The Costs of Environmental Regulation in a Concentrated Industry*

Stephen P. Ryan†

February 20, 2009

Abstract

The typical cost analysis of an environmental regulation consists of an engineering estimate of the compliance costs. In industries where fixed costs are an important determinant of market structure this static analysis ignores the dynamic effects of the regulation on entry, investment, and market power. I evaluate the welfare costs of the 1990 Amendments to the Clean Air Act on the US Portland cement industry, accounting for these effects through a dynamic model of oligopoly in the tradition of Ericson and Pakes (1995). Using the recently developed two-step estimator of Bajari, Benkard, and Levin (2007a), I recover the entire cost structure of the industry, including the distributions of sunk entry costs and capacity adjustment costs. My primary finding is that the Amendments have significantly increased the sunk cost of entry. I solve for the Markov perfect Nash equilibrium (MPNE) of the model and simulate the welfare costs of the Amendments. A static analysis misses the welfare penalty on consumers, and may even obtain the wrong sign of the welfare effects on incumbent firms. I also offer a novel constructive approach for the identification of parameters in a broad class of two-step estimators, including distributions of private information in dynamic games.

*This is a revised version of my job market paper. I would like to especially thank Pat Bajari for guidance and support. I have also benefited from conversations with Tom Ahn, Arie Beresteanu, Jane Cooley, Paul Ellickson, Han Hong, Shanjun Li, Chris Timmins, Justin Trogdon, and numerous seminar participants. All remaining errors are my own.

†MIT Department of Economics and NBER.

1 Introduction

In the United States, the Environmental Protection Agency (EPA) is responsible for setting and enforcing regulations broadly consistent with national environmental policies, such as the Clean Air Act (CAA). The CAA gives the EPA a mandate to regulate the emissions of airborne pollutants such as ozone, sulfur dioxide (SO_2), and nitrogen oxides (NO_x), in the hopes of producing a healthier atmosphere. The Clean Air Act and its subsequent Amendments require the Agency to assess the costs and benefits of a regulation before promulgating policy. The cost analysis is typically an engineering estimate of the expenditures on control and monitoring equipment necessary to bring a plant into compliance with the new regulations. However, this type of cost analysis misses most of the relevant economic costs in concentrated industries, in which sunk costs of entry and costly investment are important determinants of market structure. Shifts in the costs of entry and investment can lead to markets with fewer firms and lower production. The resulting increase in market concentration can have far-reaching welfare costs beyond the initial costs of compliance. This is a particularly acute problem for environmental regulators, as many of the largest polluting industries are also highly concentrated.[1]

In this paper, I measure the welfare costs of the 1990 Clean Air Act Amendments on the US Portland cement industry, explicitly accounting for the dynamic effects resulting from a change in the cost structure. Portland cement is the binding material in concrete, a primary construction material found in numerous applications, such as buildings and highways. The industry is typical of many heavy industries, consuming large quantities of raw materials and generating significant amounts of pollution byproducts. It is a frequent target of environmental activists and has been heavily regulated under the Clean Air Act. In 1990, Congress passed Amendments to the Clean Air Act, adding new categories of regulated emissions and requiring plants to undergo an environmental certification process. It has been the most comprehensive and important new environmental regulation affecting this industry in the last three decades since the original Clean Air Act.

My strategy for evaluating the effects of the Amendments on this industry proceeds in three distinct steps. First, I pose a theoretical model of the cement industry, where oligopolists make optimal decisions over entry, exit, production, and investment given the

[1] For example, the 1997 Economic Census reports that the Herfindahl-Hirschman Index (HHI) for many polluting industries exceeds 1,000, such as manufacturers of paper pulp, petrochemicals, soaps and detergents, tires, ceramic tiles, lime and gypsum, aluminum, and copper, among others. For comparison, the HHI for Portland cement, the industry studied in this paper, is 466. This national measure understates the effective degree of concentration since the industry is spatially segregated into regional markets.

strategies of their competitors. Second, using a unique panel data set covering two decades of the Portland cement industry, I recover parameters which are consistent with the underlying model. Third, I use the theoretical model to simulate economic environments with the cost structures recovered before and after the Amendments. I exploit a specific timing feature of the implementation of the Amendments to identify which changes in the cost structure were due to the regulations. By comparing the predictions of the model under these different cost structures, I can calculate the changes to a number of relevant policy quantities, such as producer profits and consumer surplus, that are the result of the regulation.

The backbone of my analysis is a fully dynamic model of oligopoly in the tradition of Maskin and Tirole (1988) and Ericson and Pakes (1995). I model the interaction of firms in spatially-segregated regional markets where firms are differentiated by production capacity. Firms are capacity constrained and compete over quantities in homogeneous good markets. Markets evolve as firms enter, exit, and adjust their capacities in response to variation in the economic environment. I incorporate sunk costs of entry, fixed and variable costs of capacity adjustment, and a fixed cost of exiting the industry. I assume that firms optimize their behavior conditional only on the current state vector, which results in a Markov-perfect Nash equilibrium (MPNE).

My model is similar to several other applications of the Ericson-Pakes model.[2] However, I extend the model in several ways to tailor it to the Portland cement industry. First, I allow firms to fully adjust their capacity in each period, whereas previous models have looked at investment games where capital accumulates slowly. Additionally, I introduce fixed costs of adjustment to rationalize the lumpy investment behavior seen in the data, as firms tend to make very infrequent but large capacity adjustments. Second, firms have private information about their fixed entry and exit costs. Third, I allow for multiple entry and exit in every period. Furthermore, potential entrants are not restricted to begin operation at an exogenously-imposed capacity level, but rather choose an optimal starting level given expectations about future market conditions.

The MPNE of the model leads to structural requirements on firm behavior which can be used as the basis of an estimator of the underlying primitives. As Benkard (2004) illustrates, the impediment to using these types of models for empirical work has been the computational burden of solving for the MPNE, which makes nested fixed-point estimators in the tradition of Rust (1987) impractical. However, a series of recent papers has built on the insights

[2]See, for example, Fershtman and Pakes (2000), Gowrisankaran and Town (1997), Besanko and Doraszelski (2004), Doraszelski and Satterthwaite (2005), and Benkard (2004).

of Hotz, Miller, Sanders, and Smith (1994) to circumvent this problem using a two-step approach, in which it is possible to estimate the dynamic model without solving for the equilibrium even once. Representative papers in this literature include Bajari, Benkard, and Levin (2007a), Aguirregabiria and Mira (2004), Pakes, Ostrovsky, and Berry (2005), and Pesendorfer and Schmidt-Dengler (2003).

In these two-step estimators, the econometrician first simply describes *what* the firms do at every state, and then imposes equilibrium restrictions from an underlying model to explain explain *why* the firms behave as they do. In my application, the first step includes flexibly recovering the demand curve, production costs, and the policy functions governing entry, exit, and investment. These reduced-form policy functions describe what actions the firm will undertake given any state vector. The key to understanding the estimator is that these observed policy functions have to be optimal given the underlying theoretical model. Therefore, in the second step I find the remaining unknown parameters that best rationalize the observed policies as the equilibrium outcomes of profit-maximizing firms. Following the simulation-based minimum distance estimator proposed by Bajari, Benkard, and Levin (2007a), I recover the both distribution of fixed adjustment costs and variable adjustment costs, the distribution of scrap values associated with exiting the market, and the distribution of entry costs. I recover these parameters before and after the 1990 Amendments in order to evaluate the changes in the underlying cost structure induced by the Amendments. As a testament to the flexibility and power of this approach, and the relatively clean institutional details of the cement industry, I am able to recover estimates of the cost of investment that are very close to accounting data from out-of-sample sources cited in Salvo (2005).

In section 6, I provide a novel constructive method of the identification of these parameters. To my knowledge, these results are new to the literature, and apply to a broad class of dynamic models estimated using a two-step approach. The identification results build constructively on the Euler equation underlying the MPNE concept in two-step dynamic estimators, and demonstrates identification through a linearization of the unknown parameters. I also demonstrate how two-step estimators can recover the parameters of distributions of private information in the underlying dynamic model, which generalizes the set of models considered in the two-step literature to date.

After recovering estimates for the underlying model primitives, I numerically solve for the MPNE of the theoretical model. This step provides policy functions which are used for the simulation of the welfare costs of the Amendments. I evaluate expected producer and consumer welfare, the number and size of firms, and the distribution of costs across incumbents and potential entrants before and after the regulations. I abstract away from any benefits accruing to consumers due to reduced emissions, given the difficultly of quantifying the amount of emissions from cement plants and their associated damages.[3] In the baseline case of entry into a new market, I find that overall welfare has decreased at least $1.4B as a result of the Amendments, due to an increase in the average sunk cost of entry. More importantly, as my estimates show the costs of production have not statistically changed after the regulations, the welfare effect on producers depends critically on whether or not the firm is an incumbent. While potential entrants suffer welfare losses as the result of paying higher entry costs, incumbent firms benefit from increased market power due to reduced competition. A static analysis of this industry would preclude changes in barriers to entry, and may obtain the wrong sign for the welfare costs of the Amendments on incumbent firms.

I conclude by comparing the MPNE generated by an oligopoly to the social planner's problem. The social planner pursues policies that maximize the expected sum of both consumer and producer welfare, net of the benefits associated with changes in the level of emissions, so it is the natural baseline for evaluating the efficiency of the observed equilibrium. By this metric, the social planner's solution highlights that the oligopoly markets considered here have socially-inefficient low aggregate capacities. This is due to the fact that oligopolists fail to internalize the full welfare benefits of their investments. As a result, overall market capacity is larger under the social planner in both new and existing markets. In both cases, producers suffer moderate profit losses while consumers enjoy a three- to five-fold increase in surplus which more than offsets those losses.

This paper makes several contributions. First, I recover the entire cost structure of an industry, including the sunk costs of entry and exit, production costs, and investment costs. I demonstrate the estimation and identification of distributions of private information in the context of dynamic games. The recovery of these parameters allows the counterfactual simulation of a regulation's welfare effects in the presence of dynamics and market power for the first time.[4] These welfare cost estimates allow me to determine a lower bound of the value of clean air if the Amendments are to be efficient. One of the most important implications

[3]This continues to be an area of active research, as the inclusion of cement plants into broader emissions markets post-2000 presents the opportunity to obtain better pollution estimates for at least a subset of plants.

[4]Mansur (2004) also examines the regulation of an industry with market power but is concerned with the effect of concentration on the quantity of pollution emissions. Benkard (2004) applies many of the ideas formalized in the BBL estimator in his examination of the wide-body aircraft industry but does not recover estimates of fixed costs.

of my findings is that static engineering estimates of compliance costs miss most of the economic penalties associated with a regulation when there are significant sunk entry costs. The static analysis misses the penalty to consumer welfare due to lower production, and obtains a welfare cost of the wrong sign for incumbent producers. I also make a contribution to the investment literature by applying a generalized (S, s) model, in the tradition of Scarf (1959) and built on the econometric model of Attanasio (2000), to a dynamic investment game.[5] My results suggest that both fixed and variable adjustment costs are an important determinant of investment behavior. I highlight the importance that sunk costs of entry have on industry structure and evolution, since they are a primary determinant of market structure in this industry. Finally, the identification results presented here are a novel approach to identification in dynamic games, are easy to verify, and are applicable to a wide range of dynamic models.

The paper is organized as follows: I give a brief overview of the Portland cement industry and relevant environmental regulations over the last 30 years in Section 2. I discuss the sources of the data and introduce the key variables of the model and estimation in Section 3. Section 4 introduces the theoretical model underpinning the estimation detailed in Section 5. I discuss the results in Section 7 and present the results of the counterfactual simulations in Section 8. Section 9 concludes with a summary of my results and a discussion of possible extensions.

2 Portland Cement Industry

Portland cement is a fine mineral dust with useful binding properties that make it the key ingredient of concrete. Water and cement form a paste that binds particulates like sand and stone together and makes a pourable material that hardens over time. The concrete is then used as a fill material, such as in highways and buildings, and in finished products like concrete blocks.

Producing cement requires two commodities in enormous quantities: limestone and heat. The limestone is usually obtained from a quarry located at the production site. Large chunks of limestone are pulverized before being sent to the centerpiece of cement operations: an enormous rotating kiln furnace. These kilns are the largest moving piece of industrial equipment in the world; they range in length from 450 to 1000 feet and have diameters of

[5] Hall and Rust (2000) apply a similar framework for modeling automobile purchase decisions and inventory control, respectively.

over 15 feet. The chemical process of converting limestone into cement requires temperatures equal to a third of those found on the surface of the sun, so one end of the kiln is heated with an intense flame produced by burning fossil fuels. The large scale nature of these installations is reflected in their raw materials and energy requirements: a large kiln can process up to 200 tons of limestone per hour, and cement kilns are the third largest consumer of energy in the world. These high energy requirements are what lead the cement industry, a tiny part of the US economy at under $10B a year in revenues, to have a large role in the environmental debate over emissions. Furthermore, the chemistry of the production of cement liberates carbon dioxide as a byproduct, which means that the production of cement is a major contributor of greenhouse gases globally.

Cement is a difficult commodity to store, as it will gradually absorb water out of the air, rendering it useless. As a result, producers and distributors do not maintain large stocks. Also, I treat cement as a homogeneous good since producers in the United States adhere to the American Society for Testing and Materials Specification for Portland Cement. Cement's use as a construction material means that producers are held to strict conformity with these specifications.

As a result of cement's tendency to spoil in storage, transportation costs are the most significant factors in determining Portland cement markets. Average transportation costs reported by U.S. producers for shipments within 50 miles of the plant were $5.79 per ton. These costs increased to $9.86 per ton for shipments within 51-100 miles, $14.53 per ton for 101-200 miles, and to $18.86 per ton for 201-300 miles. For shipments that are 500 miles or more from the plant, transportation costs increased to $25.85 per ton.[6] These high costs, in conjunction with cement's low unit value, are the principal reasons the majority of cement is shipped locally. Jans and Rosenbaum (1997) quote a Census of Transportation report stating that 82.5 percent of cement was shipped under 200 miles, with 99.8 percent being shipped under 500 miles.

In 2000, the domestic Portland cement industry consisted of 116 plants in 37 states, run by one government agency and approximately 40 firms. The industry produced 86 million tons of Portland cement with a raw value of approximately $8.7 billion; most of this was used to make concrete, with a final value greater than $35 billion. Domestic cement production accounted for the vast majority of the cement used in the United States. According to the USGS (2001), about 73 percent of cement sales were to ready-mixed concrete manufacturers,

[6] These figures are taken from American University's Trade and Environment Database (TED) case study on Cemex.

