

HIGHER PUNISHMENT, LESS CONTROL?

EXPERIMENTAL EVIDENCE ON THE INSPECTION GAME

Heiko Rauhut

ABSTRACT

Rational choice theory predicts for higher punishment less crime. However, many field studies could not support this conclusion. A game theoretic approach can explain these puzzling findings because it takes not only criminals' but also control agents' rationality into account. Mixed Nash equilibria predict for higher punishment less control and no effect on crime rates. A new experimental design is introduced to test game theoretic hypotheses. 196 subjects have been partitioned into 'inspectees' who can steal money from each other and 'inspectors' who can invest in control activities to catch inspectees. Static and dynamic analyses show that strategic interaction plays an important role for crime and punishment. However, effects are not as strong as predicted. Higher punishment indeed causes less control, but crime is deterred as well. Furthermore, dynamical analyses with the learning model fictitious play reveal that humans learn only slowly in inspection situations.

KEY WORDS • control • crime • game theory • learning • punishment

1. Introduction

Do severe penalties help to reduce crime? This question on the deterrent effect of punishment divides experts and the public. It seems straightforward that higher punishment deters crime. However, the empirical literature is far from finding a consistent deterrent effect of punishment. A recent overview of reviews concludes that the null

hypothesis should be accepted for the relationship between sentence severity and crime. It summarizes: 'A reasonable assessment of the research to date – with a particular focus on studies conducted in the past decade – is that sentence severity has no effect on the level of crime in society' (Doob and Webster 2003: 143). Von Hirsch et al. (1999: 47) note in an extensive review that 'present association research, mirroring earlier studies, fails . . . to disclose significant and consistent negative associations between severity levels (. . . such as duration of imprisonment) and crime rates.' More specific studies often reach similar conclusions. MacCoun and Reuter (1998: 213) state about drug control that 'severity of sanctioning has little or no influence on offending.' Zimring et al. (2001: 105) conclude on the effects of the Three-Strike laws in the US that 'the decline in crime observed after the effective date of the Three-Strikes law was not the result of the statute.'

So why is it that sentence severity has often relatively little impact on crime rates?¹ In this article, two propositions are outlined to pave the way for game theoretic laboratory experiments in the field of crime and punishment.

Proposition 1. A game theoretic perspective can explain why punishment has little effect on crime.

Most analyses of deterrence assume a decision theoretic context: criminals weigh costs and benefits for a certain criminal behavior with a given punishment level and detection probability. However, the rationality of control agents is neglected in this type of analysis. Punishment is only threatening if criminals are actually detected; however, detection takes effort. Rational control agents will try to optimize their level of control according to a specific crime level and vice versa. Game theory can provide predictions on the interaction between crime and control rates. Tsebelis (1990) pioneered the game theoretic approach of crime and punishment. Its premise rests on the assumption that criminals and control agents have opposite incentive structures. While criminals like to commit crimes, control agents like to catch them. If criminals knew that they are not watched, they would prefer to commit a crime. However, if they committed a crime and control agents knew it, control agents would prefer to take some effort to detect them (i.e. to improve their careers). As a consequence, criminals would resign from crime, which makes control again unattractive. Players want to avoid predictability in such a situation. In game theoretic terms, we say there is no solution in pure strategies. A preliminary prediction might be that criminals and control agents will

change their behavior all the time to be unpredictable. Fortunately, game theory provides a more precise answer to the problem. It is predicted that criminals and control agents will choose a certain probability for their action. This probability choice is called a mixed strategy. The chosen probability shall leave the opponent indifferent.² Thus, criminals calibrate their probability to commit a crime so that control agents are indifferent between control and no control. The other way round, control agents choose their probability of control so that criminals are indifferent between committing a crime or not. Hence, criminals are sensitive for payoffs of control agents and control agents are sensitive for payoffs of criminals. Consequently, higher punishment does not deter crime but control.

Proposition 2. Laboratory experiments provide the best validation regime to test game theoretic hypotheses for crime and punishment.

There have been field studies on punishment effects on crime, however, there are many potential confounding variables. Furthermore, previous field studies of the effects of punishment on crime could not specifically test Tsebelis' explanation, because these studies took severity of punishment as an exogenous condition and did not use information about how this affected behavior of law enforcement agencies in the field settings under study. Therefore, it is suggested to develop a new design for a laboratory experiment. In the laboratory, punishment levels can be manipulated while everything else can be held constant, so that causal effects of punishment on crime *and* control can be measured with high internal validity.³ Furthermore, utilities for crime and control can be measured with monetary incentives, so that precise point predictions can be tested.

There is a growing body of literature that uses laboratory experiments to estimate effects of informal punishment (Fehr and Gächter 2000, 2002; Falk and Fischbacher 2002; Diekmann 2003; Fehr and Fischbacher 2004; Voss and Vieth 2006; Anderson and Putterman 2006). However, there is no laboratory experiment on the inspection game, which studies formal punishment regimes that pay control agents a reward if they successfully detect those who break the norms.

This article is structured as follows. First, the decision theoretic versus the game theoretic approach of crime and control are contrasted. Game theoretic hypotheses are derived from mixed Nash equilibria. An experimental design is introduced and empirical results from two experiments are presented. The article concludes with a discussion of further research and policy implications.

2. Theoretical analysis of crime and punishment

2.1 *Crime from a decision theoretic perspective*

In his seminal article, Gary Becker (1968) was the first to develop a sound formal theory of crime and punishment. Becker (1968) regards crime as rational behavior, accessible to standard market equilibrium analysis. Criminals are regarded as utility maximizers who optimize their payoffs under restrictions and risk. Criminals have clear incentives for criminal conduct; they gain material utility from theft or burglaries, respectively immaterial gains for mayhem and the like. More specifically, criminal i receives the combined monetary and psychic payoff y from a certain crime. However, she faces the conviction probability of c to receive the punishment p . Formally, we can write the expected utility π from crime for criminal i as $\pi_i = c(y - p) + (1 - c)(y)$ (see Becker 1968: 177). Let s_i denote the likelihood that criminal i commits the crime. We can rearrange terms and obtain the payoff of criminal i as

$$\pi_i(s_i) = s_i(y - cp). \quad (1)$$

In this perspective higher punishment reduces criminal activities. However, Becker's analysis neglects strategic interaction: Outcomes of criminals, victims and police depend on the action of ego and on the actions of all other involved actors. Criminals want to commit crimes on their victims, victims do not want to be victimized and control agents want to catch criminals, which criminals want to avoid. Confounding decision theoretic problems with problems of strategic interaction can be coined as the Robinson Crusoe fallacy (Tsebelis 1989). While sociologists, especially from the labeling paradigm, stressed the dynamics of social interaction early on, game theory and labeling theory did not synthesize so far.⁴

As a consequence, I will elaborate a game theoretic perspective on crime and punishment in the next sections. For that, I will first present the original contribution of Tsebelis (1990) of the first game theoretical model of crime and punishment. Subsequently, I will demonstrate with examples from other areas of the social sciences the structural differences between the model of Tsebelis (1990) and other, so-called discoordination games. The demonstration of such differences serves as an argument why the original model is modified with ideas from social contract theory (Taylor 1976, 1987; Kavka 1983). As a result, a model is presented that conserves the structural features of the original contribution of Tsebelis

(1990) but enables more realistic tests that allow us to draw more reasonable inferences to criminal behavior outside the lab. In particular, the modified model is motivated by the idea that law enforcers in the field might also be motivated by the knowledge that criminals inflict harm on their victims, and criminals, furthermore, might constrain themselves to some extent due to this knowledge. I conclude this paper with a brief comment why I nevertheless do not incorporate aspects from behavioral game theory into this model.

2.2 Crime from a game theoretic perspective

Tsebelis (1990) was the pioneer who introduced a game theoretic model of crime and punishment. There are two different groups of actors: members of one group can decide to commit a crime and members of the other group can decide to inspect. Members of the first group are called *inspectees* and members of the second group *inspectors*. Inspectees can decide to commit a crime with payoff y and punishment costs p if caught. If they do not commit a crime, their payoffs remain unchanged. Inspectors can decide to inspect inspectees. They have to invest inspection costs k to detect the action of the inspectee. If an inspector catches an inspectee having committed a crime, the inspector receives the reward r . If not successful, inspection costs are lost. If they do not inspect, they remain at their income level. Undetected crime is assumed to be attractive and punishment to be a threat, so that $p > y > 0$. Likewise, it is assumed that inspectors gain from successful inspection so that $r > k > 0$. The situation is illustrated with the 2×2 matrix in Table 1.

