ECONOMIC RESEARCH REPORTS

ARMS-CONTROL INSPECTION STRATEGIES THAT INDUCE COMPLIANCE: A GAME-THEORETIC ANALYSIS

by

D. Marc Kilgour
and
Steven J. Brams

R.R. #: 90-40 September, 1990

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

ARMS-CONTROL INSPECTION STRATEGIES THAT INDUCE COMPLIANCE: A GAME-THEORETIC ANALYSIS

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

ABSTRACT

ARMS-CONTROL INSPECTION STRATEGIES THAT INDUCE COMPLIANCE: A GAME-THEORETIC ANALYSIS

Arms-control inspection is modeled by two game forms that subsume different games between a Detector (D) and an Evader (E). In each form,

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

Besides certain 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 verification.

In the simultaneous form, D and E make simultaneous (or independent) choices. Depending on the parameters, three inspection games, each with a uniquely determined equilibrium--two in pure strategies and one in mixed strategies--are possible. Because none of the equilibria involves <u>certain</u> compliance by E, D is not always able to deter E from violating an agreement.

In the sequential form, D has the first move, announcing--and committing itself, or an independent agency, to--a probability p of inspecting E. E, with full knowledge of p, then chooses to comply or violate the agreement. If the probability of detection is above a certain threshold and certain other conditions are met, D can announce an inspection strategy that will <u>always</u> deter E from violating, which leads to an equilibrium in the sequential form Pareto-superior to that in the simultaneous form. Paradoxically, then, by publicly revealing its inspection strategy in advance, D may help itself as well as E escape a Pareto-inferior equilibrium outcome. Policy implications of this finding, and possible applications of the model, are discussed.

ARMS-CONTROL INSPECTION STRATEGIES THAT INDUCE COMPLIANCE: A GAME-THEORETIC ANALYSIS

D. Marc Kilgour and Steven J. Brams

1. Introduction

In an age in which democracies, in one form or another, have sprung up almost everywhere, it is useful to bear in mind that there exist huge arsenals of both conventional and nuclear weapons. If arms control is to be successful, it must include procedures for verifying arms-control treaties. In this paper, we show that if a player in a two-person game publicly reveals its inspection strategy in advance, it may not only help its adversary, who is thereby better informed, but, paradoxically, it may also help itself.

We demonstrate this result by constructing two different game forms that subsume specific arms-control inspection games between a Detector (D) and an Evader (E). In each form,

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

In both forms, we postulate as parameters certain costs and benefits that the players receive at the different possible outcomes.

We also postulate another parameter--the conditional probability that if E violates and D inspects, D will detect the violation. Thus, there is uncertainty that a violation will be uncovered if there is an inspection.

The first form we analyze is specified by a game tree. In this form, which we call simultaneous, D and E make simultaneous (or independent) choices. Depending on the values of the parameters, three games, each with a uniquely determined Nash equilibrium--two in pure strategies and one in mixed strategies--can occur. Because none of the games involves <u>certain</u> compliance by E, D is not always able to deter E from violating the agreement.

In the second form, which we call <u>sequential</u>, D has the first move: to choose a probability p of inspecting E, which D is then committed to carrying out (in ways to be specified later). We assume that D announces p before E and D make their strategy choices. Thus, E makes its choice of complying or violating with full knowledge of p but without knowledge of whether or not there will be an inspection. As in the simultaneous form, there is uncertainty as to whether, if E violates and D inspects, the violation will be uncovered.

In the sequential form, if the conditional probability of detection is above a certain threshold, there is a p that D can announce that will <u>always</u> deter E from violating. Moreover, under certain conditions, the resulting outcome is an ϵ-equilibrium in the sequential form that is Pareto-superior to the Nash equilibrium in the simultaneous form, 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 Nash equilibuium. This outcome, however, is not the only Pareto-superior outcome: D 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: D must pay <u>some</u> costs (i.e., those of carrying out the probabilistic inspections) to deter cheating in equilibrium. In general, it seems, some form of self-sacrifice by D is required to make its commitment credible and deterrence, therefore, stable.