Table 1: Cement Industry Summary Statistics

Year	Production	Imports	Consumption	Price	Capacity	Capacity Per Kiln
1980	68,242	3,035	70,173	111.90	89,561	239
1981	65,054	2,514	66,092	103.70	93,203	267
1982	57,475	2,231	59,572	95.76	89,770	287
1983	63,884	2,960	65,838	91.01	92,052	292
1984	70,488	6,016	76,186	89.70	91,048	297
1985	70,665	8,939	78,836	84.71	88,600	305
1986	71,473	11,201	82,837	81.48	87,341	305
1987	70,940	12,753	84,204	78.07	86,709	314
1988	69,733	14,124	83,851	75.50	86,959	327
1989	70,025	12,697	82,414	72.04	84,515	337
1990	69,954	10,344	80,964	69.02	83,955	345
1991	66,755	6,548	71,800	66.37	84,471	352
1992	69,585	4,582	76,169	64.25	85,079	357
1993	73,807	5,532	79,701	63.58	84,869	363
1994	77,948	9,074	86,476	68.06	85,345	364
1995	76,906	10,969	86,003	72.56	86,285	367
1996	79,266	11,565	90,355	73.64	85,687	376
1997	82,582	14,523	96,018	74.60	86,465	383
1998	83,931	19,878	103,457	76.45	87,763	393

Summary statistics for the Portland cement industry 1980-1998. The data is from Historical Statistics for Mineral and Materials Commodities in the United States, an online publication of the US Geological Survey. The units on quantities are thousands of metric tons, while prices are denoted in 1998 constant dollars.

12 percent to concrete product producers, 8 percent to contractors, 5 percent to building materials dealers, and 2 percent for other uses. Cement expenditures in construction projects are usually on the order of less than 2 percent of total outlays.

Table 1 reports summary statistics for the industry over the period 1980-1998. One point of interest is that capacity utilization rates have risen since the passage of the Amendments. Production has increased while overall productive capacity has remained relatively steady. Imports grew as the production of domestic cement reached its maximum level, and firms chose to import instead of build new production facilities.[7] The industry has become slightly more concentrated over time. According to the Economic Census (1977–2002), which collects extensive information on American industry every five years, the national four-firm concentration ratio was 38.7 in 2002, 33.5 in 1997, 28 in 1992, 28 in 1987, 31 in 1982, and 24 in 1977. However, these numbers mask the regional variation in cement concentration, as national market share may not be representative of competitive conditions in any given geographic region, due to the local nature of the cement industry.

[7]Cement imports come primarily from Canada, China, Korea, Thailand, Spain, and Venezuela. Asian sources have become the dominant source of cement imports, with Thailand becoming the single-largest exporter in 2000.

The effects of imports on domestic producers are difficult to quantify due to the idiosyncracies associated with distributing cement from waterborne sources. For most markets, the economic impact is small and indirect, as few regions have the infrastructure and geography to profitably exploit the availability of imports. An examination of the import data provided in the USGS reports indicates that cement imports vary widely across markets and across time. Imported cement is actually shipped as clinker, the unground precursor of cement. In order to turn this raw material into cement, the importer must have a grinder and a supply of gypsum. Additionally, domestic cement producers have been highly successful in preventing large-scale imports through trade tariffs. For example, producers in states bordering the Gulf of Mexico have been successful in getting anti-dumping tariffs passed against imports from Mexico. This has limited the ability of importers to achieve greater penetration of local cement markets in these states. In large part, the response of potential importers has been to circumvent the tariffs through the acquisition of domestic facilities. In markets where imports do play a significant long-run role in the domestic market, such as around the Great Lakes region, I model this as a permanent shifter in the demand curve for domestically-produced cement.[8]

There have been two major regulatory events of interest to the Portland cement industry in the last 30 years: the Clean Air Act of 1970 and its subsequent Amendments in 1990.[9] The stated purpose of the Clean Air Act was to "protect and enhance the quality of the Nation's air resources so as to promote the public health and welfare and productive capacity of its population." To this end, Congress empowered the EPA to set and enforce environmental regulations governing the emission of airborne pollutants.

In 1990, Congress passed the Amendments to the Clean Air Act, which defined new categories of regulated pollutants and required major polluters to obtain a permit for operation. These Amendments mandated new monitoring, reporting, and emission requirements for the cement industry. The Amendments created a new class of emission restrictions governing

[8] I am exploring the interactions of environmental regulation and cement imports in ongoing work. The potential for emissions leakage, e.g. Fowlie (2008), to undo domestic regulations is an important question the environmental literature, and could mitigate any benefits from the Amendments.

[9] There have been other substantial changes to environment policy during the time period as well. One such change was the New Source Review (NSR) requirements instituted in the 1977 Amendments to the Clean Air Act. The NSR requires firms to obtain costly permits before substantially modifying older capital equipment or building new plants, and is believed by the EPA (2001) to have changed the pattern of investment in other industries, such as power plants and refineries. The NSR also likely created barriers to entry in this industry, as new entrants may have faced higher capital costs than incumbents, and plants located in lower pollution areas may face different emissions requirements than those in higher pollution areas. Since the NSR requirements did not change over the sample period that I am examining, I do not explicitly model the differential effects of the NSR on firms in the cement industry.

hazardous air pollutants and volatile organic compounds. One key identifying feature of this legislation is that EPA did not promulgate final requirements for these new pollutants for 12 years. Therefore, there were no changes to firms' variable costs as a result of the Amendments, as they did not require the firms to adhere to any new emissions standards.[10]

There were two components of the legislation that began to bind immediately. Under Title V of the Amendments, all firms emitting significant quantities of pollutants had to apply for operating permits. The permits require regular reporting on emissions, which necessitate the installation and maintenance of new monitoring equipment. The Amendments also required firms to draw up formal plans for compliance and undergo certification testing. Industry estimates for the costs of compliance with these operating permits is on the order of five to ten million dollars. By 1996, virtually all cement plants had applied for their permits, which they are required to renew every five years. The EPA estimated that these certification costs would not exceed $5M per establishment.

The second aspect of the Amendments which is critical to understanding their welfare implications is that they required greenfield plants to undergo an additional, rigorous environmental certification and testing procedure. These additional fixed costs involved potential entrants contracting with environmental engineering firms to produce reports on their impact on local air and water quality as a result of the construction and operation of a new plant. Industry sources estimate that these costs would add approximately $5M to $10M to the cost of building a greenfield facility. It is this change to the sunk costs of entry which is going to drive many of the results that I find below.

3 Data

I collect data on the Portland cement industry from 1980 to 1999 using a number of different sources. I require market-level data on prices and quantities to estimate the demand curve for cement. The US Geological Survey (USGS) collects establishment-level data for all the Portland cement producers in the US and publishes the results in their annual Minerals Yearbook.[11] The USGS aggregates establishment-level data into regional markets to protect

[10]To the best of my knowledge, as of 2009 no firm has made any changes to its production process as a result of the Amendments, due in part to legal opposition from the Portland Cement Association. Firms may also reasonably anticipate that changes to their marginal costs may ultimately be close to zero, as either they will be grandfathered into the legislation or the EPA may give pollution credits in return for adopting lower emissions standards.

[11]The Bureau of Mines had this responsibility prior to merging with the USGS in the 1990s. The data was collected by a mail survey, with a telephone follow-up to non-respondents. Typically the total coverage of

Table 2: Summary Statistics

Variable	Minimum	Mean	Maximum	Standard Deviation
Market-level Supply and Demand Data				
Quantity	186	2,835.84	10,262	1,565.34
Price	36.68	67.46	138.99	13.68
Plants In Market	1	4.75	20	1.94
Skilled Wage	20.14	31.72	44.34	4.33
Coal Price	15.88	26.64	42.33	8.13
Electricity Price	4.23	5.68	7.6	1.01
Natural Gas Price	3.7	6.21	24.3	2.21
Population	689,584	10,224,352	33,145,121	7,416,485
Plant-level Production Data				
Quantity	177	699	2348	335
Capacity	196	797	2678	386
Plant-level Investment				
Capacity Investment	-728	2.19	1,140	77.60

Demand data are from annual volumes of the USGS's Mineral Yearbook, 1980-1981 to 1998-1999. There are 517 observations in 27 regional markets. The unit for market quantity is thousands of tons per year, while price is denoted in dollars per ton. Labor wages are denoted in dollars per hour for skilled manufacturing workers, and taken from County Business Patterns. Population is the total populations of the states covered by a regional market. The units are dollars per ton for coal, dollars per kilowatt hour for electricity, and dollars per thousand cubic feet for gas. All prices are adjusted to 1996 constant dollars. The data on production and capacity are taken from the Portland Cement Association's annual Plant Information Summary, with full coverage from 1980 to 1999. Units on quantity and capacity are in thousands of tons per year.

the confidentiality of the respondents. The Minerals Yearbook contains the number of plants in each market and the quantity and prices of shipped cement. There is occasional irregular censoring of data to ensure the confidentiality of individual companies, although this affects only a small number of observations representing a low percentage of overall quantity. Usually the USGS merges a censored region into a larger region in subsequent years to facilitate complete reporting.

This data is the source of market definitions that I use through the remainder of the analysis. The USGS examines the set of firms which compete with one another, and aggregates their locations together into a market. While this is clearly an imperfect measure of market definition, since prices in neighboring markets almost surely have influence on the prices within a given market, it is roughly consistent with the ideas that cement is very expensive to ship long distances. If one were to draw 100 mile circles around each cement plant in the United States, the resulting areas of significant overlap look much like the market definitions from the USGS data.

the industry exceeded 90 percent; in some years, 100 percent response was indicated. The USGS attempted to fill in missing observations with data from other sources.

I collect data on electricity prices, coal prices, natural gas prices, and manufacturing wages to use as instruments in the demand curve estimation. The data for fuel and electricity prices are from the US Department of Energy's Energy Information Administration.[12] Natural gas and electricity prices are reported at the state level from 1981 to 1999. Coal prices are only available in a full series over that time span at the national average level. I impute skilled manufacturing wages at the state level from the US Census Bureau's County Business Patterns. All prices are adjusted to 1996 constant dollars.

Table 2 shows summary statistics for the demand data. Most markets are characterized by a small number of firms, with the median market contested by four firms. The size of the markets varies greatly across the sample: the smallest market is two percent of the size of the largest market. Price also varies substantially across markets, with Alaska and Hawaii generally being the most expensive markets.

Data on the plant-level capacities and production quantities are from the Portland Cement Association's annual Plant Information Summary (PIS) and cover 1980 to 1998. These trade association databooks have complete coverage of all cement producers in the United States, and give detailed information on grinding and kiln capacity. For each establishment, the PIS reports daily and annual plant capacities. I interpret the daily capacity to be a boilerplate rating, determined by the manufacturer of the kiln at the time of its manufacture, of how much the kiln produces in a given 24-hour period of operation. A critical assumption that I make is that I interpret the number listed under yearly capacity as representing how much cement that plant actually produced in that year. This assumption is supported by the fact that plants operate continuously in runs lasting most of the year except for a maintenance period, generally a month in duration, in which the plant produces nothing. If the firms are assumed to run at perfect efficiency on the days that they operate, then the boilerplate rating multiplied by the length of a year gives the theoretical maximum that a plant could have produced. These boilerplate ratings typically do not change from year to year. On the other hand, the yearly capacity numbers never achieve this bound and fluctuate from year to year. Additionally, the yearly numbers approximate the market-level quantities reported in the USGS data, which was collected through a confidential survey of cement manufacturers. Therefore, I interpret the reported annual capacity of the kiln to be the amount of cement that it actually produced in that year.

I emphasize, however, that production quantity is not exogenously set as a fixed percentage of the theoretical maximum capacity, as firms still choose how long to operate their

[12] http://www.eia.doe.gov.

Figure 1: Capacity of Cement Plants in Colorado and Wyoming

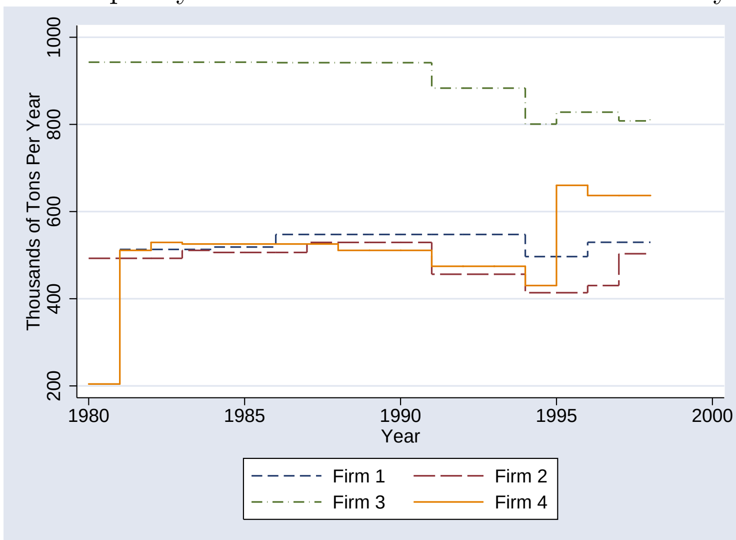

kilns before performing maintenance. Given that firms are at the edge of their maximum productive capacity during the sample period, capacity choice is clearly the most important strategic decision firms have to make, but it should be emphasized that they still face a trade-off between production and maintenance. The last two rows of Table 2 give the summary statistics for production and capacity levels.

A key empirical fact of this industry is that most firms do not make adjustments to their capacity in most periods. The modal adjustment is zero, with a mean of just 2.9 thousand tons per year (TPY). This lumpy adjustment behavior is illustrated in Figure 1, which tracks the capacity levels of firms in the Colorado and Wyoming market over the course of the sample period. While there is some noise in the data, it is clear that most firms have relatively steady levels of capacity over time, with infrequent discrete adjustments. In addition to capacity investment, there are jumps in market-level capacity due to entry and exit.

To match the market-level demand data to the establishment data from the PIS, I combine some of the markets in the USGS data to form continuously-reported metamarkets. I then group all the plants into the appropriate metamarkets for every year of establishment data. The production data consists of an unbalanced panel of 2,233 observations.

4 Model

I construct a theoretical model that captures the salient features of the cement industry to quantify the effects of the Amendments on industry structure. The industry is characterized by simultaneous entry, exit, investment, and production decisions of a small number of firms in regional markets. The firms behave strategically and anticipate the future when making decisions. The structure within each market is completely determined by the distribution of production capacities among active firms. My models builds on Maskin and Tirole (1988) and Ericson and Pakes (1995), who provide an elegant theoretical framework of industry dynamics that can account for these features.

4.1 State Space

The basic building block of the model is a regional market. I assume that each firm operates independently across markets.[13] Each market is fully described by the $\overline{N} \times 1$ state vector, s_t, where s_{it} is the capacity of the i-th firm at time t, and $\overline{N}$ is the exogenously-imposed maximal number of active firms. Firms with zero capacity are considered to be potential entrants. Time is discrete and unbounded. Firms discount the future at rate β.

4.2 Timing

Each decision period is one year. In each period, the sequence of events unfolds as follows:

- Potential entrants receive a private draw from the distribution of entry costs. Incumbents receive private draws on the fixed cost of investment/divestment and their scrap value for exiting.
- All firms simultaneously make entry, exit, and investment decisions.
- Incumbent firms compete over quantities in the product market.
- Firms enter and exit, and investments mature.

