ECONOMIC RESEARCH REPORTS

PUTTING THE OTHER SIDE "ON NOTICE" CAN INDUCE COMPLIANCE IN ARMS CONTROL

BY

D. Marc Kilgour
and
Steven J. Brams

RR # 92-07 **February, 1992**

C. V. STARR CENTER FOR APPLIED ECONOMICS

NEW YORK UNIVERSITY
FACULTY OF ARTS AND SCIENCE
DEPARTMENT OF ECONOMICS
WASHINGTON SQUARE
NEW YORK, N.Y. 10003

Journal of Conflict Resolution (forthcoming 1992)

Putting the Other Side "On Notice" Can Induce Compliance in Arms Control

D. Marc Kilgour
Department of Mathematics
Wilfrid Laurier University
Waterloo, Ontario N2L 3C5

Steven J. Brams
Department of Politics
New York University
New York, NY 10003

*The authors gratefully acknowledge the valuable comments of the editor and referees. Steven J. Brams gratefully acknowledges the support of the C. V. Starr Center for Applied Economics at New York University.

Abstract

Arms-control inspection is modeled by two games, one played simultaneously and one sequentially, between an inspector (O) and an inspectee (E). In each game,

- E may choose to comply with or violate an arms-control agreement;
- O may choose to inspect, or not, for a possible violation by E.

Besides various costs and benefits, the parameters of the games include the conditional probability that a violation will be detected if there is an inspection, reflecting the uncertainty of inspection.

In the simultaneous game, O and E make simultaneous choices. Because none of the three possible equilibria involves certain compliance by E, O is not always able to deter E from violating an agreement. In the sequential game, by contrast, O, by announcing in advance an inspection strategy and credibly commiting itself to carrying it out, can, with certainty, deter E from violating, which in general leads to an equilibrium in the sequential form Pareto-superior to that in the simultaneous form. Thus, there are evident benefits for both O and E when O "moves" first, given that its detection probability is above a certain threshold. Policy implications of this finding, especially in regional conflicts today, are briefly discussed.

Putting the Other Side "On Notice" Can Induce Compliance in Arms Control

1. Introduction

In spite of the upheavals in Eastern Europe, the former Soviet Union, and other former dictatorships, it is important to bear in mind that there still exist large arsenals of both conventional and nuclear weapons around the world. If arms control is to be successful, it must include procedures for verifying arms-control treaties, including those enforced by reginal and international organiazations. In this paper, we show that if a player in a two-person game publicly reveals in advance and credibly commits itself to an inspection strategy, it may not only help its adversary, who is thereby better informed, but, paradoxically, it may also help itself by giving away this information.

Such a consequence will not be surprising to game theorists, who are well aware of the paradoxical consequences that information in games may have through the signals it sends. As will become evident, we assume that the signals have credibility--they are backed up by actions--and are not merely "cheap talk."

In the context of arms control, one may think of an inspectee (E) as being better informed when it knows that, if it cheats, it will get caught by an inspector (O) with a certain probability than when it knows only that O will act "optimally." Without knowing the probability of getting caught, E might think that it could get away with at least some cheating when, if it possessed more information, it would know how much.

We originally developed our model with the superpower conflict in mind, thinking that the Intermediate Nuclear Forces (INF) treaty of 1988, which provides for specific numbers of annual inspections at agreed-upon sites, could become the model for future arms-control agreements. But with the disintegration of the Soviet Union and its replacement by a commonwealth, superpower treaties of this kind may have been overtaken by events.

Now we believe that other conflicts--like the ethnic strife in Yugoslavia, other Eastern European countries, and the former Soviet republics, or the continuing imbroglio in the Middle East--probably pose greater dangers for world peace. As a consequence, these arenas may become the most important future targets for arms-control agreements. With conflicts in these places as much as the former East-West conflict in mind (although the latter conflict may have subsided, the problems of how arms reductions will be made and verified remains), we next offer an overview of our model and results.

2. Overview

We shall define two different games that model arms-control inspection situations involving O and E. In each game,

- E may choose to comply with or violate an arms-control agreement;
- O may choose to inspect, or not, for a possible violation by E.

We assume that these games are played repeatedly--in many times and many places--and in each occurrence O may inspect for a violation by E but is not required to do.

In both games we shall analyze, we postulate as parameters the various costs and benefits that the players may receive at the different possible outcomes. We also postulate another parameter--the conditional probability that if E violates and O inspects, O will detect the violation. Thus, there is uncertainty that a violation that does occur will be uncovered, even if there is an inspection.[1]

The first game we analyze is specified by a game tree. In this game, which we call simultaneous, O and E make simultaneous choices. (If not simultaneous, we assume that each player chooses without prior knowledge of the other player's choice.) Depending on the values of the parameters, three cases, each with a unique Nash equilibrium--two in pure strategies and one in mixed strategies--can occur. Because none of the games involves certain compliance by E, O is not always able to deter E from violating the agreement.

In the second game, which we call sequential, O has the first move: to announce a probability p of inspecting for a violation, which E is then committed to carrying out (in ways to be specified later). We assume that O announces p before E makes its strategy choice. Thus, E decides to comply or violate with full knowledge of p but without knowledge of whether or not there will be an inspection in the actual play of the game. As in the simultaneous game, there is uncertainty as to whether, if E violates and O inspects, the violation will be uncovered.

In the sequential game, if the conditional probability of detection is above a certain threshold, there is a p that O can announce that will, with certainty, deter E from violating. Moreover, under certain conditions, the resulting outcome is a subgame-perfect ε-equilibrium that is Pareto-superior to the Nash equilibrium in the corresponding simultaneous game, which involves at least some violations (including possibly successful ones).

Thus, both players may prefer to play a game in the sequential form because of its Pareto-superior ε-equilibrium. This outcome, however, does not maximize the players' payoffs: O would prefer not to bear the costs of any inspections if E would comply. However, the latter outcome, though also Pareto-superior, is not a Nash equilibrium: O must pay some costs (e.g., the costs of carrying out occasional inspections) to deter cheating in equilibrium. In general, it seems, some form of self-sacrifice by O is required to make its commitment credible and deterrence, therefore, stable--that is, for O to ensure that E will comply with certainty.

Although both players may prefer to play a sequential rather than a simultaneous game, the allocation of Pareto-superior benefits in such a game may lead to controversy, as we illustrate later with an example. We conclude by discussing some policy implications of our results and possible applications of the model.

3. The Simultaneous Game

The game tree of the simultaneous game is given in Figure 1: E

Figure 1 about here

---

violates an arms-control agreement with probability q and complies with probability 1-q; simultaneously or independently, D inspects with probability p and does not inspect with probability 1-p. (We show these probabilites on the game tree but stress that they do not indicate chance events.) The latter moves are enclosed in the same information set to indicate that O makes its choice without knowledge of what E does, so this game is one of imperfect information.

