WORKING
PAPERS

OPTIMAL

OPTIMAL CIVIL PENALTIES

by John Nash

Federal Trade Commission

The views expressed in this paper are those of the author and do not represent the views or policies of the Federal Trade Commission or of any individual Commissioner.

I. Introduction
The purpose of this paper is to examine in some detail the optimal strategy for setting civil penalties for a regulatory agency that is charged with regulating an industry of risk-
� ~�

neutral firms that impose monetarily measurable social injury when they violate the regulations. In my model, there is no uncertainty in the minds of the firms or the regulator as to what constitutes a violation. The agency's budqet, which is outside of its control, determines the probability that it will be able to detect any particular violation, so this probability is exoqenous from the point of view of the agency.l Thus, the agency's decision variable is the size of the penalty. Its desideratum is to achieve the economically efficient violation rate. That is, it wants to set a penalty structure such that the firm will not violate unless its expected gain �f�r�o�~� the violation exceeds the expected social injury. The second section of this paper describes the standard approach to this problem, explains why this approach is unsatisfactory from the practical and the theoretical points of view, and �d�e�s�~�r�i�b�e�s� an alternative approach that IS more satisfactory. The third section describes an

1 The problem of establishinq an optimal probability of detec-tion is left to the leqislative hody that sets the hudget. For discussion of how this optimal probability might be set, see Becker, Landes and Posner, Polinsky and Shavell, and Keenan and Rubin.

empirical implementation of this alternative approach. Section IV discusses some extensions.

II. The Optimal Penalty
The standard approach in the literature to the optimal
penalty problem is to set it equal to (G/y), where G is the social injury from violation and y

One could in principle adapt the one-period model to a multi-period context by varying the penalty depending on the violation length in such a way that the standard equality is met in each

case, i.e., P(T) = �y�~�T�)�'� where y(T) is the probability of being caught at time T after violation begins. This would mean that, given any lenqth of violation, it would be profitable to continue the v i<.)la t ion one Hlv c period if and only if G* > G. Suppose, tor example, that the potential injury from violation is $100 per year, and that the probabilities of detection in the first three

years of violation are i' i, and 1. Then, the appropriate penalty
structure for this firm using this approach would be $500, $400, and $100, depending on whether the violation is caught during the first, second, or third year of violation. With such a structure, a firm that has never violated, but is contemplating a violation of one year would find positive expected profits from doing so only if G* > G. Likewise, for a firm that had violated one year or two years. It is intuitively plausible, and easily shown, that such a penalty structure also makes the expected net gain for a strateqy of 'permanent' violation positive if and only if G* > G. However, while such a variable penalty structure is optimal for a

certain class of violationl, it presents several problems in theory and at least one in practice.

1 For violations that involve a sequential "violate--do not violate" decision each period, the variable penalty structure is superior to a fixed penalty. For violations that involve a once-(footnote continued)

- 4-

One problem, 0f course, is that by recognizing that the
periodic probability of detection increases over time, this approach prescribes a penalty that is inversely related to the actual length of the violation. The firm that is caught with its hand in the cookie jar on its first attempt is thus fined an amount approachinq or equal to its wealth, while the firm that is finally cauqht after sneakinq cookies for many years is only fined the value of one cookie. The punishment, being tailored to deter prospective violations, does not seem to fit the actual crime. For this reason, it seems extremely dubious that this kind of model would ever be acceptable as a framework for public policy. In

establish the true length of the

or rP(T) for a fixed or variable penalty, respectively, where �~� is the discount rate), plus the direct gain from another period of violation (G* in either penalty regime), minus the difference between the penalty next period and the penalty this period (0 for a fixed penalty, negative for a penalty that decreases with the �l�e�n�g�t�h�~�o�f� violation). So the marginal value of evading detection is (rP* + G*) for a fixed penalty and (rP(T) + G*-dP(T)), with dP(T) < 0, for a variable penalty. For small values of I, the latter is clearly greater than the former, both because peT) is very large for T small and because of the negative dP(T) term. Thus, for at least some small values of T, the incentive to evade detection is greater under a regime of variable penalties, and more resources will be devoted to this goal at least toward the beginning of the violation than would be devoted under a regime of fixed penalties. While the incentive
to evade under a variable penalty regime will eventually (for T
large) fall below the incentive with a fixed penalty, the net
present value of resources devoted to evasion over the lifetime
of the violation under the former regime will certainly be larger
if the rate of discount is high.
On a prac 11.4 355.276ca41 0 0 1158 247.78 240 T Tm (aprac

