IMPERFECT INFORMATION AND ANTITRUST VIOLATIONS:

THEORY AND EXPERIMENTAL EVIDENCE

by

Michael K. Block*

and

Stanley S. Reynolds**

This research was funded in part by the National Institute of Justice grant 83-IJ-CX-4055. Points of view or opinions stated in this document are those of the authors and do not necessarily represent the official positions or policies of the National Institute of Justice or of the United States Sentencing Commission.

*U.S. Sentencing Commission, 1331 Pennsylvania Ave., NW, 14th Floor, Washington, D.C.

**Department of Economics, University of Arizona, Tucson, Arizona 85721.

INTRODUCTION

Economic analyses of the deterrence of criminal antitrust activity have focused on two types of public policies. The first type of policy specifies penalties imposed on violators of antitrust laws. The penalties include payment of fines to the government, imprisonment, and payment of treble damages to injured parties. The second type of policy is the intensity with which antitrust laws are enforced and with which potential violations are investigated by government agencies. This intensity of enforcement and investigation effort is viewed as having an important effect on the probability of detection of antitrust violations. The deterrence effects of such policies have typically been analyzed by positing an expected profit or expected utility calculus for potential antitrust violators. This analytic approach to deterrence of antitrust violators has been utilized in several studies, including Breit and Elzinga (1973, 1985), Landes (1983), Block, Nold and Sidak (1981) and Block and Feinstein (1986).

The point of departure for this paper is the observation that while antitrust penalty levels are often known to potential violators, there may be considerable uncertainty attached to the probability of detection. This lack of information about detection probabilities and its implications are examined in this paper. A potential violator of antitrust laws (e.g., the management of a firm) is viewed as making decisions over a number of time periods and is given the opportunity to revise his estimate of the probability of detection based on information acquired over time. The economic agent is assumed to act so as to maximize the DPV of a stream of expected utilities and to update his probability of detection using Bayes rule.[1] The key hypotheses that emerge from this theoretical model are tested by using data acquired from a series of laboratory experiments. The experiments are designed to control for crucial variables thought to affect an agent's decisions, including the information available about the detection probability.

This approach yields several types of interesting theoretical results. The first type of result concerns the dynamics of antitrust violations over time. Suppose that the only way the agent can acquire useful information about detection probabilities is through his own experience as a violator of antitrust laws. If a violation is chosen in period t and it is not detected then a violation will be chosen in t+1. Also, once the agent chooses not to violate the law the crime will not be chosen in any subsequent periods. The second type of result concerns the role of Bayesian learning. Because a crime choice permits the opportunity to learn about detection probabilities, such a choice has value for subsequent periods over and above the current expected utility of crime. As a consequence, a crime choice can be 'optimal' even though current expected utility is negative. For a simple two-period version of the model with a risk-neutral agent, we show that this positive informational value of crime decreases as the discount rate rises and as the variance of the subjective prior distribution over the detection probability falls. The third type of result deals with the policy trade-off between antitrust penalties and detection probabilities. We distinguish between the objective probability of detection (which is under the influence of policy) and the subjective detection probability as perceived by a potential violator. Changes in objective probability can alter subjective probabilities in two ways: by shifting the prior mean and by changing the likelihood of detections that influence Bayesian updating. If the prior mean shifts by only a fraction of the change in objective probability then an increase in penalties coupled with an offsetting decrease in the objective detection probability will have a short-run deterrent effect for a risk-neutral agent (for a two-period model).

This paper also reports on a series of experiments that were designed to test various hypotheses from the theory. The subjects in the experiments were students at the University of Arizona. The predictions on the dynamics of antitrust violations did not fare well in the experiments; there were numerous cases of falsifications of these predictions. The experiments provide strong empirical evidence of a deterrence effect for an increase in the penalty and an increase in the prior probability of detection. The observed average incidence of "crime" in the experiments was less than the incidence predicted by risk-neutral behavior in the dynamic model. This provides evidence of risk averse behavior in the experiments.

I. EXPECTED UTILITY AND BAYESIAN LEARNING

In this section we describe the hypothesized objective of an agent and the process of learning about crime detection probabilities. The notation is summarized in Table 1. The agent has a single period utility function, $U(\cdot)$, that is increasing in net income for the period. By choosing the no crime option the agent earns income y for the period with certainty. Choosing the crime option yields income $z>y$ if the crime is undetected and the income $z-e<y$ if the crime is detected. The variable e represents the monetary fine levied on the agent when crime is detected. The variable p represents the perceived probability of detection. If the individual is certain that p is the probability of detection then the expected utility

criteria indicates that no crime is preferred if $U(y) \geq (1-p)U(z) + pU(z-e)$ and crime is preferred otherwise (assuming indifference leads to nonparticipation). The utility function is normalized so that $U(y)=0$.

In what follows we consider a case in which p is not known with certainty and in which the individual has a known planning horizon of T discrete time periods. The objective function of the individual is the discounted sum of expected utilities,

$$(1) \qquad W = \sum_{t=1}^{T} \gamma^{t-1} E[U(m_t)],$$

where m_t is net income in period t and $\gamma = 1/(1+r)$ is the time discount factor. The individual is presumed to choose crime or, no crime in each period t=1, ..., T so as to maximize W.

A standard approach to modeling this sort of decision-making problem is to treat p as a random variable. The agent is assumed to have some prior probability distribution over different possible values of p. Bayes rule can be used to update the agent's distribution over p-values based on the agent's observations of the outcomes of illegal behavior. DeGroot (1970, pp. 174-175) shows that if the prior distribution over multinomial probabilities is a Dirichlet distribution then the posterior distribution after observing events is also Dirichlet. For our binomial case, a Dirichlet prior over (p,1-p) is equivalent to a Beta distribution over p. The Beta density function is,

$$(2) \qquad f(p;\alpha,\beta) \begin{cases} = \dfrac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} p^{\alpha-1} (1-p)^{\beta-1} & 0<p<1 \\ 0 & \text{elsewhere,} \end{cases}$$

where $\alpha>0$, $\beta>0$. The Beta distribution is very flexible and can accommodate a wide range of "shapes" by changing the values of α and β. The mean prior probability of detection is $E(p) = \alpha/(\alpha+\beta)$. If α=B=1 then the prior is simply the uniform distribution over (0,1) (i.e., a diffuse prior distribution).[2]

Suppose the prior distribution of p is $f(p;\alpha,\beta)$ and then a total of (d+n) draws are taken from the enforcement mechanism with d detections and n no-detections. The posterior distribution is Beta with parameters $(\alpha+d)/(\alpha+d+\beta+n)$. In what follows we assume the agent's prior distribution over p is Beta with parameters α and β.

