1. INTRODUCTION

As part of its mandate under the Clean Water Act, the EPA collects data on pollution discharges from major water pollution sources. These data show tremendous overcompliance with the relevant regulations for the pollutants that have been the historical focus of pollution control efforts. For the time period we examine, average BOD effluent concentrations were typically around 40 percent of the permitted levels.[1] In other words, effluent from these plants is far cleaner on average than it needs to be.

This paper uses previously unexploited data to examine the nature of and reasons for this overcompliance. We use a national sample of monthly plant-level BOD concentration discharges; these monthly observations allow us to observe discharge variability, a key characteristic of water pollution that is poorly observed from annual data. They thus give a much fuller picture of polluter behavior and allow us to test hypotheses that have not previously been examined. Discharge variability has strong implications for analysis and for water pollution policy itself.

We draw three major conclusions about point-source water pollution. First, we reiterate previous findings that have shown substantial overcompliance with the major Federal regulation governing point-source water pollution, the National Permit Discharge and Elimination System (NPDES). Such over-compliance is observed for both wastewater treatment plants and for-profit manufacturing plants.

Second, plants that have higher discharge variability have lower median discharges, relative to the permitted levels. This is strong evidence for a "safety margin" effect. The implications of discharge variability and the accompanying safety margin are legion; we discuss these througout the paper.

Third, plants in richer communities have lower probabilities of having a discharge above the permitted limit. Because these types of violations are rarely observed with these plants, we develop a novel method for measuring the plant-specific probability of violation.

One important policy implication of our findings is that moves by Congress or the EPA to make regulations more stringent by reducing permitted concentrations, for example, would not necessarily lead to equal reductions in discharges; when discharges are random there is no longer a one-to-one relationship between permitted and actual discharges. Likewise, seemingly innocuous aspects of the regulations, such as the time period over which

[1]Biochemical oxygen demand, BOD, is a measure of the amount or concentration of wastewater components that lead to depletion of oxygen in the receiving water body, with predictable effects on living organisms. We thank a referee for this succinct description.

discharges are averaged, are likely to have substantial effects on pollution levels. These and other policy implications are discussed in the paper's final section.

We begin with a brief example that shows the nature of BOD discharges, and then lay out common explanations for overcompliance. The main results are in Sections 5-7. In Section 8.2 we discuss policy implications of our findings.

1.1 EXAMPLE

By way of introduction, we first show a graph of typical pollution behavior and sketch out the research questions suggested by such behavior. Figure 1 shows monthly average BOD concentrations in wastewater from the Annapolis Water Reclamation Facility, a municipal watewaster treatment plant, for 1991-1999. In this figure, concentrations are depicted relative to their permitted value; that is, actual monthly discharges are divided by the permitted level. We call this the *discharge rate*. This plant's discharge rate tends to be around 0.3, far below the benchmark of 1.0.

Figure 1. Annapolis (Maryland) Water Reclamation Facility

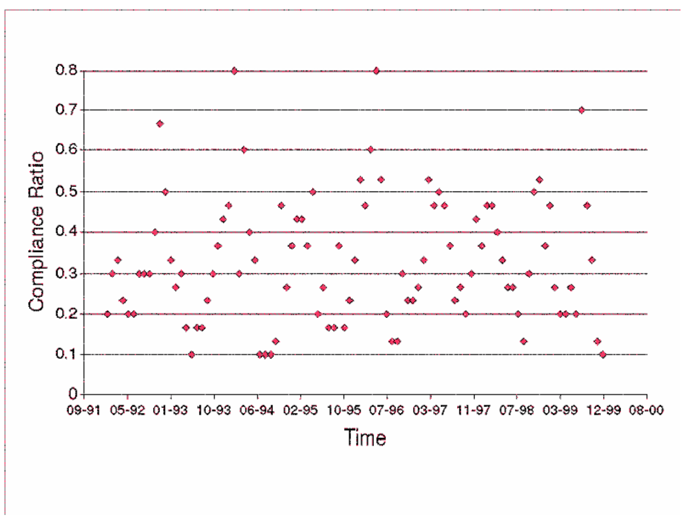

The pattern shown in Figure 1 is a typical one, common in all regions and industries. In addition to a low average discharge rate, the figure shows a high degree of month-to-month variability, variability that would be masked by any annual averaging. This variability is inherent in the control of BOD and is due to weather, human error, and mechanical breakdown as well as a basic randomness in the processes used to treat organic waste. The figure also shows no numeric violations (discharge rates above 1.0), despite the great variability of the discharges.

1.2 RESEARCH PLAN

Regulators and plant managers, when queried about relatively low average discharge rates such as those exhibited in Figure 1, claim that such behavior is warranted by the discharge variability. Plants are posited to pollute below their permitted level, on average, to provide a safety margin in the case of an unexpectedly large discharge. We call this the *safety margin* explanation of overcompliance. Plants reduce their average discharge so that they are more likely to fall below a discharge rate of 1.0 during a very bad month.

This explanation is compelling but has not been tested before to our knowledge. Thus, the intent of the first half of this paper is to devise a test of – and then test – the safety margin explanation. We find evidence consistent with the safety margin explanation.

Note that this role for discharge variability makes it more difficult to assess plants' compliance behavior because average discharges are no longer a sufficient statistic to describe how well a plant is doing. Previous results on polluter behavior, based on average discharges, may therefore be called into question.

To remedy this, we construct a plant-specific implied probability of violation that incorporates both average (median) discharges and their variability. Actual violations are not a good measure of behavior because of the large number of zeroes.

Using this probability of violation, we then re-examine the role of community characteristics. Some anecdotal evidence suggests that community factors may be important. Plant managers frequently mention the public pressure they face for clean discharges. Kagan *et al*. (2003), for example, quote a pulp and paper plant manager as saying, "We spent a lot of money to achieve this [reduction in complaints.] The driver is to pacify the community." We find that community characteristics have a strong, though noisy, effect on polluter behavior.

With our more accurate measure of polluter behavior, we can also revisit the question of "overcompliance." We find evidence suggesting that plants are indeed overcomplying with this regulation, even accounting for the safety margin effect.

2. LITERATURE

No other studies have explicitly modeled discharge variability to our knowledge. In the paper closest to ours, Brännlund and Lofgren (1996) demonstrate that water pollution regulations were binding for Swedish pulp and paper plants even though average discharges were below the permitted levels. They demonstrate this by showing that (annual) average discharges were affected by the level of the regulation. Their analysis is motivated by an appeal to the randomness of discharges. However, they do not measure the randomness and therefore are limited in the questions they can address.

A recent paper by Shimshack and Ward (2005), analyzing monthly pulp and paper plant discharges, recognizes and attempts to correct for discharge variability. They calculate a plant-specific risk-of-violation using regression residuals. This measure essentially proxies for plant variability by looking at the plant's "adjusted" mean; the mean is presumably adjusted for the safety-margin reason. The risk of violation is a right-hand side variable in their regressions rather than being the dependent variable.

A more common approach to studying water pollution has been to examine factors that influence plant-level average discharges, primarily inspections, enforcement, and related regulatory variables.[2] Except for Earnhart (2004a, 2004b), which uses monthly data, these papers use either quarterly or annual data, and none measures or explicitly accounts for plant-level variability. Note that their results, based on average discharges, must be treated with care since average discharges are not an accurate measure of plant behavior when discharge variability differs across plants. Plant fixed effects may mitigate this problem.