Figure 1

Simultaneous Game

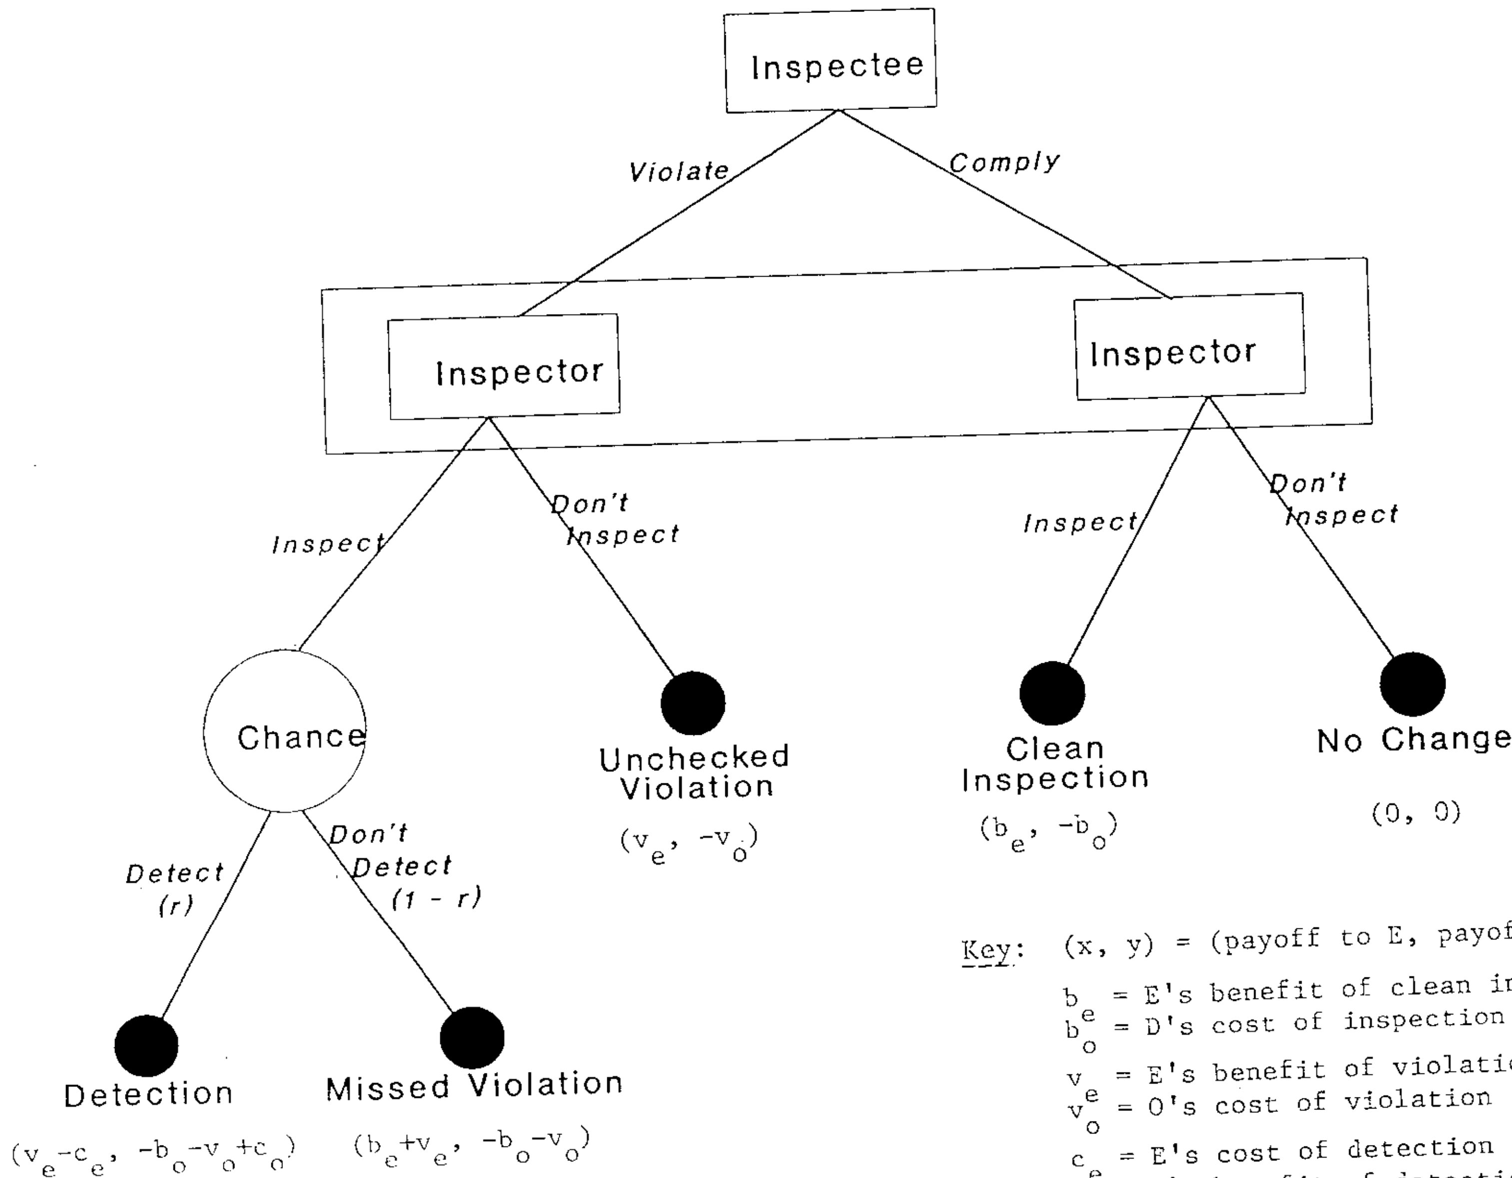

<u>Key</u>: (x, y) = (payoff to E, payoff to O)

b_e = E's benefit of clean inspection
b_o = D's cost of inspection

v_e = E's benefit of violation
v_o = O's cost of violation

c_e = E's cost of detection of violation
c_o = O's benefit of detection of violation

Reading across the next-to-lowest level of the game tree from right to left, the right-most outcome is no change, in which E complies and O does not inspect. Without loss of generality, we assume the payoffs of E and O at this outcome are (0, 0).

On the other hand, if E complies and O inspects, the outcome is a clean inspection, with payoffs of (b_e, $-b_o$) to E and O, respectively. E benefits from the fact that its compliance is certified by an inspection.[2] Because an inspection is costly to perform, we assume b_o is a cost to O and, therefore, indicate it as a negative quantity, $-b_o$ (all the parameters are assumed to be nonnegative).

Moving leftward, if E violates and O does not inspect, the outcome is an unchecked violation, with payoffs of (v_e, $-v_o$) to the players. Clearly, E benefits and O loses when O, by not inspecting, fails to uncover E's violation at this outcome.

To be sure, one can think of situations in which O might prefer not to unearth A violation. For example, it may be better for O simply to overlook a violation that is truly minor. We make the conservative assumption that this is not the case, but this assumption could be altered in the model.

The left-most node at the next-lowest level, where E violates and C inspects, is not resolved until "chance" intervenes to produce an outcome at the lowest level. We assume that either O's inspection equipment detects a violation with probability r or does not detect a violation with probability l-r. Call the former outcome a detected violation, with payoffs of (v_e-c_e, $-b_o-v_o+c_o$) to the players: E benefits from the violation (v_e) but loses from its detection ($-c_e$); O loses both from the cost of the inspection ($-b_o$) and the violation ($-v_o$) but benefits from the detection (c_o).

