

On-Line Appendix
Escalation of Scrutiny: The Gains from Dynamic
Enforcement of Environmental Regulations
Wesley Blundell, Gautam Gowrisankaran, and Ashley Langer

A1. Illustrative Simple Case of our Model

To illustrate the value of dynamic enforcement, we present a simple, special case of our general model. For this special case, we assume a two-period model with $\beta = 1$. In both periods, a single new violation occurs with probability p . Violations are costly to the regulator in that they may result in emissions, but costless to the plant. Inspections occur with probability $\mathcal{I}(\Omega) = 1$, are costless to the regulator and plant, and perfectly reveal the presence of a violation. The regulator assigns the plant to compliance if it has 0 outstanding violations; regular violator status if it has 1 outstanding violation; and HPV status if it has 2 outstanding violations. We assume that the signals for violations and fines are the same: $e_t^1 = e_t^2$, and indicate the number of outstanding violations. Note that $e_1^1 \in \{0, 1\}$ and $e_2^1 \in \{0, 1, 2\}$. The remaining signals follow from the above description of the status transitions, with $e_t^3 = \mathbb{1}\{e_t^1 = 0\}$, $e_t^4 = \mathbb{1}\{e_t^1 = 1\}$, and $e_t^5 = \mathbb{1}\{e_t^1 = 2\}$.

A period 1 investment, $X_1 = 1$, clears a period 1 violation with probability q ; violations are never cleared without investment. The pollution cost to the regulator is $c_E e_t^1$ at period t , for some marginal pollution damage parameter c_E . The regulatory state records the history of investments and violations. Thus, for example, at period 2, the regulatory state after the inspection is $\tilde{\Omega}_2 = (X_1, e_1^1, e_2^1)$. Finally, the per-period objective function to the plant is $-\theta^X X_t - \text{Fine}(\Omega_t, e_t^1)$, where θ^X is the cost of investment. The regulator minimizes the sum over the two periods of $c_E e_t^1$, $\theta^X X$, and its cost of assessing fines.

We allow the regulator to pre-commit to an enforcement strategy and focus on the case with a period 1 violation—so $e_1^1 = 1$ —as this is the only case where the regulator might want to incentivize period 1 investment. The simplest policy that a regulator could choose would be a linear fine policy $c_F e_t^1$. When θ^X is known and contractable and the cost of investment is sufficiently low relative to other costs, the regulator incentivizes period 1 investment by choosing the lowest c_F that would compel the plant to invest.

With a linear fine policy, the regulator has to issue fines for the period 1 violations even though this has no effect on investment. Thus, this fine lowers the regulator objective function. An alternative is for the regulator to choose a static escalation mechanism: it could fine only when $e_t^1 = 2$, which would remove the cost of fining when $e_t^1 = 1$ but the plant has not had a chance to invest, and would still incentivize investment in period 1. For this reason, the regulator can incentivize investment for the same values of θ^X as the linear fine policy with lower expected fines, thereby adding surplus. Although the model is dynamic, this escalation mechanism is not explicitly dynamic (since it does not depend on

the regulatory state Ω_t , but only on the current number of outstanding violations, e_t^1). It increases marginal deterrence in period 2 since it will result in no fines in period 1. Because expected fines are lower, the regulator will further choose to incentivize investment for more values of θ^X , thereby adding further surplus in some cases.

A dynamic escalation mechanism would increase surplus relative to the static escalation mechanism. In this case, the regulator could fine when $e_{t-1}^1 > 0$, $X_{t-1} = 0$, and when it wants to incentivize investment. Choosing this policy for the same set of θ^X as above will mimic the same investment incentives but with no fines paid in equilibrium (since plants whose investment does not succeed in returning the plant to compliance are not fined), and hence no fine costs. Thus, the regulator will choose to incentivize investment for even more values of θ^X .

If instead of a single θ^X the regulator faces a distribution of θ^X values and cannot contract on θ^X , dynamic enforcement also adds value by better selecting the set of plants which it incentivizes to invest. For simple investment cost type distributions, the regulator will incentivize investment for more values of θ^X with dynamic enforcement than with a static escalation mechanism or with linear fines.

Overall, our illustrative simple case shows that escalation mechanisms add value by increasing the marginal deterrence for two violations relative to one. Dynamic escalation mechanisms add more value by reducing equilibrium fines and by increasing the set of actions over which the regulator can condition.

A2. Data Construction Details

ECHO Database Overview

The ECHO database is divided into a number of components. We principally use four ECHO components: (1) the *Facility Registry Service* dataset, (2) the *Air Facility System Actions* dataset, (3) the *Air Program Historical Compliance* dataset, and (4) the *High Priority Violator History* dataset. We discuss each of these components in turn.

First, the *Facility Registry Service* dataset is a master list of plants. For our purposes, it provides address information and the six-digit North American Industry Classification System (NAICS) industrial sector for each plant. Our analyses control for the EPA region, the first two digits of the NAICS code, and the expected gravity of violations based on industry and county. We keep seven industries with high pollution damages that we believe to have plants of broadly comparable costs of investment and enforcement: the three manufacturing industries, mining and extraction, transportation, educational services (which includes school buses), and utilities.

Second, the *Air Facility System Actions* (AFS) dataset (or *Actions* dataset for short) records the history of regulatory actions taken by state, regional, and federal environmental regulators, from Q4:2006 through the Q4:2014.³³ We use this

³³The EPA transitioned to a new reporting system after 2014.

dataset to create our base list of inspections, violations, fines, and investments. Since this dataset is subject to federal minimum data requirements, we believe it provides a relatively complete description of the regulatory action history for each plant. One potential issue with our data is that some states were not reporting non-HPV violations prior to 2010.³⁴ The EPA customer support staff were not sure if the data had been corrected and suggested we review the data for anomalous changes in the violation rate. We examined the data for changes in the prevalence of violations in 2010 by performing a series of regressions of reported violations on state or region dummies interacted with a dummy for post-2010. We found no systematic evidence of an increase in reported violations, suggesting that this error had been corrected ex post.

Each record in this dataset details a regulator action, such as an inspection, a notice of violation, a fine, or the review of an investment in pollution abatement. The unit of observation is the AFS ID, which indicates a polluting source. Each record lists a calendar date and provides information on the related EPA program³⁵ and the penalty amount when the action is a fine.³⁶ For each plant, we combine EPA actions across all EPA programs to capture completely its regulatory enforcement status.

