1}{Z_{t+1}} E[R_{t+1}] + \beta(1-\delta) \left\{ E \left[\frac{c_{t+1}}{p_{t+1} A_{t+1}} \right] - \left(1 - \frac{1}{\eta}\right) \right\}.$$

The intuition behind the plant's optimal dynamic behavior can be seen in this simplified Euler equation. The first term on the left hand side is the inverse of the plant's price-cost ratio (note that the production function implies the plant's marginal cost is c_t/A_t). The second term is a function of the elasticity of demand familiar as the inverse of the optimal markup for a firm facing a residual demand elasticity of $-\eta$. Thus the left hand side of the equation, in a

¹⁰ We observe all the other state variables in our dataset, the Census of Manufactures (CM) microdata. Age, by five-year categories, is available because we have a census of all establishments every fifth year. A , total factor productivity, can be measured from the plant's reported output and inputs. Z , the plant's sales history, can be constructed (given a value for δ) from the plant's sales reported in past CMs.

completely static production/pricing optimization problem, would be zero. It is not generally so here, and that is because of the dynamics discussed above. Since the plant shifts out its demand curve tomorrow by making more sales today, it will want to markup price less over marginal cost than it would in a static world to induce extra sales. (Another way to think about this is that its marginal revenue now isn't just what's implied by the contemporaneous demand function. It also includes the effect on the discounted expected increase in future demand via growth in "demand stock" Z). With a lower markup than implied by the static markup rule, the cost-price ratio in the first term will be larger than the second term, and thus the left hand side generally positive.

The first right-hand-side term is a parameter-dependent constant multiplied by the ratio of the plant's expected next-period revenue and its operating history captured in Z_{t+1} . (Z_{t+1} is not preceded by an expectation operator because it is solely a function of period- t values; see (3).) This term is positive as long as the exogenous impact of age on demand is positive.

The second term on the right hand side is the same markup function as that on the left hand side of the Euler equation, except it is for prices and costs in the next period. Of course, being in the future, it is affected by discounting and the depreciation of Z_t , and it holds in expectation rather than ex-post. Again, this term would be zero in a static setting but is positive here.

When the first right-hand-side term is positive, the second term will be less than the current-period inverse markup expression. The future markup expression in the braces will then be lower than its current-period counterpart—i.e., the price-cost ratio will be higher in the next period. Thus the plant reduces the extent to which it underprices for future demand. As the plant ages, the ratio of its next period expected revenue $E[R_{t+1}]$ to its demand stock Z_{t+1} will fall. This will drive the first right-hand-side term toward zero. In the limit, the only stable dynamic path will be for the plant to set its markup equal to the static rule, setting all terms in the Euler equation equal to zero.

Thus we should expect young plants to have the smallest margins, and then see them rise to the short-term optimum as the plant ages. How fast margins rise is in part a function of the depreciation rate of the plant's demand stock.

4.1. Estimation

We estimate the model's parameters using two complementary mappings from the data to our model. The first is the demand equation (2). This equation alone can be used to estimate all of the model's parameters, actually. But our model also offers the additional structure of the Euler equation (7a). The benefit of using this equation to estimate the model as well is that it explicitly exploits the plant's dynamic choices, using different data variation than the demand equation estimation. As will be clear below, however, not all parameters are identified by the Euler equation alone.

A basic measurement and estimation issue for both the demand and Euler equations is to construct measures of the demand stock, Z . We observe plant revenues in every Census of Manufactures back to 1963, so R_t is directly observable. Past revenues can be used to construct the plant's demand stock Z_t as a function of past sales and the depreciation rate:

$$(3a) \quad Z_t = (1 - \delta)^\tau Z_{t-\tau} + \sum_{i=1}^{\tau} (1 - \delta)^i R_{t-i},$$

where τ is the number of periods the plant has operated.¹¹

Now consider the measurement and econometric issues specific to estimating the demand equation (2). One reason for needing to estimate the demand equation as well is that the effect of age on plant demand ϕ in the simplified Euler equation (7a) is missing. Note that while equation (7)—the version of the Euler equation with the plant's unobservable state variable θ_t —includes all of the model's parameters, equation (7a) is missing ϕ , the effect of age on plant demand. This is because substituting out for θ_t using the demand curve causes the Age_t terms to cancel. However, we can still recover ϕ as well as impose additional structure on the data to estimate the other model parameters by estimating the demand equation.

¹¹ We face a number of practical constraints in the construction of Z_t . The first is that while we are able to trace back plant revenues almost 20 years before our sample begins, several plants—about a third of our sample—had been in existence before that year. Our measures of Z are therefore left-censored for these plants. We test for sensitivity of our results to inclusion of these cases below. A second estimation issue is that we do not observe plant sales in the four years between censuses. Hence we can only build Z stocks using observed revenues. Essentially, we are assuming that sales are constant between censuses and ignoring the impact of depreciation in the intervening years. We expect the fact that the cross-sectional variation in sales swamps intertemporal variation within plants to mitigate this measurement problem. Third, our Euler and demand estimation approach is based on plants observed in years t and $t + 1$ and in an associated fashion plants that have all variables in equations (2b) and (7c) available in periods t and $t + 1$. We deal with this selection issue in the estimates discussed below. A fourth and related measurement issue is how to initialize Z for entrants. In this draft, we finesse this issue since we base the estimation on plants that exist in periods t and $t + 1$ and also have positive lagged revenue in both of those periods. As such, only plants with positive revenue in period $t - 1$ are included in the estimation so that in both t and $t + 1$, Z is well defined based on cumulative lagged revenue.

In considering estimating the demand equation (2), we must address the issue of endogeneity. The RHS variables of (2) include endogenous plant level prices as well as state variables Z_t and Age_t that, in the presence of serially correlated demand shocks, are correlated with the unobserved demand shock. To deal with these issues, we first take logs of (2) which yields:

$$(2a) \quad \ln q_{t+1} = \theta_{t+1} + \phi \ln Age_{t+1} + \gamma \ln Z_{t+1} - \eta \ln p_{t+1}$$

where without loss of generality we have dated the demand equation in $t+1$ to keep the estimated demand equation's timing consistent with the Euler equation. We assume that the unobserved demand shock follows an AR(1) process:

$$(8) \quad \theta_{t+1} = \rho \theta_t + v_{t+1}$$

where v_{t+1} is iid. We then quasi-difference the demand equation (2a) so that we have:

$$(2b) \quad \ln q_{t+1} = \rho \ln q_t + \phi \ln Age_{t+1} - \rho \phi \ln Age_t + \gamma \ln Z_{t+1} - \rho \gamma \ln Z_t - \eta \ln p_{t+1} + \rho \eta \ln p_t + v_{t+1}$$

The residual from the quasi-differenced demand equation (2b), v_{t+1} , is the innovation to the unobserved demand shock. As such it is uncorrelated with variables dated t and earlier and with instruments dated in $t+1$ that are correlated with the RHS variables of (2b) but uncorrelated with the innovation to demand shocks. As discussed (and implemented) in section 2.1, physical productivity is a valid instrument for plant-level prices in the demand equation. We use this instrument here as well.