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

2. The Simultaneous Form

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

---

Figure 1 about here

---

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. The latter moves are enclosed in the same information set to indicate that D make its choice without knowledge of what E does, so a game in this form is one of imperfect information.

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

On the other hand, if E complies and D inspects, the outcome is a <u>clean inspection</u>, with payoffs of (b_e, $-b_d$) to E and D, respectively. E benefits from the fact that its compliance is certified by an inspection.[1] Because an inspection is costly to perform, we assume b_d is a cost to D and, therefore, indicate it as a negative quantity, $-b_d$ (all the parameters are assumed to be nonnegative).

FIGURE 1

GAME TREE OF SIMULTANEOUS FORM

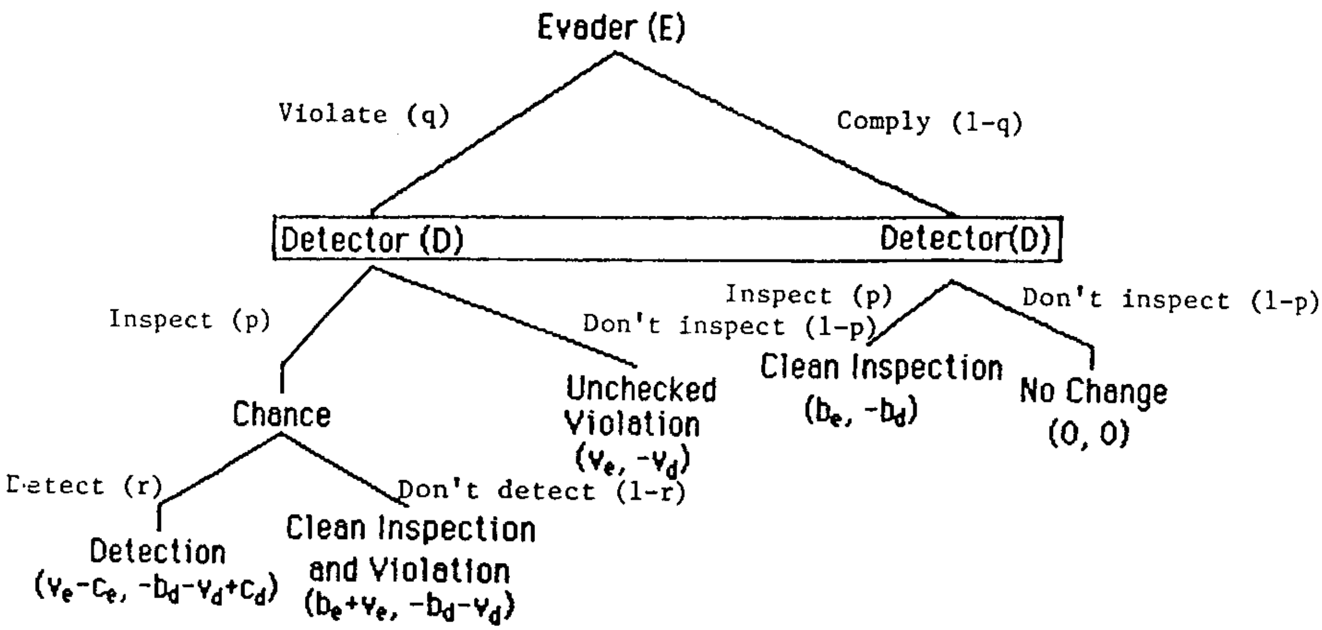

<u>Key</u>: (x, y) = (payoff to E, payoff to D)

b_e = E's benefit of clean inspection; $-b_d$ = D's cost of inspection

v_e = E's benefit of violation; $-v_d$ = D's cost of violation

$-c_e$ = E's cost of detection of violation; c_d = D's benefit of detection of violation

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

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 D's inspection equipment detects a violation with probability r or does not detect a violation with probability 1-r. Call the former outcome a <u>detected violation</u>, with payoffs of (v_e-c_e, $-b_d-v_d+c_d$) to the players: E benefits from the violation (v_e) but loses from its detection ($-c_e$); D loses both from the cost of the inspection ($-b_d$) and the violation ($-v_d$) but benefits from the detection (c_d).