Third, the *Air Program Historical Compliance* dataset records the historical compliance status for each plant and EPA program at the AFS ID and quarter level. These data derive from a combination of self-reports by plants and regulator inputs. We follow the literature (Laplante and Rilstone, 1996; Shimshack and Ward, 2005) in treating the self-reported data as accurate.³⁷ We use this dataset to determine whether a plant is in compliance or a violator in any quarter. This dataset provides a more direct measure of violator status than does the *Actions* dataset, since the *Actions* dataset does not always indicate when a violation is resolved. Since this dataset is at the plant / quarter level, we aggregate EPA actions to this level and use this as the time period for our analysis. We also use this dataset to determine whether a plant has shut down, dropping plants from the sample once they have exited.

Fourth, the *High Priority Violator History* dataset records the dates at which a plant receives or resolves a high priority violation. We use this dataset to record the quarter of entry and exit from HPV status. Analogous to the *Air Program Historical Compliance* dataset, this dataset provides the most direct measure of HPV status.

³⁴See <https://echo.epa.gov/system/files/FRVMemoandAppxFinal3.22.10.pdf>.

³⁵The CAAA include many different statutes that address different dimensions of air pollution. The EPA enforces different statutes through different programs.

³⁶It is possible for plants to contest fines in court. However, Helland (2001) finds that fewer than 4% of fines are successfully contested by plants, a number that is in keeping with our own analysis of the Integrated Compliance and Information System's (ICIS) Federal Enforcement and Case Data.

³⁷The literature makes this assumption because the expected penalty from purposefully deceiving regulators is far greater than the penalty for an emissions violation.

Regulatory Actions and Outcomes

Compliance and violator statuses. During our sample period, the EPA’s *Air Program Historical Compliance* dataset reported each plant’s compliance status for every CAAA program. Since there is a CAAA program for each major category of air pollutant, a plant can simultaneously be in violation of multiple CAAA programs. We assume that a plant is a CAAA violator if it is a violator for any CAAA programs. For each program, we classify a plant as being a violator if compliance status is equal to “1” (in violation, no schedule), “6” (in violation, not meeting schedule), “7” (in violation, unknown with regard to schedule), “B” (in violation with regard to both emissions and procedural compliance), “D” (HPV violation), “E” (federally reportable violation), “F” (High Priority Violator on schedule), “G” (facility registry service on schedule), or “W” (in violation with regard to procedural compliance).³⁸

The *Historical Compliance* dataset also reports codes indicating an unknown compliance status: “Y” (unknown with regard to both emissions and procedural compliance), “0” (unknown compliance status), “A” (unknown with regard to procedural compliance), and “U” (unknown by evaluation calculation). From our discussions with the EPA, these codes arise when a plant has not been inspected within the required time frame, but there has been no indication of a violation by the plant. Given this, we code these plants as being in compliance.³⁹ In some cases, we observe a violation at some quarter t in the *Actions* dataset and the plant is reported to be a violator at quarter $t + 1$ but not at quarter t . In these cases, we assume that the reporting that indicated that the plant was in compliance at quarter t was erroneous, and hence we record the plant as being in violator status at quarter t .

We code all other plants—except those that are listed as HPVs in the *High Priority Violator History* dataset—as being in compliance. Thus, we do not use additional information on compliance in the ECHO database for some plants and pollutants, such as continuous emissions monitoring system reports.

Inspections. The *Air Facility System Actions* dataset reports multiple types of inspections, which we collapse into a single “inspection” variable. These include on- and off-site full compliance evaluations conducted by either the federal or state EPA, partial compliance evaluations, and stack tests. We also consider an inspection to have occurred if the EPA issues a Section 114 letter for gathering information from the plant. In some cases we observe multiple inspections in the same quarter; e.g., if stack tests are conducted for multiple pollutants. Since our inspection variable is dichotomous, we consider these tests together to be equivalent to a single inspection.

Violations. The *Actions* dataset also reports violations. We define a violation to be the issuance of a “Notice of Violation” (NOV). An NOV is defined as “a

³⁸Although this list indicates both plants that are regular violators and HPVs, we determined HPV status from the *High Priority Violator History* dataset, for greater accuracy.

³⁹Evans (2016) also considers plants in unknown compliance status to be in compliance.

notice sent by the State/EPA ... for a violation of the Clean Air Act.” There are three codes that indicate an NOV in our data: “6A” (EPA NOV issued), “7A” (notice of noncompliance), and “7C” (state NOV issued).⁴⁰ In some cases, we observe a violation at some quarter t in the *Actions* dataset but the plant is not reported to be a violator in the *Historical Compliance* dataset at quarter t or $t+1$ and did not receive a fine at quarter t . We believe that these violations likely reflect minor issues that are dissimilar to other violations, and hence we exclude them from our analysis.

Plant Exits

The *Historical Compliance* dataset also allows us to understand when plants shut down. Plants may have a compliance status of “9” (in compliance: shut down). If we observe a plant in this status, we assume that it has exited. We remove it from our sample for the quarter with this status and all subsequent quarters.

Investment

Our data do not directly report investments or investment costs (unlike in the Duflo et al., 2018, study of pollution in India, for instance). Instead, we infer investments from the behavior of EPA regulators. We determine that an investment occurred if we observe any of the following three types of events: (1) the resolution of a major violation, (2) the issuance of a Prevention of Significant Deterioration (PSD) permit, and (3) exit from HPV status. We now provide detail on each of these categories.

First, the overwhelming majority of our investments come from codes that indicate the resolution of a major violation. There are three codes in the *Actions* database that we consider evidence of this type of investment: (1) “VR” or “violation resolved,” (2) “OT” or “other addressing action,” and (3) “C7” or “closeout memo issued.” According to the November, 2008 *Air Facility Systems National Action Types-Definitions* EPA document,⁴¹ “a violation is resolved when it is addressed and a closeout memo has been issued, all penalties have been collected and the source is confirmed to be in physical compliance.”⁴²

Similarly, “other addressing action” is an addressing action for HPV cases with criminal or civil action referrals. Finally, “a closeout memo is issued when a violation is resolved with all penalties collected and the source is confirmed to be in physical compliance.” Of the investments that are determined by a resolution code, we observe “VR” for the overwhelming majority (77%). An additional 14% of these investments are from “C7”, and the remaining 10% are from “OT.”

⁴⁰See https://echo.epa.gov/files/echodownloads/AFS_Data_Download.pdf.

⁴¹Downloaded September 2014.

⁴²Note that we do not always observe “VR” or other investment codes when plants return to compliance from regular violator status. Thus, we allow for the possibility that plants can return to compliance from regular violator status without an investment.

Second, a PSD permit is required for new pollution sources or for major modifications of existing sources.⁴³ While it is possible that major modifications of existing sources may occur for reasons other than a plant attempting to return to CAAA compliance, we believe that changes to a plant that were substantial enough to warrant a new PSD permit issuance likely imply a major investment in pollution abatement.