Variability plays a role in other regulated behavior besides water pollution. One example is oil spills from tankers. Anderson and Talley (1995), for example, report aggregate mean and standard deviation of oil spilled but do not measure their joint effect on the constructed spill propensity index.

We also examine the effect of community characteristics on polluter behavior. This literature is vast. Most of it, however, focuses on non-regulated behavior such as legal toxic releases, pollution control expenditures, or participation in voluntary programs.

Our results are more directly comparable to studies that focus on regulated behavior. Two studies are especially worth noting. Gray and Shadbegian (2004) undertake a detailed study of air and water pollution at U.S. pulp and paper mills. They find that plants in poor neighborhods emit more BOD quantity, a finding

---

[2]See Earnhart (2004a, 2004b); Foulon, Lanoie and Laplante (2002); Helland (1998); Laplante and Rilstone (1996); Magat and Viscusi (1990); McClelland and Horowitz (1999); Niemann (2001).

similar to ours. The data and methods are too different from ours to compare quantitative esimates. Earnhart (2004b), in a study of wastewater treatment plant discharges in Kansas, makes a strong case for community effects. He finds that community demographic variables have substantial effects on local plants' discharge rates. His study does not measure discharge variability but does include plant fixed effects.

Recent studies by Kagan *et al*. (2003) and Gunningham *et al*. (2004) look at a variety of factors to explain pollution by 14 pulp and paper plants in multiple countries. They consider management variables, "regulatory style," and community pressure. Because they have a small data set, they cannot statistically test these explanations together. They also have inadequate data to measure community pressure, although they do have a great many quotes by plant officials describing the influence that communities have on local polluters.

3. DATA

We examine all dischargers that reported to the NPDES from 1991 to 1999 and that are considered "major" sources by EPA. Most of the data are for sewage treatment plants. Discharges are measured for each outlet pipe, but most plants have just one or two outlets; we take the average over outlets for each plant with more than one outlet. Permits specify only the allowable discharges; they do not prescribe technologies that plants must adopt to comply with those permits beyond secondary treatment. Penalties for discharges above the permitted amount are approximately "smooth" in the violation: small, infrequent violations receive small penalties on average; large or frequent violations receive larger penalties. EPA officials have told us that plants are expected to comply with the NPDES limits 100 percent of the time.

Our first decision is which pollutant to examine. We chose to examine monthly average BOD concentrations. We chose this measure over other possible pollution measures for four reasons. (1) It has the greatest number of observations in our data. (2) It has the highest violation rate among the four major conventional water pollutants.[3] (3) Concentration limits are much less susceptible than quantity limits to questions about how the permit levels are set. (4) EPA officials have told us that BOD concentration is a major focus of the agency. Earnhart (2004a) provides additional justification for focusing on BOD concentration for sewage treatment plants.

---

[3]The four measures are BOD and total suspended solids measured as quantities or concentrations. Note that BOD concentration is only marginally superior on criteria (1) and (2).

Each plant faces a limit, usually 30 milligrams per litre, on its monthly average BOD concentration. For plant i in month t, define the *discharge rate* as $C_{it} = Z_{it}/R_{it}$, where Z_{it} is reported discharges and R_{it} is the permitted level. We drop subscripts when no confusion is likely. Whenever $C > 1$, the plant is in violation. There is relatively little cross-sectional or time-series variation in R_{it}, so changes in C_{it} are mostly due to changes in Z_{it}.

The raw data have a number of missing and miscoded observations. Some plants are required to report only average concentrations over a quarter. We include these plants but treat their observations as single (*i.e.*, monthly) points.[4] We discuss further missing-data problems in Sections 4.2 and 5.1. We drop observations that appear to be improperly entered, such as when discharges are either negative or unrealistically large; these most likely are improperly entered.[5]

Our analysis also includes sociodemographic data for the zipcode in which a plant is located. The Summary Tape File 3B of the 1990 U.S. Census provides aggregate demographic data for every residential zipcode. We use six indicators: (a) percent of non-white residents, (b) median household income, in thousands of dollars, (c) percent of workers who travel to work by carpool, (d) percent of workers in manufacturing, (e) total population in the zipcode, in thousands, and (f) percent of residents living in urban areas. All variables are for 1990.

4. MEASUREMENT OF VARIABILITY

Discharges exhibit considerable variability on a month-to-month basis, as Figure 1 shows. Variability arises due to natural variability in the composition of complex organic wastes (*e.g.*, Leduc, Unny, and McBean 1988)), environmental factors, and operational factors (Berthouex and Fan 1986). Weather creates both a predictable seasonal variation in discharges and a large random component. Seasonal effects are the main source of "predictable" variation; we discuss them in Section 5.2.

In order to study variability, it is necessary to: (i) derive an appropriate empirical measure; and (ii) understand how EPA has accounted for variability in permit design and enforcement. We take up these tasks next in 4.1 and 4.3.

---

[4]It is not always clear whether plants that report quarterly are reporting quarterly averages or periodic monthly averages or even simply infrequently reporting monthly averages. To the extent that some plants may be reporting quarterly averages (which will show less variability than monthly averages) and are doing so because they are good performers, then the coefficients in Table 3a and 3b will show *less* response to variability than is actually occurring.

[5]We discussed these deletions with EPA officials.

4.1 IDENTIFYING UNCONTROLLABLE VARIABILITY

At some plants, some variations in discharges are due to deliberate changes in plant design or operation. Structural changes are economically important, but the resulting changes in discharges should not be included in any measure of true discharge variability. Some of the variability observed in Figure 2, below, for example, does not represent variability for which the plant might want to compensate; indeed, it represents variability *resulting from* plant behavior. Any measure of discharge variability relevant for the safety margin explanation must exclude changes in discharges due to structural changes at the plant.

Figure 2. Mount Laurel (New Jersey) Sewage Treatment Plant

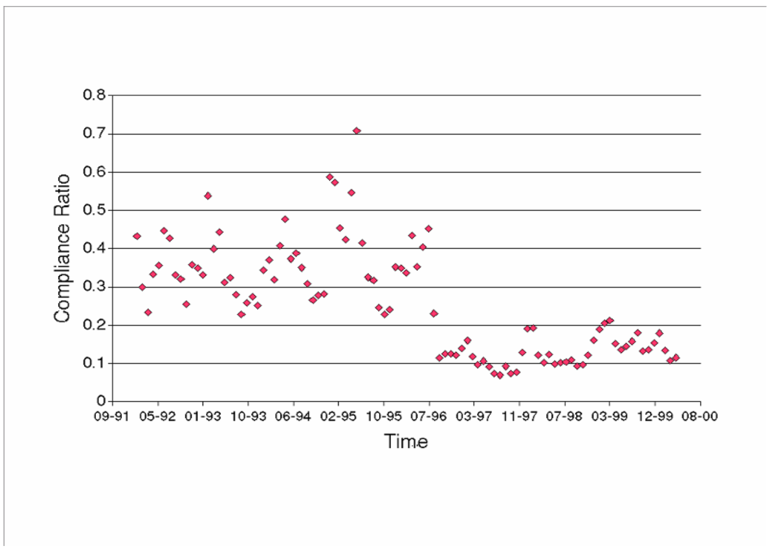

Thus, to test for the safety margin hypothesis, we first work to remove plant-level discharge variability due to structural change. A systematic qualitative method for identifying when plants undertook these "deliberate" changes cannot be reliably derived because in many cases the specific plant changes are subtle and plant-specific. Instead, we use a CUSUM test for removing such components from the data.[6]