It is best to commit a crime if not inspected and not commit a crime if inspected. Inspectors have reverse payoffs: it is best to inspect a criminal and desist from inspection if there is no crime. Thus, there is no solution in pure strategies; so actors have to choose a certain probability for their actions. Let s_i be the probability that inspectee i chooses to commit the crime and c_j the probability that inspector j will inspect inspectee i . We can write the payoff function π for inspectee i who plays against inspector j as

$$\pi_i(s_i, c_j) = s_i(y - c_j p).$$

The payoff function ϕ for inspector j who plays against inspectee i is

$$\phi_j(s_i, c_j) = c_j(s_i r - k).$$

Table 1. The inspection game

		Inspector j	
		inspect	not inspect
Inspectee i	crime	$y - p, r - k \Leftarrow$	$y, 0$
	no crime	$0, -k \Rightarrow$	$0, 0$

with $p > y > 0, r > k > 0$

The difference between the game theoretic and the decision theoretic approach sharpens: in the decision theoretic context, criminals have to adjust their criminal activity to a given conviction probability. In the game theoretic approach, the conviction probability is a product of rational choices of inspectees who want to maximize their payoffs. The best response for inspectee i can be calculated by the first partial derivative $\frac{\partial \pi_i}{\partial s_i}$

$$s_i^* (c_j) = \begin{cases} 1 & \text{if } y - c_j p > 0 \\ [0, 1] & \text{if } y - c_j p = 0 \\ 0 & \text{if } y - c_j p < 0. \end{cases}$$

The first partial derivative $\frac{\partial \phi_j}{\partial c_j}$ shows the best response for inspector j as

$$c_j^* (s_i) = \begin{cases} 1 & \text{if } s_i r - k > 0 \\ [0, 1] & \text{if } s_i r - k = 0 \\ 0 & \text{if } s_i r - k < 0. \end{cases}$$

Due to $p > y > 0$ and $r > k > 0$, there are no Nash equilibria in pure strategies. The optimal choice is to choose a mixture of strategies that leaves the opponent indifferent. If alter is not indifferent, she will take advantage and exploit ego on ego’s costs and vice versa. Hence, both actors should calibrate their chosen probability with the utility function of their opponent in such a way that she is indifferent. Given the indifference conditions above, we calculate the mixed Nash equilibrium for the decision to commit a crime for inspectee i as $s_i^* = \frac{k}{r}$ and likewise the mixed Nash equilibrium c_j^* to invest in control for inspectors as $c_j^* = \frac{y}{p}$. Results are counterintuitive: less attractive crimes would cause less inspection, not less crimes. Conclusively, higher punishment would have no impact on crime – higher punishment would cause less inspection. Likewise, more attractive inspection would cause less crime and did not affect inspection behavior. While the derived hypotheses

from the game theoretic model of crime might astonish, Tsebelis (1990) has shown that the equilibria are stable for various extensions of the model: infinite strategy spaces, unilateral and bilateral myopic behavior and imperfect information.

2.3 The meaning of mixed strategies

The model of Tsebelis (1990) is supposed to describe the strategic interaction between criminals and control agents such as policemen, public prosecutors, or judges. While the game captures the reversed interests between criminals and inspectors, it lacks a 'criminal element,' i.e. the harm implied by criminal activities. The model rather describes a 'matching-pennies' situation. One party is interested in a mismatch (i.e. commit a crime if there is no control) while the other party is interested in a match (i.e. perform control if there is crime).

Such 'matching-pennies' situations are well known in many areas of social life. In the following, we will show with examples from sport, literature, war, and education the advantage of conceptualizing conflict between two parties as so-called 'discoordination problems' (Rasmusen 2004: 77). These examples illustrate, on the one hand, the similarity between crime and other bilateral conflicts in social life. On the other hand, however, the examples illustrate that criminal behavior implies additionally negative externalities for third parties, which can be understood as the missing 'criminal element' in the inspection game. While this aspect is not so important for understanding the structural features of the inspection game, it is crucial for the construction of a design which enables a valid test of these features in the laboratory. The missing 'criminal element' in the inspection game and its importance for respective experimental designs will be discussed in sections 2.4 and 2.5.

The most intuitively understandable discoordination situations can be found in professional sports. Consider the example of penalty kicks in soccer (Chiappori et al. 2002; Palacios-Huerta 2003; Moschini 2004; Berger and Hammer 2007). The goal-keeper prefers to dive left when the penalty kick goes left and prefers the right side if the kick goes right. In contrast, the kicker has the opposite preferences. There is clearly no Nash equilibrium in pure strategies, as the goal-keeper would always want to switch away from a mismatch, and the kicker would always want to switch from a match. That is, the penalty kick is a paradigmatic example of a discoordination game with the zero-sum (or constant-sum) properties: what one player wins, the other player loses. Another example is the serve and return play of tennis players (Walker and Wooders 2001).

The serving player aims at outplaying the returning player. The best strategy for the server is to serve long-line when the return player expects a cross and vice versa. The returning player, on the other hand, wants to choose the same corner as the server. Again, one player is interested in a match, whereas the other player is interested in a mismatch.

Examples of discoordination games are not restricted to sports. Even literature is full of discoordination problems. Brams (1994) illustrates a scene in *Sherlock Holmes*, where Holmes is pursued by his opponent Moriarty.⁵ Holmes has to decide whether to get off his train at Dover or at Canterbury, an intermediate stop. He chooses Canterbury, anticipating that Moriarty will take a special faster train to Dover to try to catch him when he gets off there. With this move, Holmes has outplayed his opponent.⁶

Such hide-and-seek situations are quite frequent in wars. One classic example is the battle of the Bismarck Sea in the South Pacific in 1943 (Rasmusen 2004: 22). General Imamura has been ordered to transport Japanese troops across the Bismarck Sea to New Guinea, whereas General Kenney aims to bomb the transport. Imamura can decide between a shorter northern route and a longer southern route to New Guinea while Kenney faces the decision of where to send his planes. For Kenney, a mismatch means losing valuable days of bombings, whereas for Imamura, a match means losing his troops.

Our last example deals with education (Gintis 2000: 71–72). Suppose a mother wants to help her unemployed son financially, but she does not want to contribute to his distress by allowing him to loaf around. The problem is that the son might even have difficulties in finding a job when his mother helps but he prefers to loaf around if he can enjoy the financial help of his mother. It turns out that this example bears the typical cyclical structure of discoordination problems as well. If the mother helps the son, the son loafs around. Knowing this, the mother will cut off financial aid, the son seeks work, but the mother feels bad. Thus, she changes her strategy, supports her son, drives him again to laziness, and, finally, cuts off her financial aid again.

2.4 Game theory and the social contract

The several examples in the last section showed that discoordination problems are not only relevant in the field of crime and punishment, but in many other fields of the social sciences, too. In addition, however, I illustrated that pure discoordination problems capture the mere conflict between two parties. Consequently, we can pose the relatively abstract

research question whether humans are in general capable of finding mixed-strategy equilibria in abstract discoordination problems. When surveying the literature, we find several elementary and advanced laboratory experiments, which focus on this rather abstract research question (O'Neill 1987; Rapoport and Boebel 1992; Bloomfield 1994; Shachat 2002; Goeree et al. 2003; Palacios-Huerta and Volij 2008).

However, the scope of this article goes beyond these approaches. Our aim is to study discoordination behavior for the particular case of crime and control. Whereas games between tennis players, between soccer players, or between mothers and sons only affect the particular players in the game, the story for crime is different. Criminal behavior bears a 'criminal element,' which should be incorporated for a valid empirical test of discoordination behavior between criminals and inspectors. In particular, gains for criminals have further impact as they reduce the welfare of the society at large. Put the other way round – only because of the welfare loss from crime, inspectors are employed by society to control, catch, and punish criminals. Expressed in methodological terms, the model suggested by Tsebelis (1990) lacks *construct validity*. While payoffs y from crime capture money, social status, or psychic relief, negative externalities in the form of losses of victims are neglected.

This missing 'criminal element' is best described in social contract theory. Thomas Hobbes argued in his *Leviathan* that it follows from the nasty and brutish nature of men to establish a strong state, so that men are controlled and, in case of non-cooperative behavior, punished. In his view, only a state-driven punishment regime can promote cooperation. Later, political philosophers and social scientists reformulated and modeled Hobbes' argument with the prisoner's dilemma (Taylor 1976, 1987; Kavka 1983).

Hobbes's analysis of the state of nature as a state of war is the introduction of [the] Prisoner's Dilemma. Namely, that in certain important situations, there is a divergence between individual and collective rationality. That is, if each individual performs the act that is, in fact, in his own individual best interest, all – ironically – end up worse off than if they had all acted otherwise. Hobbes, in effect, though not in so many words, points out this problem with respect to attack behavior and promise keeping in the state of nature. Each would be better-off if they all kept their agreements and refrained from attacking one another. But there are apparent unilateral advantages to be gained by violating agreements and by conquest, and one will suffer substantial disadvantages if others do these things and one does

not. As a result, agreements are in vain and anticipatory attack is the most reasonable individual strategy. (Kavka 1983: 309–310)

We can borrow this idea from classic contract theory to model criminal behavior with a prisoner's dilemma. In particular, we can think of crime as a welfare transfer between perpetrator and victim: While criminals enjoy positive effects, victims suffer negative effects from crime. Whereas we denoted previously criminal gains with y , we can now specify y as the welfare transfer between perpetrator and victim. More specifically, we assume that the commodity x is transferred from victim to perpetrator with an efficiency factor γ , and replace y with γx . Furthermore, crime is mostly inefficient: a stolen watch will pay less to a thief on the black market than the victim would have paid for it. Similarly, the psychic and physical absolute value of mayhem or rape will be less for the perpetrator than for the victim. As a consequence, we assume that $0 < \gamma < 1$ and model criminal behavior as a prisoner's dilemma between two criminals h and i in the 2×2 matrix (Table 2).