Finally, we also infer that an investment has occurred if a plant exits HPV status, even if we do not observe one of these codes. We make this choice because we believe that a major investment would have been necessary in order to resolve the substantial violations that would have originally merited the determination of HPV status as well as all outstanding violations.

To verify that our measure of investment does indeed capture investments in pollution abatement capital equipment, we collected additional data from the Texas Commission on Environmental Quality (TCEQ). The TCEQ data provide information on the installation and removal of pollution control devices for all plants covered by Texas Administrative Code, Title 30, Rule 101.10. This regulation applies to plants with the highest emissions, which is a subset of plants in Texas that are regulated by the EPA. The installation of control devices forms a direct marker of an investment, corresponding to our definition.

We matched the Texas data manually to our base data using firm/regulated entity name, city, and address. Although the set of plants that is regulated by this statute is a subset of the set that show up in our EPA data, we are able to match 1,044 out of 2,109 of the EPA plants in Texas to a plant in the TCEQ data. In all, the TCEQ data contained 1,520 plants with a change in an emissions source or abatement device during our period, so our 1,044 matched observations represent 69% of these. (Note also that not every plant covered by this regulation will have an abatement device and that the TCEQ data cover more industries than the 7 in our study, but the TCEQ data do not report industry.) Overall, we believe that our match rate is high enough to make meaningful statements regarding the abatement device changes for larger plants in Texas.

We first investigated whether an investment in the EPA dataset correlated with the installation of an abatement device in the TCEQ data. One issue is that the timing of investment in the two datasets is somewhat different. On the one hand, the EPA data record an indirect measure of investment that only appears in the data once the EPA has confirmed that the violation has been resolved and hence we might expect the EPA measure to lag the Texas measure. On the other hand, the Texas measure of investment only occurs after TCEQ has recorded it in their system following a plant visit, which is supposed to occur within a year of the device installation. TCEQ also does not require self-reporting for abatement devices. Thus, the TCEQ measure may lag the EPA measure.

Despite these limitations, we find a strong and significant relationship between the EPA investment measure and the TCEQ abatement device installation mea-

⁴³See <https://www.epa.gov/nsr/prevention-significant-deterioration-basic-information>.

sure. Specifically, we found that 45% of EPA investments have a TCEQ abatement device installation within four quarters conditional on the plant being observed in the both datasets (and unconditionally, the figure is 29%). Similarly, a regression of EPA investment on TCEQ abatement device installation within four quarters gives a coefficient of 0.031 with a t-stat of 16.9.

We also used the TCEQ abatement device measure to determine whether additional EPA actions should be included in our measure of investment. We identified three groups of actions that could plausibly be added: (1) an indicator for whether a penalty was paid (C3); (2) an indicator for a violation being withdrawn (WD); and (3) indicators for the EPA determining that the plant was no longer deemed to be in violation due to a rule change or to the plant not being subject to the rule (2L, 2M, NM, NN). Overall, we found only 18 of these actions, compared to 1,094 EPA investments for plants in Texas. Of these 18, only 5 had a TCEQ abatement device change within 4 quarters. Thus, we decided not to add these codes to our definition of investment.

Finally, we investigated whether the installation of an abatement device in compliance in the TCEQ data predicted avoidance of violator status. Specifically, we regressed exit from compliance on recent TCEQ abatement device installation, defined as a TCEQ abatement device installation in the current quarter or within the previous four quarters. We find that, similar to EPA investment, TCEQ abatement device installation in compliance actually increases the likelihood of future violator status. Also, as with the EPA investment variable, TCEQ abatement device installation in violator status predicts a return to compliance.

Pollution and Damages Data

National Emissions Inventory data. We match 59% of observations in the ECHO data for 2008 and 2011 to the NEI data. The imperfect match is consistent with other studies that use the NEI data; e.g., Shapiro and Walker (2018) achieve a 77.4% match rate between the NEI and the Census of Manufacturing. We measure smokestack emissions for six pollutants: PM_{2.5}, NO_x, SO₂, volatile organic compounds, NH₃, and Pb. For our counterfactuals, we need the expected level of pollution by regulatory state. To obtain this, we aggregated the matched NEI data to the region, industry, gravity state, and compliance / regular violator / HPV status level. We then calculated the mean pollution for each of these states, imputing missing values. We did not use the full regulatory state here given the limited number of matching observations in the NEI data for some states.

AP3 data. The AP3 data come from an integrated assessment model that explicitly considers the impact of pollution emitted in different locations, and thereby takes into account differences in local populations and underlying pollution levels. While we consider the damages from criteria air pollutants—ozone (O₃), particulate matter (PM), carbon monoxide (CO), nitrogen oxides (NO_x), sulfur dioxide (SO₂), and Pb—the AP3 data include damages from smokestack

emissions that can lead to criteria air pollutants—PM_{2.5}, NO_x, SO₂, volatile organic compounds (a precursor to ozone), and NH₃ (a precursor to PM).

National Ambient Air Quality Standards attainment data. We consider NAAQS attainment status for each pollutant covered during this period. In particular, we use information on non-attainment for 8-hour ozone (1997 and 2008 standards), carbon monoxide (1971 standard), lead (1978 and 2008 standards), PM-10 (1987 standard), and PM-2.5 (1997 and 2006 standards) in each year from the EPA’s “Green Book.” We do not include information on the 1979 1-hour ozone standard because it was revoked on June 15, 2005; the 1971 nitrogen dioxide standard because all areas were in attainment as of September 22, 1998; or the 2010 sulfur dioxide standard because the original areas were not designated until October 4, 2013, after the end of our sample period.

A3. Details on Empirical Framework

Plant Dynamic Optimization

A plant that is not in compliance makes an investment decision in each period, knowing that the investment will reduce its expected future cost of regulatory enforcement. The plant’s optimization therefore requires evaluating the value of being in a given state, Ω , at the start of the next period.

Let $V(\Omega)$ denote the value function at the beginning of the period, $\tilde{V}(\tilde{\Omega})$ denote the value function at the point right after the regulator has moved but before the plant receives its draws of ε , and $Com(T)$ be an indicator for T designating compliance.⁴⁴ We first exposit $V(\Omega)$, the value function at the beginning of the period:

$$(A1) \quad V(\Omega) = \sum_{i \in \{0,1\}} \mathcal{I}(\Omega)^i (1 - \mathcal{I}(\Omega))^{1-i} \int [U(\Omega, e) + \tilde{V}(T(\Omega, e))] dP(e|\Omega, \mathcal{I}, i),$$

where $dP(e|\Omega, \mathcal{I}, i)$ is the integral over the density of the environmental compliance signal e given the plant state, the inspection policy, and the inspection decision. Note that the plant does not make any decision at the beginning of the period, and hence there is no maximization in (A1). However, the plant must integrate over the regulator policies and e .