---

[6]We also tried a method for removing structural variability based on plants' permit cycles. Because EPA requires new permits to be obtained after major changes in plant design or capacity, data within a single permit cycle would likely exclude the structural changes whose influence we

We use a CUSUM test to identify plants whose discharges underwent a significant structural change (Brown *et al*. 1975) (BDE). Suppose discharge rates follow the model:

$$\ln C_{it} = \alpha_i + \varepsilon_{it} \tag{1}$$

where α_i and ε_{it} are plant-specific mean and error terms. We want to test for a structural change in α_i at some time, say t_0. Following BDE, we construct the difference (i.e., the residual) at each time t between actual and predicted compliance based on estimation of (1) up to t. Under the null hypothesis of structural change at t_0, the expected value of this residual is zero until t_0 and non-zero thereafter. BDE develop a statistical test based on the series of weighted residuals. We use this CUSUM test at a 5 percent level of significance to identify plants that underwent a structural change in α_i during the sample period. We equate a structural change in the econometric sense with a "deliberate" change in the economic sense.

4.2 Data used for analysis

We identify 4,330 plants in SICs 2, 3, and 4952 (sewerage systems). In deriving the plant-level variability measure, we focus on plants that have at least 15 months of data, not necessarily consecutive; there are 2,709 such plants.[7] Plant size, as measured by the "design flow" is not available for 272 plants, thus reducing our sample to 2,437.

Of these, 1,096 plants do not fail the CUSUM test and thus had no structural changes during the sample period. These plants consitute our main sample. Summary statistics are shown in Table 1.

This sample is considerably reduced from the original sample, so we want to be clear about its validity. Our results apply to the universe of plants that have not undergone a structural change during this time period. This restrictiveness in our sample selection is necessary because it is important that our variability measure be uncorrupted by structural changes. Measurement of true uncontrollable variability (conditional on inclusion in our sample) is essential for our test of the safety margin explanation to be meaningful.

---

are trying to eliminate. This method yields fewer usable observations and is problematic because of unreliable reporting of permit renewals, so we do not pursue it.

[7] This criterion was chosen, somewhat arbitrarily, to balance greater precision in the plant-level measure of discharge variance against a larger sample of plants.

Table 1. Summary statistics: Ratio of reported discharges to permitted discharges, by region, SIC code, and month

	# Plants	Ave. obs./plant	Mean	Median	Max	σ^a
EPA Region:						
1	88	86	0.38	0.31	8.89	0.5
2	124	76	0.42	0.37	8.87	0.43
3	77	70	0.32	0.23	46.9	0.47
4	306	71	0.4	0.31	61.3	0.47
5	77	68	0.39	0.3	7.8	0.49
6	188	67	0.47	0.38	13.13	0.41
7	78	59	0.45	0.4	8.27	0.46
8	42	77	0.34	0.27	5.5	0.46
9	61	66	0.34	0.24	22.93	0.48
10	55	79	0.41	0.35	7.87	0.46
SIC 2	118	63	0.37	0.28	33.07	0.56
SIC 28	46	63	0.32	0.24	11.19	0.57
SIC 29	14	56	0.32	0.25	8.27	0.6
SIC 3	29	71	0.32	0.24	8.87	0.61
SIC 4952	949	72	0.41	0.33	61.3	0.44
January	1095	6	0.43	0.34	24.14	
February	1094	6	0.44	0.35	8.87	
March	1095	6	0.43	0.35	12.75	
April	1092	6	0.44	0.35	61.3	
May	1088	6	0.42	0.34	17.2	
June	1087	6	0.39	0.33	11.93	
July	1087	6	0.38	0.3	33.07	
August	1086	6	0.38	0.3	22.93	
September	1087	6	0.37	0.3	11.58	
October	1088	6	0.38	0.3	46.9	
November	1094	6	0.38	0.3	11.67	
December	1095	6	0.4	0.31	13.13	
Full sample (plants with 15 or more observations)	1096	71	0.4	0.32	61.3	0.46
Plants with 30 or more observations	947	79	0.4	0.32	61.3	0.46
Plants with 50 or more observations	794	87	0.4	0.32	61.3	0.47

[a]Standard deviation of logged discharge rate, per plant, weighted average over plants.

We considered trying to develop a broader model that encompasses structural change, which would then allow us to include plants that fail the CUSUM test. This task is complex and essentially would have resulted in our putting the cart before the horse: A model of endogenous structural change requires an understanding of the role of variability in plant behavior, since variability may be a key reason why plants undertake investment or hire new managers. We feel that the current paper's analysis is necessary before tackling such issues.

We also considered including data for the structural-change plants from either before or after the structural break. This procedure was problematic for several reasons, including small numbers of observations per plant and difficulty in identifying the exact "breakpoint" date. It is worth noting that about half of the structural-change plants actually have a deterioration in their performance and many omitted plants appeared to have increased variability over time. The causes and consequences of structural change clearly merit further analysis. But they cannot be tackled without first understanding how variability affects plants that have not undergone a structural change.

The other substantial source of attrition in our sample is missing data. EPA allows plants not to report discharges if water pollution is not a serious problem in that area or the plant is known to be a low polluter. In some situations, EPA does not set a limit although it may still require that discharges be reported. We discuss the implications of this rule in Section 5.1. We suspect, however, that most of these data are missing due simply to random failure to report.[8]

We can provide some corroborating evidence. We looked at the frequency of numeric violations; that is, the percentage of observations for which $C > 1$. Violation rates are 3.11 percent for our sample (n = 54,639) and 3.95 percent for all data (*i.e.*, including those with fewer than 15 observations per plant, n = 332,709); n is the number of plants multiplied by the number of observations at each plant. These violation frequencies seem sufficiently similar. As with the structural change problem, an analysis of missing observations is clearly warranted. But the complexity of the problem (lack of instrumental variables; lack of a structural model) and the *ex ante* evidence lead us to leave this for future research.

Of these 1,096 plants, 764 had zipcode level sociodemographic data. These plants constitute the sample we use for those regressions with sociodemographic data.

For each plant in our final sample, we calculate the logs of the monthly discharge rates. We then calculate the plant-specific median and standard

---

[8]For example, "[k]nowledgeable agency sources said that the Clean Water Act compliance statistics, gathered by individual states from plants' and utilities' mandatory reports, were reliable." (*Washington Post* 2003). The NPDES data have been the subject of much criticism, but primarily because of their limited coverage. Note that these claims about NPDES problems are predicated on the data being an accurate representation of what they are supposed to measure.

deviation of this series, denoted m_i and σ_i. The standard deviation of the log of monthly discharge rates is our measure of discharge variability. We drop the subscript where no confusion is possible and refer to the plant-specific median and standard deviation as m and σ.

Let Z_m be plant-level median discharges and let $z_m = \ln(Z_m)$ and $r = \ln(R)$. Since the median of the logs is the log of the median, when there is no time series variation in R we have $m = z_m - r$.

4.3 EPA POLICY

Plants face regulations that govern both monthly average and daily maximum concentrations. The monthly average is the better measure to focus on for three related reasons: (1) EPA guidelines specify that plants should design their treatment systems to be in compliance with the monthly permits (*i.e.*, not daily). (2) EPA data on daily maximum concentrations are much sparser than monthly averages. The EPA's focus on the reporting of monthly averages makes sense given reason 1 and the Agency's emphasis on long term performance, discussed below. (3) The relationship between daily and monthly limits is designed so that a plant that is in compliance at the monthly level will typically be in compliance at the daily level. (Daily limits are higher than the monthly limits.)