In this analysis we assume that the only way for the agent to acquire information about p is to actually participate in illegal behavior (to "search" for the probability of detection). There are, of course, factors other than a firm's own criminal experience that can alter its perceived probability of detection. Evidence of spillover effects from antitrust enforcement is found by Block and Feinstein (1986). Spillover effects arise when antitrust enforcement activities spread from one submarket to another. So, a firm observing indictments in a related market or, submarket may revise its estimate of the detection probability upward. There is also evidence concerning the infrequent nature of recurrences in antitrust enforcement (see Clabault and Block (1982), pp. 1053-1070). It may be that the probability of detection of antitrust violation is greater for a firm that has been indicted in the past. Landsberger and Meilijson

(1982) analyze the impact of 'state-dependent' enforcement levels in the context of a model of tax evasion. Such a policy amounts to having detection probabilities that vary depending on whether an agent has a recent conviction or, has 'passed' a recent investigation. We do not include spillover effects or, state dependent detection probabilities in our analysis of antitrust deterrrence. However, both of these aspects of antitrust enforcement could be built into an extended version of our model.

II. 'OPTIMAL' ANTITRUST VIOLATIONS OVER TIME

A decision-making strategy that maximizes W in (1) is sought. Any strategy that specifies a fixed series of choices for periods 1 through T at the beginning of the planning horizon will be suboptimal. Such a strategy ignores possible information acquired about p over time through the agent's criminal experience. An optimal strategy at time t takes the experience up to t into account and also anticipates the current value of information acquisition for subsequent decisions. The dynamic programming approach can be utilized to characterize the optimal strategy. This approach involves the computation of a value function for each time t. The value function defines the maximum possible value of discounted, expected utility for periods t through T, conditional on information acquired prior to t.

An individual's value function is defined by using the backward recursion method of dynamic programming. Initially, we must define period T value given any feasible history of criminal detections and nondetections (d,n) up to time T. (Feasibility at time T requires that $0 \leq d+n < T$.)

(3) $$V_T(d,n) = \max\ \{0, \int_0^1 [(1-p)U(z) + pU(z-e)]\ f(p;\alpha+d,\beta+n)dp\}$$

Period T value is simply the greater of the utility of the outside opportunity ($U(y)\equiv0$) and the expected utility of crime in T. The expected utility of crime is computed using the updated prior distribution for the probability of detection, p. Equation (3) may be used to calculate value functions recursively for periods $t<T$.

(4) $$V_t(d,n) = \max\ \{\gamma V_{t+1}(d,n), \int_0^1 (1-p)(U(z) + \gamma V_{t+1}(d,n+1))$$
$$p(U(z-e) + \gamma V_{t+1}(d+1,n))]f(p;\alpha+d,\beta+n)dp\}, \qquad 1\leq t<T.$$

Equations (3) and (4) define a unique sequence of value functions for periods $t=1,2,\ldots,T$.

The agent's optimal decision in a period is directly related to his value function. If the second argument inside the curly brackets in (4) exceeds the first argument then crime is optimal in period t. In this case the probability of detection in t+1 will be revised based on the agent's experience in t. The expected value of this revision is specified by the value functions $V_{t+1}(d+1,n)$ and $V_{t+1}(d,n+1)$. If the second argument inside the brackets in (4) is not larger than the first then crime is not optimal in t. The expected value of the no-crime decision is $U(y) + \gamma V_{t+1}(d,n) = \gamma V_{t+1}(d,n)$. The information (d,n) available in period t is carried over into t+1 in this case.

Proposition 1 If crime is not optimal in period t then crime is not optimal in any later period τ, $1 \le t < \tau \le T$.

Proof: Without loss of generality, we look at the case for which $\tau = t+1$. The proof proceeds in three steps.

Step 1: We first demonstrate that if d+n<t then $V_t(d,n) \ge V_\tau(d,n)$. $V_\tau(d,n)$ specifies an optimal decision rule for the $T-\tau+1$ periods τ, $\tau+1$, ..., T. If the same decision rule is followed for periods t, t+1, ..., $T-\tau+t$ then an agent can earn an expected value of $V_\tau(d,n)$ over these periods plus a nonnegative expected value for periods $T-\tau+t+1$, ..., T. $V_t(d,n)$ cannot be less than the sum of these expected values, so $V_t(d,n) \ge V_\tau(d,n)$.

Step 2: If crime is not optimal in t then $V_t(d,n) = \gamma V_{t+1}(d,n)$. It follows that $V_{t+1}(d,n) \ge V_t(d,n)$ since $\gamma \le 1$. From step 1 we also have $V_t(d,n) \ge V_{t+1}(d,n)$. These two inequalities imply that $V_t(d,n) = V_{t+1}(d,n)$.

Step 3: Next we show that $V_t(d,n) = 0$. This follows directly from step 2 and $V_t(d,n) = \gamma V_{t+1}(d,n)$ if $\gamma<1$. For the case $\gamma=1$, suppose that $V_t(d,n) = V_{t+1}(d,n) > 0$. In this case, there must be at least one period in the sequence t+1, T+2, ..., T for which the crime choice has positive expected utility, given the experience to date. Let $\bar{\phi}$ be the mean subjective probability of detection for this period. Let $\rho > 0$ be the probability that all crimes committed in this sequence of periods are undetected. By choosing crime in t, the agent has expected value greater than $V_{t+1}(d,n) + \rho[\bar{\phi}U(z) + (1-\bar{\phi})U(z-e)] > V_{t+1}(d,n)$. That is, by starting the strategy

specified by $V_{t+1}(d,n)$ one period early, the agent earns expected value in excess of $V_{t+1}(d,n)$. The supposition that $V_t(d,n) > 0$ contradicts nonoptimality of crime in t.

Step 4: If $V_t(d,n) > 0$ then by step 2 we must have $V_{t+1}(d,n) = 0$ and crime is not optimal for any periods after t. This completes the proof.