Obviously, it is the dominant choice to defect, i.e. to commit a crime, although both parties benefited from mutual resistance to commit the crime. Herewith, we can specify the above-discussed 'criminal element' as the welfare loss associated with mutual defection compared to mutual cooperation. Both lost nothing $[0,0]$ if both desisted from crime. However, the Nash equilibrium predicts that both commit a crime and face losses of $[x(\gamma - 1), x(\gamma - 1)]$.⁷

2.5 Inspection of criminals in a prisoner's dilemma

The social trap of criminals and victims creates a demand for punishment, as both actually preferred to live in a situation of mutual cooperation but end up in a situation of mutual defection. Punishment can provide a solution to this demand in adding an incentive to coordinate on a more beneficial equilibrium, as argued by earlier scholars:

Hobbes proposes a plausible solution to the problem of diverging individual and collective rationality: the creation of a power to impose sanctions that would alter the parties' payoffs so as to synchronize individual and collective rationality. (Kavka 1983: 310)

The innovative element, now, is to combine the prisoner's dilemma with the inspection game. The main argument for this innovation is to provide a model, which reflects the same structural features as the

Table 2. The crime game

		Criminal <i>i</i>	
		crime	no crime
Criminal <i>h</i>	crime	$\gamma x_i - x_h, \gamma x_h - x_i$	$\gamma x_i, -x_i$
	no crime	$-x_h, \gamma x_h$	$0, 0$

with $0 < \gamma < 1, x_h = x_i > 0$.

inspection game (Tsebelis 1990), but incorporates the reasoning of social contract theory that crime shall be captured by Pareto-inefficient strategies and inspection by a strategy to deter such Pareto-inefficiencies of the inspectees. In particular, it is suggested to model criminal behavior as a prisoner’s dilemma between two inspectees who are in a separate strategic interaction situation with the inspector. Such a concept allows to test the implications of the discoordination game for the field of crime and punishment.

More specifically, we suggest a simple four-player system. Inspectee *i* can offend against inspectee *h* and inspectee *h* can offend against inspectee *i* with criminal payoffs of γx and victimization losses of x . In addition, there are two inspectors; inspector *j* can inspect inspectee *i* and inspector *l* can inspect inspectee *h*, as shown in Figure 1.

We can combine our reasoning on the inspection game and on the prisoner’s dilemma and write our model as follows. The payoff function π for inspectee *i*, who plays against inspectee *h* and inspector *j* and the payoff function for inspectee *h* who plays against inspectee *i* and inspector *l* are

$$\pi_i(s_h, s_i, c_j, c_l) = s_i(\gamma x_h - c_j p) - s_h x_i \tag{2}$$

$$\pi_h(s_h, s_i, c_j, c_l) = s_h(\gamma x_i - c_l p) - s_i x_h. \tag{3}$$

If no inspectors are present ($c = 0$), it is a dominant choice for both inspectees to commit a crime due to $\frac{\partial \pi}{\partial s} = \gamma x$. However, both inspectees face losses of $(\gamma - 1)(x_h + x_i)$ in total. Due to $0 < \gamma < 1$, player *i* and player *h* are in a prisoner’s dilemma when no inspectors are present. Hence, both will commit a crime against each other and both will face losses due to inefficient welfare transfers. Because payoffs from victimization cancel out in the first partial derivative, the best response for inspectee *h* is independent of the choice of inspectee *i* and vice versa.

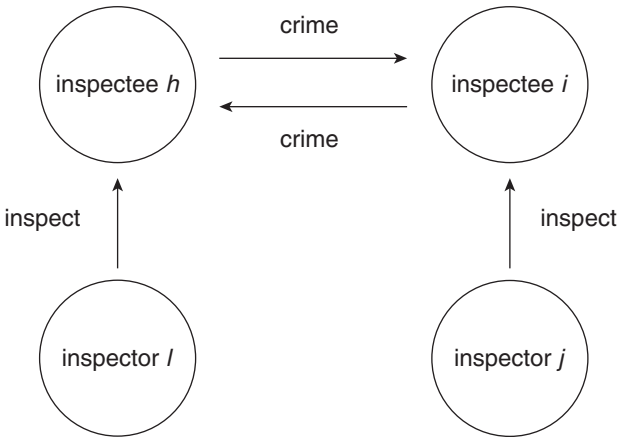


Figure 1. The combination of the inspection and the crime game

However, the third-party intervention of inspectors changes the prisoner’s dilemma to an inspection game: If caught, inspectees have to face the punishment cost p . Inspectors have to pay inspection costs k and gain rewards r if successful, as discussed in section 2.2. The payoff function ϕ for inspector j who inspects inspectee i and the payoff function for inspector l who inspects inspectee h are

$$\phi_j(s_h, s_i, c_j, c_l) = c_j(s_i r - k) \tag{4}$$

$$\phi_l(s_h, s_i, c_j, c_l) = c_l(s_h r - k). \tag{5}$$

Inspectees optimize their choices against inspectors and inspectors optimize their choices against inspectees, irrespective of the behavior of opponent inspectees. Hence, we derive the best response as in section 2.2, except that criminal payoffs are now denoted by γx instead of y . Therefore, we calculate the mixed Nash equilibrium strategy for inspectees from the indifference condition of the best response for inspectors as

$$s^* = \frac{k}{r}. \tag{6}$$

The mixed Nash equilibrium strategy for inspectors to invest in inspection is calculated from the indifference condition of the best response for inspectees as

$$c^* = \frac{\gamma x}{p}. \tag{7}$$

In conclusion, we derive *ceteris paribus* two main hypotheses about the effect of strength of punishment on crime and control:

Hypothesis 1. Strength of punishment p has no impact on crime.

Hypothesis 2. The higher the punishment p the less inspection.

We see that the derived Nash predictions are the same as in the two player system of Tsebelis (1990); thus we have succeeded in proposing a model that implements on the one hand the discoordination mechanism and on the other hand negative externalities from crime.

2.6 Some notes on behavioral game theory

It is possible to raise another argument, why the incorporation of the prisoner's dilemma into the inspection game is useful for a laboratory test of the inspection game. The inequity between a criminal and a non-criminal inspectee can invoke feelings of injustice. While the criminal might feel guilt for outsmarting her non-criminal opponent, her opponent might feel envy for being outsmarted.

We know from laboratory experiments that such emotions matter and change behavioral patterns (for an overview see Fehr and Gintis 2007). Furthermore, inspectors, as third parties, might feel with the inspectees, as documented by third-party punishment experiments of dictators games (Fehr and Fischbacher 2004). In the meanwhile, there has been progress on how these emotions can be incorporated into game theory (Rabin 1993; Fehr and Schmidt 1999; Bolton and Ockenfels 2000). It is possible to derive different equilibria for the inspection game, in line with the research direction of behavioral game theory (Camerer 2003, 2004). More specifically, the introduction of inequity aversion for inspectees would decrease the payoffs for criminal conduct. Furthermore, the affection of inspectors is expected to provide an additional inspection reward in terms of positive feelings for doing justice to the inspectees. As a consequence, both parties would change their probability mixture accordingly. It is possible to test the implications of behavioral game theory with a parallel measurement of behavior and attitudes, for example with reciprocity scales (Perugini et al. 2003).

Moreover, behavioral game theory offers even more arguments to change strategy mixtures. It is possible to consider aspects of motivation crowding theory (Frey 1994; Frey and Jegen 2001), i.e. to assume that the introduction of punishment might crowd out inequity aversion.

Alternatively, there might be other behavioral side-effects of punishment: for example, inspectees could commit more crimes in a mild punishment regime compared to regimes without punishment. The reason is that mild punishment could be regarded as a 'legitimate prize for crime' (Gneezi and Rustichini 2004), so that criminals buy their legitimation for criminal conduct.

As we see, introducing notions from behavioral game theory can drive equilibria in various directions. In contrast to such relatively arbitrary extensions, the analysis of the inspection game is founded on the solid axioms of standard game theory and can, nevertheless, explain why higher punishment does not decrease crime rates. Therefore, we regard the analysis with standard game theoretic tools as more parsimonious and better substantiated so that we leave the various extensions from behavioral game theory aside in this paper. Nevertheless, the proposed model and associated experimental design could be useful in deriving implications from behavioral game theory and respective tests in future studies.

3. Experimental design

3.1 Basic design and procedures

In the first part of the experiment, subjects earn own property by providing correct answers in a knowledge quiz. We use experimental tokens, which are transferred to euros with the exchange rates given below. The quiz includes thirty multiple choice questions, with two choices each. It covers politics, art, geography, science and mathematics. Subjects have 90 seconds to answer 5 questions within each of these knowledge fields. A quiz lets subjects attribute own property to obtained money. Previous experiments report that effort for entitlements creates stronger and more reliable incentives (Falk and Fischbacher 2002; Gächter and Riedl 2005).

In the second part of the experiment, subjects receive instructions about the rules of the game. Subjects are told that they will be randomly divided into 'players' and 'inspectors.' In the experimental instructions, we avoided value-laden terms for inspectees like criminal or technical terms like inspectee. Yet, they do not know which role they will play; inspectee or inspector. Accordingly, subjects have to comprehend rules for both parties. Instructions are structured in three sections. The first section covers the general structure of the inspection game.

The remaining two sections cover the specific monetary payoffs for inspectees and inspectors respectively. Next, subjects complete eleven multiple choice questions about the rules of the game for both parties, inspectees and inspectors. Only now, subjects are randomly divided into inspectees and inspectors and learn their roles. Thereafter, inspectees and inspectors play two practice periods in their respective roles before the experiment is started with monetary stakes.

