

Encouraging Compliance: Bonuses versus Fines in Inspection Games

by

Daniele Nosenzo^a, Theo Offerman^b, Martin Sefton^c and Ailko van der Veen^d

13 September 2012

Abstract

In this paper we examine the effectiveness of bonuses and fines in an ‘inspection game’, where costly inspection allows an authority to detect whether or not an individual complies with some standard of behavior. Standard game theoretic analysis predicts that in the inspection game non-compliant behavior is deterred by fines targeted at non-compliant individuals, but *encouraged* by bonuses awarded to compliant individuals. In an experiment we find that fines are effective in deterring non-compliance. However, in agreement with recent behavioral theories, we find that the effect of bonuses on compliance is much weaker than predicted.

Keywords: Compliance; Deterrence; Costly Monitoring; Rewards and Punishments; Bonuses and Fines; Inspection Games; Experiment.

JEL Classification Numbers: C72, C92, K42

Acknowledgements: We thank Daniel Seidmann, participants at the 2010 ESA Conference in Copenhagen, the 2010 CREED-CeDEx-CBESS Meeting in Amsterdam, and seminar audiences in Amsterdam for useful comments. Daniele Nosenzo acknowledges support from the Leverhulme Trust (ECF/2010/0636). Ailko van der Veen acknowledges support from the Dutch Science Foundation NWO (VICI nr 453-03-606). The Research Priority Area Behavioral Economics of the University of Amsterdam contributed to the costs of the experiments. We thank CREED programmer Jos Theelen for programming the experiment.

^a CeDEx, School of Economics, University of Nottingham, Nottingham, NG7 2RD, United Kingdom.

^b CREED, Department of Economics, University of Amsterdam, 1018 WB, Netherlands.

^c Corresponding Author. CeDEx, School of Economics, University of Nottingham, Nottingham, NG7 2RD, United Kingdom. Email: martin.sefton@nottingham.ac.uk. Tel: +44 (0) 115 846 6130. Fax: +44 (0) 115 951 4159.

^d CBESS, School of Economics, University of East Anglia, Norwich, NR4 7TJ, United Kingdom.

1. INTRODUCTION

There are many situations where authorities have preferences over individuals' choices. A tax authority wants taxpayers to truthfully report income, an employer wants an employee to work hard, a regulator wants a factory to comply with pollution regulations, police want motorists to observe speed limits, etc. A fundamental problem for authorities is how to induce compliance with desired behavior when individuals have incentives to deviate from such behavior. A standard approach is to monitor a proportion of individuals and penalize those caught non-complying.

To further encourage compliance, the authority may consider rewarding an individual who was inspected and found complying. In the context of tax compliance, in 2003 the National Tax Service (NTS) of Korea introduced a system of bonuses for taxpayers found to have high compliance levels: bonuses included benefits such as providing a three-year exemption from tax audit and preferential treatment from financial institutions, e.g. reduced interest rates on loans (see NTS Annual Report, 2004, p. 31).¹ As another example, in 2010 the Swedish National Society for Road Safety implemented a 'speed camera lottery' experiment, whereby a speed camera in Stockholm captured images of the license plates of all drivers, irrespectively of their speed: drivers found to obey the speed limit were entered into a lottery financed by the money collected from the fines on those caught speeding.²

Alternatively, authorities may consider increasing the sanctions on individuals who, upon inspection, are found not complying. For example, in 2004 Tesco attempted to reduce levels of unplanned absence by introducing tougher measures (including reductions in sick pay) against employees taking sick leave.³ In the tax compliance context, the Dutch government decided to increase the fine for undeclared savings from 100% to 300% in May 2009.⁴ Similarly, to deter speeding, a new traffic law in Illinois (USA) increased fines by 60% in September 2010.⁵

Fines and bonuses are also used to encourage workers in employer-employee relationships. They are particularly useful in short term relationships where other mechanisms

¹ There is some evidence of the success of rewards in a tax compliance context from both laboratory (Alm et al., 1992) and field studies (Wan, 2010). Wan (2010) shows that lottery receipt policies in China which essentially subsidize businesses for issuing transaction receipts, have led to higher rates of compliance.

