

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
haltiwang@econ.umd.edu

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

February 2012

Abstract

It is well known that new businesses are typically much smaller than their established industry competitors, and that this size gap closes slowly. We show that even in commodity-like product markets, these patterns do not reflect productivity gaps, but rather differences in demand-side fundamentals. We document and explore patterns in plants' idiosyncratic demand levels by estimating a dynamic model of plant expansion in the presence of a demand accumulation process (e.g., building a customer base). We find active accumulation driven by plants' past production decisions quantitatively dominates passive demand accumulation, and that within-firm spillovers affect demand levels but not growth.

* We thank seminar participants at UC-Berkeley, UCLA, Carnegie Mellon, Colorado, Harvard, LSE, MIT, National Bank of Belgium, Penn State, Princeton, SITE, the Cowles Conference, CAED, CEPR IO, and UBC Summer IO meetings for their comments. We thank Lauren Deason for her excellent work developing the downstream demand indicators. Contact information: Foster: Center for Economic Studies, Bureau of the Census, 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. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed.

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 (1988), Caves (1998), Cabral and Mata (2003)). 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 the sources of the size gaps between young and old plants. While other 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 an earlier paper 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 (see Foster, Haltiwanger, and Syverson (2008)). 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 offers 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.

We highlight these patterns in Table 1, which shows the evolution of physical total factor productivity (TFP) and idiosyncratic demand across plants of various ages (this table is analogous to Table 5 in Foster, Haltiwanger, and Syverson (2008)). Physical TFP is physical units of output per unit input. Plant-level demand can be thought of as the logged 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 “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 58 percent (the demand measure’s units are logged output, so $e^{-0.550} = 0.577$) of the output of a plant in the same industry that is 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. Shapiro (1982), Caminal and Vives (1999), Radner (2003), and Fishman and Rob (2003 and 2005), and Bar-Isaac and Tadelis (2008) 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.

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

We are purposefully not too specific about the particular process behind demand-stock growth among the industries in our sample. Demand growth may well have multiple sources. These could include customer learning through “word of mouth,” the firm’s own advertising efforts, or several other possibilities, and can involve expansion of downstream buyers on either the extensive or intensive margin. We refer to the process 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 (see Bartelsman and Doms (2000) and Syverson (2011) for surveys of this literature) by explicitly accounting for demand-side effects on plants’ growth and survival. Das, Roberts, and Tybout (2007); Eslava et al. (2008); Foster, Haltiwanger, and Syverson (2008); Kee and Krishna (2008); and De Loecker (2011) are examples of the new approach. Earlier heterogeneous-productivity industry frameworks captured differences among industry producers in a single index, often explicitly or implicitly taken to be producer costs/productivity (e.g., Jovanovic (1982), Hopenhayn (1992), Melitz (2003), and Asplund and Nocke (2006)). Related empirical work on business dynamics also did not make distinctions as to the forms of heterogeneity (e.g., Dunne, Roberts, and Samuelson (1989a and 1989b); Troske (1996); Ericson and Pakes (1998); Ábrahám and White (2006); Brown, Earle, and Telegdy (2006)). 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 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.

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.

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

Our results show that while almost all entrants have lower idiosyncratic demand levels than incumbents, the gap is especially large for those owned by firms that are part of new and/or small firms. These patterns might in part reflect large or established firms' pre-existing "brand capital" imparting higher initial demand levels on their new plants. Nevertheless, all new plants, regardless of the characteristics of the firm that owns them, exhibit a limited speed of convergence in catching up to their more established competitors' demand levels. 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. Conditional on plants' sales histories, age itself only accounts for a small fraction of demand accumulation over time.

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 main empirical results are discussed in Section 5. Section 6 discusses alternative explanations and provides robustness checks and Section 7 concludes.

2. Data and Measurement Issues

This paper uses essentially the same data set of homogenous goods producers we used in Foster, Haltiwanger, and Syverson (2008).³ Details on the selection of our sample and

³ We drop good producers producing one product used in Foster, Haltiwanger, and Syverson (2008): gasoline. The current study requires not only contemporaneous data but lagged data starting in 1963 to construct initial capital

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.⁴

The roughly 17,000 plant-year observations in the sample include producers of one of ten 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), block ice, processed ice, hardwood plywood (plywood), and raw cane sugar (sugar).⁵ These products were chosen because of their physical homogeneity which 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

stocks and also lagged revenue measures. We found the historical data for the gasoline refining industry was somewhat spotty, which limited the number of industry plants for which we had valid data. We also think that our learning about demand model is somewhat less well suited to gasoline products, especially since there is so little entry in gasoline to identify our learning effects. As will become clear, we think the model is especially well suited to our local products industries.

⁴ 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.

⁵ Our product definitions are built up from the seven-digit SIC product classification system. Some of our ten 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 ten 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 ten products can be found in section A.3 of the Appendix.

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.

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 ten 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 function separately for each of our ten products:

$$(1) \quad \ln q_{it} = \alpha_0 + \alpha_1 \ln p_{it} + \sum_t \alpha_t \text{YEAR}_t + \alpha_2 \ln(\text{INCOME}_{mt}) + \eta_{it},$$

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 (INCOME_{mt}). We define local markets using the Bureau of Economic Analysis' Economic Areas (EAs).⁶

Plant quantities are simply their reported output in physical units. We calculate unit prices for each producer using their reported revenue and physical output.⁷ These prices are then

⁶ EAs are collections of counties usually, but not always, centered on Metropolitan Statistical Areas. The 172 EAs are mutually exclusive and exhaustive of the land area of the United States. See U.S. Bureau of Economic Analysis (1995) for detailed information.

⁷ 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.

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 positive (negative) demand shocks η_{it} by raising (reducing) 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 with any short-run plant-specific demand shocks embodied in η_{it} . Hence they appear quite suitable as instruments for plant prices.⁸

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 results 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.

⁸ 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 1-2 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.

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 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 of firms to which plants are tied.

Consider the following example. Two new plants are built in an industry: one plant is a de novo entry by a firm with no prior experience; the other plant is opened by a large firm with considerable history, perhaps but not necessarily in the same industry and geographic area. 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 characteristics of the firms that own the plants. In the first specification, we 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 essentially a crude proxy for firm size. Such multi-plant firms account for 59 percent of the observations in our sample. The second specification uses a series of dummies for the age of *the firm*, defined as the age of the firm's oldest plant. These are 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.

The results looking at the impact of multi-plant 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 shows the evolution of idiosyncratic demand for single-unit plant/firms, while the column-wise sum of the two rows' values reflects the same evolution for plants in multi-plant firms. Note that the excluded group is different here from that in Table 1. There, it was all old plants—those having first appeared three or more CMs prior, and therefore being at least 15 years old. Here, it is only old plants in single-unit firms. Hence the age coefficients in the table show 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 is noticeable.

