EI @ Haas WP 234R

Market-Based Emissions Regulation and Industry Dynamics

Meredith Fowlie, Mar Reguant, and Stephen P. Ryan

Revised April 2014

Revised version published in the *Journal of Political Economy* 124(1), 249-302, 2016.

Energy Institute at Haas working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to review by any editorial board.

© 2014 by Meredith Fowlie, Mar Reguant, and Stephen P. Ryan. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit is given to the source.

http://ei.haas.berkeley.edu

Market-Based Emissions Regulation and Industry Dynamics

Meredith Fowlie, Mar Reguant, and Stephen P. Ryan*

April 28, 2014

Abstract

We assess the long-run dynamic implications of market-based regulation of carbon dioxide emissions in the US Portland cement industry. We consider several alternative policy designs, including mechanisms that use production subsidies to partially offset compliance costs and border tax adjustments to penalize emissions associated with foreign imports. Our results highlight two general countervailing market distortions. First, following Buchanan (1969), reductions in product market surplus and allocative inefficiencies due to market power in the domestic cement market counteract the social benefits of carbon abatement. Second, trade exposure to unregulated foreign competitors leads to emissions "leakage" that offsets domestic emissions reductions. Taken together, these forces result in social welfare losses under policy regimes that fully internalize the emissions externality. In contrast, market-based policies that incorporate design features to mitigate the exercise of market power and emissions leakage can deliver welfare gains when damages from carbon emissions are high.

1 Introduction

In the absence of a coordinated global agreement to curtail greenhouse gas emissions, regional market-based climate change policy initiatives are emerging. Examples include the

*UC Berkeley and NBER, Stanford GSB and NBER, and the University of Texas at Austin and NBER. We thank the editor, Nick Buchholz, Brian Copeland, Larry Goulder, Tom Holmes, Ryan Kellogg, Alberto Salvo, James Schmitz, numerous seminar participants, and three anonymous referees for very useful comments. We thank Erica Myers and Louis Preonas for excellent research assistance. The authors gratefully acknowledge the support of NSF grant SES-0922401.

Emissions Trading Scheme (ETS) in the European Union and California's greenhouse gas (GHG) emissions trading program. In these "cap-and-trade" (CAT) programs, regulators impose a cap on the total quantity of emissions permitted and distribute a corresponding number of tradeable emissions permits. To mitigate potentially adverse competitiveness impacts, and to engender political support for the program, it has become standard to allocate some percentage (or all) of these emissions permits for free to industrial stakeholders (Joskow and Schmalensee, 1998; Hahn and Stavins, 2010). In this paper, we explore both the static and dynamic implications of several different permit allocation mechanisms.

A particularly appealing quality of the cap-and-trade approach to regulating industrial emissions is that, provided a series of conditions are met, an emissions trading program designed to equate marginal abatement costs with marginal damages will achieve the socially optimal outcome (Dales, 1968; Montgomery, 1972).[1] Unfortunately, policy makers do not work in first-best settings where the conditions required for optimality are always satisfied. Real-world policy settings are typically characterized by several pre-existing distortions that complicate the design of efficient policy. In this paper, we focus on two distortions in particular.

First, many of the industries currently regulated under existing and planned emissions regulations are highly concentrated. In a seminal paper, Buchanan (1969) argues that a first-best policy designed to completely internalize external damages should be used only in "situations of competition" because concentrated industries are already producing below the socially-optimal level, and the loss of consumer and producer surplus induced by further restricting output can overwhelm the gains from emissions mitigation. An important counterpoint is offered by Oates and Strassmann (1984), who argue that the welfare gains from a Pigouvian tax (or a first-best cap-and-trade program) will likely dwarf the potential losses from non-competitive behavior. There has been surprisingly little work done to empirically investigate this trade-off between incentivizing pollution abatement and exacerbating the pre-existing distortion associated with the exercise of market power in concentrated industries subject to emissions regulations.

Second, regional climate change policies are textbook examples of "incomplete" regulation. When an emissions regulation applies to only a subset of the sources that contribute to the environmental problem, regulated sources can find it more difficult to compete with producers operating in jurisdictions exempt from the regulation. Shifts in production and

[1] Conditions include zero transaction costs, full information, perfectly competitive markets, and cost minimization behavior.

associated "emissions leakage" can substantially offset, or paradoxically even reverse, the reductions in emissions achieved in the regulated sector. This leakage is particularly problematic when emissions damages are independent of the location of the source, as is the case with GHGs.[2]

These complications have engendered a lively policy debate about how to design and implement climate change mitigation policies. Policy makers have been exploring several different approaches to (partially) compensating firms for their compliance costs via allocations of free emissions permits. Under a grandfathering regime, permits are freely distributed to regulated sources based on pre-determined criteria, such as historic emissions. Under so-called "dynamic updating," permits are allocated in proportion to firm's output in the previous period. Using emissions permits to incentivize production can mitigate product market surplus losses and reduce emissions leakage.[3]

Designing a policy that strikes the appropriate balance between curbing domestic GHG emissions and protecting the competitive position of emissions-intensive manufacturing sectors requires detailed knowledge of the structure and dynamics of the industries subject to the regulation. In this paper, we focus on an industry that has been at the center of the debate about U.S. climate change policy and international competitiveness: Portland cement. Cement is one of the largest manufacturing sources of domestic carbon dioxide emissions (Kapur et al., 2009).[4] The industry is highly concentrated, making the industry potentially susceptible to the Buchanan critique. Moreover, import penetration in the domestic cement market has exceeded 20 percent in recent years, giving rise to concerns about the potential for emissions leakage (Van Oss and Padovani, 2003; USGS, 2010).

A distinguishing feature of this paper is its emphasis on industry dynamics. We extend the dynamic oligopoly framework developed in Ryan (2012) as the foundation for our analysis. In our model, strategic domestic cement producers compete in spatially-segregated regional markets. Some of these markets are trade-exposed, whereas other landlocked markets are sheltered from foreign competition. Firms make optimal entry, exit, and investment decisions in order to maximize their expected stream of profits conditional on the strate-

[2]The damaging effects of greenhouse gas emissions are global; damages are a function of the level of emissions, but not the location. However, the same processes that generate GHG emissions also generate more locally-damaging co-pollutants. Accounting for the effects of these local co-pollutants is beyond the scope of this analysis.

[3]See also, Bernard et al. (2007) and Holland (2012).

[4]Carbon dioxide is the primary greenhouse gas emitted by industrial activities, but other greenhouse gases are also emitted. Because GHGs are typically measured in terms of carbon dioxide equivalents, we will use the terms "greenhouse gas" and "carbon dioxide" interchangeably.

gies of their rivals. Conditional on capital investments, producers compete each period in homogeneous quantities. Regional market structures evolve as firms enter, exit, and adjust production capacities in response to changing market conditions.

Our model is estimated using twenty five years of detailed data on the Portland cement industry. In the benchmark model we estimate, GHG emissions are unconstrained. We use this model to simulate the dynamic industry response to four counterfactual, market-based emissions policy designs: permit auctioning (isomorphic to a carbon tax), grandfathering, dynamic allocation updating, and a border tax adjustment which penalizes imports according to their average carbon content.

We begin by assuming that these policies will be designed such that the equilibrium permit price (or tax) is set equal to the assumed social cost of carbon (SCC) emissions. Under this assumption, we find that all four policy designs actually reduce net social welfare for SCC values below \$40 per ton of CO_2. Echoing Buchanan (1969), the combination of emissions leakage and welfare losses in the product market exceed the benefits of carbon mitigation. Losses are particularly acute for the auction/carbon tax scenario in which firms bear the full cost of compliance. Policy-induced disinvestment and exit further concentrates the ownership of productive capacity in the product market. The magnitude of the losses is substantial (\$18 billion under auctioning/carbon tax when the carbon value is \$30). Grandfathering helps slow the rate of firm exit, but does nothing to incentivize cement production. Consequently, grandfathering also results in substantial welfare losses at carbon values below \$60.

Policies that allocate free permits in proportion to production do substantially better because the implicit production subsidy mitigates both the exercise of market power in the product market and emissions leakage. At \$38 per ton of CO_2 (the current standard for monetizing the social costs of carbon emissions), these updating and BTA regimes are welfare improving. The BTA regime outperforms dynamic updating at high carbon values because the tax on imports more effectively mitigates emissions leakage and improves domestic terms of trade.

Consistent with the theory of the second best, policy outcomes could be improved if the social cost of carbon is only partially internalized by firms. Output-based, dynamic permit allocation updating essentially embeds this idea; firms are refunded a fraction of their compliance costs. We investigate these policy design trade-offs from an optimal taxation perspective. We solve for the optimal level of carbon prices, and the associated level of welfare gains, under the various regimes we consider. We find that these market-based policies can deliver welfare gains if the compliance costs (per ton of emissions) borne by

firms fall substantially below the true social cost of emissions.

This paper begins to address what Millimet et al. (2009) identify as a "striking gap in the literature on environmental regulation." Very little work has been done to bring recent advances in the structural estimation of dynamic models to analyses of industrial responses to environmental regulation. Our paper differs from past work in both the methods we use and the relationships we emphasize.[5] Our approach emphasizes dynamic industry responses to policy interventions, and the interplay between emissions regulations and pre-existing distortions associated with the exercise of market power in the cement market. We estimate an empirically tractable dynamic model of the U.S. cement sector in order to obtain estimates of key parameters such as investment costs. We contrast our dynamic policy simulations with a static modeling framework in which firms can alter production levels, but industry structure (i.e. technological characteristics, production capacities, etc.) is held fixed. These two modeling frameworks predict substantively different welfare effects, thus highlighting the role of dynamic processes in determining the long-run welfare effects of these environmental policies.

The paper is organized as follows: Section 2 introduces the conceptual framework for our applied policy analysis. Section 3 provides some essential background on the US Portland cement industry. We introduce the model and a detailed description of the alternative policy designs we consider in Section 4. We present the estimation and computational methodology in Section 5. Simulation results are summarized in Section 6. We conclude with a discussion of the results and directions for future research in Section 7.

2 Conceptual Framework

To build intuition for the basic economic forces at work in our empirical setting, we first present a simple, static model. Figure 1 shows a domestic monopoly producer (right panel) facing a competitive fringe of importers (left panel). The thick black, kinked line in the right panel represents the residual demand curve faced by a domestic monopolist. This curve is constructed by subtracting the import supply curve from the market aggregate demand

[5] A growing literature examines the impacts of emissions trading programs on highly concentrated, trade-exposed, and emissions-intensive industries. Several of these studies have assessed impacts of the EU ETS on European cement producers. For example, Szabo et al. (2006) and Demailly and Quirion (2006) use a bottom-up model of the cement industry to examine impacts of alternative policy designs on industry profits, emissions, and emissions leakage. More recently, Ponssard and Walker (2008) specify a static oligopoly model of a regional European cement industry to examine the short run responses of European cement producers to the ETS.

Figure 1: Emissions-Intensive, Trade-Exposed Monopoly

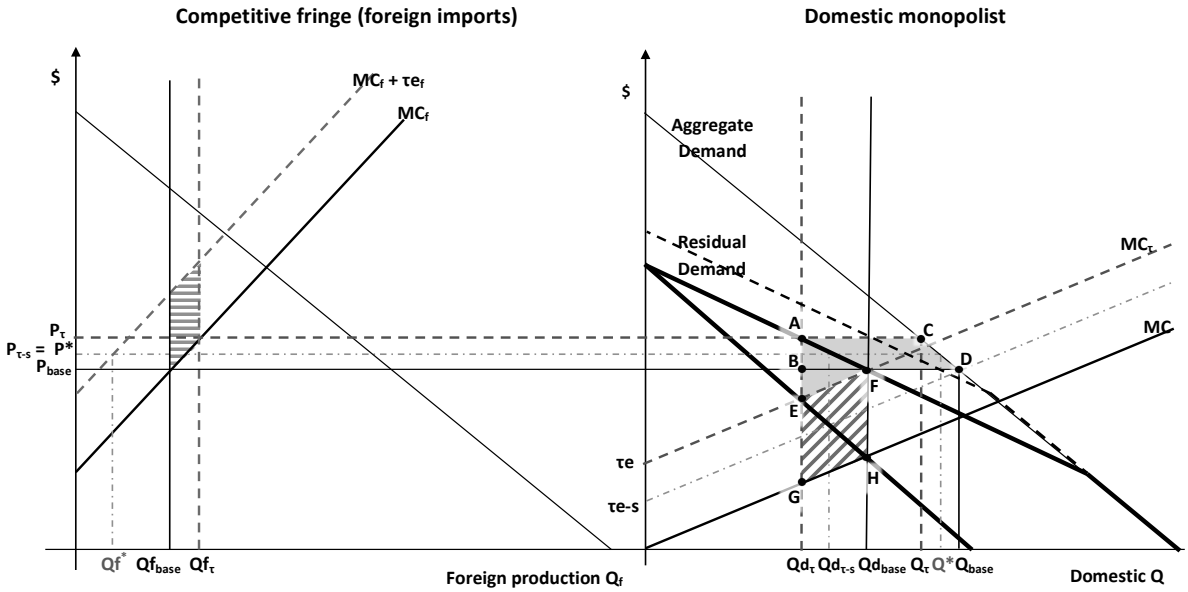

curve. The thick black line below it represents the corresponding marginal revenue curve.

Absent any emissions regulation, the domestic monopolist sets residual marginal revenue equal to marginal cost and produces output Qd_{base} at price P_{base}. Foreign producers supply Qf_{base} at this price. Total quantity, Q_{base}, is equal to $Qd_{base} + Qf_{base}$. This is the baseline against which we will compare the alternative policy outcomes. Note that the distortions associated with the exercise of market power in the domestic market manifest in two ways. First, the domestic firm restricts output in order to drive up the equilibrium product price. Second, production is not allocated optimally across domestic and foreign producers; marginal production costs differ across domestic and foreign producers.

Now suppose that production generates harmful emissions of a global pollutant. For ease of exposition, we assume a constant emissions rate per unit of output e and a constant marginal social cost of emissions τ across domestic and foreign production.[6] The curve labeled MC_τ captures both private marginal costs and the monetized value of the damages from the domestic firm's emissions: $MC_\tau = MC + \tau e$. Absent import competition, the socially optimal level of output would be defined by the intersection of MC_τ and aggregate demand.

Competition from foreign imports further complicates the picture. The broken line la-

[6]Note that this τ value is intended to capture global damages from greenhouse gas emissions (Working Group on Social Cost of Carbon, 2013).

beled $MC_f + \tau e_f$ represents the total social costs associated with foreign production. The downward sloping broken line in the right panel represents the residual demand curve that incorporates the emissions externality associated with foreign production. The intersection of this residual demand curve and MC_τ defines the socially efficient product price P^*. The socially optimal import quantity is Qf^*. The socially optimal level of domestic consumption is Q^*.