We note that demand estimation relies on variation (both across plants and within plants over time) in age, past revenues, and cost-driven price shifts for identification. A challenge in the estimation of (2b) is to obtain sufficient variation in the data to identify separately the dynamics of the unobserved demand shock, the role of plant age and the role of learning about demand through experience. It is partly for these identification challenges that we seek to also exploit the variation important for identification of the Euler equation, (7a).

For the estimation of (7a), we note that it can be further simplified by multiplying both the numerator and the denominator of the cost-price ratio by the plant's quantity. Then the ratio becomes the plant's total variable costs as a share of revenue. That is,

$$(7b) \quad \frac{C_t}{R_t} - \left(1 - \frac{1}{\eta}\right) = \frac{\beta(1-\delta)\gamma}{\eta} \frac{1}{Z_{t+1}} E[R_{t+1}] + \beta(1-\delta) \left\{ E\left[\frac{C_{t+1}}{R_{t+1}}\right] - \left(1 - \frac{1}{\eta}\right) \right\},$$

where C_t are total variable costs. Both plants' variable costs and revenues are readily observable in our data.¹² Thus we can observe in our data all of the components of the Euler equation, up to parameters.

To estimate the Euler equation, we assume that the expectation errors are additively separable, and that their mean is zero at the true parameter values. This gives us the moment condition:

$$(7c) \quad E[\varepsilon_{t+1}] = \frac{C_t}{R_t} - \left(1 - \frac{1}{\eta}\right) - \frac{\beta(1-\delta)\gamma}{\eta} \frac{R_{t+1}}{Z_{t+1}} - \beta(1-\delta) \left(\frac{C_{t+1}}{R_{t+1}} - \left(1 - \frac{1}{\eta}\right) \right) = 0.$$

We use this moment condition and instruments that are orthogonal to the expectation error to estimate the model's parameters by GMM. The instruments we use for the Euler equation are variables dated t and earlier. These include lagged cost-revenue ratios, lagged revenues, and age dummies. We note that the Euler equation identifies parameters from changes in plants' variable-costs-to-revenue and revenue-to-demand-stock ratios, different variation from that used to identify the demand equation.

One potentially important econometric issue in estimating both equations (2b) and (7c) is selection. Estimation of each of these equations requires plants that are present in both t and $t+1$ and accompanying measurement of all variables in both equations in t and $t+1$. The five-year mean exit rate for our data sample is around 20 percent, so selection may be empirically important. We also know from our earlier work (Foster, Haltiwanger and Syverson (2008)) that selection is non-random and related to plant-level fundamentals including physical productivity and demand shocks. We found, for example, that a one standard deviation increase in the demand shock (which in this paper reflects the combined influence of unobserved shocks, age and Z effects) decreases the probability of exit by 5 percent. We will therefore control for selection bias in our estimations below.

4.2. Discussion

The comparison between the estimates of ϕ and γ , which respectively parameterize the influence on demand of plant age and past sales, will be informative about the sources of the dynamics of the demand process discussed above. Age captures deterministic demand shifts that would happen regardless of the level of a plant's past activity. We think of this process as

¹² We do have to make an assumption about the size of fixed costs.

“demand accumulation by being.” Z_t , on the other hand, captures the influence of past sales activity, or “demand accumulation by doing.” Models that posit dynamic demand linkages through passive consumer learning imply that the influence of plant age—the simple existence of the plant for a period of time—will be greater. Those emphasizing endogenous demand-stock building—resulting from the active efforts of the plant—will show a large influence of Z_t . We can measure the relative importance of each in the data.

In future versions of the paper, we will allow a plant’s initial demand stock Z_0 to vary across plants in different types of firms, and estimate these Z_0 as parameters. Such specifications would allow new plants in firms that are larger, older, or who are already operating other plants in the same market, to have demand stock spillovers on their new plants that de novo firm/plant entrants would not. We could therefore directly embody the sort of distinguishing demand growth patterns shown in Tables 2 and 3 in the dynamic model.

5. Estimation Results

We jointly estimate via GMM the demand (2b) and Euler (7c) equations. For this draft, we focus on concrete plants only.¹³ In future versions, we will estimate the model on the pooled dataset, but allow the parameters to vary across industries. While we would prefer to estimate a fully interacted model, some of our 11 sample industries simply do not have enough plant-year observations to separately identify their industry’s parameters with any useful precision. (Recall that we need to observe a plant in at least two periods to identify the dynamic parameters, so the number of observations useful for estimation is even smaller than those reported in Table A.1.) We will use some parsimonious form that allows industry-level parameters to vary with a function of the industry’s attributes. Example attributes could include the concentration or turnover rates in the downstream sector, since these could influence the parameters of the demand-building process.

The variables included in our estimated model are defined as described above. However, we make one change in the specification from (2b). We allow the influence of plant age to vary non-parametrically rather than imposing the constant-elasticity form shown in the equation. We do this by including a set of plant age dummies in the estimated version of (2b): a young dummy equal to one if the plant in period t is one census period (i.e., 5-9 years) old, and a medium age

¹³ We don’t estimate β in the Euler equation but rather set it to be consistent with annual discount factor of 0.98.

dummy equal to one if the plant is two census periods (10-14 years) old. The omitted group consists of mature plants at least three census periods (15+ years) old in period t . (Remember that we have no newly born plants in the estimation sample because we need to use lagged variables to identify the dynamic parameters.)

5.1. Estimates of the Model without Depreciation of Demand Stock

We first estimate the model while imposing that the depreciation rate of the demand stock, δ , is zero. The demand stock simply reflects cumulative real revenue in this case. This is the demand-side analog to learning-by-doing models that do not allow for “forgetting” in the style of, e.g., Benkard (2000). The results of the estimation are reported in Table 4. The first numerical column reports results obtained while excluding the 1963 cohort from the sample (and so excluding left censored cases). The second column shows results for the full sample.

First, we note that the estimates of the price elasticity η , which are in the range of -4 to -4.5, are plausible for the industry. These values are close to those estimated using a different methodology and a slightly larger sample than in Foster, Haltiwanger and Syverson (2008).

In terms of the main parameters of interest, the results are consistent with the basic notion of a dynamic demand-accumulation process that we discussed earlier. We find positive and significant effects of “demand accumulation by doing.” Specifically, higher demand stock Z leads to higher demand with an estimated elasticity γ of about 0.65. Thus, as captured in the Euler equation, a plant’s output (or price) choice in the current period affects its marginal revenue not just in the present period, but in the future as well. Producing more today will shift the plant’s demand curve out tomorrow.

