CONF-9106112--1

UCRL-JC--105885

DE91 007148

TREATY VERIFICATION WITH AN UNCERTAIN PARTNER

Stein Weissenberger

This paper was prepared for submittal to the
International Conference on Game Theory
Florence, Italy
June 25-27, 1991

January, 1991

Lawrence
Livermore
National
Laboratory

Treaty Verification with an Uncertain Partner

Abstract

A simple model is used to analyze the performance of a system for verifying compliance with an arms control treaty. Blue and Red are partners in to a treaty. Blue prefers to comply, but is uncertain whether Red similarly prefers compliance (in the absence of threatened violation detection). Blue's uncertainty is modeled as a probability distribution over three different Red types: Violators, Compliers, and Deterrables. Criteria are derived to determine the level at which Blue should set his detection threshold, and when it is best for Blue not to verify at all. The results involve both game-theoretic and Bayes solutions.

1. Introduction

Arms control treaties are an important means to reduce the risk of war and the costs of maintaining armaments. However, recent experience shows that a nation considering a possible new treaty can expect to face decisions of sizable impact, complexity, and uncertainty. These problems motivate us to apply formal methods of analysis to bring greater clarity to the review of these difficult decisions.

In spite of their many complications, these treaties typically involve just two primary classes of decisions. The first class concerns what specific constraints to incorporate into the treaty. A real treaty may include many details, but the significant features typically involve limits imposed on the number, size or type of object built, or experiments performed. For example, how many missiles of what type should each party be able to deploy? How large a nuclear test should each party be able to perform? Although the methods of this paper apply to this problem, we treat it elsewhere [1], and consider it no further here.

The second class of decisions is the focus of this paper and concerns how to determine compliance with treaty constraints. Normally, certain data are taken (sites inspected, missiles counted, seismic measurements made, etc.), the data are analyzed, conclusions are drawn ("they complied", "they violated"), and actions taken (a public statement, an abrogation of the treaty, etc.) The goal of this verification process is to assure each side that the other is complying (when they are), to reduce damage by promptly detecting violations (when violations occur), and, through the threatened costs of an apprehended violation, to motivate the other to comply, i.e., to deter actual violations.

Specific verification decisions are of two principal types. One is how to configure the verification system to make it more effective, e.g., how to set the "decision threshold" in a seismic monitoring instrument (a paradigm that we will carry through the body of this work). This is essentially a "design" problem and is the focus of the present paper. The other question is how

much overall quantity (and thus quality) of verification to invest in. How much verification is enough? This is a larger resource-allocation problem, and is considered elsewhere [1].

The basic decision problems in treaty verification were described in Reference [2]. Here we consider the same general issues but take a somewhat different point of view, emphasizing "deterrence" as a central consideration in treaty verification. The basic model we use is that of "Inspector"/"Inspectee" interactions broadly explored from a game theory perspective in Reference [3] and applied to the specific verification problem of on-site inspection in [4]. (See also References [5] and [6] for related analyses of the problem.) Reference [1] established the general framework of the present paper and used it to evaluate the performance of verification systems as well as the desirability of the treaty constraints themselves.

In [3], the basic treaty verification problem was viewed as a two-person game in which the players — Blue and Red — have complete but imperfect information. Information was complete because both players knew all system parameters (utilities of both sides, detection probabilities); information was imperfect because evidence as to whether Red has complied or violated is "noisy". Reference [1] introduced several ways of dealing with a larger field of imperfection, specifically with uncertainty in Red's utilities. One approach there was ad hoc and involved the notion of maximizing "excess deterrence". The other approach, only briefly explored in [1], involved the introduction of multiple Red "types". This formulation is the focus of the present paper.

We begin the development in Section 2 by summarizing the formulation of Reference [1]. A basic Red "type" leads to a basic "standard" optimum verification system design. The basic Red type exhibits what is often considered "normal" Red behavior: while preferring to violate, this type may be deterred from doing so by an appropriately designed verification system. Next, in Section 3 two new types are introduced: one that always violates, and one that always complies. Blue is uncertain as to which type Red actually is, and this uncertainty is modelled by a (Blue) probability distribution over the three types. The remainder of the work explores how the probability

distribution over types, together with other system parameters, effects optimum verification system designs. These designs are compared with that resulting from the "base" case. In order to focus on the analysis of types, we omit a number of topics treated in [1] (maximum-deterrence designs, "maxi-min" designs, multiple violations, and an explicit qualitative treatment of various system parameters).

The analysis of types in Section 3 leads to results that dictate which of two sorts of verification system designs is best: one based on a "game theoretic" approach (the "base" case) or one based on a strict "Bayesian" approach. Unsurprisingly, if the probability of a "deterrable" type is sufficiently high, the game theory solution is dictated; otherwise the Bayes solution is best.

A concluding section summarizes results and closes with a discussion of recent trends in the literature of verification. New advocacy for decreased emphasis on verification is consistent with our results as the probability of a complier type becomes high.

2. Basic Model

We begin by considering a restricted but important problem that we later extend. A nuclear testing treaty is assumed to be in place, specifying the maximum permissible yield of nuclear tests. This limit could be the current 150 kt of the recently ratified Threshold Test-Ban Treaty (TTBT), or it could be some different limit negotiated in a future treaty, e.g., a Low-Yield Threshold Test-Ban Treaty (LYTTBT).

In this problem, described by the decision tree in Fig. 1, Red (R), the "Inspectee", decides whether to comply (C) or violate (V) the treaty. Blue (B), the "Inspector", has the problem of designing and implementing a verification system, making measurements, determining whether or not a violation has occurred, and, finally, taking the appropriate actions of challenge (C) or accept (A). In the decision tree in Fig. 1 Blue's reaction (challenge or accept) is modeled as a direct consequence of the design of the verification system, and is not itself the decision of interest to us here.

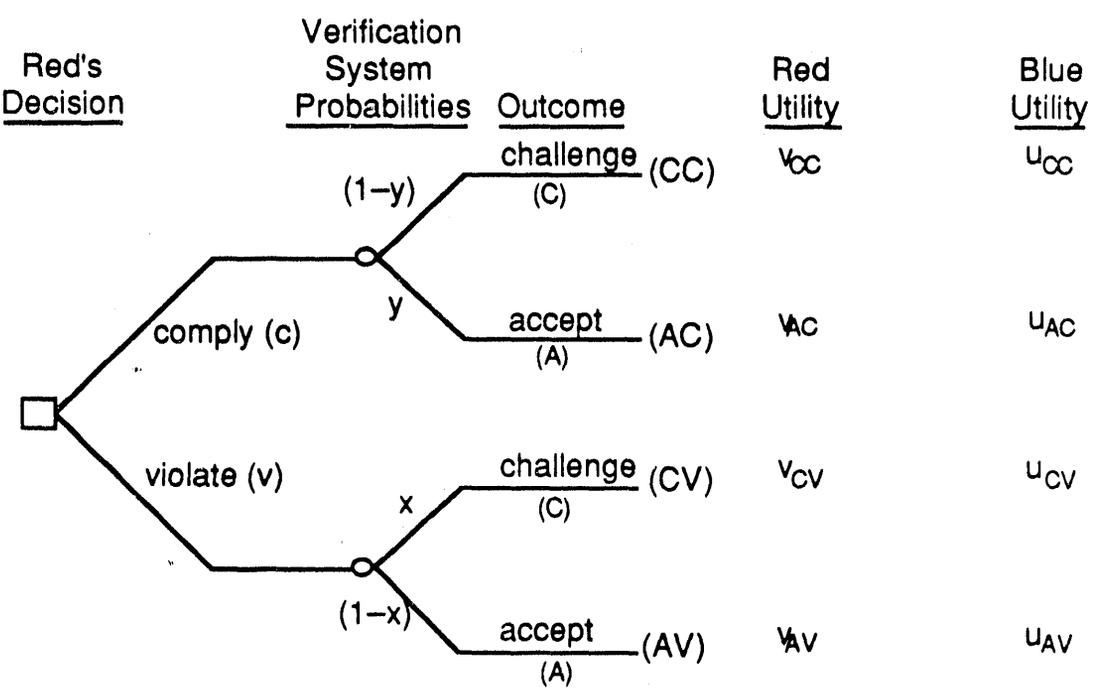

Fig. 1

Basic decision tree

Utilities of Blue are designated with a "u", those of Red with a "v". We assume at this point that the utilities satisfy the following inequalities:

$$v_{AV} > v_{AC} > v_{CV} > v_{CC} \tag{2-1}$$

$$u_{AC} > u_{CC} > u_{CV} > u_{AV} \quad . \tag{2-2}$$

This basic utility ordering is represented graphically in Fig. 2. Refs. [1] and [3] discuss the rationale for these orderings. We point out here only that they imply that Red would most of all prefer to violate (if not caught), and that Blue prefers, over anything, that Red comply. Red's preference for a challenged violation over a challenged compliance is consistent with Red's retention of some value from violation even if challenged.

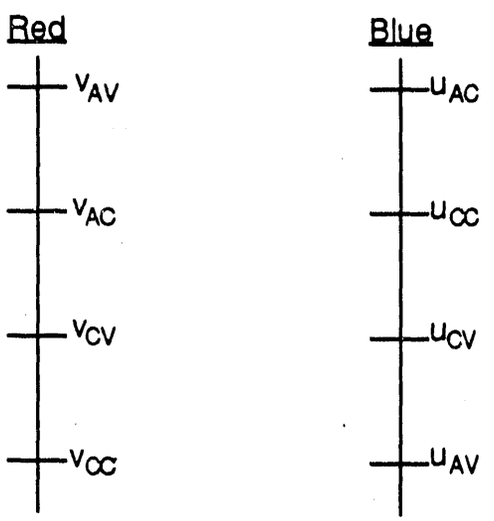

Fig. 2

Basic utility ordering for Blue and Red

We summarize Blue's "design" problem through the choice of two related probabilities, x and y, that describe the performance of the verification system. As in [2] and [3], x is the probability of correctly challenging a violation, and y is the probability of correctly accepting compliance. The probabilities x and y are not independent, but are connected through a relation we will call the detection characteristic. An example characteristic is shown in Fig. 3. Blue's problem is to select x and y in order to achieve a desirable outcome for Blue; Red's problem is to decide, from Red's point of view, whether to comply or violate. We assume that both Blue and Red are

expected-utility maximizers and (for now) that both have complete and perfect information about all utilities and the x and y probabilities.

