Euro. Jnl of Applied Mathematics (2010), *vol.* 21, *pp.* 441–458. © Cambridge University Press 2010 441
doi:10.1017/S0956792510000136

Inspection games with long-run inspectors

LUCIANO ANDREOZZI

Università degli Studi di Trento Via Inama, 5 38100 Trento, Italy
email: luciano.andreozzi@unitn.it

(*Received 8 May 2009; revised 7 April 2010; accepted 7 April 2010; first published online 12 May 2010*)

The inspection game has been widely used as a model for crime deterrence. In this article, we shall present a broad review of the literature and an original result. The original result concerns the interaction between a single, long-run inspector and a population of short-run, boundedly rational would-be criminals. This approach gives rise to an optimal control problem that can be solved explicitly. We show that a forward-looking inspector obtains the same payoff he would obtain if he could commit to a probability of inspection. The consequences of this result for the economics of crime deterrence are also explored.

1 Introduction

The large literature on crime deterrence pioneered by Gary Becker [4] is mostly based on decision theory. The standard model is based on the assumption that a rational agent can commit a crime and faces a fixed probability of being discovered and punished. He will commit the crime if its expected benefits are larger than the expected sanctions. Two factors influence such a decision: the probability with which his crime will be discovered and the severity of punishment. An increase in any of these two dimensions (a larger probability of apprehension or harsher punishment) will reduce the probability that a crime will be committed.

Over the years a small game-theoretic literature formed, which challenges this conclusion. This approach was originally developed in a series of papers published over a short time-span: Graetz *et al.* [13] Tsebelis [24, 25], Holler [18] and Wittman [27, 28]. The common theme of these papers is that decision theory offers a poor instrument to treat crime deterrence.[1] The reason is that any criminal decision involves at least two actors: a would-be criminal and a law enforcer. The former decides the probability of committing a crime, the latter the probability that the crime will be discovered. As Tsebelis [24] pointed out, both these probabilities should be treated as endogenous in crime deterrence models and this requires a game-theoretic framework.

The game-theoretic approach to crime deterrence is based on a simple simultaneous 2×2 game known as inspection game. The first player (the 'Public') decides whether to violate or not violate a law, while the second player (the 'Inspector') decides whether to

[1] For example, Graetz *et al.* [13] notice that "the introduction of the law enforcement agency into a game-theoretic analysis of tax compliance offers considerable opportunity for insights and predictions that are simply not possible in the standard economic analysis of law enforcement". (p. 29)

inspect the other player or not. Inspecting is assumed to be costly, so the Inspector prefers not to inspect if he believes that no crime is going to be committed. Respecting the law is also costly, which implies that if no inspection takes place, then the law is violated.

The inspection game has a single Nash equilibrium in mixed strategy. The early research on this game has focused on the counter-intuitive comparative static properties of this equilibrium. One can easily show that increasing the severity of penalties reduces the frequency of inspections but leaves the frequency of offences unaltered. This result is usually cited as a proof that the standard, decision-theoretic approach to crime deterrence (which predicts the opposite effect) is flawed.

The use of the mixed strategy Nash equilibrium as a solution concept requires quite restrictive assumptions. First, the game must be played simultaneously, meaning that each player chooses without knowing the choice made by the other player. Second, the players must know each other's payoffs and be rational enough to compute their equilibrium probabilities. Both these assumptions are questionable in most of the situations the inspection game is supposed to model. For example, it is almost never the case that would-be criminals know exactly the real incentives policemen have to enforce the law. A few papers published over the following years relaxed either of the two.

A first strand of the literature gets rid of the simultaneity assumption and works with the sequential version of the inspection game in which the Public chooses after having observed the strategy chosen by the Inspector. This is the so-called Stackelberg version of the game, in which the Inspector acts as the leader and can commit himself to play a certain strategy. Bianco [5] and Hirshleifer and Rasmussen [16] study the case in which the inspector can only commit to play a *pure* strategy, while Andreozzi [2] and Cox [7] allow for the possibility that he commits to play a mixed strategy.

A second strand of the literature relaxes the assumption of rationality and relies on evolutionary game theory instead. Cressman [8] and Andreozzi [1] present a model in which the inspection game is played repeatedly by agents drawn randomly from two large populations of would-be criminals and inspectors. Agents are not assumed to be rational. Rather, they adjust their behaviour over time by switching from less profitable to more profitable strategies. Both these papers show that the frequency of law infractions tends to oscillate indefinitely around its equilibrium value, something which could not be anticipated with more traditional models.

The present paper has two main objectives. We shall first propose a broad review of the results discussed so far. Section 2 reviews the early results based on Nash equilibrium. Section 3 presents the results on the Stackelberg version of the game and Section 4 presents the evolutionary approach.

Second, in Section 5 we shall propose an original model which is a combination of the two main approaches explored so far. The idea of the model is that in many situations the Inspector is a single long-run agent, such as a tax authority or a police force, who faces a large population of small and myopic would-be criminals. This implies that there is a fundamental asymmetry between the probability of inspection and the frequency of law infractions. The former is decided by a single agent who considers the future consequences of his current choices. The inspector will anticipate, for example, that if he reduces the frequency of inspections today, there will be an increase of crime tomorrow. On the contrary, the frequency of offences is the result of many decisions taken by a multitude

Table 1. *The inspection game*

	Inspect	Not Inspect
Violate	a_{11},b_{11}	a_{12},b_{12}
Not Violate	a_{21},b_{21}	a_{22},b_{22}

of dispersed actors. Since each would-be criminal cannot influence the overall rate of law violations, when deciding to commit a crime he will only consider his own current payoff (see Fudenberg and Levine [10, 11] for two classical papers in this vein).

To model this kind of situation, we imagine that a single Inspector plays repeatedly the inspection game with agents drawn randomly from a large population of identical would-be criminals. The Inspector is assumed to be perfectly rational, meaning that he maximizes the current value of his future stream of payoffs. Would-be criminals are not rational. Their behaviour is described by the same dynamic adjustment process employed in standard evolutionary models.

The inter-temporal maximization problem faced by the Inspector gives rise to an optimal control problem that can be solved explicitly. We obtain the following results. First, the optimal solution is of the bang-off variety. This implies that, at variance with the standard evolutionary approach, a rational Inspector will not generate oscillations in the enforcement of law. This result is not intuitively obvious, because one cannot exclude on *a priori* ground that a rational inspector find it convenient to alternate periods with a high frequency of inspections to periods with low frequency. This would be the case, for example, if the solution to the control problem was of the bang-bang variety.