Suppose the domestic policy maker has the authority to regulate domestic, but not foreign, producers. We first consider a policy regime in which the domestic monopolist is required to pay a fee of τ per unit of emissions. This increases the monopolist's variable operating costs by τe. The monopolist will choose to produce Qd_τ; the equilibrium product price is P_τ.

Figure 1 illustrates how this emissions regulation can *reduce* welfare (consistent with the theory of the second best). Intuitively, the costs associated with further exacerbating the exercise of market power in the domestic market can outweigh the benefits associated with the policy-induced emissions abatement. When domestic producers are required to pay τ per unit of output, domestic production drops even further below optimal levels.

The welfare effects of this policy can be decomposed into three components. Consider first the change to domestic producer and consumer surplus. The policy-induced reduction in consumer surplus that is not transferred to domestic producers is represented by area $ABCD$. In this trade-exposed market, the introduction of the emissions regulation increases the import market share. This induces "rent leakage," or transfer of surplus from domestic to foreign stakeholders. We assume that increases in foreign producer surplus do not factor into the domestic policy maker's objective function because they accrue outside her jurisdiction. Policy-induced reductions in domestic producer surplus that are not transferred to the government as tax revenue are given by $BGHF$.

Of course, the primary purpose of the emissions policy is to reduce emissions and associated damages. A second welfare component captures the value of the emissions reductions achieved domestically. This value is represented by area $EFGH$ (shaded with diagonal lines) in the right panel of Figure 1. In this case, the policy-induced loss in domestic economic surplus exceeds this value by an amount represented by the shaded area $AEFDC$.

A third welfare component accounts for the effect of the policy on foreign emissions. Here we assume that the policy-induced increase in import supply is met entirely by an increase in foreign production levels (versus a reallocation of foreign production across jurisdictions). Emissions leakage is represented by the shaded region in the left panel. Taken together, the

total welfare loss induced by the policy is represented by area $AEFDC$ plus the damages associated with emissions leakage (represented by the shaded area in the left panel).

Although a complete internalization of the carbon externality by domestic producers results in a net welfare loss in Figure 1, this will not always be the case in an industrial context characterized by both imperfect competition and exposure to competition from unregulated imports. As the marginal social cost of emissions increases and/or the import supply responsiveness attenuates, the policy can induce benefits (such as reduced emissions damages) that outweigh the costs (such as foregone producer and consumer surplus).

In the more detailed analysis that follows, we will be interested in analyzing the welfare implications of augmenting an emissions price τ with a domestic production subsidy s. This policy feature alleviates the market power distortion by incentivizing domestic output, while also mitigating, or even eliminating, emissions and rent leakage.[7] Figure 1 depicts the equilibrium outcome under a market-based emissions regulation that augments the emissions fee τ with an output-based rebate (or subsidy) s. The production subsidy incentivizes an increase in domestic production (domestic output is $Qd_{\tau-s}$). In addition to mitigating the exercise of market power, rent and emissions leakage are reduced because the subsidy acts to improve the terms of trade (vis a vis the regime that administers only the emissions fee).

Although the level of aggregate domestic consumption Q^* and the equilibrium product price P^* in this output-based rebating scenario are equal to those in the first best case, allocative efficiency is not achieved. Foreign imports still capture too much of the domestic market share; the marginal cost of domestic production is much lower than the marginal cost of importers. This highlights an important economic point: one generally needs as many policy instruments as market failures in order to achieve efficiency. While the tax on emissions and the production subsidy address the emissions externality and the exercise of market power in the domestic product market, respectively, an additional policy instrument is needed to address the asymmetry in compliance requirements across domestic and foreign producers.

[7] Policy makers have started to experiment with rebating tax revenues and allocating emissions permits on the basis of production. For example, in Sweden, revenues from an emissions tax are fully refunded to the industries that paid the tax on the basis of their energy use (Sterner and Hoglund, 2000). In existing and planned emissions trading programs in Australia, California, and Europe, permits are freely allocated to trade-exposed industries on the basis of output.

2.1 Applying the Framework

To more realistically simulate the response of domestic cement producers to alternative policy interventions, several of the simplifying assumptions that facilitate the graphical exposition must be relaxed. We highlight two of these assumptions here.

First, whereas Figure 1 features a domestic monopolist, regional cement markets in the United States are supplied by more than one domestic firm. Much of the intuition underlying the simple static monopoly case should apply in the case of a static oligopoly (Ebert, 1992).[8] A second modification pertains to industry dynamics. Figure 1 depicts static, short-run responses to market-based policy intervention. Over a longer time frame, firms can alter their choice of production scale, technology, entry, exit, or investment behavior in response to an environmental policy intervention. The welfare effects of a market-based emissions policy can look quite different across otherwise similar static and dynamic modeling frameworks.

On one hand, incorporating industry dynamics into the simulation model can improve the projected welfare effects of a given emissions regulation. Intuitively, the short run economic costs of meeting an emissions constraint can be significantly reduced once firms are able to re-optimize production processes, adjust investments in capital stock, and so forth. On the other hand, incorporating industry dynamics may result in estimated welfare effects that are strictly more negative than those generated using static models. First, in an imperfectly competitive industry, emissions regulation may further restrict already sub-optimal levels of investment, thus exacerbating the distortion associated with the exercise of market power. Second, a dynamic model can aggravate the extent of leakage to unregulated areas by accelerating exit and retirement of regulated production units in the domestic market.

3 Research context

The US domestic Portland cement industry has been at the center of the debate about domestic climate change policy and international competitiveness. Cement is one of the largest manufacturing sources of domestic carbon dioxide emissions (Kapur et al., 2009). Policies designed to internalize the social costs associated with greenhouse gas emissions could result in major changes to the industry's cost structure. For example, if we assume a cost of carbon in the neighborhood of \$40/ton, complete internalization of the emissions

[8] In certain situations, the oligopoly response to the policy could be more nuanced. For example, if firms are highly asymmetric and the inverse demand function has an extreme curvature, it is possible for the optimal tax rate to exceed marginal damage (Levin, 1985).

externality would almost double average variable operating costs.

3.1 The US Portland Cement Industry

Portland cement "clinker" is made by heating ground limestone and clay to a temperature of around 1400 degrees Celsius. Cement is then produced by grinding this clinker, along with gypsum, to produce an extremely fine powder. Concrete, an essential construction material used widely in building and highway construction, is basically a mixture of aggregates (e.g. sand and gravel), water, and Portland cement.

The US Portland cement industry is highly concentrated, making it potentially susceptible to the Buchanan critique. The top five companies collectively operate 54.4 percent of U.S. clinker capacity with the largest company representing 15.9 percent of all domestic clinker capacity. Moreover, import penetration in the domestic cement market has exceeded 20 percent in recent years, giving rise to concerns about the potential for emissions leakage (Van Oss and Padovani, 2002; USGS, 2010).

The US cement industry is fragmented into regional markets. This fragmentation is primarily due to transportation economies. The primary ingredient in cement production, limestone, is ubiquitous and costly to transport. To minimize input transportation costs, cement plants are generally located close to limestone quarries. Land transport of cement over long distances is also not economical because the commodity is difficult to store (cement pulls water out of the air over time) and has a very low value to weight ratio. It is estimated that 75 percent of domestically produced cement is shipped less than 110 miles (Miller and Osborne, 2010).[9]

Domestic demand Demand for cement comes primarily from the ready-mix concrete industry, which accounts for over 70 percent of cement sales. Other major consumers include concrete product manufacturers and government contractors. Since 1960, domestic demand has been fluctuating between 60,000 and 100,000 tons. Demand for domestic cement tends to reflect the cyclical nature of the larger economy, and construction activity in particular. In the construction sector, cement faces competition from alternatives such as asphalt, clay brick, rammed earth, fiberglass, steel, stone, and wood (Van Oss and Padovani, 2003). Another important class of substitutes are the so-called supplementary cementitious materials

[9] Most cement is shipped by truck to ready-mix concrete operations or construction sites in accordance with negotiated contracts. A much smaller percent is transported by train or barge to terminals and then distributed.

(SCMs) such as ferrous slag, fly ash, silica fume and pozzolana (a reactive volcanic ash). Concrete producers can use these materials as partial substitutes for clinker.[10]

Trade Exposure Whereas overland transport of cement is very costly, sea-based transport of clinker is relatively inexpensive. In the 1970s, technological advances made it possible to transport cement in bulk quantities safely and cheaply by barge and in large ocean vessels. Since that time, U.S. imports have been growing steadily. Over the period 1980 to 2006, the import market share has increased from below 3 percent to over 25 percent. Canada is currently the largest supplier of imported cement, followed by China, Korea, and Mexico (USGS (2012), fact sheet). Exposure to import competition in regional markets has given rise to growing concerns about unilateral climate policy.

Carbon dioxide emissions from cement production Cement producers are among the largest industrial emitters of airborne pollutants, second only to power plants in terms of the criteria pollutants currently regulated under existing cap-and-trade programs (i.e. NOx and SO_2). The cement industry is also one of the largest manufacturing sources of domestic carbon dioxide emissions (Kapur et al., 2009). Worldwide, the cement industry is responsible for approximately 7 percent of anthropogenic CO_2 emissions (Van Oss and Padovani, 2003).

Approximately half of the carbon dioxide associated with the manufacture of cement is directly released as a byproduct of the chemical process that transforms limestone to clinker. Fossil fuel combustion at cement manufacturing operations accounts for approximately 45 percent of the industrys emissions. Trace amounts of carbon dioxide are released during the grinding phase.

Carbon dioxide emissions intensities, typically measured in terms of metric tons of emissions per metric ton of clinker, vary across cement producers. Much of the variation is driven by variation in fuel efficiency. The oldest and least fuel efficient kilns are "wet-process" kilns. As of 2006, there were 47 of these wet kilns in operation (all built before 1975) (PCA, 2006). "Dry process" kilns are significantly more fuel efficient, primarily because the feed material used has a lower moisture content and thus requires less energy to dry and heat. The most modern kilns, dry kilns equipped with pre-heaters and pre-calciners, are more than twice as

[10]The substitution of SCM for clinker can actually improve the quality and strength of concrete. Substitution rates range from 5 percent in standard Portland cement to as high as 70 percent in slag cement. These blending decisions are typically made by concrete producers and are typically based on the availability of SCM and associated procurement costs (van Oss, 2005; Kapur et al., 2009).

fuel efficient as the older wet-process kilns.

Emissions Abatement Several recent studies assess the potential for carbon emissions reductions in the cement sector.[11] Using different scenarios, baseline emissions and future demand forecasts, all reach similar conclusions. Although there is no "silver bullet," there are four key levers for carbon emissions reductions.

The first set of strategies involve energy efficiency improvements. The carbon intensity of clinker production can be reduced by replacing older equipment with current state of the art technologies. Converting all existing kilns to more efficient, state-of-the art technology could achieve reductions in domestic carbon dioxide emissions on the order of 15 percent.[12]

A second set of carbon mitigation strategies involve substitution. One form of substitution increases the use of alternate construction materials such as wood or brick, thus reducing demand for cement. Alternatively, the amount of clinker needed to produce a given amount of cement can be reduced by the use of supplementary cementitious materials (SCM) such as coal fly ash, slag, and natural pozzolans.[13] It is estimated that the increased use of blended cement could feasibly reduce carbon emissions by a third over the time frame we consider (Mahasenan et al., 2005).

Fuel switching offers a third emissions abatement strategy. Less carbon intensive fuels, such as waste derived fuels or natural gas, could replace coal as the primary kiln fuel. Although there are limits to the substitutability of fuels, it is estimated that fuel switching can reduce the carbon intensity of cement production by as much as 25 percent (on Sustainalble Development, 2010).

Finally, carbon dioxide emissions can be separated and captured during or after the production process and subsequently sequestered. This abatement option is unlikely to play a significant role in the near term given that sequestration technologies are in an early stage of technical development and are relatively costly.

We explicitly model the replacement of older kiln technology with current, state-of-the-art technology. We assume all new entrants adopt new, state-of-the-art equipment. This assumption finds empirical support in the data. Our specific assumptions about the emissions intensities of old and new production equipment are described in Appendix C.

[11] A comprehensive list of studies can be found at http://www.wbcsdcement.org/pdf/technology/References%20FINAL.pdf

[12] These calculations are based on estimated emissions intensities. See Appendix C for details.

[13] When part of the cement content of concrete is replaced with supplementary cementitious materials, the extent of the emissions reduction is proportional to the extent to which SCM replaces clinker. Substitution rates as high as 75 percent are possible.

We implicitly capture the substitution of alternate materials for cement clinker in our policy simulations. Supplementary cementitious materials and substitute construction materials were widely used throughout our study period. Although we do not explicitly model the substitution of SCMs for clinker, this substitution is implicitly captured, to some extent, by our estimated demand elasticity.

Ideally, a model designed to simulate industry response to emissions regulations would capture all viable carbon abatement strategies. Unfortunately, we cannot estimate the costs associated with responses that have yet to be observed in the data. Consequently, fuel switching and carbon sequestration are not represented in our analysis. Although these options are not expected to play as significant a role as efficiency improvements or substitution, this omission will bias up our estimates of the economic costs imposed by the emissions regulations we analyze.

3.2 Market-based Emissions Regulation

We analyze both static and dynamic industry response to the introduction of market-based emissions regulation. Our primary focus is a multi-sector, nation-wide cap-and-trade program. A defining feature of the program is a cap which imposes a binding constraint on the quantity of carbon emissions released by sources in the program. A corresponding number of pollution permits are issued. To remain in compliance, regulated sources must hold permits to offset uncontrolled emissions. These permits are traded freely in the marketplace.

Having defined the emissions cap, the regulator must decide how to allocate or distribute the emissions permits. We are particularly interested in exploring the efficiency implications of alternative emissions permit allocation approaches. The first policy design we analyze is a cap-and-trade program in which permits are allocated via a uniform price auction. Within our modeling framework, this policy design is mathematically equivalent to a carbon tax.

Many industry stakeholders vehemently oppose a policy regime that would auction all permits (at least in the near term). In existing and planned emissions trading programs, the majority of permits are distributed *gratis* to regulated firms. This motivates the study of our second policy regime, "grandfathering," where permits are freely allocated according to pre-determined factors, such as historic emissions.

In recent years, a third design alternative has emerged. Emissions permits are allocated for free to eligible firms using a periodically-updated, output-based formula. This dynamic allocation updating is being used to mitigate leakage and associated competitiveness impacts in trade-exposed, emissions-intensive industries. The incentives created by this dynamic

allocation updating rule are quite different than those associated with grandfathering or auctioning because updating confers an implicit production subsidy.