Proposition 2 If the crime choice is optimal in $t<T$ and crime is not detected in t then the crime choice is optimal in period t+1.

Proof: Optimality of crime in t implies that $V_t(d,n) > \gamma V_{t+1}(d,n) \geq 0$. The proof proceeds by contradiction. Suppose crime were not optimal in t+1 after an undetected crime in t. The proof of propositon 1 establishes that t+1 value is zero in this case, $V_{t+1}(d,n+1) = 0$. It is also true that $V_{t+1}(d+1,n) \leq V_{t+1}(d,n+1)$ (this can be established by a backward induction argument on t). So, $V_{t+1}(d+1,n) = 0$ and $V_t(d,n) = \phi U(z) + (1-\phi)U(z-e)$, from equation (4), where $\phi = (\beta+n)/(a+d+\beta+n)$. The value of a crime choice in t+1 given experience (d,n+1) is greater than or equal to $\phi' U(z) + (1-\phi')U(z-e) > \phi U(z) + (1-\phi)U(z-e) > 0$ where $\phi' = (\beta+n+1)/(\alpha+d+\beta+1)$. Thus, $V_{t+1}(d,n+1) > 0$ which contradicts the supposition of nonoptimality of crime in t+1. This completes the proof.

Propositions 1 and 2 describe two testable implications of the joint hypotheses of Bayesian learning and discounted, expected utility maximization. Proposition 2 states that if crime is optimal in period t and there is a "favorable" outcome (from a criminal's point of view!) then crime continues to be optimal in period t+1. Essentially what occurs in this case is that no detection in period t causes the agent to revise the expected probability of detection downward, making the crime decision attractive in period t+1. Proposition 1 indicates that once a person stops committing crimes they should "stay stopped." If crime is not an attractive option in period t and no crime is committed then the agent has no additional information about detection probabilities in period t+1.

Thus, crime remains unattractive in period t+1. Again, it should be stressed that these predictions depend on having no sources of learning about detection probabilities other than an agent's own experience.

III. DETERRENCE EFFECTS OF ANTITRUST POLICY

The deterrent effects of changes in fines and enforcement levels are examined in this section. We will demonstrate the role that Bayesian learning plays in the dynamic analysis and show how deterrent effects of antitrust policies differ in the dynamic and static models. A simplified two-period version of the dynamic model with a risk neutral agent is utilized. The sequence of decisions for the agent is illustrated in figure 1. In each period, the agent chooses either crime (B) or no-crime (A). A crime choice is followed by a random draw from the enforcement mechanism. If crime is chosen in period 1, the agent updates his probability of detection for the next period based on the outcome in period 1.

Let $p_1 = \alpha/(\alpha+\beta)$ be the mean subjective probability of detection in period 1. Expected income from crime in period 1 is, $z-p_1e$. If crime is chosen in period 1 then the mean probability of detection for period 2 is updated to $p_2 = \bar{p} = (\alpha+1)/(\alpha+\beta+1)$ in the case of detection or, $p_2 = \underline{p} = \alpha/(\alpha+\beta+1)$ in case of no detection. The certain income from no crime is $y \equiv 0$.

The value function for period 1 is,

$$V_1 = \max\{\gamma \max\{0, z-p_1e\}, z-p_1e + \gamma p_1 \max\{0, z-\bar{p}e\} + \gamma(1-p_1)\max\{0, z-\underline{p}e\}\}. \tag{5}$$

The first term inside the outer curly brackets is the expected value of the no-crime choice; the second term is the expected value of the crime choice. Each of these expected values is calculated assuming an optimal decision will be made in period 2 based on the information available at that time. Define G to be the difference between these two terms:

(6) $$G \equiv z - p_1 e - \gamma \max\{0, z - p_1 e\} + \gamma p_1 \max\{0, z - \overline{p}e\} + \gamma(1-p_1)\max\{0, z - \underline{p}e\}.$$

G represents the net expected gain for the crime choice in period 1. Crime is optimal in period 1 if $G > 0$. So, a deterrent exists if $G \leq 0$. The <u>sign</u> of G is the determinant of the period 1 choice. As Block and Lind (1975, p. 485) point out, the magnitude and direction of changes in G with respect to changes in underlying parameters $\{z, e, \alpha, \beta, \gamma\}$ may also be important. The impact of parameter changes on G provides an indication of the strength of deterrence or, incentive effects.

First, the effect of a change in the fine on G is considered. For purposes of comparison, we also examine the effect of Δe on the simple, static expected income model. For the static model, the net expected gain from crime is

(7) $$\overline{G} = z - p_1 e,$$

assuming a mean probability of detection equal to p_1. The results are illustrated in figure 2. The net gain from crime is the same for the static and dynamic models except for an interval of fine levels. Over this interval, the net gain from crime is strictly higher for the dynamic model than for the static model.

The divergence between G and $\overline{G}$ reflects the value of information acquired through the crime choice in the Bayesian model. The presence of this informational value of crime leads to a curious result. A higher fine level is required to deter crime in the dynamic model than in the static model. Solving for the value of e that sets G equal to zero yields,

(8) $$\bar{e} = \frac{z}{p_1}\left[\frac{(\alpha+\beta+\gamma\beta)(\alpha+\beta+1)}{(\alpha+\beta)(\alpha+\beta+\gamma\beta+1)}\right].$$

The term in square brackets is greater than one. As a result, $\bar{e}$ exceeds the fine level required for deterrence in the static model, z/p_1. For fines in the interval $z/p_1 < e < \bar{e}$, crime is an optimal choice in period 1 even though the expected period 1 value of crime is negative. In this case, a favorable outcome in period 1 causes an upward revision of the future expected value of crime. An unfavorable period 1 outcome leads to the no-crime choice in period 2, with a certain payoff of zero.

The fine required to deter potential antitrust violators, $\bar{e}$, is a function of the discount factor, $\gamma \equiv 1/(1+r)$, and the parameters, α and β of the prior distribution over detection probabilities. It is straightforward to show that $\bar{e}$ is increasing in γ. As a firm discounts the future less, crime has more informational value and a larger fine is required for deterrence. A mean preserving spread of the prior distribution also increases $\bar{e}$. The greater the variance of the prior, the higher is the informational value of crime and the higher is the fine required for deterrence. On the other hand, as the variance of the prior approaches zero, $\bar{e}$ approaches z/p_1. As the prior variance gets smaller, the results from the dynamic model look more like the static model results.