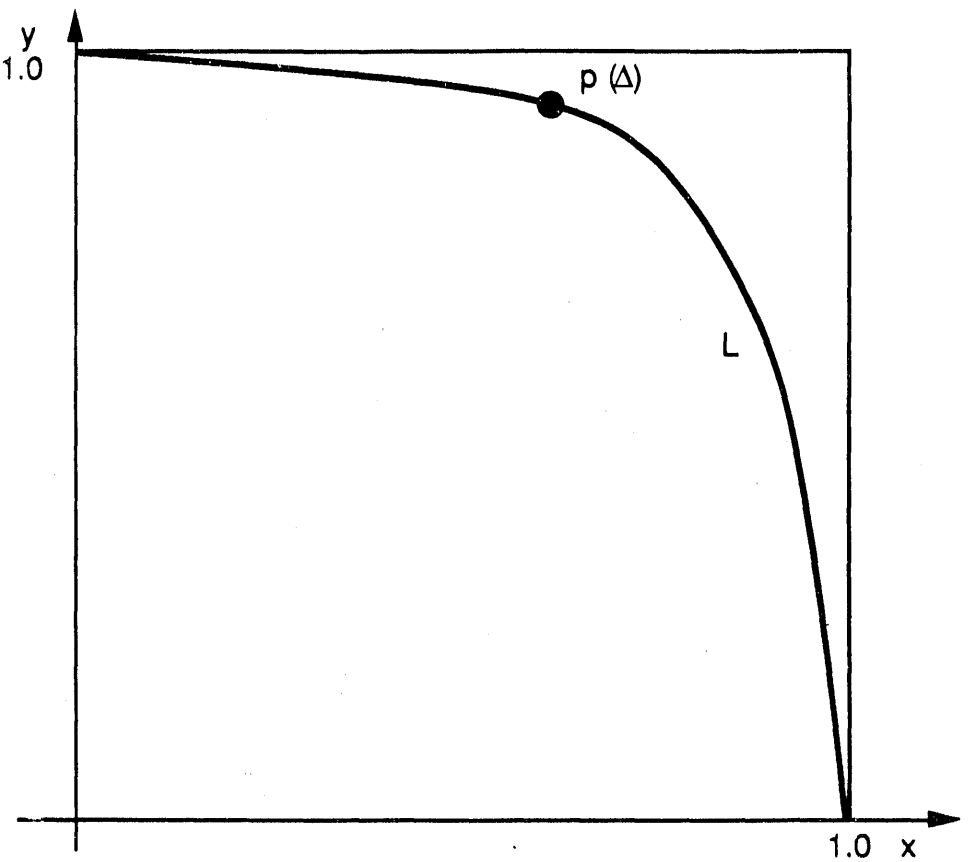

Fig. 3

A verification system detection characteristic

The detection probabilities x and y depend on a single verification-system parameter Δ, the "decision threshold". If Blue's measurement of Red's test yield exceeds Δ, then Blue reacts as if a violation occurred, i.e., "challenges"; otherwise Blue accepts the test as in compliance. As Δ varies, x, y will in general vary over the detection characteristic, shown as the curve L in Fig. 3.* The point p (Δ) represents a particular verification-system design resulting from a choice of threshold Δ, and giving rise to a detection-probability pair x, y. Some general characteristics of

* The connection between (x, y), Δ, and other system parameters is discussed in some detail in [1].

this relationship are that it passes through the points (1,0) and (0,1) (for sufficiently small and large Δ, respectively), and that it is monotone (non-strictly) decreasing and convex.

Red's decision (whether to comply or violate) is based on a comparison of Red's expected utility resulting from compliance, $EV_{R|C}$, and the expected utility resulting from violation, $EV_{R|V}$. From Fig. 1,

$$\begin{aligned} EV_{R|C} &= yv_{AC} + (1-y)v_{CC} \\ EV_{R|V} &= xv_{CV} + (1-x)v_{AV} \end{aligned} \tag{2-3}$$

It's convenient to define a quantity we will call deterrence, D_R, as the difference between these values, i.e.

$$D_R \equiv EV_{R|C} - EV_{R|V} \tag{2-4}$$

Substitution from (2-3) into (2-4) gives

$$\begin{aligned} D_R = (v_{AV} - v_{CV})x + (v_{AC} - v_{CC})y \\ -(v_{AV} - v_{CC}) \end{aligned} \tag{2-5}$$

When deterrence D_R is positive, Red will comply, when deterrence is negative Red will violate, and when it is zero, Red will be indifferent between the two.

$$\text{Red} \begin{cases} \text{complies when } D_R > 0 \\ \text{violates when } D_R < 0 \end{cases} \tag{2-6}$$

As in [3], we show graphically how Red's decision depends on x and y in Fig. 4.

A "deterrence region" in Fig. 4 exists provided

$$\text{and} \quad \begin{aligned} v_{AC} &> v_{CV} \\ v_{AV} &> v_{CV} \end{aligned} \quad , \tag{2-7}$$

both reasonable assumptions on Red's preferences.

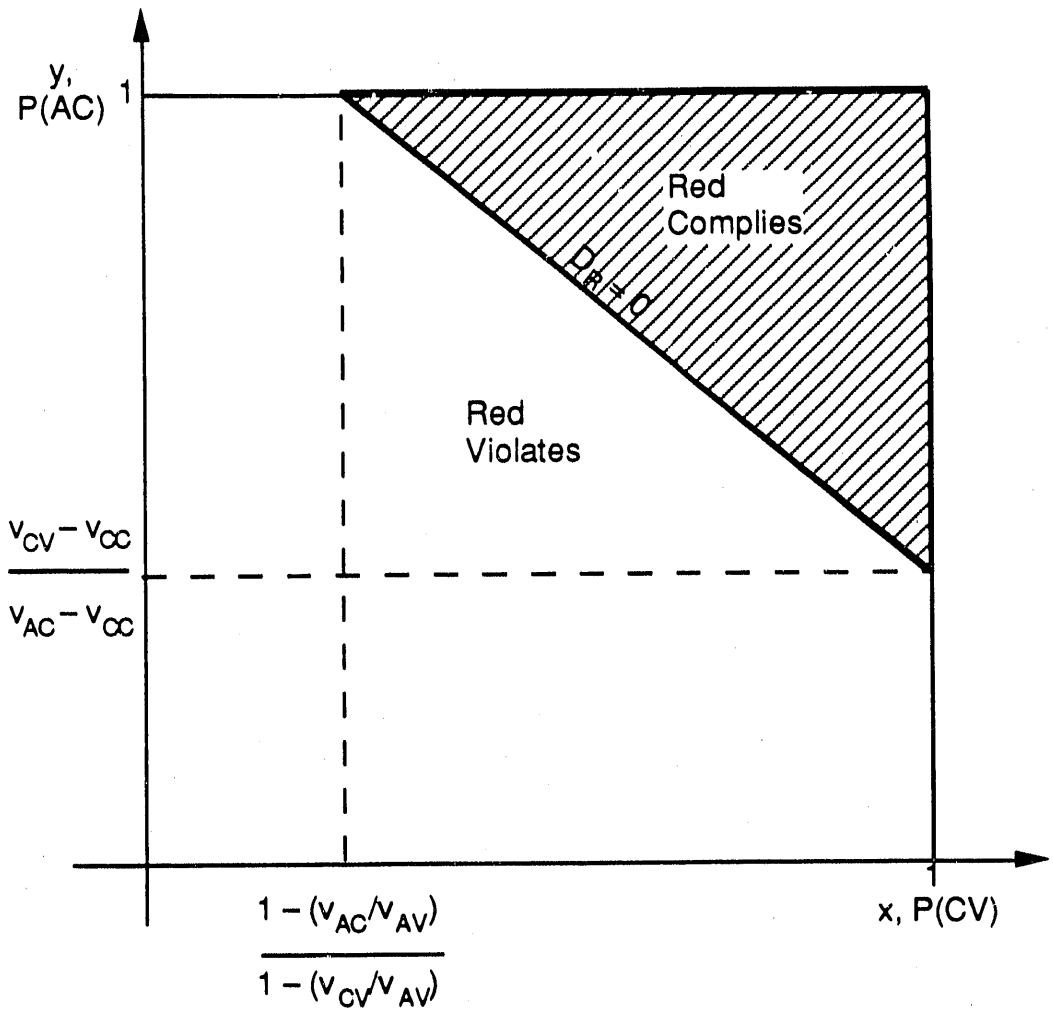

Fig. 4

Deterrence region in the detection probability plane

If Red is deterred (i.e. if x, y is in the shaded region in Fig. 3), Blue's expected utility will be

$$EV_{BIC} = u_{CC} + (u_{AC} - u_{CC})\, y \ , \quad \text{for } D_R > 0 \ . \tag{2-8}$$