When D's detection equipment does not detect a violation (with probability l-r), we have a missed violation, with payoffs of (b_e+v_e, $-b_o-v_o$): E benefits from the fact that the inspection shows it complied (b_e) when in fact it violated (v_e); O loses from both the cost of the inspection (b_o) and the violation (v_o).

We summarize the payoff parameters of the model below:

Inspectee (E)	Inspector (O)
b_e = benefit of clean inspection	b_o = cost of inspection
v_e = benefit of violation	v_o = cost of violation
c_e = cost of detected violation	c_o = benefit of detected violation

Because all the symbols represent positive numbers, a parameter described as a "benefit" contributes positively, and a parameter described as a "cost" contributes negatively, to a player's utility.

Notice that an inspection might be clean either because E complies (clean inspection) or because E violates but the inspection does not uncover the violation (missed violation). Similarly, a violation might be successful either because O does not inspect (unchecked violation) or because O inspects but does not uncover the violation (missed violation).

Our addition and subtraction of parameters at each outcome shows that, for O, a missed violation ($-b_o-v_o$) is worse than an unchecked violation ($-v_o$). In fact, the reverse might be the case. Although O incurs the cost of an inspection at the former outcome, at the latter outcome O made the attempt to check for a violation; while unsuccessful, O "did its best." Arguably, an unsuccessful attempt--at least psychologically--is better than not even trying to detect the violation.

For E, by contrast, a missed violation (b_e+v_e) seems unequivocally better than simply a clean inspection (b_e), because getting away with a violation is psychologically as well as materially beneficial. Later we note that reversing the ordering of the costs to O of an unchecked versus a missing violation has no effect on our conclusions, which illustrates the robustness of our model to different assumptions, both of which in this case are plausible.

The expected payoffs of the players in the simultaneous game will be a function of the strategic probabilities (E chooses probability q, O chooses probability p), the chance event ("nature" chooses according to probability r, a characteristic of the detection system), and the aforementioned costs and benefits:

$$\begin{aligned} V(p,q) &= qpr(v_e-c_e, -b_o-v_o+c_o) + qp(l-r)(b_e+v_e, -b_o-v_o) \\ &\quad + q(l-p)(v_e, -v_o) + (l-q)p(b_e, -b_o) + (l-q)(l-p)(0, 0) \\ &= q(v_e, -v_o) + p(b_e, -b_o) + pqr(-b_e-c_e, c_o). \end{aligned} \quad (l)$$

The right side of (l) is a linear combination of ordered pairs; multiplication of each ordered pair by its coefficient gives the ordered pair (V_e,V_o), where V_e is E's expected payoff and V_o is O's.

We prove in the Appendix that the simultaneous game subsumes three cases--each with a unique Nash equilibrium--depending on the relationship of the detection probability r to the benefit and cost parameters (asterisks signify the players' Nash equilibrium strategies and their resulting expected payoffs):

1. $\underline{r \leq b_o/c_o}$: $p^* = 0$ and $q^* = 1$; $V_e^* = v_e$ and $V_o^* = -v_o$.
2. $\underline{b_o/c_o < r \leq v_e/(b_e+c_e)}$: $p^* = 1$ and $q^* = 1$; $V_e^* = v_e + b_e - r(b_e + c_e)$ and $V_o^* = -v_o - b_o + rc_o$.
3. $\underline{r > b_o/c_o \text{ and } r > v_e/(b_e+c_e)}$: $p^* = v_e/[r(b_e+c_e)]$ and $q^* = b_o/(rc_o)$; $V_e^* = b_e v_e/[r(b_e+c_e)]$ and $V_o^* = -b_o v_o/(rc_o)$.

Observe that the case 1 and case 2 equilibria are pure, whereas the case 3 equilibrium is mixed. It is easy to see that if $b_o/c_o < v_e/(b_e+c_e)$, all three games can occur (but not simultaneously) for different values of the parameter r. If $b_o/c_o \geq v_e/(v_e+c_e)$, only games cases 1 and 3 are possible, because this inequality precludes an r that can satisfy both case 2 inequalities.

To interpret the conditions under which each case occurs, notice that these conditions indicate the relationship of the detection probability r to two ratios, b_o/c_o for O and $v_e/(b_e+c_e)$ for E. O's ratio of the cost of an inspection to the benefit of detecting a violation we call its <u>cost ratio</u>. E's ratio of the benefit of a violation to the benefit of a clean inspection plus the absolute value of the cost of being detected we call its <u>benefit ratio</u>. The following propositions describe verbally when each case occurs and its Nash equilibrium:

1. <u>Low r: inspection always worthless; violation always worthwhile</u>. If D's detection probability is less than or equal to its cost ratio, it will never inspect (dominant strategy) and E will, as a best response, always violate.

2. <u>Intermediate r: inspection always worthwhile; violation always worthwhile</u>. If O's detection probability is greater than its cost ratio and

less than or equal to E's benefit ratio, E will always violate (dominant strategy) and O, as a best response, will always inspect.

3. High r: occasional inspections worthwhile; occasional violations worthwhile. If O's detection probability is greater than both its cost ratio and E's benefit ratio, it will sometimes inspect and E will sometimes violate (mixed strategies for both players).

It is noteworthy that no equilibrium strategies call for E always to comply, which is a result obtained in a similar game-theoretic model used to study the enforcement of irrigation rules to prevent the stealing of water (Weissing and Ostrom, 1991). On the contrary, it is in E's interest to violate an agreement either some (case 3) or all (case 1 or case 2) the time. How often in the case of a case 3 equilibrium depends on $q^* = b_o/(rc_o)$, which is in direct proportion to D's cost of inspecting and in inverse proportion to O's detection probability and the benefit O receives from detecting a violation.

In case 3--the best one can hope for in a treaty--E should violate more often as O's cost of inspecting increases; and less often as D's detection probability and benefit from detection increase. For O, $p^* = v_e/[r(b_e+c_e)]$ indicates that it should inspect more often as E's benefit of a violation increases; and less often as O's detection probability, E's benefit from a clean inspection, and E's cost of detection of a violation increase. Significantly, a common feature of both equilibrium probabilities is that they are inversely proportional to r: a greater detection probability decreases both E's violation probability and O's inspection probability. These probabilities can also be interpreted in terms of levels, with higher probabilities translating into more serious violations or more thorough inspections.

It is perhaps dismaying that the Nash equilibria in none of the three cases--whatever the values of the benefit and cost parameters of each player or of D's detection probability--results in E's always complying. If total compliance is unattainable at an equilibrium strategy in the simultaneous game, can it be induced in another game?

4. The Sequential Game

Consider a game in which O has a move prior to play in the simultaneous form: it announces its choice of a nonzero inspection probability p ($0 < p \leq 1$) before E makes its choice of q in the Figure 1 game. This sequential game is shown in Figure 2, in which O

Figure 2 about here

---

announces p and then E, knowing p, chooses to comply or violate.