In the third part, subjects play 30 periods of the inspection game. We use a stranger-matching: each inspectee i is randomly matched with one different inspectee h each period. Each of these two inspectees can take simultaneously a fixed amount of money x from the personal account of the other inspectee. However, thieves only earn γx , while victims lose x with $0 < \gamma < 1$. For the matching of inspectors with inspectees, we use a stranger-matching as well. Each inspector is randomly matched with one different inspectee each period. Inspector j can choose to pay inspection costs k to reveal the decision of her matched inspectee i . If inspector j inspects inspectee i and inspectee i actually stole money from her matched inspectee h , inspector j receives a monetary reward r , while inspectee i receives punishment $p > \gamma x$. However, inspectee i keeps her stolen money γx .

There are four experimental treatments. Strength of punishment p and their order is varied. $N = 20$ subjects take part in each experimental session.⁸ In sessions 1–5, subjects were allocated to low punishment for periods 1–15. Thereafter, strength of punishment changed. The same subjects were allocated to high punishment for periods 16–30. Subjects kept their role as inspectee or inspector respectively. We denote *experiment 1* as treatment with low punishment first. In *experiment 2*, treatments were reversed: subjects started with high punishment for 15 periods and continued thereafter with low punishment. The treatment conditions are summarized in table 3. The proposed design provides high statistical power due to the ability of within and between subject comparisons for both inspectees and inspectors.

3.2 Payoffs and information conditions

The experiment is conducted in a computer laboratory and programmed with the software z-Tree (Fischbacher, 2007). Subjects interact anonymously via a computer network with each other and do not learn the identity of their interaction partners at any time. The parameters are as follows. Inspectee h can steal from inspectee i $x = 10$ experimental tokens. The welfare inefficiency factor is $\gamma = 0.5$. Hence,

Table 3. Treatment conditions

	Experiment 1	Experiment 2
	Low punishment first (session 1–5)	High punishment first (session 6–10)
15 periods low punishment	48 inspectees 48 inspectors	50 inspectees 50 inspectors
15 periods high punishment	48 inspectees 48 inspectors	50 inspectees 50 inspectors

In total: 10 sessions, 196 subjects, 5880 decisions

inspectee h only earns $\gamma x = 5$ experimental tokens for theft. Inspector j has to pay $k = 5$ tokens inspection costs and gains $r = 10$ tokens for successful inspection. Punishment is in the low punishment condition $p_l = 6$ tokens and in the high punishment condition $p_h = 25$ tokens. Note that inspection costs do not change when strength of punishment is changed. Thus, increased punishment is interpretable as increased efficiency of punishment.

One additional problem about the payoffs has to be solved. On average, inspectees will lose money whereas inspectors will stay at their income level during the course of the game. However, we like inspectees and inspectors to receive roughly equivalent payments at the end of the experiment. There are two possible alternatives. Inspectors earn less money in the knowledge quiz or inspectees and inspectors can be treated with a different exchange rate. We take the second solution because it guarantees equivalent treatments for inspectees and inspectors in the knowledge quiz. Inspectees are paid 10 cents and inspectors are paid 2 cents for each experimental token. Hence, inspectees can steal 1 euro with a gain of 50 cents each period and inspectors can invest 10 cents with the prospect of earning 20 cents for successful inspection. Note that different exchange rates do not result in different Nash equilibria. Inspectees and inspectors learn their exchange rates when they learn which role they will play. Inspectees and inspectors only know their own exchange rate; however, they are told that there is a different exchange rate due to unequal expected earnings at the end of the experiment.

Subjects are given the following information after each period. Inspectee h learns the decision of her matched inspectee i and the decision of her matched inspector l and, if eligible, whether she has been punished. Inspector l learns the decision of her matched inspectee h .

Both players learn their current income level each period. Table 4 summarizes the payoff parameters for both roles.

3.3 Predictions

Given the payoff parameters for inspectees, we can predict with (6) theft probability for low punishment as $s_{(lp)}^* = \frac{1}{2}$ and equally for high punishment $s_{(hp)}^* = \frac{1}{2}$. Given the payoff parameters for inspectors, we can predict with (6) inspection probability for low punishment as $c_{(lp)}^* = \frac{5}{6}$ and for high punishment as $c_{(hp)}^* = \frac{5}{25}$. In conclusion, we infer the following predictions (note rounding):

- Prediction 1.** The theft rate averages 50% for low and for high punishment.
- Prediction 2.** The inspection rate averages 80% for low punishment and 20% for high punishment.

4. Experimental results

4.1 Some descriptive features of the experiment

In total, we have observations from 196 subjects. These subjects were randomly chosen from an address pool of 692 students and randomly allocated to one experimental session. Each subject participated only once in the experiment. The address pool consists of students from many different fields from the University of Leipzig. We offered each subject a show-up fee of 5 euros for participating in the experiment, which was paid at the end of the experiment. The experiment lasted for about one hour. Subjects understood the instructions well, as we learned by the comments of the subjects. 72% of subjects were female, due to an over-representation of students from humanities and social sciences in the address pool.⁹

The income distribution throughout the periods illustrates that inspectees lose money over time, while inspectors roughly stay at their income level. On average, inspectees start with 23 euros and lose 19 euros during the game. Inspectors start on average with 4.50 euros and finish with 4.70 euros in the final period. However, the range of final payments for inspectees is between 0.00 and 19.30 euros and for inspectors between 2.00 and 6.60 euros. For total earnings, the show-up fee of 5 euros has to be added.

Table 4. Payoff parameters (in experimental tokens)

Inspectees		Low punishment	High punishment
p	Strength of punishment	6	25
γ	Crime inefficiency factor	0.5	0.5
x	Loss victim	10	10
γx	Gain thief	5	5
E	Exchange rate (points-euro)	0.10	0.10
Inspectors		Low punishment	High punishment
c	Inspection cost	5	5
r	Reward successful inspection	10	10
E	Exchange rate (points-euro)	0.02	0.02

4.2 Construct validity of the experimental design

We conducted several qualitative interviews after the experiments to determine the construct validity of our experimental design. These interviews reveal that subjects understood the structure of the game quite well. More importantly, however, the interviews show that the chosen design of the combination of the structural features of the inspection game with those of the prisoner's dilemma was successful. The subjects felt like being in a discoordination situation of the particular type of crime and crime control. We asked the subjects, which real-life situations they associated with the setting in the laboratory experiment. Inspectees thought of fare-dodging, shoplifting, theft, speeding, smuggling, cheating with scholarships and inspectors thought of policemen, private control agents, governmental controls, or, more playfully, of playing cops and robbers.

4.3 Punishment effects on crime and inspection

First, we test predictions 1 and 2 in a static analysis. We analyze the mean effect of strength of punishment on theft and inspection, averaged over all periods. We treat repeated measurements for theft and inspection over 15 periods as a *reliability measure*. Moreover, we treat results from our 10 different sessions as a reliability measure, so that we interpret session data as *replications* of the same experiment.

Figure 2 shows mean theft and mean inspection rates for low vs. high punishment for experiment 1, separately for each session. Each bar represents 150 decisions (except session 3 with 135 decisions). For

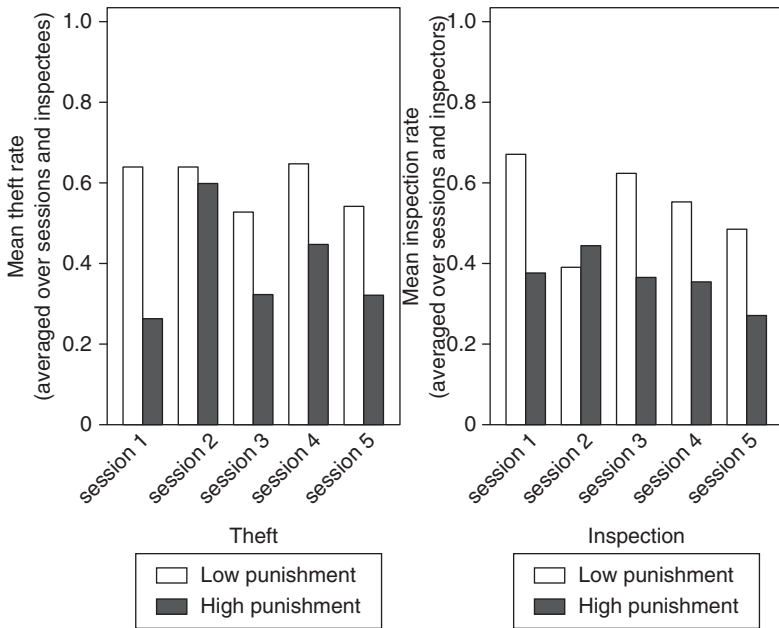


Figure 2. Mean punishment effect for each session (experiment 1)

all sessions, there is less theft for low punishment than for high punishment. The average theft rate over all sessions is for *low* punishment 60% and for *high* punishment 40%. We cannot hold hypothesis 2.5; stronger punishment does apparently deter crime.

For all sessions (except session 2) we observe for higher punishment less inspection. The average inspection rate over all sessions is for *low* punishment 54% and for *high* punishment 36%. Therefore, we can confirm hypothesis 2: Higher punishment does indeed deter inspection. With logistic regressions¹⁰ it can be shown that punishment effects on inspection are even in single sessions significant on the 5% level. This is a remarkably strong finding considering that there are only 10 inspectors per session. However, subjects do not reach the exact mixed Nash predictions; there is too much inspection for high punishment and too little inspection for low punishment.