Because Blue prefers to accept compliance rather than to challenge it, i.e.,

$$u_{AC} > u_{CC} \ ,$$

Blue's payoff is a linear, increasing function of y. Note also that when Red is violating, Blue's payoff is

$$EV_{B|V} = u_{AV} + (u_{CV} - u_{AV})\,x \quad , \tag{2-9}$$

which is a linear, increasing function of x (a decreasing function of y), provided that Blue (quite understandably) prefers to challenge rather than accept a violation,

$$u_{CV} > u_{AV} \quad .$$

Now turn to Fig. 5 to consider Blue's design problem. First, recall that Blue's payoff when Red is complying (Eqn. (2-3)) is a strictly increasing function of y. Next, observe that the strength of deterrence is proportional to the distance from the line $D_R = 0$. Now, Blue can limit himself to consider only points both in the deterrence region ($D_R > 0$) and on the line L. (We'll show later that Blue always prefers compliance to violation.) Blue can eliminate further points by noting that lines D_R = constant > 0 intersect L in general at two points, p and p′, (when there is an intersection). One of which (p) has a higher payoff to Blue than the other (p′). Thus the lower payoff point p′ can be eliminated, and Blue needs only consider the points p on the line L_0, the heavily lined "dominant subset" of L. The right end point of L_0, p_2, is the point of maximum (excess) deterrence, while the left end point (at $D_R = 0^+$), p_1, is the point of maximum expected value to Blue (while preserving deterrence). This design is the solution to the formal optimization problem:

$$\underset{\Delta}{\text{maximize}}\; EV_B(\Delta) \tag{2-10}$$

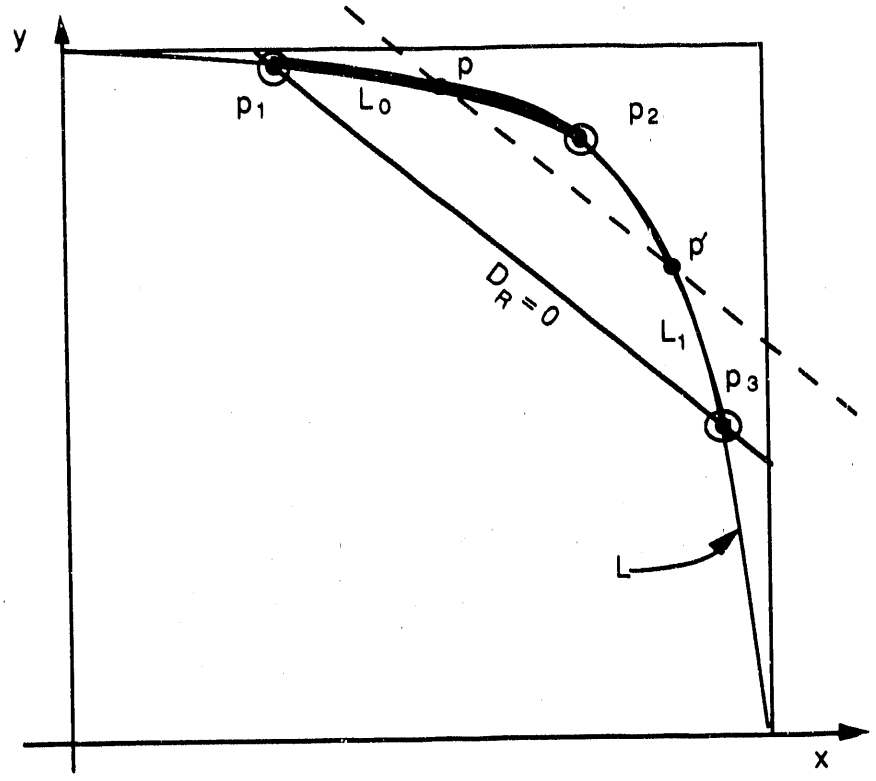

Fig. 5

Verification system
detection characteristic L and design choices

Figure 6 plots Blue's expected utility EV_B as a function of y. Below the point y_1 (corresponding to p_1 in Fig. 5), Red is deterred and $EV_B = EV_{BIC}$. Beyond y_1 (outside the deterrence region), Red violates, and $EV_B = EV_{BIV}$. The maximum of EV_B occurs at y_1^-, just inside the deterrence region. This design (p_1 in Fig. 5) is similar to the "optimal inducement strategy" of Brams and Kilgour [3].* In order to induce Red compliance, Blue must make a credible commitment to employ the specified detection probabilities (x, y). This pair of (Red,

* The model in Brams and Kilgour differs somewhat from ours. They assume a fixed detection system (i.e., single value of (x, y)) for Blue, but with modifications in the effective values of x and y produced by a "mixed strategy" of sometimes (with specified probability) not consulting the detector and simply announcing "accept" or "challenge". This option produces a set of effective x-y design possibilities that differ from ours, and in fact, because of the convexity of the detection characteristic are always inferior to ours. Brams and Kilgour also admit the possibility of mixed strategies by Red, i.e., complying or not, with some probability $s \neq 0, 1$. We, on the other hand, consider Red (like Blue) to have only pure strategies available, i.e., simply comply or violate. (Blue's "pure strategy" is to select a value of (x, y), always to "consult" the detector, and always to act on the reading of the detector as instructed by the specified value of (x, y).)

Blue) decisions is thus not a Nash equilibrium. In a Nash equilibrium neither side is motivated to attempt a unilateral change. (E.g., having induced Red compliance, Blue in fact would be tempted to change the detection probabilities to (0, 1), and thus achieve his very best possible pay-off. We assume here that Red has complete information, and that Blue must follow through on his original design.)

The shape of the expected value curve in Fig. 6 gives some support to designs with "excess deterrence". Unlike a typical performance measure with an internal, local maximum, this one is

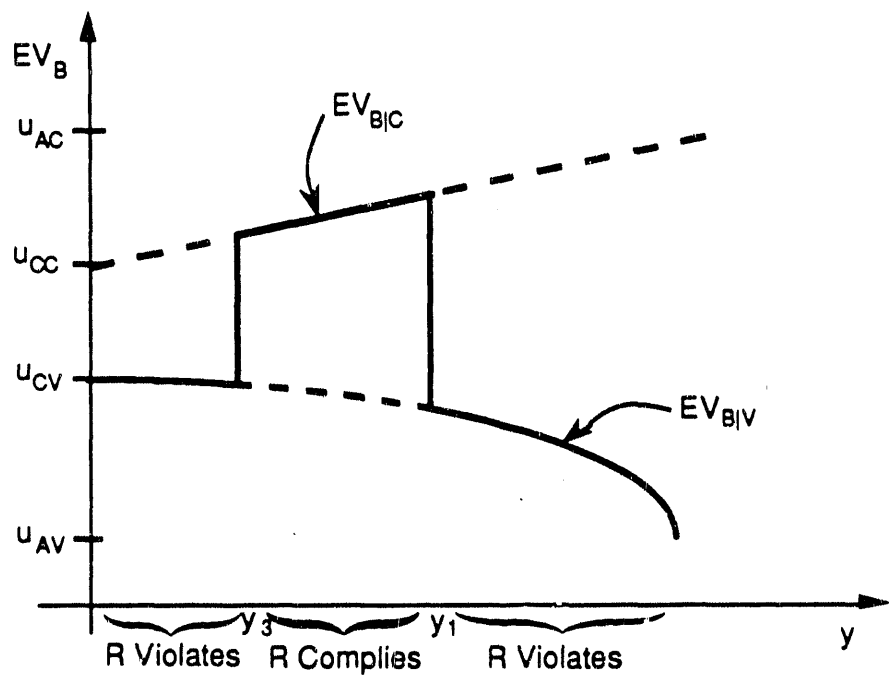

Fig. 6

Blue's expected payoff as a function of y (solid line). The maximum payoff occurs at $y = y_1$ (corresponding to the point p_1 in Fig. 5).

not smooth. Because a discontinuity occurs at the peak, small errors in selecting the value of the design-variable y can cause very large losses in performance. Hence some conservativeness, i.e., moving toward the left of y, in Fig. 6, seems quite natural. Such conservative designs as (x_2, y_2) were explored in [1] in some detail. Here we instead focus on the maximum-payoff designs (x_1, y_1). This changed focus is justified by introducing Red Types to explicitly account for Blue's

uncertainty in Red's utilities (in contrast to the ad hoc robustness achieved by maximum-deterrence designs).

A more cautious analysis of Blue's decision might fold in another factor: the cost of failure of deterrence. When deterrence fails, recall from (2-9) that EV_B is given by

$$EV_{B|V} = u_{AV} + (u_{CV} - u_{AV})\, x \quad .$$

Now, a Blue more pessimistic than in the previous design based on the criterion in (2-10), could adopt the viewpoint of balancing deterrence against payoff given deterrence failure. Such a perspective would lead to a consideration of only the points on the line L_1 in Fig. 3, i.e., the points between the maximum-deterrence point p_2 and the maximum payoff (given deterrence failure) p_3. These points are the mirror image of the points on L_0 that are selected by the (optimistic) Blue who acts as if deterrence succeeds. Formally, these designs on L_1 are obtained as solutions to a class of optimization problems that could be defined analogously to (2-10), with $EV_{B|V}$ substituted for EV_B and with the addition of an explicit constraint that $D_R > 0$.

Note also that, in the light of the meaning we attach to the segments L_0 and L_1, the point (p_2) is a kind of compromise between pessimism and optimism with regard to deterrence failure. However also note that because of the geometry of the optimality condition, for small movements from p_2, Blue can always gain significantly in payoff (either $EV_{B|C}$ or $EV_{B|V}$) at the cost of negligible loss of deterrence. Thus in "practical" designs one might want to deviate from exact maximum-deterrence, either to one side or the other, depending on whether one was "optimistic" or "pessimistic". (And if one can't decide which, then presumably stay at p_2.)