EPA also accommodates discharge variability by focusing its enforcement efforts on long term violators (GAO 1996). Long term averages exhibit considerably less randomness.[9]

Research suggests that discharges are distributed lognormal (EPA 1996, Niku *et al*. 1979). Therefore, our analysis examines the log of the discharge rate, $c_{it} = \ln(C_{it})$.

5. RESULTS I: HOW DOES DISCHARGE VARIABILITY AFFECT AVERAGE DISCHARGES?

5.1 HOW DOES DISCHARGE VARIABILITY AFFECT AVERAGE DISCHARGES?

The claim that discharges are low on average to compensate for discharge variability implies that discharges will be lower, relative to the permit level, for

---

[9]The agency's focus on long-term violation seems counter to the plants' apparent focus on monthly variability. One possibility is that plants could not "bank on" EPA's long-term focus. We thank Wally Thurman for this observation.

plants with more variable discharges. A plant with highly variable discharges should aim for lower average discharges than a plant with low discharge variability.[10] This is a straightforward hypothesis, but its import has not been recognized to our knowledge.

To test this hypothesis, we run a regression of median log compliance, m, on the standard deviation of log compliance, σ, and test whether the coefficient on σ was negative.[11,12] Separate regressions are run for the three most common permit levels (20, 30, and 45 mg/l).

5.1.1 VARIABLE DEFINITIONS

The prediction about this relationship is based on a plant's "intended" average pollution level, that is, its true mean, and its "true variability," as represented by the variance. Because only imperfect measures of these variables are available, we must consider the possibility of measurement error in both the dependent and independent variables. Measurement error is particularly troublesome when these errors are correlated. Our choices are guided by a desire to minimize this correlation.

The true mean may be captured by two possible measures: sample mean and sample median. We use the sample median as our dependent variable rather than the mean because the measurement error in the median has a much smaller correlation with measurement error in σ. To see this, consider a plant with a small number of monthly discharge observations, one of which is very large. Such a large data point causes both the sample mean and the sample standard deviation to be large. The sample median, however, will be unaffected.

---

[10]The claim that higher variability will lead polluters to reduce average pollution is intuitive but not an iron-clad theoretical result. It can fail as a theoretical prediction if the probability of violation is high, a condition that does not hold for our data.

[11]Although this discussion treats σ as exogenous, plants do have some choice over their discharge variability. Because of the widespread belief that discharge uncontrollability is the underlying factor behind low average discharges, we felt compelled to deal explicitly with the implied hypothesis, *i.e.*, to estimate the m-σ relationship.

The key result from Table 3 is to demonstrate the inverse relationship between the median and variance, which we label a safety-margin. This is the crucial finding that implies that any study of polluter behavior must take discharge variability into account. The results in Tables 5 and 6 focus on the probability of a violation and are unaffected by endogeneity of σ.

[12]A more substantive concern is whether m-σ represents a technical rather than a behavioral relationship, a long-standing identification problem in economics. This does not appear to us to be the case; for example, there is too much variation in the implied probability of violation. We thank a referee for pointing out this issue.

For σ, the expected measurement error is smaller for plants with a larger number of monthly observations. Thus, in our regressions we weight observations by the number of months of data used in constructing σ.

5.1.2 MODELED PROBABILITY OF VIOLATION

To examine how well these variables capture plant behavior, we used m and σ to predict, for each plant, the probability of a violation, $\Pr(C > 1 \mid m_i, \sigma_i)$, under the assumption that discharges are distributed lognormal relative to their permitted level.[13] We also counted for each plant the actual proportion of observations where discharges exceeded the permit, which we refer to as numeric violations. For example, if a plant had 25 observations and 2 numeric violations, then its *observed violation rate* is 8 percent.

Table 2 summarizes the predicted (*i.e.*, modeled) and observed violation rates for our sample. This comparison is a check on the accuracy of m and σ as measures of the plants' discharge behavior and on the assumption of log-normality.

Our measure does a remarkably good job in capturing violation rates. We calculate that on average plants will violate 3.57 percent of the time; the actual figure is 4.71 percent. Percentile calculations are also close. We calculate that the median violator will violate 0.8 percent of the time; the actual figure is 0.

Table 2. Observed and modeled violation rates ($\Pr(C > 1 \mid m_i, \sigma_i)$, in percent)

Percentile:	10^{th}	25^{th}	50^{th}	75^{th}	90^{th}	95^{th}	Highest	Mean
Observed	0	0	0	3.9	10.0	16.4	79.4	3.57
Modeled	≈0	≈0	0.8	5.2	14.6	21.8	76.3	4.71

[13]If discharges are distributed log normally then only average discharges where the average was constructed using geometric means should be distributed log normally. Unfortunately, most plants use arithmetic means. Since our data appear to be distributed close enough to log normal, we are not particularly concerned about this issue. We hope it is clear that *we* did not construct the monthly average concentrations; the plants did.

5.1.3 HYPOTHESIS TEST

We use data for SIC=4952. We run separate regressions for each permit level; 636 plants had one of three common permit levels. Results are in Table 3a.

Table 3a. Relationship between median compliance and standard deviation, SIC=4952

Dependent variable: Median log compliance

Permitted concentration:	R=20	R=30	R=45
Constant	-1.20**	-0.52**	-0.73**
	(6.19)	(4.78)	(3.68)
σ	**-0.05**	**-1.53****	**-1.10****
	(0.16)	**(7.30)**	**(2.87)**
ln(VOLUME)	-0.04	0.03	0.01
	(0.69)	(1.28)	(0.13)
R^2	≈0	0.13	0.07
N	106	404	126

Numbers in parentheses are *t*-ratios. **Significantly different from zero at the 1% confidence level.

The coefficients on σ are negative, as implied by the safety-margin hypothesis. An increase in variance is associated with a decrease in median discharges. This is an intuitively appealing result that has been widely believed to hold, but it has never been measured or even, as far as we can tell, empirically confirmed.

The magnitude of this relationship can be gauged by looking at the predicted discharge rate for average-size plants facing a 30 mg/l standard, the most common standard. When σ = 0.05 (the lowest observed variability), the predicted discharge rate is 0.57; when σ = 0.46 (the mean value), the predicted discharge rate is 0.30; and when σ = 1.2 (highest observed variability), the predicted discharge rate is 0.10. Thus, variability has a large effect on median discharges.

The m-σ tradeoff facing any individual plant may be different from the measured coefficient for two reasons. First, our samples include plants with "old" and plants with "new" technologies, including, of course, a large number of plants whose technology cannot be so classified. Newer technologies tend to have both lower discharges and lower variability. Therefore, the *m*-σ tradeoff facing any individual plant would be *larger* in magnitude that the observed *m*-σ relationship over all plants.

A second reason to be cautious about the measured coefficient is the possibility that plants may fail to report a high discharge rate or may falsely report $C < 1$. These would tend to be high-C plants, and their self-censoring (which is illegal) would dampen the measured variability. In this case, the observed *m*-σ relationship would be overestimated. We mention this point because readers are often bothered by its possibility. Misreporting has substantial penalties, and EPA does not appear to believe it to be a widespread problem. See footnote 8.