When D's detection equipment does not detect a violation (with probability 1-r), we have a <u>clean inspection and violation</u>, with payoffs of (b_e+v_e, $-b_d-v_d$): E benefits from the fact that the inspection shows it complied (b_e) when in fact it violated (v_e); D loses from both the cost of the inspection (b_d) and the violation (v_d).

We summarize the parameters of the model, with their signs for each player:

<u>Evader (E)</u>	<u>Detector (D)</u>
b_e = benefit of clean inspection	$-b_d$ = cost of inspection
v_e = benefit of violation	$-v_d$ = cost of violation
$-c_e$ = cost of detection of violation	c_d = benefit of detection of violation

Observe that the benefits (positive) and the costs (negative) are in terms of E's payoffs, whereas the signs are reversed for D. Also 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 (clean inspection and violation). Similarly, a violation might be successful either because D does not inspect (unchecked violation) or because D inspects but does not uncover the violation (clean inspection and violation).

Our addition and subtraction of parameters at each outcome shows that, for D, a clean inspection and violation ($-b_d-v_d$) is worse than an unchecked violation ($-v_d$). In fact, the reverse might be the case. Although D incurs the cost of an inspection at the former outcome, at the latter outcome D made the attempt to check for a violation; while unsuccessful, D "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 clean inspection and 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 shall return to the question of how reversing the ordering of the costs to D of an unchecked versus an undetected violation affects our conclusions.

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

$$
\begin{aligned}
V(p,q) &= qpr(v_e-c_e, -b_d-v_d+c_d) + qp(1-r)(b_e+v_e, -b_d-v_d) + q(1-p)(v_e, -v_d) \\
&\quad + (1-q)p(b_e, -b_d) + (1-q)(1-p)(0, 0) \\
&= q(v_e, -v_d) + p(b_e, -b_d) + pqr(-b_e-c_e, c_d). \qquad (1)
\end{aligned}
$$

We prove in the Appendix that this form subsumes three possible inspection games--each with a uniquely determined Nash equilibrium--depending on the relationship of the detection probability r to the benefit and cost parameters (the stars signify the players' Nash equilibrium strategies and their resulting expected payoffs):

1. $\underline{r \le b_d/c_d}$: $p^* = 0$ and $q^* = 1$; $V_e^* = v_e$ and $V_d^* = -v_d$.
2. $\underline{b_d/c_d < r \le v_e/(b_e+c_e)}$: $p^* = 1$ and $q^* = 1$; $V_e^* = v_e + b_e - r(b_e + c_e)$ and $V_d^* = -v_d - b_d + rc_d$.
3. $\underline{r > b_d/c_d \text{ and } r > v_e/(b_e+c_e)}$: $p^* = v_e/[r(b_e+c_e)]$ and $q^* = b_d/(rc_d)$; $V_e^* = b_e v_e/[r(b_e+c_e)]$ and $V_d^* = -b_d v_d/(rc_d)$.

Note that the game 1 and game 2 equilibria are pure, whereas the game 3 equilibrium is mixed. It is easy to see that if $b_d/c_d < v_e/(b_e+c_e)$, all three games can occur (but not simultaneously) for different values of the parameters. If $b_d/c_d \ge v_e/(v_e+c_e)$, only games 1 and 3 can occur, because this inequality precludes an r that can satisfy both game 2 inequalities.

To interpret the conditions under which each game occurs, notice that the occurrence of each game depends on the relationship of the detection probability r to two ratios, b_d/c_d for D and $v_e/(b_e+c_e)$ for E. D's ratio of the absolute value 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 game occurs and its Nash equilibrium:

1. Low r: inspection always worthless; violation always worthwhile. 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. Intermediate r: inspection always worthwhile; violation always worthwhile. If D'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 D, as a best response, will always inspect.