An objection to "pessimistic" designs between p_2 and p_3 is that they are inconsistent with the meaning of being in the deterrence region. On the one hand, our analysis says that Red will comply; on the other hand we design as if Red violated. One defense of designs approaching p_3 is the following: Consider design pairs approaching $D_R = 0$ on L_0 and L_1 respectively, i.e., points approaching p_1 and p_3. Imagine these two designs have the same level of deterrence. As the

deterrence goes to zero, the decision of Red becomes indeterminate, i.e., at exactly p_1 and p_3 we can't say what Red will do. Clearly, here it is optimistic to be at (or near) p_1 and pessimistic to be at (or near) p_3.

Yet another approach, that blends the two above (optimism and pessimism) might associate with values of D_R a probability of a violation such that this probability increases with decreasing D_R. This would provide a natural way to fold $EV_{B|V}$ into the overall analysis. Section 3 explores a different approach to this issue of uncertainty in Red's actions by considering different Red "types" that occur with specified probabilities but each of which behaves deterministically.

We can display the differences between "optimistic" and "pessimistic" designs by mapping the set of design candidates L_0 and L_1 from the probability (x, y) plane to the value (D_R, EV_B) plane, as shown in Fig. 7. This gives a clearer picture of how values are traded off for various designs. The solid line plots D_R versus $EV_{B|C}$ (i.e., given actual Red compliance) within the

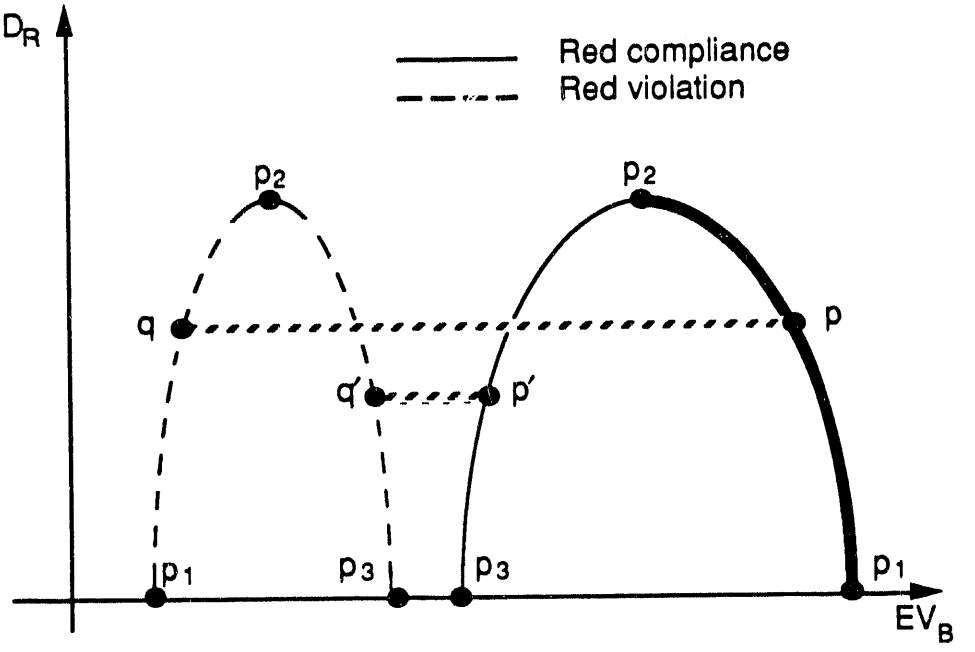

Fig. 7

Optimal designs in the D_R – EV_B value plane
p – q: "optimistic" design
p′ – q′: "pessimistic" design

deterrence region, between the points (p_1 and p_3) of Fig. 5. The heavy solid line is the image of the line of dominating solutions L_0, while the light solid line is the image of L_1. (Note that dominance shows up here explicitly, by a comparison of the values of EV_B for two points on the curve at equal values of D_R.) Similarly, the dashed line plots D_R versus $EV_{B|V}$ (i.e., given actual Red violation).

Suppose Blue picks an "optimistic" design on L_0 that produces a value at p (Fig. 7) if Red complies. If instead Red violates, then the resulting payoff point is the lower one at q. The loss of value from p to q is the cost of the loss of deterrence. In contrast, consider the "pessimistic" design at point p′ where the expected payoff to Blue is less if Red complies, but greater, at q′ (than the alternative q), if Red violates.

It is shown in the Appendix that Blue always prefers compliance to violation when

$$\begin{array}{ll} & u_{AC} > u_{CV} \\ \text{and} & u_{CC} > u_{CV} \end{array} , \qquad (2\text{-}12)$$

the ordering that we have assumed here. (This result justifies our concentration solely on designs that achieve deterrence.) The Appendix also states a weaker condition for this result. Figure 7 shows this usual case, where there is no intersection between the solid (compliance) curve and the dashed (violation) curve.

To summarize the results of this section, we have defined "good" verification system designs as those that tend to deter Red violations of the treaty. Of these systems, there is one that "maximizes deterrence". We have further defined, from among the set of deterring designs those two (at p_1 and p_3) that achieve maximum expected utility for Blue, but at the expense of marginal deterrence. The design at p_1 is based on the optimistic assumption that deterrence works, while the design at p_3 on the pessimistic assumption that it doesn't.

In the next section we look at the important question of what happens to these designs when Blue has uncertainty about Red's utilities.

3. Analysis of Types

We consider the possibility of three different Red types, distinguished from each other qualitatively by their utilities:

Deterrable Type (R_D)

R_D is a deterrable violator. There is at least one value of detection probability (x, y) such that R_D prefers compliance to violation, and in general there are other values of (x, y) such that R_D is not deterred. As shown in Fig. 8 there is an intersection between the deterrence line $D_R = 0$ and the detection characteristic. R_D is the type described in Section 2.

Violator Type (R_V)

This type is undeterrable. R_V's utilities are such that deterrence is negative for all available values of detection probabilities (x, y). The deterrence line $D_R = 0$ lies above and to the right of the detection characteristic (see Fig. 8). Consequently, R_V always violates, no matter what detection probabilities (x, y) Blue selects.

Complier Type (R_C)

R_C always complies, whatever the value of detection probabilities. The deterrence line lies below and to the left of the detection characteristic.

The utility ordering for R_D was given in Eqn. (2-1) and is repeated in Fig. 9. Also shown there is the ordering for the two new types. The utilities of the Violator, R_V, satisfy either

$$v_{AV} > v_{CV} > v_{AC} > v_{CC} \qquad (3\text{-}1)$$

or else they are those of R_D with the magnitudes such that there is no intersection with the detection characteristic. In Fig. 8 the first case is labeled "R_V", the second "R_V'"; R_V' orderings are not

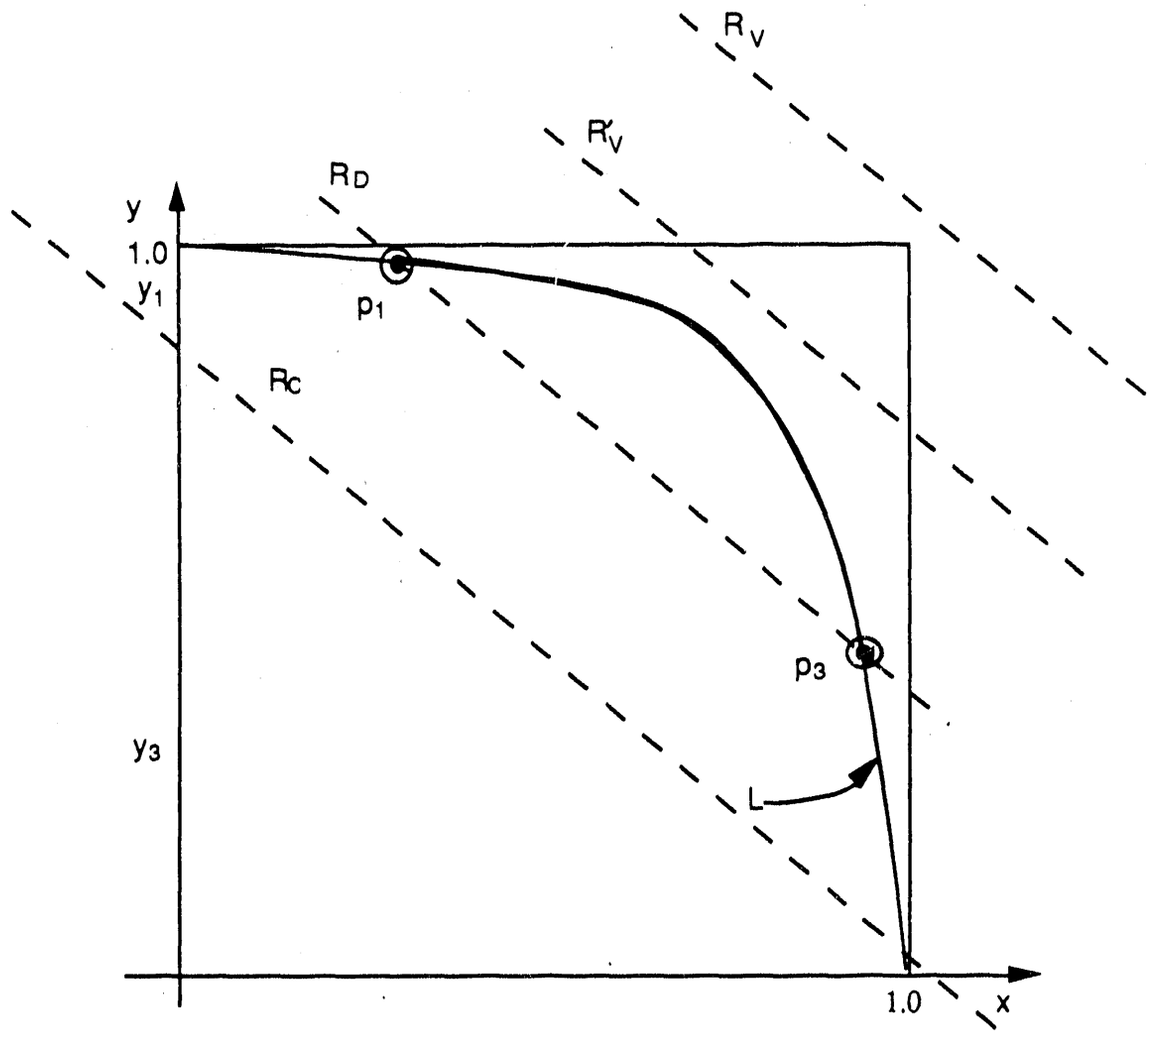

Fig. 8

Deterrence lines $D_R = 0$ (dashed), for three different Red types: deterrable (R_D), violator (R_V), and complier (R_C).