Second, we show that the solution to the optimal control problem approaches the Stackelberg solution of the game as the Inspector becomes more patient. That is, a rational, forward-looking Inspector will behave *as if* he could commit to a probability of inspection. This implies that incentives for capturing criminals can have the same perverse effects first studied by Andreozzi [2]. In particular, increasing the premium the inspector gets for each infraction he discovers might produce a larger number of infractions in equilibrium.

2 The inspection game

Consider the following situation. One agent, A, must decide whether to *Violate* or *Not Violate* a given law or regulation, knowing that there is a second agent B (the 'Inspector') who might *Inspect* him or not. An agent who has been inspected when in non-compliance suffers a penalty which is severe enough to deter the offence. Hence, if A expects B to play *Inspect*, then he would comply. However, both complying and inspecting are costly, so that if B knows that A is not going to play *Violate* anyway, he will rather play *Not Inspect*. Similarly, if A knows that B is not going to *Inspect*, he will play *Violate*. These assumptions are summarized in the bimatrix game in Table 1, in which the entries fulfill the following inequalities: $a_{21} > a_{11}$ (A prefers to play *Not Violate* if B plays *Inspect*); $a_{12} > a_{22}$ (A prefers to play *Violate* if B plays *Not Inspect*); $b_{11} > b_{12}$

(B prefers to play *Inspect* if A plays *Violate*); $b_{22} > b_{21}$ (B prefers to play *Not Inspect* if A plays *Not Violate*). Players can also employ mixed strategies. We shall indicate with $\pi^A(p,q)$ and $\pi^B(p,q)$ A and B's expected payoffs when A plays *Violate* with probability p and B inspects with probability q.

One can easily check that, given our assumptions on a_{ij} and b_{ij}, the inspection game has a single mixed-strategy Nash equilibrium (p^*, q^*):

$$(p^*, q^*) = \left(\frac{b_{22} - b_{21}}{b_{22} - b_{21} - b_{12} + b_{11}}, \frac{a_{12} - a_{22}}{a_{12} - a_{22} + a_{21} - a_{11}}\right). \tag{2.1}$$

The early research on the inspection game (Tsebelis [24, 25] and Wittman [27]) has focused on the way in which the equilibrium probabilities of inspection and violation are affected by changes in the entries of the two payoff matrices. When the severity of penalties for a discovered crime is increased, the entry a_{11} is reduced. When the premium the Inspector gets for each law infraction he discovers is increased, the entry b_{11} increases. Tsebelis [24, 25] and Wittman [27] main result was that these changes in a_{11} and b_{11} affect in a counter-intuitive way the equilibrium values p^* and q^*.

Claim 1 In the inspection game in Table 1:

(1) Increasing penalties (i.e. reducing a_{11}), leaves the frequency p^* of law violation unchanged and reduces the frequency of inspections q^*;

(2) Increasing incentives for Inspectors to play *Inspect* (i.e. raising b_{11}), leaves the frequency of inspections q^* unchanged and reduces the frequency of law infractions p^*.

Claim 1 is a consequence of a well-known fact in game theory. In any mixed strategy Nash equilibrium, each player's strategy is calculated as to make the other player indifferent between the pure strategies he employs with positive probability. In (2.1) p^* depends upon coefficients b_{ij}, while q^* depends upon coefficients a_{ij}. It follows that any change in A's payoffs will only bring about a change in B's equilibrium strategy and *vice versa*.

Claim 1 is based on the hypothesis that the appropriate concept solution for the simultaneous, one-shot inspection game is the mixed-strategy Nash equilibrium. Not everyone agrees on this point. Holler [18] resumed an old perplexity about mixed-strategy Nash equilibria (see also Holler [17]). The reason for this perplexity is that mixed-strategy Nash equilibria are never strict.[2] This means that if A expects B to play his Nash equilibrium strategy q^*, then he is indifferent between *Violate*, *Not Violate* or any mix of these two strategies. As a consequence, there seems to be no reason for him to play the particular mixed strategy p^* that makes B indifferent between *Inspect* and *Not Inspect*. The same reasoning applies to player B.

Holler [17, 18] followed Harshanyi [14, 15] in believing that in the inspection game maximin strategies where more reasonable than Nash equilibrium strategies. Player A's maximin strategy is the strategy that maximizes A's minimum payoff, that is the payoff

[2] In a two player game, a pair of strategies (p^*, q^*) is a *strict* Nash equilibrium if each player prefers strictly to play his Nash equilibrium strategy if she expects the other player to do the same. A Nash equlibrium (p^*, q^*) is weak if at least one player has an alternative strategy $p' \neq p^*$ or $q' \neq q^*$ such that he is not worse off if he unilaterally switches to p' or q'.

A obtains if B chooses the worst strategy given A's choice. In zero-sum games Nash equilibria always involve the use of maximin strategies. In non-zero sum games, however, maximin strategies do not represent Nash equilibria. This implies that either A's maximin strategy is not a best response to B's maximin strategy, or vice versa. So there seems to be no reason why players should expect the maximin strategies to be the outcome of a non-zero-sum game.

However, the inspection game belongs to a special class of games known as *unprofitable games*. To see what an unprofitable game is, let $\bar{\pi}_i$ be the payoff player i can guarantee himself by playing is maximin strategy. A game is unprofitable if each player i obtains in every Nash equilibrium a payoff which is not larger than $\bar{\pi}_i$. This implies that when a Nash equilibrium is expected to be the outcome of an unprofitable game, each player gets $\bar{\pi}_i$ either if he plays his Nash equilibrium strategy or if he plays his maximin strategy. There is an important difference between these two strategies, though. By using his Nash equilibrium strategy, player i obtains $\bar{\pi}_i$ only if also the other player plays his equilibrium strategy. By playing his maximin strategy, player i guarantees himself $\bar{\pi}_i$ irrespective of the other player's choice. This is the reason why Harshanyi [14, 15] believed that in unprofitable games maximin strategies were more reasonable than Nash equilibria.