Single-unit plants exhibit similar patterns to those seen before for the whole sample. 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 young (5-9 years old), where young single-unit plants have demand levels 16 percent below old single-unit plants. Convergence then largely stalls; medium-aged single-unit plants still have 14 percent demand deficits.

For plants in multi-plant firms, similar qualitative relationships are present, but their demand levels are significantly higher than single-unit plants at every age. That said, they are still considerably smaller than old plants in multi-unit firms, with average demand levels for new plants 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.

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.⁹ We develop a model of dynamic, endogenous demand accumulation in the next section that we will take to our sample to further investigate the nature of the accumulation process.

⁹ 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 firm. 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. In the appendix (see section A.2 of the appendix) we also show that the patterns in Tables 1 and 2 are robust to controlling for firm age. That is, there is slow growth of new plants even in large, mature firms.

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 estimate this model using producers' expansion patterns in our dataset to learn about the nature of the processes driving demand growth.

We assume the plant faces an isoelastic contemporaneous demand curve:

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

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 Z_t evolves according to the following process:

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

Thus, Z_t is a sort of operating history of the plant. It grows with past plant sales R_{t-1} (defined as $p_{t-1}q_{t-1}$; we use lagged rather than current sales only for analytical convenience), 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

$$(5a) \quad \pi_t = \theta_t^\eta \text{Age}_t^\eta Z_t^\eta (A_t x_t)^{1-\frac{1}{\eta}} - c_t x_t - f.$$

The plant manager maximizes the present value of the plant's operating profits.¹⁰ This problem can be expressed recursively as follows:

$$(6) \quad V(Z_t, A_t, \text{Age}_t, \theta_t) = \max \left\{ 0, \sup_{x_t} \theta_t^\eta \text{Age}_t^\eta Z_t^\eta A_t^{1-\frac{1}{\eta}} x_t^{1-\frac{1}{\eta}} - c_t x_t - f + \beta EV(Z_{t+1}, A_{t+1}, \text{Age}_{t+1}, \theta_{t+1}) \right\},$$

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 evolve exogenously. The plant discounts the future 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 (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 is 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

¹⁰ We abstract from any agency issues that may arise between plants' managers and the owners of these establishments (if they are different people).

interpreted through the lens of our model as demand. It is therefore important that we consider any other such forces and how they might impact the interpretation of our results. We 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.¹¹ 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 (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 completely static production/pricing optimization problem, would be zero. It is not generally so here because of the dynamics discussed above. Because the plant shifts out its demand curve tomorrow by making more sales today, it will 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 is not just what is 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

¹¹ 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 of δ) from the plant's sales reported in past CMs.

lower markup than implied by the static 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 endogenous impact of age on demand is positive (i.e., as long as γ is positive). Alternatively, if $\gamma = 0$, this is no longer a dynamic problem and there is no incentive to deviate from the optimal static markup.

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, as long as the first right-hand-side term is also positive.

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.

The remaining issue for measuring demand stocks is how to initialize Z for entrants, Z_0 . Here, we draw insights from the descriptive empirical results in Section 3. We allow a plant's

initial demand stock to be a function of the structure of the firm that owns it. Specifically, we specify the initial demand stock of plant e as

$$(8) \quad Z_{0e} = (K_{0e})^{\lambda_1} \left(\frac{K_{0s(e)} + K_{0e}}{K_{0e}} \right)^{\lambda_2},$$

where K_{0e} is the initial physical capital stock of e , $K_{0s(e)}$ is the *sum* of the physical capital stocks of plant e 's siblings (i.e., the total capital stock that year of the other plants owned by the same firm within manufacturing), and λ_1 and λ_2 are parameters. The logic behind (8) is that a plant's initial demand stock can be related to its own physical size (K_{0e}) as well as the size of its owning firm. This specification therefore incorporates the possibility, seen in the previous section's results, that entrants of larger firms start with larger idiosyncratic demand levels than do those of smaller firms. Note that (8) mechanically allows for single-plant firm entrants, where the entrant *is* the firm, because in that case $K_{0s(e)} = 0$ and the ratio in the parentheses is unity. Additionally, (8) nests the possibility that multi-plant firm entrants do not have initial demand advantages, which would be the case if $\lambda_2 = 0$. This specification lets the data tell us how important the owning firm's characteristics are in determining the initial demand stock of a new plant.¹²

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. 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. Substituting out for θ_t using the demand curve causes the Age_t terms to cancel. Estimating the demand equation lets us recover ϕ as well as harness additional data variation to identify the other parameters.

We must address the issue of endogeneity in estimating the demand equation (2), as its right-hand-side variables include endogenous plant level prices as well as state variables Z_t and

¹² We face a two other 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. Since we cannot see these plants' past sales, we cannot fully construct an initial Z for these firms. Instead, we extend the logic of our modeling of new plants' Z_0 by letting the 1963 cohort's Z_{1963} be given by the same form as (8). 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.

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:

$$(9) \quad \theta_{t+1} = \rho \theta_t + v_{t+1}$$

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 unobserved demand shock innovation. 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.

Estimation of this demand equation 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 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. Thus we can observe 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 have zero mean 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 the model's 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 average five-year exit rate for our data sample is around 20 percent, so selection could 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, is 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 "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 growth 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. This shows up in the demand accumulation by being channel. Those emphasizing endogenous demand-stock building—resulting from the active efforts of the plant—will show a large influence of Z_t , demand accumulation by doing. We can measure the relative importance of each in the data.

5. Estimation Results

We jointly estimate via GMM the demand (2b) and Euler (7c) equations.¹³ We estimate the model for the entire sample, for local products only and for concrete plants only. Local products are defined as products for which the majority of output is shipped less than 100 miles according to the Commodity Flow Survey. These products include boxes, bread, concrete, and ice. We highlight the local products subsample since it is possible that our model is better suited to such products, or the parameters of the learning dynamics might easily be different for these products. The concrete-only subsample enables us to focus on a specific product where we have many observations, permitting estimation of industry-specific parameters. We would prefer to let all parameters to vary across all products in our estimation, but some of our 10 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 the sample description reported in appendix Table A.4.) These subsamples serve as an alternative means of exploring the robustness of our findings across products. However, we do also report some results below where we permit key parameters to vary with a function of the industry's attributes.

The variables included in our estimated model are defined as 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 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 . (Recall that we have no entrants in the estimation sample because we need to use lagged variables to identify the dynamic parameters.)

We also include controls in the demand equation not explicitly referenced in the above discussion of the model. Specifically, because we are pooling data across products and years, we

¹³ We do not estimate β in the Euler equation but rather set it to be consistent with annual discount factor of 0.98. We check below the robustness of our results to alternative values.

include a set of fully interacted product and year effects. We also include measures of the local market for those products that are deemed local products. Specifically, we include a measure of local income in the local market (see Foster, Haltiwanger and Syverson (2008) for details) as well as a measure of the average price of local competitors in the same industry. These latter two variables are potentially important in accounting for shifts in demand that would otherwise be subsumed into the unobservable demand component θ . There is no reason to believe that they should be directly relevant for the Euler equation, however.