A change in the level of resources devoted to (public) antitrust enforcement should also affect the incentives to violate antitrust laws. In the standard, static expected utility model an increase in enforcement activity has been viewed as leading to a higher probability of detection and lower expected utility for crime. In our Bayesian model, the objective probability of detection need not be identical to the subjective detection probability. In order to describe the impact of increased enforcement, we must describe how a change in objective probability effects the subjective probability. The following provides a fairly general way of approaching the problem. Let ρ be the objective probability of detection. The mean of the subjective prior probability, $\underline{p}_1$, is assumed to be given by,

(9) $$p_1 \equiv \bar{\rho} + \mu(\rho - \bar{\rho}),$$

where $0<\mu<1$ measures the response of the prior mean to changes in the objective detection odds. The parameter $\bar{\rho}$ can be viewed as a fixed initial objective detection probability. In the extreme case of $\mu=0$ the prior mean is completely unresponsive to changes in ρ. The subjective probability of detection would respond to $\Delta\rho$ only after additional observations of the outcomes of crime choices. If $\mu > 0$ then the prior mean p_1 does respond to a change in ρ. This might occur, for example, if the antitrust enforcement agency announced a higher budget for enforcement activities or, if the agency announced particular industries as targets of investigations. In this case, the subjective mean probability may also be adjusted subsequent to the announcement based on the agent's criminal experience.

The initial deterrent effect of a change in the objective probability of detection is found by differentiating G with respect to ρ. G is piecewise differentiable, so the derivative is broken down into several parts.

(10) $$\frac{\delta G}{\delta \rho} = \begin{cases} -\mu e \quad , & p_1 < \dfrac{z}{\dfrac{e((\alpha+\beta)(\alpha+1))}{\alpha(\alpha+\beta+1)}},\ p_1 > \dfrac{z}{e(\frac{\alpha+\beta}{\alpha+\beta+1})} \\[2ex] -\mu\left[e(1-\gamma)+\gamma(2-\underline{p}e)+\gamma(1-p_1)\ e(\frac{\alpha+\beta}{\alpha+\beta+1})\right], & \dfrac{z}{\dfrac{e((\alpha+\beta)(\alpha+1))}{\alpha(\alpha+\beta+1)}} < p_1 < \dfrac{z}{e} \\[2ex] -\mu\left[e+\gamma(1-\underline{p}e)+\gamma(1-p_1)\ e(\frac{\alpha+\beta}{\alpha+\beta+1})\right], & \dfrac{z}{e} < p_1 < \dfrac{z}{e(\frac{\alpha+\beta}{\alpha+\beta+1})} \end{cases}$$

An increase in ρ has a deterrent effect on period one crime as long as $\mu>0$. The magnitude of the deterrent effect depends on the current level of fines and detection probabilities (because of the "kinks" in G).

Next we consider the impact of offsetting changes in the level of fines and the objective probability of detection that leave the expected fine unchanged. Such offsetting changes are given by

$$\text{(11)} \qquad \frac{d(e\rho)}{de} = \rho + e\,\frac{d\rho}{de} \equiv 0.$$

The total differential of G is,

$$\text{(12)} \qquad dG = \frac{\delta G}{\delta e}\,de + \frac{\delta G}{\delta \rho}\,d\rho.$$

Substituting for $d\rho/de$ from (11) into (12) yields,

$$\text{(13)} \qquad \frac{dG}{de} = \frac{\delta G}{\delta e} - \left(\frac{\rho}{e}\right)\left(\frac{\delta G}{\delta \rho}\right).$$

Substituting for the partial derivatives of G in (13), simplifying and evaluating the expression at $\rho = p_1$ yields,

$$\text{(14)} \qquad \left.\frac{dG}{de}\right|_{e=\bar{e}} = -p_1(1-\mu)-\gamma(1-p_1)p_1\mu\,\frac{\alpha+\beta}{\alpha+\beta+1} - \mu\gamma p_1^2\,\frac{\alpha+\beta}{(\alpha+\beta+1)(\alpha+\beta+\gamma\beta)}.$$

This expression is negative when $\mu=0$, positive at $\mu=1$ and strictly increasing in μ. So, there is a critical value $\bar{\mu}$ between zero and one such that for $\mu<\bar{\mu}$, the derivative in (14) is negative. That is, if p_1 shifts by less than a fraction $\bar{\mu}$ of the change in ρ then an increase in the fine coupled with an offsetting decrease in ρ has a deterrent effect.

IV. EXPERIMENTAL DESIGN

The preceeding sections of the paper describe testable implications of a theory of criminal activity over time with Bayesian learning. In this section we describe the design of controlled laboratory experiments that were used to test hypotheses from the theory.[3] An alternative approach to hypothesis testing would be econometric tests based on field data. Such an approach would require information about sequences of choices over time of individual criminals and noncriminals; such information may be difficult to acquire. Information about undetected crimes is also required, and again this information would be difficult to acquire. Laboratory experiments also permit much more control of the environment for potential offenders than is possible in the field.

A laboratory environment was created in which the monetary levels, y, z, e and the time horizon T were set by the experimenter and known to the subjects. A critical part of the theory is the uncertainty that an agent has concerning the true probability of detection. We utilize a procedure orginated by Grether (1980) that permits experimental control of the prior probability of detection for a crime choice. This procedure involves specifying two urns (or, bingo cages) from which enforcement draws might be made. Each urn is filled with two types of balls. One type repesents a crime detection and the other type represents an undetected crime. Draws are made from the urns with replacement. The first urn has a probability of detection q_1 and the second urn has detection probability $q_2<q_1$. Subjects are aware of the composition of the urns and the detection probability for each urn. At the beginning of a sequence of T periods, a random draw is made that determines which of the two cages is selected and used for enforcement draws. Cage 1 (or, urn 1) is selected with probability m; cage 2 is selected with probability (1-m). A subject is aware of the probabilities of selecting the 2 cages but does not know which of the cages is actually selected. This procedure induces a prior probability of detection equal to $mq_1+(1-m)q_2$. After each crime choice by

a subject, a draw is made from a cage to determine detection or nondetection. The subject learns whether or not the crime was detected but does not know which cage the draw was made from. The outcome of the draw does permit the subject to update his estimate of the probability of detection.[4]

This probability structure would have the following interpretation in the antitrust context. There are two levels of enforcement resources that the agency might utilize in any one industry. A high enforcement level yields a probability of detection of antitrust violation q_1; a low enforcement level yields detection probability q_2. A firm does not know which level of resources the antitrust agency utilizes in its industry. Instead, the firm assigns probabilities to high and low enforcement resource levels in its industry.