5.1.4 THE EFFECT OF REGULATORY STRINGENCY

We also examine how the *m*-σ relationship varies with the stringency of the regulation. The regulation is most likely to be binding when low effluent concentrations are allowed (20 mg/l) and least likely to be binding when high concentrations are allowed (45 mg/l). "More binding" regulation is expected to increase the magnitude of the effect of variability on median compliance. Therefore, the σ coefficients should be highest in magnitude at $R = 20$.

Table 3a shows a mixed pattern. At $R = 45$, the absolute value of the σ coefficient is smaller than for $R = 30$, as hypothesized, but at $R = 30$, the absolute value of the coefficient is larger than for $R = 20$, which contradicts the hypothesis.[14]

5.1.5 OVERCOMPLIANCE?

Because the variability effect is large, one cannot assess a plant's compliance behavior based only on average discharges. Our results, however, allow us to assess the degree to which plants are indeed overcomplying, if at all. To do this, we predict the discharge rate for a plant with zero discharge variability, using the results from Table 3a. We predict that even with no variability, plants pollute at just 60 percent of their permitted levels. This is rather substantial

[14]It is possible that this result reflects our sample selection criterion that plants not have undergone a structural change. The candidate explanation is that plants with R = 20 are the most likely to have undertaken structural change; therefore, those who did not undertake a structural change and who are the basis for the estimates in Table 3a will be those who, for some unobserved reason, find the regulation less binding. We thank Don Fullerton for this insight. This candidate explanation must be treated with some care, however, since our preliminary assessment of plants undertaking structural change showed many plants that did not appear to be improving their performance.

overcompliance. That is, plants are overcomplying by about 40 percentage points *even when the compensation for variability is accounted for*.

Another more robust, though less conclusive, demonstration of this "variability-compensated" overcompliance comes from the low implied probabilities of violation, shown in Table 2.

This overcompliance result is consistent with general findings by Kagan *et al*. (2003) and Gunningham *et al*. (2004) who argue that pulp and paper plants are "intentionally" overcomplying, based on discussions with plant managers. Our results give a numerical estimate of this effect. Note, we do not claim that *all* plants are overcomplying. A small group continues to have poor compliance records.

5.2 PEAK "BAD SEASON" RANDOMNESS

Variability arises from both seasonal patterns and a residual randomness. The policy implications of these two sources are somewhat different. Therefore in this section we look specifically at the role of weather randomness. Such randomness is most important in the peak "bad season," which occurs in the winter and early spring.

To examine whether plants are targeting their median discharges to the "bad season," we construct for each plant the bad-season median, m_b, and variance, σ_b, using data just from January-April; note that the number of monthly observations used to construct these are much smaller than in the previous section, of course. We use the same sample as for Table 3a. Regressions are re-weighted for the number of observations used to construct m_b and σ_b. We again expect a negative relationship between the median and the variance.

Results are shown in Table 3b. They again show a large negative relationship between m_b and σ_b. As before, the magnitude can be gauged by looking at the predicted discharge rate for average-size plants facing a 30 mg/l standard. When $\sigma = 0.41$ (the mean value), the predicted discharge rate is 0.33; when $\sigma = 1.5$ (the highest value), the predicted discharge rate is 0.11. Thus, this type of variability indeed has a large effect on discharge rates. For $\sigma = 0$ (certainty), the predicted discharge rate is 0.51. This provides a further measure of the "true" overcompliance; it is roughly the same as that found in Section 5.1.

These regressions further show that (i) even plants with very low variability tend to keep discharges far below the permitted level (the overcompliance phenomenon), and (ii) randomness (as opposed to predictable variability) is an important factor in plant behavior.

Table 3b. Relationship between median compliance and standard deviation during the "bad season," SIC=4952

Dependent variable: Median log compliance (January-April)

Permitted concentration:	R=20	R=30	R=45
Constant	-1.08**	-0.75**	-0.78**
	(-5.80)	(-7.48)	(-4.83)
σ_b	**-0.19**	**-1.02****	**-0.92****
	(-0.52)	**(-4.75)**	**(-2.70)**
ln(VOLUME)	-0.004	0.07**	-0.001
	(-0.05)	(2.59)	(-0.03)
R^2	≈0	0.08	0.06
N	106	404	126

Numbers in parentheses are *t*-ratios. **Significantly different from zero at the 1% confidence level.

6. RESULTS II: ANALYSIS OF THE IMPLIED PROBABILITY OF A VIOLATION

6.1 METHODOLOGY, HYPOTHESES, AND RESULTS

When discharges are variable and that variability differs across plants, it is necessary to focus on the probability distribution of discharges rather than simply their mean. That is, any analysis of discharges must also account for discharge variability. In this section we examine the factors that influence the implied probability that discharges will exceed their permitted levels. These factors include plant, community, and regulator characteristics. Previous research has focused primarily on the effect of inspections; we broaden that outlook here, although at the expense of being able to tease out the specific effect of a plant visit by an inspector.[15]

[15]Inspections have been excluded from our research for two reasons. First, to measure σ we have to adopt a cross-sectional rather than time-series approach. Inspections require a time series approach, however; for example, does an inspection lead to reduced discharges? A longer panel data set might allow us to answer such a question since we could presumably allow σ to vary over time. Second, the observed low levels of discharges led us to focus first on factors that have been associated with overcompliance.

Thus, we use as our dependent variable $y = m/\sigma$, called the standardized discharge rate, which is median (log) discharge rate "corrected" for variability. When discharges are distributed lognormal, then y is a sufficient statistic for the probability of a violation, $\Pr(C > 1 | m_i, \sigma_i) = \Pr(C > 1 | y_i, 1)$. The advantage of y as the dependent variable is that it is expected to be distributed normal.

We regress y on: (1) dummies for EPA-regulated vs. state-regulated plants, including EPA region; (2) a dummy for manufacturing *vs*. municipal wastewater treatment plants; (iii) size of plant; and (iv) zipcode level economic and demographic variables. Summary statistics of all variables are in Table 4.

Table 4. Summary statistics for regression variables (n=764)

	Mean	Std. Dev.	Min.	Max.
m	-1.16	0.63	-4.38	0.48
y	2.91	2.51	-0.72	35.46
r	3.21	0.59	0.92	6.03
σ	0.46	0.18	0.045	1.24
ln(VOLUME)	0.98	1.41	-4.94	6.70
STATE	0.82	0.39	0	1
MANUF	0.12	0.32	0	1
Non White	14.86	17.86	0	94.27
Income	28.80	10.94	7.76	82.45
Carpool	14.36	4.54	2.72	33.73
Manufacturing	19.70	10.18	1.61	60.34
Population	18.13	13.57	0.34	83.58
Urban	33.52	44.36	0	100

Plants are regulated either by the EPA or by state regulators. State regulation occurs in states that have applied to the EPA for enforcement power and have demonstrated that they have sufficient enforcement capability. STATE is a dummy variable equal to 1 when the plant is regulated by the state and 0 when it is regulated by the EPA.

We also examine whether manufacturing plants have higher or lower probabilities of violation than sewage treatment plants. MANUF is a dummy variable equal to 1 if the plant is a manufacturing plant (SIC 2 or 3) and 0 if it is a sewage treatment plant (SIC 4952). It is not clear whether sewage treatment plants, which are not profit-maximizers, should be expected to have higher or lower probability of violation than manufacturers. Profit-maximizing may make it less likely that plants will undertake "non-economic" overcompliance. On the other hand, municipal plants may be less nimble in responding to the forces that are leading plants to reduce discharges.