We now exposit $\tilde{V}(\tilde{\Omega})$:

$$(A2) \quad \begin{aligned} \tilde{V}(\tilde{\Omega}) &= Com(\tilde{\Omega}) \times \int [\beta V(\tilde{\Omega}, \theta) + \varepsilon_0] dF(\varepsilon_0) + (1 - Com(\tilde{\Omega})) \times \\ &\int \int \max\{\beta V(\Omega|\tilde{\Omega}, X=0) + \varepsilon_0, -\theta^X + \beta V(\Omega|\tilde{\Omega}, X=1) + \varepsilon_1\} dF(\varepsilon_0) dF(\varepsilon_1) \\ &= Com(\tilde{\Omega}) [\beta V(\tilde{\Omega}, \theta) + \gamma] + (1 - Com(\tilde{\Omega})) \times \\ &[\ln(\exp(\beta V(\Omega|\tilde{\Omega}, X=0)) + \exp(-\theta^X + \beta V(\Omega|\tilde{\Omega}, X=1))) + \gamma], \end{aligned}$$

⁴⁴For ease of notation, we are conditioning on the plant’s parameter vector θ .

where $dF(\cdot)$ is the integral over the density of the type 1 extreme value distribution. The first part of (A2) reflects the case of compliance. In this case, the plant transitions to the same state $\tilde{\Omega}$ in the next period. Since there is no plant choice here, in expectation, the plant receives the continuation value plus the mean value of the type 1 extreme value distribution which is γ , Euler's constant. The second part of (A2) reflects the case of a plant that is a violator or high priority violator. In this case, it makes a choice of whether to invest or not. Since the value is computed ex ante to the realization of the idiosyncratic draws, we can use the familiar logit aggregation. The transition state, though still not stochastic, is now potentially different from the current state, because it updates both lagged investments and depreciated accumulated violations.

Finally, having defined the value functions, we can write the probability of a plant choosing investment given a regulatory state $\tilde{\Omega}$ and its cost and utility parameters θ as:

$$(A3) \quad \Pr(X = 1 | \tilde{\Omega}, \theta) = \frac{(1 - Com(\tilde{\Omega})) \exp(-\theta^X + \beta V(\Omega | \tilde{\Omega}, X = 1))}{\exp(-\theta^X + \beta V(\Omega | \tilde{\Omega}, X = 1)) + \exp(\beta V(\Omega | \tilde{\Omega}, X = 0))}.$$

Since the probability in (A3) is used in our estimators, we have written it as a function of the structural parameter vector θ .

Computation of Bellman Equation

The plant's decision as to whether or not to invest at any state is based on dynamic optimization. As such, we solve for the Bellman equation for candidate parameter values, based on equations (A1) and (A2). Specifically, for our quasi-likelihood estimator, we perform a non-linear search for θ and hence we solve for the Bellman equation for each of the candidate values of θ that are considered in the course of the non-linear search. For our GMM estimator, we solve for the Bellman equation for each of the 10,001 values in our fixed parameter grid.

The states in Ω and $\tilde{\Omega}$ are discrete, except for depreciated accumulated violations. Our Bellman equation discretizes this latter variable, using 20 grid points that are evenly spaced from 0 to 9.5. The transition from $\tilde{\Omega}$ to Ω , given in (A2), will result in a new level of depreciated accumulated violations that does not necessarily correspond to a grid point. As such, we use linear interpolation to calculate (A2).

The transition from Ω to $\tilde{\Omega}$, given in (A1), is stochastic, as it depends on the regulatory CCP. We perform this calculation by simulating from the estimated regulator CCP. Specifically, we first calculate the inspection probability for each state from the predicted values of our estimates. We then calculate the violation probability for each state and inspection decision. Following this, we calculate the distribution of fines for each state, inspection decision, and violation decision, using 20 evenly spaced points from the estimated residual distribution—which we denote F —ranging from $R^{-1}(0.025)$ to $R^{-1}(0.975)$. Finally, we calculate the

transition probabilities between the three statuses of compliance, regular violator, and HPV, for each state, inspection decision, violation decision, and discretized fine decision.

Altogether, this gives 240 ($2 \times 2 \times 20 \times 3$) possible regulatory outcomes using our discretized method. We calculate the probability and mean fines for each one. The Bellman equation then integrates over these possibilities. We compute our Bellman equation until a fixed point, defined as a sup norm tolerance of 10^{-9} between subsequent iterations. Following Assumption 1, when we compute Bellman equations under counterfactual policy environments, the state-contingent inspection, violation, and transition probabilities remain the same as in the base computations.

Empirical Implementation of Homogeneous Coefficients Model

In addition to our main random coefficient model, we estimate a model with homogeneous coefficients θ using a quasi-likelihood nested fixed point estimator. We calculate a quasi-likelihood (and not a likelihood) because we use the regulator's estimated CCPs in the plant's dynamic optimization process. In this model, there are no serially correlated unobservables for a plant over time, and hence, we can treat each plant i and quarter t as an independent observation. The quasi-log-likelihood of a parameter vector θ is:

$$(A4) \quad \log L(\theta) = \sum_i \sum_t \log \left(\left[X_{it} Pr(X = 1 | \tilde{\Omega}_{it}, \theta) + (1 - X_{it})(1 - Pr(X = 1 | \tilde{\Omega}_{it}, \theta)) \right] \right),$$

where the $Pr(X = 1)$ values are obtained from investment probabilities at the fixed point of the Bellman equation.

Our nested fixed point estimator is similar to Rust (1987). One difference is that in Rust (1987), the state transitions conditional on actions are exogenous, while here, they derive from the regulator's CCPs, making our estimator consistent with a dynamic game.⁴⁵ We obtain inference for our parameters and counterfactuals by bootstrapping our entire estimation process including the regulator's CCPs, with resampling at the plant level.

Choice of Fixed Grid Values for GMM Estimation

Our fixed grid estimator requires the ex ante specification of potential parameter grid values. We follow Fox et al. (2016) and first estimate the quasi-likelihood model and then center our fixed grid on these estimates. This requires specifying a range for the parameter grid around the quasi-likelihood estimates. We used a range of 15 (from 7.5 below the quasi-likelihood model to 7.5 above) for investment

⁴⁵We could also estimate the plant's utility function with a CCP estimator (Aguirregabiria and Mira, 2007), which is quicker to compute, but we did not, since the computational time for the nested fixed point quasi-likelihood estimator is not excessive.

and 5 for the other parameters. We chose these ranges after experimenting to make sure that they were large enough that we did not have parameters with positive weights near the boundary.

