Productivity Dispersion and Plant Selection in the Ready-Mix Concrete Industry*

Allan Collard-Wexler †‡
Economics Department, NYU Stern and NBER

February 28, 2011

Abstract

This paper presents a quantitative model of productivity dispersion to explain why inefficient producers are slowly selected out of the ready-mix concrete industry. Measured productivity dispersion between the 10th and 90th percentile falls from a 4 to 1 difference using OLS, to a 2 to 1 difference using a control function. Due to volatile productivity and high sunk entry costs, a dynamic oligopoly model shows that to rationalize small gaps in exit rates between high and low productivity plants, a plant in the top quintile must produce 1.5 times more than a plant in the bottom quintile.

*JEL Keywords: L13, L6, D24.

†This paper is part of my Ph.D. dissertation at Northwestern University under the supervision of Robert Porter, Michael Whinston, Aviv Nevo and Shane Greeenstein. I thank Lynn Riggs, Ginger Gin, Jan de Loecker, Chad Syverson and Pat Bajari for their insights, as well as many seminar participants. Financial support from the Fonds Québecois sur la Recherche et la Société et Culture and the Center for the Study of Industrial Organization at Northwestern University is gratefully acknowledged. The research in this paper was conducted while the author was a special sworn status researcher of the U.S. Census Bureau at the Chicago and New York Census Research Data Centers. Research results and conclusions expressed are those of the author and do not necessarily reflect the views of the Census Bureau. This paper has been screened to insure that no confidential data are revealed. Support for this research at the Chicago and New York RDC from NSF (awards no. SES-0004335 and ITR-0427889) is also gratefully acknowledged.

‡wexler@stern.nyu.edu

1 Introduction

A society's ability to provide for its members is determined in large part by the productivity of its plants. A presumption shared by many economists is that plant-level efficiency is governed by common inputs such as technology and the availability of capital or educated workers. Yet there can be considerable differences in plants' productivities for an identical product within the same market. For instance, in the ready-mix concrete industry, a plant in the 90^{th} percentile of productivity has more than 4 times the value added as one in the 10^{th} percentile if both plants use the same inputs.[1]

Why do we observe such a large dispersion in productivity? Moreover, why do these differences persist over time? Indeed, productivity differences are puzzling, since selection by exit should drive out inefficient plants from the market. For instance, suppose there is an incumbent plant with low productivity, $\underline{\rho}$, then it should be replaced by an entrant with high productivity, $\bar{\rho}$. Thus, in the long run, selection eliminates low-productivity plants.

I find some support in the data for the selection mechanism. In the ready-mix concrete industry, the exit rate of a plant in the top quintile of productivity is 3%, while the exit rate of a plant in the bottom quintile of productivity is 7%. Yet it is difficult to reconcile large differences in productivity with such a small gap in exit rates.

Two forces blunt the selection effect. First, sunk costs insulate inefficient incumbents. These sunk costs create a gap between the minimum level of productivity that allows a plant to enter the market (say, $\underline{\rho}^E$) and the lower level at which an incumbent exits (say, $\underline{\rho}^I$). An inefficient producer will remain active at a productivity level for which it would never have considered entering, if productivity is $\rho \in (\underline{\rho}^I, \underline{\rho}^E)$.

Second, productivity differences do not persist over time. At the extreme, if productivity is independent across time, then a plant's efficiency has no bearing on its decision to remain in the market, since current productivity provides no information about future productivity. In general, the lower the persistence of productivity, the less information is provided by current productivity about future profits.

Thus, the question I address in this paper is: To what extent do sunk entry costs and productivity volatility rationalize large differences in plants' productivity mapping into a

[1] Productivity is defined as the residual in the regression of log value added on log salaries and log total assets with year dummies. Thus, it is a measure of the profitability of input use.

small selection effect? To answer this question for the ready-mix concrete industry, I build and estimate a dynamic model of entry and exit that incorporates competitive considerations and the evolution of plant-level productivity.

The first part of the empirical model is concerned with the measurement of productivity. In particular, this part of the model determines the extent to which productivity dispersion is real and not simply an artifact of measurement error. I use a control function technique introduced by Ackerberg, Frazer, and Caves (2006) to separate true productivity dispersion from measurement error. I find that productivity dispersion is substantially lower once I filter out measurement error: a plant in the 90^{th} percentile of productivity only produces twice the value added as a plant in the 10^{th} percentile, in contrast to the 4-to-1 ratio when I simply run OLS.

Second, the model uses entry, exit and investment decisions to estimate the effect of low productivity on plant profitability. Here I use a dynamic oligopoly model to account for entry decisions and the localized market structure of the ready-mix concrete industry. I find that to rationalize observed investment, entry and exit decisions, the profits of a plant in the bottom quintile of productivity must be $ $220,000$ lower than those of a plant in the top quintile (where the median plant has approximately $ $550,000$ of value added each year). This is a substantial effect of productivity, as it implies that a plant in the 90^{th} percentile produces 1.5 times more value added than a plant in the 10^{th} percentile (if both plants use the same inputs)—only somewhat less than the 2-for-1 difference observed in the data.

Third, I conduct counterfactuals simulation to evaluate the role of sunk entry costs and productivity volatility on equilibrium exit decisions. I find that if productivity were perfectly persistent, then plants' decisions to exit would be seven times more responsive to productivity. In contrast, a reduction in sunk costs of 20% raises entry and exit rates, but does not increase plants' responsiveness to productivity differences.

In Section 2, I review the literature on plant selection and oligopoly dynamics. Section 3 presents the model of competition used for this paper. In Section 4, I describe the data on ready-mix concrete plants, and Section 5 discusses the measurement of productivity. Section 6 provides basic empirical evidence for plant selection. Section 7 presents the estimation of the model of entry and exit, while Section 8 discuses results. Section 9 presents the counterfactuals experiments.

2 Literature

In this paper I draw on three literatures. The first concerns the evolution of productivity at the plant level, the second focuses on structural estimation of productivity, and the third addresses the estimation of models of dynamic oligopoly models.

2.1 Productivity and Plant Selection

Theoretical models of industry dynamics are explored by Jovanovic (1982) and Hopenhayn (1992), who study the effect of firms learning about their productivities on the entry and exit process and an industry's steady-state.

Syverson (2004) documents productivity dispersion in the ready-mix concrete industry using data from the U.S. Census Bureau. Productivity is defined as the residual of the regression of log output on log salaries, log assets and log materials, where the coefficients on inputs are simply the input's cost shares. The magnitude of productivity dispersion is robust to several different measures of productivity, including defining output as total cubic yards of ready-mix concrete produced. Syverson conjectures that competition plays a key role in eliminating unproductive plants, which limits the dispersion of productivity. The empirical evidence to support this conjecture is the distribution of productivity in large and small markets, where market size is determined by the density of construction activity. Productivity is both more dispersed and lower in in small markets. Moreover, there is a smaller share of low-productivity plants in large markets than small ones. Competition appears to truncate the distribution of productivity by driving out inefficient plants. Instead of focusing on the cross-sectional implications of plant selection considered by Syverson (2004), this paper explores the mechanism for plant selection in detail.

Foster, Haltiwanger, and Syverson (2008) investigate the role a plant's profitability and productivity play in its decision to exit. Profitability differs from productivity since a plant in a concentrated market experiencing high demand can make large profits without being particularly good at producing ready-mix concrete. Foster, Haltiwanger, and Syverson (2008) find that as plants age, their technical efficiency falls, but their prices rise.

Dunne, Klimek, and Roberts (2005) also look at the entry and exit decisions of several geographically differentiated producers (including ready-mix concrete) and find that plants built by firms with previous industry experience have lower exit rates than those of newer entrants. It is difficult, however, to gauge if these other characteristics of firms can explain the dispersion of productivity present in the data examined in this paper.

2.2 Structural Estimates of Productivity

To estimate production functions, it is necessary to account for two biases. First, firms may observe their productivity shock before choosing inputs, and this leads to a correlation between flexible inputs chosen by firms and productivity. This is known as the simultaneity problem. Simultaneity will lead us to overstate the importance of flexible inputs, such as labor, which will be correlated with the productivity shock. As a result, the importance of more permanent inputs, such as capital, may be understated. Second, the distribution of productivity is truncated since low-productivity firms are more likely than higher-productivity firms to shut down. Thus, firms that have a higher likelihood of exiting the industry (e.g., small firms) also could have a more selected productivity distribution. To correct for both the simultaneity and selection problems, Olley and Pakes (1996) propose a control function correction. This procedure assumes that unobserved productivity shock is a function of the firm's investment decision, conditional on the state it is in. For instance, more productive firms will invest more. If investment is put into a production function regression, it is possible to control for factors that are correlated with higher productivity but which should not lead directly to higher production. Levinsohn and Petrin (2003) extend the Olley and Pakes (1996) approach, using material inputs, rather than investment, as a proxy. These material input controls have the advantage of allowing more continuous variation in the data. In contrast, investment data are quite lumpy. Ackerberg, Frazer, and Caves (2006) propose an integrated framework for thinking about control function estimates of production functions using either material or investment controls or the literature on dynamic panel models. I use this approach in my estimates for production functions since it offers more flexibility in specifying the moments conditions that I use. Moreover, I take the control function approach very literally, in that I back out and analyze the firm's "true productivity".

2.3 Estimation of Dynamic Multi-Agent Models

Models of dynamic oligopoly pose daunting econometric challenges that require specialized solution techniques. The framework for empirical models of dynamic oligopoly was developed by Ericson and Pakes (1995), who incorporate the solution concept of the Markov-Perfect Equilibrium (Maskin and Tirole (1988)). To bring this framework to data, the econometrics of dynamic discrete choice (e.g., Rust (1987)) can be used to estimate model parameters, given the choices of firms. However, Rust's Nested Fixed Point algorithm is intractable for estimating all but the simplest dynamic game. Finding an equilibria of a dynamic game is computationally intensive, since an equilibria is a fixed point in both the agent's value function and its best-response policy, given that other players are also playing a best-response. Benkard, Bajari, and Levin (2007), Pesendorfer and Schmidt-Dengler (2008), Pakes, Berry, and Ostrovsky (2007) and Aguiregabiria and Mira (2007) develop techniques for estimating the parameters of a dynamic game without the computational burden associated with computing the solution to the game for each parameter vector.

3 Model

In this section, I describe the model of industry dynamics used both to measure productivity and gauge the effect of plant selection. This model closely follows Ericson and Pakes (1995). There are N firms in each market, and these firms can be either potential entrants, denoted $s_i^E = \{\epsilon_i^E\}$, or incumbent firms, whose state, s_i, is composed of:

$$s_i = \{k_i, \omega_i, \epsilon_i\} \tag{1}$$

where k_i is capital stock, ω_i is the firm's productivity, and ϵ_i^E and ϵ_i are full support i.i.d. private information shocks to the profitability of either entering or exiting the market.

The market level state, $s = \{s_1, \cdots, s_N, M\}$, is the composition of all firm-level states for both incumbents and potential entrants, where M denotes total demand in the market. In particular, M will be generated by demand for ready-mix concrete by the construction sector. The important feature of this model is that firms need to keep track of the states of their competitors, since other firms' productivity and capital assets will affect their own profits, $\pi(s)$.

3.1 Actions

Firms make two choices: entry/exit and investment. In each period, potential entrants choose to enter a market or not. I denote the entry decision as $\chi_i^E \in \{0,1\}$ and the entry cost as $\phi^E + \epsilon_i^E$. Likewise, incumbent firms can choose to exit or stay in the market. I denote the exit choice as $\chi_i \in \{0,1\}$ and the fee paid by firms that choose to exit as $\phi + \epsilon_i$. Firms also choose how much capital to purchase or sell, denoted as i_i, with investment costs denoted as $c(i)$. Subsequently, capital evolves to $k_i' = \delta k_i + i_i$, where δ is the depreciation rate of capital.

Figure 1 illustrates the timing assumptions of the model within each period. First, the unobserved states ϵ_i and ϵ_i^E are privately observed by each firm. Second, firms simultaneously choose whether or not to operate a ready-mix concrete plant in the next period and, if so, how much to invest in the plant.