Finally, border tax adjustments offer an alternative approach to mitigating emissions leakage in trade-exposed, emissions intensive industries. These import taxes are intended to penalize the emissions embodied in foreign imports, thus "leveling the carbon playing field." Although border tax adjustments face formidable legal challenges (see, for example, Fischer and Fox (2009)), we consider this policy design feature because it has the potential to play an important role in leakage mitigation.

4 Model

The basic building block of the model is a regional cement market.[14] Let $\overline{N}$ be the maximal number of active firms in the market. Each market is described by two $\overline{N} \times 1$ state vectors, s and e. The vector s describes the productive capacity of the firms at the market. Firms can adjust their capacity over time, by means of entry, exit, investment and disinvestment. Firms with zero capacity are considered to be potential entrants.

The vector e describes the emissions rate of each firm. We assume that there are three discrete levels of emissions rates, corresponding to the three major types of production technology (wet, dry, state-of-the-art dry) in the cement industry. Incumbents may be of any technology type, while we assume that all new entrants are endowed with the frontier technology.

Firms obtain revenues from the product market. They incur costs from production, entry, and new investment. We model timing as an infinite horizon model with each discrete decision period being one year. Firms discount the future at rate β. In each period, first, incumbent firms decide whether or not to exit the industry based on their exit cost shock. Second, potential entrants receive both investment and entry cost shocks, while incumbents who have decided not to exit receive investment cost shocks. All firms then simultaneously make entry and investment decisions. Third, incumbent firms compete over quantities in the product market. At the end of the period, firms enter and exit, and investments mature.

We assume that firms who decide to exit produce in the period before leaving the market, and that adjustments in capacity take one period to realize. We also assume that each firm operates independently across markets.[15]

[14] The model is based on Ryan (2012), to which we add imports, divestment, emissions technologies, and environmental policies.

[15] This assumption explicitly rules out more general behavior, such as multimarket contact as considered

4.1 Static payoffs

Firms compete in quantities in a homogeneous goods product market. Firms face a constant-elasticity aggregate demand curve:

$$\ln Q_m(P_m;\alpha) = \alpha_{0m} + \alpha_1 \ln P_m, \tag{1}$$

where Q_m is the aggregate regional market quantity, P_m is price, α_{0m} is a market-specific intercept, and α_1 is the elasticity of demand.

For firms in trade-exposed regional markets, residual demand is more elastic, as they also face import competition. The import supply curve is given by:

$$\ln M_m(P_m;\rho) = \rho_0 + \rho_1 \ln P_m, \tag{2}$$

where M_m measures annual import supply in market m and ρ_1 is the elasticity of import supply. Here we assume that the elasticity of import supply is an exogenously determined parameter.[16] Domestic firms in import-exposed markets face a residual demand curve formed by subtracting off the import supply curve from the market-level demand curve. For clarity, we omit the m subscript in what follows.

In the model, each firm chooses the level of annual output that maximizes its static profits given the outputs of the competitors, subject to capacity constraints that are determined by dynamic capacity investment decisions:

$$\overline{\pi}(s,e,\tau;\alpha,\rho,\delta) \equiv \max_{q_i \le s_i} \; P\big(q_i + \sum_{j\neq i} q_j^*;\alpha,\rho\big)q_i - C_i(q_i;\delta) - \varphi(q_i,e_i,\tau), \tag{3}$$

where $P(Q;\alpha,\rho)$ is the inverse of residual demand. The profit $\overline{\pi}(s,e,\tau;\alpha,\rho,\delta)$ defines the equilibrium static profits of the firm for a given level of capacity and kiln type. If all firms produce positive quantities then the equilibrium vector of production is unique, as the best-response curves are downward-sloping.

The cost of output, q_i, is given by the following function:

$$C_i(q_i;\delta) = \delta_{i1} q_i + \delta_2 1(q_i > \nu s_i)(q_i/s_i - \nu)^2. \tag{4}$$

in Bernheim and Whinston (1990) and Jans and Rosenbaum (1997).

[16] In fact, firms that own a majority of the domestic production capacity in the United States are also among the largest importers. It is possible that domestic climate policy could induce a structural shift in the supply of imports to the domestic market. We return to this issue in Section 6.5.

Variable production costs consist of two parts: a constant marginal cost, δ_{i1}, and an increasing function that binds as quantity approaches the capacity constraint.[17] We assume that costs increase with 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, 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.

The term $\varphi(q_i, e_i, \tau)$ represents the environmental compliance costs faced by the firm. The carbon cost, τ, is an exogenous parameter intended to capture the monetized damages associated with an incremental (one ton) increase in carbon emissions.[18] Importantly, we assume a constant real carbon price over our relatively short (30 year) time horizon. In our model, there is no technological innovation over time, nor is there economic growth. Thus, some of the standard justifications for implementing a policy regime in which the compliance cost per unit of emissions increases over time do not apply in our case.

The policy designs we analyze can be classified into one of four categories: auctioning/carbon tax; grandfathering; output-based rebating; and an auctioning regime augmented with a border-tax adjustment.

Emissions tax or emissions trading with auctioned permits The first policy regime we analyze is an emissions tax or an emissions cap-and-trade program in which all emissions permits are allocated via a uniform price auction. In the tax regime, regulated firms must pay a tax τ for each ton of emissions. In the emissions trading regime, the equilibrium permit price is τ; under our assumption that cement firms are price-takers in the permit market, a change in the net supply or demand for permits from the domestic cement industry does

[17] Note that we do not consider fixed costs of production and operation. The reason is that we do not observe sufficient periods of operation without production (mothballing) which are required to separately identify those parameters from the distribution of exit costs.

[18] The exogeneity assumption seems appropriate as the domestic cement industry is a relatively small player in a potential economy-wide emissions market, such that changes in industry net supply/demand for permits cannot affect the equilibrium market price. Keohane (2009) estimates the slope of the marginal abatement cost curve in the United States (expressed in present-value terms and in 2005 dollars) to be 8.0 x 10^7 \$/GT CO_2 for the period 2010–2050. Suppose this curve can be used to crudely approximate the permit supply function. If all of the industries deemed to be "presumptively eligible" for allowance rebates reduced their emissions by ten percent for this entire forty year period, the permit price would fall by approximately \$0.25/ton. This also assumes that mitigation in the cement industry is not offsetting distortionary mitigation in another industry. In Section 6.4, we calculate the welfare costs associated with achieving given levels of abatement in the cement industry alone.

not affect this price. The environmental compliance cost to the firm is given by:

$$\varphi(q_i, e_i, \tau) = \tau e_i q_i. \tag{5}$$

Grandfathering In this policy scenario, a share of emissions permits are allocated for free to incumbent firms that pre-date the carbon trading program. Firm-specific permit allocation schedules are determined at the beginning of the program and are based on historic emissions. The environmental compliance cost to the firm in this regime is:

$$\varphi(q_i, e_i, \tau) = \tau(e_i q_i - A_i), \tag{6}$$

where A_i is the total emission permits that the firm receives for free from the regulator.

Note that the first order conditions associated with static profit maximization under grandfathering are identical to those under auctioning. This highlights the so-called "independence property," which implies that firms' short run production and abatement decisions will be unaffected by the choice between auctioning permits or allocating them freely to firms in lump sum (Hahn and Stavins, 2010). Dynamically, however, both mechanisms generally generate different long-run outcomes, primarily due to the exit decision being distorted by the transfer of valuable assets to incumbent firms under grandfathering.

When permits are grandfathered in a cap-and-trade program, policy makers must decide ex ante how to deal with both firms that exit and new entrants.[19] We assume that the share of emissions allowances allocated to a firm is proportional to the installed kiln capacity at the outset of the program, s_{i0}. However, if firms divest part of their historic capacity, they give up part of their initial allocation, i.e. $A_i = \psi_g \cdot e_i \min\{s_{i0}, s_i\}$, where ψ_g is a parameter converting capacity into permits.[20] Furthermore, we assume that a firm forfeits its future entitlements to free permits when it exits the market.[21] Finally, we assume that new entrants are not entitled to free permits.[22]

[19] See Dardati (2013) for a recent contribution studying the effects of these policies.

[20] We include this feature to better represent some of the trade-offs faced when implementing grandfathering. In the EU ETS, the allocation of free permits is reduced dynamically if firms divest part of their grandfathered capacity.

[21] Note that if firms were to keep all their permits indefinitely then this mechanism would be dynamically welfare-equivalent to auctioning, although distributionally different, so the independence property would apply. In the EU ETS, most states require firms to forfeit their free permits upon closure.

[22] In practice, policies regarding free permit allocations to free entrants and former incumbents vary. In the EU ETS, policies governing the free allocation of permits to entrants vary across member states.

Output-based allocation updating/rebating The third policy regime we analyze incorporates output-based rebating. Permits are allocated (or tax revenues are recycled) per unit of production based on an industry-specific emissions intensity benchmark. The environmental compliance cost to the firm becomes:

$$\varphi(q_i, e_i, \tau) = \tau \cdot (e_i - \psi_d) \cdot q_i, \tag{7}$$

where ψ_d controls the proportion of emissions rebated to the firm. Equation 7 illustrates that output-based updating operates as a discount on the amount of permits (or tax payments) required to achieve compliance. Alternatively, one can think of this as a production subsidy.

Border tax adjustment with auctioned permits The fourth and final policy design that we consider layers a border tax adjustment (BTA) on top of the standard tax/auctioning regime. This BTA mechanism imposes a tax on emissions embodied in cement imports equal to the tax imposed on domestic emissions. This effectively levels the carbon playing field with international competitors. The BTA regime is equivalent to the auctioning regime in terms of the function $\varphi(q_i, e_i, \tau)$. However, domestic firms now face a different residual demand, as the import supply is shifted to the left as follows:

$$\ln M(P; \rho, \tau) = \rho_0 + \rho_1 \ln(P - \tau e_M), \tag{8}$$

where e_M is the emissions rate on imported cement.

4.2 Dynamic decisions

Firms have the opportunity to adjust capacity in each period. Firms can increase or decrease their capacity through costly investments, denoted by x_i. The cost function associated with these investments is given by:

$$\Gamma(x_i; \gamma) = \gamma_{i1} + 1(x_i > 0)(\gamma_2 x_i + \gamma_3 x_i^2) + 1(x_i < 0)(\gamma_4 x_i + \gamma_5 x_i^2). \tag{9}$$

Firms face both fixed and variable investment and divestment costs. The 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 change in capacity. The fixed investment cost is drawn each period from the common distribution F_γ, which is distributed normally with mean μ_γ and standard deviation σ_γ, and is private information to

the firm. Firms also face variable adjustment costs that scale with the size of the capacity change.

Firms also make market participation decisions, denoted by a_i. When a firm changes market participation status, it receives a transfer $\Phi(a)$ that varies depending on their current status and chosen action:

$$\Phi(a_i; \kappa_i, \phi_i) = \begin{cases} -\kappa_i & \text{if the firm is a new entrant,} \\ \phi_i & \text{if the firm exits the market.} \end{cases} \tag{10}$$

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 and paying for an environmental cleanup. This value may be positive or negative, depending on the magnitude of these opposing payments. The scrap value is private information, drawn anew each period from the common distribution, F_ϕ. All of the shocks that firms receive each period are mutually independent.

Collecting the costs and revenues from a given firm, the per-period payoff function is:

$$\pi_i(a, x, s, e; \theta, \tau) = \overline{\pi}_i(s, e; \alpha, \rho, \delta, \tau) - \Gamma(x_i; \gamma_i) + \Phi(a_i; \kappa_i, \phi_i). \tag{11}$$

where θ denotes the vector of parameters in the model, and the permit price is τ.

To close the dynamic elements of the model it is necessary to specify how transitions occur between states as firms engage in investment, entry, and exit. We assume that changes to the state vector through entry, exit, and investment take one period to occur and are deterministic. The first part 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. The second part 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.[23]

4.3 Equilibrium

In each time period, firm i makes entry, exit, production, and investment decisions. Since the full set of dynamic Nash equilibria is unbounded and complex, we restrict the firms'

[23] It is conceptually straightforward to add uncertainty over time-to-build in the model, but assuming deterministic transitions greatly reduces the computational complexity of solving for the model's equilibrium.

strategies to be anonymous, symmetric, and Markovian, meaning firms only condition on the current state vector and their private shocks when making decisions, as in Maskin and Tirole (1988) and Ericson and Pakes (1995). We describe the equilibrium Bellman equations in online Appendix A.

To compute the equilibrium of the model, we develop parametric approximation methods for the computation of dynamic games. In particular, we interpolate the value function using cubic splines. The equilibrium is computed separately for every market and environmental policy considered. The interested reader can find a detailed description of the methodology in online Appendix B, where we also discuss the main strengths and limitations of this methodology.[24]

4.4 Welfare measures

Within a regional market, it is useful to decompose the net welfare impact of a policy intervention into the three components introduced in Section 2.

We define the following per-period equilibrium welfare measures:

$$\begin{aligned} w_1(s,e,\tau;\theta) &= \int_0^{Q*} P(z;\alpha)dz - P(Q^*;\alpha)Q^* + \sum_i \Pi_i(a^*,x^*,s,e,\tau;\theta) \qquad (12a) \\ &\quad + \sum_i \varphi(q_i^*,e_i,\tau) + \tau_M e_M M, \\ w_2(s,e,\tau;\theta) &= w_1(s,e,\tau;\theta) - \tau \sum_i e_i q_i^*, \qquad (12b) \\ w_3(s,e,\tau;\theta) &= w_2(s,e,\tau;\theta) - \tau e_M M(P^*;\rho). \qquad (12c) \end{aligned}$$

Welfare measure w_1 captures the domestic economic surplus of cement consumption: consumer surplus, producer surplus, and government revenues. Note that government revenues include the carbon price paid by importers (τ_M), which will be zero under most mechanisms, but equal to τ in the BTA case. We assume that domestic policy makers exclude profits earned outside their jurisdiction from any welfare analysis.

Welfare measure w_2 accounts for both economic surplus changes plus the costs associated with domestic emissions. In a cap-and-trade system where aggregate emissions are fixed, any increase in emissions from the cement industry must be offset by emissions abatement in other covered sectors. Equation 12b implicitly assumes that the permit supply curve facing the cement industry is locally flat. We further assume that other covered sectors are undistorted.

[24] See also Doraszelski and Pakes (2007); Arcidiacono et al. (2012); Farias et al. (2012).

[25]

Finally, welfare measure w_3 adds a penalty for emissions leakage at the cost of carbon τ. Both domestic emissions and the emissions associated with foreign imports are penalized at the social cost of carbon.

