Online Appendix:

Market-based Emissions Regulation and the Evolution of Market Structure

Meredith Fowlie, Mar Reguant, and Stephen P. Ryan

November 26, 2012

A   Equilibrium equations

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, we 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:

$$\sigma_i : (s, \epsilon_i) \to a_i, \tag{1}$$

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 at the time of the exit decision is:

$$V_i(s; \sigma(s), \theta, \epsilon_i) = \bar{\pi}_i(s; \theta) + \max \left\{ \phi_i, E_{\epsilon_i} \left\{ \beta \int E_{\epsilon_i} V_i(s'; \sigma(s'), \theta, \epsilon_i) \, dP(s'; s, \sigma(s)) \right. \right.$$
$$\left. \left. + \max_{x_i^* \neq 0} \left[-\gamma_{i1} - \gamma_2 x_i^* + \beta \int E_{\epsilon_i} V_i(s'; \sigma(s'), \theta, \epsilon_i) \, dP(s_i + x^*, s'_{-i}; s, \sigma(s)) \right] \right\} \right\}, \tag{2}$$

1

where θ is the vector of payoff-relevant parameters, E_{ϵ_i} is the expectation with respect to the distributions of shocks, and $P(s'; \sigma(s), s)$ is the conditional probability distribution over future state s', given the current state, s, and the vector of strategies, $\sigma(s)$.

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. We 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.[1] 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:

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

(3)

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

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

(4)

for all s, ϵ_i, and all possible alternative strategies, $\tilde{\sigma}_i(s)$. As we work with the expected value functions below, we note that the MPNE requirement also holds after integrating out firms' private information: $E_{\epsilon_i} V_i(s; \sigma_i^*(s), \sigma_{-i}(s), \theta, \epsilon_i) \geq E_{\epsilon_i} V_i(s; \tilde{\sigma}_i(s), \sigma_{-i}(s), \theta, \epsilon_i)$. Doraszelksi 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 we discuss next in the context of our empirical approach.

B   Computation

Once the parameters have been estimated, the model can be computed to compare the market performance under market-based policy designs. In order to compute the equilibrium of the game, we make use of parametric approximation methods. In particular, we interpolate the value function

---

[1]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 (2012) for an example of such an optimal waiting problem.

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.[2] 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.[3]

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).[4] 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 s is approximated as,

$$\hat{V}_i^k(s) = \sum_{j=1}^{(J+2)^A} \phi_{N_A,j} B_{N_A,j}(s), \tag{5}$$

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 $B_{N_A,j}(s) = \prod_{i \in A} B_j(s_i)$.

---

[2]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).

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

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

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,

$$\max_{s_i'} \quad \pi_i(s_i, s_i'|s_{-i}) + \sum_{s_{-i}' \in S_{-i}} Pr^k(s_{-i}'; \sigma^k(s)) \hat{V}_i^k(s_i'|s_{-i}'). \tag{6}$$

We compute the best-response of a firm 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 best-response of the firm.[5] 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 $\| \sigma^{k+1} - \sigma^k \| < \epsilon_\sigma$ and $\| V^{k+1} - V^k \| < \epsilon_V$.

C 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 (**?**).

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

---

[5]Taking the actions of other firms as given, the cubic spline is defined by a cubic polynomial at each of the grid intervals, which implies that at most there will be $2(J-1)+2$ candidate local optima, where J is the number of grid points.

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 (**?**).[6] 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.[7]

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 (**?**).

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. 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 (**?**) establishes a best practice standard of 2.89 mmbtu/ metric ton of clinker (not clinker equivalent). The calculation

---

[6]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.

[7]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.

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.[8]

D   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 (**?**). 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 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. A more detailed discussion is relegated to the robustness section.

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

---

[8]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.

Table E.1: Estimation of Demand Elasticity

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

Notes: Robust standard errors in parenthesis. Unit of observation is market-year. Market fixed-effects included in all specifications. Sample 1980-2009.

alternatives, our results will over estimate compliance costs and thus should be interpreted as upper bounds.

E Estimation additional results

E.1 Demand estimation

We estimate the parameters of the demand equation using data collected annually by the USGS over the period 1981-2009. Regional average prices are reported in \$/metric ton. Cement shipped by local producers, which includes cement produced from imported clinker, is reported annually (aggregated regionally). We map these regional data into the regional market definitions described in the paper using capacity weights. For example, if a region (as defined by the USGS), encompasses two regional markets, prices and quantities are assigned to regional markets in proportion to installed capacity in each market.