In contrast, we find that, having controlled for demand stock, “demand accumulation by being” is actually *negative*. The coefficients on the young and medium dummies are both positive and significant, and the estimated young coefficient is the larger of the two. Since the omitted group is the oldest plants, this means the demand impact of age is monotonically decreasing. In other words, in the absence of production activity, a plant’s demand curve will expectedly shift in. This implies a plant must produce a certain amount of output in a given period just to “tread water” in terms of its dynamic demand accumulation process.

The fact that we see idiosyncratic demands growing over time among surviving plants (as seen in the descriptive results in the earlier sections) indicates that for survivors, the

accumulation by doing process is strong enough to overcome the (dis)accumulation by being process estimated here. A negative impact of age is also consistent where plants unable to sustain sufficient demand growth find their present discounted value of operating less than their outside option, so they exit. (Note how the coefficients for exiting plants in the idiosyncratic demand regressions in Tables 1-3 are negative and have large magnitudes.) In this sense, the age effect embodies a type of “gravity” that is always weighing down a plant’s prospects and must be overcome with sufficient demand growth to stave off exit.¹⁴

Both of these “accumulation by doing” and “accumulation by being” effects are estimated while controlling for the potential presence of serially correlated unobserved demand shocks. We parameterize the persistence of these demand shocks with the five-year AR(1) coefficient ρ , which we estimate to be about 0.45. This five-year persistence rate corresponds to an annual rate of 0.85.

The table also reports the coefficient estimates for two selection controls. As noted above, our estimation sample is selected on survivorship because we need to observe plant activity in both the current and previous periods. To account for possible selection bias, we include the estimated Mills ratio from an (unreported) probit specification on plant survival into the next period. We include this Mills ratio in both the demand and Euler equations, because both could be influenced separately by selection bias. To be able to rely on more than just functional form to identify these selection corrections, we include the plant’s logged capital stock in the survival probit. We showed in Foster, Haltiwanger, and Syverson (2008) that this predicted plant survival. But it’s also excluded from the plant’s dynamic profit maximization problem in our model, making it a candidate excluded instrument for selection.¹⁵

The coefficient estimates on the selection controls suggest that any selection bias is modest in our sample. The estimates are often insignificant. Further, they have little quantitative impact. The values of the other parameter estimates are changed little from their values in specifications, reported in earlier versions of this paper, that exclude the selection controls. We

¹⁴ This is an empirical manifestation of the type of driving force theoretically embodied in the η process of Ericson and Pakes (1995). They postulated this process (which is unrelated to the price elasticity represented by η in our model) as a stochastically negative innovation to a business’s type (e.g., productivity). It served in their model to ensure that every producer eventually exits.

¹⁵ Of course, this raises the issue of what happens if plant capital *is* a dynamic choice in reality, but we have incorrectly excluded this choice from the model. We will address this point in detail in Section 6 below when we consider alternative explanations for our findings.

have also tried more flexible ways of controlling for selection bias, such using various polynomials of the plant's estimated survival probability in place of the standard Mills ratio control. The model's main parameter estimates were robust across these alternative specifications. Hence it seems that selection issues are not skewing our reported estimates.

Comparing the model estimates across the two columns in the table, it appears that excluding the 1963 cohort (with left-censored values of Z) only really affects the coefficients on the age dummies. Other parameter estimates are statistically equivalent across the two samples. The primary effect of including these oldest plants is to increase the average value of Z among the old, excluded age group. Hence to explain the relative levels of demand among the younger plants (which of course have the same absolute value as before), the model increases the value of the younger age dummies relative to the old age dummies. The qualitative points discussed above are consistent across both samples: merely existing is not enough for a business to build its demand stock. In absence of actively making sales, its demand will shift in over time, eventually forcing it to exit.

5.2. Estimates of the Full Model

Table 5 shows the results from the specification that allows for depreciation of the demand stock.

We first discuss the estimates of this depreciation rate, δ . Our earlier imposition of $\delta = 0$ by fiat was apparently not imprudent. Both estimates of the depreciation rate, with and without the 1963 cohort, are statistically indistinguishable from zero and economically small. (One in fact is negative; we do not impose the theoretical condition $\delta > 0$ in the estimation procedure.)

Given that the estimated depreciation rate is near the imposed value in the Table 4 specification, it's not surprising that the other coefficients in Table 5 echo their Table 4 counterparts. The estimated elasticities of current demand to the plant's accumulated demand stock Z are similar, as are the price elasticities and the persistence of demand shocks. Selection does not appear to be a quantitatively important issue in this specification either. The only coefficient that appears to have changed substantially from its Table 4 version is the medium-age-plant dummy in the sample including the 1963 cohort. It is smaller, and insignificantly different from the excluded old plant group. We don't have an explanation for this mild anomaly.

5.3. Discussion

These early results are encouraging in that they indicate we should be able to identify and estimate the main parameters of interest given the available data. The early results are also quite striking in that they suggest that higher demand of mature plants shown in Section 2 is driven by a positive demand-accumulation-by-doing effect that is partly offset by a negative accumulation-by-being effect. Other parameters are largely consistent with findings from other work we and others have done on the sample industry. There is more to do, however, to extend the estimation to the entire sample and to more explicitly incorporate the types of variations in idiosyncratic demand across different plant types (and not just ages) that were seen in Section 2. We discuss some of these future plans in Section 7 below.

6. Alternative Explanations

We anticipate two basic concerns readers might have with our work to this point. The more minor regards whether our idiosyncratic demand measures—the ones used in Sections 1 and 2 to motivate our model—actually reflect a plant’s demand state in a given period rather than something else. The second concern is that we have allowed only one channel for dynamics in our model: demand stock accumulation. If there are other dynamic factors that a plant’s management takes into account when making decisions, we would mistakenly measure their influence as a response to our specified demand dynamics.

We agree that both of these concerns are theoretically valid, and they almost surely have some empirical relevance. However, we believe that the setting of the problem and the way we estimate the model substantially mitigates such concerns. We explain why here.

6.1. What Do Our Idiosyncratic Demand Stock Measures Reflect?

At their narrowest definition, our idiosyncratic demand stock measures reflect the across-plant variation in units of output sold that is, by construction, purged of the effects of plants’ physical production costs. That is, if a plant both sells more output and has a higher idiosyncratic demand measure than another, this means the first plant’s high sales aren’t simply the result of the plant having lower prices because it has low costs. It would sell more than the other plant even at the same price. Regardless of any other measurement issues with these

idiosyncratic demand measures, then, what is always true is they reflect quantities sold that are orthogonal to plants' physical production costs as captured in our TFP measures.

That said, there are other measurement issues that might lead to these demand measures to capture other factors. Primary among these is the issue of capacity utilization. The demand measure is based on the quantity (i.e., the number of units) the plant sells. Our descriptive results could be explained by an alternative story where new plants are built to be the same size (at least in terms of capital) as older plants in their industry, but they look like they have low demand because they are slow to be fully utilized. In this case, firms design plants to be “grown into”; they have the physical infrastructure to handle output levels typical of older incumbents, but are only lightly utilized at first.

