Pollution Permits and the Evolution of Market Structure

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

November 4, 2010

Abstract

We explore the long run dynamic implications of subjecting an imperfectly competitive industry to market-based pollution regulation. We are particularly interested in understanding how the allocation of emissions permits in a cap-and-trade program can influence the evolution of a trade exposed oligopolistic industry. Using two decades of panel data on the US Portland cement industry, we estimate a fully dynamic model of firms' strategic entry, exit, production, and investment decisions. We then use the model to simulate counterfactual outcomes under three general classes of allocation regimes: auctioning, grandfathering, and contingent updating. We find that the dynamic evolution of market structure can vary significantly across the policy scenarios we consider. We quantify the overall costs of achieving desired emissions reductions and the distribution of those costs.

*UC Berkeley and NBER, MIT, and MIT and NBER. The authors thank Erica Meyers for excellent research assistance.

> The California cement industry's potential cost disadvantage under AB32 is staggering...In the absence of measures that either relieve the initial cost pressure or impose equivalent costs of imports, such a substantial price increase will render the California cement industry economically unviable, will result in a massive shift in market share towards imports in the short run, and will precipitate sustained disinvestment in the California cement industry in the long run.

Letter from the Coalition for Sustainable Cement Manufacturing and Environment to Larry Goulder, Chair of the Economic and Allocation Advisory Committee. Dec. 19, 2009.

1 Introduction

As the scope and stringency of environmental regulations increase, there are mounting concerns about potentially adverse impacts on industrial competitiveness, trade flows, and emissions leakage. "Leakage" refers to any increase in emissions in one jurisdiction that occurs as a direct consequence of emissions regulation imposed in another jurisdiction. The potential for emissions leakage continues to be a major obstacle to regional climate change policies.

There are several related channels through which emissions leakage can occur. It can manifest immediately as firms adjust variable input and output decisions such that less (more) stringently regulated production assets are used more (less) intensively. It can occur gradually as firms accelerate the retirement of older production technologies in more stringently regulated jurisdictions and invest in new facilities and equipment in less stringently regulated jurisdictions. "Indirect" leakage can also occur. As demand for carbon intensive fuels decreases in more stringently regulated regions, carbon intensive fuel prices may fall, and producers in unregulated jurisdictions may substitute towards these inputs.

Ideally, an analysis of the potential for emissions leakage under an existing or proposed policy will account for all three of these channels. However, both the academic literature and policy analysis carried out by government agencies tends to emphasize the more direct and immediate effects of emissions regulations. Dynamic industry responses to market-based environmental policy interventions, which are expected to play a dominant role in the context of proposed climate change regulations, have received relatively less attention. In a recent survey, Millimet et al. (2009) suggest that "the most striking gap in the literature on environmental regulation is its inability to connect to the rich literature on dynamic industry models in order to understand the long-term impacts of environmental regulation on industries and firms."

In this paper, we use the Markov-perfect Nash equilibrium (Maskin and Tirole, 1988; Ericson and Pakes, 1995) dynamic oligopoly framework developed in Ryan (2010) as the foundation for an analysis of market-based regulations limiting industrial emissions.[1] In particular, we investigate the static and dynamic efficiency properties of alternative approaches to allocating permits in emissions trading programs. How permits are initially allocated can have significant implications for who ultimately bears the cost of achieving the mandated emissions reductions and how cost-effectively the reductions are achieved.

Our analysis will focus on the cement industry. For a number of reasons, this industry has been at the center of the debate about climate change policy and domestic competitiveness. First, cement is one of the largest manufacturing sources of domestic carbon dioxide emissions (Kapur et al, 2009). Second, the anticipated effects of proposed federal and state-level climate change policies on cement production costs are striking. Given projected carbon prices and current emissions intensities, average variable production costs would increase by more than 50 percent.[2] Finally, 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, 2003 ENV; USGS Mineral Commodity Summary 2010).

Static, short run analytics are ill-suited for a study of how proposed emissions regulations might impact the domestic cement industry. Capital stock turnover is expected to play an essential role in improving the environmental performance of this industry (Worrell 2001, Sterner 1990). This is partly due to the limited opportunities to reduce carbon intensity through process changes and disembodied capital change, and partly due to fact that some very old and inefficient kilns are still in operation. Replacing these with newer and more efficient technologies would yield emissions reduction benefits.

Our analysis begins with the specification of a theoretical model of dynamic oligopoly in

[1] A majority of emissions regulated under existing and planned cap-and-trade program come from industries that can likely be classified as imperfectly competitive. Emissions from restructured electricity markets represent the majority of emissions currently targeted by existing cap-and-trade programs in the United States and Europe. Numerous studies provide empirical evidence of the exercise of market power in these industries, such as Borenstein et al. (2002); Joskow and Kahn (2002); Wolfram (1999); Puller (2007); Sweeting (2007); Bushnell et al. (2008). Other emissions intensive industries being targeted by regional emissions trading programs, such as cement and refining, are also highly concentrated.

[2] On average, domestic cement producers emit approximately one ton of carbon for each ton of cement produced. Marginal costs of cement production are estimated to be in the range of \$30-\$40/ton. Analyses of proposed federal greenhouse gas emissions trading programs project a permit price of \$15-\$20 per ton of CO_2 by 2015. Absent cost mitigation measures, variable operating costs would increase by approximately 50 percent.

which 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 strategies of their rivals. Conditional on capital investments, producers compete each period in homogeneous quantities.[3] Regional market structures evolve as firms enter, exit, and adjust production capacities in response to changing market conditions.

To convert the theoretical framework into an estimable empirical model, we introduce additional assumptions about the properties of the structural errors. We then estimate the model using two decades of detailed data collected at both the regional market and kiln level. In the first step, we recover reduced form parameters of the policy functions that describe what actions a firm will take conditional on the prevailing market conditions. In the second step, we identify the remaining structural parameter values (such as the distribution of investment costs) that are most consistent with observed outcomes, conditional on the assumptions of the underlying model.

With parameter estimates in hand, we then turn to our investigation of the static and dynamic implications of the permit allocation design decision.The basic intuition underlying our counterfactual simulations is quite simple. In the benchmark scenario in which emissions are unconstrained, firms invest at the level where marginal costs equal expected marginal benefits subject to covering their fixed costs. The expected benefits are a function of the period payoffs, as firms with larger capacities are able to compete over a larger segment of this trade exposed market. With the introduction of an emissions cap-and-trade program, firms face a different payoff structure. Firms' operating costs and marginal revenues vary systematically across the permit allocation design alternatives we consider. This variation has potentially significant implications for industry dynamics in equilibrium.

In this working paper, we present results from a proof of concept exercise that investigates the welfare implications of the permit allocation design choice in two regional cement markets: Salt Lake City and San Francisco. Counterfactual policy simulations suggest that the long-run dynamic implications are significant. More specifically, we find that rates of

[3]A majority of emissions regulated under existing and planned cap-and-trade program come from industries that can likely be classified as imperfectly competitive. Emissions from restructured electricity markets represent the majority of emissions currently targeted by existing cap-and-trade programs in the United States and Europe. Numerous studies provide empirical evidence of the exercise of market power in these industries, such as Borenstein et al. (2002); Joskow and Kahn (2002); Wolfram (1999); Puller (2007); Sweeting (2007); Bushnell et al. (2008). Other emissions intensive industries being targeted by regional emissions trading programs, such as cement and refining, are also highly concentrated.

entry, exit, and investment in new capacity differs significantly across the policy scenarios we consider. Differences in industry structure beget differences in the cost effectiveness of emissions abatement, producer profits, consumer surplus, and emissions leakage.

To highlight the importance of properly accounting for industry dynamics, we compare the predictions of the dynamic model against a model that limits firms responses to price and quantity changes; technology operating characteristics and industry structure are held fixed. When we allow the structure of the industry to respond to the market-based emissions regulation, the welfare impacts of the policy intervention are amplified in most cases. Intuitively, the industry structure evolves differently under different policy design scenarios, and this has significant implications for welfare. Future work will extend the analysis to consider the entire domestic cement industry, and the distribution of costs and benefits associated with different policy designs across regional markets with different structural characteristics.

This paper makes substantive contributions to three areas of the literature. First, this paper is germane to the literature that considers the dynamic efficiency properties of market-based emissions regulations. By their very nature, long-run policy impacts are very difficult to identify empirically. During the time it takes for these impacts to manifest, a host of other potentially confounding factors and processes change and evolve. The conventional approach to analyzing these long run relationships has been to use either highly stylized theoretical analysis of the effects of environmental policies on market structure (Conrad and Wang, 2003; Requate, 1997) or large, deterministic, optimization-based simulation models (Jensen and Rasmussen, 2000; Fischer and Fox, 2007). Our paper is unique in its detailed representation of industry dynamics and its use of state-of-the-art econometric modeling tools to estimate key structural parameters cannot be observed directly.

Second, we are not aware of any other study of the impacts of market-based emissions regulations (and permit allocation design decisions in particular) in the domestic cement industry. This industry has an important role to play in efforts to reduce industrial CO_2 emissions. Domestic cement producers are at the center of debates about potentially adverse impacts of emissions regulations on domestic industrial competitiveness. We discuss how our results can inform this important policy discussion.

Finally, the paper makes an important methodological contribution in its application of parametric value function methods to a dynamic game. We make use of interpolation techniques to compute the equilibrium of the counterfactual simulations. This allows us to treat the capacity of the firms as a continuous state. Even though parametric methods have been used in single agent problems, its application to dynamic industry models with discrete entry, exit and investment decisions have not been very successful to date (Doraszelski and Pakes, 2007).

The paper is organized as follows. Section 2 gives an overview of the related literature. Section 3 provides some essential background on the US Portland cement industry. We introduce the model and provide a detailed description of the different pollution permit mechanisms in Section 4. We present the estimation and computational methodology in Section 5. Counterfactual simulation results with and without permit markets are reported in Section 6. We conclude with a discussion of the results and directions for future research in Section 7.

2 Background and Related Literature

Regulatory agencies have been allocating tradable emissions permits under the auspices of local, regional, and nationwide emissions cap-and-trade programs for over a decade. Historically, policy makers have chosen between two types of permit allocation approaches: auctioning and grandfathering. Under an auction regime, emissions permits are sold to the highest bidder. In contrast, "grandfathered" permits are freely distributed in lump-sum to regulated sources based on pre-determined, firm-specific characteristics.

Economists have generally argued in favor of auctioning permits when auction revenues can be used to offset factor taxes or other pre-existing distortions.[4] However, policy makers have routinely chosen to forego auction revenues in favor of handing permits out for free to regulated entities.[5] The ability to make concessions to adversely impacted and politically powerful stakeholders via grandfathering has played an essential role in securing widespread support for the adoption of emissions trading programs. A pure grandfathering approach is unlikely to be a politically feasible option in the context of a Federal GHG trading program. This is primarily due to the unprecedented value of the permits to be allocated; grandfathering all permits to industry would amount to significant overcompensation for compliance costs incurred.[6]

[4] A summary of the literature that considers the permit allocation design choice in the presence distorted factor markets is provided by Goulder and Parry (2008).

[5] A majority of permits are distributed freely to regulated entities in Southern California's RECLAIM program, the European Union's Emissions Trading Program (EU ETS), the National Acid Rain Program (ARP), and the regional NOx Budget Trading Program.

[6] The Congressional Budget Office estimates that emissions permits allocated annually under the federal cap-and-trade system proposed by the Senate in 2009 could be worth up to $300 billion a year by 2020 (CBO, 2009).

More recently, a third approach has emerged. Under "contingent allocation" regimes, rules established ex ante determine how a firm's permit allocations will be periodically updated over the course of the trading program based on lagged production decisions. By making future permit allocations a function of current production, allocation updating confers an implicit production subsidy. Contingent allocation is similar to grandfathering insofar as it can be used to transfer value to firms via free permit allocations, but different in that the implicit production subsidy reduces the marginal costs of production whereas grandfathering reduces fixed costs.

Previous work addressing permit allocation design can generally be divided into two categories. The first consists of theoretical analyses that derive general insights about the effects of different allocation mechanisms. Among the most relevant to this paper are Fischer (2003) and Neuhoff et al. (2005). Both discuss how allocation updating can be used to reduce inefficiencies resulting from the exercise of market power in imperfectly competitive product markets. They also examine the potential entry and exit effects of these mechanisms. Our model encompasses these two channels in an oligopolistic framework.

The second category includes industry-specific detailed policy simulations using large-scale linear programming models or other deterministic, optimization-based approaches. Several of these papers consider how European cement producers might be impacted under alternative permit allocation regimes (see, for example, Demailly and Quirion (2006); Ponssard and Walker (2008); Jensen, 2001; Szaboe et al, 2006; US EPA, 1996).[7] One limitation of these numerical simulation models is that they must rely on the extant econometric literature to provide "off-the-shelf" estimates of important structural parameters (such as the fixed costs of entry or the elasticity of import supply). It is often the case that the econometric literature is not up to the task; models are often parameterized using outdated values or educated guesses. We believe our paper is unique in that we use econometrically-derived values of these primitives to assess the industry response to alternative permit allocation designs.