² See <http://www.freakonomics.com/2010/12/14/paying-drivers-to-not-speed/>.

³ See <http://news.bbc.co.uk/1/hi/business/3719183.stm>.

⁴ See <https://zoek.officielebekendmakingen.nl/kst-31301-16.html>. In a meta-analysis of laboratory experiments Blackwell (2007) finds strong evidence that increasing the penalty rate leads to higher tax compliance.

⁵ See <http://abclocal.go.com/wls/story?section=news/local&id=7651607>.

that may discourage shirking, such as reputational concerns, are weak, e.g., in industries with high turnover, or in industries that make use of temporary workers as is the case with day laborers who are frequently used in sectors like construction and farming (Greenhouse, 2005; Hanson and Bell, 2007).

In this paper we present experimental evidence on which of these two mechanisms – rewards targeted at compliant versus sanctions targeted at non-compliant individuals – is most successful in encouraging compliance. Our paper contributes to a vast empirical literature on deterrence. This literature, surveyed in Freeman (1983), Cameron (1988) and Ehrlich (1996), offers convincing evidence that both negative and positive incentives have a deterrent effect on non-compliance. Freeman (1983) concludes that the deterrent effects of positive measures are weaker than the negative deterrent effects of both the probability and the severity of sanctions. However, Ehrlich (1996, p. 61) identifies many econometric problems with existing studies, and concludes that “*the present evidence does not allow one to conclude that positive incentives are either less or more potent than negatives ones*”.

As a framework for studying these issues we use an ‘inspection game’, which we describe in Section 2.⁶ In this game an authority chooses to inspect or not, and an individual chooses to comply or not. The authority and the individual have conflicting preferences: the authority’s best response to non-compliance is to inspect, the individual’s best response to inspection is compliance, the authority’s best response to compliance is not to inspect, and the individual’s best response to non-inspection is not to comply. Thus, the unique Nash equilibrium is in mixed strategies, with positive probabilities of inspection and non-compliance. We study how the inspection and non-compliance probabilities are affected by the introduction of fines for non-compliant behavior or bonuses for compliant behavior. From a theoretical point of view, we show that fines deter non-compliance. On the other hand, bonuses for compliant behavior increase the equilibrium probability of non-compliant behavior. Thus, according to standard game theoretical reasoning, fines, and not bonuses, should be used to encourage compliance in such settings.

In Section 3 we report an experiment designed to test these predictions about the comparative effectiveness of bonuses and fines in the inspection game. Our inspection game is

⁶ This game is also referred to as the ‘police-public game’, especially in the political science literature – see, e.g., Tsebelis (1989); Bianco et al. (1990); Cox (1994); Andreozzi (2004). See Avenhaus et al. (2002) for a review of the theory of inspection games.

framed as an employer-worker scenario where an employer can either inspect or not and a worker can either work (comply) or shirk (not comply). We designed three experimental treatments, each consisting of two parts. The first part was identical across treatments: subjects played a control version of the inspection game where the employer pays the worker a flat wage, unless she is inspected and found shirking in which case the wage is not paid. In the second part of the BONUS treatment, subjects played a version of the game where the employer paid an additional bonus to the worker when the employer inspected and the worker worked. In the second part of the FINE treatment, subjects played a version of the game where the worker paid a fine to the employer if the employer inspected and the worker shirked. Finally, in the second part of the CONTROL treatment, subjects continued playing the same game as in the first part. Notice that the inspection game in the first part already contains the possibility of a fine: if the employer inspects and finds the worker shirking, she refrains from paying the wage. In the second part, we consider the possibility that the fine is increased when the worker is found shirking, or the possibility that the worker receives a bonus on top of the wage when found performing well. This design allows us to examine whether the introduction of an extra fine or bonus is more effective in encouraging working/discouraging shirking in applications that are characterized by the inspection game. In addition, we are able to compare the efficiency properties of rewarding versus punishing mechanisms.