repeated in Fig. 9. It is easy to show that no deterrence region exists when (3-1) holds. The utilities of the Complier satisfy

$$v_{AC} > v_{CC} > v_{AV} > v_{CV} \tag{3-2}$$

and are also shown in Fig. 9. In (3-2) the ordering $v_{CC} > v_{AV}$ could be relaxed; this ordering might be termed a strong preference for compliance. Under this ordering, deterrence is always positive, and hence this complier type in fact always does comply. For a complier with weak

preference, $v_{CC} < v_{AV}$; in this case it's easy to show that although there are regions in the x-y plane where $D_R < 0$, they always lie below the diagonal between (0,1) and (1,0); hence any "non-perverse" detection system will also deter this type. In conclusion then, Complier types always comply, even when the middle utility pair ordering is reversed.

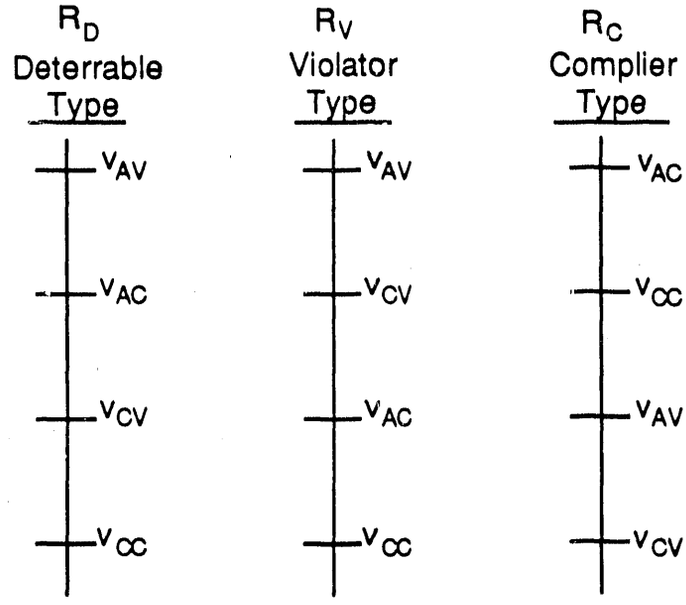

Fig. 9

Utility ordering for the three Red types

Note that the two new types in some sense *do not play the game*. R_C always complies, whatever Blue does. R_D never complies, either because Blue's particular verification system isn't good enough to induce him to, or because no verification system would be good enough.

We now assume that Blue assigns probabilities to the existence of each Red type, as shown in Fig. 10. The events of the occurrence of a particular Red type are mutually exclusive and exhaustive. Blue's probabilities over types reflects his uncertainty as to Red's utilities v. While Blue could alternately assign a probability distribution to Red's utilities, the more important issue to Blue is simply Red's type.* Fig. 11 displays Red types in the p-q plane. The "pure" types occur at the vertices of the triangular region, which defines all possible mixtures of types that can occur. Each assessment by Blue of his probability distribution of types is represented by a point in this region.

* The exact utilities of R_D do of course matter. Ref. [1] works out a case where there are sub-types of R_D. To simplify the present analysis, we consider only one type R_D.

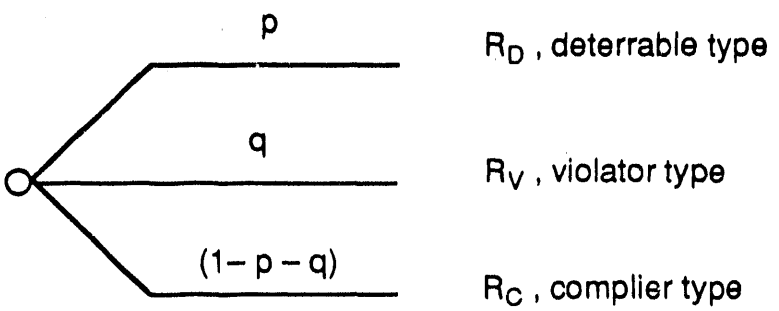

Fig. 10

Blue's probability distribution over Red types

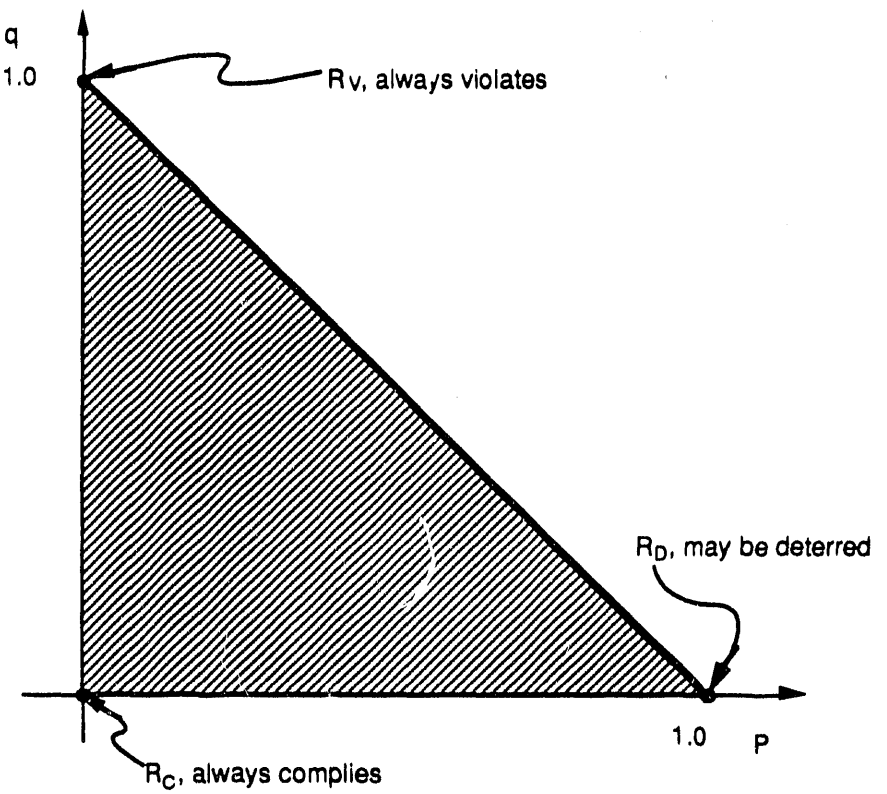

Fig. 11

Red types in the q-p plane

From Fig. 8 and 10 we can now compute Blue's expected utility with the added two types:

$$EV_B = \begin{cases} (q+p)\,EV_{B|V} + (1-q-p)\,EV_{B|C}\ , & y > y_1 \\ q\,EV_{B|V} + (1-q)\,EV_{B|C}\ , & y_1 > y > y_3 \\ (q+p)\,EV_{B|V} + (1-q-p)\,EV_{B|C}\ , & y_3 > y \end{cases} \tag{3-3}$$

To simplify the subsequent analysis, and without loss of generality,* we take $u_{AC} = 1$, $u_{AV} = 0$. Then from (2-8) and (2-9) we obtain

$$EV_{B|C} = u_{CC} + (1 - u_{CC})\,y \tag{3-4}$$

and $$EV_{B|V} = u_{CV}\,x\ . \tag{3-5}$$

Also recall from Section 2 the two alternative designs at $p_1\ (y = y_1^-)$ and $p_3\ (y = y_3^+)$, each optimum in a specific sense for R_D alone. We can evaluate Blue's expected utility for each of these designs from (3-3). Define each payoff as

$$EV_{B1}^* \equiv EV_{B|y = y_1^-} \tag{3-6}$$

$$EV_{B3}^* \equiv EV_{B|y = y_3^+} \tag{3-7}$$

Then

$$\begin{aligned} EV_{B1}^* &= q\,EV_{B|V|_{y=y_1}} + (1-q)\,EV_{B|C|_{y=y_1}} \\ &= q\,u_{CV}\,x_1 + (1-q)[u_{CC} + (1-u_{CC})\,y_1] \end{aligned} \tag{3-8}$$

and $$EV_{B3}^* = q\,u_{CV}\,x_3 + (1-q)[u_{CC} + (1-u_{CC})\,y_3] \tag{3-9}$$

when $q = 0$, these payoffs reduce to these given by (3-4) and (3-5) and are the best that Blue can achieve in the senses defined in Section 2. We want to determine now, for $q \neq 0$, $p \neq 1$, under what conditions (and with what verification system designs $y \neq y_1, y_3$) Blue can obtain a better payoff than (3-8) or (3-9).

---

* If a set of treaty constraints were being considered, as in [1], we could not make this simplification.

We will now put some structure into these questions, in order to simplify their analysis and to sharpen the interpretation of the answers. We pose the following:

A.) When is it better to "not verify at all" than to be at y_1 (y_3)?

B.) When is it strictly best to "not verify at all"?

C.) When is it strictly best to be at neither y_1/y_3 nor at a "no-verify" condition?