Lets see what maximin strategies are like in an inspection game. Player A chooses a probability of violation p such that his payoff is the same independently of the choice made by B. Similarly, B chooses a probability of inspection q such that he gets the same payoff regardless of A's choice.[3] Formally, lets first define $\hat{p}^+$ and $\hat{q}^+$ as follows:

$$(\hat{p}^+, \hat{q}^+) = \left(\frac{a_{22} - a_{21}}{a_{11} - a_{12} - a_{21} + a_{22}}, \frac{b_{22} - b_{12}}{b_{11} - b_{12} - b_{21} + b_{22}} \right). \tag{2.2}$$

The two maximin strategies p^+ and q^+ for player A and B are thus defined as

$$p^+ = \begin{cases} 0 & \text{if } \hat{p}^+ < 0, \\ \hat{p}^+ & \text{if } \hat{p}^+ \in [0,1], \\ 1 & \text{if } \hat{p}^+ > 1, \end{cases} \qquad q^+ = \begin{cases} 0 & \text{if } \hat{q}^+ < 0, \\ \hat{q}^+ & \text{if } \hat{q}^+ \in [0,1], \\ 1 & \text{if } \hat{q}^+ > 1. \end{cases} \tag{2.3}$$

One can easily see that, when $p^+ \in [0,1]$, if A plays *Violate* with probability p^+, then he gets the same payoff regardless of B's probability of inspection q. Also, this payoff is the same A would get at the Nash equilibrium (p^*, q^*).

Holler's results are summarized by the following:

Claim 2 If the two players employ their maximin strategies, and the maximin strategies are mixed (that is if $p^+, q^+ \in (0,1)$) then

(1) increasing the severity of punishment (lower a_{11}) will reduce crime;

(2) increasing b_{11} will not reduce crime but will reduce the frequency with which the Inspector plays *Inspect*.

[3] Formally, A's maximin strategy solves the following equation $\pi^A(p,1) = \pi^A(p,0)$ and B's maximin strategy solves $\pi^B(1,q) = \pi^B(0,q)$.

Table 2. *An inspection game with pure maximin strategies for both players*

	Inspect	Not Inspect
Violate	$r - f, pr - c$	$r, 0$
Not Violate	$0, -c$	$0, 0$

The differences between Nash equilibrium and maximin predictions are best illustrated by an example. Saha and Poole [22] discuss the inspection game in Table 2. An undiscovered law infraction yields the member of the public a revenue r. However, if the infraction is discovered a fine $f > r$ is levied. Similarly, the Inspector pays a cost $c > 0$ to play *Inspect*, but if he discovers a law violation he gets a premium $pr > c$. Both players get nothing if they play *Not Violate* and *Not Inspect*, respectively. This game has a single mixed-strategy Nash equilibrium $(\frac{r}{f}, \frac{c}{pr})$, in which both players get a payoff equal to 0. However, both players could guarantee themselves the same payoff by simply playing their (pure) maximin strategies *Not Violate* and *Not Inspect*. Notice that (*Not Violate*, *Not Inspect*) is not a Nash equilibrium, because if B plays *Not Inspect* A's best response is *Violate*.

This example clearly illustrates that both Nash equilibrium and maximin look suspicious as predictors of how people play the inspection game. In the rest of this paper we shall investigate some of the lines of research that have been proposed to overcome this kind of difficulty.

3 The inspection game's Stackelberg equilibrium

The results discussed in the previous section were based on the assumption that the game is played only once, by a single inspector and a would-be criminal. An early criticism to this approach (see Hirshleifer & Rasmussen [16] and Cox [7]) is that the interaction between law enforcers and the public is seldom symmetric. While law enforcers are usually organized into police forces and tax authorities, the public is made by a large number of unorganized individuals. This clearly changes the nature of interaction. An organized police force might try to build a reputation for enforcing the law. If a police force plays consistently *Inspect*, for example, the public will come to expect it will use the same strategy in the future and will eventually play *Not Violate*.

The simplest way to model this kind of interaction is to assume that the game is played sequentially. (See Mailath & Samuelson [20] for a recent textbook presentation of reputation models of this kind.) The Inspector first makes a commitment to use one of his strategies and then the would-be criminal chooses a best response to it. This is the so called Stackelberg version of the game, in which the Inspector acts as the *leader* while the would-be criminal is the *follower*. A crucial difference arises whether one assumes that the leader can commit to use a pure or a mixed strategy. Both these possibilities have been explored in the literature, and we shall deal with them in turn.

3.1 Commitment to a pure strategy

Hirshleifer and Rasmussen [16] assume that the Inspector can only commit to play a pure strategy (*Inspect* or *Not Inspect*). Since Violate is a best response to *Not Inspect* and *Not Violate* is a best response to *Inspect*, there are only two admissible outcomes of the game: (*Violate*, *Not Inspect*) and (*Not Violate*, *Inspect*). Since B chooses first, he will commit herself to play *Inspect* if $b_{21} > b_{12}$ and he will opt for *Not Inspect* otherwise. Hirshleifer and Rasmussen [16] interpret this result as follows. A rational Inspector will only enforce those laws whose cost of enforcement is worth paying (for himself). When $b_{21} < b_{12}$, the Inspector prefers to play *Not Inspect*, and allow the public to violate the law, rather than obtaining the public's compliance with the law at the cost of playing *Inspect*.

It is worth noticing that this simple result contains another version of Tsebelis' original claim that punishment has no effects on crime. To see this, consider that B's decision to commit herself to *Inspect* or *Not Inspect* depends only upon B's own payoffs. It follows that one cannot reduce law infractions by increasing penalties: if $b_{12} > b_{21}$ B will not inspect and A will violate the law, no matter how severe the punishment for violators is. In this respect, the sequential version of the game is no different from the simultaneous move version.

3.2 Commitment to a mixed strategy

Andreozzi [2] consider the sequential version of the inspection game in which the Inspector B can commit himself to play *Inspect* with probability q before the game is played.[4] The crucial assumption in this model is that when deciding whether to violate or not the law, A observes B's probability of inspection but not its actual realization.

Let $BR^A(q)$ be A's best reply correspondence, that is A's probability of violating the law when he expects B to play *Inspect* with probability q. We have that

$$BR^A(q) = \begin{cases} 1 & \text{if} \quad q < q^*, \\ [0,1] & \text{if} \quad q = q^*, \\ 0 & \text{if} \quad q > q^*. \end{cases} \tag{3.1}$$

A will will violate the law ($BR^A(q) = 1$) if $q < q^*$ and he will not violate the law ($BR^A(q) = 0$) if $q > q^*$. When $q = q^*$, A is indifferent between his two pure strategies and therefore his choice is indeterminate. However, as it is customary in this kind of literature, one can assume that in this case the follower chooses the strategy which yields the largest payoff to the leader. The reason is that that by choosing a mixed strategy arbitrarily close to q^* the leader can always make the follower strictly prefer one of his pure strategies.[5] So, the leader can always induce the follower to choose the strategy which yields the highest payoff to himself.

[4] Cox [7] contains a closely related approach, with similar results. The main focus of the paper is on situations in which the members of the public have different payoff functions. It shows that it is not, in general, true that increasing fines has no effects on crime.