VOLUME is the plant's designed water volume and is the appropriate measure of size for water pollution point-sources. We examine whether larger plants are more likely to be in compliance. Two opposing forces affect the relationship between this variable and the probability of violation. Larger plants are more prominent polluters and can potentially cause more damage than smaller plants, so they may face higher scrutiny from regulators or the community. BOD treatment may also exhibit economies of scale. Both of these factors would lead larger plants to have lower violation probabilities. On the other hand, a large plant may employ a significant part of the workforce in a community and may be more successful in resisting community pressure; this seems particularly likely for manufacturing plants. These factors would lead larger plants to have higher violation probabilities on balance. Earnhart (2004a) finds that larger plants have higher discharge rates.

The sociodemographic variables are median household income, proportion of residents that are non-white, proportion of workers who carpool, proportion of workers employed in the manufacturing sector, population, and proportion of the population that is urban. All variables are at the zipcode level.

Results are in Table 5. We find that: (1) state-regulated plants, in general, have lower probabilities of violation than EPA-regulated plants, although the estimated effects are not large. (2) Manufacturing plants have a lower violation probability than wastewater treatment plants. (3) Larger plants have lower probabilities of violation. (4) Community variables, treated as a composite, have large effects on the probability of violation. These results are conditional on the plant being in our sample; that is, on not having undertaken a significant plant change during 1991-1999.

6.2 REGULATOR EFFECTS

The regressions in Table 5 also show the effects of different enforcement regimes. Regression #1 requires all state regulators to have the same regulatory regime, but this may be different from the EPA regime. Regression #2 allows state-regulation to differ across EPA regions. In two regions, 3 (mid-Atlantic) and 6 (South-Central), state-regulators are found to be significantly more stringent than EPA; in all other regions, state regulation is not significantly different from EPA regulation. Note that almost all state-region coefficients (γ_{11} through γ_{20}) are negative or very close to zero, which implies that in those situations where states have received permission to oversee regulation, the states have fulfilled their duties no better or worse than EPA.

Table 5a. Standardized discharge rate[a] regressed on plant, regulator, and community variables (n = 764)

	#1	#2	#3
γ_1 (Intercept)	-3.18**	-3.30**	-3.41**
	(2.80)	(2.84)	(3.03)
γ_2 (*STATE*)	-0.33	--	--
	(1.42)		
γ_3 (*MANUF*)	-0.69**	-0.72**	-0.69**
	(2.16)	(2.23)	(2.13)
γ_4 (ln(*VOLUME*))	-0.20**	-0.21**	-0.19**
	(2.33)	(2.45)	(2.18)
γ_5 (Non White)	-0.019	-0.0016	-0.018
	(1.21)	(0.09)	(1.09)
δ_5 (Non-White)^2	3.2×10^{-4}	1.2×10^{-4}	2.9×10^{-4}
	(1.48)	(0.52)	(1.34)
γ_6 (Income)	-0.060	-0.059	-0.062
	(1.46)	(1.40)	(1.51)
δ_6 (Income)^2	0.00079	8.1×10^{-4}	8.1×10^{-4}
	(1.50)	(1.49)	(1.53)
γ_7 (Carpool)	0.11	0.13	0.11
	(1.04)	(1.19)	(1.00)
δ_7 (Carpool)^2	-0.0031	-0.0034	-0.0029
	(0.97)	(1.07)	(0.91)
γ_8 (Manufacturing)	0.068**	0.061*	0.066**
	(2.15)	(1.88)	(2.04)
δ_8 (Manufacturing)^2	-9.4×10^{-4}	-9.2×10^{-4}	-9.0×10^{-4}
	(1.48)	(1.41)	(1.40)
γ_9 (Population)	0.0034	-0.0016	0.0018
	(0.16)	(0.08)	(0.09)
δ_9 (Population)^2	1.3×10^{-4}	2.2×10^{-4}	1.6×10^{-4}
	(0.38)	(0.60)	(0.44)
γ_{10} (Urban)	-0.0035	0.0078	0.0033
	(0.23)	(0.51)	(0.22)
δ_{10} (Urban)^2	-4.4×10^{-5}	-8.6×10^{-5}	-4.0×10^{-5}
	(0.28)	(0.53)	(0.25)
γ_{11} (State regulated, Region 1)	--	-0.48	--
		(0.22)	
γ_{12} (State regulated, Region 2)	--	-0.23	--
		(0.73)	
γ_{13} (State regulated, Region 3)	--	-0.91**	--
		(1.98)	
γ_{14} (State regulated, Region 4)	--	-0.16	--
		(0.50)	
γ_{15} (State regulated, Region 5)	--	-0.12	--
		(0.30)	

γ_{16} (State regulated, Region 6)	--	-1.55**	--
		(3.87)	
γ_{17} (State regulated, Region 7)	--	0.25	--
		(0.62)	
γ_{18} (State regulated, Region 8)	--	-0.58	--
		(1.06)	
γ_{19} (State regulated, Region 9)	--	-0.47	--
		(1.00)	
γ_{20} (State regulated, Region 10)	--	0.27	--
		(0.56)	
γ_{21} (EPA regulated, Region 1)	--	--	0.37
			(1.16)
γ_{22} (EPA regulated, Region 5)	--	--	0.71
			(0.25)
γ_{23} (EPA regulated, Region 6)	--	--	0.44
			(1.03)
γ_{24} (EPA regulated, Region 9)	--	--	-1.12
			(1.41)
γ_{25} (EPA regulated, Region 10)	--	--	0.72
			(1.31)
R^2	0.03	0.06	0.04

Hypothesis tests

Hypothesis		F-statistic	
All demographic H_0:[γ_5=δ_5=…=γ_8=δ_8=0]	1.96**	1.37	1.85*
EPA vs. state-regulated H_0:γ_2=0	2.02	--	--
EPA vs. state-regulated H_0:[γ_{11}=…=γ_{20}=0]	--	2.33**	--
EPA vs. state-regulated H_0:[γ_{21}=…=γ_{25}=0]	--	--	1.21

*Significant at 90% confidence level. **Significant at 95% confidence level.

[a]The *standardized discharge rate* is $y = m/\sigma$, where m = median of log discharge rate; σ = standard deviation of log discharge rate; each of these is a plant-level measure.

Regression #3 looks at whether EPA regulation is the same across EPA regions. We find no significant difference in EPA regulation across regions. We also ran regression #1 separately by region (not shown), but the results are similar to regression #1.

6.3 COMMUNITY EFFECTS

Because demographic variables are highly correlated, we focus on the effects of the community variables jointly rather than individually. The joint hypothesis that all community-level variables are zero is rejected at a 5 percent level of statistical significance in regression #1, and at a 10 percent level in regression #2.

The real meat is in the demographic variables' environmental and economic significance. A picture of the predicted effects of the community variables comes from comparing predicted violation rates for different communities and types of plants. Table 6 shows four communities; two are wealthy and two are poor, of which one is highly urban and the other highly rural. These communities were selected to be representative of community "types"; they are not outliers in the set of demographic variables. Predicted violation rates are given by $\mathrm{E}[\,\Pr(C>1\,|\,\tilde{y}_i,1)\,]$ where the expectation is taken over $\tilde{y}$, a normal random variable with conditional mean $X\beta$ and variance σ_y^2, with parameters taken from regression #1.