Comments:

i) "No-verification" is taken here to mean either the design $x = 0$, $y = 1$ (where Blue always accepts) or the design $x = 1$, $y = 0$ (Blue always challenges). Although they are points on the detection characteristic, occurring as limiting cases as $\Delta \rightarrow \infty$ and 0, respectively, they are in fact fixed responses and do not truly involve a verification measurement with real information. These two designs are special cases of an entire "no-verification detection characteristic" which consists of a diagonal straight line between (0, 1) and (1, 0). It's easy to show that we can ignore all these points except the end ones. First, R_D cannot be deterred by any of these designs. Second, along this line, Blue's expected payoff will always be a linear function of y; hence, except for uninteresting special cases the maximum will occur at the end points of the interval.

ii) We will take $x = 0$, $y = 1$ as the reference no-verification system for the design at y_1, and $x = 1$, $y = 0$ as the reference for the design at y_3. These are plausible pairings because of their proximity. They also make sense because the y_1 design is optimistic in expecting compliance, while y_3 is pessimistic in fearing violation.

iii) "No-verification" at (0, 1) and (1, 0) have two quite different implications for a prospective treaty. The state (0, 1) is quite consistent with the existence of a treaty

— Blue simply always announces Red compliance. However, the state (1, 0) is antithetical to the idea of a treaty, as Blue inevitably claims Red violates. Thus, "no-verification" at (1, 0) is equivalent to saying there should be no treaty.

iv) We will devote most of our attention to the design point p_1, only briefly considering "pessimistic" designs at p_3.

We now consider each of these questions in turn.

A. When is it better not to verify at all than to be at y_1 (y_3)?

A.1) The design at y_1. To compare the two alternatives ($y = y_1$, $y = 1$) define the function

$$\Phi \equiv EV^*_{B_1} - EV_B(1) \tag{3-10}$$

From (3-3), (3-4), and (3-5),

$$EV_B(1) = 1 - q - p \quad , \tag{3-11}$$

and with $EV^*_{B_1}$ from (3-8), we obtain

$$\Phi \equiv q\, u_{CV}\, x_1 + (1-q)[u_{CC} + (1-u_{CC})\, y_1] - (1-q-p) \tag{3-12}$$

The answer to question A. is then simply that we should

$$\begin{cases} \text{verify } (y = y_1) \text{ when } \Phi > 0 \\ \text{don't verify } (y = 1) \text{ when } \Phi < 0 \end{cases} \tag{3-13}$$

We can easily explore this relationship in more detail in some special cases:

Special case: q = 0 (no pure Violator type R_V)

This case consists of combinations of types on the base of triangle in Fig. 11, and ranges from R_C to R_D as p goes from 0 to 1. Φ becomes

$$\Phi = u_{CC} + (1 - u_{CC})\, y_1 - (1 - p). \quad . \tag{3-14}$$

As a result we have the decision rule

$$\boxed{\text{If } \frac{(1 - u_{CC})(1 - y_1)}{p} \begin{cases} <1, \text{ then verify at } y = y_1 \\ >1, \text{ then don't verify } (y = 1) \end{cases}} \tag{3-15}$$

Thus, we tend <u>not</u> to verify as $p \rightarrow 0$ (the likelihood of a complier increases). For sufficiently small p it is always better not to verify at all. Also, for y_1 sufficiently close to 1 (a sufficiently good probability of correctly accepting compliance by the reference verification system), it becomes best for Blue to verify at y_1.*

Fig. 12 shows how the verify/don't verify decision depends on values of p and y_1, for several different values of u_{CC}. Note that the <u>value</u> of <u>not</u> verifying (compared to verifying) depends on the distance below the diagonal line; thus, Blue gets a greater return from not verifying at point a than not verifying at point b. It is not surprising that when Red is deterred by verification systems that approach perfection in detecting compliance, then it makes less difference whether a verification system is employed or not.

Special case: p = 0 (no pure deterrable type R_D)

This case corresponds to the left axis of the triangular region in Fig. 11 and consists of mixtures of Compliers (R_C) and Violators (R_V). From (3-12),

$$\Phi = q\, u_{CV}\, x_1 + (1 - q)\, [u_{CC} + (1 - u_{CC})\, y_1] - (1 - q) \tag{3-16}$$

---

* Note that y_1 is in turn a function of Red's utilities and the verification detection characteristic "quality". It is easy to show that as verification quality increases, i.e., as the characteristic approaches the square with vertex (1, 1), that y_1 increases; hence increased detection quality also tends to drive towards the condition of verification. (It makes sense that the verification alternative becomes more attractive as detection capability gets better.)

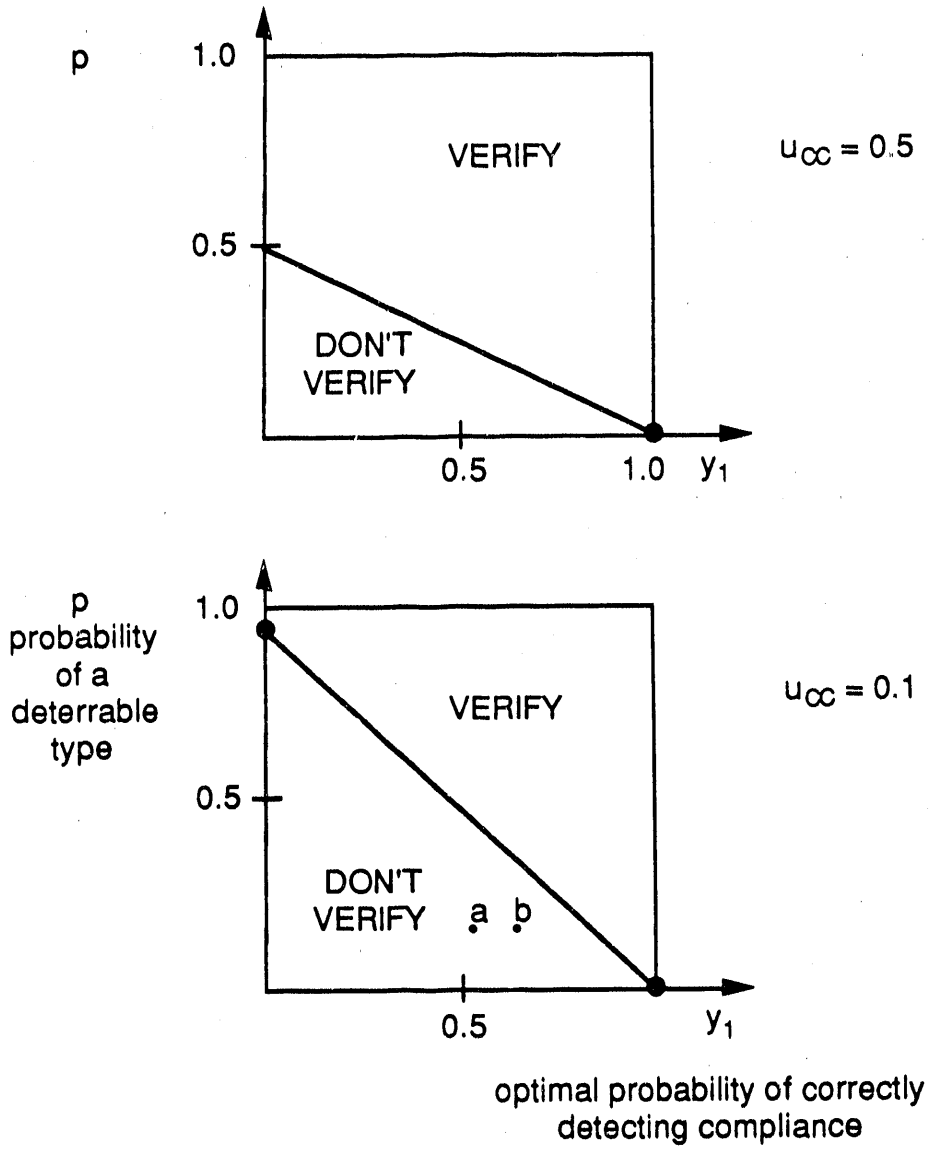

Fig. 12

Verification regions in p-y_1 plane

Hence we have the decision rule

$$\text{If } \frac{q}{1-q} u_{CV} x_1 + u_{CC} + (1 - u_{CC}) y_1 \begin{cases} >1, \text{ then verify} \\ <1, \text{ then don't verify} \end{cases} \tag{3-17}$$

Thus Blue will tend to verify as $q \rightarrow 1$, i.e., as Blue becomes more certain that Red is a Violator.

A.2 The design at y_3.

$$\Phi_3 \equiv EV^*_{B_3} - EV_B(0) \tag{3-18}$$

From (3-3),

$$EV_B(0) = (q + p)\, EV_{B|V}(0) + (1 - q - p)\, EV_{B|C}(0)$$

which becomes, using (3-4, 5),

$$EV_B(0) = (q + p)\, u_{CV} + (1 - q - p)\, u_{CC} \tag{3-19}$$

From (3-9) and (3-19) then

$$\begin{aligned}\Phi_3 \equiv\ & q\, u_{CV}\, x_3 + (1 - q)[u_{CC} + (1 - u_{CC})\, y_3] \\ & -(q + p)\, u_{CV} - (1 - q - p)\, u_{CC}\end{aligned} \tag{3-20}$$

Special case: $q + p = 1$ (no pure Complier type R_C)

Φ_3 then becomes

$$\Phi_3 = (1 - p)\, u_{CV}\, x_3 + p\,[u_{CC} + (1 - u_{CC})\, y_3] - u_{CV}$$

giving the decision rule

$$\boxed{\text{If } \frac{(1 - p)\, u_{CV}\,(1 - x_3)}{p\,[u_{CC} + (1 - u_{CC})\, y_3]} \begin{cases} <1, \text{ then verify } (y = y_3) \\ >1, \text{ then don't verify } (y = 0) \end{cases}} \tag{3-21}$$

Note that as p varies from 1 to 0 (the likelihood of a deterrable type decreases), the decision in (3-21) switches from verification to no verification. And, as pointed out in comment iii) above, the implication of no-verification here is "no treaty".