3. High r: occasional inspections worthwhile; occasional violations worthwhile. If D'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's always complying. On the contrary, it is in E's interest to violate an agreement either some (game 3) or all (game 1 or game 2) the time. How often in the case of a game 3 equilibrium depends on $q^* = b_d/(rc_d)$, which is a direct function of the absolute value of D's cost of inspecting and an inverse function of D's detection probability and the benefit D receives from detecting a violation.

Thus, E should violate more often as D's cost of inspecting increases; and less often as D's detection probability and benefit from detection increase. For D, $p^* = v_e/[r(b_e+c_e)]$ indicates that it should inspect more often as E's benefit of a violation increases; and less often as D's detection probability, E's benefit from a clean inspection, and E's cost of detection of a violation increase. Significantly, the common factor in both equilibrium probabilities is that they are inverse functions of r: a greater detection probability decreases both E's violation probability and D's inspection probability.

It is perhaps dismaying that the Nash equilibria in none of the three games--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 unatttainable as an equilibrium strategy in the simultaneous form, can it be induced in another game form?

3. The Sequential Form

Consider a game form in which D 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 form. Moreover, after having made its announcement, D is committed to carrying out an inspection with probability p--or possibly there is some agent who performs this task.

There are at least two ways of interpreting p. First, assume that there is a single inspection site. A random device might be used to determine whether an inspection would actually be carried out in, say, a certain time period. Second, if there are n inspection sites, D's commitment would be to inspect pn sites, selected at random; in this case, the benefit and cost parameters discussed in section 2 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 D's prior move, E can now choose q, knowing p in advance. In a one-site, one-inspection game in the sequential form, this is not to say that

E knows whether or not it will be inspected but does know the probability that this event will occur. Consequently, games in the sequential form are still games of imperfect information: as in the simultaneous form, E and D, when they make their choices, do not assuredly know the choices of their adversaries. E's only additional information in the sequential form is that it knows the value of p when it makes its choice of q.

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

But D's announcement does alter the players' equilibrium strategies. As shown in the Appendix, D'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 $p < \bar{p}$, D can induce E to choose $q = 1$ (always violate) to maximize V_e; if $p > \bar{p}$, D can induce E to choose $q = 0$ (always comply) to maximize V_e.

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

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

When $q = 0$, from (1)

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

which D maximizes by making p as small as possible, subject to the constraint that $p > \bar{p}$. For small positive ϵ, the strategy pair ($p = \bar{p}+\epsilon$, $q = 0$) is an ϵ-equilibrium[2] in the sequential form that is Pareto-superior to the corresponding Nash equilibrium in the simultaneous form if

$$r > b_d/c_d \text{ and } r > v_e/(b_e+c_e) < v_d/c_d; \qquad (2)$$

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

$$v_e/(b_e+c_e) < r \leq b_d/c_d \text{ and } v_e/b_e < \bar{p} < v_d/b_d; \qquad (3)$$

(3) is a substantially more restrictive set of conditions than for a game 1 equilibrium in simultaneous form.

There are no other conditions that enable D, by choosing $p > \bar{p}$, to induce E to choose $q = 0$ such that the ϵ-equilibrium in sequential form is Pareto-superior to the Nash equilibrium in simultaneous form. Hence, either (2) or (3) is necessary and 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 D to induce E always to comply.

Observe that (2) requires a "high" r, whereas (3) requires an "intermediate" r and additional conditions on $\bar{p}$. In most realistic scenarios, $v_e > b_e$ (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 game 1, while a theoretical possibility, seems highly unlikely ever to occur.

We illustrate the difference between a Nash equilibrium in the simultaneous form and an ϵ-equilibrium in the sequential form with a game 3 example. This example highlights why a minimal kind of inducement by D may be viewed as unfair by E, forcing D 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 player, except for the change in sign: $b_e = b_d = 1$; $v_e = v_d = 2$; and $c_e = c_d = 3$. In addition, assume that $r = 3/4$. Because these parameter values satisfy (2), D'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

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

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

Clearly, inducement can reap major rewards for D. In practice, however, its effective use may require that D 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 $p = 2/3^+$, D 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 fairer than allowing all the benefits to go to D. D'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 D'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 D 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 fairness is operative (e.g., that prescribes dividing the profits of inducement equally, as illustrated in our example).