I assume that firms who decide to exit produce in this period before leaving the market, and that adjustments in capacity take one period to realize.

[13] This assumption explicitly rules out more general behavior, such as multimarket contact as considered in Bernheim and Whinston (1990) and Jans and Rosenbaum (1997).

4.3 Payoffs

The payoff to each firm is a function of the current state space, s, and actions, a, undertaken in that period, $\pi(s, a)$. Firms obtain revenues from the product market and incur costs from production, entry, exit, and investment.

4.3.1 Product Market Payoffs

Firms compete in quantities in a homogeneous goods product market. I assume within a given market firms face a constant elasticity of demand curve:

$$\ln Q(\alpha) = \alpha_0 + \alpha_1 \ln P, \tag{1}$$

where Q is the aggregate market quantity and α_1 is the elasticity of demand.

Production costs are given by the following function:

$$C_i(q_i; \delta) = \delta_0 + \delta_1 q_i + \delta_2 1(q_i > \nu s_i)(q_i - \nu s_i)^2. \tag{2}$$

Fixed costs of production are given by δ_0. Variable production costs consist of two parts: a constant marginal cost, δ_1, and an increasing function that binds as quantity approaches the capacity constraint. I assume that costs increase as the square of the percentage of capacity utilization, and parameterize both the penalty, δ_2, and the threshold at which the costs bind, ν. This second term, which gives the cost function a "hockey stick" shape common in the electricity generation industry, accounts for the increasing costs associated with operating near maximum capacity, as firms have to cut into maintenance time in order to expand production beyond utilization level ν.

Each firm maximizes their static profits given the outputs of the competitors:

$$\max_{q_i} P\left(q_i + \sum_{j \neq i} q_j; \alpha\right) q_i - C_i(q_i; \delta), \tag{3}$$

where $P(Q; \alpha)$ is the inverse of Equation 1. In the presence of fixed operation costs the product market may have multiple equilibria, as some firms may prefer to not operate given the outputs of their competitors. However, if all firms produce positive quantities then the equilibrium vector of production is unique, as the best-response curves are downward-sloping.

4.3.2 Investment Costs

Firms can change their capacity through costly adjustments. The cost function associated with these activities is given by:

$$\Gamma(x_i;\gamma) = 1(x_i > 0)(\gamma_{i1} + \gamma_2 x_i + \gamma_3 x_i^2) + 1(x_i < 0)(\gamma_{i4} + \gamma_5 x_i + \gamma_6 x_i^2) \tag{4}$$

Firms face both fixed and variable adjustment costs that vary separately for positive and negative changes. Fixed costs capture the idea that firms may have to face significant setup costs, such as obtaining permits or constructing support facilities, that accrue regardless of the size of the kiln. These fixed costs are drawn each period from the common distribution F_γ and are private information to the firm. Divestment sunk costs may be positive as the firm may encounter costs in order to shut down the kiln and dispose of related materials and components. On the other hand, firms may receive revenues from selling off their infrastructure, either directly to other firms or as scrap metal.[14] These costs are also private information, and are drawn each period from the common distribution G_γ.

4.3.3 Fixed Costs of Operation, Entry, and Exit

Firms face fixed costs unrelated to production, given by $\Phi_i(a)$, which vary depending on their current status and chosen action, a_i:

$$\Phi_i(a_i;\kappa_i,\tau,\phi_i) = \begin{cases} \kappa_i & \text{if the firm is a new entrant,} \\ \tau & \text{if the firm continues as an incumbent,} \\ \phi_i & \text{if the firm exits the market.} \end{cases} \tag{5}$$

Continuing incumbents must pay an operating cost of τ. Firms that enter the market pay a fixed cost of entry, κ_i, which is private information and drawn from the common distribution of entry costs, F_κ. Firms exiting the market receive a payment of ϕ_i, which represents net proceeds from shuttering a plant, such as selling off the land or paying for an environmental cleanup. The scrap value is private information, drawn anew each period from the common distribution, F_ϕ. Denote the activation status of the firm in the next period as χ_i, where $\chi_i = 1$ if the firm will be active next period, whether as a new entrant or a continuing

[14] One online example of a used market for cement equipment is www.usedcementequipment.com. While the prices used equipment may be low, or even nominally zero, transportation and cleanup costs are typically high, occasionally into the millions of dollars depending on the size and type of equipment.

incumbent, and $\chi_i = 0$ otherwise.

Collecting the costs and revenues from a firm's various activities, the per-period payoff function is:

$$\pi_i(s, a; \alpha, \delta, \gamma_i, \kappa_i, \tau, \phi_i) = P(Q; \alpha)q_i - C(q_i; \delta) - \Gamma(x_i; \gamma_i) + \Phi_i(a_i; \kappa_i, \tau, \phi_i). \tag{6}$$

4.4 Transitions Between States

To close the model it is necessary to specify how transitions occur between states as firms engage in investment, entry, and exit. I make two assumptions governing these transitions.

Assumption 1. *Changes to the state vector through entry, exit, and investment take one period to occur.*

This is a standard assumption in discrete time models, and is intended to capture the idea that it takes time to make changes to physical infrastructure of a cement plant.

Assumption 2. *Investment is deterministic: a firm's capacity vector is always equal to last period's capacity plus (minus) last period's investment (divestment).*

This assumption abstracts away from depreciation, which does not appear to be a significant concern in the cement industry, and uncertainty in the time to build new capacity. While it is conceptually straightforward to add uncertainty over time-to-build in the model, assuming deterministic transitions greatly reduces the computational complexity of solving for the model's equilibrium.

It is worth noting that, although the transitions are deterministic conditional on the vector of actions, firms may not know which actions its rivals will undertake with certainty. Rivals do not observe shocks to investment costs, scrap values, or setup costs, and therefore are uncertain about market structure in the next period.

4.5 Equilibrium

In each time period, firm i makes entry, exit, production, and investment decisions, collectively denoted by a_i. Since the full set of dynamic Nash equilibria is unbounded and complex, I restrict the firms' strategies to be anonymous, symmetric, and Markovian, meaning firms only condition on the current state vector when making decisions, as in Maskin and Tirole (1988) and Ericson and Pakes (1995).

Each firm's strategy, $\sigma_i(s)$, is a mapping from states to actions:

$$\sigma_i : s \longrightarrow a_i. \tag{7}$$

In the context of the present model, $\sigma_i(s)$ is a set of policy functions which describes a firm's production, investment, entry, and exit behavior as a function of the present state vector. In a Markovian setting, with an infinite horizon, bounded payoffs, and a discount factor less than unity, the value function is:

$$V_i(s;\sigma(s),\theta) = \max_{\sigma_i(s)} \left\{ \pi_i(s,\sigma(s);\theta) + \beta \int V_i(s';\sigma(s'),\theta)\ dP(s';s,\sigma(s)) \right\}, \tag{8}$$

where θ is the vector of payoff-relevant parameters, $\pi_i(s,\sigma(s);\theta)$ is the per-period payoff function, and $P(s';\sigma(s),s)$ is the conditional probability distribution over future state s', given the current state, s, and the vector of strategies, $\sigma(s)$. The value function represents the expected discounted stream of payoffs accruing under optimal behavior given the strategies of competitors, denoted by $\sigma_{-i}(s)$.

Note that a firm will only undertake investment when its draw on investment costs is sufficiently low:

$$\begin{aligned} \gamma_{1i} + \gamma_2 x_i^* + \gamma_3 x_i^* + \beta \int V_i(s';\sigma(s'),\theta) dP(s_i + x^*, s'_{-i}; s, \sigma(s)) \\ \geq \beta \int V_i(s';\sigma(s'),\theta) dP(s_i, s'_{-i}; s, \sigma(s)). \end{aligned} \tag{9}$$

This induces truncation in the values of investment costs that firms will pay in equilibrium when they actually invest.[15] A similar consideration holds for the exit decision, where the draw on scrap values must be sufficiently high in order to make exit optimal:

$$\phi_i \geq \beta \int V_i(s';\sigma(s'),\theta) dP(s';s,\sigma(s)).$$

Potential entrants must weigh the benefits of entering at an optimally-chosen level of capacity against their draw from the distribution of entry costs. Firms only enter when the draw is sufficiently low. I assume that potential entrants are short-lived; if they do not enter in this period they disappear and may never enter in the future. This assumption is for computational convenience, as otherwise one would have to solve an optimal waiting

[15] A symmetric argument holds for divestment.

problem for the potential entrants.[16] Potential entrants will enter if the following inequality holds:

$$\kappa_i \leq \beta \int V_i(s'; \sigma(s'), \theta) dP(s'; s, \sigma(s)),$$

where $\sigma(s)$ is understood to contain the optimal investment for the entrant.

Markov perfect Nash equilibrium requires each firm's strategy profile to be optimal given the strategy profiles of its competitors:

$$V(s; \sigma_i^*(s), \sigma_{-i}(s), \theta) \geq V(s; \sigma_i'(s), \sigma_{-i}(s), \theta), \tag{10}$$

for all s and all possible alternative strategies, $\sigma_i'(s)$. This equilibrium concept places significant structure on the optimal behavior of firms; the inequality in Equation 10 forms the basis of the empirical estimator.

Doraszelski and Satterthwaite (2005) discuss the existence of pure strategy equilibria in settings similar to the one considered here. The introduction of private information over the discrete actions guarantees that at least one pure strategy equilibrium exists, as the best-response curves are continuous. There are no guarantees that the equilibrium is unique, a concern I discuss next in the context of my empirical approach.

5 Empirical Strategy

5.1 Overview

The empirical goal of this paper is to estimate all of the parameters in the theoretical model described above. In a direct approach to this estimation, such as in the method of simulated moments or simulated maximum likelihood, one searches over the vector of model parameters in order to minimize the divergence between the model's predictions and their empirical counterparts. However, as is well-documented in previous work such as Benkard (2004), generating predictions from dynamic models can be computationally expensive, as it is necessary to solve for a MPNE at every guess of the parameter vector. Further, the presence of multiple equilibria requires the econometrician to both compute the set of all possible equilibria and to specify how agents decide on which equilibrium will be played in the data, as in Bajari, Hong, and Ryan (2007b).[17]

[16] See Ryan and Tucker (2006) for an example of such an optimal waiting problem.

[17] Borkovsky, Doraszelski, and Kryukov (2008) outline a general approach to solving for the equilibria of Markovian games, and provide a good discussion of why it is generically hard to find all of the equilibria to

In order to sidestep these two issues, I follow the two-step empirical strategy laid out in Bajari, Benkard, and Levin (2007a), hereafter referred to as BBL. The intuition of BBL is straightforward: the econometrician lets the agents in the model solve the dynamic program, and finds parameters of the underlying model such that their behavior is optimal. The BBL estimator proceeds in two steps. In the first step, the econometrician flexibly estimates the equilibrium policy functions, $\sigma(s)$. Without imposing any structure, this step simply characterizing what firms do mechanically as a function of the state vector; these are reduced-form regressions correlating actions to states. This step also avoids the need to compute the equilibrium to the model, as the policy functions are estimated from the equilibrium that is actually played in the data.

The second step is to impose optimality on these recovered policy function by appealing to the definition of MPNE in Equation 10. By the construction of the value function in Equation 8, given an estimate for $\sigma(s)$ it is possible to construct $V_i(s; \sigma(s), \theta)$ for some guess of θ. It is possible to construct $V_i(s; \sigma_i', \sigma_{-i}(s), \theta)$ in analogous fashion by using some alternative policy function for firm i. Since the MPNE requirement holds against all possible alternative strategies, the alternative strategy can be any perturbation of the policies observed in the data, which are held to be optimal under the assumption of profit-maximizing behavior. Given a sufficiently rich set of alternative policies, the BBL estimator finds parameters θ such that profitable deviations from the optimal policies are minimized.

The first step in applying this methodology to the present context is to recover the policy functions governing entry, exit, and investment along with the product market profit function. In the second step, I take these functions and impose the restrictions of the MPNE to recover the dynamic parameters governing the costs of capacity adjustment and exit. Taken collectively, these estimates then allow me to simulate the value of a new firm entering the market, which can be used to recover the distribution of the sunk costs of entry. It is worthwhile to note that the assumption that demand and production are static implies that their corresponding parameters can be estimated directly from the data. This is useful because it improves the statistical efficiency of the estimator, as a subset of the parameters are identified independently of those which depend on the construction of the continuation value.

The approach in BBL has several regularity assumptions in order to produce valid estimates of the model primitives. Aside from functional form assumptions made below, the following assumption will allow me to group together all markets when estimating policy

these systems.

functions in the first step:

Assumption 3. *The same equilibrium in played in all markets.*

This assumption is critical to obtaining consistent estimates of the unknown parameters. For example, suppose that two equilibria are played in the data, with each equilibrium producing a distinct set of policy functions: $\sigma_1(s)$ and $\sigma_2(s)$. By grouping all markets together, the resulting estimate for $\sigma(s)$ is then a convolution of the policy functions corresponding to the two equilibria, and consistent with neither. It follows directly that the imposition of the MPNE requirement under an inconsistent estimate of $\sigma(s)$ will generically not produce consistent estimates of the underlying primitives. Under Assumption 3, it is possible to estimate the policy functions by grouping data from all markets.[18]

This assumption seems reasonable in the context of the cement industry. There does not appear to be any clear reason as to why the entry, exit, and investment behavior would be significantly different across markets in the data. The product is homogenous, players are often the same firms across different markets, the production technology has stayed essentially unchanged for a century, demand has relatively low volatility, and the industry turns over slowly.

I also make the following assumption regarding the beliefs of the firms with respect to the change in regulatory policy.

Assumption 4. *Firms assume that the regulatory environment is permanent.*

This assumption allows me to avoid having to model the beliefs of the firms regarding the distribution of future regulatory environments. In principle, it is possible to model these beliefs if there is an observable covariate which moves around beliefs of possible regulatory changes to the economic environment in the future. However, I assume that the firms behave as if the cost changes due to the Amendments were unanticipated, one-time changes that will never be repeated in the future. This assumption is standard in applications with regulatory change, such as Rothwell and Rust's (1997) study of nuclear power plant lifetimes under different operating license regimes.

[18]This is a stronger condition than the two-step approach requires: one could estimate the policy functions separately on each market, and then impose the MPNE conditions. However, the limitations of my data preclude such an approach.

5.2 Step One: Product Market Profits and Policy Functions

In the first step, I estimate the profits accruing to firms in each period and characterize the entry, exit, and investment behavior of firms conditional on the state variables.

5.2.1 Demand Curve and Production Costs

I estimate several variations on the following specification of the demand for cement in market j at time t:

$$\ln Q_{jt} = \alpha_0 + \alpha_1 \ln P_{jt} + \alpha_{2j} + \alpha'_{3t} X + \epsilon_{jt}. \tag{11}$$

The coefficient on market price, α_1, is the elasticity of demand, and X is a vector of covariates that influence demand. I assume that shocks to demand are iid. I instrument for the potential endogeneity of price with the error term using supply-side cost shifters: coal prices, gas prices, electricity rates, and wage rates. Each market has a demand shifter in the intercept, α_{2j} which is relative to the demand in Alabama.[19] I estimate several specifications of the demand function, including controls for housing permits, time trends, and population.