The theory the experiments are designed to test is a theory of criminal violation of antitrust laws over time. However, the experiments do not use terms like crime, antitrust, price-fixing, etc. Subjects are simply offered a choice between two alternatives; choice A yields a certain payoff and choice B has a stochastic payoff, with the probability structure described above. Alternatives A and B are intended to represent no-crime and crime decisions, respectively. Since the theory posits that decisions depend only on the expected utility of income, the experiments are designed to focus on this aspect of decisions exclusively. If such a theory performs poorly in controlled laboratory experiments, its ability to predict behavior in the naturally occurring ("real") world is suspect.

Three types of experiments were administered to each subject. These were the baseline (BL), treatment 1 (TR1), and treatment 2 (TR2). The parameter values used are listed in table 2. In BL the parameters were set so that expected income for the crime choice in period 1 exceeds income for no-crime y, (based on the induced prior). In TR1 the fine level is higher than in BL. Treatment 1 allows us to examine the deterrence effects of an increase in fines. TR2 is the same as BL except for a higher m-value. Treatment 2 permits examination of the deterrent effect of an increase in the mean prior probability of detection. In TR1 and TR2 the expected income from crime is less than y, but the crime choice is optimal in period one for a risk-neutral person because of the value of information. Each subject was administered the following experimental sequence: BL, TR2, BL, TR1. This sequence permits us to determine whether subjects return to baseline behavior after the first treatment.

The subjects recruited for our experiments were students at the University of Arizona. The experiments were run on a personal computer. A computer program listed the instructions (see the appendix) on the screen and led the subject through a sample 12 period run of the experiment. A subject was prompted by questions generated by the program and the subject entered his/her responses on the keyboard.[5] Each subject was paid in cash at the end of his/her experimental session. A subject was given $4.00 of working capital at the start of the experiment with which to absorb possible losses.[6] The expected payoff to a risk-neutral, value-maximizing subject in an experimental trial is given by V_1 in the last row of table 2. Thus, the total expected payoff for a subject from working capital, 2 baseline trials, a treatment 1 trial and a treatment 2 trial is $22.21 = $4 + 2($5.29) + $3.78 + $3.85. The actual average payoff to subjects was $12.12. Each experiment lasted 20-30 minutes.

V. EXPERIMENTAL RESULTS

A total of 30 subjects participated in these experiments. As a result we have 60 observations for the baseline case and 30 observations each for treatments 1 and 2. Our first use of the experimental data was to see if the data falsified propositions one and/or two. Recall that proposition one states that once an individual does not choose crime they should not choose crime in subsequent periods. A simple test of the proposition is to see how often it is violated -- i.e., to see how often an A choice (no crime) is followed by a B choice (crime). Aggregating all of the data, violations occurred in 38% of the relevant cases. Proposition two states that if crime is chosen in period $t<T$ and is undetected then crime will be chosen in $t+1$. Violations of proposition two occurred in 31% of all relevant cases. The predictions on optimal intertemporal choice do not fare well in these experiments. Note that these predictions are independent of the risk attitutdes of subjects, so observed deviations from the theory cannot be accounted for by variations in risk preference across subjects.

Even though the theory is not a very good predictor of period-by-period choices of subjects in the laboratory, it may still be useful in the sense of predicting central tendencies of behavior. To measure the average incidence of crime, we define the variable X to be the proportion of B choices out of unconstrained choices in a 12 period run of the experiment. This X variable was used to assess treatment effects (deterrence effects of increases in fines and detection probabilities) and to test the predictions of a risk neutral version of the theory.

Statistical tests were conducted to determine whether the treatment effects of fines (TR1) and prior probability of detection (TR2) were significant. These tests are reported in table 3. We appeal to asymptotic normality of the ratio of mean differences to asymptotic standard errors. We tested whether the mean X-value for each of the treatments was significantly different from the mean X-value for the baseline experiments. We found that there was no significant difference between the first and second baseline experiments, so that the baseline results were aggregated for testing treatment effects. Each treatment yields a statistically significant reduction in the mean X-value (the value of the standard normal variate $z_{.05}$ is 1.645). Thus, both a higher fine and a higher prior mean probability of detection yield a significant reduction in the average incidence of the crime choice.

A second set of statistical tests was used to examine differences between observed and predicted X-values. The predictions are derived from the value-maximizatin problem and assume risk neutrality. Each predicted choice for a subject is made conditional on the experience of detections and nondetections the subject has prior to the choice. The results are reported in table 4. In each of the cases BL, TR1 and TR2 the risk-neutral predicted mean X-value exceeds the observed mean X-value. Moreover, the difference is statistically significant in each case. The risk-neutral theory predicts a significantly higher average incidence of crime than was actually observed in the experiments. This suggests that observed behavior is more likely to be consistent with risk averse preferences than with risk neutrality.

VI. CONCLUDING REMARKS

This paper examines the incentives for firms to violate antitrust laws when firms have imperfect information about the probability of detection.

A Bayesian model of intertemporal choice was formulated that permitted an economic agent to revise his prior estimate of the detection probability based on his own criminal experience. This approach yields several interesting theoretical results. A crime choice followed by no detection of the crime was shown to lead an optimizing agent to commit the crime again in the next period. On the other hand, once a decision is made not to commit a crime, the crime choice remains suboptimal for all subsequent periods. Imperfect information about the probability of detection implies that a crime choice has informational value. A crime choice provides the opportunity to learn more about the probability of detection. As a result, deterrence in the dynamic model with Bayesian learning requires a higher fine than is necessary in a static model.