of how many violators are cauqht each year, but also of how many were not cauqht, a fact which is by definition impossible to know. Other investiqators have used proxies for detection probability such as the requlatory aqency's budqet (Block, Nold, and Sidak), but such measures are obviously inexact.
Because the standard approach to penalty calculation is �~�
based on a quantity that is not directly estimatahle, and hecause its implications in a multi-period framework are not suitahle for public policy, it seems desirahle to develop an alternative approach, one that is based on an explicitly multi-period model and that immediately

where �~� is the expectations operator, T is the number of periods the violation continues before the firm is caught and fined, G* is the gain per period from violation, E is the appropriate discount rate, and P is the penalty. Performing the integration, (l) is equal to:
(2) G*[(l - E e-rT)/r) - P E(e-rT)

the violation j3 caught has an exponential distribution with par-meter A (c.f. Mood, Graybill, and Boes, p. 121). That is:

(4) f(T:A) = A e-AT,

(The frequency of such a distribution declines exponentially with T.) The moment generating function for such an exponential distribution is:

(5) E e-rT - A - A+r
Substituting (5) into (3) gives:

(6) p* = G[l-(A:r))/r(A:r) = G/A.
Tnis is the same answer that comes from the familiar I-period problem: the length of the violation is irrelevant: the penalty is always the one-period gain divided by the (constant) proba-bility of detection. The reason this is true is that, although the problem is multi-period in nature, it is entirely repetitive. That is, the potential violator is faced with the same "game" each period, so a penalty that deters violation one period will do so in every period.
For some classes of violations, we would expect that the probability of detection really does remain more or less constant. Cases against firms that are under surveillance (for whatever reason) would not generally be triggered by complaints, but rather by the (probably random) checking or surveillance pro-cess. It would seem, therefore, that the one-period probability

- 10-

where r(y) is the gamma function, defined as r(y)
for y>O. The moment generating function for this distribution is equal

the population with respect to these characteristics bearing on the chances of getting caught, though not necessarily with respect to

orders. The sample distribution vf T looks similar to an exponential distribution, being more or less downward sloping. Neither of two qoodness-of-fit tests would allow rejection of the hypothesis that the sample of T was drawn from an exponential distribution with parameter equal to the inver.se of the sample average T.l Since an exponential distribution for T would be - �~� produced if and only if the periodic �p�r�o�b�a�b�i�]�~�~�v� of detection is independent of the length of violation, this evidence can also be viewed as support for the proposition that this independence holds. The corollary is that the' first-best solution to the optimal penalty prohlem is to make penalties equal to the periodic social injury multiplied by the averaqe lenqth of viola-tion as an estimate of (1/A).2 The fiqures in table

previous Commission orders.l The sample 1S considerably smaller, and the average violation length shorter than for order violation cases,2but once again the sample distribution resembles a drawing from an exponential distribution. Neither the Kolmogorov-Smirnov test nor the chi-square goodness-of-fit test would reject the null hypothesis that the sample was drawn from an exponential distribu- �~� tion with parameter equal

length, the distribution of the random variable !, the length of violation before detection, is not exponentially distributed, and may be unimodal. Indeed, that is the form of the sample distri-bution in Table 3. In fact, in contrast with the

discuss the implications for the theoretical calculation of the optimal penalty, and then the implications for the empirical estimation of the necessary parameters. At first blush, one might conjecture that this modification in the model would mean that the optimal penalty should be larger, since there is some chance �t�h�a�t�~� the firm will not be caught in its life-time and so will avoid paying any penalty. One might also conjecture that the method used in the empirical section to estimate the optimal penalty would give downwardly biased results. For the case of constant probability of detection, I show that the first conjecture would be wrong and the second correct, though the bias in the estimate may not be serious.
In examining the theoretical implications, it is convenient
to consider first a world populated with firms with finite and fixed lives, fixed in the sense that the life-time of each firm is the same and is known to both the firm and the regulator. Call the life-time z. The desideratum, once again, is to devise an optimal penalty �~�*� such that the expected present value expectall

the

for the firm to be caught requires that T < z. So, for a firm prospectively considering a strategy of violation, the present value of the penalty the firm will pay is:
(9) PVP = P e-rT if T < z

= 0 if T > z.