In the remainder of the analysis we focus entirely on "optimistic" designs p_1.

B. When is it <u>strictly</u> best to "not verify"?

For simplicity we consider the case where there are no Compliers, q = 0.

To answer question B. we need to undertake a broader search for possible maxima. In addition to the possibility of local maxima at $y = y_1$ and $y = 1$, we now consider the possibility of local, interior <u>smooth</u> maxima. We can find necessary conditions for the existence of such maxima by considering EV_B over the interval $y_1 < y \leq 1$:

$$EV_B = p\, u_{CV}\, x + (1-p)[u_{CC} + (1-u_{CC})\, y] \quad , \qquad y_1 \leq y \leq 1 \tag{3-22}$$

(Since EV_B is always a linear function of y for $y_3 < y < y_1$, a local maximum cannot occur there, and we can ignore this interval.)

From (3-22) we obtain the result that

$$\frac{dEV_B}{dy} = 0$$

for

$$\frac{dy}{dx} = -\frac{p}{1-p}\,\frac{u_{CV}}{1-u_{CC}} \quad . \tag{3-23}$$

It's straightforward to show that $\frac{d^2y}{dx^2} < 0$ for normal detection characteristics [1]. Hence <u>solutions to (3-23) for $y_1 < y \leq 1$ will be local maxima</u>.* Since $\frac{dy}{dx} < 0$ and $0 < (p/1\text{-}p) < \infty$, for $0 < p < 1$, it is always possible to find permissible values of p for which there are solutions to (3-21). However, even when a <u>local</u> maximum exists, it may not be a <u>global</u> one. Note that when y = 1 gives a larger value of EV_B than $y = y_1$, <u>and</u> condition (3-23) is satisfied, then the internal

* These solutions are "Bayes" in the sense that they maximize expected value for Blue based on Blue's prior probability of a Red violation, which in this case, on this interval $y_1 < y \leq 1$, is just p.

maximum is global. We return to the consideration of global, internal maxima in part C., but at this point we wish to limit ourselves to the narrower question of when "no verification" (y = 1) is better than either $y = y_1$ or the interior maximum defined by (3-21).

Since

$$\frac{dEV_B}{dy} = p\, u_{CV} \frac{dx}{dy} + (1-p)(1-u_{CC}) \tag{3-24}$$

and since $\frac{dx}{dy} \to -\infty$ as $y \to 1$ [1], clearly in general y = 1 cannot be a local maximum. (This is obviously consistent with the existence of at most one maximum (3-23) internal to the interval $y_1 < y \leq 1$.) There is, however, one special case for which y = 1 may be a global maximum. If $u_{CV} = 0$, then $\frac{dEV_B}{dy} > 0$ for $y_1 < y \leq 1$, and y = 1 is a local maximum. Eqn. (3-15) tells us when $EV_B(1) > EV_B(y_1)$, and thus we arrive at the general result:

When q = 0, it is strictly optimal for Blue not to verify at all (to set y = 1) if and only if

$$\frac{(1-u_{CC})(1-y_1)}{p} > 1 \tag{3-25}$$

and $u_{CV} = 0$.

Note that $u_{CV} = 0$ means that the cost of a challenged violation is equal to the cost of an accepted violation. This equality tends to hold to the extent that challenges have little effect on ameliorating damage resulting from a violation. It is only under this condition that it can be strictly optimal to not verify at all.

In the next sub-section we consider the remaining question, which relates to designs that are intermediate between standard ones ($y = y_1$) and no-verification ones (y = 1).

C. When is it strictly best to be neither at $y = y_1$ nor at a "no-verify" condition (y = 1)?

We continue to restrict ourselves to the case q = 0. Recall that satisfaction of the condition (3-23) is necessary and sufficient for a solution (x_0, y_0) to be a local maximum. Clearly, then, this value of (x_0, y_0) also yields a global maximum if in addition

$$p\,u_{CV}\,x_0 + (1-p)[u_{CC} + (1-u_{CC})\,y_0] > u_{CC} + (1-u_{CC})\,y_1 \qquad (3\text{-}26)$$

and $y_0 > y_1$,

i.e., if the payoff at the local maximum at y_0 is greater than at the local maximum at y_1, and if the y_0 maximum occurs ($y_0 > y_1$).

We can explore properties of such global maxima, $y_1 < y_0 < 1$, through an illustrative example. The family of functions

$$x^n + y^n = 1 \qquad (3\text{-}28)$$

are consistent with the required behavior of detection characteristics. This family in fact sweep through a full range of cases, from "no verification" for n = 1, to "perfect verification" for $n = \infty$. For a simple example, we take n = 2. Solution of (3-23) then gives

$$x_0 = \alpha\, y_0$$
$$y_0 = \left(\alpha^2 + 1\right)^{-\frac{1}{2}} \qquad (3\text{-}29)$$

where $\alpha = \dfrac{p\,u_{CV}}{(1-p)(1-u_{CC})}$.

Figure 13 plots y_0 and $EVB(y_0)$ from (3-29) and (3-22) for the case $u_{CC} = 0.7$, $u_{CV} = 0.3$. Note that as the payoff for being at y_0, $EV_B(y_0)$, increases (as $p \rightarrow 0$), so does the difference in payoff between being at y_0 and being at y = 1 go to zero. That is, the higher the payoff (and the higher the probability of compliance), the less distinction there is between being at y_0 and not verifying at all. Thus for high probability of compliance, the practical answer to question B. above is that it may not matter much whether one is at y_0 or at y = 1, even if $u_{CV} \neq 0$.

Figure 13 is not the whole picture. In order to determine whether y_0 is a global maximum, we need to make the tests (3-26). If we do this, for the case of Fig. 13, we obtain the results plotted in Fig. 14. Shown here as a function of y_1 is a probability p_C, defined as the value of p such that if $p \leq p_C$, there exists a global maxima at some $y_0 \geq y_{0C}$. Also plotted for reference is a quantity p_A, which is the critical value of p, from (3-15), based on the comparison of payoffs (in

part A) of being at $y = y_1$ versus no verification at all at $y = 1$. (The probability p_A was plotted in Fig. 12 for that comparison.) Thus, the region $p_A < p < p_C$ can be seen as defining those parameter combinations for which a deviation from $y = y_1$ to $y = y_0$ is recommended, but for which it is <u>not</u> preferable to not verify at all. Above the p_C curve it is always best to verify in the usual way at $y = y_1$.

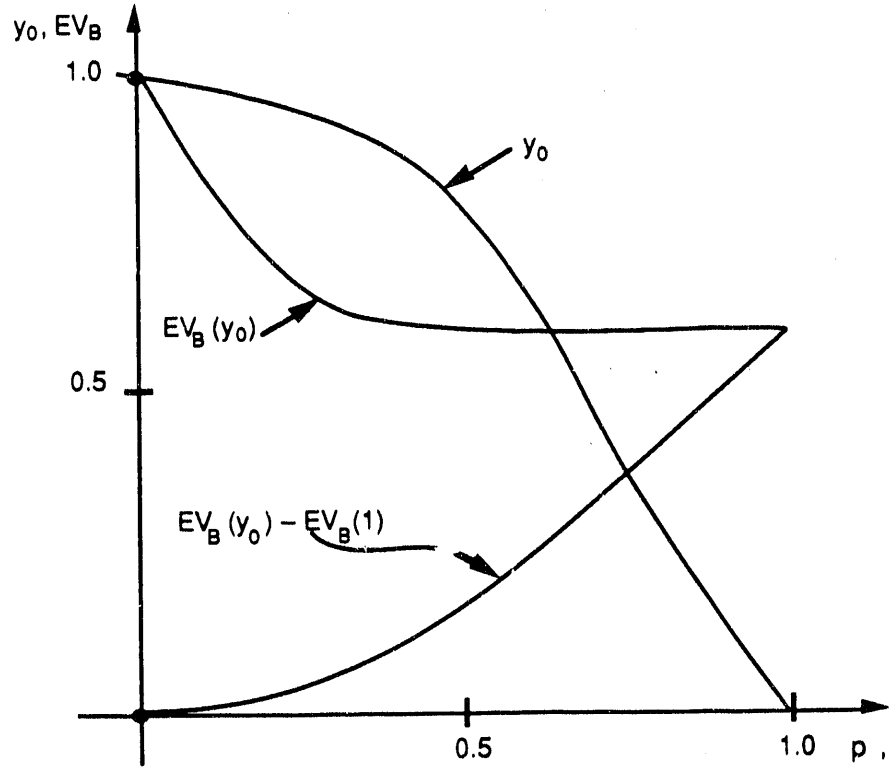

Fig. 13

Smooth maxima (solutions of 3-23) for illustrative example. Detection characteristic: $x^2 + y^2 = 1$; $u_{CC} = 0.7$, $u_{CV} = 0.3$.

In effect, the probability p_C marks a dividing line between the "game theory solution" y_1 and the "Bayes solution" y_0. For probabilities above p_C, the game theory solution dominates, while for probabilities below p_C, the Bayes solution y_0 is preferred.