We will focus on comparing the net present value of these welfare measures against the baseline case in which no emissions regulation is in place. We define $w_0(s, e, \tau; \theta)$ as the per-period welfare in the baseline case. The net present value (NPV) welfare measures that we consider are:

$$W1 = \sum_{t=1}^{T} \beta_S^t \big(w_{1t}(s, e, \tau; \theta) - w_{0t}(s, e, \tau; \theta)\big), \tag{13}$$

where β_S is social discount factor. $W2$ and $W3$ are defined analogously.

5 Data and Estimation

This section begins with a discussion of the data. We then turn to the estimation which proceeds in several steps. We first estimate the so-called static parameters: the parameters of the demand function, import supply and parameters used to characterize the cost structure. Next, we estimate the policy functions that describe firm's entry, exit, and investment choices. These policy functions are then used to find the dynamic parameter values which reconcile observed investment, entry, and exit choices with our model of profit maximization. This section concludes with a description of how we calibrate the parameters that define the counterfactual environmental policies.

5.1 Data

Our cement industry data come from two main sources: the U.S. Geological Survey (USGS) and the Portland Cement Association. The USGS collects establishment-level data from all domestic Portland cement producers. These data, aggregated regionally to protect the confidentiality of the respondents, are published in annual volumes of the *Minerals Yearbook*. Kiln-level data are available from the *Plant Information Survey* (PIS), an annual publication of the Portland Cement Association. The PIS provides information on the location, vintage, kiln-type, primary fuel, and operating capacity of each operating kiln.

[25] Some of the other industries subject to climate change regulation might also be at risk for emissions leakage and imperfect competition. We return to this assumption in Section 6.

Firm-level data on entry, exit, and capacity adjustment is an important input to our analysis. We obtain kiln-level information from the annual PIS and cross-validate this information using the annual summaries published by the USGS. Over the twenty-five year study period, we observe 11 plant entries and 51 exits, with an implied entry and exit rate of 2.2 percent and 2.0 percent, respectively.[26] We observe 144 capacity increases (i.e. investment in one or more new kilns). We observe 95 capacity decreases. Overall, the total capacity adjustment rate is 6.6 percent.[27]

We choose not to use the regional definitions adopted by the USGS in our analysis. In recent years, increased consolidation of asset ownership has led to higher levels of data aggregation in the USGS reports. Instead, we follow the EPA (2009) and use city-centered market definitions derived from industry-accepted limitations of economic transport as well as company-specific SEC 10k filings which include information regarding markets served by specific plants. We re-weight the USGS data on prices and quantities by kiln capacity in each region to form less aggregate measures of production and prices. For example, if kiln capacity in USGS market A is equally divided between EPA markets B and C, production quantities in market A are equally divided between our defined markets B and C. For computational reasons, in the counterfactual analysis we focus on markets with five or fewer firms. These markets are listed in Table 1.[28]

Table 1 reports the regional market-level summary statistics using PCA data from 2006. The table helps to highlight inter-regional variation in market size, emissions intensity, and trade exposure. Notably, the degree of import penetration varies significantly across inland and coastal areas. As expected, import penetration rates tend to be highest along the markets with direct coastal ports versus those served by inland waterways.

We also collect data on electricity rates, coal prices, natural gas prices, and wage rates to serve as instruments in our demand estimation. Energy prices are collected from the U.S. Energy Information Administration, while the wage rates are derived from the U.S. Census'

[26] To compute the entry rate, we consider that there is one potential entrant in every period, therefore we divide by twenty markets times twenty-five periods. To compute the exit rate, we divide by the number of active firm yearly observations in the sample.

[27] In the data, we periodically observe year-to-year fluctuations in kiln-level operating capacities. In particular, we often observe kiln capacities declining the year before a major capacity addition. We interpret small fluctuations of less than 10 percent as noise in the data. As such, these small, short-lived fluctuations are smoothed out of the data.

[28] In restricting our attention to those regional markets with five or fewer incumbent firms, we omit four markets from the analysis: Atlanta, Baltimore, Los Angeles, and San Antonio. Our sample covers approximately 70% percent of the market. We have repeated all of the analysis including only markets with three or fewer firms, and four or fewer firms, and the conclusions of our work are robust to the subset of markets considered. See Table F.7 in the online appendix for results under alternative subsets of the data.

Table 1: Descriptive Statistics for Regional Markets (based on 2006 data)

Market	Number of Firms	Aggregate Annual Capacity	Average Emissions Rate	Import Market Share
Birmingham	5	1288	0.94	0.35
Chicago	5	972	0.98	0.04
Cincinnati	3	875	0.93	0.21
Dallas	5	1766	1.05	0
Denver	4	998	0.95	0
Detroit	3	1749	1.02	0.19
Florida	5	1297	0.93	0.35
Kansas City	4	1661	0.95	0
Minneapolis	1	1862	0.93	0.2
New York/Boston	4	1033	1.16	0.45
Phoenix	4	1138	0.93	0.13
Pittsburgh	3	614	1.08	0
Salt Lake City	2	1336	1.01	0
San Francisco	4	931	0.93	0.18
Seattle	2	607	1.05	0.65
St Louis	4	1358	1.05	0

Notes: Capacity is measured in thousands of tons of cement. Emissions rates are defined as tons of CO_2 per ton of produced cement.

County Business Patterns. All prices are adjusted to year 2000 constant dollars.

5.2 Static Parameters

Demand Following Ryan (2012), we estimate the following demand equation:

$$\ln Q_{mt} = \alpha_m + \alpha_1 \ln P_{mt} + \alpha_2 X_{mt} + \varepsilon_{1mt}. \tag{14}$$

The dependent variable is the natural log of the total market demand in market m in year t. The coefficient on market price, α_1, is the elasticity of demand. We instrument for the potential endogeneity of price using supply-side cost shifters: coal prices, natural gas prices, electricity rates, and wage rates. The matrix X_{mt} includes demand shifters such as population and economic indicators. We estimate the parameters of the demand equation using the annual USGS data over the period 1981-2009 using limited-information maximum likelihood.

Table 2 summarizes the estimation results for several specifications with robust standard errors reported in parentheses. The first specification includes only regional market

Table 2: IV Estimation of Demand Elasticity

	(1)	(2)	(3)	(4)	(5)
Log price	-2.03	-0.89	-1.47	-0.92	-1.10
	(0.28)	(0.22)	(0.17)	(0.18)	(0.18)
Log Population		1.34			
		(0.14)			
Log Units			0.51		0.40
			(0.04)		(0.07)
Log Unemployment				-0.65	-0.29
				(0.05)	(0.09)
First stage F-test	132.19	113.73	199.75	170.47	193.11

Notes: Huber-White robust standard errors are in parentheses. The unit of observation is a market-year. Market fixed-effects are included in all specifications. Sample run from 1980 to 2009.

fixed effects. The point estimate for the elasticity of demand is -2.03.[29] This specification omits several factors that presumably shift demand, such as population, unemployment, and measures of construction activity. Subsequent specifications (2) through (5) include these factors. Our point estimate of the own-price demand elasticity is somewhat sensitive to the inclusion of these covariates, varying between -0.9 and -2.0.

We select specification (1) as our preferred specification; it is the most parsimonious and consistent with the dynamic structural estimation. Our theoretical model does not explicitly capture changes in population or building activity over time. Given the critical role that the demand elasticity plays in our analysis, we perform a series of robustness checks where we simulate policy outcomes over a range of possible demand elasticity values.

Imports For trade-exposed markets, defined as markets in which we see imports claiming some non-zero market share, we estimate the following import supply schedule using limited information maximum likelihood:

$$\ln M_{mt} = \rho_0 + \rho_1 \ln P_{mt} + \rho_{2m} + \rho_3' \ln Z_{mt} + \varepsilon_{2mt}. \tag{15}$$

[29]The estimate is higher in absolute value than some other demand elasticities reported in the literature. For example, Jans and Rosenbaum (1997) estimate a domestic demand elasticity of -0.81. Using data from 12 European countries over the period 1990-2005, Sato et al. (2008) estimate a demand elasticity of -1.2. Using USGS data from the Southwestern U.S., Miller and Osborne (2010) estimate an aggregate demand elasticity of -0.16. On the other hand, Foster et al. (2008) estimate several similar high demand elasticities for homogeneous goods industries, such as -5.93 for ready-mixed concrete, cement's downstream industry.

Table 3: IV Estimation of Import Elasticity

	(1)	(2)	(3)	(4)
Log Price	2.47	2.85	2.52	3.00
	(1.64)	(2.50)	(1.28)	(1.12)
Log Transport		0.75	0.45	0.47
		(0.26)	(0.13)	(0.12)
Log Coal Price		0.01	-0.06	0.11
		(0.12)	(0.15)	(0.14)
Log Oil Price		0.33	0.36	-10.66
		(0.25)	(0.18)	(4.40)
Regional Dummies	Yes	No	Yes	Yes
Yearly Dummies	No	No	No	Yes

Notes: Huber-White robust standard errors are reported in parenthesis. The unit of observation is a market-year. The sample runs from 1993 to 2009.

The dependent variable is the log of the quantity of cement shipped to market m in year t. The average price paid for imported cement is P_{mt}. These data are reported by Customs district, which may contain several ports of entry. Each port of entry is matched to a regional market as described above. The model is estimated using data from the period 1993-2009.[30]

We instrument for the import price using gross state product, new residential construction building starts, and state-level unemployment. The matrix Z_{mt} includes other plausibly exogenous factors that affect import supply. To capture transportation costs, we subtract the average customs price from the average cost, insurance, and freight (CIF) price of the cement shipments. This residual price accounts for the transportation cost on a per unit basis, as well as the insurance cost and other shipment-related charges. The Z_{mt} matrix also includes coal and oil prices to capture variation in production costs. Region dummy variables capture regional differences.

The most parsimonious specification includes only regional fixed effects. The estimated import supply elasticity is 2.47. This parameter is imprecisely estimated, with a standard error is 1.64. An alternative specification includes a series of supply shifters, including coal prices, oil prices, a measure of the cost of transporting the cement from the supply country to the import district in the United States. Including these controls does not significantly affect our point estimate.

Presumably, the degree of import competition varies across trade-exposed regional markets. For example, one might expect import responsiveness to vary across markets served primarily by terminals on inland waterways versus coastal markets supplied via marine ter-

[30] District-level data on imports from earlier years contain many missing values.

minals. Unfortunately, because publicly available data on cement imports are noisy and highly aggregated, we are unable to estimate market-specific supply elasticities. We can, however, allow for regional variation in the level of imports supplied at a given cement price. We obtain market-specific intercepts by fitting import supply curves at the market level while fixing the elasticity coefficient at 2.5.[31]

Production Costs We use the estimated demand and import parameters, together with the restrictions on behavior implied by the Cournot oligopolistic model, to estimate the firms' production costs in Equation 3. For each firm i in market by j at time t, the estimator minimizes two equally-weighted moments: the sum of squared differences between the observed quantities and the predictions of the model, and the sum of squared differences between marginal cost and marginal revenue at the equilibrium level of output.[32] If a firm has multiple plants in a single market, we treat that firm as having a single plant with capacity equal to the sum of capacity in each of those facilities.

There are three basic parameters in the cost function: the constant foundation of the marginal cost curve, δ_1; the increasing marginal cost parameter, δ_2, which is incurred as the firm produces close to maximum capacity; and the threshold, ν, determining when δ_2 enters the cost function. To bound the threshold as a percentage of installed capacity, we estimate $\tilde{\nu}$ in a logit transformation $\nu = \exp(\tilde{\nu})/(1.0 + \exp(\tilde{\nu}))$. The parameters are estimated using GMM. Standard errors for production costs (and all following parameters) are calculated by bootstrapping complete market histories, with replacement, 200 times. When computing standard errors, we hold elasticity of demand and import supply at their empirical means.[33]

The results from the estimation are included in Table 4. Baseline marginal costs are estimated to be \$47 per ton of cement. At an average price of \$75 per ton of cement during our sample period, this implies a gross margin of \$28 per ton, or 37 percent, over the range before the increasing marginal costs start. This markup seems reasonable for a capacity-constrained industry with extremely high sunk costs. The estimated threshold for those capacity costs is at 87 percent of annual capacity, which, combined with the high additional

[31] The US EPA (Burtraw, 2011) assumes an import supply elasticity of 3.94 when analyzing the impacts of environmental regulations on the cement sector. In the interest of understanding why our estimates differ, we obtained data and code from the EPA analysis. There are two main reasons why our elasticity estimates differ. First, the EPA analysis uses weighted 2SLS, versus LIML, to estimate a very similar import supply specification. Second, whereas we use data on all cement imports, EPA analysts use data from the 5 largest trade partners only.

[32] Experimentation with alternative weights did not change the results significantly.

[33] As part of an extensive sensitivity analysis, point estimates of production and dynamic costs for different combinations of elasticities are provided in Tables F.2 and F.3 in the online appendix.

Table 4: Marginal Cost Estimates

	Estimate	**SE**
Marginal cost ($/000 Ton)	46.99	(0.82)
Capacity cost ($/Extra % Utilization)	803.65	(60.92)
Utilization Threshold Estimate	1.889	(0.040)
Implied Utilization % Threshold	0.869	(0.005)

Notes: Bootstrapped standard errors in parenthesis.

production costs after that point, is roughly consistent with the idea that cement plants typically shut down for a month and a half for maintenance per year.

To assess the plausibility of these estimated production costs, we collected the annual financial statements for Cemex, one of the largest cement producers in our sample. Over the years 2010, 2011, and 2012, Cemex reports gross combined profits of $12.377 billion on revenues of $43.183 billion, for a profit margin of 28.6 percent.[34] Furthermore, the EPA reports engineering estimates of average production costs of $44.4 per ton of produced cement, which is close to our point estimate.[35] In sum, these marginal cost estimates lie within an economically reasonable range.

5.3 Dynamic Parameters

To estimate the dynamic parameters, we follow the two-step empirical strategy laid out in Bajari et al. (2007), and used in Ryan (2012). First, we estimate the policy functions that describe firm investment, entry, and exit behaviors as a function of economic state variables. Second, we project these policy functions onto our underlying structural model via forward simulation.

Policy Functions To estimate the investment policy function, we follow the approach in Ryan (2012) and use an (s,S) rule model. The (s,S) model is designed to capture lumpy adjustment behavior—periods of inactivity followed by large discrete changes in capacity—and consists of two latent equations: a target equation, $T(s)$, and a band equation, $B(s)$, which is defined to be non-negative. The target equation sets the level of capacity a firm

[34] While our estimated margins are higher than Cemex's reported margins, this could be partly explained by our assumption that fixed costs of production are zero and also the fact that we have evaluated the margins using the marginal cost on the flat (cheapest) part of the cost function.