5.1. Estimates of the Model on the Full Sample

We estimate two versions of the model. One imposes that the depreciation rate of the demand stock, δ , is zero. The other version allows δ to be estimated with the other parameters. In the $\delta = 0$ case, the demand stock simply reflects cumulative real revenue. This case is the demand-side analog to standard learning-by-doing models that do not allow for “forgetting” in the style of Benkard (2000). The results of the estimation are reported in Table 3. Column 1 reports the results of the cumulative learning model with no depreciation, and column 2 reports the results of the model when δ is estimated.

We find qualitatively similar results in the two alternative models. For example, we find roughly similar elasticities of demand, positive and significant estimates of γ consistent with “demand accumulation by doing” and also evidence of “demand accumulation by being.” In what follows, we focus our attention on the model with estimated depreciation because the evidence clearly rejects the hypothesis that the depreciation rate of the demand stock is zero. The estimate for the full sample for δ is 0.893. This reflects total depreciation over a five-year horizon, so the implied annual depreciation rate is 36 percent ($(1 - 0.36)^5 = 0.107$). As will become clear, the finding of an economically and statistically significant depreciation rate is a common finding in alternative specifications we consider. We focus on such specifications in the remainder of the paper.

We now turn to a more detailed discussion of the estimates of this model. First, consider the estimates of the price elasticity of demand η . The estimate for the full sample is -1.8. This value is in a similar range as those in Foster, Haltiwanger and Syverson (2008) with a significantly richer specification of the demand structure and its determinants. Also, note that we include as a control a measure of competitors’ price in the local market for those products that

are shipped locally (for national products this effect is not separately identified, as we are already including product-by-year effects). We find that the elasticity of a plant's demand with respect to a price increase by its local competitors is 0.338. This is consistent with the hypothesis that higher prices of competitors, other things equal, increase demand for the plant in question.¹⁴

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” in the elasticity of future demand to the demand stock, γ . The estimated value of γ is around 0.8. Producing more today will significantly shift the plant's demand curve out tomorrow; a 10 percent increase in a plant's demand stock corresponds to an 8 percent increase in the number of units the plant sells at any given price. As reflected in the Euler equation, a producer'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.

This parameter estimate can also help us get a feel for the potential return to a business “investing” in its demand stock by lowering prices today in hopes of shifting out its demand tomorrow. Based on the estimated price elasticity in the model with depreciation, a ten percent price cut will increase current quantity sold by about 18 percent and current revenues by eight percent. (This is a sizeable price deviation from one's competitors, but not unusual. The average within-market standard deviation of plants' logged prices is 0.18.) This increase in revenues will shift out the plant's demand in the following year by about four percent, taking into account both depreciation and γ . This means the plant will be able to sell four percent more units at a given price than it would otherwise. A ten percent price cut that persists for five years will cumulatively yield (taking into account depreciation) a 13 increase in the demand stock Z above what it would be in absence of the cut.

In addition to the endogenous demand accumulation effect, we find that, having controlled for a plant's demand stock, “demand accumulation by being” also contributes to the demand gap. The coefficient on the young dummy is negative and significant, and the coefficient on the medium age dummy is much smaller and not significant. Since the omitted group is the oldest plants, this means the demand impact of age increases monotonically. This is qualitatively consistent with the raw demand gap patterns in Table 1, but these effects are much smaller than those in Table 1. The demand gap closed from young to old plants due to age

¹⁴ We also find that local income increases demand.

effects alone is about seven percent here, only about one-fifth the Table 1 gap. This indicates that once we have accounted for endogenous demand accumulation (and other factors), the remaining “exogenous” age gap is much smaller.

Remember that 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 0.37. This five-year persistence rate corresponds to an annual rate of 0.82.

The impact of the characteristics of the owning firm on an entering plant’s initial demand stock is seen in the comparison of the estimates of λ_1 and λ_2 . The value of λ_1 , which parameterizes how a plant’s initial demand stock Z is related to its physical capital stock, is 0.651, indicating that, not surprisingly, plants with larger initial physical capital tend to have larger starting demand stocks. The parameter also indicates that the ratio between the two types of capital falls in the plant’s size. The estimated value of λ_2 , which is the elasticity of a plant’s initial demand stock to the size of the firm (in physical capital terms) relative to the entering plant, is 0.548. This indicates that, consistent with the descriptive results seen in Table 2, new plants of larger firms do in fact have higher initial demand stocks. A plant started by a firm that is twice as large as another entering plant’s firm will start with about a 38 log point ($0.548 \cdot \ln 2 = 0.38$) higher demand stock.

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 capital predicted plant survival. At the same time, it is also excluded from the plant’s dynamic profit maximization problem in our model, making it a candidate excluded instrument for selection.¹⁵

¹⁵ Of course, this raises the issue of what happens if plant capital *is* a dynamic choice in reality despite our exclusion of it. We address this point in detail in Section 6 below when we consider alternative explanations for our findings.

The coefficient estimates on the selection controls suggest that any selection bias is relatively modest in our sample. The estimates are statistically significant but small in magnitude. The values of the other parameter estimates are roughly similar in unreported specifications that exclude the selection controls. We have also tried more flexible ways of controlling for selection bias, such as using various polynomials of the plant's estimated survival probability in place of the standard Mills ratio control.¹⁶ The model's main parameter estimates are robust across these alternative specifications. Hence it seems that selection issues are not skewing our reported estimates.

5.2. Estimates Using Local Products and Concrete Plants

To explore the consistency of our parameter estimates across the industries in our sample, we estimate the model on two successively smaller subsamples. One uses only those plants in local products industries (boxes, bread, concrete, and ice), and the other uses concrete plants alone. (We choose concrete for the single-industry subsample because it has the largest number of plants in our sample of any industry.) The results are in Table 4; column 1 reports the estimates for the local products subsample, and column 2 reports the concrete results. We again focus on the specification with depreciation since in both of these subsamples the estimated rate of depreciation is far from zero.

Overall, the results for the two subsamples are qualitatively similar to the results for the full sample, suggesting it is not ridiculously restrictive to constrain the parameters to be the same across all product industries. There are some quantitative differences, however, that we discuss briefly. Demand is more own-price elastic for concrete than for the entire sample. Concrete demand is also more responsive to local competitors' prices.¹⁷ These elasticities for the all-local-products subsample are essentially the same as those in the entire sample, however.