In order to estimate the costs of production, I search over δ to match the observed quantities for each firm in each market. For each guess of δ, I solve the vector of first-order conditions derived from Equation 3. To obtain an interior solution where all firms produce positive quantities, I make the following assumption:

Assumption 5. *The fixed costs of operation, δ_0, are zero.*

This assumption is empirically driven. As I do not directly observe profits, the only way to infer the fixed costs of operation is to observe firms shutting down production in some periods.[20] In the sample period, all firms produce in all periods, so I cannot identify fixed costs of operation. This normalization does not seem too stringent in the cement industry, both by the economic reasoning that fixed costs cannot be too large as zero production is never observed, and on the a priori grounds that these plants have relatively small staffing requirements and have a production technology where output quantity is directly proportional to energy and material inputs. This system of equations has an easily-computed fixed point, as the best-response curves are downward sloping in rivals' production.

For each firm i in market by j at time t, the estimator minimizes the difference between

[19] Alabama is the first market alphabetically.

[20] The formal identification of all parameters in the model is discussed below in Section 6.

the observed quantities and the predictions of the model, as given by Equation 3:

$$\min_{\delta}(nT)^{-1}\sum_{i=1}^{N}\sum_{j=1}^{J}\sum_{t=1}^{T}\left(q_{ijt}-\hat{q}_{ijt}(\delta)\right). \tag{12}$$

There are three basic parameters in the cost function: δ_1, δ_2, and ν. I also include post-1990 dummy shifters on each of those parameters to capture any changes in the production costs arising from the passage of the Amendments. In order to restrict the threshold at which capacity costs bind to be between 0 and 1, I make use of a logit transformation: $\nu = \exp(\widetilde{\nu})/(1.0+\exp(\widetilde{\nu}))$. If a firm has multiple plants in a single market, I treat that firm as having a single plant with capacity equal to the sum of capacity in each of those facilities.

5.2.2 Investment Policy Function

Both the presence of fixed costs in the model and empirical evidence such as the time series of capacities as in Figure 1 suggest that the empirical policy function should be both a function of the state variables and flexible enough to account for lumpy investment behavior. One model that satisfies both of these requirements is the (S, s) rule of investment, such as in Scarf (1959), where firms tolerate deviations from their optimal level of capacity due to fixed adjustment costs.[21] In the language of the (S, s) rule, firms have a target level bounded on either side by an adjustment band, both of which can be functions of observable variables. When the actual level of capacity hits one of the bands, the firm will make an adjustment to the target level. The target level and bands are only observed when the firm makes adjustments, and are flexibly parameterized to be functions of the underlying state variables. This model also nests the model of continuous investment as the bands go to zero, and is thus quite flexible in its ability to capture a range of investment behavior.

I follow Attanasio's (2000) empirical model of the (S, s) rule, and focus on the investment behavior of firms with positive capacity levels at the start and end of each period. Entry and exit is treated separately, and is acceptable in this context because I am only interested in what the investment behavior of a firm will be given a specific state. Firms have a target

[21] Deriving this rule as the explicit solution to an optimization problem is involved—see Hall and Rust (2000) for an example of the optimality of this rule in an inventory setting. When I solve for an equilibrium of the model the fixed costs of investment are large enough to induce firms to engage in lumpy investment behavior.

level of capacity that they adjust to when they make an investment:

$$\ln s_{it}^* = \lambda_1' bs(s_{it}) + \lambda_2' bs\left(\sum_{j\neq i} s_{jt}\right) + u_{it}^* \tag{13}$$

where the desired (logged) level of capacity is a function of the firm's own capacity, the sum of its competitors capacities, and a mean zero error term, u_{it}^*.

Since it is desirable to be as flexible as possible in modeling a firm's behavior as a function of the state, I use the method of linear sieves, a simple semi-nonparametric approach to approximating and estimating unknown functions, to estimate the target equation. The method of linear sieves approximates an unknown function using a sequence of simple finite-dimensional polynomial approximations to those functions, where the degree of the approximating polynomial grows with the number of observations. Outside of their desirable limit properties, which ensure that one can recover the unknown function to an arbitrarily close degree with infinite data, a great benefit of linear sieves is their simple analytical form, as they represent the function as the inner product of some unknown parameters and basis functions. In this particular case, the basis functions are cubic b-splines, which are finite-dimensional piecewise polynomials, denoted here by $bs(\cdot)$.[22]

The critical aspect of the (S, s) rule that generates lumpy investment behavior is that firms only adjust s_{it} to s_{it}^* when current capacity exceeds one of the bands around the target level. The lower and upper bands are given by:

$$\underline{s}_{it} = s_{it}^* - \exp\left(\lambda_3' bs_1(s_{it}) + \lambda_4' bs_2\left(\sum_{j\neq i} q_{jt}\right) + \underline{u}_{it}^b\right), \tag{14}$$

and

$$\bar{s}_{it} = s_{it}^* + \exp\left(\lambda_5' bs_1(s_{it}) + \lambda_6' bs_2\left(\sum_{j\neq i} q_{jt}\right) + \bar{u}_{it}^b\right). \tag{15}$$

The inclusion of the exponential function ensures that the desired level of adjustment is always in between the bands. This model also nests a model of continuous adjustment as the width of the bands goes to zero. I assume that the residuals in the bands are iid normal with zero mean and equal variance, and are independent of the error in the target. Note

[22] Chen (2006) provides an exhaustive overview of the approximating and statistical properties of b-splines and other linear sieves. Throughout the paper, I used uniform b-splines with ten knots, where the range of the knots was chosen to bound the empirical data. Further implementation details of the b-splines are available from the author upon request.

that when a firm makes an adjustment, it reveals both the target level, s^*_{it}, and the size of the band, $s^*_{it} - s_{i,t-1}$, which is sufficient to identify the parameters of the adjustment policy function.

I assume that the upper and lower bands are symmetric functions of the target capacity ($\lambda_3 = \lambda_5$ and $\lambda_4 = \lambda_6$); the reason is that I observe so few downward adjustments in capacity than the upper bound is not precisely identified. Since I assume that the change in capacity simultaneously reveals the size of the band and the target level, I use a first-stage OLS estimator to recover λ

I estimate separate policy functions for the period before 1990, and the period after 1990. This will capture any differences in the firms' equilibrium investment behavior caused by a permanent shift in the cost and regulatory environment.

5.2.3 Entry and Exit Policy Functions

I characterize the probability of entry using a probit regression:

$$Pr(\chi_i = 1; s_i = 0, s) = \Phi\left(\psi_1 + \psi_2\left(\sum_{j\neq i} s_{jt}\right) + \psi_3 1(t > 1990)\right), \tag{16}$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal, and $1(\cdot)$ is the indicator function. The exit policy is also modeled analogously:

$$Pr(\chi_i = 0; s_i > 0, s) = \Phi\left(\psi_4 + \psi_5 s_{it} + \psi_6\left(\sum_{j\neq i} s_{jt}\right) + \psi_7 1(t > 1990)\right). \tag{17}$$

Explanatory variables in both policy functions are a constant, the sum of competitors' capacities, and a dummy variable for before and after 1990. I also add the firm's own capacity to the exit equation. I am assuming that there is only one possible entrant in any period; given the very low rate of entry in the cement industry, this assumption is not very important.

5.3 Step Two: Recovering the Structural Parameters

The first step provides functions that describe both how the state vector evolves over time and what product market profits are at each state. The second step is concerned with finding parameters that make these observed policy functions optimal, given the underlying theoretical model.[23]

To derive the estimator for investment costs and the distribution of scrap values, recall the firm's value function:

$$V_i(s;\sigma(s),\theta) = \max_{\sigma_i(s)} \left\{ \pi_i(s,\sigma(s);\theta) + \beta \int V_i(s';\sigma(s'),\theta)\ dP(s';s,\sigma(s)) \right\}, \qquad (18)$$

Note that it is possible to decompose $\pi_i(s,\sigma(s);\theta)$ into the inner product of its known and unknown components:

$$\begin{aligned}\pi_i(s,a;\theta) &= P(Q;\alpha)q_i - C(q_i;\delta) - \Gamma(x_i;\gamma) + \Phi_i(a_i;\kappa_i,\tau,\phi) \\ &= u_i(s)\cdot 1 - 1(x_i>0)\cdot(\widetilde{\gamma_{1i}} + \gamma_2 x_i + \gamma_3 x_i^2) + 1(x_i<0)\cdot(\widetilde{\gamma_{4i}} + \gamma_5 x_i + \gamma_6 x_i^2) \\ &\quad + 1(\chi_i = 1)\cdot\tau + 1(\chi_i=0)\cdot\widetilde{\phi}_i,\end{aligned}$$

where $u_i = P(Q;\alpha)q_i - C(q_i;\delta)$, the profit of being active in that period, which is a known quantity from the first-stage regressions. Also known are the policy functions describing investment and exit behavior.

As discussed above, the draws on investment, divestment, and exit costs, $\widetilde{\gamma_{1i}}$, $\widetilde{\gamma_{4i}}$, and $\widetilde{\phi}_i$, are not equal to the mean of their respective distributions, F_γ, G_γ, and F_ϕ, as firms will only act when those draws are sufficiently favorable. Instead of making parametric assumptions about the distribution of the fixed costs and deriving the associated truncated distributions, I instead consider direct estimation of those quantities using linear sieves. Writing the approximation for both expected truncated distributions as linear b-splines:

$$\widetilde{\gamma_{1i}}(p_i) = \theta_{\gamma,1}\cdot bs(p_i(s)),\ \widetilde{\gamma_{4i}}(p_d(s)) = \theta_{\gamma,2}\cdot bs(p_d(s)),\text{ and } \widetilde{\phi}(p_e) = \theta_\phi \cdot bs(p_e),$$

where θ_γ and θ_ϕ are finite-dimensional vectors. The basis functions are defined over the probabilities of investment, $p_i(s)$, divestment, $p_d(s)$, and exit, $p_e(s)$, as these are sufficient statistics for the state vector.[24]

As in Bajari, Benkard, and Levin (2007a), it is useful to exploit the fact that all of the

[23] This section follows the derivations in Bajari, Benkard, and Levin (2007a) closely.

[24] A simple parametric example establishes why the truncated values are a function of the probabilities of their respective actions. Suppose that the fixed cost of exit is distributed normally. A firm will exit when the draw on scrap values is greater than the continuation value of remaining in the industry, denoted by $c(s)$: $\phi_i \geq c(s)$. The probability of this event is given by the policy function, $p_e(s)$. Note that $p_e(s) = Pr(\phi_i \geq c(s)) = 1 - \Phi(c(s);\mu_\phi,\sigma_\phi)$. Inverting obtains: $c(s) = \Phi^{-1}(1-p_e;\mu_\phi,\sigma_\phi)$. It is then possible to compute the expected scrap value conditional on entering, $E[\phi_i|\phi_i \geq \Phi^{-1}(1-p_e;\mu_\phi,\sigma_\phi);\mu_\phi,\sigma_\phi]$, which is only a function of the probability of exit and the parameters of the exit cost distribution.

unknown parameters enter linearly into the payoffs of the firm in the current period and all future periods:

$$\pi_i(s, a; \theta) = \zeta(s_{it}) \cdot \theta. \tag{19}$$

Defining the following:

$$W_i(s_t; \sigma(s)) = E_{\sigma(s)} \sum_{t'=0}^{\infty} \beta^{t'} \zeta(s_{i,t+t'}), \tag{20}$$

the value function is then:

$$V_i(s_t; \sigma(s), \theta) = W_i(s_t; \sigma(s)) \cdot \theta. \tag{21}$$

Imposing the Markov perfect equilibrium condition (see Equation 10) for all alternative policies σ_i' obtains:

$$W(s_o; \sigma_i^*, \sigma_{-i}) \cdot \alpha \geq W(s_o; \sigma_i', \sigma_{-i}) \cdot \theta, \tag{22}$$

At the true parameters the above relation should hold for all alternative policies. Exploiting the linearity of the unknown parameters, I can rewrite the above equation in terms of profitable deviations from the optimal policy:

$$g(\sigma_i'; \theta) = [W(s; \sigma_i', \sigma_{-i}) - W(s; \sigma_i^*, \sigma_{-i})] \cdot \theta. \tag{23}$$

Intuitively, I want to find parameters such that profitable deviations from the optimal policies are minimized. Formally, I draw alternative policies from a distribution H over all alternative policies to generate a set of $n_k = 90,000$ inequalities. The true parameter minimizes:

$$\min_{\theta} \int 1\left(g(\sigma_i'; \theta) > 0\right) g(\sigma_i'; \theta)^2 dH(\sigma_i'). \tag{24}$$

To form the estimator, I replace the above with its sample counterpart:

$$\min_{\theta} Q_n(\theta) = \frac{1}{n_k} \sum_{j=1}^{n_k} 1\left(g(\sigma_{i,j}'; \theta) > 0\right) g(\sigma_{i,j}'; \theta)^2. \tag{25}$$

Implementing this estimator proceeds in two separate steps. In the first step, I compute W for both the observed and alternative policies for a very large set of starting states, drawn randomly.[25] I generate the alternative policies by adding noise to the observed policy

[25]The maximal number of firms in these situations are set well above the number of firms that one ever observes in the data for a similar sized market. This ensures that simulated distribution of potential entrants

functions. For example, to perturb the exit policy function I add an error drawn from the standard normal to the terms inside the exit probit. The linearity of the unknown parameters becomes useful during the minimization, as I do not have recompute separate outcome paths for each set of parameters. Note that the function is not trivially minimized at the zero vector because the profits from the product market enter in each time period. I use the Laplace-type estimator (LTE) of Chernozhukov and Hong (2003) to search over θ in Equation 25. The LTE procedure is robust to non-smooth functions and has the nice feature of jointly estimating the mean and variance of the unknown parameters.

5.3.1 Distribution of Sunk Entry Costs

Having recovered the policy functions and the parameters necessary for the construction of the period payoffs it is possible to find the distribution of sunk costs. Consider the value function of a potential entrant:

$$V_i^e(s,\kappa_i) = \max\left\{0, \max_{x_i}\left\{-\kappa_i - \gamma_1 - \gamma_2 x_i - \gamma_3 x_i^2 + \beta E(V(s')|s)\right\}\right\}. \tag{26}$$

All of the terms in Equation 26 are known or computable at this stage except for κ_i.[26] Clearly, an entrant will enter a market if and only if the draw of κ_i is low enough. From the perspective of the econometrician, the probability that a firm enters is equal to the probability of that entrant receiving a draw that is less than the value of entry:

$$Pr(\kappa_i \leq EV^e(s)) = F_\kappa(EV^e(s); \mu, \sigma^2). \tag{27}$$

The left-hand side of Equation 27 corresponds to the entry policy function estimated in Equation 16. It is possible to simulate $EV^e(s)$ through forward simulation, as above when recovering investment and exit costs. I assume that F_κ is a normal distribution with mean μ_κ and variance σ_κ^2. Drawing $s = \{1, \ldots, NS\}$ random states of the industry, I search for parameters of this distribution which match the observed probabilities of entry as well as possible:

$$\min_{\{\mu_\kappa,\sigma_\kappa^2\}} \frac{1}{NS}\sum_{i=1}^{NS}\left[Pr(\text{entry}|s_i) - \Phi\left(EV^e(s_i);\mu_\kappa,\sigma_\kappa^2\right)\right]^2. \tag{28}$$

corresponds to that faced by real firms.