The experimental results were inconsistent with the theoretical stopping and continuation decision rules for crime choice. There are several possible explanations for the failure of this part of the theory. Subjects may have made mistakes when entering their choices (i.e., pushing the wrong button on the computer) or subjects may not have fully understood the laboratory environment. The latter problem might be mitigated if subjects were experienced. All subjects were inexperienced in these experiments. One interesting possibility is that subjects failed to update the detection probability according to Bayes rule. We have collected data from the experiments that would allow us to test the hypothesis of Bayesian learning, as in Grether (1980), but we have not yet analyzed that data. The deterrence effects of fines and an increase in the prior mean detection probability on the incidence of crime are strongly significant. These deterrence effects emerge in spite of the fact that subjects faced considerable uncertainty about the true likelihood of detection. The risk-neutral theory was found to predict a higher average incidence of crime than actually occurred in the experiments. This provides support for a hypothesis of risk-averse behavior.

NOTES

1. Bayesian updating has been used in other analyses of the economics of crime. See for example Reinganum and Wilde (1983) and Salant (1984).
2. The result on Dirichlet distributions could be used to extend this analysis to cases with more than two outcomes for a criminal. For example, crime might be i) undetected or, ii) detected and given a small fine or, iii) detected and given a large fine. In this case, if the prior distribution over multinomial probabilities is Dirichlet then the posterior distribution after observing the outcome of a crime is also Dirichlet. Rothschild (1974) uses this sort of result in analysis of a consumer's search for the lowest price for a particular commodity.
3. A good explanation of laboratory experimental methods appears in Smith (1976). In a recent paper Cox and Oaxaca (1986) report on a series of experiments designed to test a stochastic control theory of individual job search behavior. The theory they test is similar to the present theory in that it requires an individual to make a series of choices over time in an uncertain environment.
4. This procedure does not yield a Beta distribution for priors and posteriors. However, it does generate well-defined prior and posterior distributions over the probability of detection that yield the same kind of theoretical results as the Beta distribution.
5. A computerized experiment has the advantage of greater control (more consistency) compared to a noncomputerized experiment. The computerized experiment has the possible disadvantage of forcing subjects to "trust" the computer to generate the random drawings as described. In Grether (1980) the random drawings were made from bingo

cages that either a subject or a monitor could see.

6. Any subject that did not have sufficient earnings to cover the fines that might result from a crime choice was constrained to choose no-crime for that period.

REFERENCES

1. Becker, Gary S., 1968. "Crime and Punishment: An Economic Approach," Journal of Political Economy, Vol 76.

2. Block, Michael K. and Jonathan S. Feinstein, 1986. "The Spillover Effect of Antitrust Enforcement," The Review of Economics and Statistics, Vol. 68, No. 1, pp. 122-131.

3. Block, Michael K. and Robert C. Lind, 1975. "An Economic Analysis of Crimes Punishable by Imprisonment," Journal of Legal Studies, 4:479-492.

4. Block, Michael Kent, Frederick Carl Nold and Joseph Gregory Sidak, 1981. "The Deterrent Effect of Antitrust Enforcement," Journal of Political Economy, 89:429-445.

5. Breit, William and Kenneth G. Elzinga, 1973. "Antitrust Penalties and Attitudes Towards Risk: An Economic Analysis," Harvard Law Review, Vol. 86, pp. 693-713.

6. Breit, William and Kenneth G. Elzinga, 1985. "Private Antitrust Enforcement: The New Learning," Journal of Law and Economics, Vol. 28, pp. 405-443.

7. Cox, James C. and Ronald L. Oaxaca, 1986. "Laboratory Experiments with a Finite Horizon Job Search Model," working paper, Department of Economics, University of Arizona.

8. DeGroot, Morris H., 1970. Optimal Statistical Decisions, McGraw-Hill Publishing Co., New York.

9. Eck, J. Ernst and L. Ricco, 1979. "Relationship Between Reported Crime and Victimization Survey Results: An Empirical and Analytical Study," Journal of Criminal Justice, Vol 7.

10. Grether, David M., 1980. "Bayes Rule as a Descriptive Model: The

Representativeness Heuristic," Quarterly Journal of Economics, Nov.

11. Kahneman, D. and A. Tversky, 1973. "Subjective Probability: A Judgement of Representativeness," Cognitive Psychology, Vol. 51, Dec.
12. Landsberger, Michael and Isacc Meilijson, 1982. "Incentive Generating State Dependent Penalty System: The Case of Income Tax Evasion," Journal of Public Economics, 19:333-352.
13. Reinganum, Jennifer F. and Louis L. Wilde, 1983. "An Equilibrium Model of Tax Compliance with a Bayesian Auditor and Some "Honest" Taxpayers," Social Science Working Paper 506, California Institute of Technology.
14. Rothschild, Michael, 1974. "Searching for the Lowest Price When the Distribution of Prices is Unknown," Journal of Political Economy.
15. Salant, Stephen W., 1984. "Litigation of Settlement Demands Questioned by Bayesian Defendants," Social Science Working Paper 516, California Institute of Technology.
16. Smith, Vernon L., 1976. "Experimental Economics: Induced Value Theory," American Economic Review, Vol. 66, May.
17. Zedlewski, Edwin, 1983. "Deterrence Findings and Data Sources: A Comparison of the Uniform Crime Reports and the National Crime Survey," Journal of Reserch in Crime and Delinquency, Vol. 20, July.

APPENDIX

INSTRUCTIONS

The experimenters are trying to determine how people make decisions. We have designed a simple choice experiment, and we shall ask you to make decisions at various times. The amount of money you make will depend on how good your decisions are.

The experiment will proceed as follows. Each run of the experiment requires you to make a series of 12 decisions. You will have two choices each time you make a decision. If you choose option A you earn $0.25 for sure. If you choose option B your earnings depend on the outcome of a random drawing. If the draw is a success, you earn a positive amount of money. If the draw is a failure, you lose money. We will tell you the $-amount of gains and losses for the B option before each run of the experiment.

The random drawings will be done by the computer. The process the computer uses to make these draws can be illustrated as follows.

-- Hit the RETURN key to continue with instructions --

There are two bingo cages on the table next to you. The cages are designated as cage X and cage Y. Inside both cage X and cage Y are six balls, some of which are marked with an N and some with a G. Cage X has 3 N and 3 G balls. Cage Y has 1 N and 5 G balls. A G draw from a bingo cage is a success. An N draw is a failure. The ball is returned to the cage after each draw. All the draws for a particular run of the experiment will have a probability of success that corresponds to either cage X or cage Y. However, you will not know whether cage X or cage Y is being used for a particular run. The cage to be used is randomly selected by the computer at the start of each run. We will tell you the odds that cage X will be chosen and the odds that cage Y will be chosen at the start of each run.