After having made its announcement, we assume that O is committed to carrying out an inspection with probability p--or possibly there is some agent who performs this task. In either event, we assume that O's announcement is credible, presumably because O (or its agent) has established a record of following through on previous announcements.[3]

There are at least two ways of interpreting p. First, at a single inspection site, a random device might be used to determine, with probability p, whether an inspection would actually be carried out in a certain time period. Second, if there are n inspection sites, O's commitment would be to inspect pn sites, selected at random. In the latter

Figure 2

Sequential Game

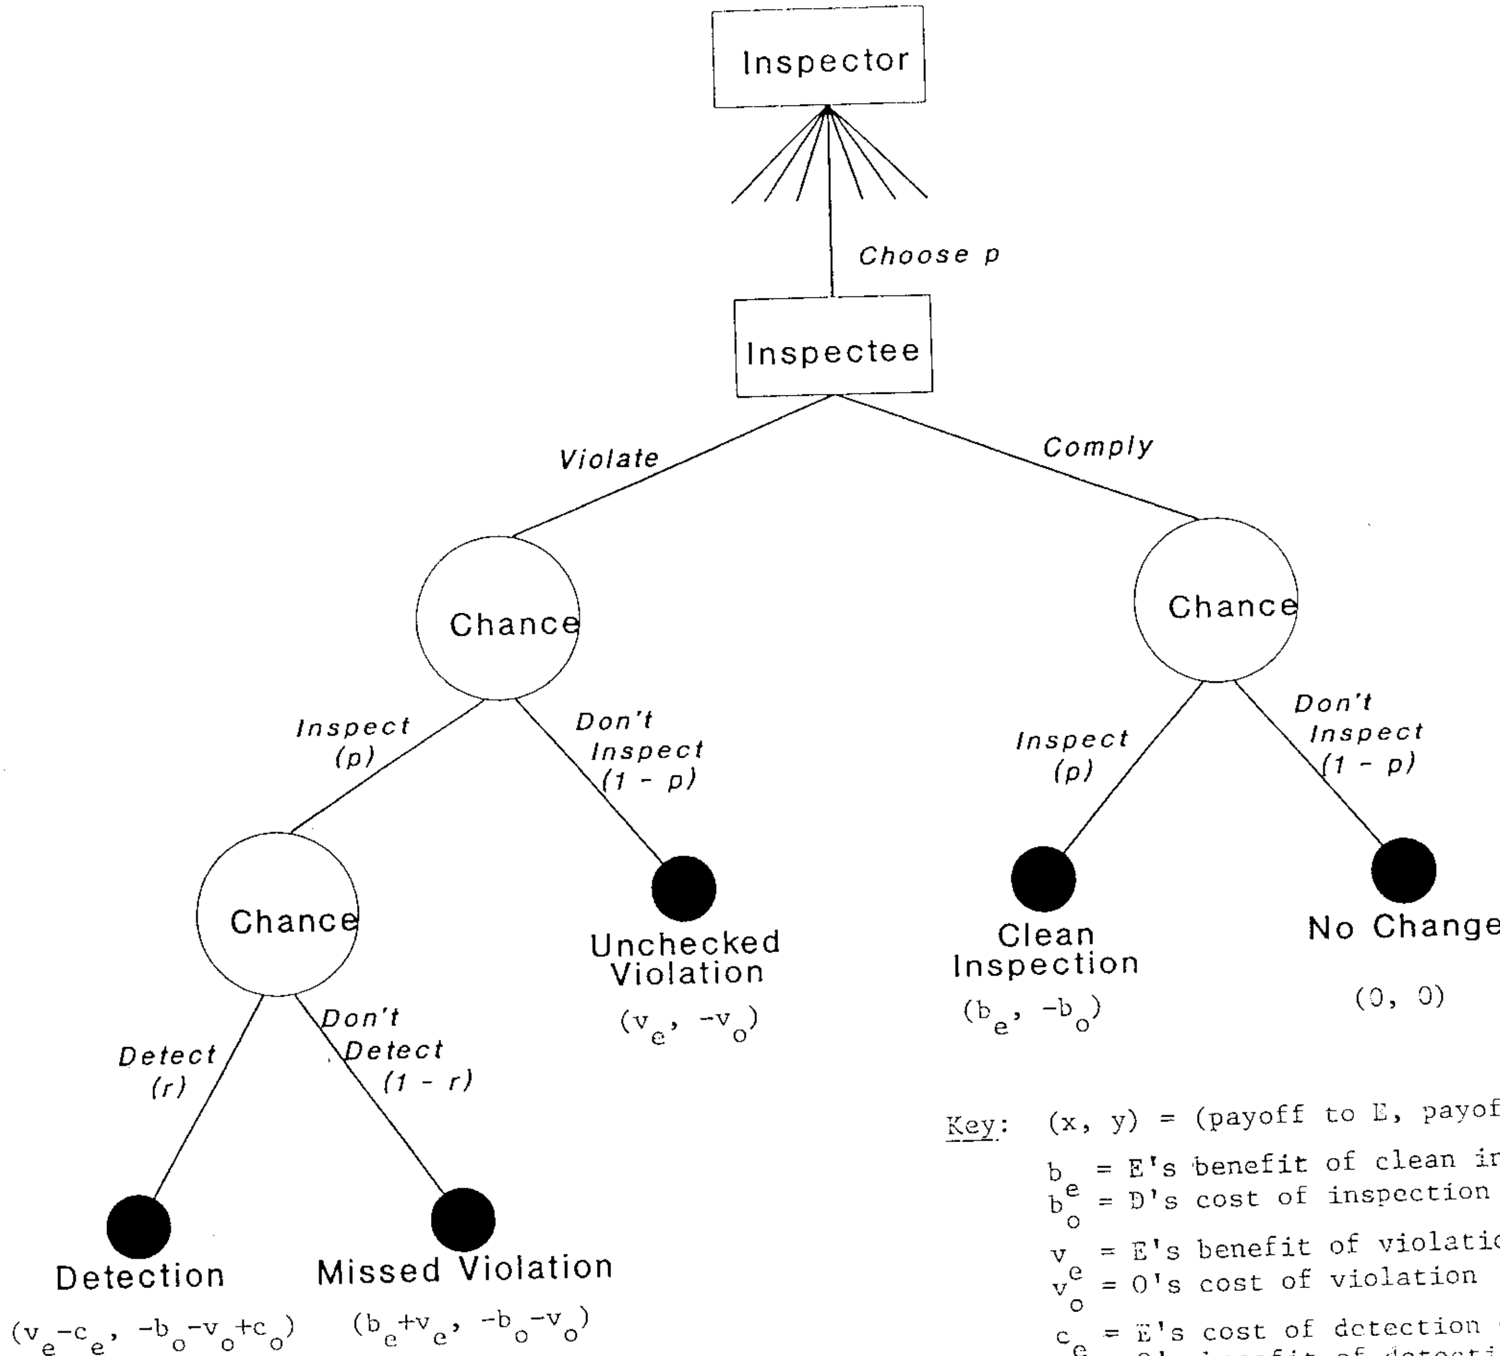

Key: (x, y) = (payoff to E, payoff to O)

b_e = E's benefit of clean inspection
b_o = O's cost of inspection

v_e = E's benefit of violation
v_o = O's cost of violation

c_e = E's cost of detection of violation
c_o = O's benefit of detection of violation

case, the benefit and cost parameters discussed in section 3 would need to be summed over the pn sites. We will return to the question of interpreting p, and applying the model, in the concluding section.