We report our results in Section 4. We find that fines are more effective than bonuses in encouraging working and in raising combined earnings. This is in line with standard game theoretic predictions. However, the prediction that bonuses discourage working receives little support: although subjects shirk on average slightly more in the BONUS treatment than CONTROL the difference is small and not statistically significant. Moreover, the prediction that introducing bonuses will reduce combined earnings is not supported: the losses to employers are almost exactly offset by gains to workers. We show that observed deviations from Nash equilibrium predictions can be explained quite well by recent behavioral theories.

In Section 5 we discuss these results in relation to the existing literature on deterrence and the effectiveness of positive or negative incentives. In contrast with much of the existing literature on deterrence, the game theoretic prediction for our framework is that bonuses increase shirking. We discuss how replacing this analysis with an aggregate market analysis changes predictions. The detrimental effect of incentives is also a feature of the literature on motivation

crowding-out (Bowles and Polania-Reyes, 2012), which we also discuss in Section 5. We note that the mechanism behind the detrimental effect of incentives is different in our setting as it is driven by standard game theoretic analysis of pecuniary rather than non-pecuniary incentives. Section 5 also includes a discussion of how our experiment relates to other experimental studies comparing the effectiveness of punishments and rewards in somewhat different contexts, including tax compliance (e.g., Alm et al., 1992), voluntary contributions to public goods (e.g., Sefton et al., 2007), and employer-employee relations (e.g., Fehr et al., 2007). Section 6 concludes.

2. INSPECTION GAMES

2.1 Nash Equilibrium

We study a simple simultaneous move inspection game. For convenience, we follow Fudenberg and Tirole (1992, p. 17) and frame this as a game between an employer who can either inspect (I) or not inspect (N), and a worker who can either work (W) or shirk (S) even though the examples in the introduction show that the inspection game is relevant for many other applications as well. In the baseline version of the game the employer incurs a cost of h from inspecting, and working results in the worker incurring a cost of c and the employer receiving revenue of v . The employer pays the worker a wage of w , unless the worker shirks and the employer inspects. The resulting payoffs are shown in the leftmost panel of Figure 1. We assume that all variables are positive and $v > c$, $w > h$, $w > c$. Note that joint payoffs are maximized when the worker works and the employer does not inspect.

Figure 1. Inspection Games*

		Baseline Inspection Game		Inspection Game with Fines		Inspection Game with Bonuses									
		W	S	W	S	W	S								
I	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="border: 1px solid black; padding: 5px;">$v - w - h$</td> <td style="border: 1px solid black; padding: 5px;">$-h$</td> </tr> <tr> <td style="border: 1px solid black; padding: 5px;">$w - c$</td> <td style="border: 1px solid black; padding: 5px;">0</td> </tr> </table>	$v - w - h$	$-h$	$w - c$	0	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="border: 1px solid black; padding: 5px;">$v - w - h$</td> <td style="border: 1px solid black; padding: 5px;">$f - h$</td> </tr> <tr> <td style="border: 1px solid black; padding: 5px;">$w - c$</td> <td style="border: 1px solid black; padding: 5px;">$-f$</td> </tr> </table>	$v - w - h$	$f - h$	$w - c$	$-f$	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="border: 1px solid black; padding: 5px;">$v - w - b - h$</td> <td style="border: 1px solid black; padding: 5px;">$-h$</td> </tr> <tr> <td style="border: 1px solid black; padding: 5px;">$w + b - c$</td> <td style="border: 1px solid black; padding: 5px;">0</td> </tr> </table>	$v - w - b - h$	$-h$	$w + b - c$	0
	$v - w - h$	$-h$													
$w - c$	0														
$v - w - h$	$f - h$														
$w - c$	$-f$														
$v - w - b - h$	$-h$														
$w + b - c$	0														
N	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="border: 1px solid black; padding: 5px;">$v - w$</td> <td style="border: 1px solid black; padding: 5px;">$-w$</td> </tr> <tr> <td style="border: 1px solid black; padding: 5px;">$w - c$</td> <td style="border: 1px solid black; padding: 5px;">w</td> </tr> </table>	$v - w$	$-w$	$w - c$	w	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="border: 1px solid black; padding: 5px;">$v - w$</td> <td style="border: 1px solid black; padding: 5px;">$-w$</td> </tr> <tr> <td style="border: 1px solid black; padding: 5px;">$w - c$</td> <td style="border: 1px solid black; padding: 5px;">w</td> </tr> </table>	$v - w$	$-w$	$w - c$	w	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="border: 1px solid black; padding: 5px;">$v - w$</td> <td style="border: 1px solid black; padding: 5px;">$-w$</td> </tr> <tr> <td style="border: 1px solid black; padding: 5px;">$w - c$</td> <td style="border: 1px solid black; padding: 5px;">w</td> </tr> </table>	$v - w$	$-w$	$w - c$	w
	$v - w$	$-w$													
$w - c$	w														
$v - w$	$-w$														
$w - c$	w														
$v - w$	$-w$														
$w - c$	w														