If T is exp-onentially distributed, the expected present value of 1. the penalty is therefore:

z 00

(lO) E(PVP) = f (P e-rT) A e-ATdT +
o z

The firm qains an amount G* each period until it is caught, or it

,

expires, whichever comes first. After it expires, its "gain" is o for each period thereafter. The present value of the firm's total gain is:
*

(11) PVG fTG* e-rtdt G (1-e- = = - rT) if T < z

0 r

fZ * e-rtd t G* (l-e- = G = - rz) if T > z. 0 r
The expected present value of the gain from violation is therefore:

(12) z G * T 'T G * ex>' E(PVG) = f (--)(l-e-r) Ae- A dT + (--)(l-e- Orr

rz) f Ae-ATdT

z

= * �(�~�)�(�l�-�e�-�A�Z� _ (_A_) (1-e-(r+A)Z) r r+A

* = (�~�)� [(1 - _A_) (1 - e - (r + A �)�~�]� r r+A

- 19-

The optimal penalty is:
(13) p*'=(f).

The optimalfty can be verified by noting that when £.* is sub-stituted for P in (10), a comparison of (10) and the last equality in (12) shows that the expected present value of the qain exceeds the expected present value of the penalty if and only if G* > G. Comparing (13) with (6), it can be seen that the optimal penalty is the same, whether the firm's life is finite or infinite. This is easily generalizable to the case when the
firm's life-time is random. Suppose that �~�,� instead of being
fixed, is random with a frequency distribution g(z). Then the expectations in (10) and (12) would need to be taken over all possible values of �~�,� as well as over all possible values of T.

the That is, the expected present value of penalty

One can easily verify that the optimal penalty in (13) assures that E(PVG) > E(PVP) if and only if Q* > G, even in this case when the lifetime if the firm (z) is random. The reason for the seemingly counter-intuitive result that the penalty for a firm with a finite lifetime is the same as if the firm were infinitely lived is as follows: Truncation of the firm's life, whether at a random or a fixed time, diminishes the probability that the firm will be fined, but also diminishes the expected value of the violation proportionately, since the firm gets no gains from violation after it expires. The fine necessary to keep it from violating unless Q* > Q, therefore, remains unchanged.
Now, consider the statistical question of how the estimator
of the optimal penalty used in Section III would be affected if the population of firms from which the estimate is made is finitely lived, with random lifetimes. The optimal penalty in Section III (for cases that are presumed to have a constant probability of detection) was estimated by multiplying the gain per period by the sample average time until detection. This

latter value is the maximum likelihood estimate of 1 (X-) , the

multiplier in (6) • The question being considered here is whether

the sample average is a good estimator of (1) A if firms are finitely lived. It would seem that this would be a downwardly biased estimate in this case, because large values of T are less likely to be observed, having been removed from the sample by attrition (expiration). This is indeed the case. If the

- 21-

probability of a firm's expiration is the same each period, so that the distribution of a firm's lifetime is exponential, then the size of the bias can be exactly derived. Suppose that the probability that a firm will expire in any given period is S, so that its lifetime is distributed as:

(15) g(z) = S e-Sz• 'A
What we seek is the unconditional expected value of the sample average time until detection, which is what we suggested using as

an estimator of (f) in the optimal penalty calculation (equation
(13) above). This expected value can be calculated by first calculating the expected value of T conditiohal U90n �~�,� and then integrating over all possible value of z. First, consider the conditional distribution of T given z. For a given value of �~�,� no value of ! greater than �~� will be observed, since the firm would have been removed from the sample. The probability of observing any value of ! less than or equal to �~�,� however, is the unconditional frequency of ! divided by the cumulative distribu-tion of T evaluated �a�t�~�.� That is, the conditional distribution of T, given �~�,� can be written:
(16)

>..e->"T

f(TI z) = (1 _ e- AZ) for T .. z

= 0 for T > z.
Thus, the

1 z e-'\ z

= .\ (1 - e-Az)

The unconditional expectation is:

(18) E(T) = f E{Tlz) Se-Sz dz
o

= J o

00

J
o

00

Since �(�~�)� is what we are trying to estimate, the last term of the right-hand side of the last equality in (18) shows the size of the bias. Obviously, the bias depends on the true values of S and.\. A realistic estimate of S seems to be about 0.004 for a sample population of all commercial and manufacturing firms in

the economy.l 1 For a true value of (r) of 3.1 (as estimated for

order violation

act in bad faith, frequently declarinq bankruptcy ano re-orqaniz-ing to avoid debtors or suits by injureo consumers or requlatory agencies), this bias could be more of a problem. For values of !