Data from experiment 2 with high punishment first reveal a similar pattern. In all sessions, crime increases for lower punishment. Averaged over all sessions, we observe 47% theft for high punishment and 69% theft for low punishment. In four out of five sessions, inspection

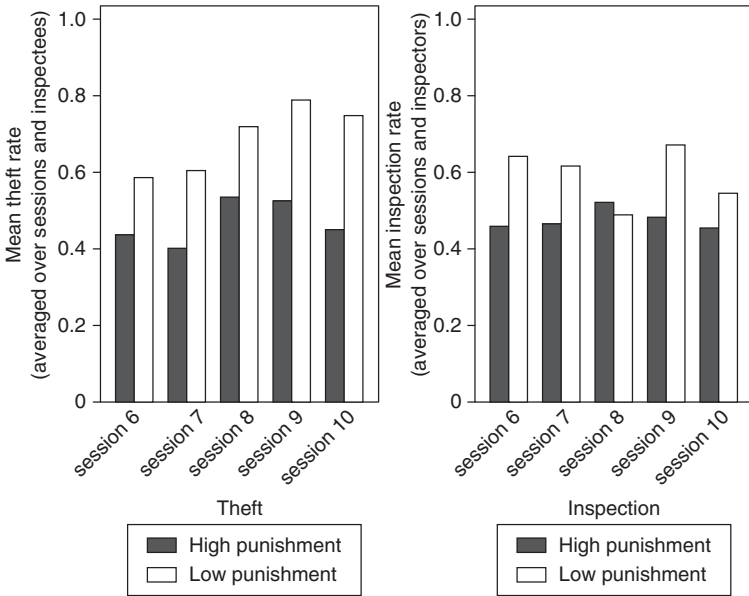


Figure 3. Mean punishment effect for each session (experiment 2)

increases for lower punishment. With logistic regressions it can be shown that punishment effects on inspection are significant on the 5% level for most sessions (except session 8 and 10). Over all sessions, inspection averages 48% for high punishment and 59% for low punishment. Hence, we can confirm hypothesis 2.5. Nevertheless, as in experiment 1, there is too much inspection for high punishment and too little inspection for low punishment.

In the following, we pool session data to analyze the statistical significance of the mean punishment effects. We estimate linear random intercept models. We have data of 2940 decisions clustered in 98 *inspectees* and 10 sessions, and 2940 decisions, clustered in 98 *inspectors* and 10 sessions. A linear random intercept model with random intercepts for subjects and sessions adjusts standard errors for clustered decisions in subjects and sessions.¹¹ We model separately theft and inspection for each treatment effect, strength of punishment and their order (low punishment as first or as second stage): for both, theft and inspection, we compare low versus high punishment, models (1) and (2), and order effects, models (3) and (4). Order effects can be interpreted as differences between increase and decrease of punishment. Models are

estimated with Stata 9.2, using the GLLAMM procedure (Rabe-Hesketh and Skrondal 2005).

Table 5 reports significant punishment effects for both, theft and inspection. *High punish* is a dummy variable, coded 0 for low punishment stages and 1 for high punishment stages (for experiment 1 and 2). Higher punishment causes 21% less theft, but also 15% less inspection. Effects are significant on the 0.1% level. Thus, hypothesis 2 can be confirmed that strategic interaction shifts effects of criminals' incentives to behavioral changes of inspectors.

However, there are also order effects (although not significant on the 5% level). *Punish increase* is a dummy variable, coded 0 for experiment 2 (high punishment first) and 1 for experiment 1 (low punishment first). There is 8% less theft and 9% less inspection, when punishment is increased compared to when it is decreased. These findings were not expected. However, we can interpret these findings with prospect theory (Tversky and Kahneman, 1981): an increase of punishment implies losses: crime becomes more expensive. In contrast, a decrease of punishment implies gains in terms of more attractive criminal opportunities. It is well known from prospect theory that actors react more sensitively to losses than to gains. Thus, punishment might deter more when it is increased. Note, however, that this is a *post hoc* interpretation of the data, which should be confirmed in future studies.

The random part of the model reveals that different subjects differ considerably in their average theft and inspection behavior: We can compare the lower 10% with the upper 10% of the population when we take the intercept \pm the z-value 1.28 times the standard deviation $\sqrt{\tau_0^2}$ for subjects. Those with little reluctance to steal (upper 10% of the criminal population) steal about 25% more than reluctant inspectees (lower 10% of the population). For inspection, the same result holds. Strongly motivated inspectors (upper 10% of the population) inspect roughly 25% more than inspectors with little motivation to inspect (lower 10%). Random session intercepts do not add much explanatory power. Subjects in different sessions do not vary much in their theft or inspection behavior. Therefore, random session intercepts are excluded from subsequent models.

4.4 Learning and behavioral dynamics

The static analysis revealed that subjects take strategic decision making into account, however less than predicted from game theory. Nevertheless, it might be the case that humans adapt over time and learn how to optimize their behavior in such inspection situations.

Table 5. Linear random intercept models for theft and inspection

	(1) Theft	(2) Inspect	(3) Theft	(4) Inspect
Fixed effects				
Intercept	0.65*** (0.027)	0.56*** (0.024)	0.58*** (0.031)	0.53*** (0.032)
High punish	-0.21*** (0.016)	-0.15*** (0.016)		
Punish increase			-0.083 (0.044)	-0.086 (0.045)
Random effects				
Decisions σ^2	0.20*** (0.0052)	0.20*** (0.0053)	0.21*** (0.0055)	0.20*** (0.0054)
Subjects τ_0^2	0.039*** (0.0069)	0.045*** (0.0074)	0.039*** (0.0069)	0.043*** (0.0071)
Sessions φ_0^2	0.0018 (0.0030)	$5.4e^{-19}$ ($6.2e^{-11}$)	0.00094 (0.0022)	$3.9e^{-18}$ ($1.1e^{-10}$)
LogLikelihood	-1873.2	-1902.8	-1953.2	-1939.5
Bic	3786.4	3845.5	3946.3	3918.9
N(decisions)	2940	2940	2940	2940

Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

2940 Decisions, 98 Subjects, 10 Sessions. The random effects denote the variances of the intercepts, clustered in subjects and sessions. High punish: Dummy, 0—low, 1—high punishment. Punish increase: Dummy, 0—experiment 2, 1—experiment 1.

For derivation of hypotheses, we use a parsimonious learning model known as *fictitious play*, described by Fudenberg and Levine (1998). Fictitious play assumes that agents start with an initial expectation of the distribution of the opponents' strategies. After the first move, agents learn about their opponents, update their expectations and meet an optimal decision, given the updated expectation and the payoffs. This learning rule is simple, forward-looking and mimics rational agents that adapt their behavior in situations of uncertainty. There are alternative learning models which can be described as backward looking (Macy 1993). Such 'Pavlovian' learning rules assume agents that have an initial aspiration level that they want to reach. If agents underscore, they change their aspiration and their strategy in a probabilistic way.¹² We use fictitious play, as it converges to mixed equilibria over time (Fudenberg and Kreps 1993), so that we can investigate how fast humans adapt.

The learning model can be described in three steps. First, criminals start with an initial expectation of being caught and inspectors with an

initial expectation of catching a criminal. While learning models predict behavioral changes for given expectations, they leave out statements about initial behavior, given no prior experience. But for the experiments, we have data to estimate these initial levels. We assume that inspectees and inspectors have an expectation that matches their initial decisions in the experiment.

Learning Phase 1. The initial belief of detection likelihood for inspectees and inspectors is between 0 and their indifference point if they committed a crime (inspectees), or did not inspect (inspectors). In the reverse case, it is between their indifference point and 1.

To determine the precise values of the initial expectation \hat{c}_i^0 that inspectee i expects to be inspected, we draw from a uniform distribution within the interval between 0 and $\frac{z_i}{p}$ if the inspectee committed a theft and between $\frac{z_i}{p}$ and 1 if she did not commit a theft. Likewise for inspectors, we determine the expectation of inspector j of catching a criminal in period 0, \hat{s}_j^0 , as a draw from a uniform distribution between 0 and $\frac{k}{r}$ if she did not perform control and between $\frac{k}{r}$ and 1 if she performed control.

In a second step, inspectees and inspectors update their prior expectation with the experienced behavior of their opponent.

Learning Phase 2. Inspectee i counts the number of experienced inspections κ_i^t until time step t and updates her prior estimated detection likelihood \hat{c}_i^{t-1} with the arithmetic mean of her full history such that $c_i^t = \frac{\kappa_i^t}{t}$. Inspectors update their estimated detection probability \hat{s}_j^t analogously.

As we have experimental data, there are two possibilities to update beliefs. The first option is to use the actual experimental data to update the agents' beliefs. Thus, inspectees and inspectors update their beliefs with their actual experiences: We define this updating rule as *unilateral learning*. Such an analysis allows us to investigate whether players follow the learning model for the case of their actual experiences. The second option is to use the experimental data only for the initialization of the prior beliefs in period $t = 0$. After the initial decision, agents update *both* simultaneously according to the learning model so that previous updates are drawn from simulated crime and inspection decisions. We call this updating rule *bilateral learning*. This analysis allows us to compare human decision making with decisions of perfectly updating agents that started with similar starting values.

In a third step, agents reach their decision by evaluating their utility with their estimated detection probabilities. Inspectees commit a crime and inspectors inspect if their expected utility is positive.