We will now walk through a complete run of the experiment to make sure that you understand the procedure. Remember, cage X has 3 N balls (failures) and 3 G balls (successes) and cage Y has 1 N (failure) and 5 G balls (successes For run # 0 a series of A and B responses were randomly chosen in order to illustrate the procedure.

-- Hit the RETURN key to continue with instructions --

RUN NUMBER 0:

CAGE X - 1 in 3 chance of selection
CAGE Y - 2 in 3 chance of selection
Option A yields $0.25
Option B yields a $1.00 gain (G) or $2.00 loss (N)

PERIOD	CHOICE	DRAW(IF B CHOSEN)	PROFIT OR LOSS
1	A or B? > A		$ 0.25
2	A or B? > B	G	$ 1.00
3	A or B? > A		$ 0.25
4	A or B? > B	N	$-2.00
5	A or B? > B	G	$ 1.00

-- Hit the RETURN key to continue with instructions --

6	A or B? > A		$ 0.25
7	A or B? > A		$ 0.25
8	A or B? > B	G	$ 1.00
9	A or B? > B	N	$-2.00
10	A or B? > A		$ 0.25
11	A or B? > B	G	$ 1.00
12	A or B? > A		$ 0.25

NOTE - Since it is possible for you to lose money in the experiment, you are given $4.00 with which to begin the experiment.

-- Hit the RETURN key to begin run number one --

TABLE 1

NOTATION

$U(\cdot) \equiv$ utility as a function of net income

$y \equiv$ income from outside opportunity

$z \equiv$ monetary gain from crime

$e \equiv$ monetary fine if crime is detected

$p \equiv$ probability that a crime is detected, $0 \leq p \leq 1$

$\gamma \equiv 1/(1+r) \equiv$ discount factor, $0 < \gamma \leq 1$.

$T \equiv$ length of time horizon, $1 < T < \infty$

TABLE 2

PARAMETER VALUES FOR EXPERIMENTS

	BL	TR1	TR2
z	\$1.00	\$1.00	\$1.00
y	\$0.25	\$0.25	\$0.25
e	\$1.40	\$2.20	\$1.40
T	12	12	12
m	1/4	1/4	1/2
q_1	1/2	1/2	1/2
q_2	1/6	1/6	1/6
p_1	1/4	1/4	1/3
$V_1(0,0)$	\$5.29	\$3.78	\$3.85

TABLE 3

STATISTICAL TESTS FOR DETERRENCE EFFECTS

	$\underline{z}$
Impact of Higher Fine (TR1/BL)	-3.42
Impact of Higher Probability (TR2/BL)	-4.07

$$z = \frac{\bar{X}^a - \bar{X}^b}{\left(\frac{\hat{\sigma}^2_{xa}}{n_a} + \frac{\hat{\sigma}^2_{xb}}{n_b}\right)^{1/2}}$$

$n_a = 30,\ n_b = 60$

a and b denote treatment and baseline, respectively

TABLE 4

STATISTICAL TESTS FOR THE RISK NEUTRAL MODEL

Experiment	$\bar{X}$	X^p	z
BL (n=60)	0.66	0.85	-6.11
TR1 (n=30)	0.48	0.73	-5.81
TR2 (n=30)	0.42	0.64	-4.43