Given O's prior move, E can now choose q, knowing p in advance. This is not to say that E knows, in a one-site, one-inspection game, whether or not it will be inspected. Rather, E knows the probability that this event will occur. Consequently, the sequential game is still one of imperfect information: as in the simultaneous game, when E and O select their actions, they do not assuredly know the choices of their adversaries. E's only additional information in the sequential game, when it chooses q and thence whether to comply or violate, is its knowledge of p.

It turns out that the players' expected payoffs in the sequential game are exactly the same as those in the simultaneous game, given by (1). The reason is that although the sequential game adds an announcement (and commitment) by O at the start of the game, the players' possible actions, and the four possible outcomes, do not change.

But O's announcement does alter the players' strategic calculations. As shown in the Appendix, O's announcement of $p = \bar{p} = v_e/[r(b_e+c_e)]$, provided $0 < \bar{p} < 1$, renders E's choice of q irrelevant: q has no effect on E's expected payoff. On the other hand, if $p < \bar{p}$, O can induce E to choose $q = 1$ (always violate) to maximize V_e; if $p > \bar{p}$, O can induce E to choose $q = 0$ (always comply) to maximize V_e. Furthermore, such inducement constitutes a subgame-perfect equilibrium point.

What is it in O's interest to do? Although one can imagine bizarre situations when it might be in O's interest to induce $q = 1$, we shall restrict our analysis to finding conditions under which it is in both players' interest for O, by announcing a p in advance, to induce $q = 0$.

For inducement of compliance to be possible, $\bar{p} < 1$ is necessary. Hence, O can induce $q = 0$ only if $r > v_e/(b_e+c_e)$. The latter inequality rules out the case 2 equilibrium in the simultaneous game, wherein "always violate" is a dominant strategy. This makes sense: if $q = 1$ is E's dominant strategy, nothing O can say will alter E's preference for always choosing it.

When $q = 0$,

$$V(p, 0) = p(b_e, -b_o)$$

from (1); thus, D maximizes V_o by making p as small as possible, subject to the constraint that $p > \bar{p}$. For small positive ε, the strategy pair ($p = \bar{p}+\varepsilon$, $q = 0$) is an ε-equilibrium[4] in the sequential form that is Pareto-superior to the corresponding Nash equilibrium in the simultaneous form if

$$v_e/(b_e+c_e) < v_o/c_o \text{ and } r > \max\{b_o/c_o, v_e/(b_e+c_e)\}. \tag{2}$$

Note that (2) is a somewhat more restrictive set of conditions than those defining the simultaneous game 3 equilibrium (involving mixed strategies). D can also induce a Pareto-superior ε-equilibrium if

$$v_e/(b_e+c_e) < r \leq b_o/c_o \text{ and } v_e/b_e < \bar{p} = v_e/[r(b_e+c_e)] < v_o/b_o, \tag{3}$$

which is a substantially more restrictive set of conditions than for a simultaneous game 1 equilibrium.

There are no other conditions that enable O, by choosing $p > \bar{p}$, to induce E to choose $q = 0$ such that the ε-equilibrium in the sequential game is Pareto-superior to the Nash equilibrium in the simultaneous game.

Hence, either (2) or (3) is sufficient for both players to prefer to play the sequential game; they have in common the condition, $r > v_e/(b_e+c_e)$, enabling O to induce E always to comply.

Observe that (2) requires a "high" r, whereas (3) requires an "intermediate" r and additional conditions on $\bar{p}$. One would normally expect that $v_e > b_e$ (i.e., E's benefit from a violation will be greater than its benefit from a clean inspection), so $v_e/b_e > 1$. Because this inequality renders the condition $\bar{p} > v_e/b_e$ in (3) impossible, inducement in case 1, while a theoretical possibility, seems highly unlikely ever to occur.

We illustrate the difference between a Nash equilibrium in the simultaneous game and an ε-equilibrium in the sequential game with a case 3 example. This example highlights why a minimal kind of inducement by O may be viewed as "unfair" by E, forcing O to concede more to E in dividing the surplus generated by the Pareto-superior ε-equilibrium.

<u>Example</u>. Assume that the benefits to one player are equal to the costs to the other, except for the changes in sign: $b_e = b_o = 1$; $v_e = v_d = 2$; and $c_e = c_o = 3$. In addition, assume that $r = 3/4$. Because these parameter values satisfy (2), O's choice of $p > \bar{p} = v_e/[r(b_e+c_e)] = 2/3$ will induce E to choose $q = 0$, yielding the players expected payoffs at the ε-equilibrium of

$$V(p, 0) = p(b_e, -b_o) = 2/3^{+}(1, -1) = (2/3^{+}, -2/3^{-}).$$

Note that E does only marginally better ($2/3^{+}$) than at the Nash equilibrium in the simultaneous game ($V_e^* = b_e v_e/[r(b_e+c_e)] = 2/3$). On the other hand, O does substantially better, reducing its expected loss ($V_o^* = -b_e v_e/[r(b_e+c_e)] = -8/9$) at the Nash equilibrium in the simultaneous game to $-2/3^{-}$ in the sequential game, an improvement of about 25 percent. (Recall that this is O's cost of a probabilistic inspection threat to deter E from cheating.)

Clearly, inducement can reap major rewards for O. In practice, however, its effective use may require that O offer E more than a marginal increase in its expected payoff (of achieving a clean inspection) over what it would receive at the Nash equilibrium. Thus, instead of choosing $p = 2/3^{+}$, O might choose $p = 3/4$, which would diminish its 25-percent cost reduction to 12.5 percent. However, it would give E the same 12.5 percent bonus over its Nash equilibrium payoff, which seems more even-handed than allowing all the benefits to go to O. O's choice of $p = 3/4$ also sends a clear signal that E can gain considerably from compliance, whereas E might not be able to ascertain, in a real-life situation, that $p = 2/3^{+}$ would in fact be beneficial.

Although O's inducement strategy of $p = 3/4$ benefits both players equally, it is not in equilibrium with E's choice of $q = 0$. The reason is that O has an incentive to let $p = 3/4$ slip toward $p = 2/3$, which in principle is sufficient for inducement, when it makes its announcement and later choice.

Of course, there is no reason to insist that ε be miniscule. This will be especially true in situations in which small gains are difficult to perceive or because some notion of even-handedness is operative (e.g., that prescribes dividing the profits of inducement equally, as illustrated in our example).

Indeed, E might threaten not always to comply <u>unless</u> there is a more equal division. This threat, if believed, would lead to strategies described would still be Pareto-superior to the Nash equilibrium in the simultaneous game yet less lopsided in favor of O.