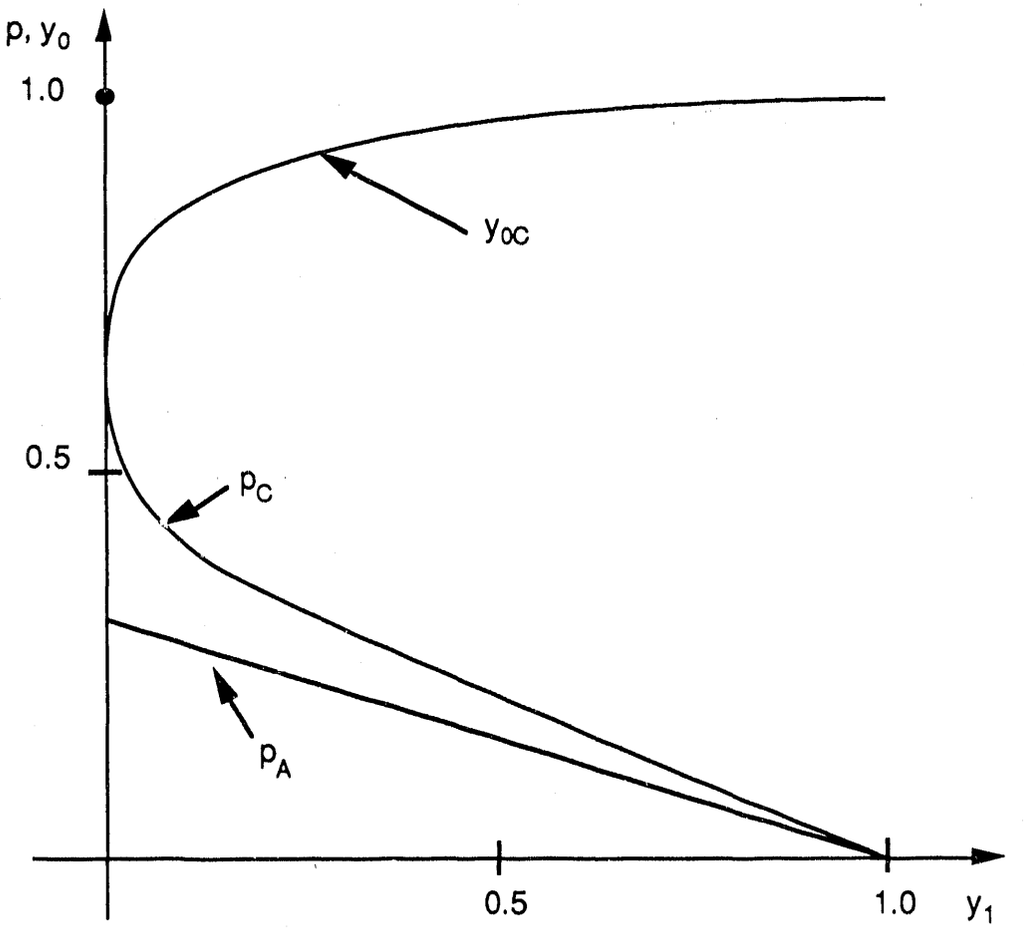

Fig. 14

Critical values of probability for the example. A design at some $y_0 > y_{0C}$ is strictly optimal for $p < p_C$. A design at $y = 1$ is better than one at y_1 for $p < p_A$. For $p > p_C$ it is best to be at $y = y_1$.

4. Conclusions

We've shown how a simple model can be used to analyze the performance of treaty verification systems when one party (Blue) is uncertain about the preferences of the other (Red). Criteria were derived that permit Blue to set optimal values of the detection threshold. Depending on system parameters and Blue's probability distribution over Red's type, the threshold is set either at the normal, game-theoretic value, or at a value that is the Bayes solution. In effect, our criteria tell Blue when to abandon the effort of deterrence, and to treat the other party as wholly characterized by prior probabilities of violation and compliance.

For limiting values of certain parameters, Blue's best policy tends towards one of no verification at all. This result, unsurprising from the perspective of the Bayes solution alone, is nevertheless of timely interest in the context of current developments in relations between the U.S. and the U.S.S.R.

Until recently it has been standard official wisdom that highly accurate verification methods are critical for a successful arms control treaty. Now, however, some [8] have changed their minds and argue that "arms control without agreement" (and hence without verification) are preferable to formal treaties with elaborate and strict verification. Others [9], not putting the case quite so starkly, describe an evolution from strict verification (and equally strict punishment of detected violations) to an emphasis on broader confidence building and dispute resolution.

A number of reasons are offered in support of the new views on verification. Among them are the costliness of obtaining high verification precision with new treaties like START and the political difficulties of enforcing "punishments" for perceived violations. However, it is clear that a dominant, underlying reason for the re-evaluation is the changed U.S. view of Soviet Union power and intentions. As Lewis Dunn points out [10], scenarios for significant Soviet military exploitation of violations of chemical and conventional weapons treaties are becoming less plausible. And in some fundamental sense, developing mutual confidence decreases the perceived

relevance and need for strict verification,* for preserving a regime of "hyper-vigilance" as Chayes and Chayes characterize our basically suspicious Cold-War view of the Soviet Union [9].

These arguments clearly reflect changing views of the role of treaty verification. In the present work we have tried to determine the extent to which existing models of verification can be used to address these new issues. We have been able to address both the "old" and "new" views of verification within the same analytical framework, and to do so in way that permits us in effect to weigh our uncertainty about the permanency of the "new" vs. the "old" Soviet Union.

Work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

* It is also true that an expansion of verification measures requiring interaction and cooperation can itself be a stimulus to developing mutual confidence. Once such confidence is firmly developed, however, the need for verification would appear to decline.

References

[1] S. Weissenberger, *Deterrence and the Design of Treaty Verification Systems*, Lawrence Livermore National Laboratory, UCRL-JC-105810DR, November 15, 1990.

[2] "A Decision Framework for Verifying Compliance with Test-Limitation Treaties", Lawrence Livermore National Laboratory, *Energy and Technology Review*, pp 15-26, March 1990.

[3] S. J. Brams and D. M. Kilgour, *Game Theory and National Security*, Basil Blackwell Inc., New York, 1988.

[4] M. J. Lippitz and R. S. Strait, *Deterrence of Arms Control Treaty Evasion by Suspect Site Inspections*, Lawrence Livermore National Laboratory, UCID-21818, November 1989.

[5] Donald Wittman, "Arms Control Verification and Other Games Involving Imperfect Detection", *American Political Science Review*, v. 83, n.3, pp. 923-945, September 1989.

[6] R. Avenhaus et. al. (eds.) *Modelling and Analysis in Arms Control*, Springer-Verlag, 1985.

[7] Eric Rasmussen, *Games and Information*, Basil Blackwell Inc., New York, 1989.

[8] Kenneth L. Adelman, "Why Verification is More Difficult (and Less Important)", in *International Security*, v. 14, **4**, (Spring 1990), pp 141-146.

[9] Antonia Handler Chayes and Abram Chayes, *From Law Enforcement to Dispute Settlement*, **op cit**, pp 147-164.

Appendix

Value of Deterrence to Blue

Define the value of deterrence to Blue, D_B, to be

$$D_B \equiv EV_{B|C} - EV_{B|V} \quad , \tag{A-1}$$

the difference in expected payoff to Blue between compliance and violation. (D_B presents an upper limit on the payment Blue would be willing to give to Red in order to guarantee Red compliance, just as D_R gives a lower bound on the payment that Red would require in order to violate.) Substituting in (A-1) from Fig. 1 we get

$$D_B = -(u_{CV} - u_{AV})\,x + (u_{AC} - u_{CC})\,y + (u_{CC} - u_{AV}) \tag{A-2}$$

Fig. A-1 shows regions of positive and negative D_B in the x-y plane. By inspection, we see that D_B is positive for all x, y in the unit square if we constrain the intersection point of the line $D_B = 0$ and the line $x = 1$ to occur for negative y. Such an intersection occurs if

$$u_{CC} > u_{CV} \quad , \tag{A-3}$$

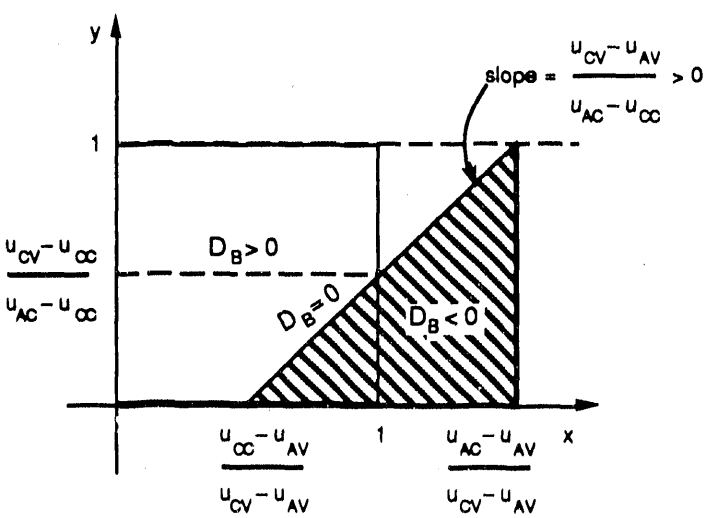

Fig. A-1

Region of positive D_B, where Blue prefers that Red complies.

in addition to the conditions already stated in (2-8) and (2-10), i.e.,

$$u_{AC} > u_{CC} \tag{A-4}$$

and $$u_{CV} > u_{AV} \ . \tag{A-5}$$

The new condition (A-3) is reasonable: Blue prefers challenged compliance to challenged violation.

We can also look for more limited conditions under which $D_B > 0$ throughout only the deterrence region $D_R > 0$. A comparison of Fig. A-1 and Fig. 2 gives

$$\frac{v_{CV}}{v_{AV}} > \frac{u_{CV} - u_{CC}}{u_{AC} - u_{CC}} \ . \tag{A-6}$$

To summarize, when acting only on the information from the verification system, Blue will always prefer that Red comply, for all values of x and y that deter Red, provided that either (A-3) or the weaker condition (A-6) holds.

END

DATE FILMED

05 / 08 / 91