Indeed, D might threaten not always to comply <u>unless</u> there is a more equal division. This threat, if credible, would stabilize strategies at $(p > 2/3^+, q = 0)$ rather than $(p = 2/3^+, q = 0)$. The resulting outcome would be not only Pareto-superior to the Nash equilibrium in the simultaneous form but also less lopsided in favor of D.

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

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 unrealistic condition on $\overline{p}$. In the absence of such an induced equilibrium, the Nash equilibrium in the simultaneous form would presumably be chosen.

To be sure, under certain conditions (not given here), noncompliance may be induced as a Pareto-superior ϵ-equilibrium in the sequential form. But the circumstances under which D would choose a p to induce E always to violate an agreement seem quite absurd. Although D might want to embarrass E on occasion by detecting a violation, it is hard to conceive of plausible scenarios when it would be in D'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, are not in general Pareto-optimal. Given compliance on the part of E, D would always prefer to choose $p = 0$ and avoid the cost of inspection entirely. But the resulting outcome, while superior for D, 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 $\overline{p}$.

But how can D 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.

4. Conclusions

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

The mixed-strategy equilibrium of game 3 is probably the norm for most arms-control agreements: E sometimes violates and D sometimes inspects. In the latter situation, E sometimes is caught cheating and suffers because of it; D suffers from the violation but benefits from its detection.

This outcome may not be ideal from either player's viewpoint. Fortunately, it may be possible to circumvent it if D announces a (mixed) inspection strategy in advance and credibly commits itself to following through on it. In particular, D 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 D benefits from a rigorously adhered to treaty.

Curiously, D could do still better by not inspecting if E complied. While Pareto-superior, however, this outcome is not in equilibrium. Without being inspected, E would always violate, which is exactly what D'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 D's inducement strategy would not only deter all violations but also lead to a Pareto-superior outcome (compared with the Nash equilibrium in simultaneous form). These conditions are generally more restrictive than the conditions for a Nash equilibrium; in particular, they require that D'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 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 involve announcing how many would be randomly selected for inspection in a certain time period. Alternatively, if there was 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 an ongoing game, D 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 D and E. Alternatively, it may be in the interest of the parties that an independent inspection agency, which is both credible and politically acceptable, carried out the inspections.

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

But if, say, D'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. In addition, the model could be used to determine how many multiple inspections at a few sites, or few inspections at multiple sites, would be needed to induce a Pareto-superior compliance equilibrium, with this number written into a treaty. Thereby a sufficiently high p could be allowed by a treaty so that compliance is ensured, benefitting the parties over what the Nash equilibrium in the simultaneous game offers them.[4]

Finally, 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 D than a clean inspection and violation. Without going through a formal analysis of this reordering of two of D's payoffs, it is not difficult to show that if the costs of not inspecting a violation are greater for D, E complies more often in equilibrium in the simultaneous form. 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 D in the sequential form can ensure total compliance by E. As in the original game, it will not only be in D's interest but also E's that, given a sufficiently high detection probablity, D announce an inspection strategy to induce compliance.

APPENDIX

If the payoff function given by (1),

$$V(p, q) = = q(v_e, -v_d) + p(b_e, -b_d) + pqr(-b_e-c_e, c_d),$$

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

$$V(p, q) = q(A_e, A_d) + p(B_e, B_d) + pq(C_e, C_d),$$

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

---

Figure 2 about here

---

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

We will show that the game in Figure 2 always has a uniquely determined outcome. 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 $B_d = -b_d < 0$. In the Figure 2 game, there are two situations in which one player has a dominant strategy, resulting in a unique Nash equilibrium in pure strategies:

1. $\underline{B_d + C_d = -b_d + rc_d \leq 0}$. $\bar{I}$ is dominant (p = 0) because $B_d < 0$. $\bar{C}$ is E's best response (q = 1), resulting in $\overline{CI}$.