We have two responses to this possibility. First, this story is not inconsistent with our theorized demand-accumulation process. New plants may operate at low utilization levels precisely because their demand stock is low. As they accumulate a customer base (or build supplier-consumer relationship capital in one form or another), their output slowly grows to fit the capacity of the plant. *Why* a firm might find it optimal to build an initially oversized plant will depend on the size of capital adjustment costs (more on this below), but our idiosyncratic demand measures could still reflect the demand accumulation process in this case.

Second, the data don't support this sort of capacity utilization pattern. We can't measure capacity utilization directly, but we can construct two good utilization proxies for each plant: the capital-stock-to-output ratio, and the energy-use-to-capital-stock ratio. The former measures whether a plant's physical size is proportional to its reported capital stock. The latter relates a common proxy in the literature for the flow of capital services—energy use—to reported capital stock measures. For capacity utilization to explain the demand patterns discussed above, younger plants would have to have systematically higher capital-to-output levels and lower energy-to-capital ratios than older plants.

Table 6 looks at utilization patterns. The table replicates the specification of Table 2, except using the capacity utilization proxies as the dependent variables (each is used in a separate regression). The results indicate mixed patterns of utilization across plant ages, but even in those cases where utilization moves in the right direction, there is not nearly enough quantitative movement to explain our patterns above. When measured by capital-to-output ratios, as in the top half of the table, utilization is actually higher at younger single-unit plants

than older ones (that is, their $\ln(K/Y)$ rises with age). This pattern is reversed among plants in multi-unit firms, but there the total utilization difference between new and old plants is about 4.5 percent. Thus it can explain only about 10 percent of the measured demand gap. Similar patterns hold, though with less monotonicity over age groups, for the results using energy-capital ratios to measure utilization. Utilization is actually higher for new single-unit plants than old ones and only about five percent lower in the case of new multi-unit plants.

6.2. Other Dynamic Forces

A more serious concern is that the demand accumulation process is the only source of dynamics in our model. If plant decisions are made in response to additional dynamic forces, our estimation will only see such actions through a demand accumulation lens, not the true economic process driving the decisions.

We see three (broadly defined) alternative dynamic factors that our plants might face. The first is a dynamic process in physical productivity—i.e., shifts in A_t over time. The second is financing constraints, and the third is capital adjustment costs. We address each of these possibilities in turn.

Physical productivity dynamics would involve predictable moves in a plant's A_t . Certainly, many have documented that plants experience persistent productivity shocks (see the papers in Bartelsman and Doms (2000), for example). Indeed, a possible source of such movements, though certainly not the only one, would be a traditional learning by doing mechanism. However, the fact that the patterns in the data are not consistent with learning by doing—this is, after all, a basic motivation of our investigation—suggests that physical productivity dynamics are less of a concern in our context.

While individual plants in our sample no doubt experience some persistent, predictable A_t shocks, the results in Table 1 indicate these do not have much of a systematic correlation with plant age. Certainly, they don't seem to hold clear patterns over the 15+ year horizons we are explaining demand movements over. Further, the quantitative movements in physical TFP that do exist are small relative to the demand variation that we focus on here. So while we agree that physical productivity dynamics exist and can play important roles in explaining certain plant-level behaviors, we do not think they are playing a major role in explaining the plant-level choices of the type and horizon that we use to identify the parameters of the demand

accumulation process.

Capital constraints can create dynamics because constrained businesses may accumulate financial capital in one period in order to loosen a constraint on expansion in the future. They would also be a reason for new businesses starting small, since if barriers to obtaining credit exist, it is plausible that new producers would be more likely to face them than would more established businesses.

We unfortunately do not have plant-specific information on credit access or costs of capital, so we cannot directly test for the presence of credit constraints. However, we are able to look at the measured demand levels and growth for different types of firms that might be expected to vary systematically in the extent to which they are credit constrained. The most applicable exercise that we have done in this regard is the breakout of demand patterns for plants of multi-unit firms in Tables 2 and 3. Plants in these larger firms expectedly face lower credit constraints than do single-unit plant/firms. And while these multi-unit plants tend to be larger, they still exhibit the slow convergence in measured demand levels seen among plants of smaller firms. This seems inconsistent with a world where the measured patterns primarily embody financing constraints instead of long-horizon demand accumulation.

Capital adjustment costs—even in the absence of any credit constraints—could produce qualitative patterns similar to those we see in the data. Plants may respond slowly to even long-run demand shocks if it is costly for them to change the size of their business. In such a case, the slow output growth we observe may not reflect gradual demand accumulation, but rather a gradual expansion in the face of persistent high demand.

We expect that capital adjustment costs do play a role in plants' decisions—after all, most capital is not rented via short-term agreements, and there are several potential frictions in capital sales markets. However, the estimates from the literature on the size of capital adjustment costs suggest that quantitatively, they cannot explain the patterns we document.

Even assuming adjustment costs at the high range of estimates, the time it would take for a plant to close the output gap (assuming capital utilization rates are constant over time) observed in Table 1 is relatively short. For example, the estimates in Cooper and Haltiwanger (2006), which were estimated using similar plant-level data to our own sample except on an annual frequency and spanning the entire manufacturing sector, suggest plant size could fully adjust in less than one year. Even some of the larger estimates of capital adjustment costs, like

those in Gilchrist and Himmelberg (1995), suggest the capacity adjustment will occur in only three years.

Hence it seems unlikely that capital adjustment costs could explain all, or even most, of the 15+ years it takes for plants in our sample to close their measured idiosyncratic demand gap. Much as with physical productivity dynamics discussed above, therefore, we expect that while capital adjustment costs are important in some contexts, they just do not have the quantitative impact necessary to explain the long-horizon demand-growth patterns we observe in the data.

7. Future Steps and Conclusion

We plan to extend our preliminary work above in several ways. We briefly describe a partial to-do list here.

We will estimate the dynamic model on the full sample, using data from all 11 industries. Given the inherently different features of each product's final demand market, we would like to allow as much flexibility as possible in the estimated demand parameters. The best way to accomplish this is to simply estimate the model separately for each industry. Unfortunately, we do not have enough data to do so for some of our industries; there are simply not enough plants that we can follow over multiple time periods to gain identification with any precision (see Table A.1 in the appendix). We see a reasonable compromise, however, in making the demand parameters parsimonious functions of industry observables, and letting the data estimate this dependence.

Our descriptive results point to important differences in demand growth depending on the size of the firm that owns the new plant and the firm's previous experience. There are several ways we would like to embed these features in the dynamic model. An obvious step would be to make plants' initial demand stocks Z_0 parameters to be estimated. We can allow Z_0 to vary with firm size and/or experience. Furthermore, we can use several different measures of firm experience besides just the age of the firm. These could include how long the firm has been making the particular product that the new plant sells in its other plants, or how long it has been operating in the plant's local market (or, for that matter, how long it has been making the plant's product in another plant in that market). The data can then tell us which feature is most important for explaining the various initial demand levels we observe across new plants of different types. Similarly, we can also allow the demand accumulation processes embodied in ϕ

and γ vary with firm size or previous experience. Our descriptive results suggest this is an important feature of the data.