[7]All of these papers assume assume Cournot competition. All but one analyze the European cement sector.

3 Regulating CO_2 emissions in the domestic cement industry

3.1 The Portland cement industry

Portland cement is an inorganic, non-metallic substance with important hydraulic binding properties. It is the primary ingredient in concrete, an essential construction material used widely in building and highway construction. Demand for cement comes primarily from the ready-mix concrete industry, which accounts of over 70 percent of cement sales. Other major consumers include concrete product manufacturers and government contractors. Because of its critical role in construction, demand for cement tends to reflect population, urbanization, and economic trends.

Cement competes in the construction sector with substitutes such as asphalt, clay brick, rammed earth, fiberglass, steel, stone, and aluminum. Wood is a particularly important substitute in housing construction (Van Oss ENV 2003). Another important class of substitutes are the so called supplementary cementitious materials (SCMs) such as ferrous slag, fly ash, silica fume and pozzolana (a reactive volcanic ash). Concrete manufacturers can use these materials as partial substitutes for clinker. The substitition 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, facts; Kapur et al, 2009).

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 CO2 emissions (Van Oss, 2003, ENV).

The cement production process involves two main steps: the manufacture of clinker (i.e. pyroprocessing) and the grinding of clinker to produce cement. In the pyroprocessing stage, a fuel mix comprised of limestone and supplementary materials is fed into a large kiln lined with refractory brick. Under intense heat (kiln temperatures reach temperatures of 1450°C), limestone undergoes a chemical transformation to the precursor of cement, known as clinker.

Once cooled, clinker is mixed with gypsum and ground into a fine powder to produce cement.[8]

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 transporation 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]

3.2 Carbon dioxide emissions from cement production

Carbon dioxide emissions intensities, typically measured in terms of metric tons of emissions per metric ton of clinker, vary considerably across cement producers. Because plants with different emissions intensities will respond differently to the policy interventions we analyze, it is important to capture this variation as accurately as possible.

In the cement production process, carbon dioxide is emitted, in approximately equal proportions, from the combustion of fuel in the kiln and the chemical reaction that transforms limestone into cement. Trace amounts of carbon dioxide can also escape in cement kiln dust.

Carbon intensity of cement production varies significantly across producers. Much of this 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 PIS, 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 fuel efficient as the older wet-process kilns.

Although data limitations prevent us from estimating emissions intensities specific to each kiln in the data set, we can estimate technology-specific emissions rates. Both the IPCC and the World Business Council for Sustainable Development's Cement Sustainability Initiative (WBC, 2005) have developed protocols for estimating emissions from clinker production.

[8]The US cement industry is comprised of clinker plants (kiln only operations), grinding-only facilities, and integrated (kiln and grinding) facilities.Almost all of the raw materials and energy used in the manufacture of cement are consumed during pyroprocessing. We exempt grinding only facilities from our analysis.

[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.

We use these protocols to generate technology-specific estimates of carbon dioxide emissions rates. Appendix 1 explains these emissions rate calculations in more detail.

3.3 Trade Exposure

Whereas overland transport of cement is very costly, sea-based transport of clinker is relatively inexpensive. The United States absorbs approximately one quarter of the total global cement trade (Van Oss, 2003 ENV). Imports are used to satisfy cyclical increases in domestic demand when domestic capacity constraints bind and to arbitrage production cost disparities across countries.

In the recent past, import penetration rates have averaged around 20 percent (USGS Mineral Commodity Summary 2010). China is currently the largest supplier of imported cement (accounting for 22 percent of imports), followed by Canada, Korea, and Thailand (USGS, 2010 fact sheet). Notably, the degree of import penetration varies significantly across regional markets. Whereas imports accounted for approximately 40 percent of California's total cement consumption in 2006, several inland markets are supplied exclusively by domestic production.

4 Model

The basic building block of the model is a regional cement market.[10] Each market is fully described by the $\overline{N} \times 2$ state vector, s_t, where s_{it} describes the productive capacity of the i-th firm at time t and its associated emissions rate. We set $\overline{N}$ to be the maximal number of firms. Firms with zero capacity are considered to be potential entrants. Time is discrete and unbounded. Firms discount the future at rate $\beta = 0.9$.

Each decision period is one year. In each period, the sequence of events unfolds as follows: first, potential entrants receive a private draw from the distribution of both investment and entry costs, while incumbents receive private draws on the fixed cost of investment/divestment and their scrap value for exiting. Second, all firms simultaneously make entry, exit, and investment decisions. Third, incumbent firms compete over quantities in the product market. Finally, firms enter and exit, and investments mature. We assume that firms who decide to exit produce in this period before leaving the market, and that adjustments in capacity take one period to realize. We also assume that each firm operates

[10]This section borrows heavily from Ryan (2010).

independently across markets.[11]

Firms in trade-exposed regional markets face an import supply curve:

$$\ln M_m = \rho_0 + \rho_1 \ln P_m, \tag{1}$$

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.[12] In future work, we hope to explore the potential implications of the strategic use of imports by dominant market players.

Firms obtain revenues from the product market and incur costs from production, entry, exit, and investment. Firms compete in quantities in a homogeneous goods product market. After netting out imports, firms face a constant elasticity residual demand curve:

$$\ln Q_m(\alpha) = \alpha_{0m} + \alpha_1 \ln P_m, \tag{2}$$

where Q_m is the aggregate market quantity, P_m is price, α_{0m} is a market-specific intercept, and α_1 is the elasticity of demand. For clarity, we omit the m subscript in what follows.

There are essentially five variable inputs used in cement production: labor, fuel (primarily coal), electricity, feedstocks, and maintenance. These factor inputs are not substitutable (Das, 1994). The majority of variable operating costs are energy related. Because frequent heating and cooling damages the firebrick lining, kilns typically operate continuously at full capacity for 24 hours a day. Annual output is adjusted by varying the length of time the kiln is shut down for annual maintenance. In the model, each firm chooses the level of annual output that maximizes their static profits given the outputs of the competitors, subject to capacity constraints that are determined by dynamic capacity investment decisions:

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

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

[12] In fact, firms that own a majority of the domestic production capacity in the United States are also among the largest importers. These dominant producers presumably use imports to supplement their domestic production as needed, and to compete in markets where they do not own production facilities. Domestic cement producers have noted that increased domestic ownership of import facilities has contributed to a "more orderly flow of imports into the U.S."

Grancher, Roy A. "U.S. Cement: Record Performance and Reinvestment", Cement Americas, Jul 1, 1999 where $P(Q;\alpha)$ is the inverse of Equation 2. In the presence of fixed operation costs the product market may have multiple equilibria, as some firms may prefer to not operate given the outputs of their competitors. However, if all firms produce positive quantities then the equilibrium vector of production is unique, as the best-response curves are downward-sloping.

We will use this model to evaluate the impacts of alternative approaches to allocating emissions permits in an emissions trading program. Firm-specific compliance costs will be determined by kiln-specific emissions rates, e_i, production quantity, and the number of permits the firm receives free of charge. While postponing the discussion of the policy designs we consider until Section 4.3, we note here that the introduction of a cap-and-trade program modifies the profit function in Equation 3 through the term $\varphi(q_i, e_i, \tau)$, where τ is the price of a permit for one metric ton of carbon dioxide. The precise nature of the modification will vary across permit allocation designs.

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

$$C_i(q_i;\delta) = \delta_1 q_i + \delta_2 1(q_i > \nu s_i)(q_i - \nu s_i)^2. \tag{4}$$

Variable production costs consist of two parts: a constant marginal cost, δ_1, and an increasing function that binds as quantity approaches the capacity constraint. I assume that costs increase as the square of the percentage of capacity utilization, and parameterize both the penalty, δ_2, and the threshold at which the costs bind, ν. This second term, which gives the cost function a "hockey stick" shape common in the electricity generation industry, accounts for the increasing costs associated with operating near maximum capacity, as firms have to cut into maintenance time in order to expand production beyond utilization level ν. We denote the profits accruing from the product market by $\bar{\pi}_i(s;\alpha,\delta)$.

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

$$\Gamma(x_i;\gamma) = 1(x_i > 0)(\gamma_{i1} + \gamma_2 x_i + \gamma_3 x_i^2) + 1(x_i < 0)(\gamma_{i4} + \gamma_5 x_i + \gamma_6 x_i^2). \tag{5}$$

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

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

$$\Phi_i(a_i; \kappa_i, \phi_i) = \begin{cases} -\kappa_i & \text{if the firm is a new entrant,} \\ \phi_i & \text{if the firm exits the market.} \end{cases} \tag{6}$$

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. Denote the activation status of the firm in the next period as χ_i, where $\chi_i = 1$ if the firm will be active next period, whether as a new entrant or a continuing incumbent, and $\chi_i = 0$ otherwise. All of the shocks that firms receive each period are mutually independent.

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

$$\pi_i(s, a; \alpha, \rho, \delta, \gamma_i, \kappa_i, \phi_i) = \bar{\pi}_i(s; \alpha, \rho, \delta) - \Gamma(x_i; \gamma_i) + \Phi_i(a_i; \kappa_i, \phi_i). \tag{7}$$

For the sake of brevity, we henceforth denote the vector of parameters in Equation 7 by θ.

4.1 Transitions Between States

To close the model it is necessary to specify how transitions occur between states as firms engage in investment, entry, and exit. 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.[13]

We also assume that the emissions rate of the firm is fixed. We assume that there are three discrete levels of emissions rates, corresponding to the three major types of production technology in the cement industry. Existing incumbents are modeled as having one of the three technologies, while new entrants are always endowed with the frontier technology. As a result, the emissions profile of an industry changes over time in response to firm turnover.

4.2 Equilibrium

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

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

$$\sigma_i : (s, \epsilon_i) \to a_i, \tag{8}$$

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

$$\begin{aligned} V_i(s; \sigma(s), \theta, \epsilon_i) &= \bar{\pi}_i(s; \theta) \\ &+ \max \Bigg\{ \max_{x_i^* > 0} \left[-\gamma_{i1} - \gamma_2 x_i^* - \gamma_3 x_i^{*2} + \beta \int E_{\epsilon_i} V_i(s'; \sigma(s'), \theta, \epsilon_i) \; dP(s_i + x^*, s'_{-i}; s, \sigma(s)) \right], \\ &\max_{x_i^* < 0} \left[-\gamma_{i4} - \gamma_5 x_i^* - \gamma_6 x_i^{*2} + \beta \int E_{\epsilon_i} V_i(s'; \sigma(s'), \theta, \epsilon_i) \; dP(s_i + x^*, s'_{-i}; s, \sigma(s)) \right], \\ &\beta \int E_{\epsilon_i} V_i(s'; \sigma(s'), \theta, \epsilon_i) \; dP(s'; s, \sigma(s)), \phi_i \Bigg\}, \end{aligned} \tag{9}$$

where θ is the vector of payoff-relevant parameters, E_{ϵ_i} is the expectation with respect to the distributions of shocks in future periods, and $P(s'; \sigma(s), s)$ is the conditional probability

[13] 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.

distribution over future state s', given the current state, s, and the vector of strategies, $\sigma(s)$.

Potential entrants must weigh the benefits of entering at an optimally-chosen level of capacity against their draws of investment and entry costs. Firms only enter when the sum of these draws is sufficiently low. I assume that potential entrants are short-lived; if they do not enter in this period they disappear and take a payoff of zero forever, never entering in the future.[14] Potential entrants are also restricted to make positive investments; firms cannot "enter" the market at zero capacity and wait for a sufficiently low draw of investment costs before building a plant. The value function for potential entrants is:

$$V_i^e(s;\sigma(s),\theta,\epsilon_i) = \max\{0, \\ \max_{x_i^*>0}\left[-\gamma_{1i} - \gamma_2 x_i^* - \gamma_3 x_i^{*2} + \beta\int E_{\epsilon_i} V_i(s';\sigma(s'),\theta,\epsilon_i)dP(s_i + x^*, s'_{-i}; s, \sigma(s))\right] - \kappa_i\right\}. \quad (10)$$

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

$$V_i(s;\sigma_i^*(s),\sigma_{-i}(s),\theta,\epsilon_i) \geq V_i(s;\tilde{\sigma}_i(s),\sigma_{-i}(s),\theta,\epsilon_i), \quad (11)$$

for all s, ϵ_i, and all possible alternative strategies, $\tilde{\sigma}_i(s)$. As I work with the expected value functions below, I note that the MPNE requirement also holds after integrating out firms' private information: $E_{\epsilon_i} V_i(s;\sigma_i^*(s),\sigma_{-i}(s),\theta,\epsilon_i) \geq E_{\epsilon_i} V_i(s;\tilde{\sigma}_i(s),\sigma_{-i}(s),\theta,\epsilon_i)$. Doraszelski and Satterthwaite (2010) discuss the existence of pure strategy equilibria in settings similar to the one considered here. The introduction of private information over the discrete actions guarantees that at least one pure strategy equilibrium exists, as the best-response curves are continuous. However, there are no guarantees that the equilibrium is unique, a concern I discuss next in the context of my empirical approach.