[5] To see this, consider that for any $\varepsilon > 0$, a probability of inspection $q^* + \varepsilon$ will induce A to play *Not Violate*, and $q^* - \varepsilon$ will induce A to play *Violate*.

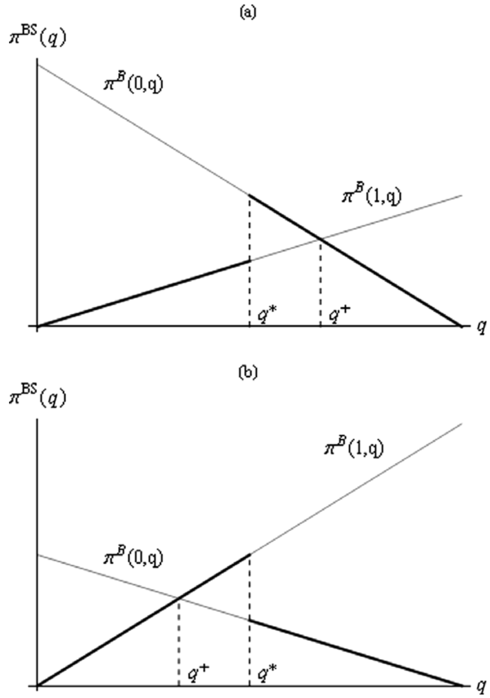

FIGURE 1. Computation of the Stackelberg equilibrium for two inspection games.

Let $\pi^{BS}(q)$ be B's payoff when he can act as a leader and choose a probability of inspection q. We have that

$$\pi^{BS}(q) = \begin{cases} q(b_{21} - b_{22}) + b_{22} & \text{if } \ q > q^*, \\ q(b_{11} - b_{12}) + b_{12} & \text{if } \ q < q^*, \\ \max(q(b_{21} - b_{22}) + b_{22},\ q(b_{11} - b_{12}) + b_{12}) & \text{if } \ q = q^*. \end{cases} \tag{3.2}$$

The Inspector chooses q as to maximize $\pi^{BS}(q)$. Andreozzi [2] main proposition is

Proposition 3 *In a general case in which $q^* \neq q^+$, the sequential version of the Inspection Game has a single subgame perfect Nash equilibrium. This equilibrium is $(0, q^*)$ if $q^* < q^+$ and $(1, q^*)$ if $q^* > q^+$.*

Figure 1 illustrates the content of this proposition. It shows the Inspector's payoff in two different games as a function of the probability q with which he inspects. The gray lines represent B's payoff when he inspects with probability q and A plays *Violate* ($\pi^B(1, q)$) and *Not Violate* ($\pi^B(0, q)$). When $q = q^+$, B gets the same payoff regardless of A's choice. q^+ is thus B's maximin strategy as defined in (2.3). The thick line represents the Inspector's payoff when A plays a best reply to q. The Inspector's payoff is discontinuous at q^* because A plays *Violate* if $q < q^*$ and *Not Violate* if $q > q^*$.

Player A has the same payoffs in both inspection games (a) and (b), so that B's Nash equilibrium strategy is the same: $q^* = \frac{1}{2}$. B's payoff differ in the two games and are choosen so that $q^+ > q^*$ in game (a) and $q^+ < q^*$ in game (b). The picture shows that in both cases the Inspector maximizes his payoff by playing q^*. However, this requires A to play *Not Violate* in game (a) and *Violate* in game (b). The graphic should also make clear the role of condition $q^* \lesseqgtr q^+$.

Proposition 3 shows that, depending on the payoff function of the Inspector, the sequential version of the inspection game admits two kinds of equilibria. In the first kind of equilibrium (represented in (a)) the public respects the law. This equilibrium has a natural economic interpretation. The Inspector plays *Inspect* with a probability slightly larger than q^*, which is the smallest probability that will induce the public not to violate the law. Notice that, since inspecting is costly, the Inspector has no reason to use this strategy with a probability which is larger than the minimum sufficient to deter crime.

In the second kind of equilibrium (represented in (b)) the public violates the law and the Inspector plays *Inspect* with a probability slightly smaller than q^*. This probability is strictly positive, but insufficient to deter crime. This equilibrium looks unreasonable at first, but it has a clear intuition. It arises in those cases in which the only reason for the Inspector to play *Inspect* is a premium he gets for each law infraction he discovers. For example, it would be the equilibrium in the game in Table 2. Clearly, in a world without law infractions the Inspector would get no premiums. Hence, he has no reason to play Inspect frequently enough to deter crime. On the other hand, if he always plays *Not Inspect*, he would not get premiums either. His optimal policy would then be to play *Inspect* with the largest probability compatible with the public to play *Violate*, that is a slightly smaller probability than q^*.

We now turn to comparative statics. In the Stackelberg equilibrium the public plays *Violate* only if $q^+ < q^*$, so we shall assume this to be the case.[6] A change in a_{11} or b_{11} will reduce crime only if it is sufficiently large to reverse this inequality. Because of Claim 1 and 2, a decrease of a_{11} will reduce q^* and leave q^+ unaltered. As a consequence, if a_{11} becomes sufficiently small (i.e. if penalties become sufficiently severe) then the public will stop playing *Violate* in equilibrium. This result is in agreement with the standard approach to crime deterrence. On the other hand, increasing the premiums for inspectors (larger b_{11}) reduces q^+ and leaves q^* unaltered. As a consequence, if initially $q^+ < q^*$, any increase in b_{11} will fail to reduce crime. This is the main result in Andreozzi [2]. One cannot control crime by giving inspectors premiums for catching criminals, if they can control the probability of inspection. In fact, they would choose a sufficiently small probability so that the crime is committed and they get the premium. Section 5 obtains a similar result in a different setting.[7]

[6] When $q^+ > q^*$ in equiibrium $p = 0$ and hence the problem of crime deterrence is already solved.

[7] Friehe [9] shows that this result does not hold in a model in which the inspector's payoffs are correlated to the public's payoffs. Such correlation could emerge, for example, if the possibility of corruption is considered. See the consequences of correlations between the payoffs of the two players is analysed also in Pradiptyo [21].

4 The evolutionary approach

All the results presented so far were based on the assumption that players were rational and had all the information required to compute their best response to other player's choices. These hypotheses are rarely satisfied in real-life situations. People are rarely aware of other people's payoff, and they are usually not rational enough to compute their Nash equilibrium strategies. The evolutionary approach to game theory is based on a set of alternative hypotheses that bypass all these difficulties. It assumes that players lack any information concerning other people's preferences and are not rational enough to compute their equilibrium strategies. However, it assumes that the same game is played repeatedly over time, which allows players to *learn* which strategies yield the largest payoffs and to *imitate* other players who employ more successful strategies. A combination of learning and imitation leads to a dynamic process such that those strategies which perform above the average will tend to be played more often than less successful alternatives. (See Weibull [26] for a thorough treatment of the evolutionary approach to game theory.)