The dependent variable in all specifications is the natural log of the cement shipments. The cost shifters that are used as instrumental variables for the cement price include energy prices (coal, natural gas, and electricity) and wages for skilled manufacturing in the cement sector as reported in the Census County Business Patterns. The independent variable for which the instruments are used is the natural log of the average cement price (in 2000\$, as are all other monetary values in the models).

Table **??** summarizes the estimation results. Robust standard errors are reported in parentheses. The first specification, which we highlight in the paper, is the most parsimonious as it includes only regional market fixed effects. The point estimate is -2.0. This specification omits several factors

that presumably shift demand (such as population, unemployment, and measures of construction activity). Subsequent specifications include these factors. Our point estimate of the own-price demand elasticity is somewhat sensitive to the inclusion of these covariates; estimates vary between -0.9 and -2.0

We select specification (1) as our preferred specification because it is the most parsimonious and most consistent with dynamic structural estimation. We do not explicitly model changes in population or building activity over time. That said, it is important to note that we cannot rule out elasticity estimates that are smaller in absolute value. In a series of robustness checks, we simulate policy outcomes over a range of demand elasticity values.

E.2   Import supply estimation

The import supply elasticity parameter is estimated with data collected annually by the USGS over the period 1993-2009. The dependent variable is the log of the quantity of cement shipped annually. Import price and quantity data are reported annually by customs district. Unfortunately, these data provide a very noisy measure of cement import prices and quantities. Reported quantities include all varieties of hydraulic cement and clicker. No information about the composition of the import flows coming through a customs district in a given year is provided. Additional noise is added to the data when we map customs districts to regional markets based on proximity.

Demand shifters used to instrument for import prices include an annual measure of construction activity (units constructed), economic activity (gross state product and unemployment rates). The most parsimonious specification includes only regional fixed effects. The estimated import supply elasticity is 2.5. This parameter is imprecisely estimated (the standard error is 3.7). An alternative specification includes a series of supply shifters, including coal prices, oil prices, a measure of the cost of transporting the cement from the supply country to the import district in the United States. To construct this last variable, we subtract the average customs price from the average C.I.F. price of the cement shipments. This residual price accounts for the transportation cost on a per unit basis, as well as the insurance cost and other shipment-related charges. Including these controls does not significantly affect our point estimate. The import supply elasticity is 2.5 (standard error 2.9).

It is important to emphasize that this important parameter is very imprecisely estimated. Although this is not surprising given how we must construct these data, it means we cannot be very confident in the elasticity estimate we choose to use (2.5). In a series of robustness checks, we simulate policy outcomes over a range of import supply elasticity values that we cannot rule out.

8

Table F.1: Environmental parameters

	(1)	(2)	(3)	(4)
Social discount rate β_S	5.0%	3.0%	2.5%	3.0%
SCC 2010[‡]	4.70	21.40	35.10	64.90
SCC 2020[‡]	6.80	26.30	41.70	80.70
SCC 2030[‡]	9.70	32.80	50.00	100.00

Source: U.S. Department of Energy (2010), "Final Rule Technical Support Document (TSD): Energy Efficiency Program for Commercial and Industrial Equipment: Small Electric Motors," Appendix 15A (by the Interagency Working Group on Social Cost of Carbon): "Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866".
[‡] Social cost of carbon in $ per metric ton of carbon dioxide in $2007.

F The Social Cost of Carbon

A deliberative interagency process was convened to generate estimates of the social cost of carbon for use in U.S. policy implementation and analysis (**?**).The four SCC schedules that were selected in this process are summarized in the table below. In light of disagreements about the appropriate choice of interest rate, three different discount rates are used (corresponding to the first three schedules). The final schedule (fourth column) corresponds to a scenario with higher than expected economic costs from climate change. The SCC increases over time because future emissions are expected to produce larger incremental damages as physical and economic systems become more stressed. In our analysis, we assume the carbon price does not change over the time horizon we consider.

References

Benitez-Silva, H., Rust, J., Hitsch, G., Pauletto, G., and Hall, G. (2000). A comparison of discrete and parametric methods for continuous-state dynamic programming problems. Computing in Economics and Finance 2000 24, Society for Computational Economics.

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

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

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

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

Habermann, C. and Kindermann, F. (2007). Multidimensional spline interpolation: Theory and applications. *Comput. Econ.*, 30(2):153–169.

Judd, K. (1998). *Numerical Methods in Economics*. MIT Press.

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

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

Ryan, S. P. and Tucker, C. E. (2012). Heterogeneity and the dynamics of technology adoption. *Quantitative Marketing and Economics*, 10:63–109.