4.3 Cap-and-trade Mechanisms

In the counterfactual policy simulations, domestic cement producers are required to participate in the greenhouse gas emissions trading program. To comply with the program, all domestic firms must hold sufficient permits to offset their emissions. We impose no spatial or sectoral restrictions on permit trading; permits can be traded freely among all program

[14] This assumption is for computational convenience, as otherwise one would have to solve an optimal waiting problem for the potential entrants. See Ryan and Tucker (2010) for an example of such an optimal waiting problem.

participants. To keep the analysis more tractable, we do not allow banking or borrowing of permits across time.

The equilibrium permit price, τ, is an exogenous parameter in our simulations. This is equivalent to assuming that the aggregate marginal abatement cost curve is flat in the neighborhood of the constraint imposed by the emissions cap. This will be an appropriate assumption if the domestic cement industry is a relatively small player in the emissions market, such that changes in industry net supply/demand for permits cannot affect the equilibrium market price.[15]

We consider the following three general classes of permit allocation mechanisms:

4.3.1 Grandfathering

Under a grandfathering regime, permits are allocated for free to incumbent firms pre-existed the carbon trading program. Firm-specific permit allocation schedules (i.e. the number of permits the firm will receive each period) are determined at the beginning of the program and are based on pre-determined operating characteristics. Firms receive the same amount of permits every period until they exit the market, in which case they lose their allocation. New entrants do not receive any permits for free in this regime.

Several studies have demonstrated that a pure grandfathering regime would grossly over-compensate industry for the compliance costs incurred under proposed Federal climate change legislation. For example, a recent paper finds that grandfathering fewer than 15 percent of the emissions allowances generally suffices to prevent profit losses among industries that would suffer the largest percentage losses of profit absent compensation (Goulder, Hafstead, and Dworsky, 2010). Under the grandfathering regime we consider, we assume that a number of permits equal to 20 percent of annual baseline emissions are grandfathered each year to incumbent cement producers. The per period profit function becomes:

$$\pi_{it} = P\left(q_{it} + \sum_{j \neq i} q_{jt}; \alpha\right) q_{it} - C_i(q_{it}; \delta) - \tau(e_i q_{it} - A_i), \qquad (12)$$

$$\text{with} \qquad \sum_i A_i = \overline{A}.$$

[15]This assumption is likely to be approximately true in the context of a federal GHG trading program that permits offsets. 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.

where τ is the permit price, e_i is the firm's emissions rate and E represents aggregate industry emissions. The number of permits the firm receives for free from the regulator is A_i; $\overline{A}$ represents the total amount of emissions allocated for free to domestic cement producers. We assume that the share of emissions allowances allocated to firm i (i.e. $A_i/\overline{A}$) is equal to its share of the installed kiln capacity at the outset of the program.

When permits are grandfathered in a cap and trade program, policy makers must decide ex ante how to deal with new entrants and firms who exit. In our simulations, we assume that a firm forfeits its future permit allocation entitlements when it exits the market. We assume that new entrants are not entitled to free permits.[16] In some existing program designs (including the EU ETS), some fraction of the permits to be allocated are set aside for new production capacity entering the market. Future work will explore these alternative policy designs that offer free permit allocations as incentive for new entrants.

4.3.2 Auctioning

In the context of an economy-wide greenhouse gas emissions trading program, an auctioning regime has its proponents. In 2007, the Congressional Budget Office Director warned that a failure to auction permits in a federal greenhouse gas emissions trading system "would represent the largest corporate welfare program that has even been enacted in the history of the United States".[17] However, important industry stakeholders have united in their opposition to a policy regime that would auction all permits (at least in the near term).[18] We include a pure auctioning design in our analysis nonetheles; it serves as a useful benchmark for other more politically palatable options.

Under the auctioning regime we consider, all emissions permits are sold in a uniform price auction. In each period, firms must purchase a number of permits equal to their emissions.

[16] 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. Most states require forfeiture of free permit allocations upon closure.

[17] "Approaches to Reducing Carbon Dioxide Emissions: Hearing before the Committee on the Budget U.S. House of Representatives", November 1, 2007. (testimony of Peter R. Orszag)

[18] The US Climate Action Partnership (USCAP) is a non-partisan coalition comprised of 25 major corporations and 5 leading environmental groups. In January 2009, the group issued its "Blueprint for Legislative Action" in which it urged Congress to use some portion of allowances to buffer the impacts of increased costs to energy consumers, and to provide transitional assistance to trade-exposed and emissions intensive industry.

The per-period production profit function becomes:

$$\pi_{it} = P\left(q_{it} + \sum_{j \neq i} q_{jt}; \alpha\right) q_{it} - C_i(q_{it}; \delta) - \tau e_i q_{it}. \tag{13}$$

Note that the first order conditions associated with static profit maximization under auctioning are identical to those under grandfathering. This highlights the aforementioned independence property, which holds 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.

4.3.3 Contingent allocation updating

Under proposed state and federal climate change legislation, output-based updating provisions are used to address concerns about near-term competitiveness impacts, job loss, and emissions leakage. Emissions permits are allocated for free to eligible firms using a continuously updated, output-based formula.[19]

Following Bushnell and Chen (2009), we adopt a closed-loop approach to modeling of an allocation updating regime. Permits are allocated based on product market shares in the current period. The per period profit function becomes:

$$\pi_{it} = P\left(q_{it} + \sum_{j \neq i} q_{jt}; \alpha\right) q_{it} - C_i(q_{it}; \delta) - \tau(e_i q_{it} - \theta_i(q_{it})\overline{A}), \tag{14}$$

where $\phi_i(q_{it})$ denotes the share of the free emissions allocated to firm i.

Implicit in Equation 14 are two simplifying assumptions. First, the quantity of permits a firm receives in the current period depends on its production level in that same period. Second, the size of the implicit subsidy per unit of output is taken to be exogenous to firms' production decisions. Together, these assumptions simplify the dynamic problem considerably, while still allowing us to capture the dynamic implications of the grandfathering

[19] Climate change legislation recently passed in the House and reported by committee in the Senate includes a provision that allocates permits to eligible industries using an output-based formula. These free allocations are intended to compensate both direct compliance costs (i.e. the cost of purchasing permits to offset emissions) and indirect compliances costs (i.e. compliance costs reflected in higher electricity prices). Under California's AB32, implementing agencies have recommended that free allocation to industry will, "to the extent feasible, be based on output-based GHG efficiency "benchmarks" and "update" to reflect changes in production each year for industry with leakage risk" (Greenhouse Gas Cap-and-Trade Regulation Status Update May 17, 2010 California Air Resources Board).

mechanism to a significant extent. Note that the independence property does not apply here. Firms production decisions will be directly affected by how the allocation updating rule is specified.

We consider two different means of updating emissions allocations. In the first "output-based" updating regime, emissions allowances are allocated according to market share:

$$\theta_i(q_i) = \frac{q_i}{\sum_i q_i}.$$

In the second, "emissions-based" regime, more polluting firms receive a larger permit allocation. This might occur if firms owning older, less efficient kilns insist that they should be entitled to a larger allowance allocation so as to compensate them for their higher compliance costs. In this case:

$$\theta_i(q_i) = \frac{e_i q_i}{\sum_i e_i q_i}.$$

In both cases, firms receive permits as long as they are producing in the market. Therefore, new entrants also receive an allocation proportional to either their output or their emissions.

5 Estimation and computation

The econometric estimation is based on the benchmark model, in which the price of emissions is set to zero ($\tau = 0$), i.e. there is no compliance cost due to emissions regulation. Once estimated, this model can be used to simulate the dynamic industry response to market-based emissions regulations that affect firms' production and investment choices primarily through operating costs provided certain assumptions are met. In particular, we will assume that firms' response to a given operating cost change is independent of whether the cost change is caused by emissions regulation or other exogenous factors (such as changes in energy prices or other inputs).

5.1 Estimation

Although our data sources and identification strategy are similar to Ryan (2010), there are some important differences in how the model is specified and estimated. In this section, significant deviations are discussed. The interested reader is referred to Ryan (2010) for additional details regarding the data and estimation.

Table 1: Descriptive Statistics for Regional Markets

Market	Number of Firms	2006 Capacity	Emissions Rate	Import Market Share
Atlanta	6	1285	0.97	0.12
Baltimore/Philadelphia	6	1497	0.99	0.12
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
Los Angeles	6	1733	0.93	0.18
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 Antonio	6	1318	0.95	0.3
San Francisco	4	931	0.93	0.18
Seattle	2	607	1.05	0.65
St Louis	4	1358	1.05	0

5.1.1 *Regional market definition*

The USGS collects establishment-level data from all domestic Portland cement producers and publishes these data in an annual Minerals Yearbook. Cement price and sales data are aggregated to the regional market level to protect the confidentiality of the respondents. In recent years, increased consolidation of asset ownership has required higher levels of data aggregation. Conversations with the experts at USGS indicate that the current regional market definitions group plants that are unlikely to compete with each other (Van Oss, personal communication).

Rather than adopt the USGS protocols, we base our regional market definitions on the industry accepted limitations of economic transport as well as company-specific SEC 10k filings which include information regarding markets served by specific plants. To merge the USGS cement prices with our data set, USGS prices are weighted by kiln capacity in each region. For example, if kiln capacity in the region we define as region A is equally divided between USGS defined markets B and C, we define the price in region A to be the average price reported in USGS markets B and C. We report some descriptive statistics for our regional markets in Table 1.

5.1.2 *Import supply and residual demand elasticities*

In this paper, the import supply response to changes in domestic operating costs is modeled separately from domestic market demand response. This separation is important given our interest in understanding how policy-induced operating cost increases could affect import penetration rates.

Regional cement prices, P_{mt}, domestic production levels, Q_{mt}, and imports, M_{mt}, are endogenous variables in our estimation framework. We jointly estimate the following system of equations using two stage least squares (2SLS):

$$\ln Q_{mt} = \alpha_0 + \alpha_1 \ln P_{mt} + \alpha_{2m} + \alpha_3' \ln X_{mt} + \varepsilon_{1mt}, \tag{15}$$

$$\ln M_{mt} = \phi_0 + \phi_1 \ln P_{mt} + \phi_{2m} + \phi_3' \ln Z_{mt} + \varepsilon_{2mt}. \tag{16}$$

Equation 15 specifies the residual demand function, as defined by total market demand less imports. The matrix X_{mt} contains the exogenous variables that affect marginal operating costs: state-level electricity prices, coal prices, and labor costs. The matrix Z_{mt} contains the exogenous import supply shifters: international freight rates and exchange rates. Each market has a demand shifter in the intercept, α_{2m}, using Atlanta as the baseline market.

Equation 16 represents the import supply schedule. This model is estimated using data from those markets accessible by water. For inland markets who are supplied entirely by domestic production, all ϕ coefficients are set to zero. The excluded instruments for cement price are exogenous factors affecting domestic demand, such as measures of building and highway construction activity, population, and prices of substitutes such as asphalt.

Here we assume that imports are supplied by a competitive fringe. In fact, a majority of imports are controlled by the same firms who own a majority of domestic production assets. If imports are actually strategic, this will affect how imports respond to to relative changes in domestic operating costs.[20]

5.2 Estimation results

Table 2 enumerates the parameter estimates used in our simulations. Overall, these estimates appear reasonable.

- The marginal cost estimate of \$30/ton of clinker falls well within the range that is

[20] In future work, we hope to modify the model so as to make imports controlled by firms with a large domestic presence endogenous.

typically reported for domestic production: \$27-\$44 per ton (Van Oss, 2003 ENV).

- The import supply elasticity point estimate is 2.7. When analyzing the impacts of environmental regulations, the US EPA assumes an import supply elasticity of 2 for the cement sector based on Broda et al (2008).

- The elasticity of aggregate demand is 2.96. This is higher in absolute value than some other demand elasticities reported in the literature. For example, Jans and Rosenbaum (1996) estimate a domestic demand elasticity of -0.81. On the other hand, using much higher-quality data, Foster, Haltiwanger, and Syverson (2008) estimate several similar high demand elasticities for homogeneous goods industries, such as -5.93 for ready-mixed concrete, cement's downstream industry. Experimentation with lower elasticities resulted in unreasonably high investment costs, on the order of several billion dollars for a median-sized plant, in order to rationalize low entry rates given the amount of foregone opportunity profits.

- Investment costs are roughly in line with the accounting costs cited in Salvo (2010), which reports a cost of \$200 per ton of installed capacity. Our numbers are slightly higher, which in line with the idea that these costs represent economic opportunity costs as opposed to accounting costs. 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 \$350M, which implies a cost of \$250 per ton of installed capacity.[21]

- The magnitudes of the fixed costs are reasonable at face value, and in conjunction with the estimated variances, are in accord with the observed rates of investment, entry, and exit in the cement industry.