While our results are preliminary, at this juncture they hint strongly that even in commodity-like product industries, entry is difficult. It takes a long time for a new business—even those of larger firms—to reach a point where they can expect (at the same price) to sell the same amount of output as do its more established competitors. These results further buttress the new work pointing the importance of idiosyncratic demand factors in explaining the fortunes of businesses.

References

- Asplund, Marcus and Volker Nocke. "Firm Turnover in Imperfectly Competitive Markets." *Review of Economic Studies*, 73(2), 2006, 295-327.
- Bahk, Byong-Hyong and Michael Gort. "Decomposing Learning by Doing in New Plants." *Journal of Political Economy*, 101(4), 1993, 561-583.
- Bartelsman, Eric J. and Mark Doms. "Understanding Productivity: Lessons from Longitudinal Microdata." *Journal of Economic Literature*, 38(3), 2000, 569-95.
- Benkard, C. Lanier. 2000. "Learning and Forgetting: The Dynamics of Aircraft Production." *American Economic Review*, 90(4): 1034-54.
- Caminal, Ramon and Xavier Vives. "Price Dynamics and Consumer Learning." *Journal of Economics and Management Strategy*, 8(1), 1999, 95-131.
- Caves, Richard E. "Industrial Organization and New Findings on the Turnover and Mobility of Firms." *Journal of Economic Literature*, 36(4), 1998, 1947-82.
- Cooper, Russell W. and John C. Haltiwanger. "On the Nature of Capital Adjustment Costs." *Review of Economic Studies*, 73(3), 2006, 611-33.
- Das, Sanghamitra, Mark J. Roberts, and James R. Tybout. "Market Entry Costs, Producer Heterogeneity, and Export Dynamics." *Econometrica*, 75(3), 2007, 837-73.
- Dixit, Avinash K. and Robert S. Pindyck. *Investment under Uncertainty*. Princeton, NJ: Princeton University Press, 1994.
- Dunne, Timothy, Mark J. Roberts, and Larry Samuelson. "Patterns of Firm Entry and Exit in U.S. Manufacturing Industries." *RAND Journal of Economics*, 19(4), 1988, 495-515.
- Ericson, Richard and Ariel Pakes. "Markov-Perfect Industry Dynamics: A Framework for Empirical Work." *Review of Economic Studies*, 62(1), 1995, 53-82.
- Eslava, Marcela, John Haltiwanger, Adriana Kugler and Maurice Kugler. "Plant Survival, Market Fundamentals and Trade Liberalization." Mimeo, 2008.
- Fishman, Arthur and Rafael Rob. "Consumer Intertia, Firm Growth and Industry Dynamics." *Journal of Economic Theory*, 109(1), 2003, 24-38.
- Foster, Lucia, John Haltiwanger, and Chad Syverson. "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review*, 98(1), 2008, 394-425.

- Gilchrist, Simon and Charles P. Himmelberg. "Evidence on the Role of Cash Flow for Investment." *Journal of Monetary Economics*, 36(3), 1995, 541-72.
- Harper, Michael, Ernst R. Berndt, and David O. Wood. "Rates of Return and Capital Aggregation using Alternative Rental Prices," in *Technology and Capital Formation*, Dale W. Jorgenson and Ralph Landau (eds.), MIT Press, 1989.
- Hopenhayn, Hugo. "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica*, 60(5), 1992, 1127-1150.
- Jovanovic, Boyan. "Selection and Evolution of Industry." *Econometrica*, 50(3), 1982, 25-43.
- Melitz, Marc J. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica*, 71(6), 2003, 1695-1725.
- Radner, Roy. "Viscous Demand." *Journal of Economic Theory*, 112(2), 2003, 189-231.
- U.S. Bureau of Economic Analysis. "Redefinition of the BEA Economic Areas." *Survey of Current Business*, February 1995, 75-81.

Table 1. Evolution of Productivity and Demand across Plant Ages

Variable	Plant Age Dummies			
	Entrant	Young	Medium	Exiter
Physical TFP	0.013 (0.005)	0.005 (0.006)	-0.004 (0.006)	-0.019 (0.005)
Demand Shock	-0.556 (0.026)	-0.406 (0.026)	-0.322 (0.026)	-0.342 (0.022)

Note: This table shows the coefficients on indicator variables for exiting, entering, and continuing plants of two age cohorts (shown by column; “young” establishments first appeared in the census five years ago, “medium” establishments first appeared in the census ten years ago) when we regress plant-level productivity and demand levels on these indicators and a full set of product-year fixed effects. The excluded category is plants that appeared three or more censuses prior. The sample includes 17,314 plant-year observations for from the 1977, 82, 87, and 92 Census of Manufactures. Standard errors, clustered by plant, are in parentheses. This table is similar to Table 5 in Foster, Haltiwanger, and Syverson (2008) but uses a measure of demand shock that is more consistent with that used in subsequent exercises.

Table 2. Evolution of Demand across Plant Ages—Interactions with Firm’s Multi-Unit Status

Variable	Plant age dummies				
	Entrant	Young	Medium	Old	Exiter
Demand shock	-0.316 (0.036)	-0.172 (0.035)	-0.149 (0.036)	Excl.	-0.179 (0.032)
Demand shock x multi-unit firm indicator	0.111 (0.039)	0.131 (0.040)	0.245 (0.044)	0.547 (0.035)	-0.293 (0.044)

Note: This table repeats the analysis of Table 1, but now allows plant age effects to vary with the multi-unit status of the plant’s owning firm. The excluded category includes plants in single-unit firms that appeared three or more censuses prior. N = 17,314 plant-year observations. Standard errors, clustered by plant, are in parentheses.

Table 3. Evolution of Demand across Plant Ages—Interactions with Firm's Age

Variable	Plant age dummies				
	Entrant	Young	Medium	Old	Exiter
Demand shock	-0.316 (0.036)	-0.175 (0.034)	-0.144 (0.035)	Excl.	-0.179 (0.032)
Demand shock x firm is multi-unit and an entrant	0.176 (0.067)	N/A	N/A	N/A	-0.179 (0.111)
Demand shock x firm is multi-unit and young or medium	0.007 (0.081)	0.131 (0.050)		N/A	-0.116 (0.084)
Demand shock x firm is multi-unit and old	0.096 (0.044)	0.126 (0.044)	0.279 (0.048)	0.555 (0.035)	-0.342 (0.047)

Note: This table repeats the analysis of Table 1, but now allows plant age effects to vary with the multi-unit status and age of the plant's owning firm. The excluded category includes plants that appeared three or more censuses prior. N = 17,314 plant-year observations. Standard errors, clustered by plant, are in parentheses.