[26] The optimal investment is given by the investment policy function; the vector of investment costs, γ, was estimated above; and forward simulation produces an approximation of the continuation value.

I estimate the parameters of the distribution of sunk entry costs separately for the time periods before and after the 1990 Amendments.

5.3.2 Standard Errors

Note that the standard errors of parameters estimated in later stages should be corrected for estimation error from the first stage. In principle, one could jointly estimate all equations simultaneously, or draw from the posterior covariance matrix of the earlier-stage parameters when performing inference on the second-stage parameters.[27] However, the computational burden of doing this has prevented me from correcting the errors; as a result the reported standard errors are biased downward from their true levels. Partially offsetting this limitation is that many of the first stage regressions are statistically independent.

6 Identification

There are two sets of parameters in the model: those that are estimable without appeal to a dynamic model, and those that depend on the continuation value. The demand curve is nonparametrically identified under much weaker monotonicity and exclusion restrictions than imposed by the linear functional form in Equation 11.[28]

The parameters of the production function are identified by functional form. The system of equations formed by Equation 3 has a single fixed point, since the best-response curves are downward-sloping in their rivals' production. As the residual demand curve facing an individual firm moves in and out, it traces out the marginal cost of production. As mentioned earlier, the fixed cost of production would be identified when the firm chooses not to produce anything in a given period. However, in the present data sample all firms produce in all periods, so this parameter must be normalized to zero.

This section provides a novel constructive approach to showing identification in two-step estimators, and demonstrate that the necessary and sufficient identification conditions are met in the present model.[29] I also show how to estimate and identify parameters of

[27]I experimented extensively with drawing from the marginal distribution of the first-stage parameters. However, it quickly became apparent that one must draw from the full joint distribution of the estimated parameters, as ignoring the covariance between parameters produced unrealistic combinations of primitives that frequently resulted in non-sensible predictions. These combinations would be ruled out by drawing the correct full joint distribution.

[28]See Newey and Powell (2003) and references therein for a general treatment of identification and estimation in nonparametric instrumental variables models.

[29]This section borrows heavily from Rust (1994).

unknown distributions in the underlying dynamic game, such as the distribution of fixed adjustment costs, which extends the class of models previously considered in the literature. The identification conditions are easy to verify, and apply to a wide class of dynamic games.

For expositional purposes, first assume that the state space consists of a finite number of states. The value function can be written as the following infinite sum:

$$V(\sigma,\theta) = \pi(\sigma,\theta) + \beta P(\sigma)\pi(\sigma,\theta) + \beta^2 P^2(\sigma)\pi(\sigma,\theta) + \beta^3 P^3(\sigma)\pi(\sigma,\theta) + \dots, \tag{29}$$

where $P^k(\sigma)$ is understood to mean the probability distribution over states in k periods given that the policies σ are followed in the intermediate stages. Since $P(\sigma)$ is a probability transition matrix, with each element bounded in the unit intervals and each row summing to one, and $0 < \beta < 1$, the terms go to zero in the limit:

$$\lim_{k\to\infty} \beta^k P^k(\sigma) = 0. \tag{30}$$

If $\pi(\sigma,\theta)$ is bounded, then this series converges, which guarantees the inverse exists and is unique in the following matrix representation of the value function:

$$V(\sigma,\theta) = \pi(\sigma,\theta) + \beta P(\sigma)V(\sigma,\theta) \tag{31}$$

$$V(\sigma,\theta)(I - \beta P(\sigma)) = \pi(\sigma,\theta) \tag{32}$$

$$V(\sigma,\theta) = \pi(\sigma,\theta)(I - \beta P(\sigma))^{-1}. \tag{33}$$

Therefore, given $P(\sigma)$, σ, and $\pi(\sigma,\theta)$, the value function is unique.

This construction is also valid if the state space is either a countably infinite or uncountably infinite set of points. If the state space is countably infinite, then $P(\sigma)$ can be represented by an infinite matrix. If the state space is uncountably infinite, then $P(\sigma)$ can be replaced by the following linear Markov operator:

$$P(\sigma(s)) = \int V(s')p(ds'|s,\sigma(s)), \tag{34}$$

where p is the probability distribution function over future states. This operator has a unit norm, and higher powers of this operator are also Markov operators. Substituting the operator in the series given in Equation 29 results in a Neumann series. This series converges by arguments analogous to the finite state case, which then implies that the inverse operator in Equation 33 exists.

The second stage in two-step estimators is to impose optimality of the observed policy functions. For continuous actions, this implies that the derivative of the value function with respect to the policies is equal to zero:

$$\frac{\partial V(\sigma,\theta)}{\partial\sigma} = \beta(I-\beta P(\sigma))^{-1}\frac{\partial P}{\partial\sigma}(I-\beta P(\sigma))^{-1}\pi(\sigma,\theta) + (I-\beta P(\sigma))^{-1}\frac{\partial\pi(\sigma,\theta)}{\partial\sigma} = 0. \quad (35)$$

Pre-multiplying by $(I-\beta P)^{-1}$ and consolidating terms results in:

$$\frac{\partial V(\sigma,\theta)}{\partial\sigma} = \beta\pi(\sigma,\theta)(I-\beta P(\sigma))^{-1}\frac{\partial P}{\partial\sigma} + \frac{\partial\pi(\sigma,\theta)}{\partial\sigma} = 0. \quad (36)$$

Equation 36 neatly summarizes the changes in the value function from a change in the policy functions: the first term is the expected change in value due to changes in the distribution of future states; and the second term is the direct change in per-period payoffs. The first term incorporates the subsequent responses by rivals to a change in the agent's state vector, although it is important to recognize that rivals' strategies do not change as a result of the agent's deviation. The second term includes the current change to payoffs, which is typically a current period cost paid by the agent in order to undertake an action which influences future states; examples are the costs of capital investment and R&D expenditures.

The econometrician does not observe shocks that the agents act upon, therefore it is necessary to work with their expected counterparts. The vectorized ex-ante profit function $\pi(\sigma,\theta)$ is:

$$\pi(\sigma,\theta) = p_i(\widetilde{\gamma}_1(p_i) + \gamma_2 x_i + \gamma_3 x_i^2) + p_d(\widetilde{\gamma}_4(p_d) + \gamma_5 x_d + \gamma_6 x_d^2) + p_e\widetilde{\phi}(p_e) + \bar{\pi}, \quad (37)$$

where x_i and x_d are vectors of investment and divestment, respectively. Substituting in the definitions for the expected truncated variables obtains:

$$\pi(\sigma,\theta) = p_i(\theta_{\gamma,1}bs(p_i) + \gamma_2 x_i + \gamma_3 x_i^2) + p_d(\theta_{\gamma,2}bs(p_d) + \gamma_5 x_d + \gamma_6 x_d^2) + p_e\theta_\phi bs(p_e) + \bar{\pi}. \quad (38)$$

All of the unknown parameters enters linearly into Equation 38. At each state in Equation 36, we obtain a system of linear equations after substituting in $\pi(\sigma,\theta)$ from Equation 38. Let X denote the matrix formed by stacking Equation 36 evaluated at $k=\dim(\theta)$ different states, and let Y denote the analogous vector of constants. Under the condition that X has full column rank, the unique solution to the system of equations is:

$$\theta = X^{-1}Y, \quad (39)$$

which establishes the identification of the parameters in Equation 37. In the present model, the full rank of X is guaranteed by nonlinearity in the policy functions.

It remains to show that the estimated truncated fixed cost functions, $\widetilde{\gamma}_1(p_i)$, $\widetilde{\gamma}_4(p_d)$, and $\widetilde{\phi}(p_e)$, identify their associated fixed cost distributions. It is necessary and sufficient to establish that the distribution function is one-to-one with the truncated fixed cost function. Consider the case of the fixed costs of investment. First, note that the probability of investment is equal the probability that the firm received a draw of investment costs lower than some bound, d:

$$p_i = Pr(\gamma_1 \leq d) = F_\gamma(d).$$

Define the inverse of the distribution function as follows:

$$d = F_\gamma^{-1}(p_i) = \inf\{x \in R : p_i \leq F_\gamma(x)\}.$$

If F is strictly increasing, d is unique; otherwise it is the smallest value x such that the inequality is satisfied. In either case, knowledge of the inverse function fully characterizes the distribution function. By the definition of conditional expected value:

$$\widehat{\gamma}_1(p_i) = E(\gamma_1|\gamma_1 \leq F_\gamma^{-1}(p_i)) = \frac{1}{p_i}\int_{-\infty}^{F_\gamma^{-1}(p_i)} x f_\gamma(x)dx.$$

Multiplying both sides by p_i and differentiating with respect to p_i results in:

$$\frac{\partial \widetilde{\gamma}_1(p_i)p_i}{\partial p_i} = F_\gamma^{-1}(p_i) f_\gamma(F_\gamma^{-1}(p_i))\frac{\partial F_\gamma^{-1}(p_i)}{\partial p_i}. \tag{40}$$

Applying the definition of a derivative of an inverse function:

$$\frac{\partial F_\gamma^{-1}(p_i)}{\partial p_i} = \frac{1}{f_\gamma(F_\gamma^{-1}(p_i))},$$

and substituting into Equation 40 obtains:

$$\frac{\partial \widetilde{\gamma}_1(p_i)p_i}{\partial p_i} = F_\gamma^{-1}(p_i),$$

which establishes that the inverse distribution function is one-to-one in $\widehat{\gamma}_1(p_i)p_i$. The desired identification result follows from the fact that the inverse distribution function completely characterizes the distribution function. It is possible to show the identification of the distri-

Table 3: Cement Demand Estimates

	I	II	III	IV	V	VI
Price	-3.21	-1.99	-2.96	-0.294	-2.26	-0.146
	(0.361)	(0.285)	(0.378)	(0.176)	(0.393)	(0.127)
Intercept	21.3	10.30	20.38	-3.41	11.6	-6.43
	(1.52)	(1.51)	(1.56)	(1.09)	(2.04)	(0.741)
Log Population		0.368		0.840	0.213	0.789
		(0.0347)		(0.036)	(0.074)	(0.033)
Log Permits					0.218	0.332
					(0.072)	(0.035)
Market Fixed Effects	No	No	Yes	Yes	No	Yes

Dependent variable is logged quantity. All market-specific fixed effects are relative to Alabama. Instruments were gas prices, coal prices, electricity prices, and skilled labor wage rates. There are a total of 517 observations.

butions of divestment and sunk exit costs in an analogous fashion.

The identification of the distribution of sunk entry costs is analogous to identification of a single-agent probit. Restating Equation 27:

$$Pr(Entry; s) = Pr(\kappa_i \leq EV^e(s)) = \Phi(EV^e(s); \mu, \sigma^2), \tag{41}$$

where μ and σ^2 are the mean and variance of the normal distribution CDF. The probability of entry is known perfectly, and is a continuous function of the state variables. The expected value of entering the market, $EV^e(s)$, is fully known from the behavior of incumbent firms. Identification requires that there exist two states, s and s', such that $EV^e(s) \neq EV^e(s')$, which would be satisfied, for example, by considering the entry of a monopolist into two markets with differing levels of demand.

7 Empirical Results

Demand Curve I estimate the parameters in Equation 1, the demand curve for Portland cement, using market-level data on prices and quantities. I use several cost-side shifters serving as instruments to account for the endogeneity of prices. The results are presented in Table 3.

The first specification is the simplest, as it is has no covariates. The price elasticity of demand is precisely estimated to be -3.21. However, one may expect that demand may vary across markets due to population or other unobservable factors. The next two models test for these factors. Specification II adds in controls for population, in logs, which is estimated to have a positive effect on quantity. The elasticity falls to -1.99, and the intercept is substantially lower. The average log population in the sample is 15.87; multiplying by the coefficient on log population shifts the constant back to 16.1, closer to the baseline model's intercept. Specification III includes market-specific fixed effects in lieu of population shifters. The results are very similar to the baseline specification, with higher elasticity and intercept than in specification II. Specification IV includes both population and fixed effects. The result is a statistically insignificant estimate of the price coefficient, a negative intercept, and a positive coefficient on log population. Specification V includes a measure of housing permits allocated in each market-year. The elasticity of demand is estimated to be -2.26, with positive coefficients on population and permits. Specification VI includes fixed effects for markets as well, which leads to a small and statistically insignificant estimate of the price elasticity.

The specification I choose to use throughout the rest of the analysis is specification III. The reasoning behind this choice is three-fold. First, it appears that market fixed effects capture much of the same cross-market variation in prices that population and permits do. A regression of quantities on prices, population, permits, and the interaction of population and permits with a time trend leads to imprecisely measured zeros on the interaction terms. This suggests that population and permits are not changing very much within market, and their explanatory power is cross-sectional. This argues that market fixed effects may reasonably proxy for these effects. Second, while the fixed effects approach is not as nuanced as the population and permits approach, which utilizes more data, it has the benefit of being the more parsimonious specification. With imperfect data, as in the present context, this is a strength. Finally, a simple plausibility check suggests that the specification with the higher elasticity results in a more reasonable estimate of plant costs. If one takes a specification with a lower elasticity as the demand curve, and works through the ensuing empirical exercise, the resulting estimates imply that firms face unreasonably large investment costs in order to rationalize their behavior. Otherwise, firms would be leaving very large amounts of money on the table; as such, the estimator predicts investment costs on the order of several billion dollars for a modern plant, which is inconsistent with anecdotal newspaper evidence and the

Table 4: Production Function Estimation Results

Parameter	Coefficient	Standard Error
δ_1	30.0	2.59
δ_2	1,157 ($\times 10^7$)	606.5
$\widetilde{\nu}$	1.896	0.083
δ_1 post-1990 shifter	2.41	3.33
δ_2 post-1990 shifter	-2.99 ($\times 10^7$)	22.0
$\widetilde{\nu}$ post-1990 shifter	0.0917	0.0801

The binding threshold at which the capacity costs become important is restricted to $[0, 1]$ by estimating a logit probability: $\nu = \exp(\widetilde{\nu})/(1.0 + \exp(\widetilde{\nu}))$. At the estimated value of 1.896, this implies that capacity costs start to bind at an approximately 87 percent utilization rate.

accounting data cited in Salvo (2005).[30] For these reasons, I proceed with specification III as the model for cement demand.