The main parameter of interest, the elasticity of demand to the plant's endogenously acquired demand capital, is roughly the same in these subsamples as for the whole sample, with γ estimated at about 0.8. While γ is similar across products, concrete has a substantially lower depreciation rate (an implied 13 percent per year as opposed to 27 percent per year) which is

¹⁶ Pakes (1994) and Aguirregabiria (1997) discuss the use of polynomials to approximate unknown functional forms for selection corrections.

¹⁷ The estimated price elasticity of demand for concrete is somewhat lower than that reported in Foster, Haltiwanger and Syverson (2008).

important for the demand accumulation dynamics. Combining these depreciation and price elasticity estimates suggests that a plant that cuts prices by ten percent to invest in future demand will raise current revenues by seven percent for local products and 13 percent in the concrete subsample. This increase in sales will in turn increase the producer's quantity demanded next year by about four percent for local products and seven percent for concrete. The same ten percent price drop sustained for five years increases the producer's demand stock Z by 15 percent in local products and 44 percent on concrete, the former reflecting the cumulative effect of higher depreciation rates.

We also find that the exogenous (age-related) demand accumulation process has similar qualitative patterns as for the entire sample. There is a positive estimated demand-accumulation-by-being effect for both local products and concrete. The quantitative effects are somewhat larger in local products than for the full sample, and larger still in concrete.

The estimated value of λ_1 is 0.937 for local products and 1.086 for concrete, which again indicates larger plants tend to have larger starting demand stocks. For these products, the ratio between initial demand and physical capital stays about the same with plant size. The influence of firm size on a plant's initial demand stock, which is embodied in λ_2 , is 0.285 for local products and 0.410 for concrete. A plant started by a firm that is twice as large as another entering plant's firm will start with about a 19 percent higher demand stock if the plant is in the local products industries and a 28 percent higher demand stock if the plant is in the concrete industry.

Again, selection does not appear to be quantitatively important. Three of the four inverse Mills ratios are statistically significant, but all are small. Estimating the model without including any selection correction terms (not shown) yielded similar estimates of the other parameters.

5.3. Interactions with Multi-Plant Firm Status

One of the most striking results of the descriptive exercises in Section 3 is that entrants that are part of larger, multi-plant firms enter with a higher demand stock than those in smaller or single-plant firms. This was confirmed in the estimated model above as well, as the elasticity of initial demand stock to the ratio of the firm's size to the entering plant's size, λ_2 , was positive. However, it was less clear in the descriptive results whether the rate of convergence of idiosyncratic demand levels was faster for young plants in multi-plant firms than those in small firms. To look for this possibility through the lens of our model, we also estimate a specification

that interacts an indicator for plants that are owned by a multi-plant firm with the model's parameters (except for λ_1 and λ_2 , which already incorporate such multi-unit firm effects). The results, for both entire sample (column 1) and the local-products-only sample (column 2), are shown in Table 5. To interpret the results in this table, the “main” effects provide estimates for single-unit plants, and the interaction effects with the MU dummy provide an estimate of whether MU plants have a significant differential from the single-unit plants (so that the total MU coefficient is the sum of the main and interaction effect).

There is some evidence that endogenous demand accumulation forces are slightly stronger among plants owned by multi-unit firms, at least for the overall sample. The interaction between the multi-unit indicator and γ is 0.084 (s.e. = 0.037) for the full sample and 0.030 (s.e. = 0.025) for the local product sample. In addition, estimated depreciation is lower for multi-unit plants.¹⁸ However, the fact that demand is slightly more inelastic for multi-unit plants implies, once the effects are worked through the model, that the implied incentives to cut price in order to build future demand capital are quantitatively not that different for single- and multi-unit plants.

We do find evidence of more of an “accumulation by being” effect for multi-unit as opposed to single-unit plants. This can be interpreted as suggesting that the residual unexplained component of the patterns observed in Table 2 is larger for multi-unit plants. We return to this issue below.

As we noted above, one of the largest differences in Table 2 is the difference in the intercepts. This is captured here by permitting the presence and size of the parent firm at the time of entry of the plant to contribute to the demand stock. Given the large estimated coefficient for λ_2 (0.397 for the full sample and 0.442 for local products), there is a large level shift in the demand curve for multi-units from the point of entry. Establishments that are part of multi-unit establishment firms have initial firm-level capital stocks that are about 1.9 times that of the median entering establishment. Using the full-sample estimates, such an establishment will start with a 25 log point ($0.397 \cdot \ln 1.9 = 0.25$) higher demand stock.

5.4. Evolution of Demand by Age: Exogenous versus Endogenous Demand Accumulation

To further quantify the contribution of exogenous versus endogenous demand

¹⁸ We note that the estimated depreciation rate, while far from zero, is substantially lower for both single- and multi-unit plants than the depreciation rates without such distinctions reported in Tables 3 and 4. Apparently, interacting other parameters with multi-unit status yields a lower implied depreciation rate.

accumulation to the observed evolution of demand across plant ages, we return to the metric used in Table 1. In particular, we use the estimated coefficients from our model along with the actual data to compute the implied levels of both demand components for every plant-year observation in our sample. We then derive the type of statistics reported in Table 1 for each of these computed components.

We compute the component of demand from the exogenous demand accumulation (“accumulation by being”) using the estimates of dummy variables for age reported in Tables 3 and 4. For the endogenous demand component (“accumulation by doing”), we first compute Z_t for every plant in the sample using our data on plants’ revenues and capital stocks along with the estimates of λ_1 , λ_2 and δ . We then combine the estimated Z_t with our estimate of γ to compute the endogenous demand accumulation component for every plant-year observation.

Table 6 reports the results of these exercises. The top panel shows the results for the full sample, the middle panel for local products plants, and the bottom panel for concrete plants. Plants are classified into three age categories: “Young” for plants that either entered in the current or prior Census (i.e., those that are less than 10 years old); “Medium” for plants that entered two Censuses ago (10-14 years old); and “Old” for plants that entered three or more Censuses ago (15 or more years old). We combine the “Entrants” and “Young” categories from Table 1 into a single “Young” category here for two reasons. First, the model only yields estimates of the exogenous demand accumulation component for these same young and medium categories relative to older plants.¹⁹ Second, this grouping of ages implies that all counterfactual estimates of endogenous demand accumulation component reflect actual past sales rather than just our estimated demand stock initialization.

Because we use somewhat collapsed age categories and capital stock data are not available for all plants used in Table 1, the first row in each of the panels of Table 6 repeats exactly the type of analysis done in Table 1 for this restricted sample. As in Table 1, these estimated coefficients are from a regression of the demand shock on age dummies and industry-year fixed effects. The demand shock patterns for each panel in Table 6 are similar qualitatively

¹⁹ While the learning by being component for the young reflects plants between 5-9 years old, one can obtain for a plant of any age an estimate of the contribution of all components of demand other than the endogenous demand accumulation at any age by taking the difference between the overall demand shock observed in the data and the endogenous demand accumulation component. This difference includes the accumulation by being component but also other components, like the unobserved idiosyncratic persistent demand component θ .

and quantitatively to those in Table 1. Young and Medium aged plants have much lower demand than old plants and convergence is slow.