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

FIGURE 2

PAYOFF MATRIX OF GAME IN SIMULTANEOUS FORM

		Detector (D)		
		Inspect (I)	Don't inspect ($\bar{I}$)	
Evader (E)	Don't comply ($\bar{C}$)	$(A_e+B_e+C_e, A_d+B_d+C_d)$	(A_e, A_d)	q
	Comply (C)	(B_e, B_d)	(0, 0)	1-q
		p	1-p	

Now assume that neither situation 1 nor situation 2 holds, so the inequalities of both 1 and 2 are reversed. 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:

$$V_e = qA_e + pB_e + pqC_e; \qquad V_d = qA_d + pB_d + pqC_d.$$

Assume q = 0. Because

$$\frac{\partial V_d}{\partial p}(p, 0) = B_d = -b_d < 0,$$

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.[5] Because there must exist at least one Nash equilibrium, 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

$$\frac{\partial V_e}{\partial q} = A_e + pC_e; \quad \frac{\partial V_d}{\partial p} = B_d + qC_d.$$

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

$$p^* = -A_e/C_e = v_e/[r(b_e+c_e)]; \quad q^* = -B_d/C_d = b_d/(rc_d),$$

give the players' (mixed) equilibrium strategies. They are the strategies (p^* for D, 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 for games 1 and 2 in section 2, correspond to those that exist in situations 1 and 2. Likewise, the game 3 equilibrium corresponds to the mixed-strategy equilibrium that we have just derived.

In the sequential game form, D makes a prior announcement of its p. We therefore must consider, for any p chosen by D, 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

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

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

Now consider the sequential form of game 3. Inducement leads to a greater expected payoff for D if its expected payoff from inducement, given in section 3, is greater than its Nash equilibrium expected payoff, given in section 2:

$$-\bar{p}b_d > -b_d v_d/(rc_d) \Rightarrow \bar{p} < v_d/(rc_d).$$

From (4) it follows that

$$v_e/[r(b_e+c_e)] < v_d/c_d.$$

[Inducement automatically leads to a greater expected payoff for E than its Nash equilibrium because

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

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

$$r > b_d/c_d \text{ and } r > v_e/(b_e+c_e) < v_d/c_d.$$

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

$$-\bar{p}b_d > -v_d \Rightarrow \bar{p} < v_d/b_d;$$

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 form, therefore, requires

$$v_e/(b_e+c_e) < r \le b_d/c_d \text{ and } v_e/b_e < \bar{p} < v_d/b_d.$$

NOTES

[1]We assume that there is no possibility of a type I 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.

[2]An ϵ-equilibrium is defined with respect to the game tree, in which D 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 D to do so if E always complies. This equilibrium is also a Stackelberg equilibrium, with D the leader and E the follower.

[3]Different effects of inducement in arms-control inspection games have been analyzed by Maschler (1966, 1967), Brams (1985), Avenhaus (1986), Fichtner (1986), Brams and Davis (1987), and Brams and Kilgour (1986, 1987, 1988); see also Wittman (1989) and O'Neill (1990) 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 detector, and (4) illustrating the issue of fairness that might arise in dividing the surplus at such outcomes. As a consequence, we are able to distinguish different possible games and their Nash equilibria in the simultaneous form, and different threshold conditions that result in Pareto-superior ϵ-equilibria in the sequential form.

---

[4]The INF (1988) treaty allows for specific numbers of annual inspections at agreed-upon sites. Game-theoretic models for the recursive allocation of cheating resources by an evader, and inspections by a detector, are given in Brams and Kilgour (1989) and Kilgour (1990).

[5]The lack of such an equilibrium can also be seen from Figure 2. 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 (1986). Safeguards Systems Analysis. New York: Plenum.

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 (1989). "Optimal Cheating and Inspection Strategies under INF." Preprint.

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.

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, forthcoming.

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.

O'Neill, Barry (1990). "Why a Better Verification Scheme Can Give More Ambiguous Evidence." International Studies Quarterly, forthcoming.

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