Some of the parameters from the model described above are not reported. Divestment is virtually never observed in the data and thus the estimates of divestment costs are estimated poorly. Fixed costs of production and operation are also not reported, as these are set to zero. 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.

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

Table 2: Simulation Parameters

Parameter	Value
Demand Parameters	
Constant	20.38
Elasticity of Demand	-2.96
Discount Factor	
β	0.9
Production Parameters	
Capacity Cost	1.157E10
Capacity Cost Binding Level	1.896
Marginal Cost	30
Investment Parameters	
Fixed Cost Mean	1,798
Fixed Cost Standard Deviation	420
Marginal Cost	233
Exit Cost	
Scrap Distribution Mean	-67,490
Scrap Distribution Standard Deviation	55,167
Entry Distribution	
Entry Cost Mean	172,680
Entry Cost Standard Variance	41,559

5.3 Computation

Once the parameters have been estimated, the model can be computed to compare the market performance under different cap-and-trade policies. In order to compute the equilibrium of the game, we make use of parametric approximation methods. In particular, we interpolate the value function using cubic splines. The reasons behind using parametric methods are twofold. First, the game has a continuous state space, given by the vector of capacities of the firms. By using parametric methods, we can allow firms to deterministically choose their capacity in a continuous space. Second, parametric approximation methods can be useful to improve computational speed. Previous work has already suggested the potential benefits of using parametric approximation methods (Pakes and McGuire, 1994).

Parametric value function methods have been explored in a single agent dynamic programming context.[22] However, they have not been widely used in dynamic games, particularly in games in which players take discrete actions, such as entry and exit (Doraszelski and Pakes, 2007). In our application, we find the method to perform well compared to a discrete value function method. In particular, this parametric method allows us to treat capacity as a continuous state, which improves the convergence properties of the game.[23]

The procedure we use is similar in spirit to the discrete value function iteration approach. In both methods, the value function is evaluated at a finite number of points. At each iteration and for a given guess of the value function, firms' strategies are computed optimally (*policy step*). Then, the value function is updated accordingly (*value function step*). This process is repeated until the value function and the policy functions do not change significantly.

The difference between the discrete value function iteration and our iterative approach is that we approximate the value function with a flexible parametric form. In particular, given a guess for the value function V^k at pre-specified grid points, we interpolate the value function with a multi-dimensional uniform cubic spline, which can be computed very efficiently (Habermann and Kindermann, 2007).[24] This interpolation defines an approximation of the value function in a continuous space of dimension equal to the number of active firms. For a given number of firms active N_A in the market, the value function at any capacity vector

[22] For a general treatment of approximation methods used in the context of dynamic programming, see Judd (1998). An assessment of these methods in a single agent model can be found in Benitez-Silva et al. (2000).

[23] This is mainly driven by the fact that firms take deterministic actions with respect to the continuous state.

[24] For a detailed treatment of splines methods, see de Boor (2001).

s is approximated as,

$$\hat{V}_i^k(s) = \sum_{j=1}^{(J+2)^A} \phi_{N_A,j} B_{N_A,j}(s), \tag{17}$$

where J is the number of grid points, $\phi_{N_A,ij}$ are the coefficients computed by interpolating the values V^k when there are A active firms, and $B_{N_A,j}(s)$ is the spline weight given to coefficient $\phi_{N_A,j}$ when the capacity state equals s. This coefficient is the product of capacity weights for each of the incumbent firms, so that $B_{N_A,j}(s) = \prod_{i \in A} B_j(s_i)$.

In the *policy step*, optimal strategies are computed over this continuous function. For a given firm, we compute the conditional single-dimensional value function, given the capacity values of the other firms, $\hat{V}_i^k(s_i|s_{-i})$. This formulation allows us to represent the single-dimensional investment problem of the firm. The following expression defines the expected value function of the firm conditional on staying in the market and investing to a new capacity s_i'. Firms maximize,

$$\max_{s_i'} \quad \pi_i(s_i, s_i'|s_{-i}) + \sum_{s'_{-i} \in S_{-i}} Pr^k(s'_{-i}; \sigma^k(s)) \hat{V}_i^k(s_i'|s'_{-i}). \tag{18}$$

We compute the optimal strategy by making use of the differentiability properties of the cubic splines, which allows us to compute the first-order conditions with respect to investment. Given that the cubic spline does not restrict the value function to be concave, we check all local optima in order to determine the optimal strategy of the firm.[25] Conditional on optimal investment strategies, we then compute the new policy function with respect to the entry, investment and exit probabilities, which gives us an updated optimal policy σ^{k+1}. This allows us to compute a new guess for the value function V^{k+1} in the *value function step*.

The process is iterated until the strategies for each of the firms and the value function in each of the possible states do not change more than an established convergence criterion, such that $\parallel \sigma^{k+1} - \sigma^k \parallel < \epsilon_\sigma$ and $\parallel V^{k+1} - V^k \parallel < \epsilon_V$.

6 Regional market simulations

In this working paper, we report results from simulating outcomes in two regional markets: Salt Lake City and San Francisco. These markets were chosen to illustrate the differences between coastal and interior markets. The Salt Lake City market is not accessible by water;

[25] Given that the cubic spline is defined by a cubic polynomial at each of the grid intervals, this implies that at most there will be $2(J-1)+2$ candidate local optima, where J is the number of grid points.

demand is met entirely by domestic supply. There is also some heterogeneity in the emissions rates of incumbent firms in this market. In contrast, San Francisco is trade exposed and incumbent producers are homogenous with respect to emissions rates.[26]

We investigate both the static and dynamic effects of alternative cap-and-trade policy designs in these two regional markets. The static analysis helps to build intuition about the relationships between policy design features and short-run equilibrium outcomes. It also serves to underscore the importance of accounting for dynamic industry response.

6.1 Short run impacts on equilibrium outcomes

The common practice in ex ante policy analysis involves simulating regulatory effects in a static setting in which industry structure is held fixed. Models are paramaterized using data from a representative year. Short run impacts of a proposed policy intervention are simulated, and these impacts are then extrapolated over a longer time horizon (OAQPS, 1999). This is the approach we adopt in our static analysis. More precisely, we simulate a single period market outcome in the unregulated baseline case and under the range of counterfactual policy designs we consider. To facilitate comparisons with our dynamic simulations, results are expressed as net present values using a social discount rate of three percent. We assume the simulated annual outcomes would be observed each year of the 30 year time horizon we consider. In the more complex, dynamic simulations, entry, exit, and capital investments are modeled explicitly over the same 30 year time horizon.

Salt Lake City: Static simulation

Figure 1 summarizes results from short-run simulations of outcomes in the Salt Lake regional market. In each figure, the horizontal axis measures the CO_2 permit price. The top left panel plots the cement price (measured in \$/ton of clinker) as a function of the permit price. The middle figure plots the relationship between cement production and the permit price. In these static simulations, the marginal effects of permit market compliance obligations are identical across the auctioning and grandfathering regimes. Thus, simulated cement prices and quantities are identical across these two market designs. In contrast, outcomes differ across the output-based and emissions-based updating regimes in the short-run. Because firms in this market differ with respect to emissions intensity, capacity, and cost structure, the size of the marginal production subsidy conferred under allocation updating

[26] In future work, the scope of the analysis will be expanded to include all domestic cement markets. This will allow us to provide a more comprehensive assessment of the industry-wide impacts of the policies we consider, and to assess the extent to which a "one-size-fits-all" policy regime can have differential impacts across heterogeneous regional markets.

will vary across firms under emissions-based updating, but will be the same for all firms under output-based updating.

Cement prices and quantities are unaffected by the introduction of a cap-and-trade program so long as the shadow value of the emissions constraint (i.e. the permit price) is relatively small. In the benchmark (i.e. unregulated) case, firms in this market are capacity constrained and earning scarcity rents. When firms are required to internalize a relatively small emissions cost, scarcity rents are reduced, but output decisions are unaffected in the short run.

The implicit subsidy conferred by allocation updating partially mitigates the effect of the market-based emissions regulation on cement prices and quantities. Cement prices rise more quickly and steeply under auctioning or grandfathering as compared to contingent allocation updating regimes.

The top right graph illustrates the relationship between industry profits and emissions permit prices under the policy designs we consider. Firm profits are most profoundly impacted under an auctioning regime wherein firms must pay for all permits. The total quantity of permits freely allocated is held constant across the grandfathering, output-based updating, and emissions-based updating. However, the value of these permits comprises a larger fraction of the value of permits demanded under grandfathering because output levels and emissions are lower. Under both updating regimes, profits are decreasing with permit price over the range of prices considered. Under grandfathering, when permit prices are relatively high, an increase in the stringency of the cap, and thus the permit price, increases the value of the firm's permit endowment by more than it decreases the firms' profits from the sale of cement.

The number of firms serving this market is unaffected by policy design changes in this static analysis, by assumption (lower left figure). The relationship between industry-level emissions and the permit price is very similar to the relationship between output and the permit price because emissions rates are constant per unit of production (lower center figure). Finally, there is no leakage in this landlocked market (lower right figure).

San Francisco: Static simulation

Figure 2 depicts these relationships between emissions permit prices and short-run equilibrium outcomes in the San Francisco regional market. This market is larger than Salt Lake City and exposed to competition from foreign imports. There are several noteworthy differences between figures 1 and 2. First, the presence of imports prevents the cement price in this market from rising above \$78. This price cap binds at a permit price of 53\$/ton

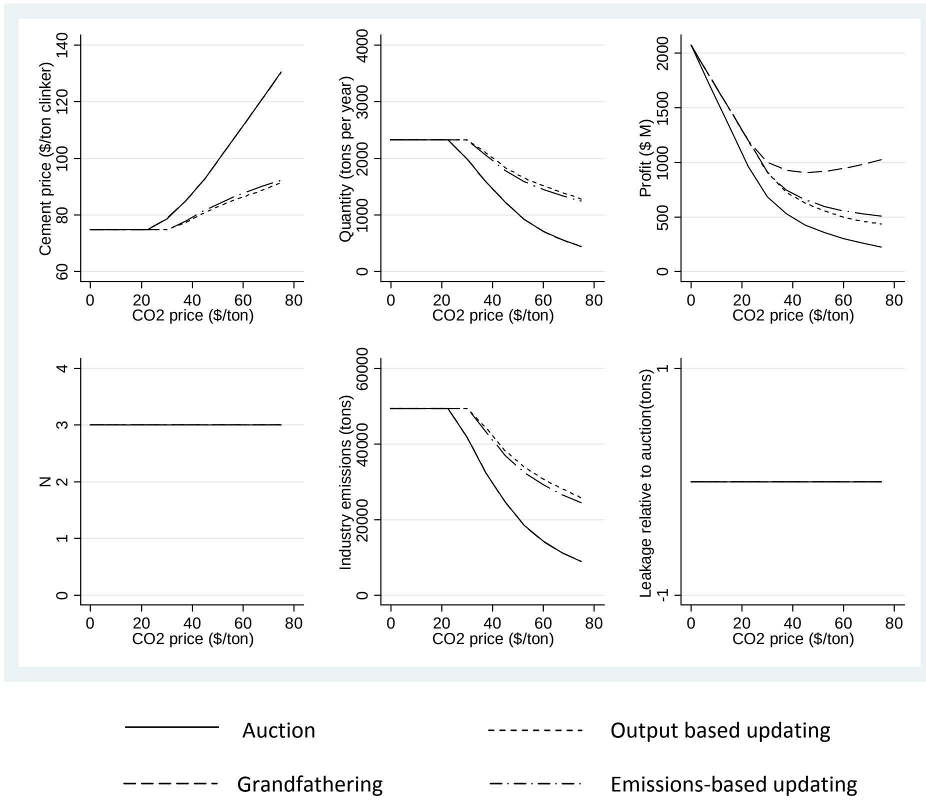

Figure 1 Static simulation results : Salt Lake City regional market

Notes: These graphs plot the results from the static simulations. Outcomes are simulated for a single period conditional on the number of incumbents in the market and the operating characteristics of the production technologies currently owned and operated by these incumbents. Each graph plots the relationship between an equilibrium outcome of interest and the exogenously set emissions permit price. To make these results more comparable to the dynamic simulations, we compute the net present value over a 30 year time horizon, assuming the per-period results we generate would not vary over years.