$\bar{X}$ and X^p are observed and predicted mean X-values

$$z = \frac{\bar{X} - X^p}{\hat{\sigma}_x} \cdot \sqrt{n}$$

Figure 1

Decision Tree for Two Period Model

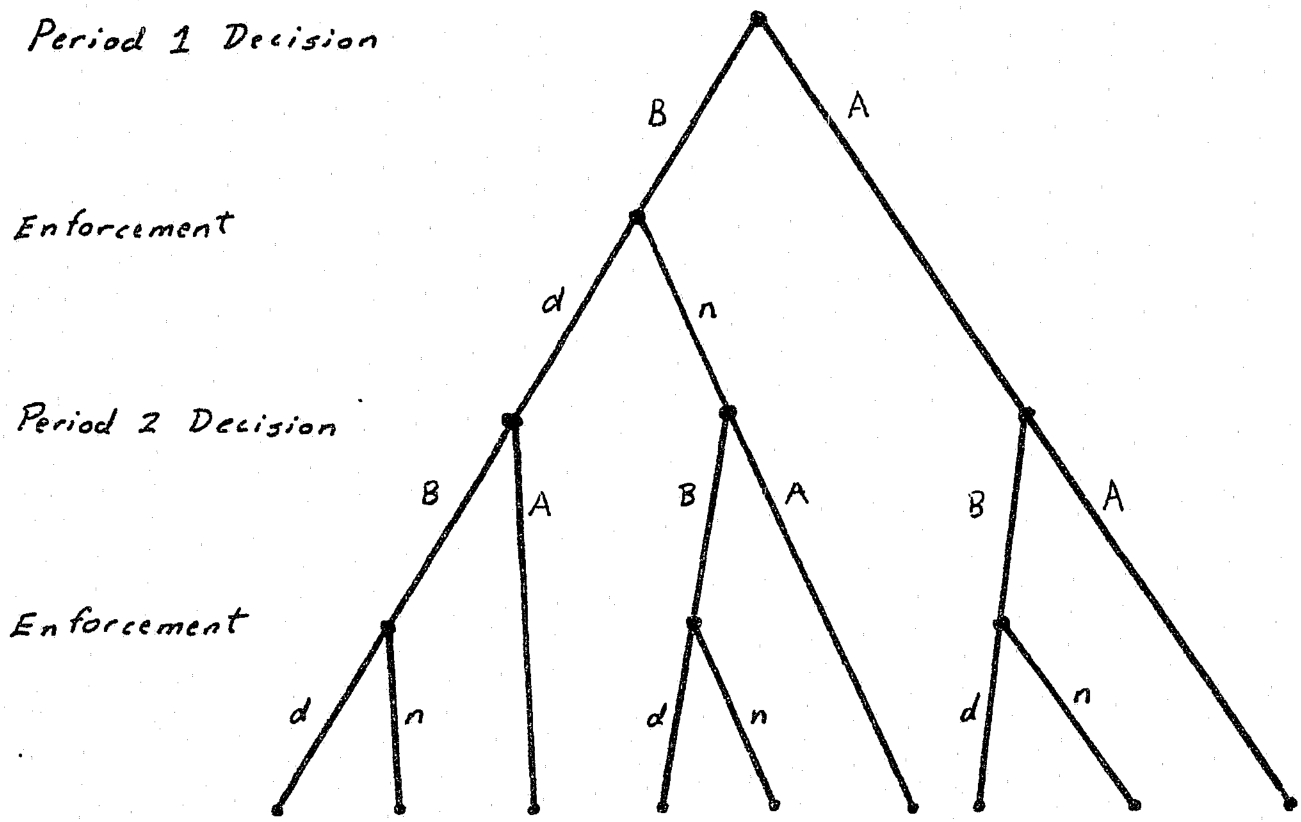

B ≡ crime decision
A ≡ no-crime decision
d ≡ detected crime
n ≡ undetected crime

Figure 2

Deterrence Effects of Fines

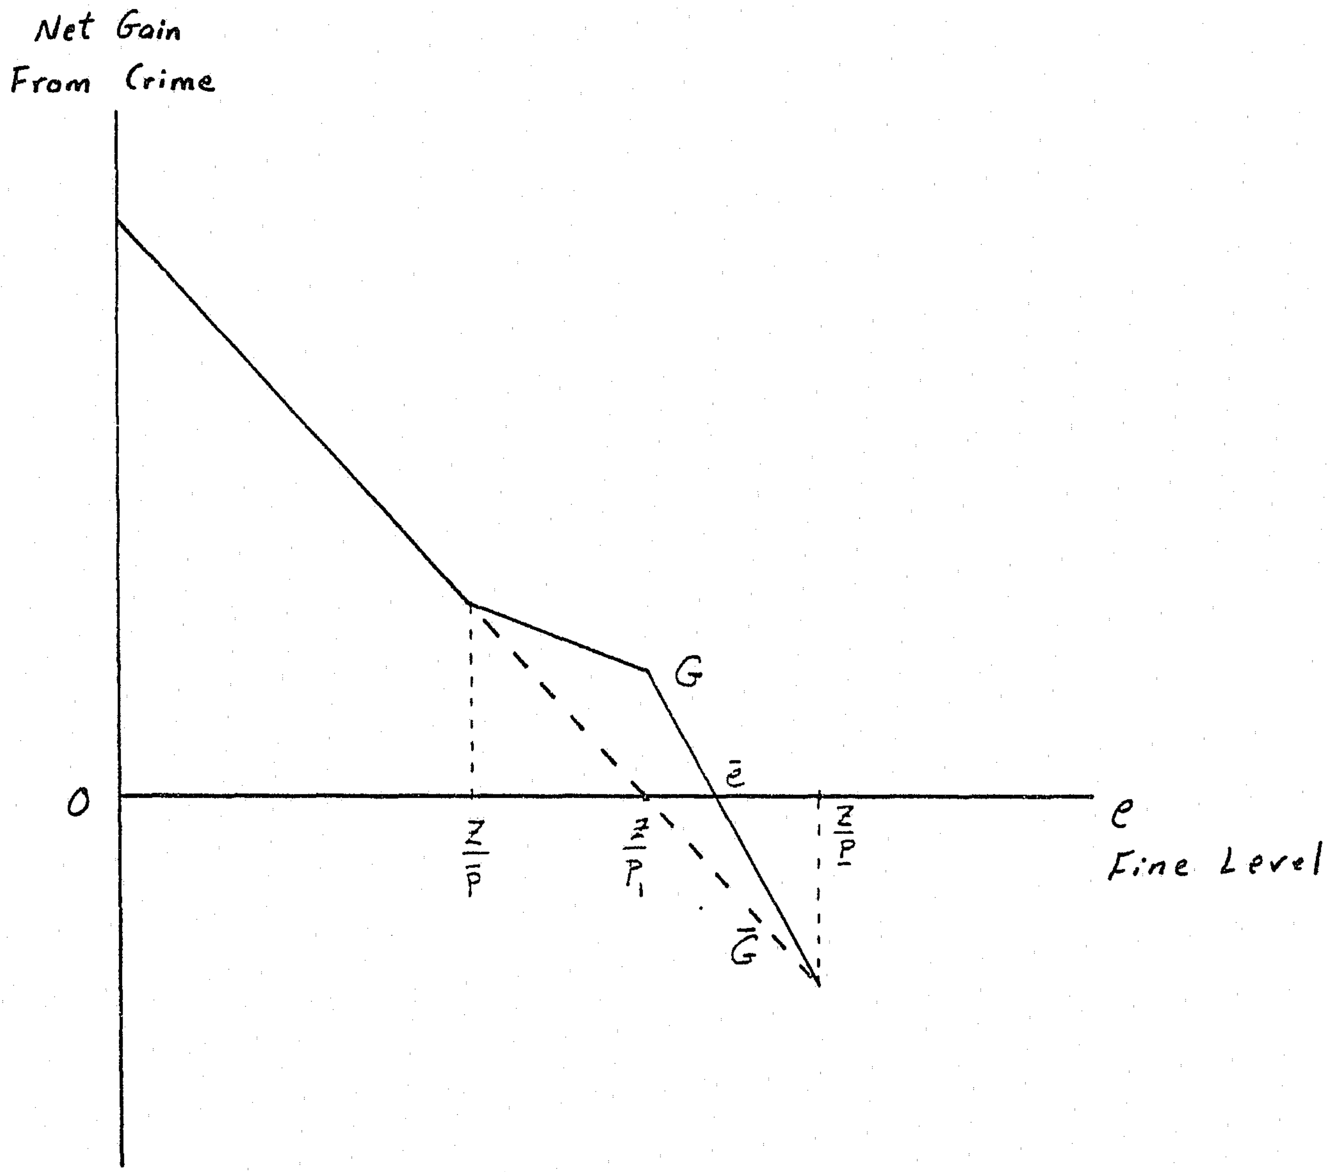

G = net gain from crime in dynamic model
$\bar{G}$ = net gain from crime in static model