To verify that the instruments used in the demand estimation are both correlated with the endogenous regressor and orthogonal to the error term, I evaluate both the fits of the instruments on the endogenous regressor and the Anderson-Rubin statistic. The F-statistic of the first-stage regression of the instruments on the endogenous regressor results is 42.99, which is significant at the 0.01 percent level and well above the rule-of-thumb threshold of 10. The Anderson-Rubin statistic is 52.54, which is also significant at the 0.01 percent level. I conclude that the tests fail to reject the hypothesis that the instruments were both well-correlated with prices and orthogonal to the error terms.

Finally, I test for the presence of time trends in each of the markets. While the F-test rejects the hypothesis that all of the coefficients are significantly different from zero, most of the market-time trends are not individually significant (22 out of 26 markets). The elasticity of demand is precisely estimated to be -2.26, with an intercept of 17.51. Saturating the model with trends and dummy variables is strong empirical evidence that the elasticity is in the range of -2 to -3, as the explanatory variables account for a wide range of market- and time-specific unobserved heterogeneity.

Production Costs Having estimated the demand curve, I recover the production cost parameters by matching predicted quantities as closely as possible to their empirical counterparts. I estimate six parameters: marginal cost, capacity cost, the capacity binding level,

[30] A table of announced plant costs is available in a previous version of the paper and from the author upon request.

Table 5: Implied Prices, Revenues, Costs, and Profits

Variable	Value	Standard Deviation
Price	57.81	16.83
Revenues	39,040	19,523
Costs	22,525	11,051
Profit	16,515	12,244
Margin	39.29 percent	18.21 percent

Implied prices, revenues, costs, and profits for every firm in the sample calculated at the estimated demand and production parameters. Prices, revenues, costs, and profits are measured in thousands of dollars. Margin is the calculated as profits divided by revenues.

and post-1990 shifters for each. The results are shown in Table 4. The estimates indicate that capacity costs become important as firms increase production beyond 87 percent of their boilerplate capacity. Once firms cross this threshold they experience large, linearly increasing marginal costs as they cut into the normal period of maintenance downtime. The penalty for cutting out your maintenance is significant, preventing most producers from exceeding 90 percent of their stated production capacity. The standard errors were calculated using the bootstrap.

I test for differences in the cost parameters before and after 1990. I find that there have been slight increases in productive efficiency after 1990. However, I fail to reject the null hypothesis that the coefficients on the dummy variables for post-1990 are all zero at the 5 percent level. This helps strengthen the argument that the Amendments did not have an influence on marginal costs. This is also a positive result in another regard, as it is necessary evidence to reject the idea that there was an additional, unobserved shock to the industry's costs structure over this time period.

The relationship between δ_2, the parameter governing how fast costs rise when approaching the theoretical capacity limit, and its post-1990 shifter is also of interest. Due to the increasing nature of costs at the margin where these parameters bind, they are poorly identified relative to each other. This tends to inflate the variance and understate the significant of those capacity costs. Restricting the late dummy for capacity costs to be zero results in the same numbers for pre-1990 costs, with much lower variances and statistical significance for δ_2. In this restricted case, I also fail to reject the null hypothesis that costs are equal across the two time periods at the 10 percent level.

As a check on the estimated parameters, I compute the market price, revenues, costs, and profit margin for every firm in my sample. The summary statistics for these values are shown in Table 5. The prices are well within the range seen in the data. The average firm grosses

Table 6: Productivity Estimates

Specification	I	II	III	IV	V
Capacity	0.8617	0.8600	0.860	0.860	0.860
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Rivals' Capacity	-0.007	-0.005	-0.002	-0.003	0.0003
	(0.001)	(0.001)	(0.001)	(0.001)	(0.0006)
Firm Entered * Capacity		0.0009	0.0002	0.0112	0.0103
		(0.0027)	(0.0027)	(0.0064)	(0.007)
Firm Exited * Capacity		-0.0154	-0.0128	-0.0173	-0.0135
		(0.0035)	(0.0036)	(0.0078)	(0.008)
Time Trend			0.671	0.681	
			(0.130)	(0.131)	
Entry Dummy				-11.66	-11.49
				(6.141)	(6.678)
Exit Dummy				3.041	0.492
				(4.810)	(5.107)
Market Fixed Effects	Yes	Yes	Yes	Yes	No
Market-Time Fixed Effects	No	No	No	No	Yes
R^2	0.9925	0.9925	0.9926	0.9926	0.9933

Number of observations = 2,233.

slightly less than \$40M a year. Profits average just over \$16M a year, which is little less than a 40 percent profit margin. This is a plausible gross return, as public financial records for major cement producer Lafarge North America report an 33 percent average gross profit margin for the three-year period 2002-2004.[31]

To test the assumption that the firms have no persistent productivity differences, I regressed output quantity on own capacity and various controls. If there are productivity differences across firms, it should be expressed in their ability to utilize their productive capacity: more productive firms produce more given the same amount of capacity. This should be an especially strong test given that most firms were capacity constrained during this time period. The controls include the capacity of rival firms, a time trend, market fixed effects, and capacity interacted with dummy variables for whether the firm entered or exited during the sample period. Market fixed effects capture variation in local demand conditions, and the capacity of rival firms shifts around residual demand facing any firms. The dummy

[31] Sales and profit data are from Hoover's Online "Annual Financials" fact sheet for Lafarge S.A., 2002-2004. http://www.hoovers.com.

variables for entry and exit capture systematic differences in productivity, as measured by production per unit of capacity. The results are presented in Table 6.

The first specification considers only own capacity and rivals' total capacity, controlling for market with fixed effects. This model does remarkably well in fitting the data, where $R^2 = 0.9925$. This suggests that unobservable productivity differences cannot be very important in the sample, as otherwise there would be significant variance in the output that variation in capacity alone could not explain. This variable is very precisely estimated across all the specifications, further supporting the idea that production is largely explained by observable capacity, controlling for common factors to all firms in the market.

To examine the question of whether firms were selecting in and out of the cement industry along these unobservables, I included dummy variables for firms that entered and exited during the sample period. The second specification allows the production rate per unit of capacity to shift for firms that entered and exited. The results suggest that new firms are no more productive per unit capacity than the average firm in the industry, while exiting firms on average produced 1.8 percent less per unit of capacity than the average firm. The third specification adds a time trend to production; if firms tended to exit earlier in the sample vis a vis new entrants, this could bias the selection effect. The addition of a time trend reduces the difference in productivity for exiting firms, although it is still significant. The last specification also adds the dummy variable for entering and exiting firms directly. If the productivity differences are explained by differences in startup times, this could show up through a level effect. Curiously the results suggest that entering firms have lower production levels than both the average firm and exiting firms, which are more productive than average firms. The production per unit capacity differential for entering (exiting) firms is still positive (negative), although neither is now significant at the 5 percent level. Finally, saturating the model with time-market fixed effects to control for any other observable factors leads to insignificant estimates for productivity differences across entering and exiting firms.

In conclusion, it appears that there is little reduced-form evidence for productivity differences across entering firms and the average firm in our sample. There is slightly more evidence for exiting firms having lower productivity than the average firm, but this is partially offset by the result that these firms have higher baseline levels of output than the average firm. Given that exiting firms tend to be smaller than the average firm, the effect of treating an exiting firm as an average firm is minimized. The most general evidence, where the specification includes time-market fixed effects, leads to statistically insignificant estimates of the differences between entering and exiting firms.

Investment Policy I model the investment policy function as an (S, s) rule. Under the assumption that the bandwidth and target level are observable when a firm makes an adjustment, it is possible to obtain consistent estimates for the parameters in Equations 13 and 15 using a simple linear regression. The bandwidth is determined by a regression of functions of a firm's own capacity and the capacity of a firm's competitors on the size of the change. The target level coefficients are determined by regressing a second, separate equation with the same state variables on the post-adjustment capacity. In order to estimate this policy function as flexibly as possible I use cubic B-splines as basis functions for the capacity variables. The model is estimated using OLS. The estimates for several specifications, run on the full set of data, are presented in Tables 7 and 8.

The first two specifications in the band equation use levels of capacity as explanatory variables, with specification II including population as a level shifter. I constructed basis functions of all three functions to allow the marginal effect to vary with the magnitude of the covariate. The base specification does a very good job of matching the observed size of adjustments. The adjusted R^2 for both regressions is almost 0.9. A regression of fitted values on the actual adjustments produces a regression with an imprecisely estimated zero intercept and a tightly estimated slope parameter of almost exactly 1; this indicates that on average the model is able to fit the observed gaps very well using the flexible basis functions over the sum of competitors' capacity and a firm's own capacity. The inclusion of population basis functions, as in specification II, slightly improves the fit of the model at the expense of greatly increasing the variance of the estimated parameters. A standard F-test fails to reject the null hypothesis that the coefficients on the population basis functions are jointly zero (p-value of 0.2149).

One may think that the inclusion of population should be through per-capita capacities to adjust for differences in market sizes. Specification III runs the same regression of log adjustment on basis functions of per-capita capacities. In this case the model does slightly less well in fitting the size of the adjustment band. The adjusted R^2 is lower, and the estimated variance in the error term is a bit larger. The signs on the coefficients for own capacity are reversed, but are not statistically significant. Competitors' aggregate capacity also enters in strongly and positively. To ensure that this isn't proxying for market-level demand shifters that make large firms uniformly more attractive, Specification IV adds in region fixed effects. The fixed effects are estimated very imprecisely, and tend to be small deviations around zero. The magnitude of the coefficients on aggregate competitor capacity are even stronger with the fixed effects, and the sign of own capacity reverts to being negative,

Table 7: Investment Policy Function Results: Adjustment Band Size

Specification	I	II	III	IV
Sum Competitors Capacity B-spline 1	6.31	7.18	3.29	5.74
	(0.973)	(5.42)	(0.827)	(1.14)
Sum Competitors Capacity B-spline 2	6.51	7.16	2.04	4.53
	(0.930)	(5.43)	(0.953)	(1.37)
Sum Competitors Capacity B-spline 3	5.66	6.41	3.57	4.89
	(0.910)	(5.40)	(0.888)	(1.28)
Sum Competitors Capacity B-spline 4	6.98	7.85	2.05	4.05
	(0.960)	(5.46)	(0.978)	(1.36)
Sum Competitors Capacity B-spline 5	5.77	6.72	2.91	5.40
	(0.939)	(5.33)	(0.994)	(1.47)
Sum Competitors Capacity B-spline 6	7.3	8.15	2.11	4.23
	(0.944)	(5.67)	(1.15)	(1.50)
Own Capacity B-spline 1	-3.79	-3.97	0.374	-2.71
	(0.923)	(0.925)	(0.880)	(1.22)
Own Capacity B-spline 2	-3.3	-3.37	0.720	-0.754
	(0.893)	(0.894)	(0.902)	(1.22)
Own Capacity B-spline 3	-2.3	-2.51	1.06	-0.325
	(0.967)	(0.969)	(1.04)	(1.28)
Own Capacity B-spline 4	-1.72	-1.76	1.87	-0.149
	(0.943)	(0.952)	(1.27)	(1.60)
Own Capacity B-spline 5	-2.63	-2.66	2.05	3.32
	(1.35)	(1.35)	(2.25)	(2.02)
Population B-Spline 1		-5.11		
		(6.78)		
Population B-Spline 2		0.886		
		(5.16)		
Population B-Spline 3		-1.39		
		(5.52)		
Population B-Spline 4		-0.008		
		(5.06)		
Population B-Spline 5		-1.60		
		(6.69)		
Capacity is Per-Capita	No	No	Yes	Yes
Region Fixed Effects	No	No	No	Yes
Adjusted R^2	0.8952	0.8955	0.8816	0.8946
Band σ^2	1.40	1.40	1.56	1.41

Dependent variable is the natural log of the change in capacity. Number of capacity changes = 774. Parameters estimated using OLS. Capacity is measured in thousands of tons per year. Population is denominated in tens of millions.

although they are estimated with significant amount of variance. While this specification fits better than the per-capita model without fixed effects, it still is not as good a fit as specification I, which hereafter is the preferred empirical specification.

Table 8 reports four specifications of the equation for the level a firm desires to adjust to given it is going to make an investment. The positive news is that the fit in all the specifications is extremely tight; the lowest adjusted R^2 is 0.9958. The estimated variance of the error term in all specifications is also very low. These statistics suggest that the model is capable of accurately fitting the capacity levels that firms adjust to quite tightly as a flexible function of competitors' aggregate capacity and a firm's own capacity. As in the band equation, I test several specifications. In the baseline specification, the explanatory variables are b-splines of aggregate competitor capacity and own capacity. The results suggest that target values are strictly increasing across the range of competitor capacities seen in the data. The function is negative with respect to firm's own capacity, although the coefficients are decreasing in the spline, which suggests that larger firms prefer to make larger adjustments.

This result may be due to the fact that larger firms operate in larger markets, and therefore the residual demand curve is larger. To test this hypothesis, specification II includes b-splines of market population. The coefficients on these variables are positive, which supports the idea that larger markets support higher investment. However, the coefficients on own capacity are virtually exactly the same, suggesting that the pattern of larger firms having larger target levels of capacity holds when controlling for market size. I also note that the coefficients on the aggregate competitor capacity decrease by roughly the size of the population variables. Specification III includes region fixed effects, to test for market-level heterogeneity not captured otherwise. The fixed effects are very imprecisely estimated. The overall effect is to push the coefficients on aggregate competitor capacity back toward the their levels in specification I while weakening the effect of own capacity. The population effects are rendered statistically insignificant, implying that market-specific variation in target levels is largely captures by the fixed effects.

To examine the effect of population further, specification IV runs the regression with per-capita capacities instead of levels. Of note is that the average population is about 1.14 when normalized to be in tens of millions of people, which helps explain why the coefficients on the capacity variables are very close to the previous specifications. The pattern of coefficients looks very similar to the first three specifications, while not fitting the data as well; the regression error an order of magnitude larger.