[35] See International (2009). An average cost of $50.30 in 2005 dollars is reported, which we convert into 2000 dollars.

adjusts to, conditional on making a change, while the band equation controls when the firm will make a change. Letting the current capacity at time t be denoted by s_t, the policy function for incumbent firms is:

$$s_{t+1} = \begin{cases} T(s_t) & \text{if } s_t < T(s_t) - B(s_t) \text{ or } s_t > T(s_t) + B(s_t), \\ s_t & \text{else.} \end{cases} \tag{16}$$

Entrants adjust to $T(s_t)$. The target and band equations are, respectively:

$$\ln T_{imt}(s) = \eta_1 + \eta_2 1(i \text{ entrant}) + \eta_3(1 - 1(i \text{ entrant})) \ln \text{Capacity}_i + \eta_4 MT_m + \epsilon_T, \tag{17}$$

$$\ln B_{imt}(s) = \eta_5 + \eta_6 \ln \text{Capacity}_i + \eta_7 MT_m + \epsilon_B. \tag{18}$$

The target capacity depends on whether the firm is an entrant to the market, the firm's current log-transformed capacity, and a "market tightness" variable MT. Market tightness is defined as the ratio of current aggregate market capacity to the maximum aggregate market capacity ever observed in our sample for that market. This measure is designed to capture deviations from the long-run sustainable size in the market. This is a market-specific measure; market tightness in a given market is measured relative to that market's maximum size.

The investment policy functions are estimated using linear regression. Information about capacity targets is revealed in the data when either a new firm enters or an incumbent makes a capacity adjustment. The band is equal to the size of the adjustment for incumbents.[36] The results from estimation are shown in Table 5.

The parameters are generally estimated with precision for both the target and band equations; parameters having the expected signs. Higher market tightness is associated with lower levels of adjustment, while new entrants are more likely to enter at higher capacity levels, all else equal. Larger firms become increasingly larger than smaller firms conditional on making an adjustment. The adjustment band increases with current capacity and decreases with market tightness. The latter parameter implies that firms will be increasingly likely to make small adjustments as the market tightness increases, which is consistent with firms viewing the gains from delaying profitable investments as declining in the competitiveness of the market.

[36]The band is not relevant for the entry decision. In our model, firms are not allowed to enter without investing, so the statistical information associated with the decision to invest upon entry is captured by the fixed cost of entry.

Table 5: Investment Policy Estimates

	Estimate	**SE**
Target Equation		
Intercept	5.16	(0.45)
Entry Dummy	1.59	(0.47)
Ln Own Capacity	0.87	(0.06)
Market Tightness (MT)	-0.67	(0.20)
Target Variance	0.14	(0.02)
Band Equation		
Intercept	-0.20	(0.81)
Ln Own Capacity	1.02	(0.13)
Market Tightness (MT)	-1.53	(0.59)
Band Variance	0.64	(0.08)

Notes: Bootstrapped standard errors in parenthesis.

To estimate the entry and exit policy functions, we use probit regressions. We assume that there is at most one potential entrant in each period, while each incumbent firm has the opportunity to exit in each period. The explanatory variables are the same as above: intercepts, market tightness, and own capacity for current incumbents. The results from the estimation are shown in Table 6.

Our estimates reflect that entry is a low-probability event under most market circumstances. Market tightness has a large, negative, and precisely-estimated coefficient, reflecting that the probability of entry declines dramatically as relative market capacity grows. For example, when market tightness is 50 percent the probability of entry is 10.5 percent, while that declines to 1.1 percent when market tightness increases to 80 percent. Exit is also a rare event, although the relatively low exit probabilities are due in part to the fact that more firms take draws from the distribution of exit costs as compared to entry costs. Own capacity is negatively related to the exit probability, while market tightness increases the probability that a firm will exit the market. To put these numbers in context, we report exit probabilities at varying levels of market tightness and firm size.

Forward Simulation To simulate the firms' strategies going forward and compute their net present value, we first set the firms' discount factor, $\beta = 0.90$.[37] We then utilize the

[37] We have investigated setting the discount factor both higher and lower by re-estimating the model on a subset of the data. Table F.1 in the online appendix reports the estimates at alternative discount factors. As expected, lowering the discount rate to $\beta = 0.85$ leads to smaller investment costs, entry costs, and exit costs, while the opposite is true for raising it to $\beta = 0.95$, but the differences in the estimated parameters in

Table 6: Entry and Exit Policy Estimates

	Estimate	**SE**
Entry Equation		
Intercept	0.47	(0.60)
Market Tightness (MT)	-3.45	(0.84)
Probability Entry, $MT = 0.5$	0.105	(0.034)
Probability Entry, $MT = 0.8$	0.011	(0.004)
Exit Equation		
Intercept	2.23	(0.99)
Ln Own Capacity (Ln 000 Ton)	-0.74	(0.13)
Market Tightness (MT)	0.76	(0.53)
Probability Exit, $Cap = 800, MT = 0.5$	0.009	(0.003)
Probability Exit, $Cap = 1500, MT = 0.5$	0.002	(0.001)
Probability Exit, $Cap = 800, MT = 0.8$	0.017	(0.004)
Probability Exit, $Cap = 1500, MT = 0.8$	0.005	(0.002)

Notes: Bootstrapped standard errors in parenthesis.

forward-simulation procedure laid out in Bajari et al. (2007). The intuition behind their estimator is, first, to use forward simulation to compute expectations about future outcomes, given all firms' equilibrium strategies, and then, in a second step, to find parameters which make the observed behavior of firms consistent with profit maximization.

We forward simulate the continuation values under both the observed policy functions and four different perturbations. The first two perturbations manipulate when a firm invests: the first requires the firm to invest with certainty in the first period regardless of the draw of the fixed costs; the second is the mirror policy, where the firm is restricted to not invest. We also consider two alternative policies with (independent) marginal perturbations of both the probability of investment and the level of investment.[38] Additionally, we also impose a rationality constraint that the expected continuation value must be positive. Finally, we complement the inequalities with equalities derived from the indifference conditions for

this range were relatively minor. It is important to note that our results would become increasingly sensitive to our assumption about the discount rate as it grows towards one.

[38]We construct a sample of inequalities based on these perturbations. We take approximately 500 firm-market configurations based on years 1985, 1990, 1995, 2000 and 2005. For each of these 500 market configurations, we compute continuation values associated with each of the six inequalities by forward simulating market outcomes 1,000 times over a period of 30 years. This results in approximately 3,000 inequalities.

Table 7: Dynamic Cost Estimates

	Estimate	SE
Investment Estimates		
Capacity Investment Cost ($/Ton)	171	(55)
Adjustment Fixed Cost ($000)	48,525	(17,081)
Adjustment Fixed Cost SD (($000)	28,536	(8,298)
Adjustment Fixed Cost, 4% Draw	-1,433	(4,894)
Adjustment Total Cost, 1.4 MTon addition, 4% Draw	237,895	(73,550)
Entry Estimates		
Entry Fixed Cost ($000)	75,032	(79,823)
Entry Fixed Cost SD ($000)	27,948	(24,508)
Entry Fixed Cost, 2% Draw	17,633	(41,025)
Entry Total Cost, Plant 1 MTon, 2% Draw	188,582	(23,152)
Exit Estimates		
Exit Scrap Value ($000)	-151,825	(61,718)
Exit Scrap Value SD ($000)	89,231	(45,130)
Exit Scrap Value, 2% Draw	31,434	(37,239)
Exit Total Scrap, Plant 1 MTon, 2% Draw	202,382	(17,469)

Notes: Bootstrapped standard errors in parenthesis.

the marginal entering and exiting firms. Since all parameters enter linearly in the profit function, we use a robust solver (IBM's ILOG CPLEX Optimizer) that ensures that we find the globally-optimal solution. The results of the estimation are shown in Table 7.

Investment costs are roughly in line with the accounting costs cited in Salvo (2005), which reports a cost of $200 per ton of installed capacity. The implied cost of a cement plant is also in line with plant costs reported in newspapers and trade journals. For example, on October 15, 2010, it was reported that the most recent expansion of the Texas Industries New Braunfels cement plant, increasing capacity from 900 thousand tons per year to 2.3 million tons per year, was pegged at a cost of $276M in 2000 dollars, which implies a cost of $197 per ton of installed capacity, which is a little higher than our estimate of $171 per ton.[39]

The distribution of fixed costs of adjustment has an estimated mean of $48.5 million and a standard deviation of $28.5 million; the expected fixed cost of adjustment below the fourth percentile (the empirical rate of investment) is about -$1.43 million. The estimated variable investment costs for a 1.4 million ton per year expansion is $238 million, very close to the costs reported by the aforementioned New Braunfels plant.

[39]Source: KGNB Radio, New Braunfels, Texas.

For entry costs, we find that the distribution of entry fixed costs has a mean of \$75.0 million with a standard deviation of \$27.9 million. This implies that the entry costs at the second percentile, which is close to the empirical probability of entry, are equal to \$17.6 million. For an entrant who invests in a one million ton per year plant, this implies that the total initial investment outlays would be on the order of \$189 million.

For exiting firms, the estimated mean of the distribution of fixed exit costs is -\$152 million. This distribution has a standard deviation of \$89.2 million, which implies that firms receiving favorable draws will be paid to exit. This makes sense, as exiting firms are predicted to have positive profits, and therefore must perceive that their outside option is relatively favorable compared to staying as an incumbent. Combined with the sell off of capacity upon exit, the value to an exiting firm would be on the order of \$202 million.

Goodness of Fit The results above suggest that our model is broadly consistent with external measurements of firms' static and dynamic costs. Table 8 presents two additional measures of dynamic fit. To compare the performance across simulations, we take the market configuration in 2005 as our baseline, as it is the one used in our simulations and policy experiments.

The first column reports the empirical moments in the data. Capacities are measured in 2005, while changes in investment, and investment, entry and exit, rates are evaluated between 1981 and 2005. The second column reports the moments generated when we forward simulate the policy functions in the BBL estimation procedure. The third column reports moments generated when we solve and simulate the dynamic programming model under the baseline policy. Both sets of moments are generated by simulating the evolution of the industry 1000 times over a 30 year horizon and averaging outcomes, taking the configuration of markets in 2005 as the starting point. The main difference between the two is that the BBL policy simulation uses the policy functions estimated in the data to simulate firms decisions, whereas the simulation column uses policy functions from the computed equilibrium of the theoretical model.

The second column measures how well the BBL approach does in capturing the essential dynamics of the industry, which is important for the forward simulation in the estimation, while the third column measures how well our theoretical model explains and replicates those dynamics at the parameter estimates. These two measures of fit complement each other, as a good fit in the policy functions is necessary, but not sufficient, for a good fit in the theoretical model.

Table 8: Comparison of Actual and Simulated Moments

	Actual Data	**BBL Policy (30 years)**	**Simulation (30 years)**
Avg. Firm Capacity	1,224	1,126	1,128
SD Firm Capacity	727	465	426
Avg. Market Capacity	4,964	4,806	4,554
SD Market Capacity	1,928	1,509	1,540
Avg. Investment	376	112	393
Investment SD	394	64	243
Avg. Divestment	-151	-130	-643
Divestment SD	214	90	528
Investment Rate	0.045	0.066	0.015
Divestment Rate	0.039	0.009	0.014
Entry Rate	0.018	0.052	0.006
Exit Rate	0.019	0.011	0.002

Notes: Investment, divestment and capacities measured in thousands of tons. Capacity levels and market structure from 2005 is used for the comparison. See text for details on how the moments are generated.

Our model performs well in fitting the market structure in the data. We come close to matching the average market size and average firm size, as well as their standard deviations, with both the empirical policy simulations and the counterfactual model. This implies that our model and estimates are not artificially introducing some long-run trends in our simulations, which could be a concern.[40] These particular moments are vitally important, as consumer welfare and producer surplus directly depend on these outcomes.

We match the size of investment and divestment adjustments relatively well, although not perfectly. On the one hand, the BBL policy function does well at matching divestment, but understates the size of the investments. On the other hand, the simulations predict investment sizes very well, but overstate some of the divestment sizes, with few firms significantly reducing their size. The empirical investment and divestment rates are also somewhat larger than those predicted by the theoretical model, while the policy functions tend to modestly overstate investment rates. Importantly, while the investment rate is higher in the policy functions, this is offset by smaller predicted movements for each change. Conversely, the

[40] We use demand and import conditions as of 2000, to capture recent market conditions outside from the construction boom. Alternatively, we could have used average demand and import parameters during the whole period (1981-2005). In such case, we find that firms intuitively reduce their capacities (from 2005 levels), bringing them closer to average capacities in the data during the whole period. We decided to use the latest years as a baseline as we are ultimately interested in policy experiments going forward.

lower investment rate in the theoretical model is partially offset by larger adjustments when firms do make investments.

Two factors are worth keeping in mind when evaluating the goodness of fit of the investment, entry and exit rates. The empirical distribution of firms starts in 1981, when there were more, smaller, dirtier, and older firms as compared to 2000. The turnover rate of the industry in this earlier period was thus greater than it was in 2000. The lower entry and exit rates in our simulations can be partially explained by the fact that the industry has consolidated into fewer, larger, cleaner, and younger firms in the last three decades. A second issue is that our model may miss some important year-to-year fluctuations in the economic environment due to dynamic factors such as directed technical change, the housing bubble, and changing economic conditions in the import sector.

5.4 Environmental Parameters

Additional parameters in the simulation model include the social cost of carbon τ, the social discount factor, β_S, technology-specific emissions rates, and the policy parameters that define each allocation mechanism.

Social Cost of Carbon and Social Discount Rate Given the uncertainty inherent in the estimation of damages from carbon emissions, it is important to consider a range of values of τ. The range of values we choose to consider, \$5 to \$65 per ton of CO_2, is informed by an ongoing interagency process designed to produce estimates of the social cost of carbon (SCC) for use in policy analysis (Working Group on Social Cost of Carbon, 2013). Appendix D discusses the outcomes of this process and the typical social discount rates used in policy evaluation.

For expositional ease, we will assume that policies are designed such that the carbon price reflects the true social cost of carbon. Thus, the carbon tax or permit price and the social cost of carbon are assumed to be one and the same. In section 6.3, we relax this assumption and hold the assumed SCC value constant across scenarios associated with different permit prices/tax levels. We set the social discount rate β_S at 0.97.