The general conclusion that we draw from our analysis of the sequential game is that O, by its announcement and commitment to an inspection strategy above a certain level, may be able to induce an equilibrium that both players would prefer over that in the simultaneous game. Under either conditions (2) or (3), it is in the interest of both players to play a sequential game in order to obtain a Pareto-superior equilibrium.[5]

These conditions depend critically on the detection probability r--it must be sufficiently high in order to induce compliance by E. If, however, r is greater than the benefit ratio but not the cost ratio, as given by (3), then an ε-equilibrium must also satisfy an unlikely condition on $\bar{p}$. In the absence of such an induced equilibrium, the Nash equilibrium in the simultaneous game would presumably be chosen.

To be sure, under certain conditions (not given here), <u>non</u>compliance may be induced as a Pareto-superior ε-equilibrium in the sequential game. But the circumstances under which O would choose a p to induce E always to violate an agreement seem quite absurd. Although O might want to embarrass E on occasion by detecting a violation, it is hard to conceive of plausible scenarios when it would be in O's interest that E always violate an agreement. Consequently, we have concentrated on finding conditions that lead players to choose Pareto-superior compliance equilibria.

These equilibria, however, do not maximize the players' expected payoffs. Given compliance on the part of E, O would always prefer to choose p = 0 and avoid the cost of inspection entirely. But the resulting outcome, while superior for O, is inferior for E because E loses the benefit of a clean inspection; moreover, this outcome is not in equilibrium.

In fact, E's best response to no inspections would be always to violate. This is why, as a necessary condition for the inducement of compliance, p must surpass the threshold $\bar{p}$.

But how can O ensure that its announcement of a mixed inspection strategy will be believed? We discuss this and other questions connected with the interpretation and application of the model in the concluding section.

5. Discussion

We have modeled arms-control inspection by inspection games in simultaneous and sequential form. In the simultaneous game, there are no equilibrium strategies that always lead to certain compliance by E, though the case 3 equilibrium probability of E's compliance increases as O's detection probability increases. On the other hand, there are cases (1 and 2) in which it is advantageous for E always to violate an agreement, even when O always inspects (case 2).

The mixed-strategy equilibrium of case 3 is probably the norm in certain arms-control agreements: E sometimes violates and O sometimes inspects or, more metaphorically, they violate and inspect at certain nonzero levels. In this situation, E sometimes is caught cheating and suffers because of it; O suffers from the violations but sometimes benefits from their detection.

This outcome may not be ideal from either player's viewpoint; a case in point is the 1919 Treaty of Versailles, whose arms-control provisions were consistently evaded (Berkowitz, 1987, p. 17). But it may be possible to circumvent this outcome if O credibly commits itself to a (mixed) inspection strategy in advance.[6] In particular, O may be able to induce E never to cheat by promising to inspect sufficiently often. Thereby, E benefits from clean inspections, never having to suffer the costs of having a violation detected, and O benefits from a rigorously adhered to treaty.

Curiously, O could do still better by not inspecting if E complied. While Pareto-optimal, however, this outcome is not in equilibrium. Without being inspected, E would always violate, which is exactly what O's continuing inspections deter. The threat of harm, as we have argued elsewhere (Brams and Kilgour, 1988), seems an inescapable feature of stable cooperative relations.

We identified conditions under which O's inducement strategy would not only deter all violations but also lead to an outcome Pareto-superior to the Nash equilibrium in the simultaneous game. These conditions are generally more restrictive than the conditions for a Nash equilibrium; in particular, they require that O's detection probability, r, be sufficiently high.

Although the value of inducement may be apparent in certain situations, it is less apparent how an announced inspection strategy can be implemented, especially one that is probabilistic in nature. In a one-shot game, a coin reflecting the appropriate odds might be used, but this approach would probably strike most practitioners as hopelessly naive--and another egregious example of ivory-tower thinking.

In fact, however, arms-control inspections are almost always part of a continuing process of treaty verification. Thus, if there were several sites to be checked over time, but not all could be, an inducement strategy might to be checked over time, but not all could be, an inducement strategy might involve announcing how many would be randomly selected for inspection in a certain time period. Alternatively, if there were one major site that needed to be inspected periodically, a strategy of how often it would be inspected might be announced, but exactly when would be determined by a random process.

Presumably, if inspection is part of an ongoing game, O would have good reason to keep its commitment; otherwise, future commitments would lose their inducement value. The need to maintain reputations of fair dealing and honesty is probably further reinforced if, over time, two parties to a treaty must play the roles of both O and E. Alternatively, it may be in the interest of the parties that an independent inspection agency, which is both credible and politically acceptable, carries out the inspections.

One interesting aspect of the two games we have analyzed is that O can decide--practically at the moment of play--which one is preferable. If the sequential form is, O can seize the initiative, announcing an equilibrium inducement strategy.

But if, say, O's analysis reveals that its detection probability is not sufficiently high to induce total compliance, it can use the sequential-form model to help determine what detection probability would be required in the future--and ask whether the cost of developing the requisite detection capability is worth the price. Conceivably, the model could be used to determine how many inspections at multiple sites, or how frequently inspections at a fixed site, would be needed to induce a Pareto-superior compliance equilibrium, with this number written into a treaty. However, determining a sufficiently high p to ensure compliance at minimum cost, are probably unrealistic.[7]

We return to the earlier question raised about how our results would be affected if an unchecked violation were more, rather than less, costly to O than a missed violation. Without going through a formal analysis of this reordering of two of O's payoffs, it is not difficult to show that if the costs of not inspecting a violation are greater for O, E complies more often in equilibrium in the simultaneous game. Nevertheless, it is always in E's interest to violate at least some of the time.

The main qualitative result of our analysis is unaffected by this reordering: only inducement by O in the sequential game can ensure total compliance by E. It will not only be in O's interest but also E's that, given a sufficiently high detection probablity, O announce an inspection strategy to induce compliance.

6. Conclusions

Taking the initiative in arms control may obviate the need for formal treaties, as is persuasively argued by Downs and Rocke (1990) using different game-theoretic models. Thus, for example, unilateral choices by the Soviets, beginning in the late 1980s, to withdraw their forces from Eastern Europe and reduce them substantially in the Soviet Union triggered reciprocal allied force reductions.

We believe that the prior announcement of inspection strategies, and presumed follow-up actions, may have similar effects in other parts of the world in which arms-control agreements are elusive yet feasible. To wit, a commitment to an inspection program may induce compliance that benefits both O and E over what they would obtain in equilibrium had O not both O and E over what they would obtain in equilibrium had O not announced its inspection program and acted consistently with its announcement.

In regions of conflict like the Middle East and Eastern Europe, a third party like the UN or the European Community is probably best equipped to play the role of O, carrying out inspections of one or more parties. Thus, multiple Es might each play an inspection game with some O, or a single E, like Iraq after its 1991 war with allied forces in Kuwait, might be forced to accede to UN inspections.

With the demise of East-West conflict--and that between the United States and the former Soviet Union in particular--we suspect that third parties, under regional or international auspices, will become increasingly important as both mediators of conflicts and agents of change. It will not be an easy task, especially in the case of ethnic conflict, for them to carry out arms-control inspections to try to ensure that the means to restart old conflicts are restrained. But our analysis indicates that O's prior announcement and commitment to a specific inspection program, given that its detection capabilities are substantial enough, may redound to the benefit of all parties.