1 = 0.05 and (r)= 3.1, for example, the bias in the estimate of <i) is about 21 percent of the true value.

The �e�~�~�d�e�n�c�e� discussed in this paper raises two interestinq

� ~�l�l�e�s�t�i�o—%�

v. Conclusion
The purpose of this paper was to develop a model for
calculating penalties that would be theoretically and practically more appropriate for multi-period violations than the one-period model that is currently the only model in the literature. The model is depigned to accommodate classes of violations for which detection probabilities change over time, and there is some evidence to sugqest that this class is important. While the model was developed specifically in the context of a violation that, once bequn, continues until the violator is caughtl, it may have more qeneral applicability to virtually all multi-period violations, even those that involve sequential decisions, given that the alternative penalty structure for such cases (a penalty inversely related to the length of the violation) is infeasible. An important implication of the model is that the size of the penalty that is socially optimal depends on characteristics of the class of case being considered, and in particular on whether the periodic probability of detection is independent of the violation length. Finally, extensions of the hasic model suqqest that it is theoretically appropriate for multi-period violations with constant prohability of detection, even when violators have finite lives, and that the hias produced by the empirical methods suggested here is not likely to be serious in this context.

I Or until some unanticipated changes occur in the qain or the penalty.

- 25-

TABLE l.--Length of Violation for �~�e�r� Violation Cases

Length Nurrt>er of Actual Theoretical Actual Theoretical
(years) Cases Frequency Frequency* Cumulative Cumulative*

0-1 12 0.273 0.276 0.273 0.276

2 8 0.182 0.200 0.455 0.476

3 6 0.137 0.144 0.592 0.620

4 8 0.182 0.105 0.774 0.725

5 �~� 5 0.114 0.076 0.888 0.801

6 1 0.023 0.055 0.911 0.856

7 1 0.023 0.040 0.934 0.896

8 0 0 0.029 0.934 0.925

9 1 0.023 0.020 0.957 0.945

>9 2 0.045 0.055 1.002 1.000

Total 44

Average: 3.10 years.

* From an exponential distribution with parameter (\ = 0.323)

TABLE 2.--Lenqth of Violation for Comrrdssion Determinations (Mainly Synopses)

Lenqth Number of Actual Theoretical Actual Theoretical (years) Cases Frequency Frequency * Ctnuulative Ctnuulative*

0-1 11 0.611 0.496 0.611 0.496

2 �~�

2 0.111 0.250 0.722 0.746
3 3 0.167 0.126 0.889 0.872
4 0 0 0.063 0.889 0.935
5 0 0 0.032 0.889 0.967
>5 2 0.111 0.033 1.000 1.000 Total 18

Average: 1.46 years.
* From an exPOnential distribution with parameter (A = 0.685) estimated by maximum likelihood.

Q = 7.146 0(4»0.10 0(5)::0.25
K-S = 0.122 0»0.20
See Table 1 for explanation of these statistics.

- 27-

TABLE 3.--Length of Violatkn for Section 5 Redress cases

Length Number of Actual Theoretical Actual Theoretical
(years) cases Frequency Frequency * Cumulative Cumulative*

0-1 9 0.134 0.105 0.134 0.105
2 �~� 7 0.105 0.189 0.239 0.294
3 16 0.239 0.184 0.478 0.478
4 12 0,179 0.151 0.657 0.629
5 10 0.149 0.107 0.806 0.736
6 2 0.030 0.083 0.836 0.819
7 4 0.060 0.058 0.896 0.877
8 3 0.045 0.039 0.941 0.916
9 3 0.045 0.026 0.986 0.942
>9 1 0.014 0.057 1.000 1.000 Total 67