Emissions Rates Both the IPCC and the World Business Council for Sustainable Development's Cement Sustainability Initiative (WBC, 2011) have developed protocols for estimating emissions from clinker production. We use these protocols to inform kiln technology-specific estimates of carbon dioxide emissions rates (denoted in tons of CO_2 per ton of

cement): 1.16 for wet process kilns, 0.93 for dry process kilns, and 0.81 for state-of-the-art kilns. Appendix C explains these emissions rate calculations in more detail. The emissions rate on imported cement, e_M, is estimated using an import volume weighted average of estimated foreign cement producers' emissions intensities (Worrell et al., 2001).

Policy parameters We begin by calibrating the policy designs we consider to match existing and proposed policy regimes. We then extend the analysis to consider design parameters that are more consistent with second-best policymaking.

In the grandfathering regime, firms receive an annual permit allocation proportional to their pre-program capacity level. We choose ψ_g such that this allocation is equal to 42.5 percent of their emissions-weighted initial capacity, which translates into approximately 50 percent of historic annual emissions.[41] In the auctioning regime, this ψ_g is set to zero.

In the policy scenario that incorporates output-based rebating, permits are allocated per unit of production based on an emissions intensity benchmark, denoted by ψ_d. We adopt the benchmark that was chosen for European cement producers in the third phase of the EU ETS (2013-2020): 0.716 permits per metric ton of clinker.[42] This translates into a reduction in compliance costs (per unit of clinker output) of between 62 percent for wet kilns to 77 percent for the most common dry-process kiln technology.[43]

6 Simulation Results

This section begins with a summary of how key market outcomes (domestic production capacity, cement prices, emissions) are affected by the introduction of market-based policies designed to reduce greenhouse gas emissions. All simulation results are summarized relative to the base case in which greenhouse gas emissions are unregulated. We then summarize the net welfare effects of the policies. The section concludes with a discussion of optimal carbon pricing, abatement curves and a series of robustness checks.

We report simulation results graphically, for the range of SCC values that have been deemed policy relevant (Working Group on Social Cost of Carbon, 2013).[44] However, our in-

[41] The utilization rate of cement kilns is around 85% in our sample and very homogeneous across plants.

[42] Available from http://eurlex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2011:130:0001:0045:EN:PDF (accessed 6/21/2013). For comparison, in California's Greenhouse Gas Trading Program a more generous benchmark of 0.786 allowances per metric ton of clinker is used.

[43] Importantly, permits are rebated based on clinker (versus cement) production, thus eliminating incentives to reduce carbon intensity through increased use of supplementary cementitious materials.

[44] See Table F.4 for numerical values for a sample of markets and carbon price levels, which include

ferences at high carbon prices are quite far from historical experience.The higher the assumed carbon price, the less plausible our partial equilibrium approach and all of the implications that come with it, such as the exogenous demand parameters, capital costs, and productive technology. This caveat notwithstanding, evaluating outcomes over this range of SCC values serves to illustrate the countervailing forces that shape interactions between market structure and carbon regulation.

To highlight the importance of accounting for industry dynamics, we contrast the results of our dynamic simulations with a simulation exercise that holds industry structure fixed. We take a standard approach to constructing the static benchmark (U.S. Environmental Protection Agency Office of Air Quality Planning and Standards Innovative Strategies and Economics Group, 1999). We simulate equilibrium outcomes in a single period and assume that these simulated static outcomes would be observed each year of the 30 year time horizon. In this static model, firms can alter production levels, but production capacity, technology operating characteristics, etc. are held constant at baseline levels.

6.1 Simulated Market Outcomes

Production capacity Figure 2a plots total domestic production capacity (summed across markets and averaged across years) as a function of the exogenous permit price, τ. The left panel, which corresponds to the static simulations, highlights the fact that domestic production capacity is held fixed at baseline levels in the static model.

The right panel shows how domestic production capacity varies with the carbon price once industry dynamics are introduced. Policy-induced reductions in installed capacity are most pronounced under the auctioning/tax regime. As τ increases, a growing number of firms elect to disinvest or exit the market completely. Augmenting this policy with a border tax adjustment mitigates the loss of domestic market share to foreign producers, thus slowing the rate of exit and disinvestment.

One important result, highlighted by this and subsequent figures, is that equilibrium outcomes under the grandfathering and auctioning regimes differ substantively. In other words, the so-called independence property fails to hold when industry dynamics are accounted for.[45] Under the grandfathering regime, an incumbent firm receives a lump sum transfer each pe-

bootstrapped standard errors on the simulation results.

[45]The independence property states that the market equilibrium in a cap-and-trade system is independent of the initial distribution of emissions permits. This is closely related to the more general principle that when markets are complete, outcomes remain efficient even after lump-sum transfers among agents (Hahn and Stavins, 2011)

Figure 2: Market Outcomes

(a) Capacity

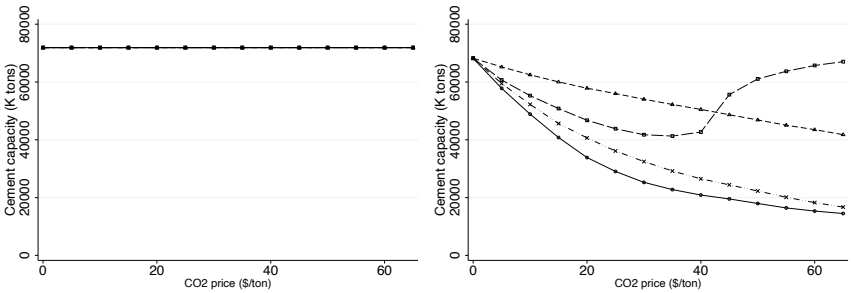

(b) Cement Prices

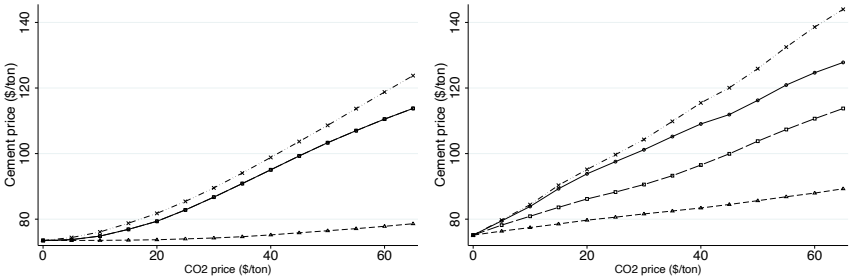

(c) Domestic Emissions

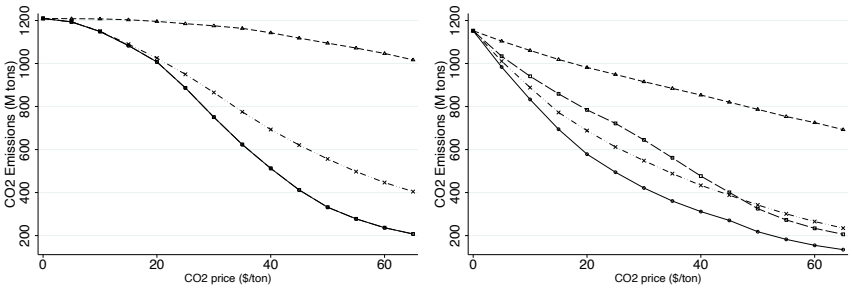

(d) Emissions Leakage

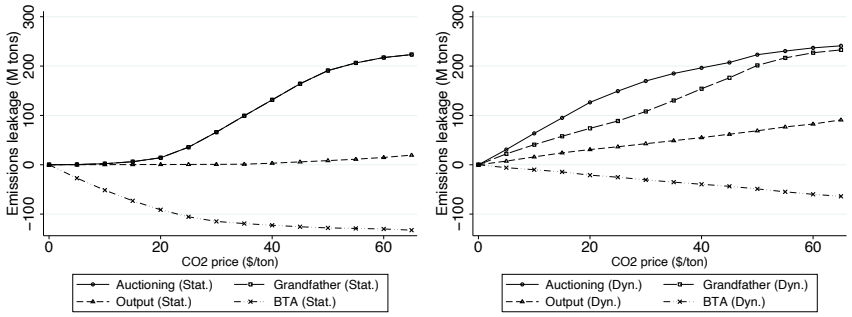

These figures show the static (left panel) and discounted present values (right panel) of industry-wide capacity, average prices, total domestic emissions, and total emissions leakage, respectively. The dynamic figures were calculated by simulating the industry forward 30 years and discounting using a rate of three percent per year.

riod in the form of a free permit allocation. The firm forfeits this entitlement if it chooses to exit or disinvest. This lowers the exit and disinvestment thresholds for incumbents vis a vis the auctioning regime. As carbon prices increase, the allocation of grandfathered permits becomes increasingly valuable, and firms find it increasingly profitable to forego production and simply sell their allocations on the open market; this explains the non-monotonic level of capacity under the grandfathering regime. At very high values of τ, permit endowments are so valuable that domestic production capacity remains at baseline levels.

Another noteworthy result pertains to the policy that incorporates the output-based subsidy. When compared to the auctioning regime, output-based updating induces much smaller reductions in domestic production capacity. This is due to the fact that contingent rebating confers an implicit subsidy of 0.716 permits per unit of production. For a firm of an average emissions intensity of 0.97, the net cost to the firm is only 25 percent of the carbon price. Thus, the equilibrium production capacity under the output-based rebating regime at any given carbon price is roughly the capacity level observed under auctioning at a carbon price four times smaller.[46]

Cement prices Figure 2b plots quantity-weighted average cement prices as a function of τ. In both the static and dynamic simulations, cement price increases are most pronounced under the auction/tax regime that incorporates a border tax adjustment. Under this policy, both foreign and domestic firms bear the complete cost of compliance; no compensation in the form of contingent rebates or lump sum transfers is offered.

A notable feature in the static left panel of Figure 2b is that the cement price is virtually unaffected at carbon prices below \$15. In the benchmark case, many domestic firms are capacity constrained and earning scarcity rents. An increase in variable operating costs reduces scarcity rents, but does not affect domestic production levels or equilibrium prices. In contrast, when firms have the ability to disinvest in response to an increase in operating costs, we observe price impacts even at low levels of τ.

Cement price increases are more significant in the dynamic simulations. As firms reduce production capacity through divestment and/or exit in response to policy-induced increases in operating costs, regional cement markets become more concentrated, and the distortions associated with the exercise of market power became more pronounced. Whereas there is a distinct increase in production capacity under grandfathering at higher carbon prices, there is no associated decrease in cement prices. The reason is that capacity is relatively

[46]Note that, if all firms had the same emissions rate e, the output-based updating regime at a price τ would be exactly equivalent to the auctioning regime at a price $\frac{e-0.716}{e}\tau$.

under-utilized at high carbon prices in the grandfathering regime. Grandfathering creates an incentive to remain in the market so as to maintain the permit entitlement.

Domestic emissions Figure 2c shows how the emissions from domestic cement production decrease with τ. The vertical axes measure domestic CO_2 emissions summed across regional markets and averaged across time periods. Domestic emissions are lowest under the auctioning regime which provides domestic producers no compensation for the costs they incur to comply with the regulation. This drives down levels of domestic cement production and associated emissions. Augmenting the auctioning regime with a border tax adjustment mitigates impacts on domestic competitiveness, thus increasing both domestic production levels and emissions.

In the static simulations, emissions outcomes are identical across the grandfathering and auctioning regimes. In the dynamic simulations, domestic emissions levels are higher under grandfathering. Intuitively, regional cement markets have a higher expected number of active firms under grandfathering, leading to higher levels of domestic production and associated emissions.

Emissions leakage Figure 2d summarizes policy-induced changes in emissions from foreign producers. To compute these emissions leakage measures, we assume the increase in demand for cement imports represents purely additional production at foreign suppliers, rather than a reallocation of foreign production that once supplied their local markets. We revisit this assumption in the next sub-section.

Focusing on the dynamic simulations (right panel), emissions leakage is most significant under the auctioning regime. Domestic producers are required to fully internalize the externality with no compensation, whereas the operating cost structure of foreign producers is unaffected. As foreign producers gain market share, emissions from foreign cement production increase vis a vis the baseline. Grandfathering slows the rate of domestic capacity reduction vis a vis auctioning, mitigating emissions leakage. Similarly, output-based rebating significantly reduces the net cost of compliance per unit of output, also reducing leakage.

Notably, we find negative leakage rates under the regime that incorporates a border tax adjustment. In other words, the introduction of this policy reduces emissions among foreign producers relative to the unregulated baseline. This occurs because we assume complete pass-through of environmental compliance costs by foreign producers whereas pass through of environmental compliance costs among strategic domestic producers is incomplete. Consequently, when emissions from domestic and foreign producers are penalized at the same

Figure 3: Market Outcomes over time

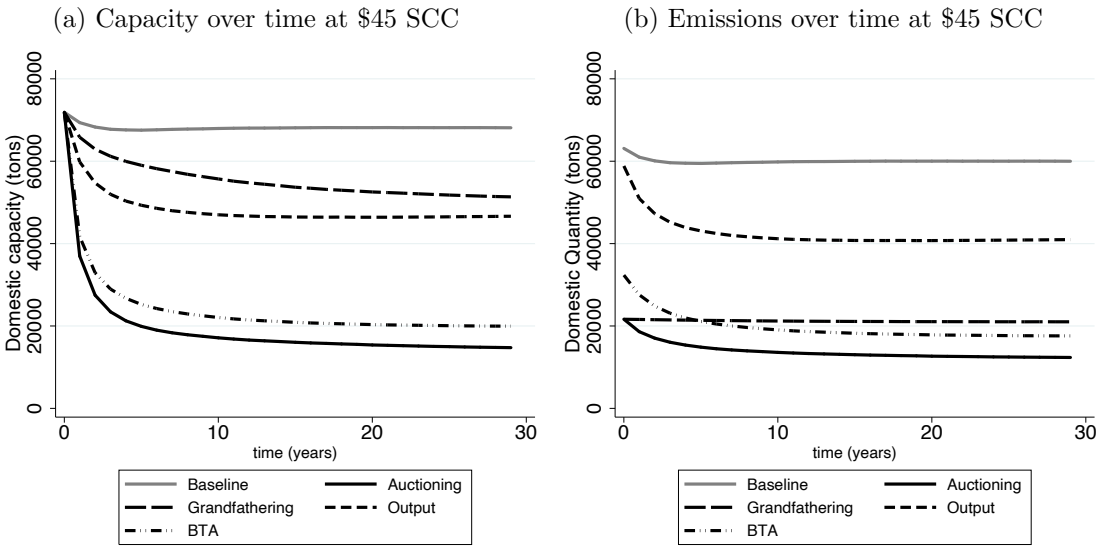

These figures show the evolution of capacities over time (average across markets). Initial market structure is the same across mechanisms (see initial capacity at the same level). The figures show the transition over a period of 30 years.

rate, we see a decrease in cement imports. The extent of negative leakage is reduced when dynamic industry responses are accounted for. This is because policy-induced increases in the cement price are larger, resulting in higher import supply levels at any given carbon price.