We choose our actual grid values by again following Fox et al. (2016) and using co-prime Halton sequences for each parameter, using the first five prime numbers, since each plant has five parameters. We scale the Halton sequences over the range between the minimum and maximum values. Co-prime Halton sequences better cover the set of parameters than would taking the interaction of the same grid points for each component (Train, 2009).

We dropped the first 20 elements of each Halton sequence as recommended in the literature (Train, 2009). We use the next 10,000 elements of the Halton sequences plus the quasi-likelihood estimates themselves as our fixed grid; hence $J = 10,001$. We also experimented with $J = 8,001$ (using the first 8,000 elements of the Halton sequence) and found similar results.

Inputs to Moments

As noted in Section III.C, we have three sets of moments. In order to explain our moments, order the states $1, \dots, K$ and let ω_k^1 denote the fixed component of state k and ω_k^2 denote the variable component of state k . Then, let $\pi_k(\theta)$ be the steady state share of plants at ω_k^2 given ω_k^1 . For a given ω_k^1 , we recover the associated $\pi(\theta)$ values by solving the Bellman equation for ω_k^1 , generating the transition matrix between variable states, and finding the vector that is invariant when transformed by this matrix.

As in (3), each moment is constructed from some m_k^d and $m_k(\theta_j)$. We now denote these terms m_k^1 , m_k^2 , and m_k^3 , and m_k^{d1} , m_k^{d2} , and m_k^{d3} , corresponding to our three sets of moments. Our first set of moments indicates differences in the steady state share of plants π_k between the model and the data. Specifically, for any moment $G_k(\eta) = m_k^{d1} - \sum_{j=1}^J \eta_j m_k^1(\theta)$, we let:

$$(A5) \quad m_k^1(\theta_j) = \pi_k(\theta_j),$$

and

$$(A6) \quad m_k^{d1} = \frac{\sum_i \sum_t \mathbb{1}\{\tilde{\Omega}_{it}^2 = \omega_k^2, \tilde{\Omega}_{it}^1 = \omega_k^1\}}{\sum_i \sum_t \mathbb{1}\{\tilde{\Omega}_{it}^1 = \omega_k^1\}}.$$

We note a few points about these moments. This first set of moments follows closely from Nevo et al. (2016), although we use the steady state distribution of our infinite-horizon dynamic problem, while they use the actual distribution of their finite-horizon problem. While in principle we could construct a moment from every $\tilde{\Omega}$, this would be difficult in practice given that we have over 50,000 states. Hence, we create moments for the 5,000 states which have the highest expected number of steady state observations at our estimated quasi-likelihood

parameters and given our data on $\tilde{\Omega}^1$.

Our second set of moments also follows closely from Nevo et al. (2016). The m_k values for these moments are constructed from the conditional steady state share of plants at any variable state times the conditional share having an investment at that state:

$$(A7) \quad m_k^2(\theta_j) = \pi_k(\theta_j) \times Share[X = 1 | \tilde{\Omega}, \theta_j],$$

and

$$(A8) \quad m_k^{d2} = \frac{\sum_i \sum_t \mathbb{1}\{\tilde{\Omega}_{it}^2 = \omega_k^2, \tilde{\Omega}_{it}^1 = \omega_k^1, X_{it} = 1\}}{\sum_i \sum_t \mathbb{1}\{\tilde{\Omega}_{it}^1 = \omega_k^1\}}.$$

We compute these moments for every state for which we compute our first set of moments, except for states that reflect compliance, as there is no investment in these states.

Our final set of moments explicitly captures the panel data aspect of investment. The m_k values for these moments are constructed from the conditional steady state share of plants at any variable state times the conditional share having an investment at that state times the sum from 1 to 6 of the product of the number itself and the conditional share with that many investments in the next six periods:

$$(A9) \quad m_k^3(\theta_j) = \pi_k(\theta_j) \times Share[X = 1 | \tilde{\Omega}, \theta_j] \times \left(\sum_{s=1}^6 s \times Share[s \text{ investments within 6 periods} | X = 1, \tilde{\Omega}, \theta_j] \right),$$

and

$$(A10) \quad m_k^{d3} = \frac{\sum_i \sum_t \left[\mathbb{1}\{\tilde{\Omega}_{it}^2 = \omega_k^2, \tilde{\Omega}_{it}^1 = \omega_k^1, X_{it} = 1\} \times \left(\sum_{s=1}^6 X_{i,t+s} \right) \right]}{\sum_i \sum_t \mathbb{1}\{\tilde{\Omega}_{it}^1 = \omega_k^1\}}.$$

These moments seek to match the extent of repeated investments by plants in the data to the model. A more traditional correlation moment would simply multiply investment at time t with investment at time $t + 1$ rather than with investment over the following six periods. We chose this formulation because we worry that investment in two subsequent quarters might partly reflect measurement error. We compute these moments for every state for which we compute our second set of moments.

To calculate the investment in the 6 periods ahead in (A9), we integrate over all potential paths conditioning on the initial state and investment decision. Each period there are ten potential paths: every interaction of (1) investment or not, (2) violation or not, and (3) regular violator and HPV statuses; plus the cases

of compliance with and without violations, but without investment.⁴⁶ Over 6 periods, this then implies $10^6 = 1,000,000$ possible paths for each parameter vector in our fixed grid θ_j . Thus, calculation of m_k for this set of moments is time consuming.

Overall, our estimator for our base specification has 14,374 moments, composed of 5,000 of the first set and 4,687 each of the second and third set. Our computation of $m_k(\theta_j)$ results in a $14,374 \times 10,001$ matrix and takes approximately eight days on an iMacPro with eight processors, with code written in C with MPI, or two days on the University of Arizona high performance cluster, using 28 processors.

Weighting Matrix and Estimation of GMM Parameters η_j

We follow the standard approach in GMM estimation of weighting by an estimate of the inverse of the variance-covariance matrix to improve the efficiency of our estimates.⁴⁷ We proceed in two stages. In stage 1, we estimate the model with a weighting matrix that does not reflect an asymptotic approximation to the variance-covariance matrix. Then, we use our stage 1 estimates to compute an approximation to the variance-covariance matrix.⁴⁸ In stage 2, we reestimate our parameters using this weighting matrix. We now detail our computation of the variance-covariance matrix for both stages.

In stage 1, we calculate the variance-covariance matrix of the moments inputs m_k , at the quasi-likelihood estimates θ_Q .⁴⁹

We calculate the diagonal elements of this matrix as:

$$(A11) \quad \text{Var}(m_k(\theta_Q)) = \frac{E[m_k(\theta_Q)m_k(\theta_Q)] - E[m_k(\theta_Q)]^2}{N_k},$$

where N_k is the number of plant / quarter observations from the region, industry, and gravity state for moment k . This is the general formula for the variance for the mean of N_k repeated *i.i.d.* draws from a random variable.