Learning Phase 3. Inspectee i commits a crime in period t if $\gamma x - \hat{c}_i^t p > 0$.
Inspector j inspects if $\hat{s}_j^t r - k > 0$.

We repeat the learning phases 2 and 3 for the same number of periods as the experiment so that we can compare experimental data with the learning model, more specifically with both versions of updating rules, the unilateral and the bilateral learning.

We employ linear growth curve models to compare the learning model with our empirical data. These models provide information whether crime and inspection increases or decreases over time for different treatment conditions. We estimate random intercepts and random periods to adjust standard errors for clustered decisions within subjects. Hence, we assume that subjects differ in their baseline motivation to steal and to inspect and in their learning behavior throughout the periods. We do not include session random intercepts because they did not explain much variation in the previous static analysis. Because the static analysis revealed differences for increasing versus decreasing punishment, we estimate separate models for experiment 1 and experiment 2. We use the same dummy variable *high punish* as reported in table 5. We include the baseline *period* variable. It is coded with 0 as starting value to make the intercept interpretable. Furthermore, it is divided by 15. Thus, it reports period effects for a whole stage, consisting of 15 periods, which facilitates interpretation. To test whether subjects adapt differently for low versus high punishment, the interaction *period* \times *low punish* is included. It is coded as *period* \times $(1 - \textit{high punish})$. Thus, the period variable reports period effects for high punishment and the interaction period \times low punish refers to the difference between high punishment and low punishment. Models are estimated with Stata 9.2 using the XTMIXED procedure. Table 6 reports estimation results.

For experiment 1, regression results show that subjects start with 50% theft and 50% inspection in the first period of low punishment. In contrast, the same subjects start with 34% less theft and 35% less inspection in the high punishment stage. Hence, we can report a considerable within subjects effect for strength of punishment. We see that inspectees start on average at the Nash prediction in the low punishment stage. Nevertheless, they increase their theft behavior for 16% in the high punishment stage, as can be seen from the effect of *period*.

Likewise, inspectors increase their inspection behavior for 14% for high punishment. Both increases are significant on the 5% level. Surprisingly, the period interaction $period \times low\ punish$ reveals no difference between low and high punishment stages: Theft and inspection increases for low and for high punishment. Thus, subjects start to steal and inspect at a moderate level and increase their behavior throughout periods, regardless of strength of punishment.

Experiment 2 starts with high punishment and concludes with low punishment. We might expect, from the static analysis, that effects are weaker for this experiment because it introduces ‘punishment gains’ instead of ‘punishment losses.’ Results show that subjects start experiment 2 with $79 - 35 = 44\%$ theft and with $55 - 12 = 43\%$ inspection in the first period of the high punishment condition. Strength of punishment has weaker effects on inspection, as already suspected. We see that subjects steal and inspect at roughly constant levels throughout the high punishment and the low punishment treatment: neither the period effect nor the period interaction with strength of punishment is significant. So subjects are less sensitive and less adaptive, when punishment is decreased compared to when it is increased, supporting prospect theory.

The random part of the models reveals that subjects differ in their baseline level of theft and inspection. These differences are about the same as in table 5. Random period effects support that subjects differ according to their learning behavior throughout the periods. When we calculate period effects for the lower 10% and for the upper 10% of the population, as we did for the static analysis for random intercepts, we observe differences for different subjects from 41% in model (3) to 54% in model (1) in theft and inspection increases throughout the periods. Taking model (1) as an example, at the top 10%, subjects increase their theft behavior of 42% throughout the periods in the high punishment stage. In contrast, at the bottom 10%, subjects decrease their theft behavior of 11% throughout the high punishment stage. In addition, there is a strong correlation between starting values of the first period and learning throughout the experiment: the negative correlation between τ_1^2 and τ_0^2 is strong and significant at the 0.01% level, as can be seen from the $\chi_2^2 LR$ value which reports the likelihood ratio test for the random period effect.¹³ Thus, originally conformist subjects increase theft behavior more than subjects who started with high crime levels. Control averse subjects increase their inspection activity more than ‘control freaks.’ This correlation supports the notion of strategic decision making as subjects adapt according to their social circumstances.

Table 6. Linear growth curve models for theft and inspection

	(1)	(2)	(3)	(4)
	Theft	Inspect	Theft	Inspect
Fixed effects	Experiment 1		Experiment 2	
Intercept	0.50*** (0.045)	0.50*** (0.051)	0.79*** (0.086)	0.55*** (0.090)
High punish	-0.34*** (0.091)	-0.35*** (0.089)	-0.35*** (0.083)	-0.12 (0.088)
Period	0.16* (0.065)	0.14* (0.060)	0.061 (0.057)	0.085 (0.061)
Period × low punish	0.051 (0.081)	-0.062 (0.079)	-0.13 (0.074)	-0.061 (0.078)
Random effects				
Random Period $\sqrt{\tau_1^2}$	0.21*** (0.031)	0.16*** (0.028)	0.16*** (0.027)	0.18*** (0.029)
Random Intercept $\sqrt{\tau_0^2}$	0.22*** (0.034)	0.28*** (0.037)	0.25*** (0.034)	0.24*** (0.034)
Correlation τ_1^2, τ_0^2	-0.72*** (0.095)	-0.68*** (0.112)	-0.41* (0.167)	-0.54*** (0.142)
Decisions $\sqrt{\sigma^2}$	0.44*** (0.0086)	0.43*** (0.0083)	0.41*** (0.0078)	0.44*** (0.0083)
LogLikelihood	-937.6	-904.3	-876.6	-962.5
χ^2_2 LR random period	47.9***	23.1***	27.4***	31.1***
BIC	1933.3	1866.8	1811.6	1983.6
N(decisions)	1440	1440	1500	1500

Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$
 $\sqrt{\tau_1^2}$ and $\sqrt{\tau_0^2}$ express the subjects' differences in their period effects and intercepts. More precisely, $\sqrt{\tau_1^2}$ and $\sqrt{\tau_0^2}$ denote the respective standard deviation. The correlation between the magnitude of the intercepts and the period effects is denoted with τ_1^2, τ_0^2 in the third row of the random effects. $\sqrt{\sigma^2}$ denotes the remaining random noise of the subjects' decisions, irrespective of their random intercepts and slopes. The χ^2_2 LR denotes the likelihood ratio test for the random period effect. This test compares with the respective Log-Likelihoods the additional explanatory power of considering different period effects for different subjects. BIC denotes the Bayesian information criterion. 1440 (1500) decisions, 98 (100) subjects, 10 sessions. High punish: dummy, 0–low, 1–high punishment. Period: 0 as starting value, divided by 15 to report period effect for whole stage. Period × low punish: period interaction with (1-high punish); reference category high punish.

For comparison of empirical data with the learning model, we plot the trajectories of crime and inspection decisions from (1) the experiments, (2) the unilateral learning model and (3) the bilateral learning model together in figure 4.

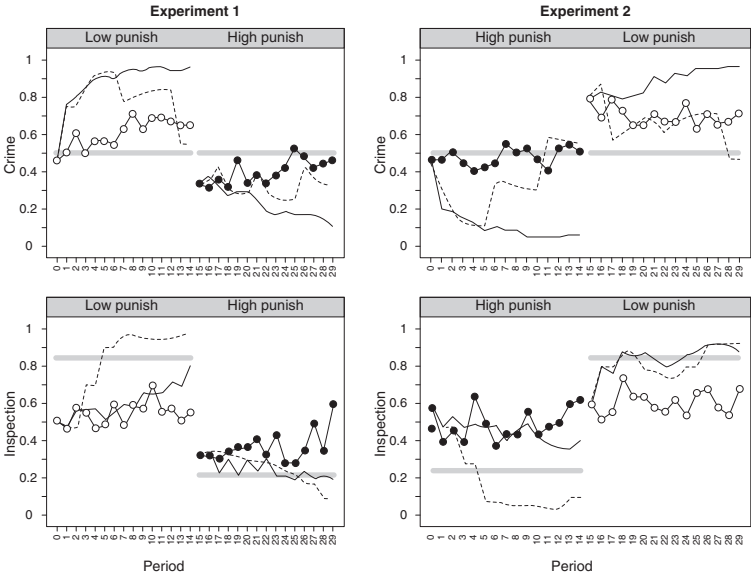


Figure 4. Humans learn slowly in the inspection game. The connected points in the upper figures show crime and in the lower figures inspection rates per period. The thick grey bar in each subfigure specifies the mixed Nash equilibrium. The solid and the dashed lines represent two different calibrations of the learning model fictitious play. Both models assume that agents update their expectation of crime detection and start with expectations that match their initial behavior. The solid lines specify *unilateral learning*, which assumes that agents update their expectations with the actually observed behavior of their opponent. The dashed lines specify *bilateral learning*, which assumes that agents update their expectations with the simulated, perfectly adapted behavior of their opponents. The figures show consistently that humans adapt slower than rational learning models predict. For low punishment, crime and inspection *increases* less than expected and for high punishment, crime and inspection *decreases* less than expected.

In both experiments, inspectors start with moderate inspection activities – not very high and not very low. However, the mixed Nash equilibria predict more extreme inspection activities – higher inspection rates for low punishment and lower inspection rates for high punishment. This implies that *both* inspectees and inspectors should *increase* their activities for low punishment and *decrease* their activities for high punishment.