Market outcomes over time Our dynamic simulation model can also be used to generate trajectories of market outcomes over time under alternative policy regimes. Figures 3a and 3b chart the evolution of domestic production capacity and domestic quantity, respectively, assuming a carbon price of $45 per ton of CO_2.

In our model, there is no technological innovation over time (except through entry of new efficient plants), nor is there growth in domestic cement demand over time. In other words, aside from policy-induced changes in market structure, economic operating conditions are stable over the 30 year time horizon we consider. Consequently, most of the industry response to a counterfactual policy intervention occurs in the years immediately following the policy change. This adjustment is not immediate due to year-to-year variation in firms' draws from the distributions of investment, entry, and exit costs. It is also notable that the adjustment takes longer in the grandfathering case, where incentives to divest are attenuated

by the payoffs of keeping free allowances. These graphs also show that these outcomes are very stable in the baseline case, which is reassuring and suggestive that our simulations are internally consistent with our assumption that the economic environment is unchanging in the baseline.[47]

The graphs in Figure 3a illustrate how the distributional differences arising in the short-run between auctioning and grandfathering can have product market implications in the longer run. At period zero, both production and capacity are the same, as the static production incentives of the two mechanisms are equivalent. However, as time passes, firms disinvest and exit at faster rates under the auctioning regime. Capacity constraints bind and production in the auctioning case falls below production in the grandfathering case. Note that differences in capacity across the auctioning and grandfathering regimes are far more stark than differences in production. Production capacity is under-utilized in the grandfathering case; firms have an incentive to keep their grandfathered investments in operation, even if they are not fully utilized.

6.2 Decomposing Changes in Welfare

Having considered the effects of counterfactual emissions regulations on specific market outcomes, we next consider the related welfare implications of these policies. Policy-induced welfare changes are decomposed into the three component parts introduced in Section 2.

W1: Domestic Economic Surplus As a point of departure we consider policy-induced changes in domestic producer surplus, domestic consumer surplus, and any revenues raised by the government through emissions taxation or permit sales. The left panel of Figure 4a corresponds to the static case. Because short run production incentives are identical under grandfathering and auctioning, impacts on domestic economic surplus are identical. The addition of a border tax adjustment improves terms of trade, generates border tax revenues, and reduces policy impacts on cement prices. On balance, this mitigates losses in domestic economic surplus at high carbon prices. Because the policy that incorporates output-based rebating has only minor impacts on domestic production across the range of prices we consider, impacts on domestic economic surplus are minimal.

The right panel of Figure 4a summarizes the corresponding dynamic results. Reductions in domestic economic surplus are most significant under the auctioning regime where we

[47] This is not necessarily the case; misspecification bias in our model could imply that firms should systematically be larger or smaller than their empirical counterparts, for example.

Figure 4: Welfare Measures across Mechanisms

(a) W1: Domestic Industry Surplus + Government Revenues

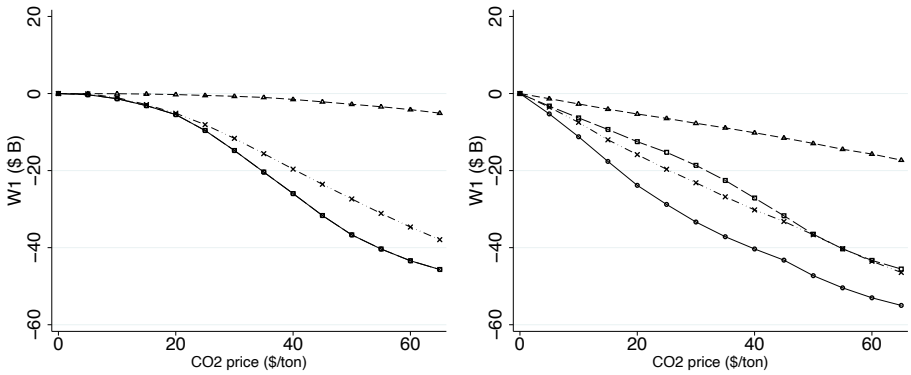

(b) W2: W1 + Domestic Emissions Reduction

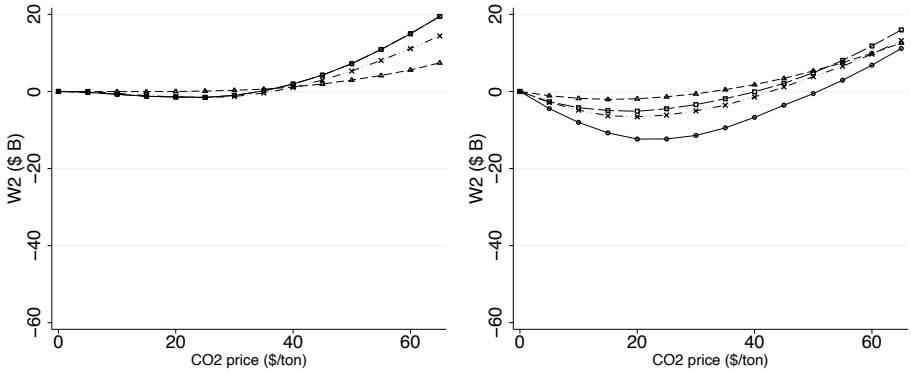

(c) W3: W2 + Emissions Leakage

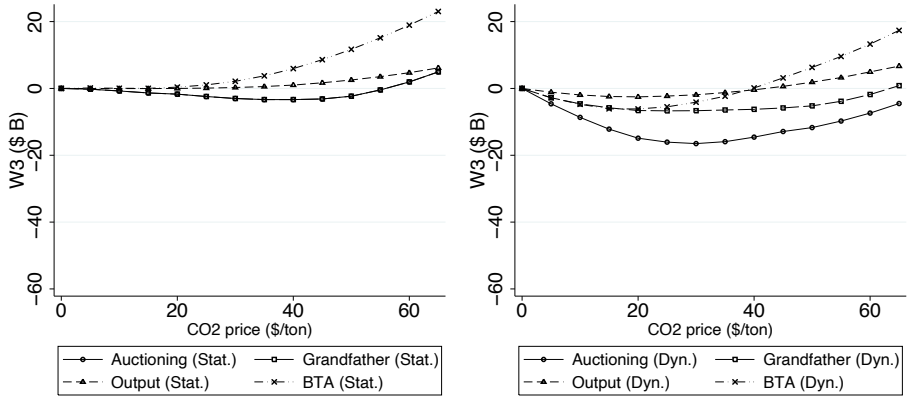

observe the highest rates of exit and disinvestment, the highest cement prices, and the most significant adverse impacts on domestic competitiveness. Under the grandfathering regime, higher levels of domestic production and lower cement prices deliver a relative increase in domestic producer and consumer surplus.

In contrast to the static case, reductions in economic surplus manifest even at low carbon prices. As discussed above, when firms have the ability to disinvest in response to a policy-induced increase in operating costs, we observe impacts on cement prices, domestic production, and thus domestic economic surplus, across the range of carbon prices we consider.

W2: Domestic Economic Surplus + Domestic Emissions Figure 4b plots changes in our second welfare measure which incorporates the value of domestic CO_2 emissions reductions. The value per ton of emissions avoided is assumed to be equal to the prevailing permit price or tax. Thus, the monetary value of domestic emissions reductions is constructed by multiplying the emissions reductions summarized in Figure 2c by the corresponding permit price.

In the static simulations (left panel), benefits associated with reduced domestic emissions do not offset the costs of a policy that incorporates grandfathering or auctioning at carbon prices below \$40. In contrast, the value of domestic emissions reductions more than offsets the economic costs under the policy regimes that incorporate a border tax adjustment or the output-based rebate.

The dynamic simulations yield quite different results (right panel). As compared to the static case, the dynamic mechanisms of divestment and exit result in much smaller levels of production; at low carbon prices, the loss in domestic economic surplus is increasing faster than the gain in benefits. However, as τ increases, the gains from emissions abatement begin to offset losses in economic surplus. All policy regimes yield welfare gains at high carbon prices.

W3: Domestic Economic Surplus + Domestic Emissions + Leakage Figure 4c plots the policy induced reductions in this most comprehensive welfare measure. In the static simulations (left panel), accounting for the significant levels of emissions leakage observed at values of τ greater than \$20 exacerbate welfare costs of the grandfathering and auctioning regimes. In contrast, accounting for negative leakage amplifies the welfare gains under the policy regime that incorporates a border tax adjustment.

In the dynamic simulations (right panel), accounting for leakage decreases welfare in most cases. Output-based updating is the least-worst (but still negative) policy for the majority of carbon prices, being eclipsed by border tax adjustments only at prices exceeding $45 per ton. Grandfathering generates marginally greater surplus relative to border tax adjustments for low to moderate carbon prices. The auctioning/carbon tax regime generates large and negative welfare effects over the entire range of carbon prices we evaluate. Notably, the highest welfare losses, in the range of $15-18 billion, correspond to carbon prices in the middle of the range of expected carbon prices for a US-wide carbon trading scheme.

As noted above, we assume that policy-induced changes in demand for cement imports translate directly into changes in the levels of foreign cement production. This assumption will exaggerate the impacts of these policies on emissions leakage if foreign producers accommodate changes in domestic demand for cement imports by reallocating their output. In this respect, Figures 4b and 4c can be viewed as upper and lower bounds on the welfare effects of these policies.

6.3 Policy Comparisons Under Optimal Carbon Prices

Simulation results summarized in the previous section suggest that the negative welfare effects of fully internalizing the emissions externality outweigh the benefits over a range of carbon prices. As a result, a policy maker looking to maximize welfare will want to set a permit price that falls below the true social cost. This insight helps explain why a regime that dynamically updates permit allocations to domestic producers based on output welfare dominates a regime that allocates permits to domestic producers in lump sum. As perceived by domestic firms, dynamic allocation updating lowers the effective cost per unit of emissions below the social marginal cost.

Across the four policy regimes we consider, we compute the permit price that maximizes our most comprehensive welfare measure (W3) for a given value of the true social cost of carbon. We first impose the constraint that all domestic cement producers must be treated symmetrically under the regulation. In the debates over carbon policy design and implementation, it is typically assumed that firms within a sector will face the same policy incentives. Given the structural differences across regional markets, as well as the differences in trade exposure, allowing policy incentives to vary across regional markets could be welfare improving. We therefore extend the analysis to consider policy designs that levy different carbon prices for trade-exposed coastal and trade-insulated inland markets.

Table 9 reports welfare maximizing carbon prices and associated welfare changes at two

Table 9: Optimal carbon prices for different mechanisms

	Federal τ_f^*	Coastal τ_c^*	Inland τ_i^*	Welfare Δ at τ_f^*	Welfare Δ at $\{\tau_c^*, \tau_i^*\}$	Welfare Δ at τ = SCC
SCC = \$ 20						
Auctioning	0	0	0	0	0	-14,886
Grandfather	0	0	0	0	0	-6,609
Output	0	0	0	0	0	-2,519
BTA	0	0	0	0	0	-6,141
SCC = \$ 45						
Auctioning	5	5	15	905	1,316	-12,890
Grandfather	10	5	35	1,357	2,259	-5,839
Output	25	15	60	1,047	1,628	619
BTA	20	25	15	5,991	6,269	3,150

Notes: Carbon prices in \$. Welfare in M\$. Optimal carbon prices computed on a grid including {0, 5, 10, 15, 20, 25, 30, 35, 40, 45, 50, 55, 60, 65}.

medium-range SCC values (\$20 and \$45). In Column 1, we impose the constraint that all cement producers face the same price. Columns 2 and 3 report the optimal prices for coastal and inland regional markets, respectively. The top panel considers the case in which the social cost of carbon is \$20 per ton of CO_2. At this value, there is no positive carbon price at which the benefits from emissions reductions exceed the costs. This is true in inland markets and in coastal markets when the emissions externality has been internalized by foreign producers. This implies that the social costs of exacerbating the exercise of market power exceeds any social gains from reducing emissions.

The bottom panel of Table 9 conducts the same analysis setting the SCC to \$45. Under the auctioning regime, the optimal permit price falls well below the true cost of carbon in order to strike the right balance between incentivizing abatement and exacerbating the distortions associated with the exercise of market power and the asymmetric treatment of domestic and foreign emissions. When this price is allowed to vary across inland and coastal markets, the price is much lower in trade-exposed markets in order to address the welfare effects of emissions leakage.

Augmenting the auctioning regime with a border tax adjustment efficiently internalizes the emissions externality associated with foreign production, but leaves the distortions associated with the exercise of market power unaddressed. In coastal markets, augmenting the auctioning regime with a border tax adjustment increases the optimal carbon price from \$5/ton to \$25/ton. Note that this is higher than the optimal price in inland markets because

coastal markets tend to be relatively more competitive.

Under the regime that incorporates dynamic allocation updating, there is substantial heterogeneity in the optimal permit price between coastal and inland markets. The implicit production subsidy appears to be too low in coastal markets, as output-based updating plays a crucial role in attenuating rents and emissions leakage. On the contrary, in a regime in which all domestic firms must be treated symmetrically, this subsidy may be overly generous as suggested by the optimal inland price of $60, which is near the upper bound on the range that we consider.

The welfare change that results if carbon is priced optimally and uniformly within the cement sector is reported in Column 4. Column 5 reports the welfare change that results under the differentiated carbon price. Finally, as a basis for comparison, Column 6 reports the welfare change that results if the carbon price is constrained to equal the assumed SCC. In general, moving from complete internalization of the emissions externality to a regime that implements the optimal uniform carbon price confers sizeable welfare gains. The additional gains from differentiating carbon prices across inland and coastal markets are not as large, but are non-trivial.

6.4 Abatement Cost Curves

Throughout the analysis, we have assumed that the the permit price τ reflects the marginal cost of abatement in other competitive sectors covered by the regulation, and that the permit supply curve is flat in the neighborhood of the imposed cap. An alternative approach to summarizing the relative welfare effects eschews these assumptions. We compute the average cost per ton of CO_2 emissions abatement in the cement sector across different policy regimes. This provides a measure of the efficiency with which cement sector emissions abatement is achieved under alternative policy designs.

Figure 5 presents the average abatement cost for the mechanisms that we consider. Each point on the graph corresponds to a specific policy regime and carbon price. The graph on the left divides reductions in domestic producer and consumer surplus (W1) by the reduction in domestic emissions. The graph on the right conducts a similar exercise, although the denominator is adjusted to reflect the policy-induced change in foreign emissions.