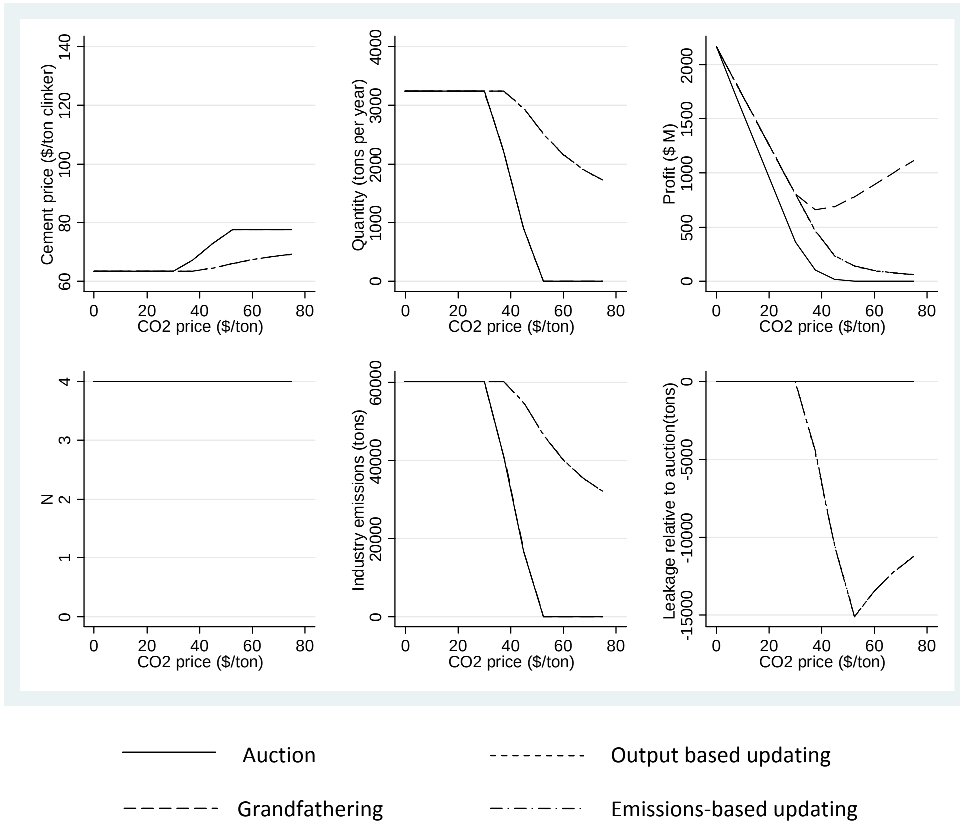

Figure 2 Static simulation results: San Francisco regional market

Notes: These graphs plot the results from the static simulations. Outcomes are simulated for a single period conditional on the number of incumbents in the market and the operating characteristics of the production technologies currently owned and operated by these incumbents. Each graph plots the relationship between an equilibrium outcome of interest and the exogenously set emissions permit price. To make these results more comparable to the dynamic simulations, we compute the net present value over a 30 year time horizon, assuming the per-period results we generate would not vary over years.

CO_2. Beyond this point, import penetration reaches 100 percent and domestic emissions are driven to zero.

In contrast to Salt Lake City, all incumbents in San Francisco use very similar production technologies and thus have the same imputed CO_2 emissions rates. Consequently, the implicit production subsidy conferred by allocation updating is symmetric across all firms under both emissions-based and output-based allocation updating in the short run when production technology is held fixed. In San Francisco, as in Salt Lake, relatively low permit prices have no short-run impact on equilibrium prices and quantities; only profits are affected. As the stringency of the emissions trading program increases, domestic cement prices rise and quantities fall. In this market, as prices rise, so to does import penetration. Contingent allocation updating significantly mitigates the effect of the emissions regulation on cement prices, domestic cement production, and emissions leakage.

The lower right figure illustrates how the dynamic updating of emissions permits significantly reduces emissions leakage relative to the auctioning or grandfathering cases. Since leakage is one-to-one in the amount of domestic production supplanted by imports, the V-shaped pattern is due to the fact that domestic firms produce much more under updating rules. After import penetration becomes 100 percent in the auctioning/grandfathering case, the relative gap between the two levels of productions begins to close as production declines in the updating case with higher permit prices. For a sufficiently high permit price, this gap would close to zero as firms would not find it profitable to produce anything under updating rules.

6.2 Long run impacts on equilibrium outcomes

Salt Lake City: Dynamic simulation

Figure 3 summarizes the results from the simulations of dynamic industry response over a 30 year time horizon in Salt Lake City. Results are expressed as net present values (using a social discount rate of three percent). There are several noteworthy differences between the static and dynamic simulation exercises.

When we allow the structure of the industry to respond to the introduction of market-based emissions regulation, the independence property ceases to hold; equilibrium outcomes differ across auctioning and grandfathering regimes. This is due to the asymmetric treatment of new entrants and existing incumbents in the grandfathering regime. In the dynamic model; the grandfathered permits create a permanent advantage to incumbent firms as long as they

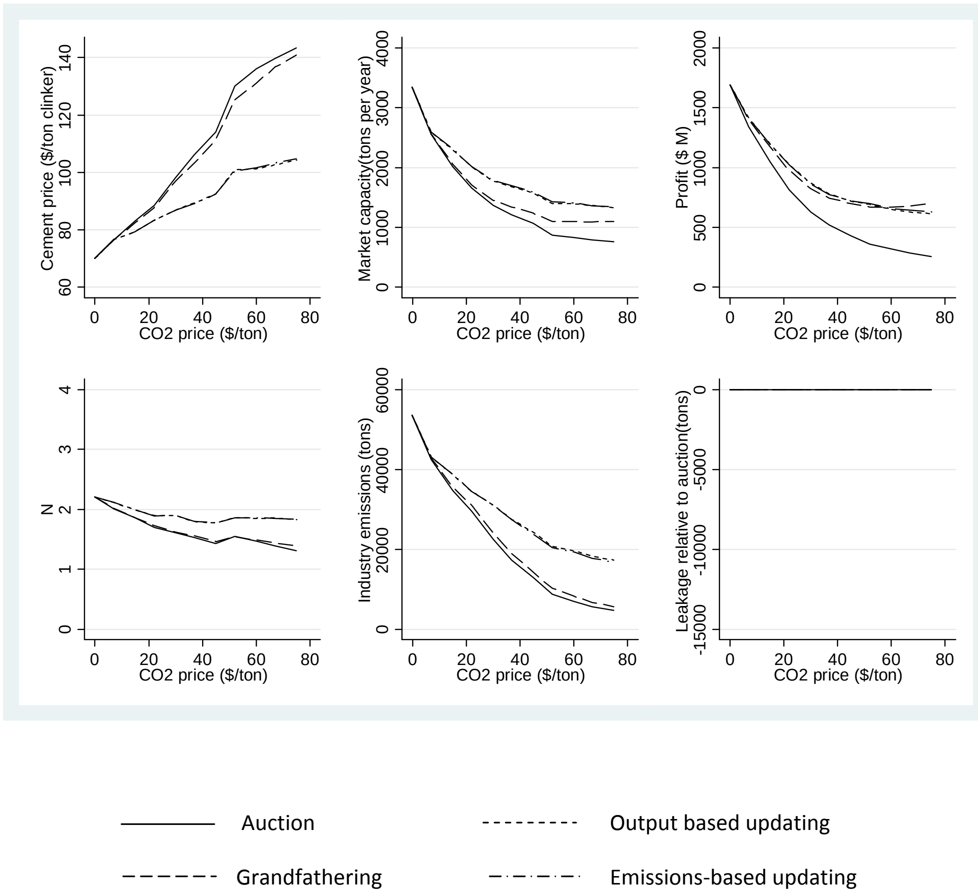

Figure 3 Dynamic simulation results: Salt Lake City regional market

Notes: These graphs plot the results of simulated dynamic industry response to counterfactual policy scenarios over a 30 year period. With the exception of the number of firms (N), these results are expressed as net present values.

stay in the market. This distorts the evolution of the market structure as grandfathered firms have a higher exit threshold as compared to incumbents in an auctioning regime.

In contrast to the static case, low permit prices have immediate impacts on equilibrium cement prices and quantities when firms are able to adjust their entry, exit, and investment decisions. Moreover, cement prices are higher for any given permit price when dynamics are properly accounted for. This is due to the large reduction in capacity induced by the introduction of the emissions regulation which is missing from the static case. As a corollary, emissions are lower for any given permit price when we account for dynamics. As expected, the reduction in equilibrium levels of capacity is most accelerated under auctioning; firms are more likely to exit when they do not have a steady stream of freely allocated permits to look forward to. In general, auctioning has the highest equilibrium prices, lowest levels of production, and lowest emissions.

Profits are correspondingly lowest under auctioning; the contingent updating and grandfathering regimes are moderately more profitable. As permit prices increase, the option value of selling grandfathered permits dominates the product market production, making grandfathering increasingly profitable in both absolute and relative terms. This is reflected in the bottom middle panel, where emissions under auctioning and grandfathering decline at roughly equal rates; this implies that firms in the grandfathering regime prefer to sell their permits in equilibrium rather than use them to offset production. As in the static case, dynamic allocation updating mitigates the effect of the emissions regulations on cement prices and domestic emissions.

San Francisco: Dynamic simulation

The simulated, long-run industry responses to emissions policy interventions in San Francisco are summarized by Figure 4. The difference in the cement price response to a given permit price across the static and dynamic cases is not as pronounced as in Salt Lake City; competition from imports dampens these equilibrium impacts. Here again there are differences in equilibrium outcomes for even the smallest permit prices due to adjustments in market structure at the extensive margin.

Accounting for dynamics reverses the profitability of grandfathering relative to contingent updating. In the static case, grandfathering's profitability increased at high permit levels. After accounting for dynamics, grandfathering is dominated by the contingent updating schemes, although it is still increasing. This is partially due to the selection effects; conditional on being active at high permit prices under contingent updating, it is likely

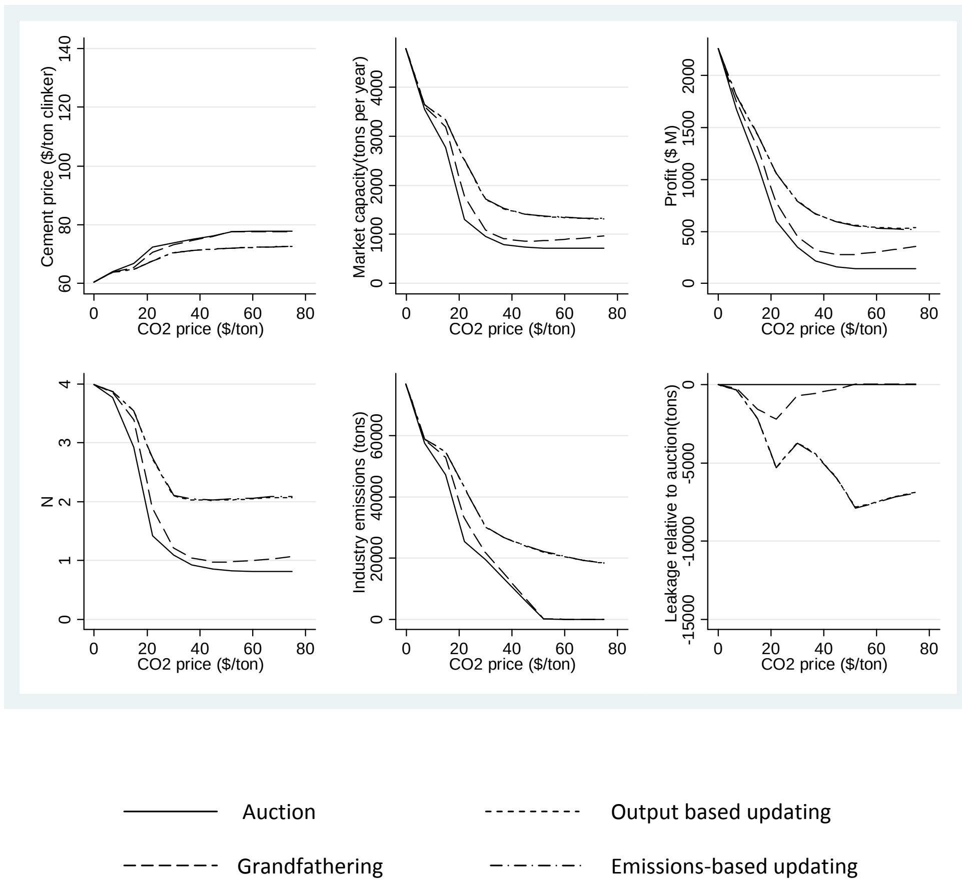

Figure 4 Dynamic simulation results: San Francisco regional market

Notes: These graphs plot the results of simulated dynamic industry response to counterfactual policy scenarios over a 30 year period. With the exception of the number of firms (N), these results are expressed as net present values.

that the firm faces fewer rivals than in the grandfathering case. To this point, at permit prices above \$25/ton CO_2, the average number of firms across simulations drops from four to approximately two under updating regimes, and approximately one under grandfathering and auctioning. A more concentrated industry implies higher cement prices in the dynamic simulations under permit allocation updating regimes, controlling for permit price. We do not observe the same price increase in the grandfathering or auctioning simulations at high permit prices because prices are constrained by presence of imports.