Since the first three specifications all provide an excellent fit to the data, and specification

Table 8: Investment Policy Function Results: Investment Target

	I	II	III	IV
Sum Competitors Capacity B-spline 1	7.74	5.80	7.26	7.094
	(0.124)	(0.714)	(0.927)	(0.282)
Sum Competitors Capacity B-spline 2	7.70	5.67	7.06	6.96
	(0.123)	(0.715)	(0.929)	(0.326)
Sum Competitors Capacity B-spline 3	7.76	5.80	7.17	7.50
	(0.120)	(0.711)	(0.936)	(0.303)
Sum Competitors Capacity B-spline 4	7.64	5.65	6.60	6.50
	(0.127)	(0.719)	(0.964)	(0.334)
Sum Competitors Capacity B-spline 5	7.88	5.96	6.82	7.31
	(0.124)	(0.701)	(0.987)	(0.340)
Sum Competitors Capacity B-spline 6	7.59	5.52	6.36	6.71
	(0.124)	(0.746)	(0.992)	(0.391)
Own Capacity B-spline 1	-2.24	-2.24	-2.15	-0.912
	(0.121)	(0.121)	(0.124)	(0.301)
Own Capacity B-spline 2	-1.36	-1.36	-1.31	-0.136
	(0.118)	(0.118)	(0.124)	(0.308)
Own Capacity B-spline 3	-0.752	-0.753	-0.702	-0.762
	(0.128)	(0.128)	(0.130)	(0.354)
Own Capacity B-spline 4	-0.182	-0.186	-0.120	1.27
	(0.124)	(0.125)	(0.134)	(0.432)
Own Capacity B-spline 5	0.074	0.0096	0.063	-0.831
	(0.179)	(0.178)	(0.181)	(0.767)
Population B-Spline 1		1.43	0.482	
		(0.892)	(2.30)	
Population B-Spline 2		2.08	0.483	
		(0.679)	(0.876)	
Population B-Spline 3		1.95	0.656	
		(0.727)	(1.04)	
Population B-Spline 4		1.98	0.015	
		(0.666)	(0.802)	
Population B-Spline 5		2.37	-0.566	
		(0.881)	(1.21)	
Capacity is Per-Capita	No	No	No	Yes
Region Fixed Effects	No	No	Yes	No
Adjusted R^2	0.9994	0.9994	0.9995	0.9958
Estimated σ^2	0.0244	0.0242	0.0235	0.184

Dependent variable is log of capacity level after adjustment. Number of capacity changes = 774. Parameters estimated using OLS. Capacity is measured in thousands of tons per year. Population is denominated in tens of millions.

II and III have a number of statistically insignificant coefficients, I hereafter proceed using specification I as the empirical model for the target level of adjustment. In the estimation that follows I estimate the policy functions separately for before and after 1990. The results are not notably different than those reported in Tables 7 and 8.

Entry and Exit Policy Several estimated specifications of the entry and exit policy function results are presented in Table 9. I estimated both policy functions both in absolute levels and in per-capita levels, to control for unobserved market-level variation in demand that could change the policy functions. Specifications I and II estimate the exit policy in levels, with and without controls for population. The sign and magnitude of the estimated coefficients are very close in both specifications. Own capacity decreases the probability of exit, and an increase competitors' capacity increases that probability. Both explanatory variables proxy for the level of residual demand facing the firm when it makes an exit decision. The constant is near -1 in both regressions and the dummy variable for post-1990 is -0.595 and -0.607, respectively. To put these numbers in context in specification I, moving a firm into the post-1990 environment has the same effect as either increasing its own capacity by 368 thousand tons per year. The variable with the most explanatory power is the post-1990 dummy variable; the marginal effect in specification I, evaluated at the means of the other explanatory variables, is to decrease the probability of a given firm's exit from 2.1 percent to 4 tenths of a percent, a fivefold decrease. This directly reflects the overall exit rates in the industry: there were 51 exits in the period before 1990 and six exits after, a difference of 81 percent.

The effect of competitor's capacity is estimated to be positive and small under both specifications; the marginal effect is so small as to be economically unimportant. The addition of population as a control to the exit equation improves the fit of the model only marginally; the likelihood ratio test fails to reject the null hypothesis that the coefficient on population is equal to zero, as it estimated very imprecisely. As a result, the marginal effects in specification II are very similar to specification I.

One could be concerned that these specifications fail to adequately capture factors that influence the residual demand curve. To guard against this, I also estimated exit policies that were functions of per-capita capacity, with and without region fixed effects. The results are shown in columns III and IV. The per-capita results roughly mirror the ones above—the post-1990 dummy still dominates the effects of the other two variables. The relative marginal effect of the post-1990 shifter is stronger, own-capacity weaker, and competitor's capacity

Table 9: Entry and Exit Policy Results

Specification	I	II	III	IV
Exit Policy				
Own Capacity	-0.0015661	-0.0015795		
	(0.000268)	(0.0002712)		
Competitors Capacity	0.0000456	0.0000379		
	(0.0000173)	(0.0000249)		
Population		0.0590591		
		(0.1371835)		
After 1990	-0.5952687	-0.606719	-0.6328867	-0.4623664
	(0.1616594)	(0.1639955)	(0.157673)	(0.1910193)
Own Capacity per Capita			-0.0005645	-0.0010199
			(0.0001255)	(0.0002164)
Competitors Capacity per Capita			0.0000744	0.0002379
			(0.00000286)	(0.0001023)
Constant	-1.000619	-1.019208	-1.664808	-1.529715
	(0.1712286)	(0.176476)	(0.1475588)	(0.3526938)
Region Fixed Effects	No	No	No	Yes
LLH	-227.21	-227.12	-238.54	-217.38
Entry Policy				
Competitors Capacity	0.0000448	-0.0003727		
	(0.0000365)	(0.0002351)		
After 1990	-0.6089773	-0.8781589	-0.602279	-1.003239
	(0.2639545)	(0.3229502)	(0.2651052)	(0.337589)
Constant	-1.714599	-0.454613	-1.665322	-0.3434765
	(0.2152315)	(0.7086509)	(0.2642566)	(0.6624767)
Competitors Capacity per Capita			0.000026	-0.0003633
			(0.000038)	(0.0001766)
Region Fixed Effects	No	Yes	No	Yes
LLH	-70.01	-56.47	-70.491	-55.53
Prob > χ^2	0.0177	0.4516	0.0287	0.3328

Sample size for exit policy function = 2233; sample size for entry policy function = 414. Capacity is measured in thousands of tons of cement per year. Population is normalized to be measured in tens of millions. Per capita capacity is measured as thousands of tons per year per tens of millions in population.

stronger in the per-capita model. The inclusion of region fixed effects improves the fit of the model, although the likelihood ratio test fails to reject the hypothesis that the coefficients on the fixed effects are jointly zero. I also estimated a model where the competitors' per-capita aggregate capacity was expanded to a fourth-degree orthogonal polynomial. The results of this model were almost identical to those from specification III. The inclusion of a time trend did not change the results. The fit of specification III, as measured by the improvement in the likelihood function after including covariates, is not as good than specification I, so we take specification I to be our preferred empirical model.

I also estimated several specifications of the entry policy functions, shown in the bottom panel of Table 9. The baseline rate of entry is low, as accounted for by the constant, which estimated to be negative in all specifications. The post-1990 dummy is negative and significant in all specifications. Analogous to the exit policy function, this reflects with the empirical trends for entry; there were 15 entries in the period before 1990 and four entries after the passage of the Amendments in 1990. The signs on incumbent capacity, whether in levels or per-capita, is positive, small, and statistically insignificant in specifications without region fixed effects. The inclusion of fixed effects flips the sign on incumbent capacity negative, as expected, but a likelihood ratio test fails to reject the hypothesis that all coefficients are equal to zero. I have reported the p-values for this test below each specifications log-likelihood value in the last row. The only model that is not rejected at the two percent level is specification I, which estimates entry as a function of incumbent capacity. The per-capita specification is modestly less significant than specification I; therefore I use the model in levels for the remainder of the analysis.

Dynamic Parameters The results of the second step estimation described above are presented in Table 10. The b-spline coefficients and their projection onto underlying distributions are both reported. The fixed costs of investment are significant, reported at $847,000 in the pre-1990 period and more than doubling to $1.8M in the post-1990 period. The fixed costs of adjustment are relatively small next to the variable costs of investment for a typical plant. The early and late estimates for the marginal cost of adjustment are very close; $235 per ton before 1990 and $233 per ton after 1990. These costs imply that a 1.5M plant would cost about $350M, which is a reasonable figure. Both of these estimates are in the same range as the accounting estimates of $200 per ton reported in Salvo (2005). This is fairly remarkable and a testament to the power of the MPNE framework give that these costs are inferred without any direct observation of investment costs in the data. It is also interesting

Table 10: Dynamic Parameters

Parameter	Early	Late
Investment B-Spline 1	-725	-937
Investment B-Spline 2	68	77
Investment B-Spline 3	761	1,243
Investment B-Spline 4	768	1,655
Investment B-Spline 5	863	2,022
Investment B-Spline 6	2,662	4,258
Investment Cost	235	233
Investment Cost Squared	0	0
Divestment B-Spline 1	-4	-15
Divestment B-Spline 2	39,078	38,092
Divestment B-Spline 3	41,271	38,092
Divestment B-Spline 4	149,364	169,672
Divestment B-Spline 5	371,211	302,092
Divestment B-Spline 6	1,135,363	1,082,398
Divestment Cost	-98	-238
Divestment Cost Squared	3,091	4,394
Exit B-Spline 1	263,442	261,375
Exit B-Spline 2	238,384	236,217
Exit B-Spline 3	2,536	1,942
Exit B-Spline 4	-75,180	-72,254
Exit B-Spline 5	-84,802	-72,254
Exit B-Spline 6	-93,061	-77,138
Implied Distributions		
Investment distribution μ	847	1,798
Investment distribution σ	130	420
Divestment distribution μ	235,351	222,524
Divestment distribution σ	212,707	171,569
Exit distribution μ	-70,881	-67,490
Exit distribution σ	58,000	55,167

Point estimates and confidence intervals were obtained using 125 simulated outcomes of 4 firms with 200 year lifetimes each. The initial state was held constant across all simulations. Each simulation path was replicated 5000 times and averaged to obtain expected values. The estimated mean variable cost of a 1,500,000 TPY investment is $352M before the Amendments and $349.5M after.

to note that these costs are essentially unchanged across the two periods, which is in line with the theoretical treatment that the Amendments only changed the sunk costs of entry during the period of time I observe.

The fixed and variable costs of divestment are estimated to be very large, in the sense that firms will almost never find it reasonable to sell off capacity. This reflects the paucity of downward capacity adjustments observed in the data. The numbers for before and after 1990 are very close, suggesting that the Amendments did not have a significant influence on divestment costs, as expected.

Finally, the distribution of exit costs is estimated to have a very low mean and fairly large standard deviations. This combination means that most firms will not find it worthwhile to exit unless they receive a very favorable draw from this distribution, as the estimated profits of remaining active are significantly positive even in the most contested markets. This is to be expected given the low exit rate of firms in this industry, particularly after 1990. Although the mean of the exit cost distribution shifts up after 1990, the standard deviation decreases. The combination of these two factors implies that exit continues to be a rare event after 1990. For example, the probability that a firm will receive a draw on the exit costs greater than $75M is 0.59 percent before 1990 and 0.49 percent after 1990. The corresponding probabilities for draws greater than $100M are 0.16 percent and 0.11 percent. These numbers are relatively close to each other, which helps support the notion that the changes in the cost structure due to the Amendments primarily influenced entry costs, although this finding suggests that there may have been a slight increase in scrapping costs.

The square of investment costs in the variable adjustment costs were set to zero after much experimentation. The reason is that the square tends to dominate the costs of investment for large changes; this implies that an entering firm building a reasonably large plant would face unreasonably large investment costs. This is an artifact of the fact that the quadratic adjustment costs are global, and most of the adjustments observed in the data do not span such large changes. I found the linear adjustment costs to give a much more reasonable results.

Distribution of Entry Costs I assume that the sunk costs of entry are independent draws from a normal distribution that is common across markets. I match the empirical probability of entry for a given state, given by the probit policy function, against the cumulative distribution function evaluated at the expected value of entering at that state. States

Table 11: Sunk Cost of Entry Distribution Results

Parameter	Mean (000 $)
Before 1990	
Mean (μ)	163,055
Standard Deviation (σ)	64,306
After 1990	
Mean (μ)	172,680
Standard Deviation (σ)	41,559

Parameters were estimated by matching the cumulative distribution function of a normal distribution to the empirical probabilities of entry. States were varied by the capacity of incumbent firms from 500,000 TPY to 3M TPY in 5,000 TPY increments. The expected value of entry was computed using 250 replications at each state.

were varied by the capacity of incumbent firms from 500,000 tons per year to 3M tons per year in 5,000 tons per year increments. The expected value of entry was computed using 250 replications at each state. The results of the estimation are presented in Table 11. One of the primary results of this paper is that I find the Amendments increased the sunk costs of entry. The mean of the entry cost distribution increased by a little less than 6 percent while the variance decreased by approximately 35 percent. These two shifts work together to significantly decreases the chance of a firm receiving a small enough draw on the sunk cost of entry to warrant building a new facility. For example, the probability that a firm receives a draw on the entry costs below $50M in the first period is about 4 percent; after the Amendments, the probability drops to 0.16 percent. If the threshold is raised to $100M, the corresponding probabilities are 16 percent and 4 percent, respectively. This is relevant because at the margin, the last entering firm can expect to make a relatively small amount of money in present value terms. Even when the firm is facing a large expected surplus conditional on entering, the shift in entry costs after 1990 greatly reduces the probability that a firm will find it optimal to do so. To emphasize, this change in the cost structure is the single most important determinant of the shift in market structure after 1990. As I show in the counterfactual simulations below, the increase in entry costs greatly reduces the chances that marginal firms enter a market, and this has significant effects on product market competition.

8 Policy Experiments

The benefit of estimating a structural model is the ability to simulate counterfactual policy experiments once a researcher knows the underlying primitives. My primary interest is to evaluate the welfare costs of the Amendments, so a natural investigation is to determine the differences across policy regimes for quantities of economic interest, including welfare measures for both producers and consumers. To achieve this, I compute the MPNE of the theoretical model with two sets of parameters: the observed post-Amendments cost structure, and the post-Amendments cost structure with the distribution of sunk entry costs taken from before the regulation.[32] It should be emphasized that these welfare calculations ignore the primary intended benefits of the Amendments: improved social welfare through cleaner air and it associated benefits. The results here should be interpreted in that light as a view on the changes in welfare due to changes in market structure as a result of the environmental regulation, and as such is only a part of the overall welfare changes of the 1990 Amendments to the Clean Air Act.[33]

With policy functions from these equilibria it is possible to simulate hypothetical markets given some starting configuration. I compute the distribution of producer profits and consumer surplus under two different starting states: a new market with no incumbent firms and four potential entrants, and a market with two incumbents and two potential entrants. I take the parameters of demand for this market from Alabama, which is close to a representative market. Ideally, one would solve out for the equilibrium of every market in the US and simulate welfare changes for each one. Computational constraints, however, prevent this approach, and I have to restrict the number of active firms to be four, which is the

[32]It is well known that these models potentially have many equilibria, some of which are not discoverable without sophisticated methods (see Doraszelski, et. al. (2008) for details). In the absence of any guarantees that the MPNE found here is unique, I report the solution method used to find the equilibrium so that my results will be reproducible. I used a modified policy iteration scheme: in a first step the value function is iterated to convergence given the vector of policy functions, and in a second step the policy functions are updated. This process was stopped when the change in the norm of the value and policy functions was less than 1E-5—a sufficiently small level that changes from step to step only occurred in the fifth decimal place in the policy functions. I found that dampening the policy update step, which restricted each element of the new policy vector to change by at most one percent, ensured the model always converged. I discretized the state space into capacity increments of 750,000 tons per year, the typical size of a new kiln in this industry, and capped the state space at 2.25 million tons per year, a level at which no firms find it profitable to invest beyond. Experiments with smaller capacity increments did not materially change the results when the number of firms was smaller; I chose 750,000 as this allowed me to solve the model with four firms in a reasonable amount of time.