Consistent with our findings above, cement-sector emissions abatement is least efficient under the auctioning regime. Under this regime, with firms bearing the full brunt of compliance costs, distortions associated with market power are exacerbated and any reductions in emissions that do occur come at the cost of significant surplus reductions. Average abate-

Figure 5: Abatement Curves

(a) Abatement Average Cost (leakage ignored) (b) Abatement Average Cost (leakage-corrected)

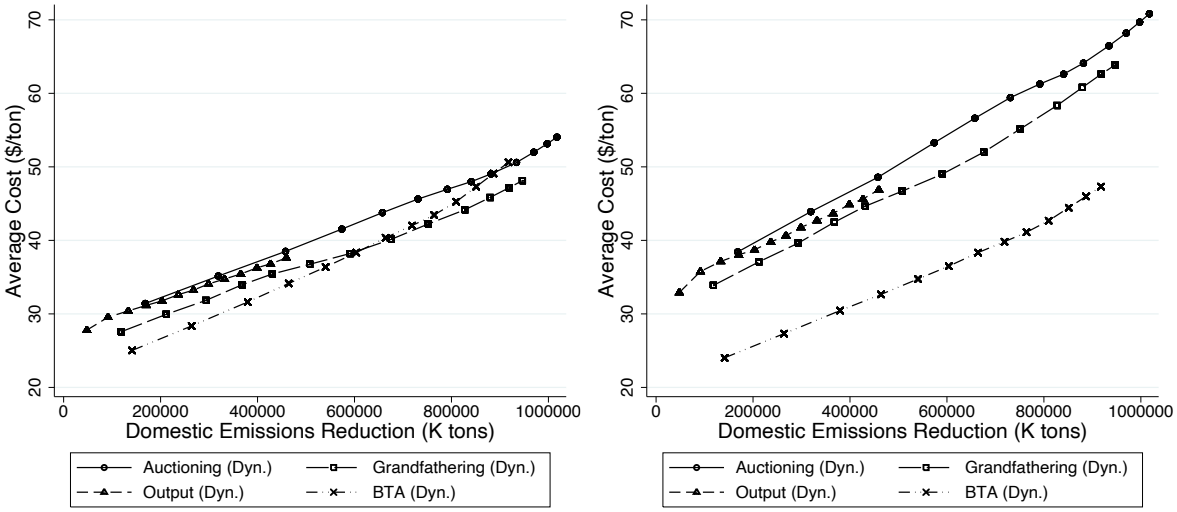

ment costs start at close to $40 per ton once leakage is accounted for. Abatement under the grandfathering regime is somewhat more cost effective because the lump sum transfer provides an incentive for firms to remain in the market (reducing market power distortions vis a vis auctioning). The relatively low average abatement costs associated with the border tax adjustment are striking once the effects of the policy on foreign emissions are accounted for. This reflects both terms of trade improvements and negative leakage.

6.5 Sensitivity Experiments

In the previous analysis, we have maintained fixed demand and import supply elasticities. These estimates vary across econometric specifications and can be quite noisy. Furthermore, one could imagine these elasticity values could change endogenously as carbon policy becomes more stringent.

In order to assess the robustness of our estimates to different elasticities, we re-compute all the calculations of market equilibria for a range of elasticity values. For computational reasons, we focus our attention to the smaller markets. We analyze how the welfare metric W3 changes with these parameters. To summarize, we find that the main results and comparative statics across allocation mechanisms are robust to changes in demand and import

elasticities.

Demand elasticities The demand elasticity plays an important role in determining, among other outcomes, gross consumer surplus, the extent of the distortion arising from the exercise of market power, and the extent to which leakage occurs under a given emissions policy. Table F.5 in the online appendix presents welfare changes (W3) for a combination of carbon prices and demand elasticities. In the columns, we report equilibrium outcomes for the different sensitivities. In the rows, we present welfare outcomes for different allocation mechanisms and carbon price values. Our baseline results are reported in the middle column ($\eta = 2$).

For low carbon prices, welfare effects of the policies we consider are more negative when demand is relatively more elastic because a given permit price has a larger impact on economic surplus. At higher carbon prices, negative welfare effects are attenuated, or turn positive, when demand is more elastic. Intuitively, as reductions in emissions play a more significant role in determining welfare at higher carbon prices, welfare effects start increasing with the elasticity of demand.

Aside from these changes in levels, it is important to emphasize that relative welfare comparisons across mechanisms do not generally change with the demand elasticity. As before, for low levels of carbon prices, an output-based updating allocation dominates, whereas for larger carbon prices the BTA mechanism becomes more attractive. The auctioning mechanism, on the other hand, is the least favorable across specifications.

Table F.5 can also be used to address concerns about the effects of carbon policy on the structure of domestic cement demand. Whereas our model effectively holds demand shifters constant, we might expect that the emissions policies we consider would affect the prices of cement substitutes. One can use Table F.5 to get a rough idea of how our estimates of welfare effects within the cement sector may change as the structure of demand changes.

If cement will become differentially more expensive (as compared to substitutes) as carbon prices rise, one can imagine that demand elasticity will become larger. Therefore, one can simply start the baseline elasticity at the zero carbon price and trace down the table, letting the elasticity increase with the carbon price. While this approach does not explicitly model interactions between climate policy and markets for cement substitutes, it provides a simple way to examine the sensitivity of our results to our partial equilibrium modeling assumptions. Intuitively, if demand becomes more elastic as carbon prices increase due to increased substitution, the benefits from the policy are larger.

Import supply elasticities The import supply elasticity parameter is another key parameter in our model that is not precisely estimated. Similar to the own-price elasticity of domestic cement demand, there is also the possibility that importing firms could respond to the policy by expanding investment in import terminals, foreign production capacity, or improved transport practices. By allowing for a more or less responsive supply curve, we proxy for these kinds of responses.

Table F.6 in the online appendix recomputes estimated welfare effects for a range of import supply elasticity values. Our baseline results are reported in the middle column ($\eta = 2.5$). Changing the import supply elasticity has two important implications. First, in trade-exposed markets, an increase in the import supply elasticity increases the elasticity of the residual demand curve faced by domestic producers, all else equal, which can be beneficial for competition. Second, the more responsive is import supply to a change in the cement price, the greater the emissions and rent leakage. We find that the latter effect dominates for most mechanisms and price ranges, and thus a more elastic import supply reduces welfare. Notably, this is not the case for the BTA mechanism, which is most effective at mitigating leakage. The qualitative findings in the paper are robust to these effects.

7 Conclusion

We use an empirically tractable dynamic model of the domestic Portland cement industry to evaluate the welfare impacts of incomplete, market-based regulation of carbon dioxide emissions. We assess the implications of several alternative policy designs, including those that incorporate both an emissions disincentive, in the form of a tax or an obligation to hold an emissions permit, and a production incentive.

We find that both the magnitude and the sign of the welfare effects we estimate depend critically on how the policy is implemented and what we assume for the social cost of carbon (SCC). Under market-based policy regimes that incorporate neither a border tax adjustment nor an implicit production subsidy, our results echo Buchanan (1969). At SCC values below \$40/ton CO_2, market-based emissions regulation that requires domestic producers to fully internalize the emissions externality exacerbates the distortions associated with the exercise of market power in the domestic product market to such an extent that reductions in domestic economic surplus exceed the benefits of emissions reductions. Emissions leakage in trade-exposed regional markets further undermines the benefits of these programs.

Notably, we find that policy designs that incorporate both an emissions penalty and

a production incentive in the form of a rebate welfare dominate more traditional policy designs at SCC values below $40. Intuitively, the production incentive works to mitigate leakage in trade-exposed cement markets and the distortion associated with the exercise of market power. A policy that penalizes emissions embodied in foreign imports induces negative leakage given our assumption that imports respond competitively, whereas domestic producers behave strategically. Consequently, this policy delivers sizeable welfare gains at high carbon prices.

Policy makers are very interested in understanding how proposed climate change policies would impact highly concentrated, emissions-intensive sectors such as the cement industry. The scale and scope of these policy interventions are unprecedented, making it difficult to anticipate how industry will respond and what that response will imply for social welfare. This paper illustrates important forces that shape the interaction of industry structure, trade flows, and proposed carbon regulations. Our results provide valuable insights into the efficiency and distributional properties of leading policy design alternatives.

References

Arcidiacono, P., Bayer, P., Bugni, F., and James, J. (2012). Approximating high-dimensional dynamic models: Sieve value function iteration. NBER Working Papers 17890, National Bureau of Economic Research.

Bajari, P., Benkard, C. L., and Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, 75:1331–70.

Bernard, A. L., Fischer, C., and Fox, A. (2007). Is there a rationale for output-based rebating of environmental levies? *Resource and Energy Economics*, 29(2):83 – 101.

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

Buchanan, J. M. (1969). External diseconomics, corrective taxes, and market structure. *The American Economic Review*, 59(1):174–7.

Burtraw, D. (2011). personal communication.

Dales, J. H. (1968). *Pollution, Property, and Prices*. University of Toronto Press. Toronto.

Dardati, E. (2013). Pollution permit systems and firm dynamics: Does the allocation scheme matter? Technical report, Facultad de Economa y Negocios - ILADES, Universidad Alberto Hurtado, Chile working paper.

Demailly, D. and Quirion, P. (2006). CO2 abatement, competitiveness and leakage in the european cement industry under the eu ets: grandfathering versus output-based allocation. *Climate Policy*, 6(1):93–113.

Doraszelski, U. and Pakes, A. (2007). A framework for applied dynamic analysis in io. In Armstrong, M. and Porter, R., editors, *Handbook of Industrial Organization*, volume 3, chapter 30, pages 1888–1966. Elsevier.

Ebert, U. (1992). Pigouvian tax and market structure: The case of oligopoly and different abatement technologies. *FinanzArchiv / Public Finance Analysis*, 49(2):154–166.

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

Farias, V., Saure, D., and Weintraub, G. (2012). An approximate dynamic programming approach to solving dynamic oligopoly models. *RAND Journal of Economics*, 43:253–282.

Fischer, C. and Fox, A. K. (2009). Comparing policies to combat emissions leakage: Border tax adjustments versus rebates. Discussion Papers dp-09-02, Resources For the Future.

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

Hahn, R. and Stavins, R. (2010). Forthcoming.the effect of allowance allocations on cap-and-trade system performance.. *Journal of Law and Economics*.

Holland, S. P. (2012). Emissions taxes versus intensity standards: Second-best environmental policies with incomplete regulation. *Journal of Environmental Economics and Management*, 63(3):375–387.

International, R. (2009). Regulatory impact analysis: National emissions standards for hazardous air pollutants from the portland cement manufacturing industry.

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

Joskow, P. and Schmalensee, R. (1998). Political economy of market-based environmental policy: The us acid rain program, the. *JL & Econ.*, 41:37.

Kapur, A., van Oss, H., Keoleian, G., Kesler, S., and Kendall, A. (2009). The contemporary cement cycle of the united states. *Journal of material cycles and waste management*, 11(2):155–165.

Keohane, N. O. (2009). Cap and trade, rehabilitated: Using tradable permits to control u.s. greenhouse gases. *Review of Environmental Economics and Policy*, 3(1):42–62.

Levin, D. (1985). Taxation within cournot oligopoly. *Journal of Public Economics*, 27(3):281–290.

Mahasenan, N., Dahowski, R., and Davidson, C. (2005). *Greenhouse Gas Control Technologies, Volume I*, chapter The Role of Carbon Dioxide Capture and Sequestration in Reducing Emissions from Cement Plants in North America. Elsevier Science.

Maskin, E. and Tirole, J. (1988). A theory of dynamic oligopoly I: Overview and quantity competition with large fixed costs. *Econometrica*, 56:549–69.

Miller, N. H. and Osborne, M. (2010). Competition among spatially differentiated firms: An estimator with an application to cement. EAG Working Paper 10-2.

Millimet, D. L., Roy, S., and Sengupta, A. (2009). Environmental regulations and economic activity: Influence on market structure. *Annual Review of Resource Economics*, 1:99–117.

Montgomery, D. W. (1972). Markets in licenses and efficient pollution control programs. *Journal of Economic Theory*, 5(3):395–418.

Oates, W. E. and Strassmann, D. L. (1984). Effluent fees and market structure. *Journal of Public Economics*, 24(1):29–46.

on Sustainalble Development, W. B. C. (2010). Cement Technology Roadmap 2009. Technical report, WBCSD.

PCA (1980-2006). U.S. and Canadian Portland Cement Industry: Plant Information Summary. Technical report, Portland Cement Association.

Ponssard, J. P. and Walker, N. (2008). Eu emissions trading and the cement sector: a spatial competition analysis. *Climate Policy*.

Ryan, S. P. (2012). The cost of environmental regulation in a concentrated industry. *Econometrica*, 80(3):1019–61.

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

Sato, M., Neuhoff, K., and Neumann, A. (2008). Demand response from CO_2 prices in energy-intensive industries. TU Dresden Working Paper.

Sterner, T. and Hoglund, L. (2000). Output-based refunding of emission payments: Theory, distribution of costs, and international experience. Discussion Papers dp-00-29, Resources For the Future.

Szabo, L., Hidalgo, I., Ciscar, J. C., and Soria, A. (2006). Co2 emission trading within the european union and annex b countries: the cement industry case. *Energy Policy*, 34(1):72–87.

U.S. Environmental Protection Agency Office of Air Quality Planning and Standards Innovative Strategies and Economics Group (1999). OAQPS economic analysis resource document. Technical report, U.S. Environmental Protection Agency.

USGS (2010). Mineral commodity summaries 2010. Technical report, U.S. Geological Survey.

USGS (2012). Mineral commodity summaries 2012. Technical report, U.S. Geological Survey.

Van Oss, H. and Padovani, A. (2002). Cement manufacture and the environment. *Journal of Industrial Ecology*, 6(1):89–106.

Van Oss, H. and Padovani, A. (2003). Cement manufacture and the environment, part ii: Environmental challenges and opportunities. *Journal of Industrial Ecology*, 7(1):93–126.

van Oss, H. G. (2005). Background facts and issues concerning cement and cement data. Technical report, U.S. Department of the Interior, U.S. Geological Survey.

WBC (2005,2011). CO2 Accounting and Reporting Standard for the Cement Industry. Technical report, World Business Council for Sustainable Development.

Working Group on Social Cost of Carbon, I. (2013). Technical support document: Technical update of the social cost of carbon for regulatory impact analysis under executive order 12866. Technical report, United States Government.

Worrell, E., Price, L., Martin, N., Hendriks, C., and Meida, L. O. (2001). Carbon Dioxide Emissions from the Global Cement Industry. *Annual Review of Energy and the Environment*, 26(1):303–329.