In San Francisco, the mitigating effects of allocation updating on emissions leakage are smaller in the dynamic simulations vis a vis the static simulations. We see this in the bottom right graph of Figure 4. The W-shaped relationship of emissions-based updating relative to the baseline is due to two factors: the first price at which the gap narrows is at which static production begins to be influenced by the permit price. The second turning point is the point where domestic production under grandfathering/auctioning hits zero. After this point, the relative gap begins to close, eventually reaching zero for very high permit prices, as in the static case. In sum, the level of emissions is much lower when accounting for dynamics than in the static case, which also implies that leakage is also relatively much larger.

6.3 Welfare metrics

The preceding results suggest that allocation updating has the potential to significantly mitigate emissions leakage and adverse competitiveness impacts in trade-exposed markets. In regional markets that are sheltered from import competition, allocation updating can mitigate the distortions associated with the exercise of market power. In this section, we evaluate the overall welfare impacts of the alternative permit market designs we consider. A simple example helps to lay the foundation for this welfare analysis.

Figure 5 illustrates the static welfare consequences of introducing an emissions trading program into an inland (and thus autarkic) regional cement market that is monopolized by a single domestic producer. We maintain our assumption that this industry is small relative to the universe of sources subject to the cap, equivalent to assuming that the abatement supply curve facing the monopolist is locally flat. The curve labeled MPC measures the marginal private costs of production (i.e. fuel costs, labor costs, etc.). The curve labeled MSC captures both private marginal costs and the monetized value of the damages from the firm's emissions: $MSC = MPC + \tau e$. This assumes a constant emissions rate per unit of output (e). The social welfare maximizing level of output is Q^* at the corresponding price

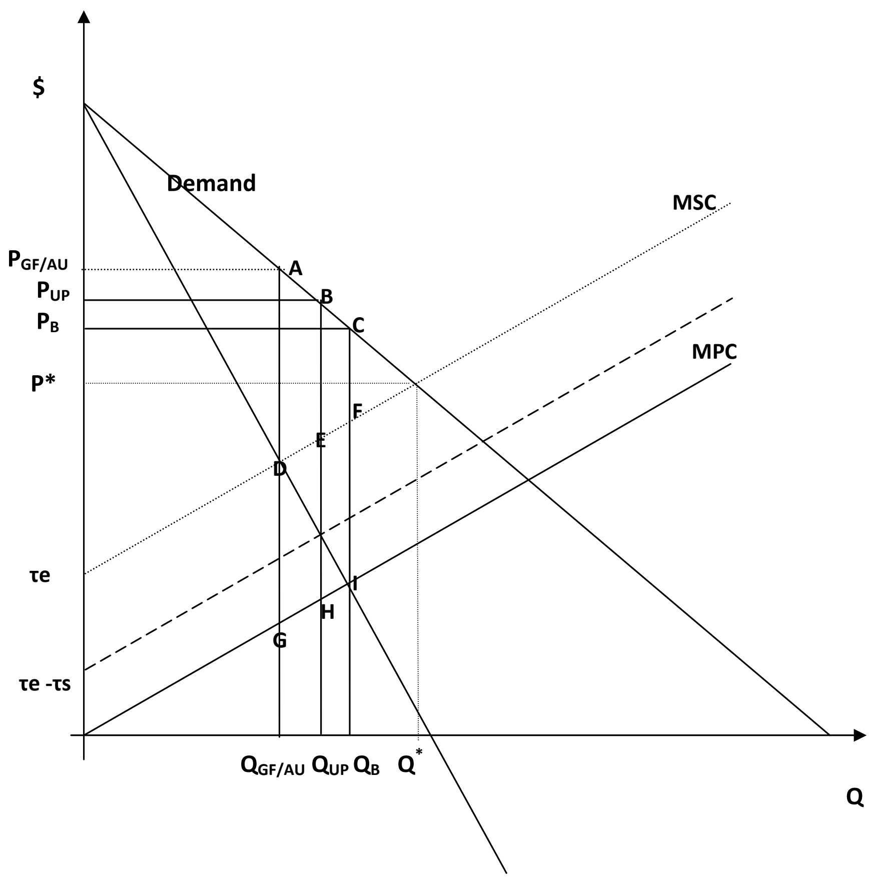

Figure 5 : Welfare impacts of allocation updating in an autarkic monopoly

P^*.

In the base case, emissions are unregulated and the profit maximizing monopolist will set marginal revenue equal to marginal cost. The equilibrium price and quantity are P_B and Q_B, respectively. In this static example, output levels under grandfathering (GF) or auctioning (AU) regime are equivalent. The emissions regulation increases variable operating costs by τe and the monopolist scales back production to $Q_{GF/AUC}$.

Under the allocation updating regime, the monopolist is allocated s permits per unit of output. The monopolist's profit maximizing choice of output under updating is Q_{UP}. Consistent with the static simulation results summarized above, allocation updating mitigates the impacts of a given permit price on equilibrium output prices and quantities.

Expositionally, it is useful to decompose the net welfare impacts of the emissions regulations vis a vis the unregulated benchmark into three component parts:

1. $W1$: The first component captures producer and consumer surplus generated in this industry plus any government revenues earned through the auctioning of emissions permits. In figure 5, the welfare change associated with the introduction of the grandfathering or auctioning regime, in terms of $W1$, is the area $AGIC$. The introduction of this emissions regulation reduces $W1$ because the pre-existing distortion associated with the exercise of market power is exacerbated by the policy intervention. This ill effect is mitigated, to some extent, by the subsidy implicit in the allocation updating regime. The introduction of the policy design that incorporates allocation updating reduces $W1$ by only $BHIC$.

2. $W2$: The second welfare component accounts for the damages from industry emissions. In figure 5, the value of the emissions reductions achieved under a policy that incorporates auctioning or grandfathering is equal to area $DGIF$. Once these emissions reductions are accounted for, the remaining welfare loss relative to the unregulated benchmark is represented by the area $ADFC$. The net welfare impact of the policy intervention is still negative because the gain from reducing emissions $W2$ does not fully offset the reduction in $W1$. Note that if the damages from emissions are sufficiently large such that the socially optimal output is less than Q_B, the introduction of the emissions regulation will be welfare improving.

   Under the allocation updating regime, emissions of the monopolist increase vis a vis auctioning or grandfathering by area $DGEH$. Although system-wide emissions will not increase (because the system-wide emissions cap is binding), a greater share of the compliance burden must be borne by other sources in the cap-and-trade program. This additional cost must be accounted for when making welfare comparisons across policy designs. The net welfare effect of the policy that incorporates allocation updating is -$BEFC$.

3. $W3$: The third welfare component accounts for any emissions associated with imports. This third component has no role to play in the autarkic example in Figure 5. However, many regional cement markets in the United States are accessible by water and are thus exposed to competition from foreign imports. The simulation results from the San Francisco market illustrate how the implicit production subsidy conferred in an updating regime can mitigate the effect of the emissions regulation on import penetration and emissions leakage. All else equal, the welfare properties of allocation updating regimes vis a vis auctioning or grandfathering will be more favorable in trade exposed markets when $W3$ is properly accounted for.

The counterfactual policy simulations we consider are considerably more complex than the stylized case depicted in Figure 5. First, all regional cement markets are served by more than one firm. Second, in several of the regional markets we analyze, there is intra-market variation in emissions rates. If relatively clean producers are not capacity constrained, this makes it possible to achieve short run reductions in emissions intensity through a reallocation of production activity. Finally, and perhaps more importantly, our analysis considers both static and dynamic efficiency. When we allow firms to make entry, exit, and investment decisions over several periods, the welfare implications of these permit allocation alternatives can look quite different.

In the static analysis, per period welfare is defined as follows:

$$w(s,a;\alpha,\delta_i,\gamma,\tau,e) = \int_0^Q P(x;\alpha)dx - \sum_i \int_0^{q_i} C(x;\delta_i)dx - P(Q;\alpha)M(P;\gamma) - \tau \sum_i e_i q_i - \tau e^M M(P;\gamma), \tag{19}$$

where e^M denotes the emissions intensity of imports. This is an exogenous parameter constructed using an import volume weighted average of estimated foreign cement producers' emissions intensities (Worrell (2001)) The vector e includes the emissions intensity measures of both domestic producers and foreign imports.

In specifying this welfare function, we assume that marginal damages from carbon dioxide emissions are approximately constant and equal to the assumed equilibrium permit price τ.

Because damages from greenhouse gases are independent of where in the world the emissions occur, we penalize both domestic and foreign emissions at a rate of τ per unit. Equation [19] implicitly assumes that the cement sector is small relative to the larger emissions trading program, such changes in cement industry emissions do not affect the equilibrium permit price. Under a binding cap, $\tau \sum_i e_i q_i$ can be interpreted as the cost of offsetting emissions from the cement industry through emissions reductions at other sources under the cap.[27] Finally, we do not account for producer surplus captured by the producers of domestic imports in our welfare measures; domestic policy makers presumably ignore the economic impacts on producers and consumers outside their jurisdiction.

We compare outcomes across permit market designs, holding the aggregate level of permitted emissions (and thus the prevailing permit price τ) constant. We assume that the cement sector is small relative to the larger emissions trading program such that changes in cement industry emissions do not affect the equilibrium permit price. Although a change in $\sum_i e_i q_i$ will not affect τ, it will have implications for the abatement activities and costs in at other sources under the cap.

The welfare implications of the permit market designs we consider are characterized using Equation 19. More precisely, $W1$ captures effects on the private economic surplus accruing from domestic cement consumption (i.e. the first three terms in Equation 19. $W2$ accounts for impacts on costs associated with emissions from the industry $\tau \sum_i e_i q_i$. Finally, $W3$ accounts for changes in emissions leakage across policy designs. Changes in emissons leakage are measured as changes in $\tau e^M M(P; \gamma)$.

Welfare analysis in the dynamic simulations sums Equation 19 and subtracts any entry, exit and investment costs accruing over the time horizon we consider. This measure of dynamic efficiency is somewhat unconventional insofar as it rules out innovation and technological change. For our purposes, a dynamically efficient outcome maximizes social welfare subject to the constraints imposed by existing and proven production technologies.

6.4 Welfare comparisons across policy designs

We begin our analysis using the results from the static simulations that hold market structure and production technology characteristics fixed. Figure 6 illustrates the effects of these market-based emissions regulations relative to the benchmark case of unconstrained (i.e. unregulated) emissions. The top row of figures measures welfare changes in terms of W1;

[27] If net permit demand from the cement sector can affect the equilibrium permit price, our estimates of the costs of allocation updating, vis a vis auctioning or grandfathering, will be too low.

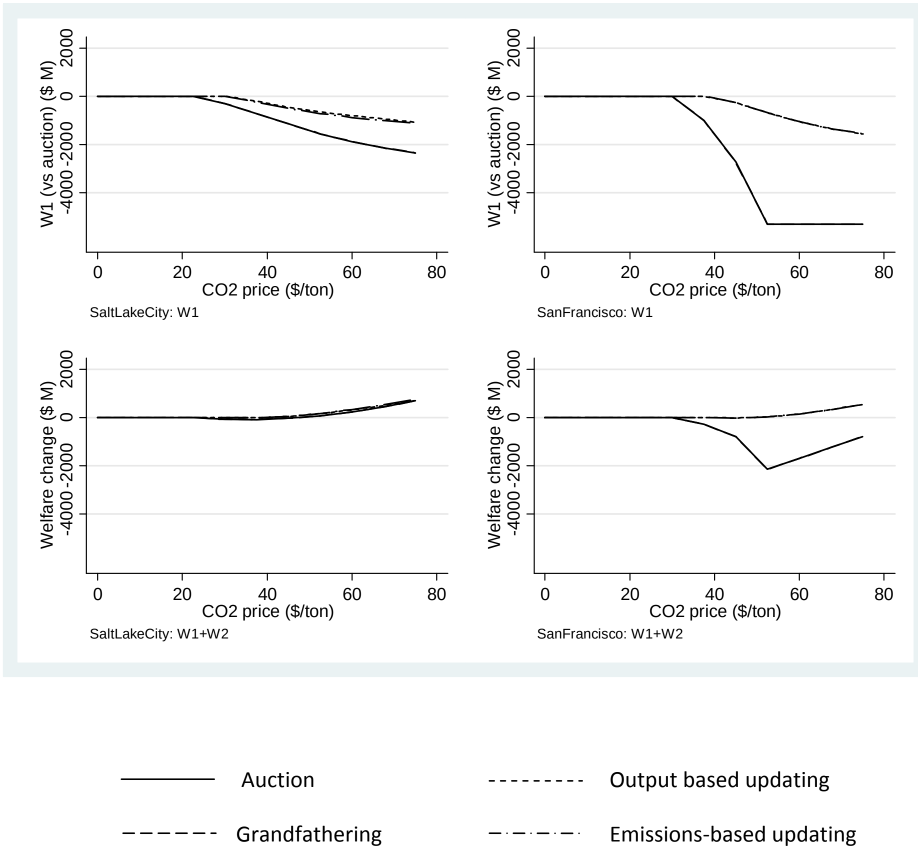

Figure 6 : Welfare comparison (W1 and W2) : Static simulation

Notes: Welfare component 1 (W1) captures the effect on consumer surplus, producer surplus, and government revenues (i.e. auction revenues). Welfare component 2 (W2) measures the damages associated with emissions from the regional cement market. Incremental changes in emissions are valued at the assumed permit price. These figures use results from the static simulations to estimate welfare impacts relative to an unregulated (i.e. emissions unconstrained) benchmark. In this static case, the welfare impacts of grandfathering and auctioning are identical. Welfare measures are expressed as net present values over a 30 year time horizon. The outcomes we simulate in a single year are assumed to reoccur each year.

this captures changes in producer surplus, consumer surplus, and government revenues exclusively. The bottom row of figures uses a more comprehensive welfare measure that adjusts for costs associated with industry emissions, $W2$. In comparisons across policy designs, changes in $W2$ capture changes in abatement expenditures accruing outside the cement sector. The right (left) panel corresponds to the Salt Lake City (San Francisco) market.