Our main finding is that plants in poorer, nonwhite communities have substantially higher violation rates. These violation rates are two to three times greater than those of their richer counterparts, based on the representative communities in Table 6. A plant in one of the two wealthy communities we selected is predicted to violate about once every two to four years. A plant in a poorer community is predicted to violate as frequently as once every eight months. Thus, communities exhibit economically and environmentally important differences in compliance, even when we take into account discharge variability, plant size, and type of regulator.[16]

Our results carry three caveats. First, the explanatory power of the regressions is somewhat low, so aspects of plant behavior remain to be explained. Second, while the differences across communities are *suggestive* of community pressure, they are not proof. We discuss this issue further below. Third, these results are conditional on a plant appearing in our sample and, in particular, not having undertaken any major changes during the study period.

[16] This relationship has been hinted at before, although shown in terms of mean discharges, which we argue may not give an accurate picture of behavior. Earnhart (2004b) examined community variables separately and found that a one-standard deviation increase in income per capita leads expected average discharge rates to fall from 0.34 to 0.27 (his Table 6b.ii). This finding cannot be compared directly to ours.

Table 6. Predicted violation rates for four communities (from regression #1)				
	Elk Grove CA	Crofton MD	Taft LA	Prichard AL
Non-White (percent)	25.0	7.7	84.8	94.3
Income (thousands)	51.8	53.9	9.3	8.0
Carpool (percent)	17.7	13.6	33.7	23.9
Manufacturing (percent)	5.6	7.0	36.7	15.3
Population (thousands)	13.6	15.5	0.8	24.9
Urban (percent)	73.4	100	0	100
Predicted violation rate: (percent)				
1-1-1 plant[b]	2.3	2.5	5.9	7.5
0-0-0 plant[b]	5.2	5.7	11.2	13.4

[a]Predicted violation rate is E[Pr(C>1|y,1], where the expectation is taken over y using re-gression #1.
[b]0-0-0 is a large volume, EPA-regulated, sewage treatment plant. 1-1-1 is a small volume, state-regulated, non-sewage treatment plant. Ln(*VOLUME*) = 0 is a little less than one standard devia-tion below the mean. The mean ln(*VOLUME*) is approximately 1.

6.4 DISCUSSION

Plants located in some communities have distinctly lower predicted probabilities of violation. This suggests either that these plants have lower abatement costs or a higher value for clean water, or face tighter scrutiny from regulators. A higher value for water quality implies – or is implied by – greater community pressure for lower wastewater discharges.

Neither the abatement cost nor regulator scrutiny argument is particularly compelling absent community pressure, given the very low levels of violations that we observe.[17] Plants that are overcomplying have no reason for abatement costs or regulator scrutiny to cause them to overcomply *more* unless some form of community pressure (or other unidentified non-regulatory effect; see McClelland and Horowitz 1999) is present. Lower abatement costs would, however, make community pressure more effective, since plants that have lower costs would find it easier to accommodate community activists. Managers at plants that have lower abatement costs would also likely find it easier to convince owners or regulatory boards to allow them to run a clean plant.

Additional non-statistical evidence that community pressure affects plants' discharge behavior comes from various statements by plant managers.

[17]Earnhart (2004b) examines how community variables affect regulator scrutiny, a phenomenon he labels "indirect pressure."

Kagan *et al.* (2003) report that "virtually every paper mill [they] visited reported significant environmental pressures from their host communities." McClelland and Horowitz (1999) quote an industry press article in which a pulp and paper mill manager describes its "good neighbor policy," which leads it to operate well below the regulated limits. Our findings provide concrete, numerical evidence of the relationship.

We must be circumspect in our claim, however, since we identify only a correlation between community characteristics and violation probabilities. We have not yet shown that this relationship is due to "community pressure."

7. RESULTS III: HOW WOULD A TIGHTENING OF THE REGULATION AFFECT DISCHARGES?

We wish to predict how much discharges would decrease, if at all, if permit levels were lowered. We refer to the z-r relationship, where z here is shorthand for z_m, median log discharges, and r is the median log permitted level. We examine dz/dr. If the regulation is not binding, then $dz/dr = 0$; if the regulation is "just binding," then $dz/dr = 1$. When plants are over-complying, dz/dr may take on any value.

We estimated the z-r relationship separately for each SIC. Results are shown in Table 7. A one-percent decrease in r is associated with a decrease in z_m between a 0.6 and 1.0 percent.[18] The coefficients on r for SICs 2 and 3 are significantly different from 1.

The Table 7 regressions must be treated with some care. Most plants face a permitted concentration of 30 mg/l, so cross-sectional variation in R is small. Some of the observations at R = 45 may be miscoded daily concentration limits. The effects on our estimates are unknown.

[18] Brännlund and Löfgren (1996) estimate the relationship between actual and permitted BOD quantity discharges as 0.7. McClelland and Horowitz (1999), arguing that U.S. regulations are binding and plants overcomplying, find that the derivative of annual BOD quantity with respect to the pollution limit is 0.65. Neither estimate is directly applicable to the z-r concentration relationship.

Table 7. Relationship between median discharges and permitted limits
Dependent variable: z_m (Plant-level median ln(discharge))[a]

	SIC=2	SIC=3	SIC=4
Constant	-.57*	-0.27	-0.61**
	(-2.39)	(-0.34)	(-4.53)
r	**0.83****	**0.59***	**0.97**
	(2.64)[b]	**(1.66)**[b]	**(0.99)**[b]
σ	-0.41	0.02	-1.04**
	(-1.78)	(0.04)	(-8.31)
ln(*VOLUME*)	0.05	0.01	0.01
	(1.92)	(0.13)	(0.90)
R^2	0.13	0.10	0.07
N	118	29	949

[a]To compare results with Table 3, recall that $m = z_m - r$.
[b]Numbers in parentheses in this row are *t*-ratios for the test that the coefficient on *r* is 1. All other numbers in parentheses are *t*-ratios for the test that the coefficient is 0.

7.1 IS THE REGULATION BINDING?

It may be tempting based on aggregate discharges to speculate that this regulation is simply not binding. "Binding" must be defined differently when discharges are random. One definition is that a regulation is binding when $dz/dr > 0$. This clearly holds for our data. The negative m-σ relationship also shows that the regulation is binding: If the regulation were not binding, plants would have no reason to compensate for discharge variability by reducing average discharges. This conclusion is consistent with other evidence that shows a positive marginal abatement cost for average discharges (Pittman 1981; McClelland and Horowitz 1999), and also with informal sentiment that this regulation has "bite."

8. CONCLUSIONS

8.1 FINDINGS

We find strong evidence that uncontrollable discharge variability leads water-polluting plants to reduce their average discharges. Plants pollute below –

sometimes far below – their permitted levels to reduce the chance of a violation. We further find evidence that suggests that a large proportion of plants are overcomplying with their BOD concentration permits even when accounting for this variability. We do not claim that *all* plants are overcomplying. A small subset continues to have poor compliance records.

Previous authors have speculated that community pressure is leading plants to overcomply. We find strong community characteristics effects. Richer-type communities have substantially lower violation rates than poorer-type communities.

8.2 Policy implications

We discuss implications for three regulatory issues: (i) efficacy of current regulations and enforcement strategies; (ii) introduction of tradeable permits; and (iii) informal regulation.