Appendix

If the payoff function given by (1),

$$V(p, q) = q(v_e, -v_o) + p(b_e, -b_o) + pqr(-b_e-c_e, c_o),$$

is rewritten in terms of summary parameters A, B, and C,

$$V(p, q) = q(A_e, A_o) + p(B_e, B_o) + pq(C_e, C_o),$$

the simultaneous game can be represented as a 2 x 2 matrix game (see Figure 3), in which E may either comply (C) or not comply ($\overline{C}$),

---

Figure 3 about here

---

and O may either inspect (I) or not inspect ($\overline{I}$). We assume that O and E choose their strategies of I and $\overline{C}$ with probabilities p and q, respectively.

We will show that game-theoretically there is always a unique outcome for the game in Figure 3. Specifically, either at least one player has a dominant strategy--resulting in a pure-strategy Nash equilibrium--or, if there are no dominant strategies, there is a unique Nash equilibrium in mixed strategies.

By assumption, $A_e = v_e > 0$ and $B_o = -b_o < 0$. In the Figure 3 game, there are two situations in which one player has a dominant strategy, resulting in a unique Nash equilibrium in pure strategies:

1. $\underline{B_o + C_o = -b_o + rc_o \leq 0}$. $\overline{I}$ is dominant (p = 0) because $B_o < 0$. $\overline{C}$ is E's best response (q = 1), resulting in outcome $\overline{CI}$.

Figure 3

Payoff Matrix of Simultaneous Game

		Inspector (O)		
		Inspect (I)	Don't inspect ($\bar{I}$)	
	Don't comply ($\bar{C}$)	($A_e+B_e+C_e$, $A_o+B_o+C_o$)	(A_e, A_o)	[q]
Inspectee (E)				
	Comply (C)	(B_e, B_o)	(0, 0)	[l-q]
		[p]	[l-p]	

Key: (x, y) = (payoff to E, payoff to O)

Probabilities of strategy choices by E (q and l-q) and O (p and l-p) shown in brackets

2. Assume situation 1 does not apply; instead, $\underline{B_o + C_o > 0}$. If $\underline{A_e + C_e = v_e + r(-b_e - c_e) \geq 0}$, $\overline{C}$ is dominant (q = 1) because $A_e > 0$. I is O's best response (p = 1), resulting in outcome $\overline{C}I$.

Now assume that neither situation 1 nor situation 2 holds, so $B_o + C_o > 0$ and $A_e + C_e < 0$. To determine whether there is a situation in which q = 0 (always comply) is a Nash equilibrium, we distinguish the payoff functions for each player:

$$V_e = qA_e + pB_e + pqC_e; \quad V_o = qA_o + pB_o + pqC_o.$$

Assume q = 0. Because

$$\frac{\partial V_o}{\partial p}(p, 0) = B_o = -b_o < 0,$$

any Nash equilibrium with q = 0 also has p = 0. But now assume p = 0. Then

$$\frac{\partial V_e}{\partial q}(0, q) = A_e = v_e > 0,$$

implying E's best response to p = 0 is q = 1. Hence, there is no Nash equilibrium with q = 0.

Similarly, it is possible to demonstrate that there is no Nash equilibrium with q = 1, except in situations 1 and 2, which proves that there are no other pure-strategy Nash equilibria.[8] Because there must exist at least one Nash equilibrium (Nash, 1951), it must be in mixed strategies when these two situations do not obtain.

What are the players' mixed strategies that support a Nash equilibrium? Differentiating each player's payoff function with respect to the strategic variable it controls yields

$$\frac{\partial V_e}{\partial q} = A_e + pC_e; \qquad \frac{\partial V_o}{\partial p} = B_o + q_o.$$

Setting these derivatives equal to 0 and solving for p and q,

$$p^* = -A_e/C_e = v_e/[r(b_e+c_e)]; \qquad q^* = -B_o/C_o = b_o/(rc_o),$$

give the players' (mixed) equilibrium strategies. They are the strategies (p^* for O, q^* for E) that render the other player's strategy choice irrelevant, robbing each player of any incentive to switch its strategy.

The pure-strategy equilibria, as given in the text for cases 1 and 2 in section 3, correspond to those that exist in situations 1 and 2. Likewise, the case 3 equilibrium corresponds to the mixed-strategy equilibrium that we have just derived. The latter equilibrium is <u>weak</u> because, although each player cannot do better by departing from its equilibrium strategy, it never does worst, either; the other player's equilibrium strategy makes each player indifferent among all its strategy choices.

In the sequential game, O makes a prior announcement of its p. To find a subgame-perfect equilibrium, we must therefore consider, for any p chosen by O, what q maximizes V_e. Because

$$\frac{\partial V_e}{\partial q} = A_e + pC_e = v_e - pr(b_e+c_e),$$

it follows that

$$\frac{\partial V_e}{\partial q} = \begin{cases} > 0 \text{ if } p < \dfrac{v_e}{r(b_e+c_e)} \Rightarrow q = 1 \text{ is the best response of E} \\ = 0 \text{ if } p = \dfrac{v_e}{r(b_e+c_e)} \Rightarrow \text{E is indifferent among all q} \\ < 0 \text{ if } p > \dfrac{v_e}{r(b_e+c_e)} \Rightarrow q = 0 \text{ is the best response of E.} \end{cases}$$

Consequently, by choosing

$$p > \bar{p} = v_e/[r(b_e+c_e)], \tag{4}$$

O can induce E to choose $q = 0$. For this to be possible, however, $\bar{p} < 1$ is necessary. Hence, O can induce $q = 0$ only if $r > v_e/(b_e+c_e)$, which rules out the case 2 sequential game as a candidate for inducement.

Now consider the case 3 sequential game. Inducement leads to a greater expected payoff for O if its expected payoff from inducement, given in section 4, is greater than its Nash equilibrium expected payoff, given in section 3:

$$-\bar{p}b_o > -b_o v_o/(rc_o) \Rightarrow \bar{p} < v_o/(rc_o),$$

or

$$v_e/[r(b_e+c_e)] < v_o/(rc_o).$$

from (4). For E, inducement automatically leads to a greater expected payoff than its Nash equilibrium because

$$pb_e > b_e v_e/[r(b_e+c_e)] = \bar{p}b_e \Leftrightarrow p > \bar{p},$$

as assumed in (4). Hence the ε-equilibrium for case 3 is Pareto-superior to the Nash equilibrium if

$$v_e/(b_e+c_e) < v_o/c_o \text{ and } r > \max\{b_o/c_o, v_e/(b_e+c_e)\}.$$

For a game 1 equilibrium, inducement will lead to a greater expected payoff for O if

$$-\bar{p}b_o > -v_o \Rightarrow \bar{p} < v_o/b_o;$$

it leads to a greater expected payoff for E if

$$\bar{p}b_e > v_e \Rightarrow \bar{p} > v_e/b_e.$$

For inducement to be Pareto-superior to the corresponding Nash equilibrium in the simultaneous game, therefore,

$$v_e/(b_e+c_e) < r \leq b_o/c_o \text{ and } v_e/b_e < \bar{p} = v_e/[r(b_e+c_e)] < v_o/b_o$$

is required.