In the static case, welfare component $W1$ is unaffected by the emissions regulation at low permit prices because production prices and quantities are unaffected. At permit prices, reductions in $W1$ are greater under grandfathering and auctioning as compared to allocation updating. Intuitively, updating mitigates the exercise of market power in each market, thus limiting the negative welfare impacts associated with higher permit price. In San Francisco, these patterns are more pronounced.

In the Salt Lake City market, the benefits associated with avoided emissions ($W2$) more than offset the costs associated with production decreases when permit prices are high. Moreover, the relative welfare gains from allocation updating are eliminated because the relative benefits associated with mitigating the exercise of market power through updating are offset by the relative costs of shifting more of the compliance burden to other sectors subject to the cap. In the San Francisco market, the benefits captured by the $W2$ component outweigh the costs captured by $W1$ at higher permit prices under allocation updating. In contrast, under grandfathering and auctioning, the emissions regulation reduces welfare relative to the unregulated benchmark over the range of permit prices considered.

Figure 7 illustrates the static welfare comparison using the most comprehensive measure that accounts for impacts on economic surplus ($W1$), domestic emissions ($W2$), and emissions from imports ($W3$). Accounting for emissions leakage has no impact in Salt Lake City where there is no trade exposure. In San Francisco, allocation updating confers the added benefit of mitigating leakage. Accounting for emissions leakage has a relatively minor effect on the net welfare impacts of the emissions regulation that incorporates allocation updating but this significantly amplifies the welfare advantages of allocation updating vis a vis grandfathering and auctioning at higher permit prices.

With Figures 8 and 9, we conduct a similar exercise using the results from the dynamic simulations. There are several important differences between the static and the dynamic welfare analysis analysis. First, in the long run, welfare impacts vary across grandfathering and auctioning regimes which can have differential impacts on entry, exit, and investment.

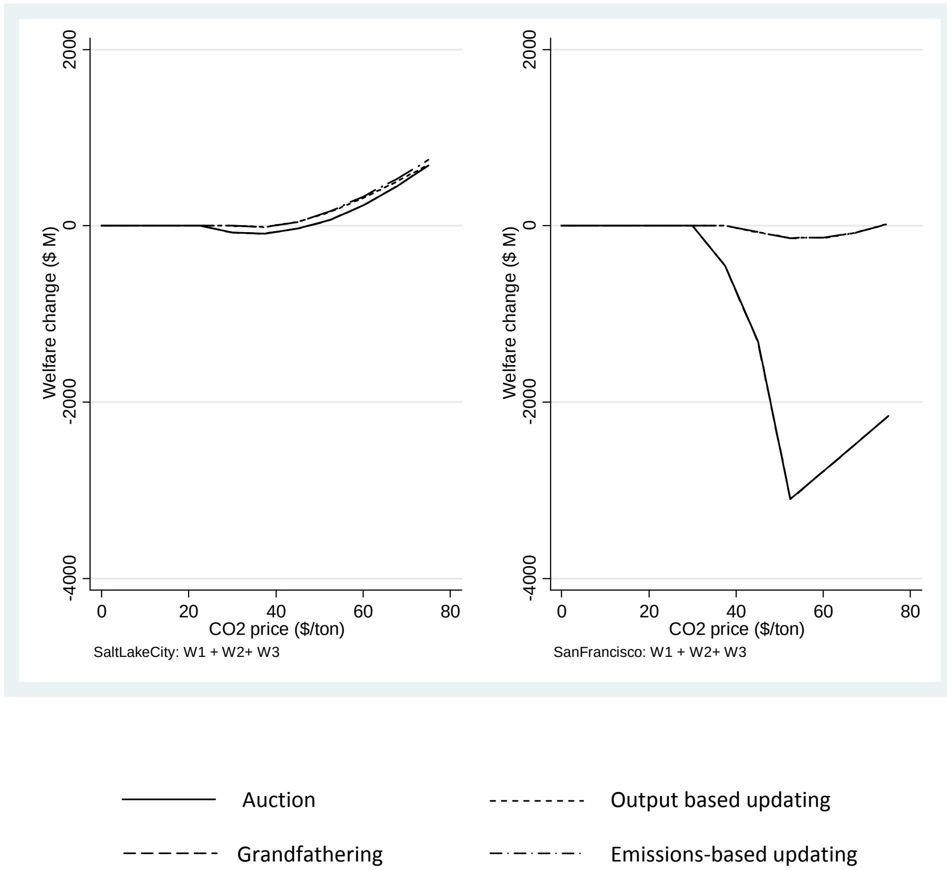

Figure 7 : Welfare comparison (W3) : Static simulation

Notes: The most comprehensive welfare measur captures consumer surplus, producer surplus, and government revenues (W1), damages associated with domestic emissions (W2) plus a penalty for any emissions associated with imports (W3). Emissions from imports are penalized at the assumed permit price. These figures illustrate how the market-based emissions policies we consider affect this measure of welfare relative to the unregulated benchmark. These results are generated using the static simulation results. Welfare measures are expressed as net present values over a 30 year time horizon.

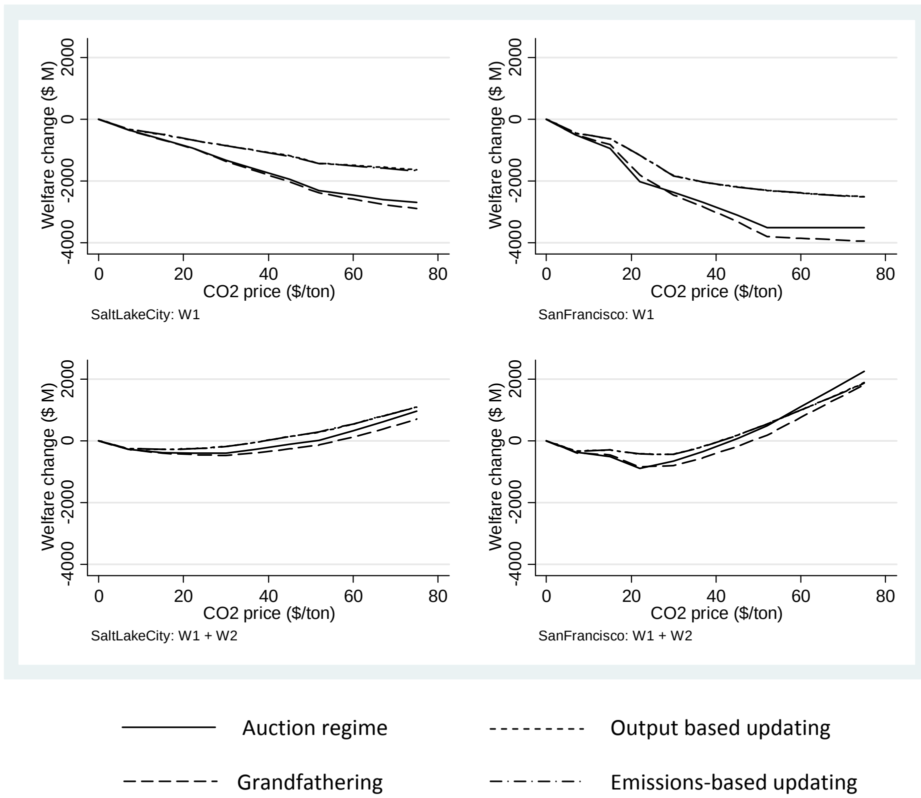

Figure 8 : Welfare comparisons of allocation updating relative to auctioning benchmark (static)

Notes: Welfare component 1 (W1) captures the effect on consumer surplus, producer surplus, and government revenues (i.e. auction revenues). Welfare component 2 (W2) measures the damages associated with emissions from the regional cement market. Incremental changes in emissions are valued at the assumed permit price. These figures use results from the static simulations. Welfare measures are expressed as net present values over a 30 year time horizon. The outcomes we simulate in a single year are assumed to reoccur each year

This is particularly true of San Francisco. When dynamics are properly accounted for, a reduction in capacity . The top panel of graphs n Figure 8 illustrate how the simulated effects on the economic surplus component, $W1$, are generally more pronounced as compared th the static case. In San Francisco, the extent to which this the emissions regulation can drive up cement prices is constrained by import pressures, This constraint binds in the auctioning and grandfathering case. The costs of the emissions regulation outweigh the benefits associated with avoided emissions over a range of permit prices.

Figure 9 offers the most comprehensive picture of the dynamic welfare implications of the policy designs we consider. Similar to the static case, allocation updating dominates grandfathering and auctioning designs over the range of permit prices in both markets. In contrast to the static case, the welfare costs of the policy intervention exceed the benefits except at high permit prices. For example, at a price of \$20 per ton of CO_2, the static analysis projects that the introduction of an emissions trading program will have no net impacts on welfare. In contrast, in the dynamic simulations, the estimated welfare impacts at a permit price of \$20 range from a NPV cost of approximately \$250 B under an allocation updating regime in Salt Lake City to over \$1 B San Francisco.

7 Conclusion and future work

We present a dynamic model to evaluate the counterfactual effects that different environmental policies would have to the US cement industry. We asses the implications of several mechanisms that have been discussed in policy circles. We examine the effects of a global emissions permit market and contrast the different dynamic implications that auctioning and grandfathering have in this context.

We find is that the dynamic effects of different allocation mechanisms are very substantial in this industry. The dynamic welfare consequences of various allocation mechanisms are complicated by the equilibrium reactions of firms along their entry, exit, and investment dimensions, which can result in a complex set of outcomes. Our results highlight the competing tensions facing policymakers in these markets: one on hand, there are concerns of serving the product market, and some allocation schemes result in more capitalized markets with more firms and lower prices, which helps obtain this goal. On the other hand, these allocation schemes may result in substantially more emissions than in more strict regimes that result in lower producer and consumer surplus. In some extreme cases, grandfathering

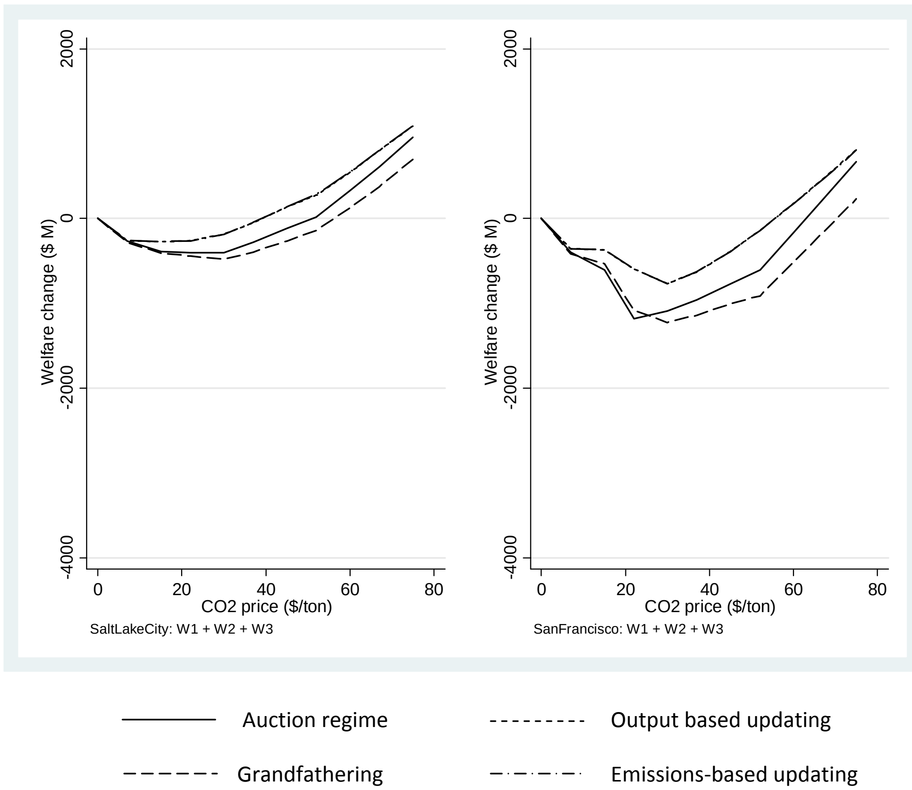

Figure 9 : Welfare comparisons of allocation updating relative to auctioning benchmark (dynamic)

Notes: The most comprehensive welfare measur captures consumer surplus, producer surplus, and government revenues (W1), damages associated with domestic emissions (W2) plus a penalty for any emissions associated with imports (W3). Emissions from imports are penalized at the assumed permit price. These figures illustrate how the market-based emissions policies we consider affect this measure of welfare relative to the unregulated benchmark. These figures illustrate how the market-based emissions policies we consider affect this measure of welfare relative to the unregulated benchmark. These graphs are generated using results from the dynamic simulations of industry response over a 30 year time horizon.

acts as a net transfer of permits to the firms, which can be viewed as a government subsidy to pollute. Our analysis shows the complex interplay between these schemes and equilibrium market structure, emissions, and welfare. We think that this paper is an important first step to understanding the interplay between these forces.