Our model lets us decompose this overall demand residual into multiple components.²⁰ The age patterns for the endogenous accumulation component are reported in the second row of each panel. For the full sample, the endogenous accumulation component is much lower for young and medium aged plants relative to old plants in a manner that quantitatively is very similar to the first row. The accumulation by being component is relatively small in magnitude, although it does exhibit modest growth over time. These results imply that most of the overall demand shock patterns for our sample plants are accounted for by endogenous accumulation of demand rather than exogenous components.

Results for local and concrete plants are similar. The endogenous accumulation component in all cases is much lower for young plants (e.g., 32 log points lower for local product plants and 24 log points lower for concrete plants) than old plants, and there is only slow convergence. Accumulation by being accounts for a larger share of demand growth in the local products and concrete subsamples, but in all cases the majority of demand accumulation occurs via the endogenous channel.

6. Alternative Explanations and Robustness Checks

In this section, we attempt to address two basic concerns that we anticipate readers might have with our work and provide some additional robustness checks. The first basic concern is relatively minor and is addressed in the first subsection. It 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. In the second subsection we address the more serious concern that we have allowed only one channel for dynamics in our model, demand stock accumulation. If a plant’s management takes into account other dynamic factors when making decisions, we would mistakenly measure these other factors’ 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.

²⁰ The two components we report do not add up to the total because there are other factors—in particular, the serially correlated demand component θ —that enter into the demand equation.

6.1. What Do Our Idiosyncratic Demand Stock Measures Reflect?

Our idiosyncratic demand stock measures reflect the cross-plant variation in units of output sold that is, by construction, purged of the effects of plants' physical production costs. If plant A both sells more output and has a higher idiosyncratic demand measure than plant B, plant A's high sales are not simply the result of plant A having lower prices because it has low costs. Plant A would sell more than plant B even if it were charging the same price. Regardless of any other measurement issues with these idiosyncratic demand measures, 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 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 do not support this sort of capacity utilization pattern. We cannot 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 plants' production quantities are proportional to their reported capital stocks. 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 7 presents the utilization patterns for our sample. 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 capital-output ratio 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 patterns in the data are not consistent with supply-side learning by doing—this is, after all, a basic motivation of our investigation. This 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 do not hold clear patterns over the 15+ year horizons we are explaining demand movements. Further, the quantitative movements in physical TFP that do exist across ages are very 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 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 Table 2. 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 shifts 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 they cannot quantitatively 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 capital adjustment costs 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 do not have the quantitative impact necessary to explain the long-horizon demand-growth patterns we observe in the data.

6.3. Robustness Checks

We discuss briefly a number of robustness checks that we conducted on our analysis in this subsection. Relevant estimates for these robustness checks are reported in the appendix.

First, we investigated the sensitivity of our results to permitting a different discount factor β . The results reported in Tables 3-6 reflect an assumed annual discount factor of 0.98. Figure A.1 shows how the key parameters of γ and δ vary for discount factors ranging between 0.96 and 0.98. The estimates are quite robust across this range. We focus on these two parameters because they are the critical estimates for the endogenous demand accumulation, but we also found (results available upon request) that other estimated parameters are also robust over this range of β .

Second, we explored refinements of the role of being part of a multi-unit firm upon entry. The main results imply that plants entering as part of a multi-unit firm have significantly higher initial demand stocks. To explore this mechanism further, we considered whether this is tied to the nature of the overlap of a new plant's activity with its parent firm. Specifically, we tested whether the entering plant's initial demand was higher if its owning firm had experience operating other plants in the same geographic area (using the BEA Economic Area definition of geographic areas) and industry (using 4-digit SIC industries). We found some evidence in favor of this conjecture for plants in the local products subsample, but not much for concrete alone or

for national product plants. We report in Table A.1 results that consider this refinement for local product plants. Even for this subsample we note that the contribution of initial demand capital from the firm's plants that are not part of the same geography or industry remains significant. It might seem surprising that we did not find evidence in support of this conjecture for concrete plants, but we note that for concrete plants in particular there is not much variation to exploit on these dimensions (i.e., concrete plants that belong to a multi-unit firm upon entry are largely from firms that operate in the same industry and geographic area). While we recognize that we only explored this refinement in a limited fashion, we think the results imply that the demand stock advantage of being part of a multi-unit does not stem simply from the activity of the firm's other plants in the same industry or geographic market.

A third set of refinements we conducted was to allow some systematic variation in the model's parameters across industries by permitting the estimated parameters to vary with observed characteristics of our sample products. We already saw in Tables 3 and 4 that while the results are qualitatively similar across our sample products, there is some quantitative variation in the parameter estimates across the full sample, the subsample of local product producers, and concrete plants. We considered allowing the parameters to vary with several alternative product characteristics, and in particular the nature of the downstream industries that purchase our sample products (we identified these downstream industries using the input-output matrix). We conjectured, for example, that producers who sell to downstream industries with more turnover of producers would face weaker incentives for demand stock accumulation, while those selling to more concentrated downstream industries would face stronger incentives. We found only modest evidence in favor of these conjectures (see Table A.2 for details). For example, the elasticity of demand with respect to the demand stock (γ) tends to increase with the concentration of downstream producers when using the full sample. However, this result was not observed in the local product plant subsample, where we found no statistically significant variation in γ on these dimensions. Also, while the point estimates in both the full and local products samples indicated that γ decreases when downstream turnover rises, neither of these estimates were statistically significant.

7. Conclusion

We have used a unique dataset to examine the sources of the persistent gap in size across

young and old establishments. Building upon our earlier work which had suggested that supply-side factors were not primarily responsible for this gap, we focused on demand-side factors in this paper. Our results imply that, even in commodity-like product industries, entry is difficult. It takes a long time for new businesses, even those owned by large firms, to reach a point where they have built enough relationship-specific capital with their potential customers to expect (at the same price) to sell the same amount of output as do their more established competitors. These results further buttress the recent literature pointing towards the importance of idiosyncratic demand factors in explaining the fortunes of businesses.

To explore the qualitative and quantitative nature of the demand accumulation process, we built and estimated a model of a producer's dynamic growth problem that allowed for expansion through both endogenous ("demand accumulation by doing") and exogenous ("demand accumulation by being") processes. While both processes play a role in new businesses' expansions, the majority of demand accumulation is the product of the producers' own efforts to build their demand stocks.