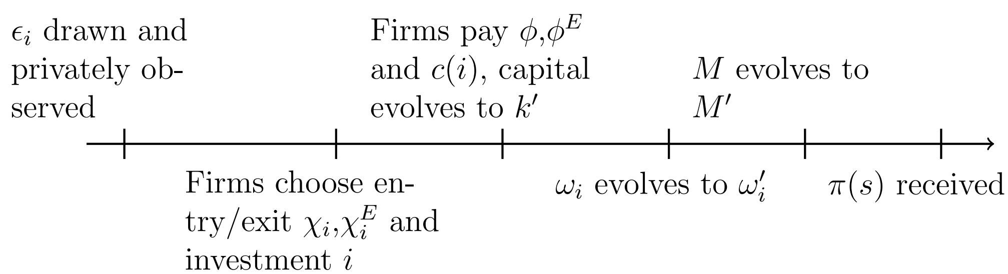

Figure 1: Timing of the Game within Each Period

Subsequently, productivity and demand evolve to their new levels. Productivity, ω_i, follows a first-order Markov process, i.e., $\omega_i' \sim f^{\omega}(\cdot|\omega_i)$, or $\omega_i' \sim f^{E\omega}(\cdot)$ for potential entrants. I make two important assumptions regarding the evolution of productivity. First, I assume that firms *cannot* control their productivity. In my model, productivity refers to total factor productivity, not to output per worker. Thus, the purchase of better machines increases the firm's capital stock, not its productivity residual. Second, in contrast to the model used by Jovanovic (1982), I assume that firms do not learn about their productivity as they age.[2] Overall market demand, M, evolves following a first-order Markov Process, i.e., $M' \sim f^M(\cdot|M)$, due to changes in the demand for ready-mix concrete from the construction

[2] For the ready-mix concrete industry, the limited effect of age on exit rates discussed in more detail in Appendix D indicates a limited role of learning.

sector. Finally, firms receive period profits, $\pi(s)$. The value function for incumbent firms is thus:

$$V(s) = \max_{\chi \in \{0,1\}} \chi(\psi + \epsilon_i) + (1-\chi)\left(\max_i \pi(s) - c(i) + \beta \int_{s'} V(s') f(s'|s,i) ds'\right) \tag{2}$$

where $f(s'|s,i)$ is the transition density of the state, as determined by the evolution of capital and productivity for all firms, the entry and exit decisions of firms, and the evolution of construction activity, M:

$$f(s'|s,i) = f^M(M'|M) \prod_{i=1}^{N} f^\omega(\omega_i'|\omega_i) f(k_i'|k_i, i_i, \chi_i) f^\epsilon(\epsilon_i') \tag{3}$$

Firms choose investment, i, to solve:

$$i^*(s) = \text{argmax}_i \pi(s) - c(i) + \beta \int_{s'} V(s') f(s'|s,i) ds' \tag{4}$$

Thus, the exit rule for firms is:

$$\chi(s) = 1\left(\pi(s) - c(i^*) + \beta \int_{s'} V(s') f(s'|s,i^*) ds' > \phi + \epsilon_i\right) \tag{5}$$

Likewise, the value function for potential entrants is:

$$V^E(s) = \max_{\chi^E \in \{0,1\}} \chi^E \left(\phi^E + \epsilon_i^E + \int_{s'} V(s') f(s'|s,i) ds'\right) \tag{6}$$

Thus, the entry rule for firms is:

$$\chi^E(s) = 1\left(\int_{s'} V(s') f(s'|s,i^*) ds' > \phi^E + \epsilon_i^E\right) \tag{7}$$

Note that the policy functions $\chi^E(s)$, $\chi(s)$ and $i^*(s)$ define a Markov Perfect Equilibrium in that these policy functions are optimal, given the transition density, $f(s'|s)$. But of course the transition density of s' is determined by the entry, exit and investment rules of all other firms in the market. Thus, the equilibrium is defined by two conditions: (1) the set of policies that are optimal, given $f(s'|s,i)$ and (2) the transition density, $f(s'|s,i)$, that occurs, given the policies $\chi^E(s)$, $\chi(s)$ and $i^*(s)$.[3]

3.2 Period Profits

Firms compete in quantities as in Cournot, given capital stock, k_i, and productivity, ω_i. There is no simple closed form for the solution to this game. Furthermore, estimating a Cournot or Bertrand model requires that I take a stance on what generates the considerable dispersion in prices for ready-mix concrete within a given market and year. For these reasons, I will use a reduced-form profit function to approximate the equilibrium profit of the period game.

4 Data

4.1 Entry and Exit

Data on Ready-Mix Concrete plants is drawn from three different data sets provided by the Center for Economics Studies at the United States Census Bureau.[4] Table 1 illustrates the datasets used. The first is the Census of Manufacturing (henceforth CMF), a complete census of manufacturing plants, every five years from 1963 through 1997. The second is the Annual Survey of Manufacturers (henceforth ASM) sent to a sample of manufacturing plants (about a third for ready-mix) every non-Census year since 1973. Both the ASM and the CMF involve questionnaires that collect detailed information on a plant's inputs and outputs. The third data set is the Longitudinal Business Database (henceforth LBD) compiled from data used by the Internal Revenue Service to maintain business tax records. The LBD covers all private employers on a yearly basis since 1976. The LBD only contains employment and salary data, along with sectoral coding and certain types of business organization data such as firm identification. Construction data is obtained by selecting all establishments from the LBD in the construction sector (SIC 15-16-17) and aggregating them to the county level.

Production of ready-mix concrete for delivery predominantly takes place at establishments in the ready-mix sector corresponding to either NAICS (North American Industrial

[3] Doraszelski and Satterthwaite (2010) show conditions under which an equilibrium exists for this model of industry dynamics.

[4] In Collard-Wexler (2008) I discuss the construction of entry and exit data in further detail.

	CMF	ASM	LBD
Collection	Questionnaire	Questionnaire	IRS Tax Data
Years	Every 5 years	1972-2000	1976-1999
Entry/Exit/Payroll	70%	30%	X
Input and Output Data	70%	30%	X

Table 1: Description of Census Data Sources

Classification) code 327300 or SIC (Standard Industrial Classification) code 3273.

4.2 Longitudinal Linkages

To construct longitudinal linkages, I use the Longitudinal Business Database Number developed by Jarmin and Miranda (2002). To identify plant entry and exit, I use plant birth and death measures, also developed by Jarmin and Miranda (2002), which identify entry and exit based on the presence of a plant in the I.R.S.'s tax records.

Over the sample period from 1963 to 2000, there are about 350 plant births and 350 plant deaths each year. In comparison, during the same period, 5000 incumbent plants ("continuers") continued their operations. Turnover rates and the total number of plants in the industry remained fairly stable over the last 30 years. As shown in Table A4 in Appendix C, the average ready-mix concrete plant employed 26 workers and sold about 3.2 million dollars of concrete in 1997. Revenue was split evenly between material costs and value added. However, these averages mask substantial differences between plants. Most notably, the distribution of plant size is heavily skewed, with few large plants and many small ones. This is indicated by the fact that more than 5% of plants have 1 employee, while less than 5% of plants have more than 82 employees. Moreover, as Table A3 (also in Appendix C) shows, continuing firms are twice as large as either entrants ("births") or exitors ("deaths"), as measured by capitalization, salaries or shipments.

4.3 Measuring Productivity

To measure productivity, I use information about a plant's inputs and outputs from the Annual Survey of Manufacturing (ASM) and the Census of Manufacturing (CMF). However, these do not provide annual data for all ready-mix concrete plants. The ASM questionnaire samples only about $\frac{1}{3}$ of ready-mix concrete producers each year.[5]

In contrast, the CMF questionnaire, which is sent to all manufacturing plants every 5 years, provides data on all plants in each market. These data can be used to look at the relationship between competition and productivity. However, since the CMF samples plants every 5 years, I can only look at a firm's entry/exit decision in the year following a Census year.

A large fraction of input and output data in the ASM and the CMF is imputed by the Census Bureau. (I discuss the full details of the imputed data in Appendix A). There are two important points to note about the imputed data. First, the data primarily are missing due to the sampling design of the Economic Census and the ASM; i.e. the probability a plant is sampled depends on it's size. Thus, unlike many settings in economics, I know that the selection into the sample is a function of employment alone and is not correlated with other characteristics of the plant.[6] Second, the goal of the imputation techniques used by the Census Bureau is to produce correct statistics for the aggregate economy. Thus, while we expect these imputes to yield the correct cross-sectional average for firm productivity, these imputes were never meant to study either plant-level productivity and exit, or plant-level productivity dynamics.

Plant efficiency plays a large role in the decision to continue operating, but plants do not report productivity directly. Instead, productivity has to be inferred based on a plant's reported outputs and inputs. I estimate total factor productivity (TFP) as the residual from

[5]My model of productivity and competition requires the productivities of *all* plants in a market, since the productivities of a plant's competitors are an important component of the model. Since the ASM samples one-third of plants in the ready-mix concrete industry, the probability that I have data on all plants in a market is decreasing in the number of plants in a market. Thus, the sample of markets is severely truncated in ASM years. For these reasons, I do not use ASM data in the estimation of the dynamic model; instead I rely on data from the CMF.

[6]To correct for the selection bias inherent to the imputation of smaller firms, I have run regressions where I adjust the estimates using propensity score weights. Since I observe all plants in the LBD, I compute the weights via a Naradaya-Watson kernel regression of the probability an observation has missing data on employment for each year:

$$\beth_{it} = \hat{w}_t(\text{Missing}, l_{it}) \tag{8}$$

There is very little difference in the estimates calculated using propensity score weighting, with the exception that exit rates are somewhat higher due to the fact that small plants, which are more likely to exit, are undersampled.

the log-linear production function OLS regression:

$$y_{it}(\text{value added}) = \beta_l l_{it}(\text{salaries}) + \beta_k k_{it}(\text{capital}) + \delta_t + \rho_{it} \tag{9}$$

where a lower-case variable that denotes the logarithm of the actual variable, δ_t, is the intercept of the production function for each year (so that year to year changes in technology do not affect the dispersion of productivity), and ρ_{it} is a plant's TFP. I deflate all items measured in dollars by the producer price index (PPI).[7]

In Appendix E, I use different measures of output (e.g., value added, cubic yards of concrete or total shipments) to estimate the total dispersion of productivity. I use a revenue measure of output, since I want to explain why unprofitable firms do not exit the ready-mix concrete industry, rather than why firms that produce little concrete per unit of input stay in the industry. Firms that can charge a higher markup, due to either better quality or market power, are more likely to stay in the industry.[8]

5 Control Function Estimates of Productivity

Productivity is typically estimated using the log-linear or Cobb-Douglas production function in value added terms:[9]

$$y_{it}(\text{value added}) = \beta_0 + \beta_l l_{it} + \beta_k k_{it} + \rho_{it} \tag{10}$$

[7]It is important to deflate data measured in dollars since the log-linear production function is not linearly homogeneous if the sum of the capital and labor coefficients differs from 1.0, and, thus, is sensitive to rescaling variables.

[8] De Loecker (2007) and Foster, Haltiwanger, and Syverson (2008) argue about the importance of separating technical efficiency from markups in the measurement of productivity. I am conflating these two sources of profitability into a single index ω. Note that this assumption is incorrect to the extent that the sum of efficiency and markups does not itself follow a Markov Process, which can be the case even if both efficiency and markups each follow a Markov process. For instance, Das, Roberts, and Tybout (2007) estimate a dynamic model with multiple components of profitability.

[9]Researchers typically measure output in terms of the value added by total shipments. In a previous version of this paper, I used shipments as my measure of output. However, the assumption that materials, such as sand, cement and gravel are Leontieff in the production function (i.e., perfect complements) makes more sense than assuming that these inputs have a constant elasticity of substitution with labor and capital, which, for example, is the case when one uses a Cobb-Douglas production function in shipments. In any case, I expect to find *more* dispersion when I use production functions in shipments than when I use production functions in value added. As well, I have estimated equation (10) using a translog production function and find similar results to what is presented in this paper.

where estimated firm productivity is ρ_{it}, the "unexplained" component of total sales. Estimated productivity, ρ_{it}, conflates true differences in productivity with errors in the measurement of either inputs or outputs. In particular, since by definition measurement error is uncorrelated with "true productivity" differences, then TFP will overestimate the degree of productivity dispersion in an industry.

However, it is also incorrect to assert that measurement error is responsible for all measured productivity dispersion. Since productivity and exit decision are correlated with measured TFP, this rejects the assertion that TFP is pure measurement error.

The production function that both Olley and Pakes (1996) and Ackerberg, Frazer, and Caves (2006) consider is the following:

$$y_{it} = f(l_{it}, k_{it}) + \omega_{it} + \epsilon_{it} \tag{11}$$

The goal is to separate "true" productivity differences between firms, denoted as ω_{it}, from measurement error denoted, as ϵ_{it}.

5.1 First Stage

To identify the extent of measurement error, I impose structure on the way firms make their investment choices. Suppose that the firm's state, s_{it}, is composed of both the firm's capital stock, k_{it}, the firm's "true" productivity level, ω_{it}, and other states, such as the level of demand in the market or the number of competitors in a market, which I refer to as x_{it}.[10]

$$s_{it} = \{k_{it}, \omega_{it}, x_{it}\} \tag{12}$$

Suppose that either investment (i_{it}) or labor input demand functions are strictly increasing in ω_{it}, conditional on the rest of the state ($\{k_{it}, x_{it}\}$). I can rewrite the investment function as:

$$i_{it} = i(s_{it}) = i(k_{it}, \omega_{it}, x_{it}) \tag{13}$$

[10]Note that I cannot accommodate the productivity of other plants, ω_{-i}, into the state x_{it}.

Under the assumption that $i(\cdot)$ is strictly increasing in ω, then this function can be inverted:[11]

$$\omega_{it} = i^{-1}(i_{it}|k_{it}, x_{it}) = h(i_{it}, k_{it}, x_{it}) \tag{14}$$

It is then possible to replace ω_{it} in equation (11) with $h(i_{it}, k_{it}, x_{it})$. This yields:

$$\begin{aligned} y_{it} &= f(l_{it}, k_{it}) + h(i_{it}, k_{it}, x_{it}) + \epsilon_{it} \\ &= \phi(l_{it}, k_{it}, i_{it}, x_{it}) + \epsilon_{it} \end{aligned}$$

I can identify the extent of measurement error in the first stage by performing a non-parametric regression of the log of sales on labor, materials and capital. Table 2 presents the first-stage, non-parameteric regression. I find that that the capital coefficient is far higher when I only use Census years, rather than both Census and ASM years. Since I find far more volatility in the capital stock in ASM years than Census years, I will only use Census years to estimate the production function, as the CMF suffers from fewer data quality issues.

I denote the estimated $\phi(\cdot)$ function as $\hat{\phi}(\cdot)$, which can be though of as the component of output that can be accounted for by the production function and "true productivity". The gap between the output and the production function is the measurement error component of TFP, which can be computed as:

$$\epsilon_{it} = y_{it} - \hat{\phi}(l_{it}, k_{it}, i_{it}, x_{it}) \tag{15}$$

Notice that the essential difference between the $\phi(\cdot)$ function and the production function $f(\cdot)$ is that ϕ includes not only inputs into the production function but also other variables that should be correlated with higher productivity but not with unmeasured inputs into the production process.

5.2 Second Stage

In the second stage, I recover a plant's "true productivity." Suppose that a plant's true productivity follows a first-order Markov process. Then next year's productivity is generated

[11] Olley and Pakes (1996) provide conditions under which the investment policy will be strictly increasing in productivity. If the market is imperfectly competitive, then it is no longer true that the investment function must be strictly increasing.

Dependant Var: Log Value Added	Production Function		First Stage ACF Regressions		
	Census Years	All Years	Census Years		All Years
	I	II	III	IV	V
Constant	1.148	1.342	1.039	1.021	1.143
	(0.05)	(0.05)	(0.05)	(0.07)	(0.05)
Log Salaries	0.675	0.747	0.660	0.659	0.733
	(0.01)	(0.01)	(0.01)	(0.03)	(0.01)
Log Assets	0.258	0.162	0.229	0.243	0.127
	(0.01)	(0.01)	(0.01)	(0.03)	(0.01)
Log Investment			0.042	0.039	0.055
			(0.01)	(0.01)	(0.01)
Zero Investment			0.162	0.179	0.201
			(0.03)	(0.04)	(0.02)
1 Competitor			-0.104	-0.139	-0.072
			(0.04)	(0.06)	(0.03)
2 Competitors			-0.124	-0.169	-0.087
			(0.04)	(0.06)	(0.03)
3 Competitors			-0.074	-0.127	-0.076
			(0.04)	(0.06)	(0.04)
4 Competitors			-0.081	-0.117	-0.047
			(0.04)	(0.06)	(0.04)
5 Competitors			-0.043	-0.06	-0.041
			(0.03)	(0.04)	(0.03)
More than 5 competitors			-0.007	-0.022	0.015
			(0.01)	(0.02)	(0.01)
Multi-Unit Firm			0.178	0.162	0.177
			(0.01)	(0.02)	(0.01)
Construction Employment			-0.003	-0.001	-0.003
			(0.00)	(0.00)	(0.00)
Squared and Cube Interactions of all above variables				X	
Year Effects	X	X	X	X	X
Observations	11097	15637	8499	8499	11814
R^2 Adjusted	0.83	0.81	0.85	0.89	0.83

Standard Errors clustered by plant.

Table 2: Ackerberg, Frazer, and Caves (2006) -First Stage and Production Function Estimates

by:

$$\omega_{it} = g(\omega_{it-1}, \hat{\chi}_{it}) + \xi_{it} \tag{16}$$

where ξ_{it} is the innovation to today's productivity and $\hat{\chi}_{it}$ is the exit propensity score computed by running a probit of the exit choice on x_{it} and k_{it} (shown in column I of Table 6 on page 23). To compute ξ_{it}, I perform a second-order series regression of ω_{it} on ω_{it-1} and $\hat{\chi}_{it}$, i.e. $\omega_{it} = \hat{g}(\omega_{it-1}, \chi_{it}) + \xi_{it}$ (shown in column VII of Table 5 on page 20).

Since ξ_{it} is unobserved by the firm at the time at which it decides how much to invest, then ξ_{it} and k_{it} should be uncorrelated. Likewise, since a firm chooses labor based on today's ω_{it}, rather than next period's productivity draw, then ξ_{it} should be uncorrelated with today's labor input. When I combine these moment conditions together, I obtain:

$$\mathbf{E} \;\; \xi_{it} \begin{pmatrix} l_{it} \\ k_{it+1} \end{pmatrix} = 0 \tag{17}$$

This, in turn, allows us to form an analogue estimator for this moment condition using a GMM criterion. First, stack the data as:

$$\mathbf{X}(\beta) = \begin{pmatrix} \vec{\xi}_{1t}(\beta) \\ ... \\ \vec{\xi}_{Nt}(\beta) \end{pmatrix}$$

$$\mathbf{Z} = \begin{pmatrix} \vec{l}_t \\ \vec{k}_{t+1} \end{pmatrix}$$

The GMM criterion using the weighting matrix $(\mathbf{Z}'\mathbf{Z})^{-1}$ is:

$$Q(\beta) = (\mathbf{X}(\beta)'\mathbf{Z})(\mathbf{Z}'\mathbf{Z})^{-1}(\mathbf{X}(\beta)'\mathbf{Z})' \tag{18}$$

I find that β_l and β_k minimize the GMM criterion, $Q(\beta)$. These parameters are presented in Table 3.[12]

[12] I find little difference in the estimated capital coefficient between the OLS and the ACF estimates, which indicates that the simultaneity problem is not a large problem in the ready-mix concrete data used here. However, I still find that "true productivity," or ω dispersion, is much lower once I eliminate measurement error.

Dependent Variable: Log Value Added	All Years	Census Year
	I	II
Log Salaries ($\hat{\beta}_l$)	0.938	0.682
	(0.270)	(0.059)
Log Assets ($\hat{\beta}_k$)	0.015	0.260
	(0.09)	(0.038)
Observations	11862	8521

Standard Errors for the ACF estimator are computed via 1000 bootstrap replications. Notice that to compute the bootstrap replications, I need to (1) compute the probit on exit and the associated exit propensity score, $\hat{\chi}_{it}$, (2) perform the first-stage regression and (3) minimize the GMM-Criterion, $Q(\theta)$. I also do block-bootstrap by resampling a plant's entire history to cluster the standard errors against serial correlation.

Table 3: Ackerberg, Frazer, and Caves (2006) Estimates of Productivity

Then "true productivity" can be computed as:

$$\omega_{it} = \hat{\phi}_{it} - \hat{\beta}_l l_{it} - \hat{\beta}_k k_{it}$$

5.3 Measures of Dispersion

My main goal is to explain the dispersion of plant-level efficiency, i.e., the fact that plants that use the same bundle of inputs have different levels of output. Dispersion can be gauged by the values for R^2 in the production function regressions in Table 2. To express the relative dispersion of productivity, I look at the dispersion that would occur if all plants used the same bundle of inputs but brought their own productivity residual. This technique allows me to look at the sources of TFP dispersion, while eliminating differences in output that are due solely to differences in the use of inputs.[13]

The predicted output of a plant that uses the median levels of capital, and labor (denoted as k_{50} and l_{50} respectively), but its q^{th} *quantile* of the productivity residual can be computed as:

$$\hat{y}^{\rho_q} = \beta_l l_{50} + \beta_k k_{50} + \rho_q$$

[13] I implicitly am foreclosing the discussion on productivity differences that are due to the use of inefficient bundles of inputs.

Likewise, I can obtain the predicted output of a plant that brings its q^{th} quantile of "true productivity" but has the median level of measurement error:

$$\hat{y}^{\omega_q} = \beta_l l_{50} + \beta_k k_{50} + \omega_q + \epsilon_{50}$$

Finally, I obtain the contribution of measurement error to productivity dispersion by looking at the quantiles of output when I change measurement error:

$$\hat{y}^{\epsilon_q} = \beta_l l_{50} + \beta_k k_{50} + \omega_{50} + \epsilon_q$$

Table 4 presents the dispersion of value added due to TFP dispersion, true productivity and measurement error. The top panel shows only the dispersion of TFP (ρ), measurement error (ϵ) and true productivity (ω). Notice that the standard deviation of TFP (ρ) is 1.5 times higher than the dispersion of productivity (ω). The bottom panel shows that the interquartile range for TFP is about \$ 350 000, or 63 % of median value added. However, when we look at the dispersion due to true productivity, the interquartile range is only \$ 140 000, or 25% of median value added. Thus, the interquartile dispersion of measurement error is \$ 270 000, or 50 % of median value added. This indicates that the dispersion of TFP is due to an equal mixture of true productivity differentials and measurement error. (Note that the sum of productivity and measurement error dispersion will be higher than TFP dispersion if these two are negatively correlated.)

5.4 Persistence of Productivity

Current productivity provides two pieces of information to a ready-mix concrete plant. Low productivity reduces current profits, since the plant produces less concrete for a given level of inputs. Moreover, if productivity is persistent, low productivity also implies lower profits in the future. Thus, greater persistence of productivity makes current productivity a better indicator of future profits. Table 5 shows the 1 and 5 year autocorrelation of productivity, along with the one and five year autocorrelations of TFP and measurement error.

First, note that both productivity and TFP have fairly low autocorrelations, with coefficients on last year's variables of 0.58 and 0.55, respectively, which correspond to autocorrelations of approximately 70%. Measurement error also is serially correlated, but less so than

Variable	Observations	Mean	Std.
Log Value Added	11814	6.36	1.39
Predicted Output $\hat{\phi}$	11862	6.33	1.28
TFP (ρ)	11814	0.13	0.71
Productivity (ω)	11862	1.06	0.46
Measurement Error (ϵ)	11814	0.03	0.60

Note: Productivity, Measurment Error and Predicted Output computed using the Ackerberg, Frazer, and Caves (2006) technique. TFP is the residual from a Cobb-Douglas production function estimated with OLS.

	Value Added in thousands of dollars due to		
Percentile	Productivity (ω_q)	Measurement (ϵ_q)	TFP (ρ_q)
10%	440	310	330
25%	490	440	440
50%	550	550	550
75%	630	710	790
90%	930	950	1400

Note: Computed using the distribution of residuals, median labor and capital.

Table 4: Dispersion of Predicted Output due to TFP Dispersion, True Productivity and Measurement Error

true productivity. These low correlations of TFP are not unique to the ready-mix concrete industry; Abraham and White (2006) have documented autocorrelations across manufacturing of around 80%, where plant size (and hence serial correlation) is much larger than in the ready-mix concrete sector. Second, the exit propensity score shows that, conditional on survival, ready-mix concrete plants that are more likely to exit have higher productivity growth. Third, the 5-year autocorrelation of TFP is much lower than the 5-year autocorrelation of productivity, with coefficients of 0.13 versus 0.38, respectively. This indicates that true productivity is more persistent than TFP. If we ignore the problem of measurement error, we would obtain far lower estimates of the persistence of productivity, and this would bias estimates of the forward looking decision to exit the market conditional on the expected net present value of productivity.[14]

Dispersion of plant productivity may be due to different vintages of ready-mix plants. Older plants, though built with less efficient technology, compete with newer plants. Competing vintages can result in the wide dispersion of productivity observed in the data. However, the ready-mix concrete industry is unusual, in that it has not experienced substantial technological change during the past 50 years. For instance, the machines and trucks used to produce ready-mix concrete 50 years ago are remarkably similar to those in use today. Moreover, as Table A5 in Appendix C shows, the volume of concrete produced per worker hour increased by less than 10% between 1967 and 1997. Because aggregate productivity has increased very slowly, the assumption that the ready-mix concrete industry is near its steady-state is plausible. Thus productivity dispersion is an outcome of an industry's equilibrium rather than a product of a transitory change.

6 Evidence for Plant Selection

Before presenting the structural model of entry/exit and investment, I provide evidence of the mechanisms of plant selection. These results inform which features of the data can identify parameters of the structural model. The process for plant selection can be seen to operate via three channels: exit, growth, and competition. In this section, I provide some

[14] In Appendix **??**, I discuss decomposing the components of productivity volatility. Even if I use a much less volatile version of capital, based on constructing capital stock using investment flows and depreciation, I still obtain comparable autocorrelations of productivity.

Dependent Variable:	TFP ρ		Measurement Error ϵ		True Productivity ω			
	1 Year Lag	5 Year Lag	1 Year Lag	5 Year Lag	1 Year Lag	1 Year Lag	1 Year Lag	5 Year Lag
	I	II	III	IV	V	VI	VII	VIII
TFP ρ	0.547	0.129						
	(0.012)	(0.013)						
Measurement Error ϵ			0.459	0.097				
			(0.016)	(0.014)				
True Productivity ω					0.581	0.554	0.612	0.383
					(0.014)	(0.018)	(0.065)	(0.015)
Exit Propensity Score $\hat{\chi}$						0.663	-0.080	
						(0.291)	(0.832)	
ω^2							-0.020	
							(0.020)	
$\hat{\chi}^2$							1.560	
							(7.384)	
$\hat{\chi} \times \omega$							0.240	
							(0.656)	
Constant	0.046	-0.026	0.059	-0.019	0.463	0.465	0.442	0.598
	(0.011)	(0.007)	(0.011)	(0.008)	(0.020)	(0.020)	(0.050)	(0.014)
Observations	4090	6369	2816	4824	2843	2843	2843	4838
R^2 Adjusted	0.330	0.016	0.230	0.010	0.390	0.390	0.390	0.110

Table 5: Autocorrelation of Productivity, TFP and Measurement Error

descriptive evidence.

For the remainder of the paper, I confine my attention to plants located in counties that never have had more than 6 plants. The reason for this restriction is purely technical—to compute counterfactuals simulations I need to store the state of the market in computer memory. The model I use has 5 productivity and 4 capital states (including staying out of the market), along with ten states for market demand. This yields 2.5 million possible states for each market, which pushes the limit of computer memory.[15]

6.1 Inefficiency Encourages Exit

The first mechanism of plant selection is the exit of inefficient producers. Plant exit provides the cleanest evidence of selection, since an incumbent's productivity can be measured in the year before exiting. A plant's exit is determined both by productivity within the plant and competition in the market, but plant-level factors typically provide greater explanatory power.

Table 6 reports the marginal effects on plant size, productivity, competition and market demand from probit regressions of the decision to exit. Column I shows only the effect of plant size, the number of competitors and demand, and, thus, I can use the entire sample from the LBD. Columns II and III use data from the CMF, which are not imputed, and, thus, have a far fewer observations. Column III also includes "market category effects," denoted μ, which are a rough proxy for market fixed effects. Indeed, Table A7 in Appendix C shows that the coefficients from the regression with these market category effects are similar to the market fixed-effect conditional logit. These market category effects will be important since they substantially raise the estimated effect of competition. This is because the number of plants in a market is positively correlated with unobserved demand shifters (as in Collard-Wexler (2010) for instance). Thus these market category proxies will show up in many of the entry, exit and investment regressions that follow.

The effect of plant size is substantial. The mean exit rate for all plants in the sample is 6%, but having a plant with more than 7 employees lowers the exit rate by between 4.4% and 5.3%. Higher construction employment in the county in which a plant is located is associated with a far lower probability of exit in Columns I and II. However, the presence of at least

[15] Specifically I need to store 2.5 million states, and for each state I need to store about 100 numbers to characterize the action specific payoff as a function of the parameter vector.

1 other plant in the county raises exit rates by between 2% and 3.4%, indicating a strong competitive effect. The presence of additional competitors above 1 has a much smaller effect on exit rates.

In Columns II and III of Table 6, a plant with mean productivity in the bottom quintile has an exit rate of 7%, while a plant in the top quintile has an exit rate of 3.3%. This effect of productivity is fairly large—comparable to the effect of going from a monopoly to a duopoly market, or the effect of going from a small to a medium-sized plant.

At first glance, the relationship between productivity and exit is somewhat disappointing. As a benchmark, consider a model in which plants compete in a perfectly competitive market. A plant below a certain threshold of productivity certainly will exit. However, the weakness of the relationship between exit and productivity in our data explains the wide dispersion of productivity between plants: inefficient producers are slowly selected out. Another reason for the modest effect of productivity in these data is the role of reallocation, which means that low-productivity plants are more likely to shrink. Table 7 illustrates this effect: plants which have productivity below the median are 2% less likely to end the next period as large plants (more than 15 employees) rather than as small plants (less than 15 employees). Moreover, over two periods, smaller plants have exit rates that are twice as high as large plants. Thus, the increased risk of exit in two periods is attributed to small plant size, despite the fact that it is the plant's low efficiency that causes it to shrink in the first place.

6.2 Competitive Markets Are More Productive

Competition lowers profitability for all firms in a market. In a market with many competitors, less productive producers should be more likely to exit. This link between productivity and market size is thoroughly investigated by Syverson (2004), and I confirm his results in Table A8, which shows that there is a higher fraction of productive plants in large markets.

7 Dynamic Choice: Empirical Model

In this section, I describe the econometrics used to estimate the dynamic entry, exit and investment model. Estimation follows Collard-Wexler (2008) quite closely, and, thus, I will give only a brief overview of the details of the estimation algorithm. The main adaptation

Dependent Variable: Jarmin-Miranda Exit	I	II	III
Size			
Medium††	-0.044	-0.053	-0.052
	(0.002)	(0.008)	(0.008)
Large† † †	-0.051	-0.054	-0.053
	(0.002)	(0.008)	(0.008)
Productivity			
2nd Quintile		0.002	0.004
		(0.012)	(0.011)
3rd Quintile		-0.022	-0.020
		(0.011)	(0.011)
4th Quintile		-0.033	-0.032
		(0.010)	(0.010)
5th Quintile		-0.037	-0.037
		(0.010)	(0.010)
Log County Construction	-0.078	-0.006	0.000
Employment	(0.015)	(0.003)	(0.003)
1st Competitor	0.032	0.020	0.034
	(0.002)	(0.007)	(0.007)
Log of more than one competitor	0.009	0.005	0.018
	(0.001)	(0.003)	(0.004)
Average Plants in County Rounded			
2			-0.060
			(0.015)
3			-0.082
			(0.017)
4			-0.094
			(0.018)
Observations	64482	4627	4627
χ^2	1089	109	133
Log Likelihood ($\mathcal{L}$)	-12656	-826	-807

†: Small is 0 to 6 employees (omitted category), ††: Medium is 7 to 17 employees, † † †: Big is more than 17 employees. Standard Errors clustered by plant.

Table 6: Marginal Effect from a Probit: More Productive Plants Are Less Likely to Exit

From		To		
		Out	Small	Big
Out		99.1%◇	0.9%	0.0%
Small+	Low Productivity*	8.5%	86.2%	5.3%
	High Productivity**	3.8%	89.9%	6.3%
Big++	Low Productivity	2.3%	15.2%	82.4%
	High Productivity	1.8%	13.2%	84.9%

+ Small: Plant with fewer than 15 employees
++ Big: Plant with at least 15 employees
*Low Productivity: Productivity below the median for the year
**High Productivity: Productivity above the median for the year
◇Number of Entrants: 6 minus the number of active firms in the county

Table 7: Low-Productivity Plants Are Less Likely to Grow than High-Productivity Plants

in this paper is the addition of a firm-specific state which characterizes it's productivity.

First, I need to discretize the state space, since the solution techniques for solving dynamic oligopoly games with more than 2 firms all use discrete state spaces.[16] Thus, I place productivity (ω_i) into 5 bins corresponding to the quintiles of the productivity distribution. I use employment as my capital state, k_i, and place employment into 4 bins. I use employment, rather than assets, as a capital state variable primarily because employment data are available for all plants in the data from the LBD, while total assets are frequently missing in the CMF and ASM data. However, given the coarseness of the discretization of k_i, classifying

[16]There are essentially three reasons why it is difficult to solve dynamic oligopoly games with continuous state spaces. The first is that in this game with six players, the state will be twelve-dimensional. To my knowledge, there are few techniques that work well for solving dynamic decision problems in continuous state spaces with more than three or four dimensions. The second problem is that, if we try to use aggregate states, such as total demand or the number of firms, we can find value functions that are very nonlinear and, therefore, are difficult to approximate using common basis functions. For instance, when Abbring and Campbell (2010) graph the value function against demand, they find a saw-shaped pattern. As demand increases, the firm's value goes up until demand reaches a high enough level that an extra firm will want to enter the market, pushing the value function down to zero. The third problem is that common approaches for solving continuous-state space problems via basis function approximations—for instance, least-squares policy iteration—suffer from serious convergence issues: Approximating the value function introduces a source of error in the bellman iteration, making it difficult to find a solution. The work of Farias, Saure, and Weintraub (2008) makes progress toward solving dynamic oligopoly games with large state spaces.

firms by capital stock or employment does not a make a large difference. The demand state, M, is taken as the number of construction employees in the county, which is placed into 10 bins.[17]

I will denote the observed part of the state as x_i—i.e., the state vector does not include shocks ϵ_i or ϵ_i^E and is just capital and productivity, $x_i = \{k_i, \omega_i\}$. Moreover, the observed state vector for the market, denoted x, is the collection of each firm's observed state, x_i, and market demand, M.

A firm's action, a_i, is the choice of its capital in the next period: $a_i = k_i'$, which includes a capital stock, $k_i = \emptyset$, that indicates the decision to exit the market. This action, a_i, replaces i, χ^E and χ in this section of the paper, as the exit decision, χ, is just the decision to choose $k_i = \emptyset$ in the next period, and the entry decision is a firm's choice of positive capital in the next period, if its $k_i = \emptyset$. Likewise, a firm's investment or disinvestment choice is just the decision to choose a larger or a smaller k_i in the next period.

Finally, the shocks ϵ_i and ϵ_i^E are implemented as i.i.d. logit shocks to the payoffs of each action, a_i. Having a full support shock for each action insures that we will see each action, a_i, played with positive probability and helps ensure to the existence of a Nash Equilibrium to the game. Notice that I implicitly have assumed that each investment level also will be played with positive probability.

7.1 Period Profits

I parametrize the firm's reward function as:

$$\begin{aligned} r(a_{it+1}, x_{t+1}|\theta) \quad &= \textstyle\sum_{j\in A} \theta_{1j} 1(a_{it+1} = j) \text{ (Fixed Cost)} \\ &+ \textstyle\sum_{j\in A} \theta_{2j} 1(a_{it+1} = j)\omega_{it+1} \text{ (Productivity Effect)} \\ &+ \textstyle\sum_{j\in A} \theta_{3j} 1(a_{it+1} = j) M_{t+1} \text{ (Demand)} \\ &+ \textstyle\sum_{j\in A} \theta_{4j} 1(a_{it+1} = j) \ln(\sum_{-i} 1(k_{-it+1} > 0)) \text{ (Competition)} \end{aligned} \tag{19}$$

which is a simple, first-order approximation of the firm's profits, given the state vector, x_{t+1}, and the action chosen, a_{it+1}, which is k_i in the next period as well. Note that, in this model, higher productivity does not have a direct effect on competitors. Instead, productivity has

[17] The following employment and productivity bins are used:

ω_i Bin	1	2	3	4	5	k_i Bin	1	2	3
Mean Productivity	0.67	0.79	0.89	0.98	1.57	Mean Employment	3.3	11.3	33.6

an indirect effect on the profits of competitors because more productive firms are more likely to stay in the industry.

The are also adjustment costs for entering the market and for changing the size of a plant. These transition costs, denoted $\tau(a_i, k_i)$, will be estimated flexibly—i.e., I will estimate a parameter for the cost of moving from each capital state k_i bin to another capital state, $\tilde{k}_i$.

7.2 CCP Estimation

To estimate the parameters of the profit function using data on entry and investment choices, I need to compute the expected net present value of profits from taking an action, a_i, in state x, which I denote as $W(a_i, x)$. This choice-specific value function, $W(a_i, x)$, is given by:

$$W(a_i, x) = \mathbf{E}_{a_{0i}, x_0} \sum_{t=0}^{\infty} \beta^t \left[r(a_{it}, a_{-it}, x_t) - \tau(a_{it+1}, k_{it}) + \epsilon_{a_{it}} \right] \tag{20}$$

where $\epsilon_{a_{it}}$ is the value of the choice-specific unobservable, conditional on having taken action, a_{it}.

I can approximate equation (20) using forward simulation given by:

$$W(a, x) \approx \frac{1}{K} \sum_{k=1}^{K} \sum_{t=0}^{T} \beta^t \left(r(a_{it}^k, a_{-it}^k, x_t^k) - \tau(a_{it+1}^k, k_{it}^k) + E(\varepsilon | P[\cdot | x_t^k]) \right) + \beta^T \xi \tag{21}$$

where it is assumed that the private information unobservable, ε_i^t, is distributed as an i.i.d. logit. Thus the expected value of ε_i^t, given that firms have conditional choice probabilities—i.e., probability $\mathbf{P}[a_i|x]$ of taking action a_i in state x, can be computed as:

$$E[\varepsilon | P(\cdot | x_t^k)] = \gamma - \sum_{j \in A} \ln(P[j | x_t^k]) P[j | x_t^k] \tag{22}$$

The term ξ is the error from truncating the simulation value at T, where T will be chosen to be large enough so that this β^T is small.[18]

However, to compute this forward simulation, I need to be able to simulate the evolution

[18] I use $K = 10\,000$ and $T = 1000$ to compute $W(a_i, x)$ via forward simulation.

of the state, x, as well as the actions taken by firms in the future. The state evolves according to the transition density:

$$x' \sim D[M'|M] \prod_{i=1}^{N} (\mathbf{P}[a_i'|x] P^{\omega}[\omega_i'|k_i]) \tag{23}$$

Actions then evolve according to $a_i \sim \mathbf{P}[a_i|x]$, where $\mathbf{P}$ are the conditional choice probabilities—i.e., the probability that a firm takes action a_i in state x, D is the process for demand, and P^{ω} is the transition density for productivity. The key insight from the modern literature on estimating dynamic oligopoly games is that if we can estimate these three objects ($\hat{D}$, $\hat{P}^{\omega}$, $\hat{\mathbf{P}}$) directly from the data and, in particular, if we can estimate $\mathbf{P}$, then it is not necessary to solve for an equilibrium of the game, since the equilibrium strategies are directly observable.[19]

The transition density for demand $\hat{D}$ is estimated using a bin estimator and is described in more detail in Collard-Wexler (2008).[20] Likewise, the productivity transition process, $P^{\omega}[\omega_i'|k_i]$, is given by:

$$P^{\omega}[\omega_i'|k_i] = \begin{cases} \hat{P}^{\omega}[\omega_i'|\omega_i] \text{ if } k_i \neq \emptyset \\ \hat{P}^{\omega E}[\omega_i'] \text{ if } k_i = \emptyset \end{cases}$$

This process is estimated using a non-parametric bin estimator without reference to the structural model and is described in Table 5 on page 20. Finally, I estimate the conditional choice probabilities, $\mathbf{P}$, using a multinomial logit on k_{it+1} (i.e., size in the next year), as shown in Table A6 in Appendix C.

A final rewriting of the W function is now in order to aid with the estimation of the model. The rewards and transition costs are linear in parameters θ, so the profit function can be rewritten as $r(a_i, x|\theta) - \tau(a_i, x_i|\theta) = \theta \cdot \vec{\rho}(a_i, x)$ where $\vec{\rho}$ is a function that returns a vector. This implies that the W function will be separable in dynamic parameters, as in

[19] Notice that, for this CCP approach to work, it is necessary to be able to estimate $\mathbf{P}$ in a first stage.

[20] Demand is placed into 10 discrete bins, $B_i = [b_i, b_{i+1})$, where the b_i's are chosen so that each bin contains the same number of demand observations. Making the model more realistic by increasing the number of bins above 10 has little effect on estimated coefficients, but lengthens computation time significantly. The level of demand within each bin is set to the mean demand for observations in this bin, i.e. $\text{Mean}b(i) = \frac{\sum_{l=1}^{L} M_l 1(M_l \in B_i)}{\sum_{l=1}^{L} 1(M_l \in B_i)}$, where L indexes observations in the data and the D matrix is estimated using a bin estimator, $\hat{D}[i|j] = \frac{\sum_{(l,t)} 1(M_l^{t+1} \in B_i, M_l^t \in B_j)}{\sum_{(l,t)} 1(M_l^t \in B_j)}$.

Bajari, Benkard, and Levin (2007), since

$$\begin{aligned} W(a_i, x|\theta) &= E\sum_{t=1}^{\infty}\beta^t\left(r(a_{it}, x_t|\theta) - \tau(a_{it}, k_{it}|\theta)\right) \\ &= \theta \cdot E\sum_{t=1}^{\infty}\beta^t\vec{\rho}(a_{it}, x_t) \equiv \theta \cdot \Gamma(a_i, x) \end{aligned} \tag{24}$$

so $W(a_i, x) = \theta \cdot \vec{\Gamma}$.[21]

Given the logit shocks to the value of taking each action, a_i, I can compute the probability of taking an action as:

$$\begin{aligned} \Psi(a_i|x,\theta) &= \frac{\exp\left(W(a_i,x)\right)}{\sum_{j\in A_i}\exp\left(W(j,x)\right)} \\ &= \frac{\exp\left(\theta\cdot\vec{\Gamma}(a_i,x)\right)}{\sum_{j\in A_i}\exp\left(\theta\cdot\vec{\Gamma}^P(j,x)\right)} \end{aligned} \tag{25}$$

7.3 Second Stage: Indirect-Inference Estimation

The model is estimated via Indirect Inference. A maximum likelihood estimation strategy, using the choice probabilities, Ψ^P, is quite practical since the log-likelihood function, $\mathcal{L}$, will be globally concave:[22]

$$\mathcal{L}(\theta) = \sum_{t,i}\ln\left(\Psi(a_{it}|x_{it},\theta,\Gamma)\right) \tag{26}$$

However, since there is simulation error embedded into the Γ function, a maximum-likelihood estimator will be biased (see, e.g., McFadden (1989); and Pakes, Berry, and Ostrovsky (2007)). Γ has simulation error both because the choice probabilities, $\mathbf{P}$, have sampling error and because the forward simulation procedure is an approximation of the true value function.

[21] Note that the choice probabilities for different values of θ in equation 25 can be evaluated using multiplication.

[22] To see this, note that the assumption of linearity in dynamic parameters gives a utility function of the form $u_a = \theta\cdot\Gamma(a,x)$, which is linear. Along with the assumption that the ϵ_a's are logit, this implies a globally concave likelihood function. I use Maximum Likelihood estimates as starting values for the Indirect Inference procedure.

The Indirect Inference (II) procedure, like many GMM estimators, will be consistent in the presence of simulation error in Γ. Essentially, the reason for this is that the II procedure matches conditional means in the data and in the simulated data, and the simulation error in Γ washes out in expectation of this. As an auxiliary model, I choose a multinomial linear probability model because it is simple to estimate and it is a close analogue to the dynamic multinomial logit model used here. Note that the auxiliary model does not need to have an interpretation of any sort. Its sole responsibility is to provide a rich description of the patterns of a data set and to be simple to estimate.

It is necessary to define the choice to be a small plant as $y^s_{it} = 1(a_{it} = \text{small})$. I run a linear probability regression of y^s_{it} on covariates z_{it} and get the OLS coefficients $\hat{\beta}^s$. The covariates of the auxiliary model, z_{it}, are indicators for the firm's current state (including the firm's productivity and size), the number of competitors in a market, the log of construction employees in the county, and dummies for the market category μ.

I also run the same linear probability regression using covariates z_{it} and the dependent variable $\tilde{y}^s_{it}(\theta) = \Psi(\text{small}|x_{it}, \Gamma, \theta)$ to obtain the predicted probability given by the dynamic model for the decision to have a small plant. This regression yields OLS coefficient $\tilde{\beta}^s(\theta)$. The idea behind Indirect Inference is to minimize the discrepancy between these two regression coefficients.[23]

I repeat this procedure for the decision to have a medium-sized plant, $y^m_{it} = 1(a_{it} = \text{medium})$, and for the decision to have a large plant, $y^l_{it} = 1(a_{it} = \text{large})$.

The criterion function evaluates the distance between the regression coefficient in the data and in the model's predictions:

$$\begin{aligned} \mathcal{Q}(\theta) = & \left(\hat{\beta}^s - \tilde{\beta}^s(\theta)\right)' \mathbf{W}^s \left(\hat{\beta}^s - \tilde{\beta}^s(\theta)\right) \\ & + \left(\hat{\beta}^m - \tilde{\beta}^m(\theta)\right)' \mathbf{W}^m \left(\hat{\beta}^m - \tilde{\beta}^m(\theta)\right) \\ & + \left(\hat{\beta}^l - \tilde{\beta}^l(\theta)\right)' \mathbf{W}^l \left(\hat{\beta}^l - \tilde{\beta}^l(\theta)\right) \end{aligned} \tag{27}$$

where $\mathbf{W}^s$ is a positive definite weighting matrix, where I use $\mathbf{W}^s = V[\hat{\beta}^s]^{-1}$, the inverse of

[23] In Collard-Wexler (2008), I show that using the choice probabilities Ψ as predicted actions gives the same θ's as drawing action $a_{it} \sim \Psi(\cdot|x_{it}, \Gamma, \theta)$ from the predicted choice probabilities using an infinite number of simulation draws. However, using the probabilities, rather than a simulator, yields a smooth criterion as a function of θ.

the covariance matrix from the OLS regression, and likewise for $\mathbf{W}^m$ and $\mathbf{W}^l$.

8 Dynamic Results

I estimate the dynamic model of entry and exit with exogenous productivity using the CCP style estimator presented in the previous section. I use a discount rate of 0.95.

There is one final problem that needs to be settled in order to estimate the model. In the estimation of productivity and of entry, exit and size choices, I have deleted plants for which the data was imputed. However, for the estimation of the dynamic oligopoly model, eliminating all imputed data is, to say the least, problematic. If there are 5 plants in a market and a single firm has imputed data, then the entire market would need to be dropped from the sample. Since even in census years, 30% of plants have missing data, almost all large markets would be eliminated. To get around this problem, I will use multiple imputation software to impute the data for missing plants. Specifically, if plant i has missing productivity ω_i, this plant will be dropped. However, if plant i has a competitor plant that is missing productivity data (i.e., if ω_{-i} is missing), then I use multiple imputation software to fill in the missing value. Section A.1 in Appendix A describes the exact procedure I use to impute the data.[24]

Table 8 presents estimates of the dynamic entry-exit and investment model. As reported in Collard-Wexler (2008), interview data indicate that the sunk costs of starting a ready-mix concrete plant are on the order of 2 million dollars. Thus, I can convert the coefficients expressed in variance units into dollars by normalizing the cost of entering as a large plant at 2 million dollars. This, in turn, indicates that the unobserved state, ϵ, has a variance of about $200,000 per year. Columns I and II present estimates using maximum likelihood, while columns III and IV show results estimated via indirect inference. The results in column I, estimated using maximum likelihood, and the results in column III, estimated using indirect inference, are very similar with only a moderate economic difference between them. This suggests that the inconsistency of maximum likelihood due to simulation error does not seriously bias the estimates.[25] I find the fixed costs of either small (less than seven

[24]Unfortunately, using multiple imputation software to fill in ω_i does not accurately replicate the exit probits shown in Table 6.

[25]In a previous paper, Collard-Wexler (2008), I find substantial differences between indirect inference and maximum likelihood estimates. I do not know why the Maximum Likelihood and Indirect Inference estimates

employees) medium (7 to 17 employees) or large (more than 17 employees) plants are fairly similar, at around \$400,000 per year. Market size, as measured by the log of construction employment in the county, has a positive effect on firm profits, and, as is documented in Collard-Wexler (2008), larger firms benefit disproportionately from higher demand—i.e., in bigger markets, we find not only more firms, but also larger firms. Competition, as measured by the log of the number of ready-mix concrete plants in the county, has a negative effect on profits. Column II shows estimates that include market category effects. On average, markets that have more firms have higher profitability. Notably, when market category effects are included, the estimated effect of competition rises considerably, which is to be expected if an unobservably more profitable market attracts more entrants, leading to a positive bias on the effect of competition.

Adjustment costs are also important. The cost of entry for a small plant is estimated to be \$1.4 million, which is smaller than the cost of entry for either a medium or large plant, calibrated to \$2 million. As well, there are significant costs involved in raising the size of a plant from small to medium and from small to large—estimated at approximately \$600,000 and \$950,000, respectively. Finally, I also find that there are significant, if smaller, costs involved in reducing plant size. On average, it costs around \$250,000 to shrink a plant's size.

Productivity increases enhance a plant's profits—be it a small, medium, or large plant. The mean productivity of a plant in the lowest quintile is 0.67, while the mean productivity of a plant in the top quintile is 1.57. Thus, a medium-size plant in the top quintile of productivity has profits that are \$220,000 higher than the profits of a plant in the bottom quintile of productivity. Using Indirect Inference estimates, I find that the effect of productivity on a medium-size plant is \$283,000. I find similar, if somewhat smaller, effects of productivity on profits for both small and large plants.

8.1 Net Present Values

To understand why productivity has such a large impact on profits, recall that the reduced-form multinomial logits on plant size described in Table A6, as well as the probit regressions on exit in Table 6 show a fairly important positive effect of productivity. However, this productivity effect is substantially smaller than the effect of competition and is roughly the

are similar in this application, but different in the previous paper.

		Maximum Likelihood Criterion				Indirect Inference Criterion			
		I		II		III		IV	
Fixed Costs	Small†	-338	(51◇)	-322	(49◇)	-458	(66)	-456	(65)
	Medium††	-476	(60)	-468	(59)	-530	(80)	-534	(80)
	Large†††	-398	(69)	-382	(69)	-388	(87)	-384	(87)
Market Type 2				40	(4)			-67	(78)
Market Type 3				75	(7)			-138	(134)
Market Type 4				92	(9)			-168	(209)
Log Construction	Small	4	(4)	-9	(4)	-11	(7)	-8	(6)
Employment (M)	Medium	48	(9)	39	(9)	46	(16)	44	(15)
	Large	75	(11)	66	(12)	62	(18)	58	(18)
Productivity (ω)	Small	180	(47)	200	(46)	326	(64)	317	(63)
	Medium	225	(54)	246	(54)	276	(75)	283	(77)
	Large	100	(62)	113	(63)	125	(78)	129	(79)
Log Competitors	Small	10	(4)	-54	(7)	2	(10)	7	(8)
($\log(N)$)	Medium	-57	(9)	-123	(11)	-41	(17)	-53	(13)
	Large	-49	(13)	-113	(14)	-10	(19)	-18	(17)
Adjustment Costs	Out → Small	-1392	(40)	-1338	(40)	-1334	(59)	-1341	(59)
$\tau(a_i, x)$	Out → Medium	-2000	(67)	-2000	(68)	- 2000	(240)	-2000	(254)
	Out → Large	-2023	(78)	-2040	(80)	-2345	(999)	-2342	(982)
	Small → Medium	-594	(46)	-635	(47)	-677	(58)	-673	(57)
	Small → Large	-919	(59)	-981	(61)	-1064	(80)	-1063	(78)
	Medium → Small	-57	(46)	-36	(46)	19	(54)	15	(53)
	Medium → Large	-312	(58)	-334	(59)	-290	(80)	-293	(79)
	Large → Small	-281	(48)	-252	(50)	-342	(70)	-343	(69)
	Large → Medium	-416	(59)	-415	(60)	-448	(80)	-447	(79)
Variance of ϵ		195		201		199		199	
Observations		14278		14278		14278		14278	
χ^2		73968		76543					
Log-Likelihood $\mathcal{L}$		-4137		-4077					

Note: All estimates are in thousands of dollars. Calibration to dollars, by assuming the entry cost of a medium-size plant, are 2 million dollars. †: Small is 0 to 6 employees. ††: Medium is 7 to 17 employees. †††: Big is more than 17 employees. ◇: Standard Errors are computed by bootstrapping the first-stage estimates of $\hat{P}$, $\hat{P}^{\rho}$ and ω, and then reestimating the model. I use 100 bootstraps to compute the Standard Error.

Table 8: Dynamic Entry Model with Exogenous Productivity

same magnitude as the effect of log construction employment. In the dynamic discrete choice model, I do not estimate the effect of current productivity on the net present value of profits in the future, but instead measure the effect of current productivity on current profits.

Next, I unpack the mapping between the current state and the expected net present value of profits in the future. To do this, Table 9 runs regressions of the net present value of competition on the current competion, the NPV of productivity on current productivity, and the NPV of market size on current market size. To facilitate the interpretation of Table 9, I present summary statistics of these variables in Table A9 in Appendix C. Column I in Table 9 regresses the net present value of years of activity on plant size. On average, a plant can expect almost 11 years of activity in net present value (equivalent to 16 years), which is about right given that I have chosen a discount rate of 5% per year and that there is approximately a 5% chance that a plant will exit in any given year. Medium and large plants expect to live about one more year in net present value (equivalent to living 19 years), which makes sense, given that larger plants have much lower exit rates.

Column III in Table 9 regresses the net present value of productivity on current productivity. If productivity were perfectly persistent, then the coefficients in this regression should equal 10, since this is calculated by multiplying today's productivity by the expected net present value of the years that the firm will remain in activity. This calculation yields a coefficient of 1.8, which indicates that the current year's productivity is a fairly poor predictor of the net present value of productivity in the future. The fact that current productivity has a big impact on both the decision to exit the market and the determination of plant size is even more striking than it may seem, since firms are responding to fairly short-lived changes in productivity.

In contrast, column IV in Table 9 regresses the net present value of log construction employment on this year's log construction employment and finds a coefficient of 6.3. Thus, current demand is a much better predictor of future demand than it is of future productivity. This explains why, in the reduced form, I find a large effect of demand on either size or exit decisions—an effect that is mitigated by the fact that demand is very persistent, as compared with productivity.

Dependent Variable:	NPV† Activity Years I	NPV Log Competitors II	NPV Productivity III	NPV Log Construction Employment IV
Medium	0.82			
	(0.06)			
Large	1.09			
	(0.07)			
Log Competitors		1.70		
		(0.01)		
Productivity (ω)			1.77	
			(0.10)	
Log Construction Employment				6.35
				(0.06)
Constant	10.77	1.64	10.5	16.25
	(0.04)	(0.02)	(0.10)	(0.60)

Note: There are 3896 observations for all regressions. †NPV's refer to the net present value of the variable (with a 5% discount rate) using the forward simulation procedure.

Table 9: Demand and Competition Show Far More Predictability than Productivity

9 Counterfactuals

In the previous section, I explained how a large effect of current productivity on profits translates into a small effect on a firm's exit decisions. Two forces were conjectured for this effect. First, sunk costs have an insulating effect on firms—i.e., a low-productivity incumbent may remain in the market because it is protected from potentially higher-productivity entrants by sunk entry costs. Second, the large volatility of productivity makes it optimal for a firm's response to current producivity to be fairly mild.

To separate the quantitative importance of these two forces, I compute counterfactuals where I change either a) the sunk costs of entry, or b) the time-series volatility of productivity. I then will look at the equilibrium strategies in the equilibrium of the resulting game. I do this in order to determine to what extent selection based on productivity is strengthened by lower sunk costs of entry or, alternatively, by greater persistence of productivity.

It is important to look at these effects on equilibrium policies because changing the parameters in the game may have effects on selection that are difficult to intuit. For instance, lower entry costs induce entry, and, if the market is competitive, productivity may influence their exit decisions. Thus, it is difficult to evaluate which features of the industry are responsible for the very small effect of productivity on exit without recomputing equilibria to the game.

The three equilibrium computations I perform are:

a) Baseline (No Changes)

b) Increasing the Persistence of Productivity.

   A key finding in this paper is that productivity is very volatile from year to year. To see how this volatility slows the selection process, I compute the entry and exit policies that would exist if productivity were fully persistent.[26]

c) Reducing Entry Barriers

   To see how entry barriers slow the replacement of less productive incumbents by potentially more productive entrants, I alter the cost of building a new ready-mix concrete plant. Specifically, I reduce all transition costs, $\tau(a_i'|k_i)$, by 20% and recompute the model.

I compute an equilibrium policy functions for the game, denoted $\mathbf{\Psi} = \{\Psi[a_i|x]\}_{a_i \in A_i, x \in X}$, by adapting the Stochastic Algorithm given in Pakes and McGuire (2001) to the discrete action setup used in this paper. Then, I use the parameter estimates from the indirect inference procedure with market category effects, denoted $\hat{\theta}^{II,M}$, that were presented in column IV of Table 8. The main variations on the Stochastic Algorithm are the discrete action setup, where firms choose their size in the next period and decide to enter or exit the market, and the randomness of the change in productivity. The details of this algorithm are similar to those used in Collard-Wexler (2008). I discuss the exact algorithm in Appendix B on page 44.[27]

[26] This is the assumption used in many models of industry dynamics such as the canonical Jovanovic (1982) model.

[27] There are more than 2.5 million states in this application, along with 7 firms, 16 possible states per firm and 10 demand states. I reduce the size of the state from 1.6 billon to 2.5 million by using the assumption of exchangeability described by Gowrisankaran (1999). In addition, there are 10 million policies to evaluate. I use a Stochastic Algorithm coded in C to compute the equilibrium to this game.

To illustrate the characteristics of the computed equilibrium policies for all three counterfactuals, I look at the role of productivity in a probit regression on the decision to exit, which is an analogue to the probit on exit in Table 6. Specifically, I first simulate actions a_{it}^k from the computed equilibrium policies; i.e. $a_{it}^k \sim \Psi[\cdot|x_{it}]$. I then run a probit regression on the decision to exit using the simulated dataset, $\mathcal{D}^k = \{a_{it}^k, x_{it}\}_{i,t}$. Since I have a large amount of data, I will only use one simulation per observation, which implies $k = 1$.

Table 10 shows probits on the decision to exit estimated on the data simulated from the 3 different counterfactuals: base (column A), no productivity volatility (column B), and lowered entry barriers (column C). [28] The constant term in column C is far larger than in columns A and B. This is to be expected since lowered entry barriers will lead to higher turnover rates in the industry, for both entry and exit. Likewise the effect of size, as measured by the number of employees, is also lower in column C versus columns A and B because I have lowered the transition costs across all states in this counterfactual.

The most important results concern the effect of productivity on the decision to exit. Notice that productivity has an effect of -0.087 in column A versus -0.672 in column B. In other words, shutting off the volatility of productivity would increase the effect productivity on exit by a factor of 8. Thus, the intuition that there should be a very strong selection effect is correct only if productivity itself does not change much over time. However, given the large volatility of productivity, these selection effects are not strongly represented in the data. Other coefficients are fairly similar among the 3 counterfactuals, with the exception of log construction employment, which is smaller in column B than in either column A or C.

In contrast, reducing entry barriers in column C leads to a coefficient on productivity that is actually higher than in the baseline model in column A. Surprisingly, higher productivity increases the probability of exit, but not significantly. Thus, lowering entry barriers would not increase the selection effect by itself. Indeed, given the estimated importance of the i.i.d. shocks ϵ_i in Table 8, lowering entry barriers makes firms more sensitive to these ϵ_i shocks, rather than increasing their sensitivity to productivity.

[28] While it is a bit odd to show Standard Errors on simulations, they indicate how difficult it would be to identify these effects in real data.

Dependent Var: Simulated Exit	A Base	B No Volatility of Productivity	C Lowered Entry Barriers
Productivity	-0.087	-0.672***	0.077
	(0.111)	(0.121)	(0.108)
7-17 employees	-0.464***	-0.671***	-0.479***
	(0.084)	(0.090)	(0.088)
More 17 Employees	-0.674***	-0.635***	-0.372***
	(0.110)	(0.097)	(0.095)
Log Construction	-0.235**	-0.035	-0.266***
Employment	(0.075)	(0.073)	(0.073)
Log Competitors	0.261***	0.319***	0.151***
	(0.047)	(0.045)	(0.046)
Market Effects			
Average 2 plants	-0.983***	-0.817***	-0.573***
	(0.094)	(0.092)	(0.089)
Average 3 plants	-1.323***	-1.176***	-0.773***
	(0.143)	(0.137)	(0.130)
Average 4 plants	-1.322***	-1.130***	-0.953***
	(0.183)	(0.166)	(0.177)
Constant	-0.391*	-0.351	-0.668***
	(0.179)	(0.181)	(0.177)
χ^2	271.78	263.43	143.10
Log-Likelihood	-805.0	-838.9	-802.5
Observations	4089	4089	4089

Table 10: Exit and Productivity in 3 Different Counterfactuals

10 Conclusion

In the ready-mix concrete industry, plants using the same bundles of inputs produce substantially different amounts of concrete. A plant in the 90^{th} percentile of productivity produces four times the value added as a plant in the in the 10^{th} percentile. However, the exit rate of plants in the bottom quintile of productivity is 6.8% on average. In comparison, plants in the top quintile of productivity had an average exit rate of 3.3%. Why do these enormous differences in productivity translate into such small differences in exit rates?

First, there is measurement error, which overstates the true dispersion of productivity.

After correcting for this using a control function technique developed by Ackerberg, Frazer, and Caves (2006), I find that the average plant in the 90^{th} percentile of productivity only produces twice as much value added as an average plant in the in the 10^{th} percentile.

Second, I build and estimate a dynamic oligopoly model of the industry. Plants with low productivity are more likely to exit and less prone to invest. In order to rationalize these observed entry, exit and investment decisions, plants in the bottom quintile of productivity must have per-period profits that are \$ 220,000 lower than plants in the top quintile. This corresponds to a difference of 1.5 in value added between plants in the bottom and top quintiles—which is quite large, though somewhat smaller than the 2 for 1 difference found in the data. To understand why this large difference in profits translates into small differences in exit rates, there are two critical factors to consider. First, in this industry, productivity has little persistence. Thus, current productivity does not provide much information on the net present value of productivity over the plant's expected lifetime. Second, the industry's high sunk costs slow the exit of unproductive producers.

I simulate the counterfactuals of either making productivity fully persistent or lowering adjustment costs (e.g., sunk costs of entry) by 20%, I find that lowered adjustment costs have little effect on the relationship between exit and productivity. In contrast, fully persistent productivity yields exit rates that are 7 to 8 times more responsive to productivity than those with the current volatility of productivity.

This paper concludes that large dispersion in productivity is quantitatively consistent with the standard model of oligopoly entry and exit once the correct volatility of productivity is accounted for. A question for future research is, thus, not what causes cross-sectional differences in productivity, but, rather, how we account for the large time-series volatility of productivity itself.

References

Abbring, J. H., and J. R. Campbell (2003): "A Firm's First Year," *Federal Reserve Bank of Chicago Working Paper*, 11.

——— (2010): "Last-in first-out oligopoly dynamics," *Econometrica*, 78(5), 1491 – 1527.

Abraham, A., and K. White (2006): "The Dynamics of Plant-level Productivity in US Manufacturing," *Center for Economic Studies Working Paper. CES-WP-06-20.*

Ackerberg, D., G. Frazer, and K. Caves (2006): "Structural Estimation of Production Function," *Working Paper UCLA*.

Aguiregabiria, V., and P. Mira (2007): "Sequential Estimation of Dynamic Games," *Econometrica*, 75(1), 1–54.

Bajari, P., L. Benkard, and J. Levin (2007): "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 75(5), 1331–1370.

Benkard, C., P. Bajari, and J. Levin (2007): "Estimating dynamic models of imperfect competition," *Econometrica*, 75, 1331–1370.

Collard-Wexler, A. (2006): "Demand Fluctuations and Plant Turnover in the Ready-Mix Concrete Industry," Working Paper, New York University.

——— (2008): "Demand Fluctuations in the Ready-Mix Concrete Industry," *Working Paper, New York University.*

——— (2010): "Mergers and Sunk Costs: An Application to the Ready-Mix Concrete Industry," Working Paper, NYU.

Das, S., M. J. Roberts, and J. R. Tybout (2007): "Market Entry Costs, Producer Heterogeneity, and Export Dynamics," *Econometrica*, 75(3), 837 – 873.

De Loecker, J. (2007): "Product differentiation, multi-product firms and estimating the impact of trade liberalization on productivity," *NBER working paper*.

Doraszelski, U., and M. Satterthwaite (2010): "Computable Markov-Perfect Industry Dynamics," *RAND Journal of Economics*, 41(2), 215–243.

Dunne, T., S. Klimek, and M. J. Roberts (2005): "Exit from regional manufacturing markets the role of entrant experience," *Journal of Industrial Organisation*, 23, 399–421.

Ericson, R., and A. Pakes (1995): "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *The Review of Economic Studies*, 62(1), 53–82.

Farias, V., D. Saure, and G. Weintraub (2008): "The Linear Programming Approach to Solving Large Scale Dynamic Stochastic Games," Working Paper, Columbia University.

Fershtman, C., and A. Pakes (2009): "Oligopolistic Dynamics with Asymmetric Information: A Framework for Empirical Work," Working Paper, Harvard University.

Foster, L., J. Haltiwanger, and C. Syverson (2008): "Reallocation, firm turnover, and efficiency: Selection on productivity or profitability?," *American Economic Review*, 98(1), 394–425.

Gowrisankaran, G. (1999): "Efficient Representation of State Spaces for Some Dynamic Models," *Journal of Economic Dynamics and Control*, 23(8), 1077–1098.

Hopenhayn, H. A. (1992): "Entry, Exit, and firm Dynamics in Long Run Equilibrium," *Econometrica*, 60(5), 1127–1150.

Jarmin, R. S., and J. Miranda (2002): "The Longitudinal Business Database," Working Paper, Center for Economic Studies, US Census Bureau.

Jovanovic, B. (1982): "Selection and the Evolution of Industry," *Econometrica*, 50(3), 649–670.

Levinsohn, J., and A. Petrin (2003): "Estimating Production Functions Using Inputs to Control for Unobservables," *Review of Economic Studies*, 70(2), 317–341.

Maskin, E., and J. Tirole (1988): "A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs," *Econometrica*, 56(3), 549–569.

McFadden, D. (1989): "A Method of Simulated Moments for Estimation of the Multinomial Probit Without Numerical Integration," *Econometrica*, 57(5), 995–1026.

Olley, G. S., and A. Pakes (1996): "The dynamics of productivity in the telecommunications equipment industry," *Econometrica*, 64(6), 1263–1297.

Pakes, A., S. Berry, and M. Ostrovsky (2007): "Simple estimators for the parameters of discrete dynamic games, with entry/exit examples," *RAND Journal of Economics*, 38(2), 373–399.

Pakes, A., and R. Ericson (1998): "Empirical Applications of Alternative Models of Firm and Industry Dynamics," *Journal of Economic Theory*, 79(1), 1–45.

Pakes, A., and P. McGuire (2001): "Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the 'Curse' of Dimensionality," *Econometrica*, 69(5), 1261–1281.

Pesendorfer, M., and P. Schmidt-Dengler (2008): "Asymptotic least squares estimators for dynamic games," *Review of Economic Studies*, 75(3), 901–928.

Rust, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55(5), 999–1033.

Syverson, C. (2004): "Market Structure and Productivity: A Concrete Example," *Journal of Political Economy*, 112(6), 1181–1222.

A Imputed Data

The U.S. Census Bureau imputes data in three ways:[29]

1. **Administrative Records**:

   Plants with fewer than 5 employees are deemed to be administrative records (henceforth AR)—i.e., they do not have to respond to Census questionnaires. The Census of Manufacturing (CMF) imputes data for AR plants, which constitute about 30% of ready-mix concrete plants. However, there is lower fraction of missing data in Census years, when the CMF is used, than in non-Census years, where Annual Survey of Manufacturers (ASM) only samples $\frac{1}{3}$ of plants.

   AR data is imputed based on the number of employees at the plant, with larger plants being sampled with a higher likelihood than smaller plants. Moreover, producers with fewer than 5 employees are rarely sampled by the CMF. Thus, it is difficult to know if very small producers are productive or not since we rarely see information on these producers. To estimate the dynamic model presented in this paper, I use imputed values of productivity for AR plants.

2. **Cold Deck Imputes**:

   If a plant does not respond to a particular question on the ASM or CMF, its response can be imputed by taking the response for the average plant and scaling that response to the number of employees at the non-responding plant. This imputation technique is known as cold deck imputation.

3. **Hot Deck Imputes**:

   Another way to impute data is to assign a plant the same level of capital, labor and output as another plant with similar characteristics, such as the same number of employees. This imputation technique is known as hot deck imputation.

Administrative Records are flagged in Census data sets, but hot and cold deck imputes are not. Therefore, I need to identify plants whose data contains hot and cold deck imputes. To eliminate hot deck imputes, I flag all plants in a given year that have identical capital, salaries and value added. For each plant that I classify as potentially a hot deck impute, I cannot identify the original plant whose data was used for imputation; i.e., I cannot distinguish the true plant from it's imputed counterpart. Likewise, to pinpoint cold deck imputes, I flag all plants with the same capital-labor ratio as the mode for that year.[30] Falsely classifying an plant as a cold deck impute is unlikely, since a plant with real data would have to have *exactly* the same capital-labor ratio as the mode for the year.

Table A1 shows the number of observations in CMF and ASM data that are Administrative Records, hot deck imputes and data collected in non-Census years by the ASM. About 40 % of the data used in the analysis for this paper is imputed.

[29] The discussion of stripping imputes from Census data draws heavily on Syverson (2004).

[30] An observation is classified as a cold deck impute if $|K_i^t/L_i^t - mode(K/L)^t| < 0.001$ where K_i^t and L_i^t denote plant i's capital and salaries in year t.

Impute Flag	Number of Observations
Administrative Records	7231
Hot Imputes	6277
Non-Census years	8217
Total Observations with capital, salary and shipments data	37559

Note: An observation can be both an Administrative Record, a hot impute and an non-Census year; thus, the total number of observations is smaller than the sum of the imputed and non-Census year data.

Table A1: A Large Fraction of Census of Manufacturing and Annual Survey of Manufacturers Data Are Imputed

Table A2 shows production function estimates using value added as a plant's output. For these estimates, I dropped Administrative Records, cold and hot deck imputes and non-Census years. Dropping imputed data from the sample has little effect on estimated coefficients but increases the variance of the productivity residual. However, if I include data from non-Census years, the coefficient on labor rises from 0.6 to over 0.8 and the capital coefficient falls from 0.3 to 0.04. The large decline in the capital coefficient, due a very low capital coefficient in non-Census year, is the reason for which I will mainly focus on non-Census years.

In addition, plant fixed-effects regressions are similar to estimates that OLS. Thus, selection does not seem to be a problem in the measurement of productivity—i.e., more productive plants are more likely to survive and increase their capital stock, since OLS and plant fixed-effect regressions are similar. This indicates that the persistent component of plant-level productivity is uncorrelated with capital stock.

A.1 Imputes for the Dynamic Oligopoly Model

The data used in the dynamic model is imputed using the following procedure:

1. I use the `STATA` missing observation (mi) code to impute productivity. Note that the demand state, M, the action, a_i, and the plant's own size state are never imputed since these are always observed in the Longitudinal Business Database (LBD), which is derived from administrative data gathered as part of the Internal Revenue Service's payroll tax collection activities. So only the ω_i variable must be imputed.

2. For incumbent firms, I only use the data from firms that *do not have imputed data*. However, the firm's state vector, $s = \{s_i, s_{-i}\}$, also contains information on other firms' productivities, and I allow these data to be imputed. In sum, ω_i is never imputed, but ω_{-i} may be.

3. For potential entrants, the data on these firms will be selected with the probability corresponding to the frequency of non-imputed data among incumbents. This can be thought of as

Dependant Var: Log Value Added	I	II	III	IV	IV (Fixed Effects)
Log Salaries	0.896	0.866	0.862	0.642	0.671
	(0.002)	(0.003)	(0.003)	(0.005)	(0.007)
Log Assets	0.041	0.033	0.040	0.323	0.265
	(0.002)	(0.002)	(0.002)	(0.005)	(0.006)
Constant	1.172	1.408	1.390	0.751	0.947
	(0.012)	(0.018)	(0.017)	(0.013)	(0.028)
No Administrative Record		X			
No Hot Imputes			X		
No ASM Years				X	
Observations	37559	30328	31282	29342	29342
R^2 (within)	83%	74%	76%	86%	65%

Note: The number of observations differs due to changes in the sample included in each regression.

Table A2: Production Function Regressions with Different Sample Selection Criteria and Output Measured as Value Added

a rough version of propensity score weighting, since it avoids having over 95% of observations come from the pool of potential entrants.

B Discrete Action Stochastic Algorithm

Algorithm Discrete Action Stochastic Algorithm (**DASA**)

1. Start in a location $l_0 = \{a_0, x_0\}$.
2. Draw an action profile $a|a_i \sim 1(a_i = a_0) \prod_{-i} P[a_{-i}|x]$ and a state in the next period x' given action profile a:
$$x'|a \sim \hat{D}[M'|M] \prod_i \iota(x'_i|a_i, x_i) \tag{28}$$
where $\iota(x'_i|a_i, x_i)$ is the *updating* function, which corresponds to the size state tomorrow (which is just the action I choose) and a draw from the productivity transition:
$$\omega'_i \sim \begin{cases} \hat{P}^{\omega I}[\omega'_i] \text{ if } x_i = 0 \text{ (initial productivity distribution)} \\ \hat{P}^{\omega}[\omega'_i|\omega_i] \text{ if } x_i > 0 \text{ (continuers productivity distribution)} \end{cases} \tag{29}$$
And finally $x'_i = \{\omega'_i, a_i\}$ is just the composition of the productivity and size state.
3. Increment the hit counter (how often you have visited the state-action pair): $h(l) = h(l) + 1$.
4. Compute the payoffs R of the action as:
$$\begin{aligned} R =& r(a_i, x) - \tau(a_i, x_i) \\ &+ \beta \sum_{j \in A} W(j, x')P[j|x'] + \beta E(\varepsilon|x', P) \end{aligned} \tag{30}$$
where $E(\varepsilon|x', P) = \left(\gamma - \sum_{j \in A} \ln(P[j|x'])P[j|x']\right)$ (where γ is Euler's Constant).
5. Update the W-function:
$$W'(l) = \alpha R + (1 - \alpha)Q(l) \tag{31}$$
where $\alpha = \frac{1}{h(l)}$.
6. Update the policy function:

   The optimal policy function given choice specific value functions $W(a_i|x)$ is just given by the logit formula:
$$\Psi(a_i|x, W, \theta) = \frac{\exp\left(W(a_i, x)\right)}{\sum_{j \in A} \exp\left(W(j, x)\right)} \tag{32}$$
The policy function is updated for state x, $P[a_i|x] = \Psi(a_i|x, W, \theta)$ for all action $a_i \in A$.
7. Draw a new action $a'_i \sim P[\cdot|x']$.
8. Update current location to $l' = \{a'_i, x'\}$.

9. The stopping rule for this algorithm is based on Fershtman and Pakes (2009) which compares the W-function to a simulated average based on rewards from steps 2 and 4 for states that are recurrent. If the Q-function is exact, then the difference between between these two objects can be accounted for by simulation error. The stopping rule is presented in below in subsection B.1.

B.1 Discrete Action Stochastic Algorithm: Termination Criteria

The stopping rule is based on the fact that if I have the "correct" W function, then it will satisfy the Bellman equation. However, it is computationally expensive to calculate the W-function exactly. Instead we can approximate the value function using forward simulation. Consider the locations $R \subset S \times A$ defined as the state-action pairs visited in the last 10 million iterations (keep a hit counter that tracks the last 10 million iterations denoted $rh(l)$).

Algorithm Fershtman-Pakes Stopping Rule (FPStop)

For all locations $l \in L$ which have been visited in the last 1 million iterations:

1. Compute the W-function using a one step forward simulation. For $k = 1, ..., K$ (I use $K =$ 10 000):

   (a) Draw an action profile a^k and a state tomorrow $x^{k\prime}$ given location tommorow $x^{k\prime}$ given location l.

   (b) Get rewards:

$$\begin{aligned} R^k =& r(a^k|x^{k\prime}, \theta) + \tau(a_i^k|x_i, \theta) \\ &+ \beta \sum_{j \in A} W(j, x^{k\prime}) P[j|x^{k\prime}] \\ &+ \beta \left(\sum_{j \in A} -ln(P[j|x^{k\prime}]) P[j|x^{k\prime}] + \gamma \right) \end{aligned} \tag{33}$$

   (c) Compute the approximation to the W-function (denote $\mathcal{W}$):

$$\mathcal{W}(l) = \frac{1}{K} \sum_{k=1}^{K} R^k \tag{34}$$

2. Compute the difference in value functions weighted by the recent hit counter rh:

$$\gamma = \frac{1}{\sum_l rh(l)} \sum_l rh(l) * (\mathcal{W}(l) - W(l))^2 \tag{35}$$

If the test statistic γ is small enough, then we can argue that we have a good approximation. In practice I have used the fact that the recent hit counter weighted R^2 between $\mathcal{W}(l)$ and $W(l)$ is greater than 0.999. This usually happens after as little as 50 million iterations, and it is usually more efficient to run the DASA for 150 million iterations (i.e. 15 minutes) which will lead to a W function which satisfy the FPStop criteria. Furthermore, in this application there are only about 3 000 state-action pairs (where the action is not 0) that are visited in the last 1 million iterations. Thus the ergodic class R is quite small compared to the size of the entire state space.

C Additional Tables and Figures

Average Shipments (in thousands)	Birth	Continuer	Death
1977	461	1,164	402
1982	1,045	1,503	520
1987	1,241	2,307	601
1992	1,509	2,218	1,417
1997	1,559	3,293	1,358

Average Capital (in thousands)	Birth	Continuer	Death
1977	217	491	185
1982	403	598	187
1987	549	1,050	270
1992	565	1,131	632
1997	728	1,992	770

Average Salaries (in thousands)	Birth	Continuer	Death
1977	83	211	83
1982	185	269	83
1987	205	413	101
1992	257	428	267
1997	243	567	241

Table A3: Characteristics of Plants that are Births, Deaths and Continuers

	Observations	Mean	Standard Deviation	5th Percentile	95th Percentile
Total Employment	70566	26	147	1	82
Administrative Record Flag	70622	0.13	0.34	0	1
Jarmin-Miranda Entry	120055	6.2%		0	1
Jarmin-Miranda Exit	119790	6.0%		0	1
Total Value of Shipments (in 000's)	70566	3380	25643	41	11000
Total Value of Inventory (in 000's)	11598	116	3702	0	140
Building Assets Ending (in 000's)	51246	153	1885	0	420
Machinery Assets Ending (in 000's)	51246	754	4463	0	2700
Machinery Depreciation (in 000's)	51246	55	478	0	220
Multi-Unit Flag	70622	0.51	0.50	0	1
Total New Expenditures (in 000's)	70566	148	1625	0	510

Table A4: Summary Statistics for Plant Data

Survey Year	Median	Median Cubic Yards of Concrete		
	Employees	Per Plant	Per Worker	Per Worker Hour
1963	8	15000	1900	1.4
1967	14	26000	2100	1.6
1972	15	35000	2200	1.6
1977	13	33000	2300	1.7
1982	13	25000	2000	1.4
1987	15	36000	2700	1.7
1992	13	32000	2600	1.7
1997	13	40000	3000	1.7

Table A5: The ready-mix concrete sector has experienced little productivity growth.

Action Chosen	I		II	
Small				
Size				
Small	4.99	(0.30)	4.71	(0.30)
Medium	4.87	(0.35)	4.53	(0.35)
Large	3.85	(0.36)	3.47	(0.36)
Productivity	1.36	(0.31)	1.35	(0.30)
Log Construction Employment	-0.12	(0.10)	-0.01	(0.11)
Log Competitors	-0.25	(0.13)	-1.62	(0.17)
Constant	-3.53	(0.26)	-3.38	(0.27)
Average Plants Rounded				
2			1.06	(0.13)
3			1.92	(0.18)
4			2.52	(0.26)
Medium				
Past Size				
Small	4.74	(0.39)	4.40	(0.39)
Medium	7.98	(0.41)	7.55	(0.41)
Large	6.01	(0.41)	5.52	(0.41)
Productivity	1.50	(0.32)	1.51	(0.32)
Log Construction Employment	0.34	(0.13)	0.50	(0.13)
Log Competitors	-0.64	(0.15)	-2.37	(0.20)
Constant	-6.25	(0.37)	-6.19	(0.39)
Average Plants Rounded				
2			1.35	(0.17)
3			2.42	(0.23)
4			3.32	(0.31)
Large				
Size in the Past				
Small	4.05	(0.48)	3.70	(0.49)
Medium	7.37	(0.47)	6.92	(0.47)
Large	9.15	(0.46)	8.61	(0.46)
Productivity	0.97	(0.35)	0.96	(0.35)
Log Construction Employment	0.74	(0.15)	0.90	(0.15)
Log Competitors	-0.39	(0.16)	-2.41	(0.23)
Constant	-8.00	(0.47)	-7.95	(0.49)
Average Plants Rounded				
2			1.56	(0.20)
3			2.82	(0.27)
4			3.89	(0.35)
Observations	14278		14278	
Log-Likelihood $\mathcal{L}$	-4066		-4166	
χ^2	17754		17555	

Note: †: Small is 0 to 6 employees, ††: Medium is 7 to 17 employees, † † †: Big is more than 17 employees. Out is the omitted category.

Table A6: Multinomial Logit Estimates of the Size and Entry Decision. Used as CCP's **P** in the dynamic model.

Dependent Variable	I	II	III	IV
Jarmin-Miranda Exit				(county fixed effects)
Size				
Medium	-0.918	-1.210	-1.209	-1.250
	(0.050)	(0.201)	(0.202)	(0.224)
Large	-1.174	-1.276	-1.260	-1.163
	(0.063)	(0.234)	(0.232)	(0.260)
Log County Construction	-0.078	-0.145	0.005	0.095
Employment	(0.015)	(0.062)	(0.067)	(0.217)
1st Competitor	0.791	0.511	0.915	0.801
	(0.058)	(0.201)	(0.220)	(0.260)
Log of competitors (above 1)	0.174	0.114	0.407	0.557
	(0.019)	(0.077)	(0.095)	(0.138)
Productivity				
2nd Quintile		0.031	0.063	0.083
		(0.190)	(0.193)	(0.221)
3rd Quintile		-0.429	-0.394	-0.093
		(0.221)	(0.222)	(0.243)
4th Quintile		-0.748	-0.720	-0.587
		(0.235)	(0.237)	(0.271)
5th Quintile		-0.890	-0.913	-0.793
		(0.245)	(0.249)	(0.284)
Average Plants in County Rounded				
2			-0.899	
			(0.194)	
3			-1.476	
			(0.269)	
4			-2.004	
			(0.362)	
Constant	-2.810	-1.746	-2.338	
	(0.084)	(0.349)	(0.373)	
Observations	64482	4627	4627	1747
χ^2	1089	109	133	102
Log-Likelihood $\mathcal{L}$	-12656	-826	-807	-396

Table A7: Logit and Conditional Logit estimates of the Decision to Exit

Market Size Category	Median Number of Plants	Productivity	Share of Plants
1	1	Low Productivity*	41%
		Medium Productivity**	33%
		High Productivity***	24%
2	3	Low Productivity	37%
		Medium Productivity	32%
		High Productivity	31%
3	5	Low Productivity	29%
		Medium Productivity	34%
		High Productivity	37%
4	16	Low Productivity	20%
		Medium Productivity	34%
		High Productivity	46%

* Lowest tercile of productivity
** Medium tercile of productivity
***Highest tercile of productivity

Table A8: In large markets a higher fraction of plants are productive.

Variable	Mean	Std.	10th Percentile	90th Percentile
NPV of periods as a small plant	7.96	1.62	5.60	10.00
NPV of periods as a medium plant	2.27	1.06	1.10	3.90
NPV of periods as a large plant	1.04	1.15	0.17	3.00
Log Construction Employment	2.20	0.51	1.60	2.70
NPV of Log Construction Employment	31.25	9.45	19.00	44.00
Productivity	0.93	0.38	0.68	1.60
NPV of Productivity	12.26	2.01	9.60	15.00
Log Competitors	1.31	1.16	0.00	3.00
NPV of Log Competitors	9.51	6.24	1.50	19.00

Note: Net Present Values (i.e. NPV) are simulated. The number of observations is $N = 4113$ in all regressions.

Table A9: Summary Statistics for Net Present Values

D Effects of Age

In the ready-mix concrete industry there is no evidence that a firm's age has an impact on its exit decision.[31] Figure A1 shows that the exit hazard decreases gently with age. Younger firms could be less likely to exit if they delay shutting down in order to accumulate information on their underlying productivity. In a Bayesian learning model, older firms use a higher productivity cutoff for exiting than younger firms. If younger firms have less information about their true productivity level, a younger firm has greater option value than an older firm, for a given realization of productivity,. Figure A2 shows that the productivities of the 25^{th}, 50^{th} and 75^{th} percentile plants do not increase with age, indicating that firms face little uncertainty on the permanent component of their productivities. [32]

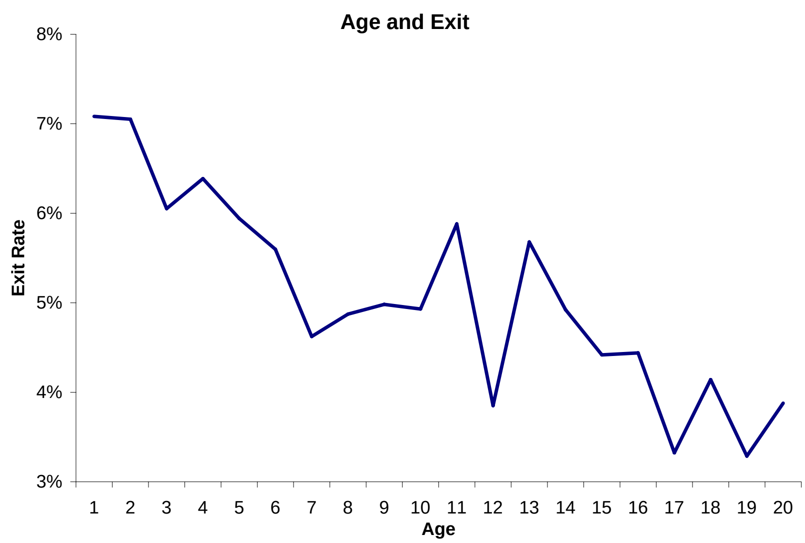

Figure A1: Older firms are slightly less likely to exit.

[31]In contrast, Abbring and Campbell (2003) find evidence that bars in Texas face considerable uncertainty about their profitability in their first year in operation.

[32]Moreover, Figure **??** shows the average number of employees at a plant rises dramatically in a plant's first year, and subsequently grows slowly. Pakes and Ericson (1998) discuss the empirical content of the passive learning models in the Jovanovic (1982) tradition. They show that one of the few empirical implications of the passive learning model is that the expected firm size is increasing in the previous size of the firm. Pakes and Ericson (1998) do not have data on plant level productivity, so their tests of the passive and active learning models are not based on plant-level productivity.

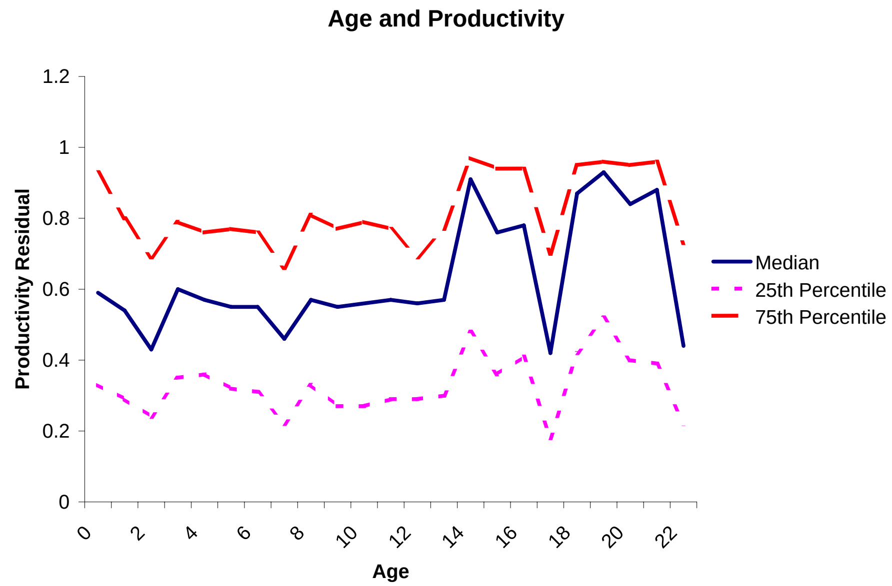

Figure A2: Plant Productivity shows little change as it ages.

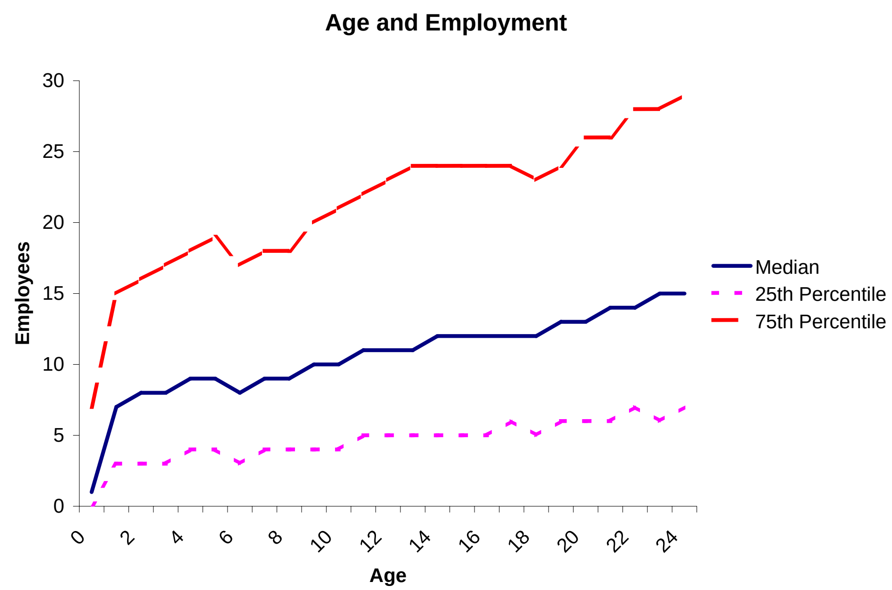

Figure A3: Average plant employment rise slowly after the first year in operation.

E Different Measures of Output in Estimating Production Functions

I use three different definitions of output which generate corresponding measures of productivity:

1. Value Added (henceforth VA).
2. Total Value of Shipments (TVS).
3. Cubic Yards of Concrete (CYC).

Productivity residuals generated by these measure of output are somewhat correlated. Measures based on revenues (VA and TVS) are highly correlated. The measure based on quantity (CYC) has a weak correlation with the revenue measures (VA and TVS).

The first measure of productivity is the residual from the log-linear production function OLS regression:

$$y_i^t(\text{value added}) = \beta_l l_i^t(\text{salaries}) + \beta_k k_i^t(\text{capital}) + A^t + \rho_i^t \tag{36}$$

where a lower case variable denotes the logarithm of the actual variable, A_t is the intercept of the production function for each year (so that year to year changes in technology do not affect the dispersion of productivity) and ρ_i^t is a plant's productivity. I deflate all items measured in dollars by the producer price index (PPI). [33]

The second measure of productivity is based on total value of shipments, the KLEM production function:

$$\begin{aligned} y_i^t(\text{total value of shipments}) = \quad & \beta_l l_i^t(\text{salaries}) + \beta_k k_i^t(\text{capital}) \\ & + \beta_m m_i^t(\text{cost of materials}) + A^t + \rho_i^t \end{aligned} \tag{37}$$

Finally, I measure productivity based on cubic yards of concrete produced by a plant. This is an effective measure of output for the ready-mix concrete industry since ready-mix concrete is a homogeneous good and ready-mix concrete is a plant's sole output:[34]

$$\begin{aligned} q_i^t(\text{cubic yards of concrete}) = \quad & \beta_l l_i^t(\text{salaries}) + \beta_k k_i^t(\text{capital}) \\ & + m_i^t(\text{materials}) + A^t + \rho_i^t \end{aligned} \tag{38}$$

Table A10 presents production function regressions with output defined as value added (VA), total shipments (TVS) or volume of ready-mix concrete (CYC). Notice that the TVS and CYC regressions have similar coefficients for materials, capital and labor. One might worry that the total value of shipments mismeasures output. Prices are higher in more concentrated markets, and thus the value of shipments will also be higher.[35] The total volume of concrete produced can also

[33] It is important to deflate data measured in dollars since the log-linear production function is not linearly homogeneous if the sum of the capital and labor coefficients differs from one, and thus is sensitive to rescaling variables.

[34] See Collard-Wexler (2006) for evidence on the preponderance of ready-mix concrete in a plant's output.

[35] Figure A2 in Collard-Wexler (2006) presents illustrates higher prices in markets with fewer plants.

mismeasure output, since ready-mix concrete trucks do not need to drive as far in denser markets to make a delivery. A plant located in a dense market can economize on trucks, drivers and fuel giving the (mistaken) appearance that it is a more efficient operation. The "fit" of the regression as measured by the R^2 is much higher for the total shipments regression than the cubic yards of concrete regression. Moreover, capital and labor coefficients are higher in the shipments regression than in the volume regression. This could indicate either that different types of concrete require different levels of inputs, or that the price level for concrete is correlated with the price level for labor and capital (perhaps in the form of higher land prices for the latter).

	Output Measure		
	Log Value Added	Log Shipments	Log Cubic Yards of Concrete
Log Salaries	0.633	0.270	0.138
	(0.006)	(0.003)	(0.012)
Log Assets	0.269	0.116	0.084
	(0.006)	(0.003)	(0.010)
Log Materials		0.587	0.689
		(0.003)	(0.011)
Constant	1.163	1.170	4.366
	(0.022)	(0.011)	(0.042)
Observations	22114	21941	15636
R2	74%	94%	58%

Table A10: Production function regressions with different output measures.

E.1 Log Shipments as an Output Measure

Table A11 reports the results if one use a production function where shipments are the firm's output rather than value added:

$$y_{it} \text{ (shipments) } = \beta_0 + \beta_l \text{ (salaries) } + \beta_k \text{ (assets) } + \beta_m \text{ (materials)} \tag{39}$$

	Dispersion due to		
Percentile	TFP (ρ_q)	Productivity (ω_q)	Measurement (ϵ_q)
10%	1.7	1.8	1.8
25%	1.8	1.9	1.9
50%	2.0	2.0	2.0
75%	2.3	2.2	2.3
90%	2.9	2.3	2.6

Table A11: Dispersion of Predicted Output due to TFP dispersion, true productivity, and measurement error (in millions of dollars) using shipments measure of output.