Table 4. Estimated Coefficients for Cumulative Revenue Learning Model (Imposing $\delta = 0$)

Parameter	Estimated Coefficient	
γ	0.642 (0.034)	0.671 (0.025)
η	4.510 (0.366)	3.980 (0.258)
Young dummy	0.443 0.100	0.647 (0.077)
Medium age dummy	0.136 (0.048)	0.257 (0.041)
ρ	0.448 (0.044)	0.451 (0.032)
Inverse Mills Ratio (Demand)	-0.069 (0.026)	-0.026 (0.018)
Inverse Mills Ratio (EE)	0.005 (0.005)	-0.002 (0.004)
Including 1963 cohort?	No	Yes
N	2015	3195

Notes: Joint Demand and Euler Estimation is based on joint estimation of equations (2b) and (7c). Demand equation also includes year dummies (not reported) and control for local demand (local BEA economic area income). Young plants here refer to plants that are present in the current and prior Economic Census. Medium age plants are those that have been present for at least two Economic censuses. The omitted age group is mature plants that have been present for at least three Economic Censuses. The instruments for demand equation include $\log(\text{TFPQ})$, lagged revenues (up to six lags), lagged price, local income, age and year dummies. Instruments for Euler equation include lagged revenue (up to six lags), lagged cost/revenue ratios (up to two lags), lagged price (up to two lags), and age dummies. Standard errors are in parentheses.

Table 5. Estimated Coefficients for Learning with Depreciation Model

Parameter	Estimated Coefficient	
γ	0.592 (0.048)	0.660 (0.034)
η	5.000 (0.562)	4.190 (0.360)
Young dummy	0.419 (0.108)	0.619 (0.079)
Medium age dummy	0.132 (0.051)	0.025 (0.042)
ρ	0.470 (0.045)	0.444 (0.033)
δ	-0.048 (0.032)	0.015 (0.022)
Inverse Mills Ratio (Demand)	-0.075 (0.027)	-0.031 (0.019)
Inverse Mills Ratio (EE)	0.004 (0.006)	0.001 (0.004)
Including 1963 cohort?	No	Yes

Notes: Joint Demand and Euler Estimation is based on joint estimation of equations (2b) and (7c). Demand equation also includes year dummies (not reported) and control for local demand (local BEA economic area income). Young plants here refer to plants that are present in the current and prior Economic Census. Medium age plants are those that have been present for at least two Economic censuses. The omitted age group is mature plants that have been present for at least three Economic Censuses. The instruments for demand equation include $\log(\text{TFPQ})$, lagged revenues (up to six lags), lagged price, local income, age and year dummies. Instruments for Euler equation include lagged revenue (up to six lags), lagged cost/revenue ratios (up to two lags), lagged price (up to two lags), and age dummies. Standard errors are in parentheses.

Table 6. Capacity Utilization Patterns Across Plant Ages and Multi-Unit Status

Capacity Utilization Measure	Variable	Plant age dummies				
		Entrant	Young	Medium	Old	Exiter
ln(K/Y)	Utilization	-0.110 (0.019)	-0.086 (0.020)	-0.042 (0.022)	Omitted	-0.022 (0.018)
	Utilization x multi-unit firm	-0.077 (0.021)	-0.087 (0.023)	-0.099 (0.025)	-0.111 (0.015)	0.055 (0.024)
ln(E/K _{eq})	Utilization	0.090 (0.029)	-0.021 (0.031)	-0.057 (0.033)	Omitted	0.074 (0.027)
	Utilization x multi-unit firm	-0.070 (0.033)	0.012 (0.035)	-0.003 (0.039)	-0.022 (0.023)	-0.017 (0.036)

Note: This table estimates the same specification as Table 2, except now uses as the dependent variable two different plant-level proxies for capacity utilization, hence showing patterns of plant utilization over age and plant multi-unit status. The two proxies are the log of the capital stock to output ratio, and the log of energy use to equipment capital ratio. N = 17,314 plant-year observations. Standard errors are in parentheses.

Appendix

A.1. Defining Our Products

As background to how we define our products, it is first necessary to understand the product coding scheme that Census uses. There are three types of codes that we highlight. First, Census codes flags products from administrative records (AR) sources. We exclude all of these AR products from our analysis. (Including in our measures of PPSR since it is obviously not possible to assign these AR products to a single 7-digit code.) Second, Census uses balancing codes to correct cases in which the sum of the total value of shipments of reported individual products does not sum to the reported total value of shipments. Census identified these balancing codes using special suffixes for the product codes in every census year except in 1987. Where balancing codes are identified, they have been deleted. Finally, Census collects data on receipts for contract work, miscellaneous receipts, and resales of products. These products are excluded from our calculations of PPSR (again, because it is obviously not possible to assign these AR products to a single 7-digit code). As a final exclusion, we did not include any products in that have a negative value since these are presumably balancing codes. The precise definitions of our eleven products are listed below (with 7-digit product codes in parentheses).

Boxes is defined as the sum of boxes classified by their end use and boxes classified by their materials. Boxes classified by end use are: food and beverages (2653012), paper and allied products (2653013), carryout boxes for retail food (2653014 category starts in 1987) glass, clay, and stone products (2653015), metal products, machinery, equipment, and supplies except electrical (2653016), electrical machinery, equipment, supplies, and appliances (2653018), chemicals and drugs, including paints, varnishes, cosmetics, and soap (2653021), lumber and wood products, including furniture (2653029), and all other ends uses not specified above (2653029 in 1977 and 1982, 2653030 in 1987). Boxes classified by their materials are: solid fiber (2653051), corrugated paperboard in sheets and rolls, lined and unlined (2653067), and corrugated and solid fiber pallets, pads and partitions (2653068). The physical data for boxes is measured in short tons.

Bread is defined as one 7-digit product, white pan bread (2051111), until 1992 when it was split into two products white pan bread, except frozen (2051121) and frozen white pan bread (2051122). The physical data for bread is measured in thousands of pounds.

Carbon Black is defined as one 7-digit product, carbon black (2895011 in 1977, 2895000 thereafter). The physical data for carbon black is measured in thousands of pounds.

Coffee is the sum of whole bean (2095111), ground and extended yield (2095117 and 2095118 in 1982 and 2095115 thereafter), and ground coffee mixtures (2095121). The physical data for coffee is measured in thousands of pounds.

Concrete is defined as one 7-digit product, ready-mix concrete (3273000), over our entire sample. Some of the products coded as 327300 in 1987 were in fact census balancing codes and thus were deleted from our sample. The physical data for concrete is measured in thousands of cubic yards.

Flooring is defined as one 7-digit product, hardwood oak flooring (2426111), over our entire sample. The physical data for flooring is measured in thousands of board feet.

Gasoline is defined as one 7-digit product, motor gasoline (2911131), over our entire sample. The physical data for gasoline is measured in thousands of barrels.

Block Ice is defined as one 7-digit product, can or block ice (2097011), over our entire sample. The physical data for block ice is measured in short tons.