Besides speaking to the literature on the role of idiosyncratic demand factors in shaping businesses' trajectories, we believe our results can offer insights to several lines of research that have focused on various aspects of demand-side influences on firm growth, including those in the fields of industrial organization, macro, and trade (e.g., Klepper (2002), Cabral and Mata (2003), Arkolakis (2010), Dinlersoz and Yorukoglu (2010), Gourio and Rudanko (2011) and Luttmer (2011)).

References

- Ábrahám, Árpád and T. Kirk White. "The Dynamics of Plant-Level Productivity in U.S. Manufacturing." Center for Economic Studies Working Paper 06-20, 2006.
- Aguirregabiria, Victor. "Estimation of Dynamic Programming Models with Censored Dependent Variables." *Investigaciones Economicas*, 21(2), 1997, 167-208.
- Arkolakis, Costas. "Market Penetration Costs and the New Consumers Margin in International Trade." *Journal of Political Economy*, 2010, 118(6), 1151-1199.
- 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.
- Bar-Isaac, Heski and Steven Tadelis. "Seller Reputation." *Foundations and Trends in Microeconomics*, 4(4), 2008, 273-351.
- Bartelsman, Eric J. and Mark Doms. "Understanding Productivity: Lessons from Longitudinal Microdata." *Journal of Economic Literature*, 38(3), 2000, 569-95.
- Benkard, C. Lanier. "Learning and Forgetting: The Dynamics of Aircraft Production." *American Economic Review*, 90(4), 2000, 1034-54.
- Brown, J. David, John S. Earle, and Almos Telegdy. "The Productivity Effects of Privatization: Longitudinal Estimates from Hungary, Romania, Russia, and Ukraine." *Journal of Political Economy*, 114(1), 2006, 61-99.
- Cabral, Luís M B, and José Mata. "On the Evolution of the Firm Size Distribution: Facts and Theory." *American Economic Review*, 93(4), 2003, 1075-1090.
- 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.

- De Loecker, Jan. "Product Differentiation, Multiproduct Firms, and Estimating the Impact of Trade Liberalization on Productivity." *Econometrica*, 79(5), 2011, 1407-51.
- Dinlersoz, Emin and Mehmet Yorukoglu. "Information and Industry Dynamics." US Census Bureau Center for Economic Studies Paper No. CES-WP-10-16.
- 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.
- Dunne, Timothy, Mark J. Roberts, and Larry Samuelson. "Firm Entry and Postentry Performance in the U.S. Chemical Industries." *Journal of Law and Economics*, 32(2) Part 2, 1989a, S233-71.
- Dunne, Timothy, Mark J. Roberts, and Larry Samuelson. "The Growth and Failure of U.S. Manufacturing Plants." *Quarterly Journal of Economics*, 104(4), 1989b, 671-98.
- Eslava, Marcela, John Haltiwanger, Adriana Kugler and Maurice Kugler. "Plant Survival, Market Fundamentals and Trade Liberalization." Mimeo, 2008.
- Fishman, Arthur and Rafael Rob. "Consumer Inertia, Firm Growth and Industry Dynamics." *Journal of Economic Theory*, 109(1), 2003, 24-38.
- Fishman, Arthur and Rafael Rob. "Is Bigger Better? Customer Base Expansion through Word-of-Mouth Reputation." *Journal of Political Economy*, 113(5), 2005, 1146-1162.
- 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.
- Gourio, Francois and Leena Rudanko. "Customer Capital." NBER Working Paper No. 17191, 2011.
- 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.
- Kee, Hiau Looi and Kala Krishna. "Firm-Level Heterogeneous Productivity and Demand Shocks: Evidence from Bangladesh." *American Economic Review Papers and Proceedings*, 98(2), 2008, 457-62.

- Klepper, Steven. "Firm Survival and the Evolution of Oligopoly." *The RAND Journal of Economics*, 33(1), 2002, 37-61.
- Luttmer, Erzo G. J. "On the Mechanics of Firm Growth." *Review of Economic Studies*, 78(3), 2011, 1042-1068.
- Melitz, Marc J. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica*, 71(6), 2003, 1695-1725.
- Pakes, Ariel. "Dynamic Structural Models: Problems and Prospects. Mixed Continuous Discrete Controls and Market Interactions" in C. Sims (ed.) *Advances in Econometrics. Sixth World Congress*. Vol 2. 1994, Cambridge University Press.
- Pakes, Ariel and Richard Ericson. "Empirical Implications of Alternative Models of Firm Dynamics." *Journal of Economic Theory*, 79(1), 1998, 1-45.
- Radner, Roy. "Viscous Demand." *Journal of Economic Theory*, 112(2), 2003, 189-231.
- Shapiro, Carl. "Consumer Information, Product Quality, and Seller Reputation." *The Bell Journal of Economics*, 13(1), 1982, 20-35.
- Syverson, Chad. "What Determines Productivity?" *Journal of Economic Literature*, 49(2), 2011, 326-65.
- Troske, Kenneth R. "The Dynamic Adjustment Process of Firm Entry and Exit in Manufacturing and Finance, Insurance, and Real Estate." *Journal of Law and Economics*, 39(2), 1996, 705-35.
- 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)	0.004 (0.006)	-0.004 (0.006)	-0.018 (0.005)
Demand Shock	-0.550 (0.022)	-0.397 (0.024)	-0.316 (0.026)	-0.339 (0.021)

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 “old” plants (those that appeared three or more censuses prior). The sample includes roughly 17,000 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 and a smaller sample (we exclude plants manufacturing gasoline from the analysis).

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	-0.318 (0.034)	-0.176 (0.035)	-0.150 (0.038)	Excl.	-0.183 (0.031)
Demand Shock x MU firm	0.106 (0.038)	0.132 (0.041)	0.237 (0.045)	0.530 (0.026)	-0.283 (0.042)

Note: This table repeats the analysis of Table 1, but now allows plant age effects to vary with the multi-unit (MU) status of the plant’s owning firm. The excluded category is “old” plants (those that appeared three or more censuses prior) that are part of a single-unit firm. N is roughly 17,000 plant-year observations. Standard errors, clustered by plant, are in parentheses.