The solid lines in figure 4 show the predicted learning curves of the unilateral learning model. More specifically, they show the optimal adaptation of ego for the actually experienced behavior of alter in the

experiment. The analysis of low punishment scenarios illustrates that humans only adapt hesitantly in the inspection game. Inspectors increase their behavior too little so that the two lines for observed and predicted trajectories divert more and more over time. Likewise, inspectees increase their criminal behavior too little, compared to the unilateral learning curves. In contrast, rationally updating agents would increase their criminal behavior much faster than observed in the experiment.

Our results hold for high punishment scenarios as observed and simulated trajectories divert more and more over time. More specifically, inspectors start with higher inspection activities than predicted by mixed Nash equilibria. Therefore, inspectors and inspectees are expected to decrease their activities, which is illustrated by the results of the unilateral learning model, indicated by the solid lines. Clearly, the differences increase over time, suggesting that humans are hesitant in updating their expectations.

The dashed lines represent the bilateral learning model. It shows the scenario, if inspectees and inspectors started with their actually observed behavior in the first period and afterwards update and decide simultaneously according to the learning model. The bilateral learning model therefore assumes higher rationality. In the early periods, the bilateral learning model follows closely the unilateral learning model. After a while, however, the simulated behavior diverts and approximates Nash predictions. It can be shown for longer simulation runs that the bilateral learning model generates oscillations around Nash predictions and differences between simulation results and Nash predictions decrease over time. Nevertheless, the results of the bilateral learning model are already quite close to Nash predictions after 15 periods, indicating that the number of periods in the experiment was sufficient, if only humans were rational updaters – but this is only partially the case.

We can draw similar conclusions from the static analysis of punishment effects and the analysis of learning dynamics: strategic decision-making affects criminal behavior *and* inspection activities. However, the power of strategic interaction is not as strong as implied by game theoretic reasoning. First, the strength of punishment affects both, criminals and inspectors. Second, criminals and inspectors are reluctant to change their prior expectations of detection probabilities. They rather stay close to their initial behavior than rapidly adapt according to a Bayesian updating process of detection probabilities.

5. Conclusion

Most rational choice approaches assume that higher punishments cause lower crime rates. However, many field studies cannot confirm such punishment effects on crime. A game theoretic approach offers a solid theoretical explanation why we do not find such effects. The approach relies on the opposite incentive structure between criminals (insp-ectees) and control agents (inspectors). Mixed Nash equilibria predict that higher punishment does not deter crime; higher punishment deters control.

A new laboratory experiment is presented, which manipulates the level of punishment while holding everything else constant. So far, there are virtually no laboratory experiments available on the game theoretic approach to crime which gives additional support for a laboratory test before testing hypotheses in the field.

Results show that stronger punishment causes lower inspection rates, which supports the main game theoretic implication. However, not only inspection is affected by punishment, but crime as well. Such mixed effects of strategic interaction are confirmed by dynamical analyses. Humans react strategically to punishment incentives, but not as strong as predicted by game theory. The behavioral dynamics in the experiments show that humans adapt slowly in inspection situations. The data is confronted with the learning model 'fictitious play.' Here, criminals and inspectors form an individual detection probability, perform their action if the expected utility is positive and update their detection probability with their experience from the last move. The comparison of experimental data with the results from the learning model suggest that humans are quite hesitant to adapt strategically to their social environment. While the simulation results for perfect bilateral updating show that simulated agents reach Nash equilibria at the end of the experiment quite closely, human subjects are less adaptive.

The inflexibility of humans to change their initial beliefs might not hold over a longer time, though. Crime and inspection rates of agents that use fictitious play oscillate around Nash predictions and get closer and closer. First, comparably low inspection rates motivate criminals to commit crimes. Inspectors will follow soon because high crime rates make inspection attractive. High inspection rates cause falling crime rates, which in turn crowds out inspection again. These oscillations, however, might take longer for human subjects than for perfectly adapting simulated agents so that oscillations might follow longer curves in empirical data. It might be the case that if we observed subjects just long

enough, we find stronger confirmation of learning theory. Therefore, future studies are advised to increase the number of periods.

For future studies, it might be possible to test further implications of strategic interdependence between dyads, which are motivated by the literature on crime and punishment. Here is a brief sketch of ideas, which provide possibilities for future laboratory experiments and analytic studies. (1) One might assume that criminals are unable to commit crimes when they are victims. (2) Inspectors could inspect all criminals. (3) Inspectors might compete to avoid the costs of inspection. (4) Inspectors could exchange information on criminals. (5) Inspectors could inspect other inspectors. (6) There is only one inspector, just like a tax agency with a randomized inspections procedure. (7) The overall crime rate affects payoffs of all players. (8) Victims are saved from inspection (tolerated gang wars). (9) Criminals have superstrategies to reciprocate victimization.

Our results have implications for political decision makers as they should take the rationality of control agents more strongly into account. It might be more efficient to offer more rewards for inspectors than using stronger punishments to reduce crime. Also, the finding of the differential impact of increasing compared to decreasing punishment can be used to design public policy programs of crime deterrence.

Acknowledgments

I gratefully acknowledge support from the DFG for financing experimental payoffs, research assistants (VO 684/5-1) and a research stay at the University of Michigan (VO 684/6-1). Many valuable comments and insightful discussions with Thomas Voss, Karl-Dieter Opp, Michael Berbaum, Andre Casajus, Holger Rauhut, Kurt Mühler, Ivar Krumpal, Roger Berger, and the insightful, constructive and detailed comments of two anonymous referees improved various parts of the work. Comments from the RCS conference (VIU), the ModSim conference (Kassel) and the ecological modeling workshop (UFZ) are gratefully acknowledged. I am grateful for research assistance in programming and subject recruitment from Isabel Kuroczka, Fabian Winter and Jana Adler.

Notes

1. For more optimistic reviews on deterrent effects of punishment compare Levitt (2002); Nagin (1998); Cameron (1988).

2. The situation of crime and control is comparable to penalty kicks in soccer, where we naturally accept that players should mix their strategies; see for example Chiappori et al. (2002); Berger and Hammer (2007).
3. For a more critical statement on experimental designs, especially about construct validity for measuring social preferences compare Levitt and List (2007). In contrast to this critique, the present study does not focus on measuring social preferences but on treatment effects for strength of punishment so that the proposed argumentation should hold. For a field study on control behavior see Rauhut and Krumpal (2008).
4. For an overview of contrasting propositions between rational choice approaches for crime and the labeling paradigm see Opp (1989).
5. Note that this example was first given by Morgenstern (1928: 98).
6. Note that it follows from a more strict game theoretic analysis, that Moriarty should have anticipated Holmes' move. Furthermore, Holmes should have anticipated Moriarty's anticipation and so forth. In conclusion, the only stable equilibrium is a probability mix of both stations.
7. If we assume for simplicity γ and x to be the same for both players.
8. Except session 3, where only 16 subjects showed up.
9. These characteristics, however, are of little importance, as it is known from other game theoretic experiments that demographic characteristics have little influence on strategic decision making in humans.
10. Regressions are estimated on the decision level. Robust standard errors, clustered for subjects, are used. The dummy-variable strength of punishment is included as the only predictor.
11. Using linear models despite categorical data is appropriate as we saw in figures 2 and 3 that mean theft and inspection behavior is bounded within 20% and 80%. Moreover, linear random intercept models were compared with logit random intercept models, revealing only minor differences. The main reason to choose linear models is that coefficients are far easier to interpret.
12. For a more detailed description of such backward-looking learning models see Macy (1991) and also Macy and Flache (2002).
13. We take halved p -values, as recommended by Snijders and Bosker (1999; 90) because variances cannot become negative.

References

- Anderson, Christopher M., and Louis Putterman. 2006. 'Do Non-Strategic Sanctions Obey the Law of Demand? The Demand for Punishment in the Voluntary Contribution Mechanism.' *Games and Economic Behavior* 54: 1–24.
- Becker, Gary S. 1968. 'Crime and Punishment: An Economic Approach.' *Journal of Political Economy* 76(2): 169–217.
- Berger, Roger, and Rupert Hammer. 2007. 'Die doppelte Kontingenz von Elfmeterschüssen.' *Soziale Welt* 4: 397–418.
- Bloomfield, Robert 1994. 'Learning a Mixed Strategy Equilibrium in the Laboratory.' *Journal of Economic Behavior and Organization* 25(3): 411–436.
- Bolton, Gary E., and Axel Ockenfels. 2000. 'ERC: A Theory of Equity, Reciprocity, and Competition.' *American Economic Review* 90(1): 166–193.