The first treatment of the inspection game in terms of evolutionary game theory is due to Cressmann *et al.* [8].[8] To see how their model works, imagine two large populations of identical agents which, with a small abuse of notation we shall denote as A and B. The inspection game is played by pairs of agents drawn randomly, one from population A (the Public) the other from population B (the Inspector). Each agent only employs a pure strategy (*Violate* or *Not Violate* if a member of population A, *Inspect* or *Not Inspect* if a member of population B). The *state* of the two populations is a pair (p, q), where p is the fraction of agents who play *Violate* in population A and q is the fraction of agents who play *Inspect* in population B.

Consider now an agent in population A who plays *Violate*. Since he is equally likely to be matched with any agent in population B, his expected payoff is $a_{11}q + a_{12}(1-q) = \pi^A(1, q)$. In other words, he gets the same payoff he would get by playing against a single B agent who uses a mixed strategy putting probability q on *Inspect*. A similar definition can be given for $\pi^B(p, 1)$. Within populations A and B, the *average* payoff will be given by

$$p\ \pi^A(1, q) + (1-p)\ \pi^A(0, q) = \pi^A(p, q),$$
$$q\ \pi^B(p, 1) + (1-q)\ \pi^B(p, 0) = \pi^B(p, q).$$

Cressmann *et al.* [8] model the evolution of the state of the two populations by means of the so-called replicator dynamics:

$$\dot{p} = p(\pi^A(1, q) - \pi^A(p, q)), \tag{4.1}$$
$$\dot{q} = q(\pi^B(p, 1) - \pi^B(p, q)). \tag{4.2}$$

In the replicator dynamics the fraction of agents playing a pure strategy grows at a rate proportional to the difference between the payoff obtained by that strategy and the average payoff within the population.[9] With simple algebra (see e.g. Andreozzi [1]) one

[8] Arce and Gunn [3] discuss an interesting version of this model, in which there is a single population and roles (inspector and would-be criminal) are randomly assigned.

[9] The replicator dynamics was originally proposed as a model of biological evolution, but is has frequently been employed within the social sciences (see Fudenberg and Levine [12]).

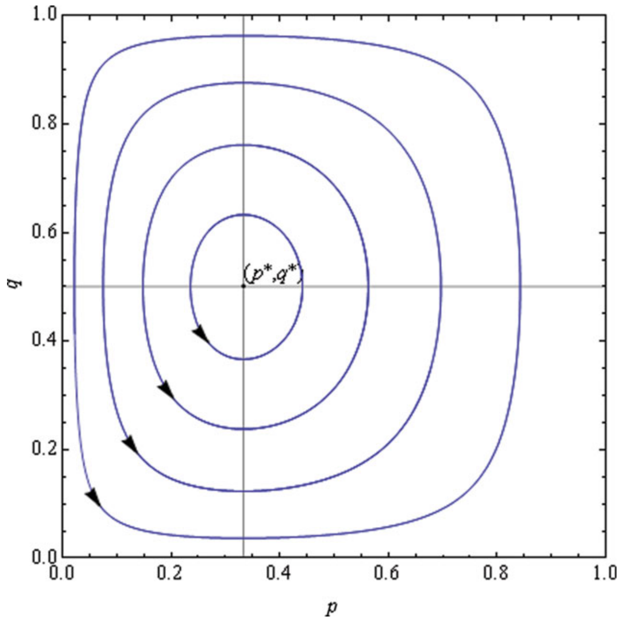

FIGURE 2. (Colour online) Oscillations in law enforcement.

can put these equations into the more convenient form:

$$\dot{p} = p(1-p)(\pi_1^B(q) - \pi_2^B(q)) = \alpha\ p\ (1-p)(q^* - q), \tag{4.3}$$

$$\dot{q} = q(1-q)(\pi_1^A(p) - \pi_2^A(p,q)) = \beta\ q\ (1-q)(p - p^*), \tag{4.4}$$

where $\alpha \equiv -a_{22} + a_{21} + a_{12} - a_{11} > 0$ and $\beta = b_{22} - b_{21} - b_{12} + b_{11} > 0$.

A *steady state* is a pair (p,q) such that $\dot{p} = \dot{q} = 0$. From (4.3) and (4.4) it is clear that the only non-trivial steady state is the Nash equilibrium (p^*, q^*).[10] It is a well known fact in evolutionary game theory that mixed strategy Nash equilibria are not asymptotically stable under the replicator dynamics. Figure 2 shows some orbits for an inspection game in which $(p^*, q^*) = (\frac{1}{3}, \frac{1}{2})$.

This result might be taken as a proof that comparative static exercises based on the Nash equilibrium values are unwarranted because the frequency of law infractions will in general be far away from its equilibrium value. However, it can be shown that *on average* the frequency with which each pure strategy is played equals its Nash equilibrium value. Let T be the period of oscillations. One can easily prove that[11]

$$\frac{1}{T}\int_0^T p(\tau)d\tau = p^* \quad \text{and} \quad \frac{1}{T}\int_0^T q(\tau)d\tau = q^*. \tag{4.5}$$

It follows that evolutionary approach lends some justification to the use of the mixed-strategy Nash equilibrium for comparative statics. Two elements must, however, be kept

[10] Every state in which in both populations only one strategy is played is trivially a rest point. This reflects the fact that the replicator dynamics cannot introduce new strategies into a population.

[11] See Weibull [26], chapter 5 for a proof of this fact.

in mind. First, evolutionary models predict that the Nash equilibrium frequencies will only be observed *on average*. The actual frequencies of law infractions will oscillate over time. This creates a kind of problem which is absent from more standard approaches based on stronger assumptions of rationality. For example, a policy maker might be interested in stabilizing the level of law infractions to its equilibrium value by reducing oscillations. To the best of my knowledge, no research has been done so far on how such a result could be obtained. Second, the evolutionary approach presented here assumes a symmetric interaction in which both the public and the law enforcers are represented by large populations of myopic and boundedly rational agents. As we showed in Section!3, this hypothesis is not realistic in many applications. The next section addresses this problem.

5 A mixed model

Consider the following situation. There is a large population A of would-be criminals and a single law enforcer B. The inspection game in Table 1 is played repeatedly by player B and an agent drawn randomly from population A. As in the evolutionary model discussed in the previous section, each agent in population A adopts a pure strategy. The state of population A at time t is represented by the fraction $p(t)$ of agents who play *Violate*, and changes over time according to the differential equation (4.3).