In the future, we would like to extend this work in several directions. One interesting avenue of research is to disentangle the effects of firm turnover from the substitution of production among emissions-differentiated firms. Our current results present the composite picture of these two effects, but understanding their separate roles would help clarify our understanding of the way that the allocation mechanisms influence the evolution of equilibrium market structure.

A more substantive undertaking will be the extension of our approach to a larger scale. The ultimate goal of our research is to build a model of the US Portland cement industry which is as close to reality as possible. Using heterogeneous market structures in the different regional markets in the United States, we plan to assess the potential differential impacts across heterogeneous markets of the implementation of a seemingly homogeneous policy such a carbon tax. This undertaking will require us to find ways to simulate very large markets, on the order of 8 or more strategic actors, which is currently a computational challenge that we are addressing. The approximation methods used in this paper may be useful in that regard, as they are capable of reducing the complexity of the state space to a smaller number of relevant variables. The approximation methods of Benkard et al. (2008) may also be fruitful in this regard.

In this case, we believe it is going to be very interesting to model the effects of asymmetric grandfathering, in which only incumbent firms receive permits for free. Not granting entrants free permits, contributes to the internalization of the cost of the emissions, and reduces the long-run average emissions in the sector. However, if entrants have cleaner technology, this can delay adoption of new technology, as entrants are relatively worse off than incumbents. We would like to examine the implications of this trade-off in our model.

Finally, we believe one of the major contributions of this project is going to be the inclusion of imports and exports. If there is the possibility of international trade, a reduction of cement production in the affected region can result in increased imports from elsewhere, as the permit price harms the competitiveness of the producers affected by the new environmental regulation. Modeling the fluxes of trade will allow us to assess the relevance of emissions leakage in this industry, which is a concern that has been brought up by industrial associations.

References

Hugo Benitez-Silva, John Rust, Gunter Hitsch, Giorgio Pauletto, and George Hall. A comparison of discrete and parametric methods for continuous-state dynamic programming problems. Computing in Economics and Finance 2000 24, Society for Computational Economics, July 2000.

C. Lanier Benkard, Gabriel Weintraub, and Ben Van Roy. Markov perfect industry dynamics with many firms. *Econometrica*, pages 1375–1411, November 2008.

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

Severin Borenstein, James B. Bushnell, and Frank A. Wolak. Measuring market inefficiencies in california's restructured wholesale electricity market. *American Economic Review*, 92 (5):1376–1405, December 2002.

James Bushnell and Yihsu. Chen. Regulation, allocation, and leakage in cap-and-trade markets for co2. Technical report, EI@Haas WP-183, Washington, DC, 2009.

J.B. Bushnell, E.T. Mansur, and C. Saravia. Vertical arrangements, Market Structure, and competition: an analysis of restructured US electricity markets. *American Economic Review*, 98(1):237–266, 2008.

Carl de Boor. *A Practical Guide to Splines. Revised Edition.* Springer-Verlag New York Inc., 2001.

Damien Demailly and Philippe Quirion. 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, 2006.

Ulrich Doraszelksi and Mark Satterthwaite. Computable markov-perfect industry dynamics. *RAND Journal of Economics*, 41(2):215–243, Summer 2010.

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

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

Carolyn Fischer. Output-based allocation of environmental policy revenues and imperfect competition. Discussion Papers dp-02-60, Resources For the Future, Jan 2003.

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

Christian Habermann and Fabian Kindermann. Multidimensional spline interpolation: Theory and applications. *Comput. Econ.*, 30(2):153–169, 2007. ISSN 0927-7099.

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

Paul L. Joskow and E. Kahn. A quantitative analysis of pricing behavior in california's wholesale electricity market during summer 2000. *The Energy Journal*, 23:1–35, 2002.

Kenneth Judd. *Numerical Methods in Economics.* MIT Press, 1998. ISBN 0-262-10071-1.

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

Daniel L. Millimet, Santanu Roy, and Aditi Sengupta. Environmental regulations and economic activity: Influence on market structure. *Annual Review of Resource Economics*, 1: 99–117, 2009.

K. Neuhoff, M. Grubb, and K. Keats. Impact of the allowance allocation on prices and efficiency. Cambridge Working Papers in Economics 0552, Faculty of Economics, University of Cambridge, November 2005.

Ariel Pakes and Paul McGuire. Computing markov-perfect nash equilibria: numerical implications of a dynamic differentiated product model. *RAND Journal of Economics*, 25 (4):555–589, Winter 1994.

Jean Pierre Ponssard and N Walker. EU emissions trading and the cement sector: a spatial competition analysis. *Climate Policy*, 8(6):467–493, 2008.

S.L. Puller. Pricing and firm conduct in California's deregulated electricity market. *The Review of Economics and Statistics*, 89(1):75–87, 2007.

Stephen P. Ryan. The cost of environmental regulation in a concentrated industry. Technical report, MIT Working Paper, 2010.

Stephen P. Ryan and Catherine E. Tucker. Diversification and the dynamics of technology adoption. *Conditionally accepted at Quantitative Marketing and Economics*, 2010.

Andrew Sweeting. Market power in the england and wales wholesale electricity market 1995–2000. *Economic Journal*, 117(520):654–85, 2007.

C.D. Wolfram. Measuring duopoly power in the British electricity spot market. *The American Economic Review*, 89(4):805–826, 1999.

Ernst Worrell, Lynn Price, Nathan Martin, Chris Hendriks, and Leticia Ozawa Meida. Carbon dioxide emissions from the global cement industry1. *Annual Review of Energy and the Environment*, 26(1):303–329, 2001. doi: 10.1146/annurev.energy.26.1.303. URL http://www.annualreviews.org/doi/abs/10.1146/annurev.energy.26.1.303.

A Construction of Emissions Rates

Over half of the emissions from clinker production come from the chemical reaction that occurs when the calcium carbonate in limestone is converted into lime and carbon dioxide. To measure carbon dioxide emissions from calcination accurately, emissions factors can be determined based on the volume of the clinker produced and the measured CaO and MgO contents of the clinker. In the absence of this detailed plant-level information, we assume a default rate of 0.525 metric tons of carbon dioxide/metric ton of clinker (WBC, 2005).

The other major source of carbon dioxide emissions from clinker production is fossil fuel combustion. The preferred approach to estimating CO_2 emissions from fuel combustion requires data on fuel consumption, heating values, and fuel specific carbon dioxide emission factors. Although the Portland Cement Association (PCA) does collect plant level data regarding fuel inputs and fuel efficiency (i.e. BTUs per ton of cement), these data are disaggregated data are not publicly available. We do have data aggregated by kiln type and vintage. We use these data (reported in 2006), together with average carbon dioxide emissions factors, provided by the U.S. Department of Energy, to estimate kiln technology specific emissions intensities.

We consider three classes of kilns in particular: wet process kilns (i.e. older, less efficient technology), dry process kilns with preheater/precalciner, and a best practice energy intensity benchmark (Coito et al., 2005)[28] Because of the dominant role played by coal/pet coke, our benchmark emissions calculations are based on coal/petcoke emissions factors. We assume an emissions factor of 210 lbs carbon dioxide/mmbtu.[29]

Our technology-specific emissions rate calculations are explained below. To put these numbers in perspective, the national weighted average emissions rate was estimated to be 0.97 tons carbon dioxide/ton cement in 2001 (Hanle et al, 2005).

Wet process In 2006, there were 47 wet process kilns in operation. On average, wet kilns produced 300,000 tons of clinker (per kiln) per year. The PCA 2006 Survey reports an average fuel efficiency of 6.5 mmbtu/metric ton of clinker equivalent among wet process kilns.

[28] The industry has slowly been shifting away from wet process kilns towards more fuel-efficient dry process kilns. On average, wet process operations use 34 percent more energy per ton of production than dry process operations. No new wet kilns have been built in the United States since 1975, and approximately 85 percent of U.S. cement production capacity now relies on the dry process technology.

[29] Fuel-specific emissions factors are listed in the Power Technologies Energy Data Book, published by the US Department of Energy (2006). The emissions factors (in terms of lbs CO2 per MMBTU) for petroleum coke and bituminous coal are 225 and 205, respectively. Here we use a factor of 210 lbs CO2/MMBTU. This is likely an overestimate for those units using waste fuels and/or natural gas.

The relevant conversion is then 0.095 metric tons carbon dioxide/mmbtu * 6.5 mmbtu/metric ton of clinker equivalent = 0.62 tons carbon dioxide/ton clinker. When added to process emissions, we obtain our estimate of 1.16 tons carbon dioxide/ton clinker.

Dry process In 2006, there were 54 dry kilns equipped with precalciners with an average annual output of 1,000,000 tons of clinker per year. The PCA 2006 Survey reports an average fuel efficiency of 4.1 mmbtu/metric ton of clinker equivalent among dry process kilns with precalciners. Thus, 0.095 metric tons carbon dioxide/mmbtu * 4.1 mmbtu/metric ton of clinker equivalent = 0.39 tons carbon dioxide/ton clinker. Adding this to process emissions results in the estimate for dry-process kilns: 0.93 tons carbon dioxide/ton clinker.

Frontier technology To establish estimates for new entrants, a recent study (Coito et al, 2005) establishes a best practice standard of 2.89 mmbtu/ metric ton of clinker (not clinker equivalent). The calculation is then: 0.095 metric tons carbon dioxide/mmbtu * 2.89 mmbtu/metric ton of clinker equivalent= 0.275 tons carbon dioxide/ton clinker. Adding this to process emissions obtains in 0.81 tons carbon dioxide/ton clinker for new kilns.[30]

B Abatement response

In the simulation exercise, the state space is modified such that emissions rates vary systematically across plants of different vintages and technology types. Incumbent firms are classified as either wet-process, dry-process, or dry-process with precalciner/preheaters. New kilns are assumed to be state-of-the-art. This modification allows us to crudely capture changes in embodied emissions intensity as the industry evolves.

There are four main strategies for reducing the carbon intensity of domestic cement industry. First, it is anticipated that capital stock turnover will be a major driver of emissions intensity reductions (Worrell, 1999). Replacing old wet-process kilns with state-of-the-art dry kilns could deliver significant reductions in combustion-related emissions.

Second, the carbon intensity of clinker production can also be reduced via fuel switching. Currently, coal and petroleum coke are overwhelmingly the dominant fuel used in pyroprocessing and electricity is used to grind raw materials into kiln feed. Most domestic kilns are

[30]This is very similar to the CO2 emissions rate assumed in analyses carried out by California's Air Resources Board in 2008 under a best practice scenario that does not involve fuel switching. If fuel switching is assumed, best practice emissions rates drop as low as 0.69 MT CO2/ MT cement. See NRDC Cement GHG Reduction Final Calculations.

capable of burning a variety of fuels in principle, although fuel switching can adversely affect plant performance.

Third, concrete manufacturers have the capacity to partially substitute SCMs for clinker inputs. The advantage of this emissions reduction strategy is that, by reducing the use of clinker, carbon emissions from both fuel combustion and calcination are eliminated. Finally, cement manufacturers have some capacity to substitute less carbon intensive raw materials for limestone.

Data limitations will prevent us from being able to model input and fuel substitution capabilities accurately at the plant level. In our model, these two abatement options are ignored. In the policy simulations, carbon dioxide emissions from the domestic cement industry can be reduced via four channels: accelerated capital turnover (i.e. retirement of older kilns and investment in newer, more efficient operations), a reallocation of production from more to less emissions intensive incumbents, an increased reliance on imports, and a decrease in domestic clinker consumption. To the extent that fuel and input substitution are economically viable and cost effective compliance alternatives, our results will over estimate compliance costs and thus should be interpreted as upper bounds.