Table 3. Estimated Coefficients for Entire Sample, Cumulative Learning and Depreciation Models

Parameter	Cumulative Learning	Learning with Depreciation
γ (elasticity of future demand to the demand stock)	0.287 (0.068)	0.795 (0.014)
$-\eta$ (price elasticity of demand)	-2.576 (0.235)	-1.808 (0.082)
Young dummy (demand shift for entering and young plants)	-0.179 (0.073)	-0.066 (0.031)
Medium age dummy (demand shift for medium-aged plants)	-0.051 (0.055)	-0.025 (0.026)
ρ (persistence of exogenous demand shocks θ)	1.188 (0.035)	0.366 (0.085)
λ_1 (elasticity of initial demand to plant's own K)	1.803 (0.177)	0.651 (0.051)
λ_2 (elasticity of initial demand to ratio of firm's K to plant's K)	0.103 (0.340)	0.548 (0.063)
Competitor's Price (local products only)	0.315 (0.334)	0.338 (0.073)
δ (demand depreciation rate)		0.893 (0.026)
Inverse Mills Ratio, Demand (selection correction, demand equation)	0.052 (0.020)	-0.022 (0.009)
Inverse Mills Ratio, EE (selection correction, Euler equation)	0.002 (0.003)	0.026 (0.005)

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 4. Estimated Coefficients for Local Industry and Ready Mix Concrete Sample, Depreciation Model

Parameter	Local Products	Concrete
γ (elasticity of future demand to the demand stock)	0.843 (0.010)	0.751 (0.023)
$-\eta$ (price elasticity of demand)	-1.705 (0.077)	-2.321 (0.141)
Young dummy (demand shift for entering and young plants)	-0.102 (0.027)	-0.211 (0.043)
Medium age dummy (demand shift for medium-aged plants)	-0.027 (0.022)	-0.066 (0.033)
ρ (persistence of exogenous demand shocks θ)	-0.142 (0.043)	0.277 (0.061)
δ (demand depreciation rate)	0.787 (0.023)	0.500 (0.048)
Competitor's Price	0.317 (0.064)	1.416 (0.225)
λ_1 (elasticity of initial demand to plant's own K)	0.937 (0.030)	1.086 (0.029)
λ_2 (elasticity of initial demand to ratio of firm's K to plant's K)	0.285 (0.055)	0.410 (0.033)
Inverse Mills Ratio, Demand (selection correction, demand equation)	-0.035 (0.010)	-0.018 (0.013)
Inverse Mills Ratio, EE (selection correction, Euler equation)	0.030 (0.004)	0.017 (0.003)

Notes: See notes to Table 3.

Table 5. Estimated Coefficients for Learning with Depreciation Model, Interactions with Multi-Plant Firm Status

Parameter	Entire Sample	Local Products
γ	0.705 (0.030)	0.770 (0.023)
η	-2.507 (0.173)	-2.140 (0.188)
Young dummy	-0.026 (0.075)	0.053 (0.044)
Medium age dummy	0.038 (0.051)	0.036 (0.042)
ρ	0.571 (0.095)	0.240 (0.093)
λ_1	1.106 (0.035)	1.028 (0.025)
λ_2	0.397 (0.050)	0.442 (0.031)
δ	0.584 (0.069)	0.673 (0.043)
Competitor's Price	0.733 (0.171)	0.338 (0.186)
$\gamma*MU$	0.084 (0.037)	0.030 (0.025)
$\eta*MU$	0.407 (0.158)	-0.055 (0.209)
Young dummy* MU	-0.305 (0.105)	-0.347 (0.062)
Medium age dummy* MU	-0.169 (0.068)	-0.123 (0.053)
$\rho*MU$	0.451 (0.096)	0.137 (0.115)
$\delta*MU$	-0.104 (0.079)	-0.144 (0.049)
Competitor's Price* MU	-0.475 (0.161)	0.042 (0.207)
Inverse Mills Ratio (Demand)	0.002 (0.013)	-0.020 (0.010)
Inverse Mills Ratio (EE)	0.010 (0.003)	0.016 (0.003)

Notes: See notes for Table 3 above. "MU" is an indicator variable equal to one if the plant is owned by a multi-unit (multi-plant) firm.

Table 6. Evolution of Demand across Plant Ages— Endogenous Learning vs. Learning By Being Effects

Variable	Young	Medium	Old
	All Plants		
Demand shock	-0.575 (0.020)	-0.287 (0.029)	Excl.
Endogenous Accumulation (Accumulation by Doing)	-0.617 (0.017)	-0.271 (0.025)	
Exogenous Accumulation (Accumulation by Being)	-0.066 (0.031)	-0.025 (0.026)	
	Local Product Plants		
Demand shock	-0.573 (0.020)	-0.287 (0.029)	Excl.
Endogenous Accumulation (Accumulation by Doing)	-0.321 (0.017)	-0.276 (0.025)	
Exogenous Accumulation (Accumulation by Being)	-0.102 (0.027)	-0.027 (0.022)	
	Concrete Plants		
Demand shock	-0.467 (0.023)	-0.228 (0.035)	Excl.
Endogenous Accumulation (Accumulation by Doing)	-0.236 (0.019)	-0.159 (0.028)	
Exogenous Accumulation (Accumulation by Being)	-0.211 (0.043)	-0.066 (0.033)	

Notes: The results in this table are based on the estimates from Tables 3 and 4, respectively. The Demand Shock is computed as the difference between (log) output and the price determinants of demand. The endogenous learning effect is computed from the evolution of the demand capital for each plant using the estimated parameters for γ and δ . The learning by being effects are repeated from Tables 3 and 4 from the estimated Young and Medium age dummies. Recall “young” establishments first appeared in the census five years ago, “medium” establishments first appeared in the census ten years ago), and the excluded group “old” are establishments that first appeared in the census 15 or more years ago.

Table 7. Capacity Utilization Patterns Across Plant Ages and Multi-Unit Status

Capacity Utilization Measure	Variable	Plant age dummies				
		Entrant	Young	Medium	Old	Exiter
Capital	Utilization	-0.110 (0.019)	-0.086 (0.020)	-0.042 (0.022)	Excl.	-0.022 (0.018)
	Utilization x MU firm	-0.077 (0.021)	-0.087 (0.023)	-0.099 (0.025)	-0.111 (0.015)	0.055 (0.024)
Energy	Utilization	0.090 (0.029)	-0.021 (0.031)	-0.057 (0.033)	Excl.	0.074 (0.027)
	Utilization x MU 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 (“Capital”) and the log of energy use to equipment capital ratio (“Energy”). N = roughly 17,000 plant-year observations. Standard errors are in parentheses.

Appendix

A.1. Robustness Checks

Figure A.1 reports the estimates of the two key parameters for endogenous demand accumulation as the discount factor varies. The results reported are for the full sample but similar patterns hold for local plants and for concrete plants only (i.e., the parameter estimates are not very sensitive to the discount factor over this range).

Table A.1 reports the estimates when the impact of being part of a multi-unit firm upon entry is allowed to vary depending on whether the multi-unit firm has activity in the same industry or same geography. The results presented are for local product plants. The specification of (8) is modified as follows for this estimation:

$$Z_{0e} = (K_{0e})^{\lambda_1} \left(\frac{K_{0s(e)} + K_{0e}}{K_{0e}} \right)^{\lambda_2} \left(\frac{K_{0s(e)}^{Same}}{K_{0s(e)}} \right)^{\lambda_3}$$

where the “same” refers to same industry in the first column and same geographic area (BEA Economic Area) of Table A.1.