* Employer is ROW player, Worker is COLUMN player. Within each cell, the Employer's payoff is shown at the top and the Worker's payoff at the bottom.

The baseline game has a unique Nash equilibrium where the employer inspects with probability $p_c = c/w$ and the worker shirks with probability $q_c = h/w$. In this equilibrium the employer receives a payoff of $\pi_c^{employer} = v - w - hv/w$, the worker receives a payoff of $\pi_c^{worker} = w - c$, and joint payoffs are $\pi_c = v - c - hv/w$.

We now compare two possibilities for deterring shirking relative to the baseline version of the game: imposing an additional fine on workers inspected and found shirking, versus paying a bonus to workers inspected and found working.

Suppose an additional fine f is imposed on a worker caught shirking, resulting in the payoff matrix shown in the middle panel of Figure 1. Note that the fine is a transfer between the worker and the employer. Now the unique Nash equilibrium has the employer inspect with probability $p_f = c/(w + f)$ and the worker shirk with probability $q_f = h/(w + f)$. Thus, according to Nash equilibrium, fines discourage both inspections and shirking. In Nash equilibrium expected payoffs are $\pi_f^{employer} = v - w - hv/(w + f)$, and $\pi_f^{worker} = w - c$, and so the employer benefits from the introduction of fines, while the worker's payoff is independent of fines. According to Nash equilibrium, fines enhance efficiency because joint payoffs are reduced by shirking and/or inspecting, and both of these are discouraged by a fine on workers caught shirking.

Next, we examine the case where the employer pays a bonus b to a worker who is inspected and found to have worked. The payoff matrix for this game is shown in the rightmost panel of Figure 1. Now in equilibrium the employer inspects with probability $p_b = c/(w + b)$ and the worker shirks with probability $q_b = (h + b)/(w + b)$. According to Nash equilibrium bonuses reduce the probability of inspection and *increase* the probability of shirking. The worker's equilibrium payoff is $\pi_b^{worker} = w - c + cb/(w + b)$, increasing in b , while the employer's is $\pi_b^{employer} = v - w - v(h + b)/(w + b)$, decreasing in b . Overall, bonuses reduce joint payoffs because the beneficial effect of less frequent inspection is outweighed by the detrimental effect of increased shirking.

As is well known, comparative static predictions based on mixed strategy Nash equilibrium can often be counter-intuitive. This is because a player's equilibrium probability must keep her opponent indifferent among actions, and so a player's *own* decision probabilities are determined by the *opponent* payoffs and not by *own* payoffs. Consider, for example, how the introduction of a bonus affects own-payoffs from the perspective of the worker. Introducing the bonus has no effect on the expected payoff from shirking, but increases the expected payoff from

working (for a given inspection probability). Based on this ‘own-payoff effect’, one might expect the worker to shirk less frequently following the introduction of bonuses. However, the Nash equilibrium prediction goes in the opposite direction: bonuses lead to an *increase* in the equilibrium shirking probability.

Figure 2 presents the games that we actually used in the experiment. For these parameters the Nash equilibrium inspection probabilities are 0.75 in the Baseline Inspection Game and 0.375 in the games with Fines and Bonuses. The Nash equilibrium shirking probabilities are 0.4 in the Baseline Inspection Game, decreasing to 0.2 in the Game with Fines and *increasing* to 0.70 in the Game with Bonuses.⁷

Figure 2. Parameterization of the Inspection Games Used in the Experiment*

		Baseline Inspection Game		Inspection Game with Fines		Inspection Game with Bonuses	
		W	S	W	S	W	S
I	W	52	12	52	32	32	12
	S	25	20	25	0	45	20
N	W	60	0	60	0	60	0
	S	25	40	25	40	25	40

* Employer is ROW player, Worker is COLUMN player. Within each cell, the Employer’s payoff is shown at the top and the Worker’s payoff at the bottom.

2.2 Alternative Equilibrium Concepts

Previous experimental work (e.g., Ochs, 1995; Goeree and Holt, 2001; Goeree et al., 2003) shows that counterintuitive Nash equilibrium predictions are often rejected by the data: changing a player’s *own* payoff does have an impact on that player’s decision probabilities. Goeree and Holt (2001) observe own-payoff effects in one-shot games; Ochs (1995) and Goeree et al. (2003) observe own-payoff effects even after players have had ample opportunities to learn. Note that own-payoff effects may either reinforce or counteract equilibrium forces. Introducing fines into the inspection game generates an own-payoff effect that pulls workers’ behavior in the same

⁷ Point earnings were derived from the game described in Figure 1 with $v = 60$, $c = 15$, $h = 8$, $w = 20$, and with 20 points added to all outcomes to ensure that subjects could not make losses in any of the games used in the experiment. These parameters were chosen so that Nash equilibrium probabilities are not too close to 0, 0.5 or 1 (all probabilities lie in the intervals [0.2, 0.4] or [0.6, 0.8]). We also sought separation between games with and without bonuses or fines so that, where a change in behavior is predicted by standard theory, the predicted change in probabilities across games is at least 20 percentage points.

direction as Nash equilibrium predictions: introducing fines does not change the expected payoff from working but does reduce the expected payoff from shirking. Thus the own-payoff effect discourages shirking, and this is consistent with the Nash equilibrium comparative static prediction. Similarly, own-payoff effects reinforce Nash equilibrium predictions about inspection probabilities in the inspection game with bonuses, but counteract Nash equilibrium predictions in inspection games with fines.

Two recent behavioral theories can account for own-payoff effects: Impulse Balance Equilibrium (IBE) and Quantal Response Equilibrium (QRE).⁸ Selten and Chmura (2008) provide a more general discussion for IBE and Brunner et al. (2011) for QRE. For the interested reader, the technical details of the models are reported in Appendix A. Here we just briefly discuss how each concept incorporates own-payoff effects.

In QRE players' choices are stochastic and decision probabilities depend on a parameter λ .⁹ The employer's inspection probability is given by $p = \frac{e^{\lambda E_{\text{employer}}[I]}}{e^{\lambda E_{\text{employer}}[I]} + e^{\lambda E_{\text{employer}}[NI]}}$ where

$E_{\text{employer}}[I]$ represent the employer's expected utility from inspecting (which depends on the own payoffs associated with inspecting and on the worker's probability of shirking). Thus, increasing either own payoff associated with inspecting will increase the probability of inspecting, *ceteris*

paribus. Similarly, the worker's shirking probability is given by $q = \frac{e^{\lambda E_{\text{worker}}[S]}}{e^{\lambda E_{\text{worker}}[S]} + e^{\lambda E_{\text{worker}}[W]}}$.

Again, increasing either own payoff from shirking will increase the probability of shirk, *ceteris paribus*. The quantal response equilibrium probabilities solve this pair of simultaneous equations.

In IBE players' choices are also stochastic and are governed by 'impulses' that players experience when their choices lead to lower payoffs than those they would have had by choosing differently. Thus, also in IBE a player's own payoffs affect the choice probabilities. An important feature of the model is that the strength of impulses depends on whether these come from payoffs above or below a player's reference point (the maximum payoff a player can

⁸ These models are particularly suited to a setup like ours where subjects get a lot of experience and feedback. Several other behavioral models can account for own-payoff effects, notably level-k models (Nagel, 1995; Stahl and Wilson, 1995) or cognitive hierarchy models (Camerer et al., 2004), although these are not so well-suited to our repeated play environment.

⁹ When $\lambda = 0$ players choose actions equi-probably and in the limit as λ approaches ∞ players always choose their best-response.

guarantee herself). Payoffs above this level are interpreted as ‘gains’, whereas payoffs below this level are interpreted as ‘losses’, and players weight losses twice as much as gains.

The fact that IBE is augmented by loss-aversion whereas Nash equilibrium and QRE are not has generated a recent debate about whether the incorporation of loss-aversion is what drives the observed differences in performance across these equilibrium concepts (Selten and Chmura, 2008; Brunner et al., 2011; Selten et al., 2011). To examine this possibility, we also consider predictions made by Nash equilibrium and QRE when these concepts are augmented with loss-aversion. Full details on QRE, IBE and how loss aversion can be included in these models are reported in Appendix A.

Table 1 presents the predictions made by the alternative concepts. For QRE predictions are based on a value of λ estimated from our experimental data, as we shall explain in Section 4. For the estimated value of λ QRE predictions are generally close to Nash equilibrium predictions. IBE predictions differ markedly from Nash equilibrium when own payoff and Nash equilibrium effects are in conflict: the IBE predicted probability of shirking in the inspection game with bonuses is 43% (versus the 70% Nash prediction) and the predicted probability of inspecting in the inspection game with fines is 61% (versus 37.5%).

Table 1. Predicted Choice Probabilities

	<i>Probability of Shirking</i>			<i>Probability of Inspecting</i>		
	Baseline Inspection Game	Inspection Game with Fines	Inspection Game with Bonuses	Baseline Inspection Game	Inspection Game with Fines	Inspection Game with Bonuses
Nash	0.40	0.20	0.70	0.75	0.375	0.375
QRE ($\lambda=0.989$)	0.46	0.19	0.68	0.76	0.41	0.35
IBE	0.41	0.16	0.43	0.68	0.61	0.40
Nash ^{with loss-aversion}	0.25	0.11	0.54	0.60	0.23	0.33
QRE ^{with loss-aversion} ($\lambda=0.289$)	0.42	0.10	0.46	0.69	0.47	0.36

Table 1 also reports predictions made by Nash equilibrium and QRE when these concepts are augmented with loss-aversion. Incorporating loss-aversion into QRE makes the predictions closer to those made by IBE. Incorporating loss-aversion into Nash generally reduces inspection and shirking probabilities, although the predictions are quite different from the alternative models.

3. EXPERIMENTAL DESIGN AND PROCEDURES

The experiment consisted of fifteen sessions at the University of Nottingham. Ten subjects participated in each session. Subjects were recruited from a campus-wide distribution list and no subject participated in more than one session.¹⁰ No communication between subjects was permitted throughout a session. At the beginning of a session subjects were randomly assigned to computer terminals and were informed that the experimental session would consist of two parts, during each of which they could earn ‘points’. Subjects were also told that their cash earnings for the session would be based on all points accumulated in both parts of the experiment.

Instructions for Part One were then distributed and read aloud. At the end of these subjects had to answer a series of questions to test their comprehension of the instructions. A monitor checked the answers and dealt with any questions in private. We did not continue with the experiment until all subjects had correctly answered all the questions. Part One then consisted of 40 rounds. At the beginning of the first round subjects learned their role: five subjects were assigned the role of ‘Employer’ and five the role of ‘Worker’. Subjects kept these roles for the entire session (i.e. for both Part One and Part Two). Across rounds subjects were randomly matched in pairs consisting of one Employer and one Worker, and in each round each pair played the Baseline Inspection Game, with Employers choosing “Inspect” or “Not Inspect” and Workers choosing “High Effort” (i.e., work) or “Low Effort” (shirk). At the end of each round subjects were informed of their own and their opponents’ choices and point earnings. Subjects were also shown their accumulated point earnings and a table with the distribution of choices across all subjects in the session for the previous twenty rounds.

At the end of Part One subjects were given instructions for Part Two, which were then read aloud. These explained that the second part consisted of another 80 rounds, again with pairings randomly determined at the beginning of each round. In our five CONTROL sessions these rounds continued using the Baseline Game payoffs. In our five FINE sessions, earnings in each round of Part Two were based on the Game with Fines, and in our five BONUS sessions Part Two earnings were based on the Game with Bonuses (see Figure 2 in the previous section).

At the end of Part Two subjects were paid in cash according to their accumulated point earnings from all rounds using an exchange rate of £0.004 per point. Sessions took about 40


¹⁰ Subjects were recruited using the online recruitment system ORSEE (Greiner, 2004). See Appendix B for instructions.

minutes on average and earnings ranged between £10.2 and £23.1, averaging £14.9 (approximately US\$24 at the time of the experiment).

4. RESULTS

Figure 3 displays the smoothed proportions of inspecting and shirking decisions across all the rounds of the experiment. For some cases there is a clear change in behavior in round 41, following the transition from Part One to Part Two and the introduction of fines or bonuses, but otherwise the observed proportions appear quite stable across rounds. Table 2 reports treatment-level averages and Figure 4 session-level averages of the proportions of shirking and inspecting over the last 20 rounds of each Part of the experiment. The Nash equilibrium predictions for choice probabilities are also reported for comparison.

Figure 3. Proportions of Shirking (left panel) and Inspecting (right panel) across Treatments*




*In each round, the average is displayed of the proportions of (max) 5 previous rounds, the current round and (max) 5 future rounds.

Table 2. Choice Proportions across Treatments, Average by Treatment*

	Part One			Part Two		
	CONTROL	FINE	BONUS	CONTROL	FINE	BONUS
Proportion of Shirking	0.39	0.52	0.45	0.44	0.23	0.50
<i>Nash</i>	0.40	0.40	0.40	0.40	0.20	0.70
Proportion of Inspecting	0.80	0.77	0.78	0.81	0.62	0.45
<i>Nash</i>	0.75	0.75	0.75	0.75	0.375	0.375

*Table shows the proportion of shirking/inspecting decisions in the last 20 rounds of each Part of the experiment.

Figure 4. Choice Proportions across Treatments, Average by Session



The first 40 rounds of the experiment (Part One) are common to the three treatments, and we do not find any significant differences in the proportions of shirking or inspecting across treatments (a Kruskal-Wallis test comparing workers’ shirking behavior across the three treatments in Part One yields a p-value of 0.37; a Kruskal-Wallis test comparing employers’ inspecting behavior across the three treatments in Part One yields a p-value of 0.78).¹¹ Averaged across all sessions the observed proportion of shirking decisions is 45% and the observed proportion of inspecting decisions is 78% (see Table 2) However, we observe substantial heterogeneity in shirking and inspecting across sessions (see Figure 4): the proportion of shirking varies from 25% to 80% across sessions, and the proportion of inspecting decisions varies from 62% to 92%.

In Part Two of the experiment the proportions of shirking and inspecting diverge significantly across treatments (Kruskal-Wallis test: $p = 0.02$ for shirk, and $p = 0.01$ for inspect). Clearly, the changes in payoff matrices introduced in Part Two of the different treatments caused subjects to adjust their behavior. For pair-wise statistical comparisons between treatments we use Mann-Whitney rank-sum tests. As predicted, we find less shirking in FINE (23%) than in CONTROL (44%), and the difference is statistically significant ($p = 0.02$). Although Nash equilibrium predicts workers will shirk considerably more in BONUS than in CONTROL (70%

¹¹ Our non-parametric analysis is based on two-tailed tests applied to 5 independent observations per treatment. We consider data from each session as one independent observation. Tests are applied to averages based on the last 20 rounds of each Part of the experiment. The data analysis does not lead to different results if we focus on all rounds.

vs. 40%), shirking in BONUS is only slightly higher than in CONTROL (50% vs. 44%), and the difference is not statistically significant ($p = 0.55$). As for the frequencies of inspections, these are lower in any of the sessions of FINE and BONUS than in any of the five sessions of CONTROL (see Figure 5). Thus, we detect statistically significant differences in inspecting between FINE and CONTROL ($p = 0.01$), and between BONUS and CONTROL ($p = 0.01$). We also note, however, that the proportion of inspections in FINE is considerably higher than predicted (62% vs. 37.5%), while the proportion of inspections in BONUS is closer to the theoretical level (45% vs. 37.5%). In fact, whereas Nash equilibrium predicts that introducing bonuses or fines have exactly the same effect on inspection probabilities, we find a statistically significant difference in the proportions of inspections between FINE and BONUS ($p = 0.01$).¹²

Overall, whereas Nash equilibrium predictions seem to capture well the comparative static effects of fines on aggregate shirking behavior and bonuses on aggregate inspecting behavior, they do not capture observed effects of fines on inspections or bonuses on shirking. It is notable that the instances where Nash predictions fail are those where own-payoff effects, as discussed in Section 2, work in the opposite direction to Nash equilibrium effects.

A natural question is then whether the alternative equilibrium concepts described in Section 2.2 can account for the observed deviations from Nash predictions. Figure 5 shows the observed and predicted effects of the introduction of bonuses and fines on the probability of shirking and inspecting using the prediction of Section 2. For QRE, as in Selten and Chmura (2008) and Brunner et al. (2011), we calculate the best fitting overall estimate for λ in our data by minimizing the sum of mean squared distances of the predicted QRE probabilities from the observed session-averaged choice probabilities in the experiment. This yields an estimated λ of 0.989 for the standard QRE model, and 0.289 when the QRE model is augmented with loss-aversion. These estimated values of λ were obtained using data from Part One as this allows us to make out-of-sample predictions for behavior in the games used in Part Two of the experiment.

¹² What are the implications of these findings for the distribution of earnings and efficiency? Relative to CONTROL, the introduction of bonuses increases the worker's payoff from 24.2 to 32.7 points ($p = 0.01$), and decreases the employer's payoff from 34.5 to 26.1 points ($p = 0.02$). The net effect is that joint earnings are slightly higher than in CONTROL, although the effect is not statistically significant ($p = 0.85$). Introducing fines decreases the worker's payoff by 1.7 points relative to CONTROL ($p = 0.06$), and increases the employer's payoff by 12.6 points ($p = 0.01$). Joint earnings are significantly higher in FINE than in CONTROL ($p = 0.01$).

Figure 5. Changes in Shirk (left) and Inspect (right) after introduction of Bonuses and Fines.




Figure 5 shows that the comparative static effects observed in our experiment are generally better captured by IBE and QRE with loss-aversion than by Nash equilibrium analysis or by QRE without loss-aversion. When Nash equilibrium effects and own-payoff effects work in the same direction (i.e. for the impact of fines on shirking and the impact of bonuses on inspections) there is little to choose among the various solution concepts. When Nash equilibrium effects and own payoff effects work in opposite directions (i.e. for the impact of fines on inspecting and the impact of bonuses on shirking), Nash equilibrium (with or without loss-aversion) is outperformed by the alternative concepts. Among these, IBE and QRE augmented by loss-aversion perform better than QRE without loss-aversion. Nash equilibrium predicts that bonuses increase shirking by 30% relative to CONTROL, whereas shirking only increases by about 6% in our data. This observed effect compares quite favorably with the comparative static predictions made by IBE (a predicted 2% increase in shirking) and QRE augmented by loss-aversion (a predicted 4% increase), but not with the comparative static predictions made by QRE without loss-aversion (a predicted 22% increase). