Chapter 1

Heterogeneity, Productivity and Reallocation: Evidence from the U.S. Cement Industry

Abstract

This paper develops a structural model of spatial competition to estimate plant-level total factor productivity in the U.S. Portland cement industry, where market power makes standard production function approaches inconsistent.[1] Using plant-level data from 1995 to 2007, I implement a simulation-based estimation framework that allows for both observable and unobservable sources of productivity heterogeneity. I find that there is notable productivity dispersion: plants in the top decile are on average 25% more productive than those in the bottom decile. Decomposing the sources of this dispersion, I find that observable factors, particularly kiln technology, account for approximately 95% of the variance in total factor productivity, leaving little role for unobserved plant-specific differences. A counterfactual carbon tax analysis exploits

[1]I am grateful to Nathan Miller and Mathiew Osborne for generously sharing the PCA data, which they collected, cleaned, and compiled. I also thank Heski Bar-Isaac and Ambarish Chandra for helpful comments and suggestions on this chapter, as well as the Competition Bureau of Canada.

this technology heterogeneity to show that environmental regulation reallocates output toward cleaner plants, while disproportionately burdening carbon-intensive producers.

1.1 Introduction

Understanding heterogeneity in productivity among competing firms is crucial for assessing market efficiency, resource allocation, and economic growth. When highly productive firms are unable to expand or when less productive firms continue to operate, aggregate efficiency suffers, leading to lower total factor productivity (TFP). Misallocation, where inputs are not allocated to their most efficient use, can substantially reduce productivity in many industries and countries. Hsieh and Klenow (2009) find that eliminating misallocation could raise aggregate TFP by 30–50% in China and 40–60% in India, highlighting the economic significance of firm heterogeneity. Restuccia and Rogerson (2008) similarly argue that policy distortions and financial frictions prevent productive firms from expanding, leading to persistent productivity dispersion. These findings highlight the importance of understanding productivity dispersion in different industries.

One standard approach to estimating firm-level productivity in the industrial organization literature is to estimate the firms' production function, as developed by Olley and Pakes (1992). They propose a control function approach to consistently estimate the production function parameters. This method has since been widely used in the literature.

However, the production function approach has also been criticized. For example, Bond et al. (2021) find that if the researcher does not have access to firm-level quantities and prices separately and only has data on revenue—which is a common limitation in available datasets—the production function cannot be identified when

firms in the industry have market power.

Related to this issue, Foster et al. (2008) distinguish between revenue-based and quantity-based measures of TFP. They argue that when firm-level prices are not observed, TFP measures may conflate market power with true production efficiency. Therefore, to accurately identify and estimate TFP, it is necessary to explicitly model imperfect competition while allowing for heterogeneity in firms' productivity.

This paper addresses this challenge by developing a novel structural model of imperfect competition in the Portland cement industry, where firms differ in both observable and unobservable productivity. The Portland cement industry is a well-suited setting for this analysis, as firms are known to exercise significant regional market power due to high transportation costs, making standard production function approaches unreliable. I examine how firms make production decisions within a Cournot framework for a homogeneous product. The model incorporates regional-level data from the United States Geological Survey (USGS), plant-level data from the Portland Cement Association (PCA), and additional regional data on wages and energy costs, enabling the estimation of a model in which firms compete by setting output quantities. The market price for cement is determined by regional demand and the aggregate output of firms, reflecting local variations in both cost conditions and competitive structures.

This paper estimates plant-level TFP in a setting where market power would otherwise confound standard approaches, and uses these estimates to shed light on the sources of productivity dispersion and its implications for environmental policy in the U.S. Portland cement industry. The main findings are threefold. First, this paper documents the dispersion in productivity in the U.S. cement industry. The findings on productivity dispersion indicate that plants in the top decile of the productivity distribution are, on average, more than 25% more productive than those in the bottom decile. Similarly, plants in the top quintile are, on average, more than 20%

more productive than those in the bottom quintile. This degree of dispersion, while modest relative to other industries, is consistent with some degree of misallocation across plants, in the sense that output is not entirely concentrated among the most productive producers. This is generally consistent with the literature that has found significant dispersion in productivity even within industries, though the magnitudes documented elsewhere tend to be somewhat larger. There are several reasons why lower levels of dispersion might be expected in this context; Portland cement is a standardized product, with almost no room for product differentiation or innovation in quality. Moreover, the industry examined in this paper is defined narrowly, which naturally leads to lower heterogeneity. For instance, Syverson (2004) finds that the average 90th-to-10th percentile ratio in productivity is 1.92 for industries defined at the 4-digit NAICS level. Portland cement, a subset of the 5-digit NAICS code 32731, is defined more narrowly than this, and lower dispersion is therefore to be expected.

Beyond documenting the extent of productivity dispersion, the structural framework developed in this paper allows for a decomposition of its sources into observed and unobserved factors. Olley and Pakes (1992) show that there is a significant dispersion in unobserved productivity which is firm-specific and persistent over time. Since then, a growing body of literature has attempted to open the black box of "unobserved" productivity by identifying concrete drivers of persistent differences, including managerial quality, technology adoption, human capital, and market frictions such as misallocation. However, this paper finds that, after controlling for observable factors, only 5% of the variation in total factor productivity in the Portland cement industry is attributable to unobserved factors. Instead, the primary driver of productivity differences is the choice of technology with which plants operate. One plausible interpretation is that this pattern reflects the substantial fixed costs of switching kiln technology, combined with the regional market power documented in this paper, rather than differences in management quality or other unobserved plant

characteristics. In geographically segmented markets with limited local competition, less efficient plants may face weaker competitive pressure to adopt newer technology than they would in a more contested market. Section 6 discusses how this finding can be understood in light of the institutional features of the cement industry.

Third, because productivity differences are driven primarily by technology rather than unobserved plant characteristics, policy that changes the relative cost of technologies, such as a carbon tax, can be expected to reallocate output and affect producer surplus in a predictable way across plants, rather than in a way confounded by unmeasured firm-level advantages. In addition to its economic importance, the cement industry is a major source of carbon emissions, with production accounting for nearly 5% of global CO_2 emissions (Miller et al. (2017)). Environmental regulations such as carbon pricing and emissions standards may impose uneven costs across firms, potentially affecting the distribution of producer surplus and the reallocation of inputs across firms. Using a counterfactual analysis, I examine the effects of a carbon tax on cement, with a focus on reallocation effects among firms based on their technology use. I find that such an intervention reallocates output toward firms using more "green" technology. A flat output tax equal to the estimated social cost of carbon in 2007 disproportionately penalizes plants that rely on more carbon-intensive technologies: the total market share of these plants falls from 18% to 14%, with nearly all of that share absorbed by plants using the most carbon-efficient technologies. Moreover, despite holding only 14% of the total market share, these firms bear 26% of the total loss in producer surplus. Although this paper does not model long-term investment decisions, the results suggest that the application of a carbon tax may incentivize firms to upgrade to cleaner technologies.

The remainder of the paper is structured as follows: Section 1.2 reviews the relevant literature. Section 1.3 examines the institutional background of the cement industry. Section 1.4 describes the datasets used in this paper. Sections 1.5 and

1.6 present the model and estimation strategy. Section 1.7 reports the estimation results, including estimates of markups and productivity distributions across regions. Section 1.8 runs a counterfactual analysis, focusing on the effects of a carbon tax on total welfare and the reallocation of output among firms. Section 1.9 concludes by summarizing the key insights and potential policy implications of the research.

1.2 Literature

This section begins with a discussion of the production function estimation approach, which plays a central role in understanding firm-level productivity. I review its standard implementation, discuss its limitations which are raised in the literature, and explain how the model developed in this paper aims to address these concerns. The section then turns to the related empirical literature.

Suppose firm i's production function is as follows:

$$q_i = \alpha v_i + \beta k_i + (\omega_i + \epsilon_i), \tag{1.1}$$

where q_i is the logarithm of quantity, v_i is the logarithm of the firm's variable input, k_i is the logarithm of the firm's capital, ω_i is logarithm of a firm's TFP, and ϵ_i is an independent random shock. ω_i is unobservable by the researcher but observed by the firm. A simple OLS yields inconsistent estimates because the firm's input choices are correlated with the error term ω_i. Olley and Pakes (1992) develop a control function approach to consistently estimate the production function parameters.

However, the production function estimation approach has also been criticized in the literature. For example, Bond et al. (2021) find that if the researcher does not have access to firm-level quantities and prices separately and only has data on revenue—which is a common limitation in available datasets—the production function cannot be identified when firms in the industry have market power. To illustrate this,

rewrite the production function equation using revenue instead of physical output as the dependent variable:

$$r_i = \alpha v_i + \beta k_i + (p_i + \omega_i + \epsilon_i), \tag{1.2}$$

where r_i is the logarithm of revenue, and p_i is the logarithm of price. Note that if firm i has market power, its choice of variable input v_i influences the market price, making it an endogenous variable. In this case, any firm-specific shock that leads to an increase in v_i, results in higher production and lower prices. Consequently, any instrument correlated with v_i will be correlated with the error term. Therefore, we cannot identify the parameters using only revenue data. This paper develops a structural model to estimate the production function while explicitly accounting for imperfect competition.

This paper is closely related to a body of research documenting within-industry productivity dispersion in various industries. Salter and Reddaway (1966) find significant productivity differences among pig-iron plants from 1911 to 1926, where the most efficient firm produced twice as much per labor hour as the industry average. Chew et al. (1989) document substantial intra-firm productivity differences across plants within a commercial food operation, while Womack et al. (2007) report that productivity in U.S. automobile assembly plants ranges from 13.2 to 78.7 labor hours per unit of output. More recently, Syverson (2004) shows that the top 10% of firms in terms of productivity are, on average, 1.92 times more productive than the bottom 10% at the 4-digit NAICS code level. Collard-Wexler (2011) finds that top-quintile firms in the ready-mixed concrete industry are 1.5 times more productive than the bottom quintile. Hubmer et al. (2024) document that, in Canada, firms at the 90th percentile are 36.2% more productive than those at the 10th percentile. This paper contributes to and extends this literature by examining the productivity distribution in the Portland cement industry and documenting that substantial dispersion exists

in this sector as well.

Further, this paper relates to a large literature on the drivers of firm-level productivity differences. Identified factors include management quality (Bloom et al. (2007)), labor quality (Ilmakunnas et al. (2004), Fox and Smeets (2011)), capital quality (Sakellaris and Wilson (2004), Van Biesebroeck (2003)), research and development (Doraszelski and Jaumandreu (2013), Aw et al. (2011)), learning-by-doing (Benkard (2000), Thornton and Thompson (2001)), product innovation (Lentz and Mortensen (2008), Balasubramanian and Sivadasan (2011)), and technology (Miller et al. (2023), Collard-Wexler and De Loecker (2015)). Yet not all productivity dispersion is accounted for after controlling for these variables. For example, Bloom et al. (2007) find that management quality explains only 10% to 23% of the interquartile range in TFP in their data, and Fox and Smeets (2011) find that labor quality explains just 11% of TFP variation, leaving a substantial portion unexplained. This paper contributes to the literature by decomposing the drivers of productivity dispersion into both observable and unobservable factors. Once controlling for technology, I find that unobserved factors play only a minor role in explaining productivity differences across firms in the U.S. cement industry.

This paper also contributes to the literature on environmental regulation and firm behavior. Cement production is highly energy-intensive, and environmental regulations impose cost shocks that may not be evenly distributed across firms. Ryan (2012) estimates the costs of environmental regulation in the cement industry and examines how emissions policies influence firm entry and exit. Fowlie et al. (2016) study the welfare effects of different environmental policy instruments. More broadly, Greenstone et al. (2012) analyze the impact of environmental regulation on U.S. manufacturing and find that regulatory costs affect both productivity and competitiveness across industries. Finally, Miller et al. (2017) estimate the price pass-through of fuel cost changes in the cement industry, showing that firms absorb approximately 11% of the

total regulatory burden. This paper contributes to this strand of research by focusing on how environmental policy affects the reallocation of output among heterogeneous firms in the short run. I find that policy intervention imposes a smaller burden on firms with more environmentally friendly technologies compared to others.

Finally, this paper is related to the literature on the evolution of market power. De Loecker et al. (2020) document a long-term rise in markups across U.S. industries. A growing body of research has since examined the evolution of markups in specific sectors. Miller et al. (2023) develop a structural model of spatial competition and show an upward trend in markups in the cement industry. The markup estimates reported in this paper are consistent with these findings, indicating that markups have indeed increased between 1995 and 2007.

1.3 Industry Background

The cement industry plays a crucial role in infrastructure and development. Cement is a critical input in the construction of roads, bridges, and buildings. Cement production is an energy-intensive process; the primary material, limestone, undergoes a process of heating and chemical transformation to create clinker, the key ingredient in cement. This transformation occurs in rotary kilns at extremely high temperatures, typically around 1400°C. Kilns are generally powered by coal and petroleum coke, though some plants incorporate alternative fuels to reduce costs and emissions.

The technology used in production varies, with older "wet" kilns requiring higher energy inputs than "dry" kilns, which are designed for energy efficiency. Technological upgrades from wet to dry kilns have led to significant cost reductions in fuel usage and production, though the capital investments required for these transitions are substantial. The choice of technology directly impacts both operational costs and emissions, with dry kilns producing lower emissions per unit of cement than their wet

counterparts. During the period of my analysis, four main technologies are observed in use: the wet process, dry process, dry process with preheater, and dry process with precalciner. The wet process, the oldest and least efficient, involves mixing raw materials with water to form a slurry before entering the kiln, requiring high energy consumption to evaporate the water before clinker formation. Over time, many plants transitioned to the dry process, where raw materials are ground into a fine powder before entering the kiln, reducing fuel usage and improving efficiency. Further advancements introduced preheater kilns, which use a series of cyclones to heat raw materials before they reach the kiln, allowing for better energy recovery from exhaust gases. The most advanced system, the precalciner kiln, adds a separate combustion chamber where most of the raw material calcination occurs before entering the kiln, further lowering energy requirements and CO_2 emissions. These technological improvements have played a crucial role in shaping production costs, plant-level efficiency, and environmental impact, influencing the competitive dynamics within the cement industry.

The U.S. cement industry is characterized by a homogeneous product and a relatively concentrated market structure, with a limited number of producers serving large geographic regions. The high transportation costs associated with shipping heavy cement materials limit competition across distant regions, leading to the development of localized markets where a few firms hold substantial market power.

Because of this regional market structure, firms often exercise market power to enjoy a higher surplus. Miller and Osborne (2014) discuss how firms' pricing strategies are shaped by these regional constraints, noting that spatial competition in the industry is influenced not only by physical distance but also by the distribution of customers. The absence of product differentiation in cement production makes plant-level productivity a key determinant of profitability and market share.

Cement manufacturing is one of the largest industrial sources of greenhouse gas

emissions, contributing around 5% of global CO_2 emissions (Miller et al. (2017)). The environmental impact arises both from the combustion of fossil fuels required to achieve the high temperatures in kilns, and from the chemical reactions that occur when limestone is converted to clinker.

1.4 Data

This paper uses data from several sources to examine production, pricing, plant characteristics, and input costs in the U.S. cement industry. The primary source for market-level production and pricing is the United States Geological Survey (USGS), which provides data spanning from 1932 to 2022. The analysis focuses on the period 1995 to 2007, during which the USGS reports data for 26 regions each year, and consistent plant-level and market-level data are available across all sources used in this paper. The dataset includes total cement production, regional consumption, average shipment prices, and the number of operating plants in each region. To preserve confidentiality, the USGS defines regions to include at least three independent plants, ensuring that individual plant-level data cannot be identified. I define each USGS region as a distinct market. These market-level data are used to estimate demand functions and to construct the moment conditions necessary for estimating the supply-side parameters of the model. Table 1.1 presents summary statistics for production, consumption, and average shipment prices.

Second, I use data from the Portland Cement Association (PCA), which provide plant- and kiln-level information. Specifically, I use the geographic location of each plant to match it to its corresponding USGS region (market). The dataset also includes the technology type of each kiln and the annual production capacity of each plant, which I use as a proxy for capital. During the period covered in this analysis, four kiln technologies are observed: wet kilns, dry kilns, dry kilns with preheater,

Table 1.1: Summary Statistics of the USGS data.

Variable	Obs	Mean	Std. Dev.	Min	Max
Production	325	3465.055	1648.509	550.3547	9271.874
Consumption	325	3784.337	1985.845	581.981	11575
Price	325	68.251	6.431	50.871	90.059
Number of Operating Plants	325	4.21	1.38	2	8

Unit of observation is region-year.
Production and consumption are reported in thousands of metric tonnes.
Price is reported in dollars per metric tonnes.
Alaska is excluded from the analysis.
26 regions observed over 13 years (1995–2007), for 325 region-year observations.

and dry kilns with precalciner. Table 1.2 reports the frequency of each technology in the dataset. Older wet kilns account for 26% of total kilns, while the most modern technology—dry kilns with precalciner—is used in 17.9% of kilns.

Table 1.2: Frequency of Different Technologies in Portland Cement Industry

Technology Type	Frequency
Dry	295
Dry Preheater	561
Dry Precalciner	275
Wet	406

Unit of observation is each individual kiln.

Third, I obtain input prices from two sources. Average wage data are sourced from the Bureau of Labor Statistics (BLS), which reports regional averages for the cement and concrete industry. Energy prices—including coal and electricity—are obtained from the U.S. Energy Information Administration (EIA) and are used as variable input costs in the production function. Table 1.3 presents summary statistics for these input prices. To ensure consistency across datasets and over time, all prices are adjusted to real 1995 dollars, allowing for accurate cost and price comparisons.

Table 1.3: Summary Statistics of the Variable Input Prices in dollars

Variable	Obs	Mean	Std. Dev.	Min	Max
CoalPrice (per metric tonne)	318	1.88	.21	1.61	2.19
ElectricityPrice (per 100 kWh)	318	15.2	1.2	14.1	18.3
Wage (annual salary)	318	32307.4	3423.6	24136.5	39785.7

All prices are expressed in constant 1995 dollars.

Unit of observation is region-year.

Wage data are sourced from the U.S. Bureau of Labor Statistics (BLS).

Coal and electricity prices are from the U.S. Energy Information Administration (EIA).

1.5 Model

This paper builds on the existing literature in modeling competition and demand in the Portland cement industry, while introducing novel elements on the supply side. The nature of competition is modeled as Cournot, and market demand follows a constant price elasticity specification, as in Ryan (2012). The key innovation lies in embedding a production function with a dynamic, plant-level productivity process that allows for persistent unobserved heterogeneity within a fully solved Cournot equilibrium, estimating both jointly via simulated method of moments, in contrast to standard proxy-variable approaches, which do not model strategic interdependence between firms' output and pricing decisions. This specification is motivated by the capital-intensive nature of cement production, where technology choices are long-lived and productivity differences across plants are likely to be persistent, and is central to the paper's aim of decomposing productivity dispersion into observed and unobserved components.

1.5.1 Supply

Consider a market structure with N distinct markets,[2] each containing I competing plants that produce Portland cement. Since transportation costs of cement are quite

[2]By market, I mean each USGS region.

significant, I make the assumption that buyers do not purchase cement from other markets. In each market n, plants compete by selecting their production quantities in a Cournot model, assuming products are perfect substitutes.[3] In this setup, each plant independently chooses its output level, with the aggregate output of all plants in the market determining the market price according to a constant elasticity demand function.[4] Input prices for coal, electricity, and labor are treated as exogenous, with plants taking them as given.

To produce cement, each plant requires both variable and fixed inputs. Variable inputs include labor, electricity, and coal, while fixed inputs consist primarily of capital, which sets each plant's production capacity. The production function for each plant i in year t is specified as

$$q_{it} = \omega_{it} V_{it}^{\alpha} K_{it}^{\beta} \quad , \tag{1.3}$$

where q_{it} is the amount of cement produced, V_{it} is the "aggregated" variable inputs, ω_{it} is the productivity level, and K_{it} is the capacity of production for plant i. V_{it} is a Leontief function of labor, electricity, and coal which are the main variable inputs in the cement industry, specified as follows:

$$V_{it} = Min\{\alpha_{coal} Coal, \alpha_{Elect} Elect, \alpha_L L\}.^{5} \tag{1.4}$$

[3]This is a simplifying assumption, as transportation costs within a USGS region may still give more distant plants some market power relative to their regional rivals. For example, Miller and Osborne (2014) show that accounting for such intra-regional transportation costs can lead to significant differences in the predicted market power compared to those generated by the model used in this paper. While this limitation could affect the estimated productivities of plants, the approach in this paper has the advantage of allowing for plant-level unobserved heterogeneity in productivity which is much more difficult to incorporate in models that take spatial differentiation into account.

[4]In 88% of the observations, different plants in the same market and year do not have the same owner. And almost in all of the remaining observations, at most 2 plants in the same market-year have the same owner. Hence, the assumption of plants operating independently seems plausible.

[5]The use of a Leontief structure for variable inputs is common in production-function estimation approaches (e.g., De Loecker and Scott (2016), Ackerberg et al. (2015), and Gandhi et al. (2011)). Cement production is a highly standardized, process-driven technology, making a fixed-proportions assumption for variable inputs a plausible short-run approximation. An additional advantage of the Leontief specification is that it allows the estimation of the production function using observed input

At the beginning of year t, each plant first sets its production capacity K_{it}, then it observes it productivity ω_{it}, and finally decides on its variable inputs.

It is worth noting that as plants are run more intensively, the likelihood of a breakdown increases. The rationale for choosing a Cobb-Douglas production function is that it ensures an upward-sloping variable cost function, which is important for capturing the cost structure of the cement industry. In particular, as plants approach their capacity limits, the cost of maintenance and wear-and-tear increases, leading to rising marginal costs. To reflect this feature of the industry, the production function must generate increasing variable costs as output expands. The Cobb-Douglas specification satisfies this property and is therefore appropriate for modeling production in this context.

In order to model ω_{it}, I assume productivity follows a Markov process, implying that current productivity depends on its previous value, plus a random shock with standard deviation σ_ϵ:

$$\omega_{it} = \eta_0 + \eta_1 \omega_{it-1} + \epsilon_{it} \tag{1.5}$$

This equation is used to draw productivity in year t provided that ω_{it-1} is already known. This means that the productivity of the first period in which a plant operates should be specified independently of equation 1.5. Since the initial capital of a plant might be correlated with its productivity level, drawing first-period productivity from the same distribution for all plants might cause endogeneity issues. To address that, I use the method suggested by Wooldridge (2005) by modeling the initial condition of plants' productivities as a function of observable plant features. Formally, the first-year productivity ω_{i0} is specified based on technology, age and the capacity of each plant.

prices, even when plant-level quantities of individual variable inputs are unavailable.

$$\omega_{i0_i} = e^{\beta_1 K_{i0_i} + \beta_2 Tech_i + \beta_3 Age_i + \sigma_w . z_{it}}, \tag{1.6}$$

where ω_{i0_i} is the productivity of plant i at its first period of operation, K_{i0_i} is the capacity, $Tech_i$ is a categorical variable which indicates the technology of each plant, and Age_i is the plant age.[6] σ_w is a term which determines the productivity dispersion among plants that are not observable by the researcher.

1.5.2 Demand

Each market n has a constant elasticity demand function for Portland cement. The demand curve for market n in year t is then defined as:

$$ln(Q_{nt}) = A_{0n} + A_1 ln(p_{nt}), \tag{1.7}$$

where Q_{nt} is the total production of cement by all plants in market n in year t. Here, A_{0n} captures the demand shifter specific to market n, while A_1 reflects the price elasticity of demand.

1.5.3 Equilibrium

To determine the equilibrium in each market, I first derive the variable cost function of each plant as a function of its output and capacity.

$$\begin{aligned} C(q_{it}) &= p_{coal_{nt}} Coal_{it} + p_{elect_{nt}} Elect_{it} + wage_{nt} L_{it} \\ &= V_{it} \left(\frac{p_{coal_{nt}}}{\alpha_{coal}} + \frac{p_{Elect_{nt}}}{\alpha_{Elect}} + \frac{wage_{nt}}{\alpha_L} \right) \\ &= p_{v_{nt}} \omega_{it}^{-1/\alpha} q_{it}^{1/\alpha} K_{it}^{-\beta/\alpha}, \end{aligned} \tag{1.8}$$

[6] Age of plants are included in case plant capitals, such as kilns, become depreciated over time and therefore gain less efficiency.

where $p_{v_{nt}}$ is the price of the "aggregated" variable inputs, with $p_{v_{nt}} \equiv \frac{p_{coal_{nt}}}{\alpha_{coal}} + \frac{p_{elect_{nt}}}{\alpha_{Elect}} + \frac{wage_{nt}}{\alpha_L}$.

Each plant maximizes its profit, taking the output levels of competing plants as given. The profit maximization problem for plant i is:

$$\pi_{it} = Max_{q_it}\{p_{nt}q_{it} - C(q_{it})\} = Max_{q_{it}}\{\gamma Q_n^{1/A_1} q_i - p_{v_{nt}}\omega_{it}^{-1/\alpha} q_{it}^{1/\alpha} K_{it}^{-\beta/\alpha}\}, \quad (1.9)$$

where $\gamma_n = e^{A_{0n}^{-1/A_1}}$. Equation 1.9 is derived from replacing p_{nt} from equation 1.6 and replacing the variable cost function with equation 1.8.

Taking the first order condition with respect to q_{it} we have:

$$\frac{\partial \pi_{it}}{\partial q_{it}} = \gamma_n Q_{nt}^{1/A_1} + \gamma_n \frac{1}{A_1} Q_{nt}^{\frac{1}{A_1}-1} q_{it} - p_{v_{nt}}\omega_{it}^{-1/\alpha} K_{it}^{-\beta/\alpha}(\frac{1}{\alpha} q_{it}^{\frac{1}{\alpha}-1}) = 0, \quad (1.10)$$

where $Q_{nt} = q_{it} + q_{-it}$ is the total market quantity. Equation 1.10 will give a system of equations for each market n that will determine the equilibrium quantities of all plants.

1.6 Empirical Strategy

1.6.1 Demand

To estimate the parameters of the demand equation, I utilize region-level data on quantity and price from the United States Geological Survey (USGS). Since cement demand may vary across regions due to unobserved factors like local infrastructure projects or economic conditions, I include region-specific fixed effects to control for this unobserved regional heterogeneity. However, estimating demand parameters presents

a challenge due to the potential endogeneity of prices: price variations may reflect both supply-side cost changes and shifts in demand.

To address this endogeneity, I follow Ryan (2012) to estimate the demand elasticity: I employ instrumental variables (IVs) that affect the supply side of the market but are assumed to be exogenous to cement demand. Specifically, I use regional coal and electricity prices, as well as regional average wages, as instruments. These cost shifters influence production costs, thereby affecting prices, but they are assumed to be exogenous to the demand for cement itself. This IV approach enables me to obtain consistent estimates of the demand elasticity, represented by the parameter A_1 in the demand function.

1.6.2 Supply

Estimating the supply-side parameters is more involved due to unobserved plant-level productivity, which varies across plants and influences production decisions and cost structures. To overcome this issue, I employ the Simulated Method of Moments (SMM), which allows for the estimation of supply-side parameters by simulating market outcomes and iteratively adjusting parameters to minimize discrepancies between simulated and observed quantities. The SMM procedure has the following steps:

1. Initial Guess of Parameters: I begin with an initial guess for all supply-side parameters, including productivity process parameters and production function coefficients.

2. Simulating Productivities: For each plant and each simulation draw:

    - In the first period the plant appears, productivity is drawn from a log-normal distribution. The mean of this distribution depends on observable plant characteristics (equation 1.6), and the standard deviation is given by the parameter σ_ω.

- In subsequent periods, productivity evolves according to a Markov process (equation 1.5), where current productivity depends on the previous period's value and a random shock $\epsilon_{it} \sim \mathcal{N}(0, \sigma_\epsilon^2)$.

3. Solving for Equilibrium Quantities: Given the simulated productivities and the current parameter values, I solve the system of equations (equation 1.10) for each market-year to determine the predicted equilibrium quantities Q^*_{itk} for all plants.

4. Computing the Objective Function: The model's fit is assessed by aggregating the simulated plant-level quantities up to the market level, and computing the sum of squared differences between the predicted and observed market-level quantities across all markets, years, and simulation draws:

$$\Sigma_n \Sigma_t \Sigma_k (Q^*_{ntk} - Q_{nt})^2,$$

where $Q^*_{ntk} = \Sigma_i Q^*_{itk}$ is the simulated total quantity for market n in year t under simulation draw k, and Q_{nt} is the observed market-level (USGS) quantity of Portland cement produced.

5. Optimization: Steps 2 through 4 are repeated for each candidate set of parameters. The parameter estimates are then obtained by minimizing the objective function.

1.7 Results

1.7.1 Demand-Side parameters

I use the instrumental variable technique to deal with endogeneity issues in the demand equation. The final specification used in this paper uses gas prices, electricity

prices, and coal prices as the instrumental variable. The first-stage partial F-statistic is 20.12. Table 1.4 depicts the parameter estimates of equation 1.7. The price elasticity is equal to -2.334, consistent with Ryan (2012) which uses the same approach to estimate the demand curve.

Table 1.4: Demand Curve Estimation Results

Variables	Estimates
Price	-2.334*** (0.480)
Region fixed effects	Yes

Results from estimating Equation 1.7. The analysis includes 325 observations. Alaska is excluded due to insufficient data.

Robust standard errors in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

Full region fixed effect estimates are reported in Table 1.6 in the Appendix.

1.7.2 Supply-Side parameters

Table 1.5 presents the supply-side parameters of the model as defined in equations 1.3 to 1.6. In this section, I attempt to interpret the estimated parameters. The parameter β_1, which measures the sensitivity of productivity to total capacity, has an estimate close to zero and statistically insignificant. This suggests that total capacity has little impact on productivity and that there is likely no major endogeneity issue related to the first-period choice of capacity.

The coefficient on plant age, β_3, is estimated at -0.00243, which is very close to zero and statistically insignificant. Given its small magnitude, this suggests that plant age has essentially no effect on TFP in this industry, conditional on the other variables.

The parameters α_{coal}, α_{elect}, and α_L, which represent the proportional usage of coal, electricity, and labor in the production process) are estimated at 0.162, 0.0244, and 420.5, respectively. Among these, only α_{elect} is statistically significant. These

values mean that each 1000 tonne of coal, 660 MWh electricity, and 0.39 unit of full-time equivalent worker per year are used in proportion to produce Portland cement.

The standard deviation parameter estimates for productivity shocks, $\sigma_w = 0.012$ and $\sigma_\epsilon = 0.005$, are both estimated to be statistically insignificant. These small values suggest that unobserved shocks and measurement error contribute only minimally to productivity variation across plants. In other words, unobserved heterogeneity is not a significant factor in explaining productivity dispersion in this industry.

Finally, the parameters η_0 and η_1 capture productivity persistence and any drift in productivity over time, respectively. Consistent with the previous literature, the persistence parameter is estimated to be 0.98, indicating that a plant's productivity level is highly persistent year to year. The intercept in the productivity evolution equation is about 0.0322. This suggests a very slight upward drift in productivity over time, perhaps due to gradual industry-wide improvements or learning-by-doing. Both parameters are statistically significant.

Figure 1.1 shows the model's prediction of national cement production versus the actual values over the sample period. The model closely tracks actual production data, suggesting it captures the main dynamics of the cement industry with reasonable precision.

1.7.3 Productivity

This model allows me to recover the productivity level of each plant-year.[7] Figure 1.2 depicts the distribution of TFP across the plants in 2000. To quantify dispersion, I calculate the ratio of 90^{th} to 10^{th} percentile plant TFP, a standard measure in this literature (e.g., Syverson (2004); Hubmer et al. (2024)), which is 1.27, lower than the 1.92 ratio Syverson (2004) documents at the 4-digit NAICS level.[8] Similarly, the top

[7] Productivity of each plant year is calculated as the average productivity of each plant in each year across the number of simulations.

[8] However, there are exceptions. For instance, Hubmer et al. (2024) find a productivity dispersion ratio of 1.36 between firms at the 90^{th} and 10^{th} percentiles. They attribute this relatively low

Table 1.5: Supply Side Estimation Results

Corresponding Equation	Parameters	Estimates
Parameters in the production function (equation 1.3)	α	.806***
		(.112)
	β	.287**
		(.104)
Parameters in the variable inputs' function (equation 1.4)	α_{coal}	.162
		(.285)
	α_{elect}	.0244***
		(.0072)
	α_L	420.5
		(302.5)
Parameters in the Markov process for productivity dynamics (equation 1.5)	σ_ϵ	.005
		(.046)
	η_0	.0322**
		(.0138)
	η_1	.98***
		(.00955)
Parameters in the specification of initial productivity (equation 1.6)	β_1	.0000036
		(.0000571)
	β_2^{dry}	1.486
		(.935)
	$\beta_2^{precalciner}$	1.634*
		(.97)
	$\beta_2^{preheater}$	1.582***
		(.349)
	β_2^{wet}	1.714*
		(1.04)
	β_3	-.00243
		(.0384)
	σ_w	.012
		(.127)

Estimation result of supply-side parameters. The analysis includes 1538 observations.
Standard errors in parentheses
* $p < .1$, ** $p < .05$, *** $p < .01$

Figure 1.1: Predicted vs Actual Cement Production

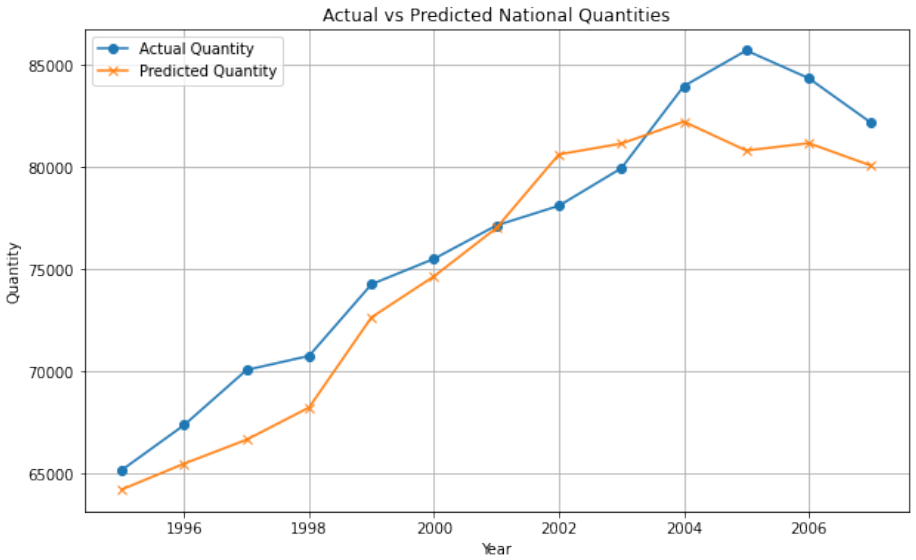
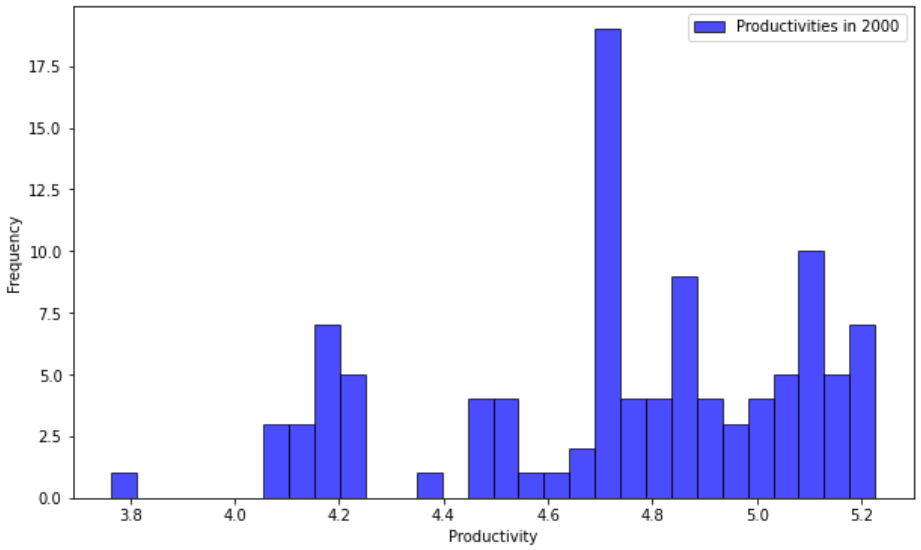


Figure 1.2: Productivity Dispersion in 2000



quartile of plants is about 1.23 times more productive than the bottom quartile. These dispersion measures are lower than what has been found in many other manufacturing industries. Portland cement is a homogenous and standardized product, which allows no room for product differentiation, innovation and quality differences. Therefore, more homogeneity in firm productivity is to be expected.

Next, we analyze the sources of dispersion in productivity. First, the model allows for studying the extent of unobserved heterogeneity in productivity. To do so, I calculate the variance of the logarithm of productivity levels with and without random shocks. Specifically, I simulate two productivity series using the estimated parameters: one including the stochastic terms ϵ_{it} and z_{it} from equations 1.5 and 1.6, and one setting both to zero so that each plant's productivity path is driven only by its observable characteristics (capacity, technology, and age). The results indicate that 95.27% of the variance is explained by observable factors, while only 4.73% is attributable to unobservables—suggesting that unobserved factors play a limited role in explaining productivity dispersion in this industry. While previous studies

dispersion to their flexible, non-parametric specification of the production function, which allows for heterogeneity in returns to scale.

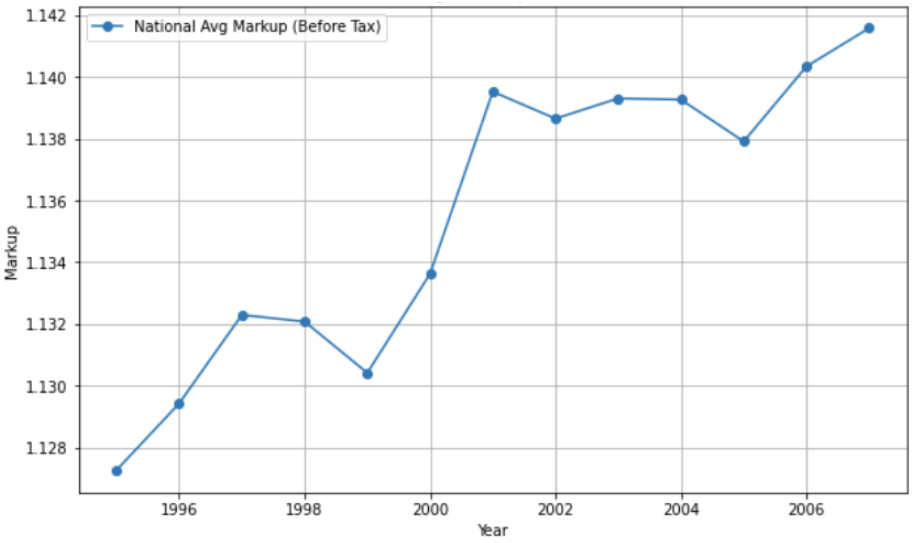
often find a larger role for unobserved productivity, the Portland cement industry has characteristics that help rationalize this result. Cement is a highly standardized commodity, produced under strict guidelines set by the American Society for Testing and Materials (ASTM). As a result, one plant's output is largely substitutable for another's within the same market, leaving little scope for unobserved quality variation. Unlike in many other industries, a cement plant cannot gain a competitive edge through product differentiation or quality enhancements. Moreover, the production process (crushing limestone, heating it in kilns to produce clinker, and then grinding it to cement) is broadly similar across firms, with the main variation coming from kiln technology, which is observed in this study. Therefore, observed factors such as kiln technology, scale, and fuel usage capture most of the relevant productivity differences across plants. Finally, although cement's high transportation costs lead to geographically segmented markets, the product remains a commodity. In regions with multiple plants or access to imports, buyers tend to switch to the lowest-priced source, as the product is virtually identical. This high substitutability on the demand side places strong competitive pressure on inefficient plants, further reducing the role of unobserved heterogeneity in explaining performance differences.

To examine the role of technology choice in driving productivity heterogeneity, I held all other observable and unobservable factors constant and calculated the standard deviation of productivity in this hypothetical case. The resulting standard deviation is 0.340, compared to 0.347 for the estimated productivities. This suggests that nearly all of the observed dispersion can be attributed to differences in technology choice.

1.7.4 Markups

Using the supply-side parameter estimates, I compute marginal costs for each plant and also the price-cost markups. Figure 1.3 shows the evolution of average national

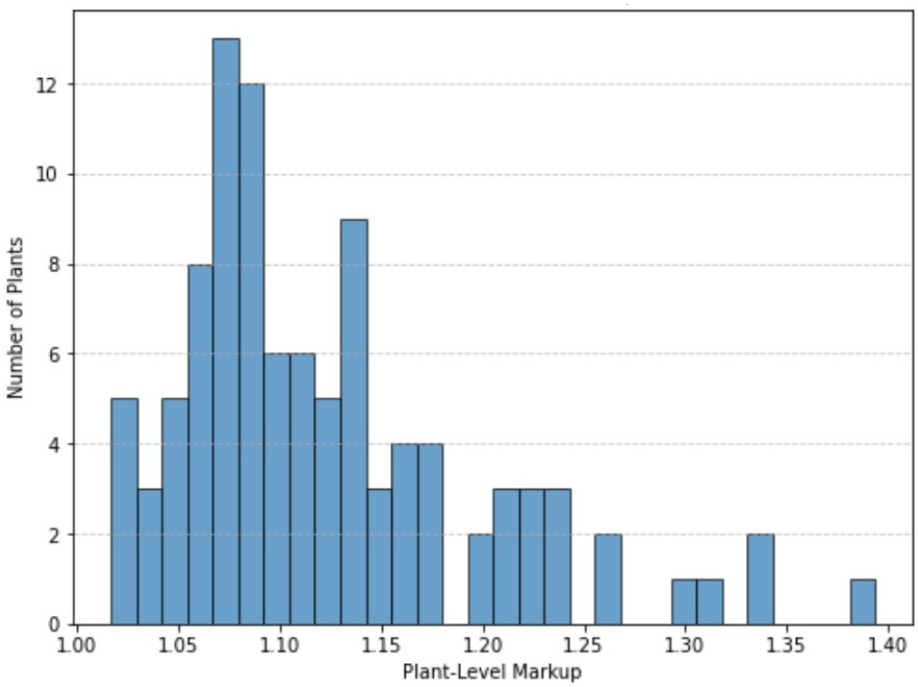
Figure 1.3: Average National Markups



markups in the cement industry from 1995 to 2007, and figure 1.4 shows the distribution of plant-level markups in 2007. The estimated average markup (price over marginal cost minus 1) is around 13%, which is relatively low compared to findings in other industries.[9] One potential reason for moderate markups is the Cournot competition assumption in the model combined with the presence of multiple regional competitors: as the number of rivals in a local market increases, markups are driven down quickly. However, the trend in markups is upward, consistent with existing literature documenting increasing market power over time in the U.S. economy (De Loecker et al. (2020)). This upward trend suggests that market power in the industry grew over the period of analysis, allowing firms to price further above marginal cost by 2007 than they did in the mid-1990s. On the other hand, it is also possible that these rising markups help cover the high fixed costs associated with adopting modern kilns and upgrading technology (Miller et al. (2023)). As plants incur substantial capital expenses to modernize and improve efficiency, higher markups may be needed to sustain profitability in this capital-intensive industry.

[9] Miller et al. (2023) report the average median markups to be 1.23

Figure 1.4: Distribution of Plant-Level Markups (2007)

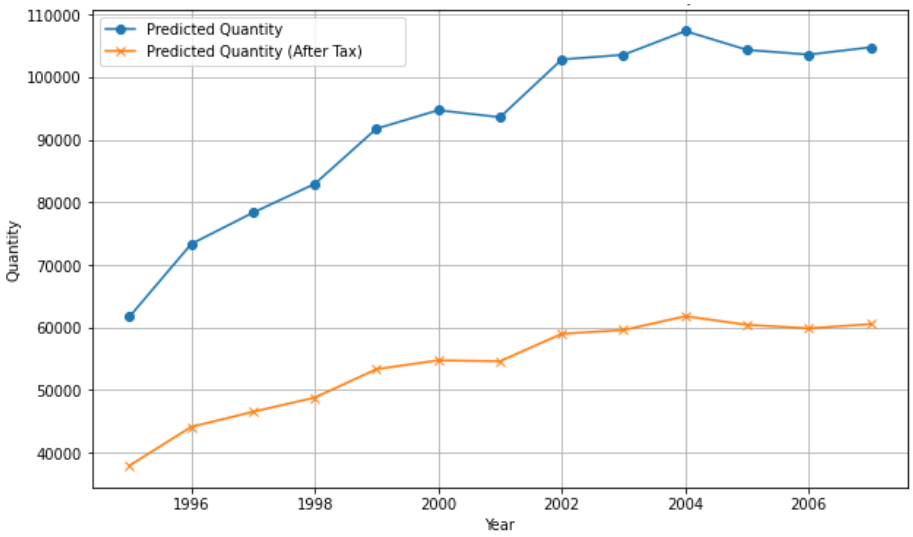


1.8 Counterfactual Analysis

The cement industry accounts for a significant portion of global CO_2 emissions, which makes this industry a key focus for environmental policies aimed at reduction of carbon print. In this section, I run a counterfactual analysis of a carbon tax to examine the welfare effects as well as the reallocation of output among heterogeneous plants. To do so, I use an estimate for the social cost of carbon and derive its implied value per tonne of cement produced.

Estimating the Social Cost of Carbon (SCC) entails considerable uncertainty. The Interagency Working Group on Social Cost of Carbon (2013) report SCC estimates for emissions in 2010 ranging from \$11 to \$89 per tonne of CO_2. For this analysis, I use their moderate estimate of \$32 per tonne (in 2007 dollars). Additionally, Fowlie et al. (2016) report that wet kilns, dry kilns, and dry kilns with precalciner emit approximately 1.16, 0.93, and 0.81 tonnes of CO_2 per tonne of cement, respectively. Using these figures, I calculate the per-unit carbon tax for each plant based on its

Figure 1.5: Predicted National Quantities before and after output tax



technology.

Figure 1.5 shows the predicted national quantities before and after applying the output tax. The model predicts a significant decline in total production—output drops by over 40% under the carbon tax. This reduction in quantity implies an approximate 17% increase in the market price of cement.

Evaluating the welfare consequences of this output tax, I compute the changes in producer and consumer surplus. To calculate the change in the producer surplus, I find the change in the profit of each plant and then sum over all plants. The profit function for plant i before the tax is as follows:

$$\pi_{0i} = p_{0n}q_{0i} - \int_0^{q_0} mc_i(q_0)dq = p_0 q_{0i} - \int_0^{q_0} \frac{1}{\alpha} p_v w_i^{-1/\alpha} K_i^{-\beta/\alpha} q^{\frac{1}{\alpha}-1} dq = p_{0n}q_{0i} - p_v w_i^{-1/\alpha} K_i^{-\beta/\alpha} [q_{0i}]^{1/\alpha}, \tag{1.11}$$

where the subscripts 0 and 1 for p and q indicate the quantity and price before and after tax, respectively. The loss in the total producer surplus is then equal to:

$$\Delta P.S = -\Sigma_{i,n}[(p_{0n}q_{0i} - p_{1n}q_{1i}) - \alpha p_v w_i^{-1/\alpha} K_i^{-\beta/\alpha}(q_{1i}^{1/\alpha} - q_{0i}^{1/\alpha})] \tag{1.12}$$

To calculate the change in the consumer surplus, I use the estimated demand curve for each market and integrate over the total quantities before and after the output tax. The change in consumer surplus in market n is then given by:

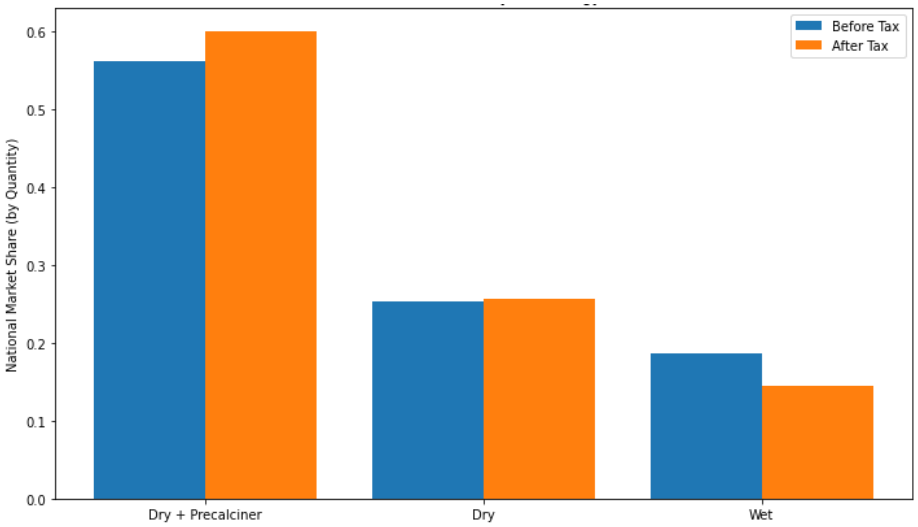
$$\Delta C.S = e^{-A_{0n}/A_1}\left(\frac{A_1}{A_1+1}\right)[Q_{0n}^{(A_1+1)/A_1} - Q_{1n}^{(A_1+1)/A_1}] - (p_{1n}-p_{0n})(Q_{0n}-Q_{1n}) \quad (1.13)$$

I find that the total producer surplus loss is about \$4.14 billion, while consumers suffer a loss of about \$1.49 billion. This suggests that most of the tax burden is borne by producers. Miller et al. (2017) find an opposite pattern empirically, showing pass-through rates above one in the cement industry (i.e., cost increases are mostly passed on to consumers). The difference in results here is likely due to the constant-elasticity demand specification in the model and the relatively high elasticity estimate I obtain.[10] In this model, consumers are somewhat more price-sensitive, limiting the ability of firms to pass through the tax fully.

I also investigate whether the carbon tax causes any reallocation effects among plants with different technologies. Figure 1.6 depicts the market share of plants by technology in the last year of the dataset (2007) before and after the tax. The hypothetical carbon tax indeed punishes plants with more carbon-intensive technology more than others. The total market share of such high-emission (wet-kiln) plants decreases from 18% to 14%, with almost all of that share being taken up by plants with the most carbon-efficient technology (precalciner dry kilns). Moreover, despite having only 14% of the total market share, the wet-kiln firms bear about 26% of the total loss in producer surplus. Nonetheless, given the large value of the carbon tax used in this simulation, the results suggest that the output shift towards more efficient technology, while directionally clear, is not extremely large in magnitude.

[10] In their paper, Miller et al. (2017) note that demand curvature plays an important role in pass-through rate in the industry.

Figure 1.6: Market share by technology before and after tax - 2007



1.9 Conclusion

This paper finds that productivity dispersion in the U.S. Portland cement industry is present but relatively modest. Plants in the top decile of the productivity distribution are only about 25% more productive than those in the bottom decile, a smaller gap than in previous literature. This lower dispersion is attributed to the standardized, homogeneous nature of cement and its production process, which leaves little room for quality variation or innovative differentiation.

After accounting for observable factors—particularly the kiln technology each plant uses—very little of the remaining productivity variation is due to unobserved plant-specific factors. Only about 5% of the logarithm of TFP variance is unexplained by observables, implying that unobserved firm-specific productivity differences play a limited role once technology, scale, and other inputs are considered. The counterfactual carbon tax analysis indicates that environmental regulation would induce a modest reallocation of output toward greener plants, while imposing a heavy burden on carbon-intensive producers. Implementing a flat carbon tax causes the market share of high-emission "wet kiln" plants to fall from 18% to about 14%, with nearly all of that output picked up by plants employing the most carbon-efficient technol-

ogy. Despite holding just 14% of total output, the least green technology plants bear roughly 26% of the total loss in producer surplus under the tax.

Bibliography

Ackerberg, D. A., Caves, K., and Frazer, G. (2015). Identification properties of recent production function estimators. *Econometrica*, 83(6):2411–2451.

Aw, B. Y., Roberts, M. J., and Xu, D. Y. (2011). R&d investment, exporting, and productivity dynamics. *American Economic Review*, 101(4):1312–1344.

Balasubramanian, N. and Sivadasan, J. (2011). What happens when firms patent? new evidence from us economic census data. *The Review of Economics and Statistics*, 93(1):126–146.

Benkard, C. L. (2000). Learning and forgetting: The dynamics of aircraft production. *American Economic Review*, 90(4):1034–1054.

Bloom, N., Dorgan, S., Dowdy, J., and Van Reenen, J. (2007). Management practice and productivity: Why they matter. *Management Matters*, 10.

Bond, S., Hashemi, A., Kaplan, G., and Zoch, P. (2021). Some unpleasant markup arithmetic: Production function elasticities and their estimation from production data. *Journal of Monetary Economics*, 121:1–14.

Carbon (2013). Technical support document:-technical update of the social cost of carbon for regulatory impact analysis-under executive order 12866. *Environmental Protection Agency*.

Chew, W. B., Clark, K. B., and Bresnahan, T. F. (1989). *Measurement, coordination and learning in a multi-plant network.* Division of Research, Harvard Business School.

Collard-Wexler, A. (2011). Productivity dispersion and plant selection in the ready-mix concrete industry. *US Census Bureau Center for Economic Studies Paper No. CES-WP-11-25.*

Collard-Wexler, A. and De Loecker, J. (2015). Reallocation and technology: Evidence from the us steel industry. *American Economic Review*, 105(1):131–171.

De Loecker, J., Eeckhout, J., and Unger, G. (2020). The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics*, 135(2):561–644.

De Loecker, J. and Scott, P. T. (2016). Estimating market power evidence from the us brewing industry. Technical report, National Bureau of Economic Research.

Doraszelski, U. and Jaumandreu, J. (2013). R&d and productivity: Estimating endogenous productivity. *Review of economic studies*, 80(4):1338–1383.

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

Fowlie, M., Reguant, M., and Ryan, S. P. (2016). Market-based emissions regulation and industry dynamics. *Journal of Political Economy*, 124(1):249–302.

Fox, J. T. and Smeets, V. (2011). Does input quality drive measured differences in firm productivity? *International Economic Review*, 52(4):961–989.

Gandhi, A., Navarro, S., and Rivers, D. (2011). On the identification of production functions: How heterogeneous is productivity? Technical report, CIBC Working Paper.

Greenstone, M., List, J. A., and Syverson, C. (2012). The effects of environmental regulation on the competitiveness of us manufacturing. Technical report, National Bureau of Economic Research.

Hsieh, C.-T. and Klenow, P. J. (2009). Misallocation and manufacturing tfp in china and india. *The Quarterly journal of economics*, 124(4):1403–1448.

Hubmer, J., Chan, M., Ozkan, S., Salgado, S., and Hong, G. (2024). *Scalable versus Productive Technologies*. Penn Institute for Economic Research, Department of Economics, University of Pennsylvania.

Ilmakunnas, P., Maliranta, M., and Vainiomäki, J. (2004). The roles of employer and employee characteristics for plant productivity. *Journal of productivity analysis*, 21:249–276.

Lentz, R. and Mortensen, D. T. (2008). An empirical model of growth through product innovation. *Econometrica*, 76(6):1317–1373.

Miller, N. and Osborne, M. (2014). Spatial differentiation and price discrimination in the cement industry: evidence from a structural model. *The RAND Journal of Economics*, 45(2):221–247.

Miller, N., Osborne, M., and Sheu, G. (2017). Pass-through in a concentrated industry: empirical evidence and regulatory implications. *The RAND Journal of Economics*, 48(1):69–93.

Miller, N., Osborne, M., Sheu, G., and Sileo, G. (2023). Technology and market power: The united states cement industry, 1974-2019. *Georgetown McDonough School of Business Research Paper*, (4041168).

Olley, S. and Pakes, A. (1992). The dynamics of productivity in the telecommunications equipment industry.

Restuccia, D. and Rogerson, R. (2008). Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic dynamics*, 11(4):707–720.

Ryan, S. P. (2012). The costs of environmental regulation in a concentrated industry. *Econometrica*, 80(3):1019–1061.

Sakellaris, P. and Wilson, D. J. (2004). Quantifying embodied technological change. *Review of Economic Dynamics*, 7(1):1–26.

Salter, W. E. G. and Reddaway, W. B. (1966). Productivity and technical change. *(No Title)*.

Syverson, C. (2004). Market structure and productivity: A concrete example. *Journal of political Economy*, 112(6):1181–1222.

Thornton, R. A. and Thompson, P. (2001). Learning from experience and learning from others: An exploration of learning and spillovers in wartime shipbuilding. *American Economic Review*, 91(5):1350–1368.

Van Biesebroeck, J. (2003). Productivity dynamics with technology choice: An application to automobile assembly. *The Review of Economic Studies*, 70(1):167–198.

Womack, J. P., Jones, D. T., and Roos, D. (2007). *The machine that changed the world: The story of lean production–Toyota's secret weapon in the global car wars that is now revolutionizing world industry*. Simon and Schuster.

Wooldridge, J. M. (2005). Simple solutions to the initial conditions problem in dynamic, nonlinear panel data models with unobserved heterogeneity. *Journal of applied econometrics*, 20(1):39–54.

1.A Additional Tables

Table 1.6: Demand Curve Estimation Results: Full Region Fixed Effects

Variables	Estimates
Price	-2.334*** (0.480)
Alabama	18.20*** (2.025)
Arizona and New Mexico	18.35*** (2.096)
Arkansas and Oklahoma	17.67*** (2.014)
California, northern	18.04*** (2.047)
California, southern	18.89*** (2.014)
Colorado and Wyoming	17.76*** (2.057)
Florida	18.73*** (2.023)
Georgia, Virginia, West Virginia	17.84*** (2.033)
Idaho, Montana, Nevada, Utah	17.98*** (2.063)
Illinois	17.76*** (2.017)
Indiana	17.56*** (1.976)
Iowa, Nebraska, South Dakota	18.37*** (2.046)
Kansas	17.43*** (2.026)
Kentucky, Mississippi, Tennessee	17.77*** (2.023)
Maine and New York	17.91*** (2.017)
Maryland	17.36*** (1.977)
Michigan and Wisconsin	18.55*** (2.029)
Missouri	18.38*** (2.005)
Ohio	16.81*** (2.026)
Oregon and Washington	17.61*** (2.025)
Pennsylvania, eastern	18.14*** (2.000)
Pennsylvania, western	17.17*** (2.024)
South Carolina	17.69*** (1.999)
Texas, northern	18.45*** (2.017)
Texas, southern	18.33*** (1.987)

Results from estimating Equation 1.7. The analysis includes 325 observations. Alaska is excluded due to insufficient data.

Robust standard errors in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$