For the off-diagonal elements, the covariance will be zero for moments with different values of $\tilde{\Omega}^1$. We can write the covariance between moments k and l from the same $\tilde{\Omega}^1$ as:

$$(A12) \quad \text{Cov}(m_k(\theta_Q), m_l(\theta_Q)) = \frac{E[m_k(\theta_Q)m_l(\theta_Q)] - E[m_k(\theta_Q)]E[m_l(\theta_Q)]}{N_k}.$$

⁴⁶To save computational time, we use the higher probability point for depreciated accumulated violations, rather than linear interpolation.

⁴⁷Our GMM estimator is non-standard in that it includes the constraints in (2), which limits our ability to prove asymptotic efficiency of this estimator.

⁴⁸We base our approximation on the stage 1 parameters with weights of 0.01 or greater.

⁴⁹For some robustness specifications, we had collinearity issues with inverting this variance-covariance matrix. We dropped moments with zero variance in one specification and used the diagonal of the matrix for another specification.

The first term in (A12) will be non-zero only for the three moments that pertain to the same state. In this case, the first term in the numerator of the covariance between the first and second set of moments will equal the second moment, while the first term in the numerator between the first and third set of moments or between the second and third set of moments will equal the third moment. The reason for this is that the moment from the second set will only be non-zero when the moment from the first set is non-zero, while the moment from the third set will only be non-zero when the moment from the second set is non-zero. The second term in (A12) is simply the product of the means.

In stage 1, we invert and take a Cholesky decomposition of this estimated variance-covariance matrix. We then pre-multiply $m_k(\theta_j)$ for each θ_j and m_k^d by this matrix and obtain stage 1 estimates of the weights η_j by minimizing the linear system of equations in (3) subject to the constraints in (2), via constrained least squares. We use the Matlab package `lsqin` to perform this minimization process, which takes approximately 10 minutes on an iMacPro. The process generates consistent estimates of η that we use to construct a weighting matrix.

We then estimate the variance-covariance matrix of $G(\eta)$ using our stage 1 GMM estimates of η . From (3), the variance of $G(\eta)$ is simply the squared weighted sum of the variance conditional on the individual parameters, since the probability of each individual parameter occurring is independent across observations.

We again take a Cholesky decomposition of the inverse of this revised variance-covariance matrix, pre-multiply the matrix of moments $m_k(\theta_j)$ across all θ_j values, and re-run our estimation of the η_j weights. This provides our stage 2 estimates of η_j , which are the ones that we report.

Bootstrap Procedure for Inference

We bootstrap to obtain standard errors for both our quasi-likelihood and GMM estimates. For our GMM estimates, we provide standard errors on the counterfactual estimates only rather than also on the structural parameters.

Our bootstrap for the GMM estimator proceeds with the following repeated procedure:

- 1) We first draw an alternative dataset by sampling with replacement at the plant level. The new dataset has the same number of plants as the original data, though not necessarily the same number of plant / quarter observations.
- 2) We then use this new dataset to recalculate the regulatory CCPs.
- 3) Using these functions, we calculate the inputs to the moments, $m_k(\theta_j)$ and m_k^d . We limit the moments to those based on the 5,000 states which have the highest expected number of steady state observations at our estimated quasi-likelihood parameter. Note that the exact number of moments, m_k ,

varies across iterations of the bootstrapping procedure, depending on how many of those 5,000 states are in compliance.

- 4) We then calculate our initial weighting matrix and estimate our first-stage GMM structural parameters η using this weighting matrix.
- 5) We then calculate the second stage weighting matrix for the moments based on these first-stage estimates, and use this weighting matrix to re-estimate the structural parameters.
- 6) Finally, we use these estimates to calculate all of the outcomes for each counterfactual. We report the standard deviation of the outcomes across the bootstrap iterations as the standard error of our counterfactual outcomes.

We report results from 100 bootstrap draws, using the University of Arizona high performance cluster to perform the computations simultaneously. Our bootstrap for the quasi-likelihood process is similar: it uses the output created in steps 1 and 2 above. It then estimates the structural parameters with a non-linear search and performs the counterfactual computation with the new structural parameters, regulator CCPs, and dataset (analogous to step 6).

A4. Extra Figures and Tables

Table A1—: Investment and Resolution of Violations

Dependent variable: return to compliance		
Current investment	-0.115	(0.002)
One quarter lag of investment	0.380	(0.006)
Two quarters lag of investment	0.083	(0.007)
Three quarters lag of investment	-0.012	(0.005)
Four quarters lag of investment	-0.051	(0.005)
Number of observations	103,338	


Note: regression includes region, industry, and gravity state dummies. Regression uses the estimation sample restricted to plants not in compliance at the start of the period. Standard errors, which are clustered at the plant level, are in parentheses.

Table A2—: State Transitions After Investment in Compliance

Outcome: transition to regular violator status		
One quarter lag of investment	1.29	(.09)
Two quarters lag of investment	1.21	(.17)
Outcome: transition to HPV status		
One quarter lag of investment	0.48	(.12)
Two quarters lag of investment	1.11	(.17)


Note: table shows estimates from a multinomial logit regression. Regression includes region, industry, and gravity state dummies. Regression uses the estimation sample restricted to plants in compliance at the start of the period. Standard errors, which are clustered at the plant level, are in parentheses.

Figure A1. : Mean Inspection Probabilities and Fines by EPA Region



Note: authors' calculations based on estimation sample. States in each EPA region are indicated next to value.

Figure A2. : Mean Inspection Probabilities and Fines by Industrial Sector




Note: authors' calculations based on estimation sample. Industrial sector measured by 2-digit NAICS code.

Table A3—: Regressions of Regulatory Actions on Depreciated Accumulated Violations

Dependent variable:	Inspection	Fine amount	Violation
Accumulated violations with no depreciation	0.004 (0.007)	-0.014 (0.004)	-0.000 (0.001)
Accumulated violations with 10% depreciation	0.132 (0.025)	0.128 (0.016)	0.008 (0.006)
Accumulated violations with 20% depreciation	-0.031 (0.022)	-0.059 (0.013)	-0.006 (0.004)
HPV status at start of period	0.115 (0.006)	0.032 (0.002)	0.006 (0.001)
Number of observations	103,338	103,338	103,338

Note: regressions include region, industry, and gravity state dummies. Regression uses the estimation sample restricted to plants not in compliance at the start of the period. Standard errors, which are clustered at the plant level, are in parentheses.

Figure A3. : Depreciated Accumulated Violations and Monitoring and Enforcement



Note: authors' calculations based on estimation sample.

Table A4—: Percent of Observations With Gravity State by Regulatory State

Gravity	Actual damage	Potential damage	NAAQS attainment	In compliance	Regular violator	HPV
1	Low	Low	Either	37.19	36.29	38.98
2	Low	High	Either	2.89	2.44	2.08
3	High	Low	Either	4.07	4.16	3.64
4	High	High	Yes	28.22	29.34	26.58
5	High	High	No	27.63	27.77	28.72
Total:				100	100	100

Note: authors' calculations based on the estimation sample. Regulatory actions and outcomes are based on start of period regulatory status.

Table A5—: Regulatory CCPs Marginal Effects: Inspections

	In compliance	Regular violator	HPV
Plant time-varying state			
Lag investment (0 to 1)	—	0.050	0.012
2nd lag investment (0 to 1)	—	0.100	0.043
Deprec. accum. vio. (mean to mean + 1)	—	0.126	0.110
Plant fixed state			
Non-attainment (given highest gravity)	−0.028	−0.022	0.006
Highest gravity and attainment (versus lowest)	−0.000	−0.022	−0.022
SE EPA region (versus SW)	−0.101	−0.026	0.040
Utility sector (versus manuf. food)	0.107	0.193	0.134
Mean	0.086	0.272	0.428
Pseudo R^2	0.085	0.091	0.075

Note: table shows marginal effects from probit regressions. Regressions include region, industry, and gravity state dummies. We run each regression separately by start of period regulatory status (compliance, a regular violator, or HPV). Each entry reports a marginal effect as described in the table.

Table A6—: Regulatory CCPs Marginal Effects: Violations

	In compliance	Regular violator	HPV
Regulator actions			
Inspection (0 to 1)	0.021	0.063	0.085
Plant time-varying state			
Lag investment (0 to 1)	—	−0.007	−0.026
2nd lag investment (0 to 1)	—	−0.001	0.029
Deprec. accum. vio. (mean to mean + 1)	—	0.026	0.041
Plant fixed state			
Non-attainment (given highest gravity)	0.001	0.001	0.010
Highest gravity and attainment (versus lowest)	−0.000	0.006	−0.010
SE EPA region (versus SW)	−0.002	−0.010	−0.026
Utility sector (versus manuf. food)	−0.001	−0.003	−0.013
Mean	0.000	0.102	0.156
Pseudo R^2	0.182	0.152	0.099

Note: table shows marginal effects from probit regressions. Regressions include region, industry, and gravity state dummies. Most regressions also include inspection \times gravity state interactions. We run each regression separately by start of period regulatory status (compliance, a regular violator, or HPV). Each entry reports a marginal effect as described in the table.

Table A7—: Regulatory CCPs Marginal Effects: Fines

	In compliance	Regular violator	HPV
Regulator actions			
Violation (0 to 1)	0.000	0.020	0.279
Inspection (0 to 1)	0.000	0.024	0.176
Plant time-varying state			
Lag investment (0 to 1)	—	0.002	−0.592
2nd lag investment (0 to 1)	—	0.002	0.139
Deprec. accum. vio. (mean to mean + 1)	—	0.000	0.000
Plant fixed state			
Non-attainment (given highest gravity)	0.000	0.005	0.196
Highest gravity and attainment (versus lowest)	0.000	−0.001	−0.117
SE EPA region (versus SW)	0.000	−0.150	0.125
Utility sector (versus manuf. food)	0.000	−0.005	0.025
Mean	0.035	0.637	8.268
Pseudo R^2	0.187	0.245	0.108

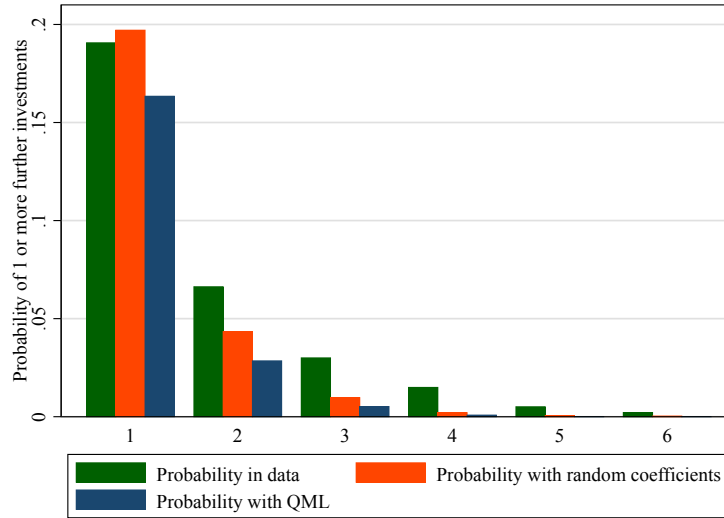
Note: table shows marginal effects from tobit regressions. Regressions include region, industry, and gravity state dummies. Most regressions also include inspection \times gravity state interactions. We run each regression separately by start of period regulatory status (compliance, a regular violator, or HPV). Each entry reports a marginal effect as described in the table.

Table A8—: Regulatory CCPs Marginal Effects: Status Transitions

Beginning State:	Compliance		Regular violator		High priority violator	
	Into regular violator	Into HPV	Into compliance	Into HPV	Into compliance	Into regular violator
Regulator actions						
Fines (mean to mean + std. dev.)	0.000	0.000	−0.048	0.001	−0.018	−0.001
Violation (0 to 1)	0.676	0.166	−0.123	0.132	−0.118	−0.017
Inspection (0 to 1)	0.006	0.004	−0.007	0.013	−0.013	−0.002
Plant time-varying state						
Lag investment (0 to 1)	—	—	0.313	−0.004	0.461	0.248
2nd lag investment (0 to 1)	—	—	0.136	0.007	−0.046	−0.008
Deprec. accum. vio. (mean to mean + 1)	—	—	0.032	0.004	−0.030	0.013
Plant fixed state						
Non-attainment (given highest gravity)	0.000	0.000	0.004	0.002	0.007	−0.005
Highest gravity and attainment (versus lowest)	−0.000	−0.000	−0.012	−0.000	0.000	−0.001
SE EPA region (versus SW)	0.002	−0.004	0.186	−0.152	−0.044	0.044
Utility sector (versus manuf. food)	−0.000	0.000	−0.011	0.011	−0.004	−0.006
Pseudo R^2	0.502		0.175		0.307	

Note: table shows marginal effects from multinomial logit regressions. Regressions include region, industry, and gravity state dummies. Most regressions also include inspection \times gravity state interactions. We run each regression separately by start of period regulatory status (compliance, a regular violator, or HPV). Each entry reports a marginal effect as described in the table.

Figure A4. : Model Fit: Further Investments in the Six Periods After Initial Investment



Note: authors' calculations based on estimation sample and estimated models evaluated at steady state.

Table A9—: Estimates of Plants' Structural Parameters: More Interactions in CCPs

	Quasi-likelihood estimates	GMM random coefficient estimates					
		(1)	(2)	(3)	(4)	(5)	(6)
Negative of investment cost ($-\theta^X$)	-2.856 (0.023)	-2.856	-2.318	-2.482	-1.906	-1.778	4.404
Inspection utility (θ^I)	-0.083 (0.028)	-0.083	-0.228	-0.130	0.106	-2.553	-2.323
Violation utility (θ^V)	0.039 (0.074)	0.039	0.260	0.767	-0.362	-1.356	-0.870
Fine utility (millions \$, θ^F)	-5.328 (0.225)	-5.328	-4.529	-6.114	-5.993	-7.055	-7.238
HPV status utility (θ^H)	-0.081 (0.007)	-0.081	-0.045	-0.094	-0.168	-2.564	0.377
Weight on parameter vector	1	0.273	0.265	0.213	0.175	0.049	0.008

Note: standard errors for quasi-likelihood estimates, which we calculate via an outer product formula, are in parentheses. GMM estimates are for a one-step estimator, unlike main results. For GMM estimates, we report the 6 parameter vectors with the highest weight. The CCPs used in these estimates include region-by-industry fixed effects instead of region and industry fixed effects.

Table A10—: Estimates of Plants' Structural Parameters for Mining and Extraction Only

	Quasi-likelihood estimates	GMM random coefficient estimates					
		(1)	(2)	(3)	(4)	(5)	(6)
Negative of investment cost ($-\theta^X$)	-2.316 (0.074)	-1.175	-2.219	-2.189	-0.964	-5.324	-8.918
Inspection utility (θ^I)	-0.129 (0.121)	-1.111	-0.993	-0.938	-0.201	2.320	-0.496
Violation utility (θ^V)	-0.218 (0.657)	-1.490	-2.481	-2.225	-1.449	-1.609	-2.616
Fine utility (millions \$, θ^F)	-5.891 (1.155)	-3.505	-6.039	-4.307	-3.728	-7.091	-8.272
HPV status utility (θ^H)	-0.058 (0.018)	-0.205	-0.074	-0.333	-0.341	-0.821	0.215
Weight on parameter vector	1	0.603	0.209	0.131	0.022	0.012	0.010

Note: standard errors for quasi-likelihood estimates, which we calculate via an outer product formula, are in parentheses. GMM estimates are for a one-step estimator, unlike main results. For GMM estimates, we report the 6 parameter vectors with the highest weight. Estimation uses only data from mining and extraction (2-digit NAICS code 21). Within this, the estimation uses the 6-digit NAICS codes with the most plant / quarters, 211111, 211112, 212312, and 212321, and EPA regions 3-8. Estimation replaces 2-digit NAICS code fixed effects in the CCPs with 6-digit NAICS code fixed effects.

Table A11—: Estimates of Plants' Structural Parameters for 10 Most Populous States

	Quasi-likelihood estimates	GMM random coefficient estimates					
		(1)	(2)	(3)	(4)	(5)	(6)
Negative of investment cost ($-\theta^X$)	-3.354 (0.036)	-2.843	-3.856	-6.458	-3.689	-0.813	-3.838
Inspection utility (θ^I)	-0.038 (0.042)	0.572	-1.195	-0.070	-0.286	-1.539	0.311
Violation utility (θ^V)	0.827 (0.076)	0.041	1.209	-0.359	0.467	3.314	-0.447
Fine utility (millions \$, θ^F)	-7.139 (0.271)	-8.967	-9.615	-8.258	-5.384	-4.934	-5.670
HPV status utility (θ^H)	-0.184 (0.009)	-0.181	-0.129	-0.155	-0.020	-2.466	-0.257
Weight on parameter vector	1	0.417	0.222	0.144	0.053	0.051	0.048

Note: standard errors for quasi-likelihood estimates, which we calculate via an outer product formula, are in parentheses. GMM estimates are for a one-step estimator, unlike main results. For GMM estimates, we report the 6 parameter vectors with the highest weight. Estimation uses only data from CA, TX, NY, FL, IL, PA, OH, MI, GA, and NC and replaces region fixed effects in the CCPs with state fixed effects.

Table A12—: Counterfactual Results With the Quasi-Likelihood Estimates:
Changing the Escalation Rate of Fines

	(1)	(2)	(3)	(4)	(5)
	Data	Baseline	Same fines for all violators; fines constant	Same fines for all violators; pollution damages constant	Fines for HPVs doubled relative to baseline
Quasi-likelihood estimates					
Compliance (%)	95.62	94.66 (0.12)	91.45 (2.84)	94.81 (0.15)	95.06 (0.12)
Regular violator (%)	2.88	3.91 (0.11)	3.78 (0.13)	3.49 (0.11)	3.91 (0.11)
HPV (%)	1.50	1.43 (0.04)	4.77 (2.91)	1.70 (0.14)	1.03 (0.03)
Investment rate (%)	0.40	0.44 (0.01)	0.43 (0.02)	0.51 (0.02)	0.45 (0.01)
Inspection rate (%)	9.65	9.43 (0.06)	10.60 (1.36)	9.52 (0.09)	9.31 (0.05)
Fines (thousands \$)	0.18	0.32 (0.04)	0.32 (0.04)	1.51 (0.29)	0.38 (0.05)
Violations (%)	0.55	0.54 (0.01)	1.08 (0.85)	0.60 (0.06)	0.50 (0.01)
Plant utility	—	−0.007 (0.004)	−0.003 (0.006)	−0.013 (0.004)	−0.008 (0.004)
Pollution damages (mil. \$)	1.65	1.54 (0.02)	1.87 (0.26)	1.54 (0.02)	1.50 (0.02)

Note: each statistic is the long-run equilibrium mean, weighting by the number of plants by region, industry, and gravity state in our data. Plant utility reports the average flow utility across types and states including ε except for Euler's constant. Column (1) presents the value of each statistic in our data. Column (2) presents the results of our model given the estimated coefficients and the existing regulatory actions and outcomes. Other columns change the state-contingent fines and HPV cost faced by plants. Columns (3) and (4) impose the same fines for all regular and high-priority violators for a given fixed state. Column (5) doubles the fines for plants in HPV status. All values are per plant / quarter. Bootstrapped standard errors are in parentheses.