Processed Ice is defined as one 7-digit product, cubed, crushed, or other processed ice (2097051), over our entire sample. The physical data for processed ice is measured in short tons.

Plywood is defined as one 7-digit product, hardwood plywood (2435100), over 1977-1987. Starting in 1992, plywood is the sum of veneer core (2435101), particleboard core (2435105), medium density fiberboard core

(2435107), and other core (2435147). The physical data for plywood is measured in thousands of square feet surface measure.

Sugar is defined as one 7-digit product, raw cane sugar (2061011), over our entire sample. The physical data for sugar is measured in short tons.

A.2. Measurement of input levels and input elasticities in the TFP indexes.

This section reports details on the measurement of input levels and elasticities in the TFP measures described in Section 3.

Labor inputs are measured as plants' reported production-worker hours adjusted using the method of Baily, Hulten and Campbell (1992). This involves multiplying the production-worker hours by the ratio of total payroll to payroll for production workers. Prior work has shown this measure to be highly correlated with Davis and Haltiwanger's (1991) more direct imputation of nonproduction workers, which multiplies a plant's number of nonproduction workers by the average annual hours for nonproduction workers in the corresponding two-digit industry calculated from the CPS. Capital inputs are plants' reported book values for their structure and equipment capital stocks deflated to 1987 levels using sector-specific deflators from the Bureau of Economic Analysis. The method is detailed in Foster, Haltiwanger and Krizan (2001). Materials and energy inputs are simply plants' reported expenditures on each deflated using the corresponding input price indices from the NBER Productivity Database.

To compute the industry-level cost shares that we use to measure the input elasticities α_j , we use the materials and energy expenditures along with payments to labor to measure the costs of these three inputs. We construct the cost of capital by multiplying real capital stock value by the capital rental rates for the plant's respective two-digit industry. These rental rates are from unpublished data constructed and used by the Bureau of Labor Statistics in computing their Multifactor Productivity series. Formulas, related methodology, and data sources are described in U.S. Bureau of Labor Statistics (1983) and Harper, Berndt, and Wood (1989).

A.3. Rules for Inclusion in the Sample

While the Economic Census data we use is very rich, it still has limitations that make necessary three restrictions on the set of producers included in our sample. First, we exclude plants in a small number of product-years for which physical output data are not available due to Census decisions to not collect it or obvious recording problems. Second, we exclude establishments whose production information appears to be imputed (imputes are not always identifiable in the CM) or suffering from gross reporting errors. Third, we impose a product specialization criterion: a plant must obtain at least 50% of its revenue from sales of our product of interest. This restriction reduces measurement problems in computing physical TFP. Because plants' factor inputs are not reported separately by product but rather at the plant level, we must for multi-product plants apportion the share of inputs used to make our product of interest. Operationally, we make this adjustment by dividing the plant's reported output of the product of interest by that product's share of plant sales. This restriction is not very binding in seven of our products whose establishments are on average quite specialized. Bread, flooring, gasoline, and block ice producers are less specialized, however, so care must be taken in interpreting our sample as being representative of all producers of those products. We test below the sensitivity of our results to the inclusion of less specialized producers.¹⁶ Characteristics of the final sample can be seen in Table A.1.

Census reports physical product data for only a subset of the 11,000 products reported in the Census of Manufactures. While we use only products for which physical output is reported, the collection of this data has changed over time for two of our products. (See Table A.1.) Census did not collect physical output for ready-mix concrete in 1997, and the unit of measurement for boxes changed over our sample period in a way that makes the

¹⁶ This input-adjustment method in effect assumes inputs are used proportionately to each product's revenue share. For example, a plant producing 1000 cubic yards of ready-mixed concrete accounting for 80% of its shipment revenues will have the same physical TFP value as a completely specialized plant producing 1250 cubic yards of concrete, assuming they employ the same measured inputs. Without adjusting the output, the first plant would appear less productive because the inputs it uses its other products would be instead attributed entirely to ready-mixed production. The average share of our sample plants' values of shipments accounted for by the corresponding product is given in parentheses: boxes (93), bread (39), carbon black (96), coffee (86), concrete (92), flooring (46), gasoline (49), block ice (37), processed ice (76), plywood (64), and sugar (90).

1992 and 1997 data incomparable to the earlier periods. Additionally, there are recording flaws in the 1992 quantity data for processed ice that make using it unfeasible.

The Census Bureau relies on administrative record data for very small establishments (typically with less than five employees). In these cases all production data except total revenues and the number of employees are imputed, and production operations are classified only up to the four-digit industry level. Since our unit of analysis is more detailed than the four-digit industry, we cannot determine whether a particular administrative record establishment actually produces the product of interest. For these reasons, we exclude administrative records cases from our sample. While about one-third of CM establishments are administrative records, their output and employment shares are much less because they are such small plants.

We also exclude establishments whose data appear to be imputed or suffer from reporting or recording errors. The Census Bureau imputes physical quantities when product-level data are not fully reported. Unfortunately, imputed data are not explicitly identified. To distinguish and remove imputed product-level data from the sample, we use techniques similar to those employed by Roberts and Supina (1996, 2000). To minimize the influence of reporting and recording errors, we also remove a small number of plants reporting physical quantities that imply prices greater than ten times or less than one-tenth the median price in a given year. In order to maintain the same sample over all exercises, we delete observations that are missing any one of the main regression variables. We also delete observations when the plant's labor or materials cost share is less than one-tenth of the corresponding industry's average cost share for that year, or when the cost share is more than one. Finally, we still find a relatively small number of obvious outliers in physical quantity measures, so we trim the one-percent tails of the physical productivity (TFPQ) distribution.

Our product specialization criterion requires that plants obtain at least 50% of their revenue from our product of interest. The text discusses the measurement reasons for imposing this restriction as well as describing a robustness check with respect to this product specialization cutoff.

A.4. Characteristics of Establishments by Product

In this section we briefly characterize some of the relevant properties of the establishments that produce our products. Table A.1 shows characteristics of the sample by product. The first five columns show the number of establishments in our sample by year for each product. The second to last column shows the real revenue shares of each product. Real revenue is the weight used in our weighted regressions. Concrete clearly dominates our sample in terms of the number of establishments while gasoline dominates in terms of the revenue share. The table's last column shows mean logged income (income is taken from Census reports for the county in which the plant is located) for each product in our sample. Concrete has the highest mean log income while carbon black has the lowest.

Table A.2 shows the entry and exit rates by product for the data pooled over all available years. Entry rates range from a low of 3.9 for sugar to a high of 26.6 for concrete, while exit rates range from a low of 9.0 for gasoline and to a high of 27.7 for processed ice. Some products appear to be in a period of retrenchment or consolidation. Sugar for example, has a very low entry rate (3.9) but a high exit rate (17.0). The number of plants in the sugar and confectionary products industry (SIC 2061) has fallen from 66 in 1977 to 39 in 1997. Other products appear to simply have a high degree of churning. For example, concrete and both types of ice products all have entry rates and exit rates that exceed 20 percent. The number of establishments in ready-mixed concrete (SIC 3273) industry increases over our sample period, while the number of establishments in the block and processed ice industry (SIC 2097) falls somewhat over our sample, from 675 establishments in 1977 to 582 establishments in 1997.