Table A.2 reports estimates when γ is permitted to vary with the product characteristics – specifically downstream product characteristics. The specification of γ is in this case:

$$\gamma = \gamma_0 + \gamma_k \text{downstream}_k$$

where downstream_k is the herfindahl index (based upon employment concentration of firms) of downstream industries in the first panel and is the firm turnover rate (sum of firm entry and exit rates) in the second panel. These measures were constructed using the input-output matrix to identify the downstream industries and then using the Longitudinal Business Database to measure concentration and firm turnover rates for these industries.

A.2 Demand Patterns by Firm and Establishment Age

Estimating the interactions between firm and plant age yields the results in Table A.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 as such are missing from the estimation (e.g., there cannot be a medium-aged plant in an entering or young firm). Old plants in single-unit firms are again the excluded group.

We focus on the multi-unit plant results in the bottom three rows of Table A.3. Starting with the bottom row, we see that among old firms (those that are at least 15 years old), the basic convergence patterns seen before hold here. 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 A.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.

A.3. 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 ten 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 3273000 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.

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.4. 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 Bailly, 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.5. 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, 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.²¹ Characteristics of the final sample can be seen in Table A.4.

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.4). 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 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

²¹ 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), block ice (37), processed ice (76), plywood (64), and sugar (90).

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.6. 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.4 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 boxes slightly 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.5 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 10.3 for plywood 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.

Appendix References

- Baily, Martin N., Charles Hulten, and David Campbell. "Productivity Dynamics in Manufacturing Establishments." *Brookings Papers on Economic Activity: Microeconomics*, 1992, 187-249.
- Davis, Steven J. and John Haltiwanger. "Wage Dispersion between and within U.S. Manufacturing Plants, 1963-1986." *Brookings Papers on Economic Activity, Microeconomics*, 1991, 115-180.
- Foster, Lucia, Haltiwanger, John, and C.J. Krizan. "Aggregate Productivity Growth: Lessons from Microeconomic Evidence." In *New Developments in Productivity Analysis*, eds. Edward Dean, Michael Harper and Charles Hulten, 303-363. Chicago: University of Chicago Press, 2001.
- 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.
- Roberts, Mark J. and Dylan Supina. 1996. "Output Price, Markups, and Producer Size." *European Economic Review*, 40(3-5), 1996, 909-921.
- Roberts, Mark J. and Dylan Supina. "Output Price and Markup Dispersion in Micro Data: The Roles of Producer Heterogeneity and Noise." In *Advances in Applied Microeconomics, Vol. 9, Industrial Organization*, ed. Michael R. Baye, 1-36. Greenwich: JAI Press, 2000.
- U.S. Bureau of Labor Statistics. "Trends in Multifactor Productivity: 1948-81." Bulletin 2178. Washington D.C.: Government Printing Office, 1983.

Table A.1. Contribution of Owning Firm in Same Industry or Geography: Local Product Plants

Parameter	Same Industry	Same Geography
γ (elasticity of future demand to the demand stock)	0.846 (0.010)	0.838 (0.011)
$-\eta$ (price elasticity of demand)	-1.701 (0.081)	-1.782 (0.090)
Young dummy (demand shift for entering and young plants)	-0.172 (0.029)	-0.159 (0.030)
Medium age dummy (demand shift for medium-aged plants)	-0.018 (0.022)	-0.049 (0.023)
ρ (persistence of exogenous demand shocks θ)	-0.058 (0.044)	-0.056 (0.045)
δ (demand depreciation rate)	0.777 (0.022)	0.772 (0.022)
Competitor's Price	0.324 (0.068)	0.308 (0.067)
λ_1 (elasticity of initial demand to plant's own K)	1.035 (0.023)	1.039 (0.022)
λ_2 (elasticity of initial demand to ratio of firm's K to plant's K)	0.379 (0.077)	0.470 (0.050)
λ_3 (elasticity of initial demand to ratio of firm's K in same industry or geography to firm's total K)	1.964 (0.625)	0.867 (0.128)
Inverse Mills Ratio, Demand (selection correction, demand equation)	-0.027 (0.010)	-0.028 (0.010)
Inverse Mills Ratio, EE (selection correction, Euler equation)	0.027 (0.004)	0.028 (0.004)

Note: Both columns report results for the joint estimation of demand and Euler equations using plant-year observations for local products. The only difference in specifications is the inclusion of a term in initializing Z_0 reflecting the ratio of firm's parent/sibling capital in the year of entry in the same industry (column 1) or same geography (column 2) to the overall firm's parent/sibling capital. See text of the appendix for details.

Table A.2. Allowing Elasticity of Future Demand to Current Demand Stock (γ) to Vary with Product Characteristics

Parameter	Full Sample	Local Products
Downstream Concentration Interactions		
γ_0 (elasticity of future demand to the demand stock)	0.734 (0.015)	0.847 (0.010)
γ_H (Interaction with downstream demand concentration)	0.098 (0.020)	-0.005 (0.005)
Downstream Turnover Interactions		
γ_0 (elasticity of future demand to the demand stock)	0.770 (0.015)	0.837 (0.010)
γ_T (Interaction with downstream demand concentration)	-0.011 (0.009)	-0.001 (0.006)

Note: The reported estimates are from specifications where γ is specified to vary with downstream product characteristics. Each column and panel represents a separate estimation of the joint demand and Euler equation. The interaction with product characteristics is specified so that the reported γ holds for a product with mean product characteristics and the interaction effect captures any changes in the parameter as a product characteristic deviates from the mean. Downstream concentration is measured by the herfindahl index and downstream turnover is based on the sum of the entry and exit rate of the downstream industries. The latter were identified using the input-output matrix. See text of appendix for details.

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

Variable	Plant Age Dummies				
	Entrant	Young	Medium	Old	Exiter
Demand Shock	-0.317 (0.034)	-0.178 (0.036)	-0.147 (0.036)	Excl.	-0.183 (0.031)
Demand Shock x MU firm and <i>entrant</i>	0.168 (0.066)	N/A	N/A	N/A	-0.167 (0.110)
Demand Shock x MU firm and <i>young or medium</i>	0.004 (0.074)	0.139 (0.044)		N/A	-0.120 (0.077)
Demand Shock x MU firm and <i>old</i>	0.091 (0.042)	0.122 (0.044)	0.267 (0.048)	0.538 (0.026)	-0.332 (0.045)

Note: This table repeats the analysis of Table 1, but now allows plant age effects to vary with the multi-unit (MU) status and age of the plant’s owning firm. The excluded category is “old” plants (those that appeared three or more censuses prior) that are part of a single-unit firm. N = roughly 17,000 plant-year observations. Standard errors, clustered by plant, are in parentheses.

Table A.4: 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
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 BEA Economic Areas (see text for details) based on data pooled over all years.

Table A.5: Entry and Exit Rates by Product

Products	Entry Rates	Exit Rates
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
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.

Figure A.1