8.2.1 Current regulation

Because of discharge variability and community pressure, together and possibly separately, plants do not exhibit a one-to-one relationship between discharges and permitted levels. This feature must be taken into account when predicting the effect of any change in the regulation. A move to make the regulations more stringent by reducing the permitted concentrations would not necessarily lead to an equal reduction in discharges.

Discharge variability also highlights the role of EPA's enforcement policy. A move to enforce standards over a shorter time (say, daily limits) would likely reduce discharges further; while a move to enforce standards over a longer time (say, annual quantity) would likely raise discharges. These predictions arise because of differences in the variability of these measures.

8.2.2 WATER QUALITY TRADING

Tradeable permits, a commonly proposed direction for pollution regulations, would likely, in the case of random discharges, substantially *raise* overall pollution closer to the permitted level, R. If a fully tradeable, bankable pollution permit system were in place, plants could pollute exactly at their permitted level because whenever low discharges occurred plants would sell the excess permits and whenever high discharges occurred plants would buy permits from other plants. Thus, mean discharges would rise to the permitted level. Non-bankability or correlation across plants would temper this effect. Of course, many other issues are involved in the design of water quality trading schemes.

8.2.3 INFORMAL REGULATION

The large role that might be attributable to community pressure may lead readers to wonder whether the formal legal system of regulation could be modified or somehow "relaxed." This question has two answers. First, between 10 and 25 percent of the plants in our sample are *not* overcomplying, and some of them are rather seriously out of compliance. Overcompliance is a feature primarily of well-to-do communities. Our evidence, preliminary though it is, does not suggest that community pressure would be a successful regulatory tool in poorer communities.

Second, the interaction between community pressure and formal regulation is unknown. Do plant managers, in trying to convince the community (or themselves!) that they are doing a good job, point to their permitted levels as a benchmark against which their success can be measured? If so, do alternative non-regulatory benchmarks exist that could play this role? Answers to these questions would seem to be required before greater reliance on "informal" regulation were prescribed.

Discharge variability creates an agency problem for engineers and plant managers; it makes it more difficult for them to show that they are doing a good job in properly balancing treatment costs with the pecuniary and non-pecuniary costs of violations.[19] Excessively low probabilities of violation may be their optimal response in this situation. For regulators and the public, variability means they must judge, from the highly variable discharge pattern, the abatement efforts being made by plants. Indeed, we believe that managers may not realize the degree to which they are overcomplying because their "true" compliance is

[19]If discharge patterns are the result of an agency problem between plant engineers and managers or regulatory boards, then any of the changes described above, such as changes in EPA enforcement, would have complex consequences, including changes in the equilibrium of the agency game. We thank a referee for pointing out this implication.

masked by variability, which makes inference difficult. The policy implications of this problem have not yet been fully explored.

8.3 Future research

EPA's and the public's attention are beginning to shift away from BOD and toward nitrogen, phosphorus, and heavy metals. Future regulations are likely to focus on these pollutants. It would have been desirable to say something about current discharges and compliance, but we found the data for analyzing these discharges to be inadequate.

Cost data would have allowed us to deepen our analysis. Better data on plant characteristics and the inclusion of plants that we could not match to a zipcode would also allow us to measure relationships more precisely. The characteristics of downstream communities may also be important, not just the communities the plants are located in.

As argued in Section 4, we also hope to examine long-term trends in discharges and the causes and consequences of investments in new control technologies.

9. References

Anderson, E.E., and Talley, W.K. (1995). "The Oil Spill Size of Tankers and Barge Accidents: Determinants and policy implications," *Land Economics* 71(2): 216-228.

Berthouex, P.M., and Fan, R. (1986). "Evaluation of Treatment Plant Performance: Causes, Frequency and Duration of Upsets," *Journal of the Water Pollution Control Federation* 58: 368-375.

Brännlund, R., and Löfgren, K.G. (1996). "Emission Standard and Stocastic Waste Load," *Land Economics* 72(2): 218-230.

Brown, R.L., Durbin, J., and Evans, J. M. (1975). "Techniques for Testing the Constancy of Regression Relationships over Time," *Journal of Royal Statistical Society* B 37: 149-163.

Earnhart, D. (2004a). "Regulatory Factors Shaping Environmental Performance at Publicly-Owned Treatment Plants," *Journal of Environmental Economics and Management*, 48(1): 655-681.

Earnhart, D. (2004b). “The Effects of Community Characteristics on Polluter Compliance Levels,” *Land Economics* 80(3): 408-432.

Environmental Protection Agency, (1996). *Draft Framework for Watershed-based Trading*, Washington DC.

Foulon, J., Lanoie, P., and Laplante, B. (2002). “Incentives for Pollution Control: Regulation or Information?” *Journal of Environmental Economics and Management* 44(1): 169-187.

General Accounting Office (GAO). (1996). *Water Pollution: Many Violations Have Not Received Appropriate Enforcement Attention*, Washington DC.

Gray, W.B., and Shadbegian, R.J. (2004). “‘Optimal’ Pollution Abatement – Whose Benefits Matter, and How Much?” *Journal of Environmental Economics and Management* 47(3): 510-534.

Gunningham, N., Kagan, R.A., and Thornton, D. (2004). “Social License and Environmental Protection: Why Businesses Go Beyond Compliance,” *Law and Social Inquiry* 29(2): 307-341.

Helland, E. (1998). “The Enforcement of Pollution Control Laws: Inspections, Violations and Self-Reporting,” *The Review of Economics and Statistics* 80(1): 141-153.

Kagan, R.A., Gunningham, N., and Thornton, D. (2003). “Explaining Corporate Environmental Performance: How Does Regulation Matter?” *Law and Society Review* 37(1): 51-90.

Laplante, B., and Rilstone, P. (1996). “Environmental Inspections and Emissions of the Pulp and Paper Industry in Quebec,” *Journal of Environmental Economics and Management* 31(1): 19-36.

Leduc, R., Unny, T.E., and McBean, E.A. (1988). “Stochastic Models of First Order Kinetics of Biochemical Oxygen Demand with Random Initial Conditions, Inputs, and Coefficients,” *Applied Mathematical Modelling* 12(6): 565-572.

Magat, W., and Viscusi, W.K. (1990). “Effectiveness of EPA’s Regulatory Enforcement: the Case of Industrial Effluent Standards,” *Journal of Law* *and Economics* 33(2): 331-360.

McClelland, J., and Horowitz, J.K. (1999). "The Cost of Water Pollution Regulation in the Pulp and Paper Industry," *Land Economics* 75(2): 220-232.

Niemann, S.W. (2001). "Essays in Public Economics: Reduction of Pollution through Enforcement of Emissions Limits and Reduction of Household Energy Use," University of Wisconsin, unpublished Ph.D. dissertation.

Niku, S., Schroeder, E.D., and Samaniego, F.J. (1979). "Performance of Activated Sludge Processes and Reliability Based Design," *Journal of Water Pollution Control Federation* 51(2): 2841-2857.

Pittman, R.W. (1981) "Issues in Pollution Control: Interplant Cost Differences and Economies of Scale," *Land Economics* 57(1): 1-17.

Shimshack, J., and Ward, M. (2005). "Enforcement and Overcompliance," Tufts University, unpublished manuscript.

Washington Post, (June 6, 2003). "EPA: Few Fined For Polluting Water; Agency Says It Must Do Better Job Of Monitoring."