A.5. Robustness of Demand Estimates

There are two potential problems with using TFPQ as an instrument. We have explored both and found our results to be robust. We briefly describe these potential issues here.

The first regards how selection on profitability impacts the assumption that TFPQ is uncorrelated with demand shocks at the plant level. This assumption, which is necessary for TFPQ to be a valid instrument in the demand estimation, strikes us as quite reasonable for *innovations* to the dynamic plant-level TFPQ and demand processes. Hence for entering plants, which get their initial idiosyncratic productivity and demand levels as draws from the respective innovation processes, the assumption of no correlation is likely to hold. However, TFPQ and demand levels may be correlated for *continuing* plants (those operating in both the current and previous periods) because of selection. Here is why. A producer continues operation if it is profitable, which depends on idiosyncratic productivity and demand. Since a producer with a higher TFPQ draw can tolerate lower demand draws (and vice versa) while remaining profitable, those producers that chose in the previous period to continue operations into the current period—i.e., the continuing plants—will tend to have negatively correlated lagged TFPQ

and lagged demand levels. The correlation arises in the lagged values because the plants' decisions to continue operating into the current period were conditioned on that information. But because both of these processes are persistent, their current values could be correlated as well, making the orthogonality assumption necessary for consistent demand elasticity estimation questionable.

To explore the sensitivity of our results to this issue, we estimated product demand curves (10) using an alternative instrument for price that is based only upon innovations to TFPQ. For entrants, this was simply the observed level of TFPQ (after taking out product-year fixed effects). For continuers, on the other hand, we regressed the plants' current TFPQ levels on their lagged values (again including product-year fixed effects), and used the residual from this to instrument for price. In this way, we use the innovation to producers' productivity levels—information that was plausibly not included in their decision in the previous period to continue operations, but that should affect current costs—to gain exogenous variation in current prices. (The fact that we only observe plants every five years makes necessary some assumptions of timing here. The first is that the period when plants have received their idiosyncratic profitability draws but have yet to decide whether to exit, which is instantaneous in our two-stage model, can in reality actually correspond to a period of a few years. The entering plants in our sample are assumed to be in this stage. The second is that the TFPQ innovation does not impact producers' exit decisions in intercensal years.)

When we use this alternative TFPQ-innovation-based instrument, we find that the pattern of elasticities and even more importantly the pattern of demand shocks to be very robust. We also find the results that depend upon demand shocks in Tables 3-6 to be virtually unaffected. These results are available upon request.

A second potential problem is measurement error. We compute prices by dividing reported revenue by quantity and any measurement error in physical quantities will overstate the negative correlation between prices and TFPQ. Since the first stage of the IV estimation regresses plants' prices on their TFP levels, measurement error would yield biased estimates of the fitted prices used in the second stage, possibly leading in turn to biased price elasticities and idiosyncratic demand measures. To address this concern, we estimated a specification that should be robust to measurement error, implemented as follows. Rather than using current plant TFP directly as an instrument, we use a "fitted" value constructed as a projection of current TFP on several values expected to be correlated with plant productivity but orthogonal to measurement error. These include the plant's lagged TFP, lead TFP, the average TFP of the *other* plants in the same industry owned by the plant's firm (these three variables are used as available, depending on the plant), a set of birth-cohort-by-industry dummies, and a set of "survival-cohort"-by-industry dummies denoting how many more CMs a plant exists before exiting (or if it survives until the end of our sample).

All of these variables should have predictive power over a plant's current TFP level but be orthogonal to classical measurement error in quantity. Lag and lead TFP values are relevant to current TFP because of productivity persistence, the average TFP of other plants owned by the same firm because of management-driven productivity spillovers or assortative matching of plants, and the birth and survival cohort indicators because of age-specific productivity evolutions and because future survival is determined by current productivity. Second, besides this fitted TFP value, we also include the local average wage as an instrument in the first stage. This factor price is of course an additional cost shifter for the plant.

As with the other robustness check above, the results of this alternative demand estimation procedure yield demand estimates that closely track those obtained in our benchmark specification. Moreover, the subsequent results using the estimated demand shocks from this alternative estimation are very similar to those reported in the paper. Again, we do not report these here for space reasons.

In considering these two robustness checks, we note that they employ essentially opposite strategies: one uses the transitory variation in TFPQ while the other uses the permanent component. We are reassured by the fact that the results are robust to using either component of TFPQ variation. If either potential problem was a major driver of our results, we would expect the demand estimates to be quite sensitive to these distinctions.

Table A.1: Characteristics of the Sample by Product

Product	Number of Observations					Real Revenue Share (%)	Mean (log) Income
	1977	1982	1987	1992	1997		
Boxes	936	905	1045	NA	NA	7.9	17.4
Bread	195	142	110	92	92	2.4	17.0
Carbon Black	31	23	22	21	18	0.7	16.2
Coffee	61	84	79	77	77	4.7	18.0
Concrete	2184	3316	3236	3427	NA	7.0	17.1
Hardwood Flooring	8	10	16	25	24	0.2	16.7
Gasoline	99	99	94	95	71	75.0	17.1
Block Ice	40	43	26	23	10	0.0	16.9
Processed Ice	87	155	144	NA	NA	0.1	16.8
Plywood	71	68	42	42	37	0.6	16.5
Sugar	40	36	30	35	26	1.3	16.6

Note: This table shows the number of establishments in our sample by product and year, as well as each product's share of total real revenue in the sample (pooled across all years). Mean log income, used in our demand estimation procedure, is for plants' corresponding Economic Areas (see text for details) based on data pooled over all years.

Table A.2: Entry and Exit Rates by Product

Products	Entry Rates	Exit Rates
All Products	22.3	19.6
By Product:		
Boxes	12.4	12.2
Bread	7.6	18.9
Carbon Black	4.8	13.4
Coffee	9.1	15.6
Concrete	26.6	21.8
Hardwood Flooring	18.7	11.9
Gasoline	4.2	9.0
Block Ice	24.5	26.5
Processed Ice	23.1	27.7
Plywood	7.4	10.3
Sugar	3.9	17.0

Note: This table shows the plant entry and exit rates (averaged across all years in the sample). Entry (exit) is determined by plants' first (last) appearance in a CM. See text for details. Entry rates for gasoline reflect a small number (less than five) births in the 1987, 1992, and 1997 CMs, even though the Energy Information Administration reports no new refineries were built during that period. It is not clear why there is a discrepancy in the Census data, but for the sake of consistency and given the low entry rate in the industry overall, we defined these new plants as entrants.