Notes

[1]For reasons to be given later (note 2), we do not postulate the possibility of an apparent detection when no violation occurs.

[2]We assume that there is no possibility of a type 1 error (false alarm), which would occur if D "detected" a violation when E had in fact complied. This assumption is one of convenience and could be relaxed in an extension of the model. It seems most applicable to on-site inspections, which are relatively robust against type 1 errors. This simplifying assumption, and others we make about relative costs are benefits, can certainly be challenged, but we emphasize the game-theoretic methodology we propose can be reapplied with alternative assumptions that may seem more plausible to some.

[3]This is where the role of reputation comes into play. Here we do not model how reputation might be established in repeated play but simply assume it in the present one-shot game. The question of reputation formation has been much investigated in the recent game-theoretic literature (see, e.g., Rasmusen, 1989), but in this paper we choose to focus on the conditions under which an existing reputation, which makes O's announcements credible, may lead to Pareto-superior outcomes in arms-control inspection games. It is worth noting that a player's reputation may also be supported by structural features of a situation (e.g, the existence of an independent inspection agency), which may be quite independent of its past behavior.

[4]An ε-equilibrium is defined with respect to the game tree, in which O is assumed not to be able to renege on its commitment to inspect with probability p, even though, as we shall show later, it would be optimal for O to do so if E always complies. This equilibrium is also a Stackelberg equilibrium, with O the leader and E the follower.

[5]Different effects of inducement in arms-control inspection games have been analyzed by Maschler (1966, 1967), Brams (1985), Avenhaus (1986, 1990), Fichtner (1986), Brams and Davis (1987), and Brams and Kilgour (1986, 1987, 1988); see also Wittman (1989) and O'Neill (1992) for related game-theoretic models. Our present model differs from most previous models in (1) unpacking certain benefits and costs rather than ordering outcomes that define specific games; (2) allowing for the possibility of noninspection, as well as nondetection, of a violation; (3) focusing on Pareto-superior inducement outcomes that help both players, not just the inspector, and (4) illustrating the issue of even-handedness that might arise in dividing the surplus at such outcomes. As a consequence, we are able to distinguish different possible cases and their Nash equilibria in the simultaneous game, and different threshold conditions that result in Pareto-superior subgame-perfect ε-equilibria in the sequential game. This distinction makes clearer than have other game-theroetic models the possible benefits of inducement to both players.

[6]As noted earlier, O's credibility will in part depend on its reputation, which we assume is sufficient to render its announcement believable. But "following through" on its announcement raises new issues if inspections

are probabilistic, because the number of inspections O commits itself to over some time period is an expected value, not a certain quantity. Thus, enforcing an announced probability of inspection is hedged by uncertainty, but statistical methods can be applied to this question.

[7]Game-theoretic models for the recursive allocation of cheating resources by an inspectee, and inspections by an inspector, are given in Brams, Davis, and Kilgour (1991), von Stengel (1991), and Kilgour (1990), but they, too, are probably difficult to apply to real-world situations.

[8]The lack of such an equilibrium can also be seen from Figure 3. If the inequalities in situations 1 and 2 fail, the players always have an incentive to depart from the pure-strategy outcomes, moving in a counterclockwise direction.

References

Avenhaus, Rudolf (1990). "Decision Procedures for Safeguard Systems." Information and Decision Technologies 16, no. 3: 229-247.

Avenhaus, Rudolf (1986). Safeguards Systems Analysis. New York: Plenum.

Berkowitz, Bruce D. (1987). Calculated Risks: A Century of Arms Control, Why It Has Failed, and How It Can Be Made to Work. New York: Simon and Schuster.

Brams, Steven J. (1985). Superpower Games: Applying Game Theory to Superpower Conflict. New Haven, CT: Yale University Press.

Brams, Steven J., and Morton D. Davis (1987). "The Verification Problem in Arms Control: A Game-Theoretic Analysis." In Interaction and Communicaton in Global Politics, ed. Claudio Cioffi-Revilla, Richard L. Merritt, and Dina A. Zinnes. London: Sage, pp. 141-161.

Brams, Steven J., Morton D. Davis, and D. Marc Kilgour (1991). "Optimal Cheating and Inspection Strategies under INF." In Defense Decision Making: Analytical Support and Crisis Management, ed. Rudolf Avenhaus, Hassame Karker, and Michel Rudnianski. Berlin: Springer-Verlag, pp. 318-335.

Brams, Steven J., and D. Marc Kilgour (1986). "Notes on Arms-Control Verification: A Game-Theoretic Analysis." In Modelling and Analysis in Arms Control, ed. Rudolf Avenhaus, Reiner K. Huber, and John D. Kettelle. Berlin: Springer-Verlag, pp. 409-419.

Brams, Steven J., and D. Marc Kilgour (1987). "Verification and

Stability: A Game-Theoretic Analysis." In Arms and Artificial Intelligence:Weapon and Arms Control Applications of Advanced Computing, ed. Allan M. Din. Oxford, UK: Oxford University Press, pp. 193-213.

Downs, George W., and David M. Rocke (1990). Tacit Bargaining, Arms Races, and Arms Control. Ann Arbor, MI: University of Michigan Press.

Fichtner, J. (1986). "On Solution Concepts for Solving Two Person Games which Model the Verification Problem in Arms Control." In Modelling and Analysis in Arms Control, eds. Rudolf Avenhaus, Reiner K. Huber, and John D. Kettelle. Berlin: Springer-Verlag, pp. 421-441.

Kilgour, D. Marc (1990). "Optimal Cheating and Inspection Strategies under a Chemical Weapons Treaty." INFOR 28, no. 1: 27-39.

Maschler, Michael (1966). "A Price Leadership Method for Solving the Inspector's Non-Constant-Sum Game." Naval Research Logistics Quarterly 13, no. 1 (March): 11-33.

Maschler, Michael (1967). "The Inspector's Non-Constant-Sum Game: Its Dependence on A System of Detectors." Naval Research Logistics Quarterly 14, no. 3 (September): 275-290.

Nash, John (1951). "Non-cooperative Games." Annals of Mathematics 54: 286-295.

O'Neill, Barry (1992). "Why a Better Verification Scheme Can Give More Ambiguous Evidence." International Studies Quarterly, forthcoming.

Rasmusen, Eric (1989). Games and Information: An Introduction to

Game Theory. New York: Basic Blackwell.

von Stengel, Bernhard (1991). “Recursive Inspection Games.” Preprint, Armed Forces University Munich.

Weissing, Franz, and Elinor Ostrom (1991). "Irrigation Institutions and the Games Irrigators Play: Rule Enforcement without Guards." In Reinhard Selten (ed.), Game Equilibrium Models: Methods, Morals, and Markets, vol. II. Berlin: Springer-Verlag, pp. 188-262.

Wittman, Donald (1989). "Arms Control Verification and Other Games Involving Imperfect Detection." American Political Science Review 83 no. 3 (September 1989): 923-945.