[33]The integration of both views of the Amendments is an interesting research question that I am pursuing in related work.

Table 12: Counterfactual Policy Experiments

	Post-Amendments (High Sunk Costs)	Counterfactual (Low Sunk Costs)	Social Planner (Low Sunk Costs)
New Market			
Producer profit	5,724	18,600	-394,333
Consumer welfare	1,913,794	2,007,275	4,991,057
Periods with no firms	1	1	1
Periods with one firm	60	7	0
Periods with two firms	88	35	0
Periods with three firms	45	100	0
Periods with four firms	5	57	199
Total welfare	1,919,518	2,025,875	4,596,724
Profits of firm 1	6,916	14,731	-394,576
Average size of active firm	1,576	1,169	2,239
Average market capacity	2,744	3,355	8,955
Average market quantity	2,385	2,915	7,786
Average market price	70	65	47
Market with Two Incumbents			
Producer profit	200,865	197,909	-159,022
Consumer welfare	2,627,188	2,682,207	5,212,515
Periods with no firms	0	0	0
Periods with one firm	0	0	0
Periods with two firms	197	69	1
Periods with three firms	3	130	0
Periods with four firms	0	1	199
Total welfare	2,828,052	2,880,116	5,053,493
Profits of firm 1	506,009	497,605	179,582
Average size of active firm	1,494	1,334	2,244
Average market capacity	3,008	3,490	8,966
Average market quantity	2,615	3,035	7,796
Average market price	68	65	47

Industry distributions were simulated along 25,000 paths of length 200 each. All values are present values denominated in thousands of dollars. The new market initially has no firms and four potential entrants. The incumbent market starts with one 750,000 TPY incumbent and one 1.5M TPY incumbent and two potential entrants. Counts of active firms may not sum to 200 due to rounding off.

median size of a cement market in the United States. While this is restrictive, the results with four firms indicate that the possibility of a fifth firm entering this market is very low. It is therefore reasonable to conclude that this restricted specification captures many of the essential dynamics of the median market.

Table 12 presents the results of the counterfactual simulations. In the case of a new market, where the initial state vector is four empty slots waiting for entrants, overall welfare has decreased significantly due to the Amendments, declining by approximately $100M in present value, a little over five percent. While artificial, the new market serves as a natural bound for the upper limit of welfare damages; as it is the market configuration that would be most affected by a change in sunk entry costs. Indeed, the driving factor for changes in welfare across both simulated markets is the change in entry rates. With the higher sunk costs of entry, the distribution of the number of active firms is shifted down. The market is now much more likely to have just one or two active firms (74 percent), as compared to the counterfactual scenario with the pre-1990 cost structure (21 percent). Market configurations with four firms, in particular, are much more rare after 1990, appearing in only 2.5 percent of simulated periods versus 27.5 percent under pre-1990 costs. This compression of the firm distribution has significant effects on market outcomes. Although the average active firm is larger by about 34 percent, reflecting the higher rate of return on investment given that firms expect softer competition in the future, prices are 7 percent higher and quantities are 18 percent lower. The lower number of firms does not translate into better outcomes for producers, for whom profits decline from $18.6M to $5.72M. This is an interesting result, as one may expect that higher costs would drive a wedge between the firms that actually enter and those that do not, increasing rents, but that is not the case in equilibrium. This is due to both the increased direct cost of entering the industry and the fact that firms are willing to enter at higher draws of entry costs knowing exactly that they will face fewer competitors in the future. In this sense, they compete away the potential projected oligopoly surplus induced by higher costs. Loss of welfare in this market is driven almost exclusively by losses of the buyers: consumer surplus decreases by almost five percent in this setting, a loss of $93.4M.

The second market I consider has two incumbents with capacities of 750,000 TPY and 1.5M TPY. While the new entry market is an extreme case of what could have happened under the Amendments, a market with incumbents of over 2M TPY capacity is a close approximation of a mature, fully capitalized cement market of average size in the United States. As such, this should provide a lower bound to welfare penalties, as this market will be least affected by a change in entry rates. The primary effect, as in the new market, is the marked decrease in entry. The number of periods with three or four firms decreases by almost 98 percent under the higher costs. The average size of an active firm is bigger in this case, again reflecting the fact that firms can recoup more of their investment costs with reduced product-market competition. However, this expansion in size is not sufficient to offset the competitive effects of additional firms: average prices increase by 4.6 percent and average quantities decrease by 13.8 percent. As a result, consumers surplus is reduced by approximately $55M, a 2 percent decrease. As predicted by theory, producers enjoy a modest surplus increase of 1.5 percent, or about $3M per firm, under the higher entry costs. Note that this implies that if the costs of obtaining the operating permits was lower than $3M for the average firm, then incumbents are actually better off under the Amendments than before 1990. In this case, the static analysis of the engineering costs would not only ignore the dynamic costs to consumers, but also obtain welfare costs to suppliers of the wrong sign. The overall change in surplus in this market is a decrease of $52M, a little under two percent.

Extrapolating these costs to the entire US, under the assumption welfare losses can be summed equally across all 27 markets, leads to an estimate of over $1.4B as a lower bound. The corresponding upper bound, $2.7B, clearly has little merit when extrapolated to the entire US, as it would be an estimate of the costs of starting the entire industry over from scratch under the two different sunk cost distributions. However, both numbers suggest that the welfare costs of the Amendments were significant, primarily through the reduction in product market competition. This result should be viewed carefully, however, as the reduction in output also reduces emissions in the short-run. In this sense, the negative consequences of environmental regulation through restricted competition in the product market are at least partially (and potentially more than) offset by reductions in emissions and their resulting welfare improvements.[34]

An interesting application of the structural model is to examine the differences between the oligopolist's MPNE to that of the social planner. The social planner sums the profits of all firms and consumer welfare and finds the optimal evolution of the industry given those state values, subject to the caveat that the benefits of reduced pollution are not incorporated in that calculation. I calculate the social planner's MPNE with the same cost structure as the oligopoly counterfactual. This solution gives an upper bound for the welfare losses under the regulation, as this would have been the best possible welfare outcome in the absence of the Amendments. The third column compares the social planner's solution to the oligopoly solution.

The key characteristic of the social planner's solution relative to the oligopolies is that it exploits lost welfare gains due to inefficiently low investment to increase overall welfare. The social planner is willing to inflict losses on the firms, through costly investments, in order to drastically increase consumer surplus. As a result, the average market size in both starting market configurations is approximately three times larger under the social planner than oligopoly. Prices are a third as much as under the oligopoly, and overall welfare roughly doubled. In the new market, the social planner solves the maximization problem by having four extremely large firms. This is an interesting consequence of the distribution of fixed

[34]Demonstrating the magnitude of these offsets is a complicated question beyond the scope of the present paper that I am pursuing in ongoing related research.

costs—if the entry costs had no variance, the social planner builds one extremely large firm to minimize the fixed costs of entry and investment. However when entry costs are stochastic the social planner finds it more worthwhile to set a threshold at which it is efficient for four firms to usually enter immediately. These results suggest that the oligopoly solution is not very close to the social optimum, even given some starting capacity, due to firms failing to internalize the social benefits of their capacity investments.

One very strong assumption made here is that demand is not growing over time. It is difficult to assess what the effects of demand growth would be in this dynamic setting. As we saw with the new entrant market, strategic competition in entry and investment may undo some of the intuition that we have about the effects of changing part of the cost structure. Growing demand makes the future more valuable relative to the world where it is not growing. One would expect this to increase the intensity of competition in entry, all else equal. On the other hand, it may be that firms that actually enter the market will do so at such a large size that this more than offsets the increased incentives for entry. This is an interesting case that I plan to explore in more detail in future research, as computational techniques and raw computing horsepower allow us to explore more complex state spaces than those considered here. In any case, the long-run effects may not be particularly pronounced in the United States, as many domestic cement markets appear to be relatively stable with respect to growth, as opposed areas of the world such as China where cement demand is booming.

9 Conclusion

In this paper, I have estimated the welfare costs of the 1990 Amendments to the Clean Air Act on the Portland cement industry. My principal finding is that a static analysis of the costs of the regulation will not only underestimate the costs to consumers, but will actually obtain estimates of the wrong sign for incumbent firms. Exploiting the timing structure of the implementation of the Amendments, I identify that the most significant economic change in the Portland cement industry was a large increase in the sunk costs of entry. As a result of lower entry rates, overall welfare decreased by at least \$1.4B. These results highlight the importance of estimating the welfare consequences of regulation using a dynamic model to account for all relevant changes to the determinants of market structure. A static model would also be incapable of calculating the counterfactual benefits to producers of paying higher entry costs but facing lower ex-post competition. The estimates that the certification process would at most cost \$5M per installation would underpredict the welfare costs by at least \$900M.

I find that the (S, s) investment rule is a flexible and powerful method for characterizing the lumpy investment behavior of firms, as these choices are partially governed by significant fixed adjustment costs. I am able to recover estimates of the investment costs which are consistent with out-of-sample accounting estimates. This is a natural omnibus check on the fit of the model, as the estimates of the investment costs are dependent on every aspect of the model save the fixed costs. The interplay between market power, investment, and production choice is particularly interesting. For smaller markets, firms find it optimal to produce near the socially efficient level due to capacity constraints. However, private investment is too low relative to the social planner. As a result, the oligopoly outcomes are far inside the socially optimal frontier.

An interesting extension of the present work would be to examine the effects of a "cap-and-trade" market-based emissions control program, similar to the trading program for SO_2 in the electricity industry. In this environment the regulatory authority removes all specific point-source control requirements and instead places an overall cap on the level of emissions in a regional area. Firms are endowed with pollution rights that they are free to trade among each other. This type of policy has the benefit of achieving the most efficient configuration of production within the industry for a given level of pollution. However, it may have other consequences with respect to market power and the concentration to pollution to a subset of firms within the market. By coupling emissions data, reported in many states at the yearly level, to production data one can back out a pollution production function. One question I can then address is whether efficiency obtains in this environment, as some of the more inefficient firms may buy pollution rights in return for additional market power. There are clearly a number of other interesting dynamic questions in this framework, from the nonlinear health effects of pollution concentration to the investment incentives of heterogeneous firms in a region, that are left for future research.

References

United States Environmental Protection Agency. NSR 90-day review background paper, June 2001. Docket A-2001-19, Document II-A-01.

Victor Aguirregabiria and Pedro Mira. Sequential simulation-based estimation of dynamic discrete games. Boston University Working Paper, 2004.

Orazio Attanasio. Consumer durables and inertial behavior: Estimation and aggregation of (S, s) rules for automobile purchases. *Review of Economic Studies*, 67:667–96, 2000.

Patrick Bajari, C. Lanier Benkard, and Jonathan Levin. Estimating dynamic models of imperfect competition. *Econometrica*, 75:1331–70, September 2007a.

Patrick Bajari, Han Hong, and Stephen P. Ryan. Identification and estimation of discrete games of complete information. University of Minnesota Working Paper, 2007b.

C. Lanier Benkard. A dynamic analysis of the market for widebodied commercial aircraft. *Review of Economic Studies*, 71:581–611, 2004.

B. Douglas Bernheim and Michael D. Whinston. Multimarket contact and collusive behavior. *RAND Journal of Economics*, 21(1):1–26, Spring 1990.

Steve Berry, Ariel Pakes, and Michael Ostrovsky. Simple estimators for the parameters of dynamic games (with entry/exit examples). Harvard University Working Paper, June 2005.

David Besanko and Ulrich Doraszelski. Capacity dynamics and endogenous asymmetries in firm size. *RAND Journal of Economics*, 35(1):23–49, Spring 2004.

Ron. N. Borkovsky, Ulrich Doraszelski, and Yaroslav Kryukov. A user's guide to solving dynamic stochastic games using the homotopy method. Harvard University Working Paper, 2008.

United States Census Bureau, 1977–2002.

X. Chen. Large Sample Sieve Estimation of Semi-Nonparametric Models. *Handbook of Econometrics*, 6, 2006.

Victor Chernozhukov and Han Hong. A MCMC approach to classical estimation. *Journal of Econometrics*, 115(2):293–346, 2003.

Ulrich Doraszelksi and Mark Satterthwaite. Foundations of markov-perfect industry dynamics: Existence, purification, and multiplicity. Harvard University Working Paper, 2005.

Richard Ericson and Ariel Pakes. Markov perfect industry dynamics: A framework for empirical work. *Review of Economic Studies*, 62(1):53–82, 1995.

Chaim Fershtman and Ariel Pakes. A dynamic game with collusion and price wars. *RAND Journal of Economics*, 31(2):207–36, 2000.

Meredith Fowlie. Incomplete environmental regulation, imperfect competition, and emissions leakage. University of Michigan Working Paper, 2008.

G. Gowrisankaran and R. Town. Dynamic equilibrium in the hospital industry. *Journal of Economics and Management Strategy*, 6(1):45–74, 1997.

George Hall and John Rust. The (S,s) rule is an optimal trading strategy in a class of commodity price speculation problems. Yale University Working Paper, 2000.

J. Hotz, R. Miller, S. Sanders, and J. Smith. A simulation estimator for dynamic models of discrete choice. *Review of Economic Studies*, 61:256—289, April 1994.

I. Jans and D. I. Rosenbaum. Multimarket contact and pricing: Evidence from the U.S. cement industry. *International Journal of Industrial Organization*, 15:391–412, 1997.

Erin Mansur. Environmental regulation in oligopoly markets: A study of electricity restructuring. Yale SOM Working Paper, 2004.

Eric Maskin and Jean Tirole. A theory of dynamic oligopoly, i: Overview and quantity competition with large fixed costs. *Econometrica*, 56:549–69, May 1988.

Whitney K. Newey and James L. Powell. Instrumental variable estimation of nonparametric models. *Econometrica*, 71(5):1565–1578, September 2003.

Martin Pesendorfer and Philip Schmidt-Dengler. Identification and estimation of dynamic games. LSE Working Paper, 2003.

G. Rothwell and J. Rust. On the optimal lifetime of nuclear power plants. *Journal of Business and Economic Statistics*, 15:195–208, 1997.

John Rust. Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica*, 55(5):999–1033, September 1987.

John Rust. Numerical dynamic programming in economics. In H. Amman, D. Kendrick, and J. Rust, editors, *Handbook of Computational Economics*. North Holland, 1994.

Stephen P. Ryan and Catherine E. Tucker. Diversification and the dynamics of technology adoption. MIT Working Paper, 2006.

Alberto Salvo. Inferring conduct under the threat of entry: The case of the Brazilian cement industry. LSE Working Paper, 2005.

H. E. Scarf. The optimality of (s,S) policies in the dynamic inventory problem. In K. J. Arrow, S. Karlin, and P. Suppes, editors, *Chapter 13 in Mathematical Methods in the Social Sciences*. Stanford University Press, 1959.

Hendrik G. van Oss. *Mineral Commodity Survey*. U.S. Geological Survey, January 2001. Cement section.