* From a gamma distribution with parameters (y = 1.940, A = 0.517) estimated by maximum likelihood.
Q = 10.836 0(7»0.10 0(9»0.25
K-S = 0.120 0>0.20
For the null

TABLE 4.--Lenqth of Violation for Truth-in-Lendinq cases

Length Number of Actual Theoretical Actual Theoretical
(years) Cases Frequency Frecruency* Cumulative Cumulative*

0-1 8 0.320 0.416 0.320 0.416

2 �~� 9 0.360 0.243 0.680 0.659
3 2 0.080 0.142 0.760 0.801
4 2 0.080 0.083 0.840 0.884
5 3 0.120 0.048 0.960 0.932
>5 1 0.040 0.069 1.000 1.000 Total 25

Average: 1.86 years.
* From an exponential distribution with parameter (A = 0.538) estimated by maximum likelihood.

Q = 4.772 8(4»0.25 8(5)::0.50
K-S = 0.10 <5 »0. 20
See Table 1 for an explanation of these statistics.

- 29-

REFERENCES

Altrogge, P. and W.F. Shughart II, "The Regressive Nature of Civil Penalties," International Review of Law and Economics, 1984, 4:55-66.

Becker, G.S., "Crime and Punishment: an Economic Approach,"
Journal of Political Economy, March/April 1968, 76:169-217.

Block, �M�.�K�~�,� F.C. Nold, and J.G. Sidak, "The Deterrent Effect of
Antitrust Enforcement," Journal of Political Economy, June
1981, 89:429-445.

Fischer, F.E., Fundamental Statistical Concem.r7.9/0 0672 0 0 7.4 139.85 5

-

e = the maximum ob::;ervable in the sample (4 years in
the above example).

While the potential exists for this kind of bias to be a
problem for data of this kind, it is unlikely to seriously affect the conclusion of the current paper for the classes of cases considered herein. The largest class of cases, both in terms of the �F�T�~�'�s� �c�a�~�e�-�l�o�a�d� and the sample used here, comprises the "Section 5" cases. These cases are based on Section 5 of the FTC act, which has been in effect since 1914.2 The bias introduced by such a long limit on the maximum waiting time is negligible. The second la.::-gest class of cases are those brought against
a firm for violating an outstanding order which proscribes or prescribes certain conduct on its part. Here, there is more potential for bias, since the maximum observable length for a violation is the period between the issuance of the order and 1984. However, a close examination of the data not only indicates that the bias is not likely to be large, but also allows us to approximate its size and adjust the parameter estimate accordingly. First, most (about 58% of the 38 cases where proper order dates were recorded) of the cases dealt with orders issued in 1973 or earlier, meaning that the maximum observable waiting time was at least 10 years. (Only 1 of the 38

2 It should be

had an order date after 1976.) Second, when the pre-1973 sample is considered alone, the estimate of the mean waiting time is not dramatically different from that of the whole sample. For the pre-1973 sample, it is 3.33 years; for the whole sample of 38, it is 2.75.3 It is probably reasonable to believe that 3.33 is a better estimate of the true mean than is 2.75 (or 3.1, for that matter), but it may still have some bias. How much? Since the average order in the pre-1973 sample was issued about 15.1 years before 1984, we could substitute 15.1 for e and 3.33 for e in (AI) and solve for e The result is e = 3.6. Doing this exercise using the numbers from the full sample (e = 2.75
and e = 12.7) gives e = 3.0. Taken together, all and

1979, meaning that the maximum observable waiting time is relatively short. Because the cases are so close together in time, it is not feasible to perform the same kind of anaysis on this class of cases as I did for the order violation cases (i.e., comparing a sub-sample to the full sample) to judge the severity of this bias. However, there is one good reason to believe that the bias is not serious. Generally, when synopses are sent out, they are sent to firms that staff believes to be engaging in a certain practice. These firms are then checked soon after being sent the synopsis to determine if they are in violation. (Sometimes synopses are sent to a large number of firms in a particular industry because staff believes that most of the industry is engaging in the practices, even if they do not suspect each individual firm. Even in such cases,

of the Truth in Lending Act, which was enacted in 1969. For
these cases, the maximum waiting time in the sample would be on
the ord7ahi7 0 0 0 0 c0fa15.7 49543l0.41 Tm 504ord7ahi7 506391170fa15.7 4