Player B is both rational and forward-looking. He will then choose a path for the frequency $q(t)$ with which he plays *Inspect* that maximizes the actual value of the future stream of payoff:[12]

$$\pi^B(p(t), q(t)) = b_{11}\ p\ q + b_{12}\ p\ (1-q) + b_{21}(1-p)\ q + b_{22}(1-p)(1-q).$$

Formally, player B solves the following optimal control problem:

$$\max_q \left(\int_0^\infty e^{-rt}\ \pi^B(p(t), \quad q(t))dt, \right. \tag{5.1}$$

$$s.t. \quad \dot{p} = \alpha p(1-p)(q^* - q), \tag{5.2}$$

$$\left. p(0) = p_0 \in (0,1), \qquad q(t) \in [0,1], \qquad \forall t \in [0, \infty \right),$$

where r is B's time discount rate.[13]

There are two reasons why this optimal control problem is of interest. First, from the evolutionary approach discussed in the previous section we know that when a population of myopic inspectors faces a population of would-be criminals, the frequency of law violations oscillates over time. One cannot exclude that similar oscillations would be observed even when the Inspector is a rational forward-looking player. Technically, this would be the case if the Inspector's optimal control problem had a bang-bang solution.

[12] For the sake of a simpler notation, we replace $p(t)$ and $q(t)$ with p and q whenever this creates no ambiguities.

[13] Notice that the initial fraction of law infractions p_0 has been restricted to the open interval $(0, 1)$. This is to avoid that the population gets locked in a steady state in which $\dot{p} = 0$ even if the two pure strategies yield different payoffs.

Second, one would want to know whether a more forward-looking Inspector (smaller r) would induce fewer law violations than a more myopic one. Intuitively, one should expect that a forward-looking Inspector will not stop playing *Inspect* when the crime rate is low, because he anticipates that a lower frequency of inspections will bring about more crime in the future. If this where the case, reducing the inspector's myopia would reduce the frequency of crime. We shall show that this is not necessarily the case and that, in fact, the opposite effect might be observed.

The current-value Hamiltonian for the Inspector's optimal problem is

$$H(p(t), q(t), m(t)) = \pi^B(p(t), q(t)) + m(t)[\alpha p(1-p)(q^* - q)]. \tag{5.3}$$

The costate variable $m(t)$ satisfies:[14]

$$\begin{aligned} \dot{m} &= rm - \frac{\partial H}{\partial p} = rm - \frac{\partial \pi}{\partial p} - m\alpha(1-2p)(q^* - q) \\ &= m(r - \alpha(q^* - q)(1-2p)) - \beta(q - q^+), \end{aligned} \tag{5.4}$$

A *stationary* optimal path is a triple $(\bar{p}, \bar{q}, \bar{m})$ such that $\dot{p} = \dot{m} = 0$ for $p = \bar{p}$, $q = \bar{q}$ and $m = \bar{m}$, while for every $q \neq \bar{q}$, $H(\bar{p}, \bar{q}, \bar{m}) > H(\bar{p}, q, \bar{m})$. If initially $p_0 = \bar{p}$, a rational inspector B maximizes his payoffs by setting $q = \bar{q}$, so that $\dot{p} = \dot{m} = 0$.

Lemma 4 *The optimal control problem (5.1) has a unique stationary optimal path:*

$$\bar{q} = q^*, \tag{5.5}$$

$$\bar{m} = \frac{\beta(q^* - q^+)}{r}, \tag{5.6}$$

$$\bar{p} = \begin{cases} \dfrac{-(\bar{m}\alpha - \beta) + \sqrt{(\bar{m}\alpha - \beta)^2 + 4\bar{m}\alpha\beta p^*}}{2\bar{m}\alpha} & \textit{if} \quad q^+ < q^*, \\ \dfrac{-(\bar{m}\alpha - \beta) + \sqrt{(\bar{m}\alpha - \beta)^2 + 4\bar{m}\alpha\beta p^*}}{2\bar{m}\alpha} & \textit{if} \quad q^+ > q^*. \end{cases} \tag{5.7}$$

Proof To see that $(\bar{q}, \bar{m}, \bar{p})$ is a stationary optimal path for the control problem (5.1) consider first that for any $p \in (0, 1)$, $\dot{p} = 0$ iff $q = q^*$. If $q = q^*$, then (5.4) reduces to $\dot{m} = rm - \beta(q^* - q^+)$, and hence

$$\dot{m} = 0 \Longleftrightarrow m = \frac{\beta(q^* - q^+)}{r} = \bar{m}. \tag{5.8}$$

Since $\bar{q} = q^*$ is an interior solution, we must have

$$\frac{\partial H(p, q, m)}{\partial q} = \beta(p - p^*) - \bar{m}\alpha p(1-p) = 0, \tag{5.9}$$

[14] Consider that it is easy (if tedious) to show that $\frac{\partial \pi}{\partial p} = \beta(q - q^+)$.

https://doi.org/10.1017/S0956792510000136 Published online by Cambridge University Press

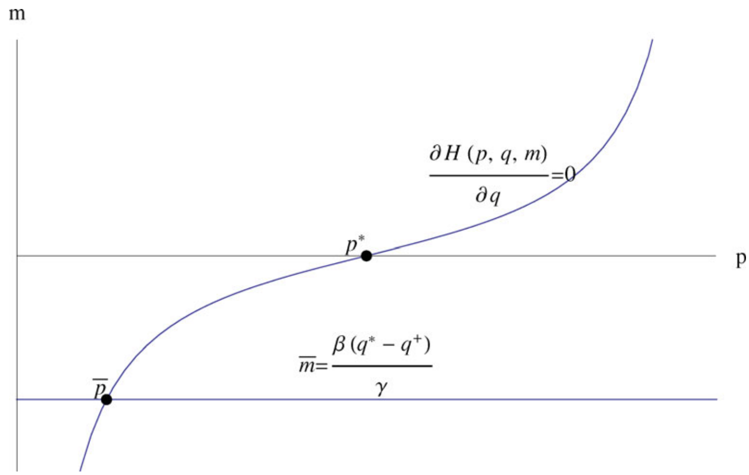

FIGURE 3. (Colour online) Computation of the stationary optimal path.

that is, $\bar{m}\alpha p^2 + p(\beta - \bar{m}\alpha) - \beta p^* = 0$. The roots of this equation are

$$p_{12} = \frac{-(\bar{m}\alpha - \beta) \pm \sqrt{(\bar{m}\alpha - \beta)^2 + 4\bar{m}\alpha\beta p^*}}{2\bar{m}\alpha}. \tag{5.10}$$

It can be easily shown that one of these two roots lies in the interval $[0, 1]$. This is the only value of p that satisfies the optimality condition (5.9) and the restriction on the state variable $p \in (0, 1)$. This completes the first part of the proof.

To see that $(\bar{p}, \bar{q}, \bar{m})$ is the only stationary optimal path, consider that for $q \neq q^*$ $\dot{p} = 0$ iff $p = 0$ or $p = 1$. However, if $p_0 \in (0, 1)$, then p cannot reach either 0 or 1 in a finite time (because $\dot{p} \to 0$ as p approaches the borders of the interval $[0, 1]$), so that if $q \neq q^*$ then $\dot{p} \neq 0$ for all t. This completes the proof. □

Figure 3 shows a graphic representation of the optimal stationary path $\bar{p}$. The horizontal line is the locus of points in the $p - m$ space such that $\dot{m} = 0$. The curved line is the locus where $\frac{\partial H}{\partial q} = 0$, which requires (see (5.9)) that $m = \frac{\beta}{\alpha}\frac{p-p^*}{p(1-p)}$. The optimal stationary value of p and m is where these two lines cross. Notice that $\bar{m} > 0$ ($\bar{m} < 0$) if $q^* > q^+$ ($q^* < q^+$). Also, the distance between the horizontal line $\dot{m} = 0$ and the x-axis increases as r approaches zero, and approaches zero as r goes to infinity.

The Hamiltonian (5.3) is linear in the control variable q and has a single interior stationary optimal path.[15] This implies that the optimal control problem 5.1 in the general case in which $p_0 \neq \bar{p}$, has a particularly simple solution: B must choose q in such a way that the state variable p approaches its optimal stationary value $\bar{p}$ in the shortest time. This is the content of the following:

[15] To see that the Hamiltonian in linear in q consider that the inspector's payoff function $\pi^B(.,.)$ is linear both in p and in q, and that RD is linear in q, although not in p.

Proposition 5

(1) *The optimal strategy S for the long run inspector B is:*

(a) *if* $p < \bar{p}$, *then* $q = 0$,

(b) *if* $p > \bar{p}$, *then* $q = 1$,

(c) *if* $p = \bar{p}$, *then* $q = q^*$.

(2) *Under the optimal strategy S, from any initial condition* $p_0 \in (0,1)$ *population A reaches its stationary optimal path level* $\bar{p}$ *in a finite time* $\bar{t}$. *After that time, B sets* $q = q^*$ *so that p remains fixed at* $\bar{p}$.

Proof (1) Because of the linearity of the Hamiltonian (5.3), the optimal path is the one that minimizes the time spent out of the optimal stationary path $(\bar{p}, \bar{q}, \bar{m})$. (Proofs of this fact can be found in Kamien and Schwartz [19], Section 16, Takayama [23] and Clark [6].) Since $\arg\max_q(\dot{p}) = 0$ and $\arg\min_q(\dot{p}) = 1$, if $p < \bar{p}$ $(p > \bar{p})$, the optimal policy for B is to set $q = 0$ $(q = 1)$. On the other hand, if $p = \bar{p}$, then $q = q^*$ so that $\dot{p} = 0$.

(2) Suppose that $p_0 < \bar{p}$ (the case $p_0 > \bar{p}$ can be treated similarly). In this case B will set $q = 0$, so that (5.2) reduces to a simple logistic equation

$$\dot{p} = p(1 - p)q^* \tag{5.11}$$

that can be integrated

$$p(t) = \frac{1}{1 - \exp(k - q^* t)}, \tag{5.12}$$

where $k = \log \frac{(p_0 - 1)}{p_0}$. It follows that the time it takes to bring p to its optimal level $\bar{p}$ is

$$\bar{t} = \frac{1}{q^*}[k - \log(\bar{p} - 1)]. \tag{5.13}$$

which is bounded away from ∞. After a time $\bar{t}$, B will set $q = q^*$ so that $\dot{p} = 0$. □

Proposition 5 shows that population A will spend most of the time at $\bar{p}$, because the Inspector will minimize the time it spends outside this state. Hence, it is interesting to see how changes in the police's time discount rate r affect $\bar{p}$.

Proposition 6 *(a) The optimal stationary value* $\bar{p}$ *converges to the stage game Nash equilibrium value* $\bar{p} \to p^*$ *as* $r \to \infty$; *(b)* $\bar{p} > p^*$ *and* $\bar{p} \to 1$ *as* $r \to 0$ *if* $q^* > q^+$; *(c)* $\bar{p} < p^*$ *and* $\bar{p} \to 0$ *as* $r \to 0$ *if* $q^* < q^+$.

Proof (a) If $r \to \infty$, $\bar{m} = \frac{\beta(q^* - q^+)}{r} \to 0$. Recall that $\bar{p}$ is a root of $\bar{m}\alpha p^2 + p(\beta - \bar{m}\alpha) - \beta p^* = 0$, which for $\bar{m} \to 0$ reduces to $p\beta - p^*\beta = 0$, whose only root is $p = p^*$.

(b) If $r \to \infty$, $\bar{m} = \frac{\beta(q^* - q^+)}{r} \to +\infty$ if $q^* > q^+$. Let us rewrite condition (5.9) as

$$\bar{m}\alpha = \frac{-\beta(p - p^*)}{p(1 - p)}. \tag{5.14}$$

https://doi.org/10.1017/S0956792510000136 Published online by Cambridge University Press

The left-hand side approaches ∞ as $\bar{m} \to +\infty$ and hence $p \to 1$ (If $p \to 0$, then the right hand side approaches $-\infty$). Similarly, the left hand side approaches minus infinite as $\bar{m} \to +\infty$, so that $p \to 0$. □

Proposition 6 is the core of this paper. It proves that if the Inspector is rational, but myopic ($r \to \infty$), he will keep the frequency of law violations around its Nash equilibrium value p^*. This is not surprising: if $p > p^*$, the Inspector gets a larger immediate payoff if he plays *Inspect*. In so doing, however, he reduces the frequency of law infractions p until it reaches the Nash equilibrium level p^*. At this point *Inspect* and *Not Inspect* yield the same payoff for him. Similarly, if initially $p < p^*$, a myopic Inspector will play *Not Inspect* because it yields a larger immediate payoff than *Inspect*. However, this will increase the frequency of violations p until it reaches p^*.

When the Inspector becomes more forward looking, ($r \to 0$) he will take into consideration the future effects of his current choices. Since he will not play the strategy that yields the largest immediate payoff, the optimal level of crime violations $\bar{p}$ will be different from p^*. The interesting result of this analysis is that increasing r will not necessarily reduce $\bar{p}$ below p^*. Proposition 6 shows that this will happen only provided that $q^* < q^+$. When the opposite inequality holds, a more forward looking Inspector (larger r) will tolerate a larger frequency of law infractions than a short-sighted one. Notice that the condition $q^* < q^+$ is the same condition that appears in Proposition 3. This means that a rational and forward-looking inspector will behave as if he could commit to a probability of inspection acting as a Stackelberg leader. The same conclusions we obtained in Section 3 hold also for the model presented in this section and will not be repeated here.

6 Conclusions

Many factors influence the frequency of law infractions. They range from the severity of punishment to the incentives inspectors have to discover crimes and the rationality of the parties involved. The inspection game proved to be a particularly sharp tool to investigate the intricate ways in which these magnitudes interact with each other to determine the equilibrium level of crime.

This paper has explored two classes of models. In the first class, the Inspector and the Public play the simultaneous version of the inspection game. These models show that higher penalties not necessarily translate into lower crime rates. They might bring about a reduction in the frequency of inspections instead.

Models belonging to the second class introduce an asymmetry between the Inspector and the Public. The Inspector can either be able to commit to a probability of inspection (Section 3) or be a forward-looking, rational agent (Section 5) who faces a population of boundedly rational would-be criminals. These models produce remarkably similar results. In both cases, the incentives for the inspector to discover law infractions can have perverse effects. Rewarding inspectors for discovering crimes might bring about more crime rather than less.

There are of course many questions all the models discussed so far leave unanswered. For example, they all assume that the inspector is in charge of preventing a single type of crime. In many circumstances, however, a single inspector must devote his scarce resources to the prevention of several crimes. He will thus face a trade-off, because an increase in the probability of inspection for one type of crime will necessarily bring about a reduction in the probability of inspection for other kinds of crime. If more policemen control tourists at the airport, there will be less policemen controlling drug-dealers in the streets. Varying the parameters of the models (e.g. the severity of punishment for different types of crimes) is bound to have even less predictable (and more perverse) effects in this type of setting than in the standard inspection game. This is likely to be a fertile ground for further research on this fascinating topic.

Acknowledgement

I would like to thank M. Holler, G. Tsebelis, L. Franckx, D. Arce, G. von Wangenheim and R. Cressmann for useful discussions and suggestions. Usual caveats apply.

References

[1] Andreozzi, L. (2002) Oscillations in against-pollution policies. An evolutionary analysis. *Homo Oeconomicus* **18**(3/4), 1–23.
[2] Andreozzi, L. (2004) Rewarding policemen increases crime: Another surprising result from the inspection game. *Public Choice* **121**(1), 69–82.
[3] Arce, D. & Gun, L. B. (2005) Working well with others: The evolution of teamwork and ethics. *Public Choice* **123**(1–2), 115–131.
[4] Becker, G. S. (1968) Crime and punishment: An economic approach. *J. Pol. Econ.* **76**, 169–217.
[5] Bianco, W. T., Ordeshook, P. C. & Tsebelis, G. (1990) Crime and punishment: Are one-shot, two person games enough?. *Am. Pol. Sci. Rev.* **84**(2), 569–586.
[6] Clark, C. W. (1990) *Mathematical Bioeconomics: The Optimal Management of Renewable Resources*, John Wiley and Sons, New York.
[7] Cox, G. (1994) A note on crime and punishment. *Public Choice* **78**, 115–124.
[8] Cressman, R., Morrison, W. G. & Wen, J.-F. (1998) On the evolutionary dynamics of crime. *Can. J. Econ.* **31**, 1101–1117.
[9] Friehe, T. (2008) Correlated payoffs in the inspection game: Some theory and an application to corruption. *Public Choice* **37**, 127–143.
[10] Fudenberg, D. & Levine, D. K. (1989) Reputation and equilibrium selection in games with a patient player. *Econometrica* **57**(4), 759–778.
[11] Fudenberg, D. & Levine, D. K. (1992) Maintaining reputation when strategies are imperfectly observed. *Rev. Econ. Stud.* **59**, 561–579.
[12] Fudenberg, D. & Levine, D. K. (1998) *The Theory of Learning in Games*, MIT Press.
[13] Graetz, M., Reinganum, J. & Wilde, L. (1986) The tax compliance game: Toward and interactive theory of law enforcement. *J. Law Econ. Organ.* **2**(1), 1–32.
[14] Harsanyi, J. C. (1966) A general theory of rational behavior in game situations. *Econometrica* **34**, 613–634.
[15] Harsanyi, J. C. (1977) *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*, Cambridge University Press, Cambridge, UK.
[16] Hirshleifer, J. & Rasmussen, E. (1992) Are equilibrium strategies unaffected by incentives? *J. Theor. Pol.* **4**(3), 353–367.
[17] Holler, M. J. (1990) The unprofitability of mixed-strategy equilibria: A second folk-theorem. *Econ. Lett.* **32**(4), 319–323.
[18] Holler, M. J. (1992) Fighting pollution when decisions are strategic. *Public Choice* **76**, 347–356.

[19] Kamien, M. & Schwartz, N. L. (1981) Dynamic optimization. *The Calculus of Variations and Optimal Control in Economics and Management*, North-Holland, New York.

[20] Mailath, G. J. & Samuelson, L. (2006) *Repeated Games and Reputations: Long-Run Relationships*, Oxford University Press, New York.

[21] Pradiptyo, R. (2007) Does punishment matter? A refinement of the inspection game. In: *Review of Law and Economics*, Vol. 3(2).

[22] Saha, A. & Poole, G. (2000) The economics of crime and punishment: An analysis of optimal penalty. *Econ. Lett.* **68**, 191–196.

[23] Takayama, A. (1985) *Mathematical Economics*, Cambridge University Press, Cambridge, UK.

[24] Tsebelis, G. (1990) Penalty has no impact on crime: A game theoretic analysis. *Ration. Soc.* **2**(3), 255–286.

[25] Tsebelis, G. (1990) Are sanctions effective? A game theoretic analysis. *J. Conflict Resol.* **34**(1), 3–28.

[26] Weibull, J. W. (1995) *Evolutionary Game Theory*, MIT Press, Cambridge, MA.

[27] Wittman, D. (1985) Counter-intuitive results in game theory. *Eur. J. Pol. Econ.* **1**, 77–89.

[28] Wittman, D. (1993) Nash equilibrium vs. maximin: A comparative static analysis. *Eur. J. Pol. Econ.* **9**, 559–565.