- Brams, Steven J. 1994. 'Game-Theory and Literature.' *Games and Economic Behavior* 6(1): 32–54.
- Camerer, Colin. 2003. *Behavioral Game Theory: Experiments in strategic interaction*. Princeton University Press.
- Camerer, Colin, George Loewenstein, and Matthew Rabin. 2004. *Advances in Behavioral Economics*. The Roundtable Series in Behavioral Economics. Princeton, NJ: Princeton University Press.
- Cameron, Samuel. 1988. 'The Economics of Deterrence: A Survey of Theory and Evidence.' *Kyklos* 41(2): 301–323.
- Chiappori, Pierre-Andre, Steven D. Levitt, and Timothy Groseclose. 2002. 'Testing Mixed-Strategy Equilibria When Players are Heterogenous: The Case of Penalty Kicks in Soccer.' *American Economic Review* 92: 1138–1151.
- Diekmann, Andreas. 2003. 'The Power of Reciprocity. Fairness, Reciprocity and Stakes in Variants of the Dictator Game.' *Journal of Conflict Resolution* 48(4): 487–505.
- Doob, Anthony N., and Cheryl M. Webster. 2003. 'Sentence Severity and Crime: Accepting the Null Hypothesis.' *Crime and Justice. A Review of Research* 28: 143–195.
- Falk, Armin, and Urs Fischbacher. 2002. 'Crime in the Lab. Detecting Social Interaction.' *European Economic Review* 46: 859–869.
- Fehr, Ernst, and Urs Fischbacher. 'Third-Party Punishment and Social Norms.' *Evolution and Human Behavior* 25(2): 63–87.
- Fehr, Ernst, and Simon Gächter. 2000. 'Cooperation and Punishment in Public Goods Experiments.' *American Economic Review* 90(4): 980–994.
- Fehr, Ernst, and Simon Gächter. 2002. 'Altruistic Punishment in Humans.' *Nature* 415(10): 137–140.
- Fehr, Ernst, and Herbert Gintis. 2007. 'Human Motivation and Social Cooperation: Experimental and Analytical Foundations.' *Annual Review of Sociology* 33: 43–64.
- Fehr, Ernst, and Klaus M. Schmidt. 1999. 'A Theory of Fairness, Competition, and Cooperation.' *Quarterly Journal of Economics* 114(3): 817–868.
- Fischbacher, Urs. 2007. 'Z-Tree. Zurich Toolbox for Ready-made Economic Experiments.' *Experimental Economics* 10(2): 171–178.
- Frey, Bruno S. 1994. 'How Intrinsic Motivation is Crowded Out and In.' *Rationality and Society* 6(3): 334–352.
- Frey, Bruno S., and Reto Jegen. 2001. 'Motivation Crowding Theory.' *Journal of Economic Surveys* 15(5): 589–611.
- Fudenberg, Drew, and David M. Kreps. 1993. 'Learning Mixed Equilibria.' *Games and Economic Behavior* 5(3): 320–367.
- Fudenberg, Drew, and David K. Levine. 1998. *The Theory of Learning in Games*. Cambridge, MA: MIT Press.
- Gächter, Simon, and Arno Riedl. 2005. 'Moral Property Rights in Bargaining with Infeasible Claims.' *Management Science* 51(2): 249–263.
- Gintis, Herbert. 2000. *Game Theory Evolving: A Problem-Centered Introduction to Modeling Strategic Behavior*. Princeton, NJ: Princeton University Press.
- Gneezi, Uri, and Aldo Rustichini. 2004. 'Incentives, Punishment and Behavior.' In *Advances in Behavioral Economics*, Roundtable Series in Behavioral Economics, edited by Colin Camerer, George Loewenstein, and Matthew Rabin, pp. 572–589. Princeton, NJ: Princeton University Press.
- Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey. 2003. 'Risk Averse Behavior in Generalized Matching Pennies Games.' *Games and Economic Behavior* 45(1): 97–113.

- Kavka, Gregory S. 1983. 'Hobbes War of All-Against-All.' *Ethics* 93(2): 291–310.
- Levitt, Steven D. 2002. 'Deterrence.' In *Crime: Public Policies for Crime Control*, edited by James Q. Wilson, and Joan Petersilia, pp. 435–450. Institute for Contemporary Studies Press.
- Levitt, Steven D., and John A. List. 2007. 'What Do Laboratory Experiments Measuring Social Preferences Reveal about the Real World?' *Journal of Economic Perspectives* 21(2): 153–174.
- MacCoun, Robert, and Peter Reuter. 1998. 'Drug Control.' In *The Handbook of Crime and Punishment*, edited by Michael Tonry, pp. 207–238. New York: Oxford University Press.
- Macy, Michael W. 1991. 'Learning to Cooperate: Stochastic and Tacit Collusion in Social Exchange.' *American Journal of Sociology* 97: 808–843.
- Macy, Michael W. 1993. 'Backward-Looking Social-Control.' *American Sociological Review* 58(6): 819–836.
- Macy, Michael W., and Andreas Flache. 2002. 'Learning Dynamics in Social Dilemmas.' *Proceedings of the National Academy of Sciences USA* 99: 7229–7236.
- Morgenstern, Oskar 1928. *Wirtschaftsprognose: Eine Untersuchung ihrer Voraussetzungen und Möglichkeiten* (Economic prediction: An examination of its conditions and possibilities). Wien: Verlag von Julius Springer.
- Moschini, Gian C. 2004. 'Nash Equilibrium in Strictly Competitive Games: Live Play in Soccer.' *Economic Letters* 85(3): 365–371.
- Nagin, Daniel S. 1998. 'Criminal Deterrence Research at the Outset of the Twenty-First Century.' *Crime and Justice. A Review of Research* 23: 1–42.
- O'Neill, Barry 1987. 'Nonmetric Test of the Minimax Theory of Two-Person Zerosum Games.' *Proceedings of the National Academy of Sciences of the United States of America* 84(7): 2106–2109.
- Opp, Karl-Dieter. 1989. 'The Economics of Crime and the Sociology of Deviant Behavior. A Theoretical Confrontation of Basic Propositions.' *Kyklos* 42(3): 405–430.
- Palacios-Huerta, Ignacio 2003. 'Professionals play Minimax.' *Review of Economic Studies* 70(2): 395–415.
- Palacios-Huerta, Ignacio, and Oscar Volij. 2008. 'Experientia docet: Professionals play Minimax in Laboratory Experiments.' *Econometrica* 76(1): 71–115.
- Perugini, Marco, Marcello Gallucci, Fabio Presaghi, and Anna Paola Ercolani. 2003. 'The Personal Norm of Reciprocity.' *European Journal of Personality* 17: 251–283.
- Rabe-Hesketh, Sophia, and Anders Skrondal. 2005. *Multilevel and Longitudinal Modeling Using Stata*. Stata Press.
- Rabin, Matthew. 1993. 'Incorporating Fairness into Game Theory and Economics.' *American Economic Review* 83(5): 1281–1302.
- Rapoport, Amnon, and Richard B. Boebel. 1992. 'Mixed Strategies in Strictly Competitive Games: A Further Test of the Minimax Hypothesis.' *Games and Economic Behavior* 4(2): 261–283.
- Rasmusen, Eric. 2004. *Games and Information: An Introduction to Game Theory*, 3rd edition. Cambridge, MA: Blackwell Publishing.
- Rauhut, Heiko, and Ivar Krumpal. 2008. 'Die Durchsetzung sozialer Normen in Low-Cost und High-Cost Situationen.' *Zeitschrift für Soziologie* 5: 380–402.
- Shachat, Jason M. 2002. 'Mixed Strategy Play and the Minimax Hypothesis.' *Journal of Economic Theory* 104(1): 189–226.
- Snijders, Thomas A. B., and Roel J. Bosker. 1999. *Multilevel Analysis: An Introduction to Basic and Advanced Multilevel Modeling*. London: Sage Publications.

- Taylor, Michael. 1976. *Anarchy and Cooperation*. New York: John Wiley and Sons.
- Taylor, Michael. 1987. *The Possibility of Cooperation*. Cambridge: Cambridge University Press.
- Tsebelis, George. 1989. 'The Abuse of Probability in Political Analysis: The Robinson Crusoe Fallacy.' *American Political Science Review* 1: 77–91.
- Tsebelis, George. 1990. 'Penalty has No Impact on Crime. A Game Theoretic Analysis.' *Rationality and Society* 2: 255–286.
- Tversky, Amos, and Daniel Kahneman. 1981. 'The Framing of Decisions and the Psychology of Choice.' *Science* 211(30): 453–458.
- Von Hirsch, Andrew, Anthony E. Bottoms, Elizabeth Burney, and Per-Olof Wikström. 1999. *Criminal Deterrence and Sentence Severity – An Analysis of Recent Research*. Oxford: Hart Publishing.
- Voss, Thomas, and Manuela Vieth. 2008. 'Kooperationsnormen und vergeltende Sanktionen – Experimentelle Untersuchungen.' *Arbeitsbericht des Instituts für Soziologie* 50.
- Walker, Mark, and John Wooders. 2001. 'Minimax Play at Wimbledon.' *American Economic Review* 91(5): 1521–1538.
- Zimring, Franklin E., Gordon E. Hawkins, and Sam Kamin. 2001. *Punishment and Democracy: Three Strikes and You're Out in California*. New York: Oxford University Press.

HEIKO RAUHUT is a post-doc researcher at the ETH Zurich, Chair of Sociology, in particular of Modeling and Simulation. His general interest concerns how social norms shape our lives. He is employing quantitative methodologies, such as laboratory experiments, survey data and agent based models for a thorough understanding of how social norms promote coordination and cooperation on one hand, but conflict and the collapse of social order on the other. Moreover, his interests expand to methodological questions such as quantitative social research methods, rational choice theory and empirical tests of game theoretic models. His current peer-reviewed publications deal with the enforcement of social norms, *Zeitschrift für Soziologie* (2008), learning dynamics in normative behavior, *JASSS* (2009), and the measurement of normative behavior, *Methoden, Daten, Analysen*, 2008. He is currently working on a book entitled 'Crime and punishment from a game theoretical perspective.'

ADDRESS: ETH Zurich, Chair of Sociology, in particular of Modeling and Simulation, UNO C15, Universitätsstrasse 41, 8092 Zürich, Switzerland [email: rauhut@gess.ethz.ch]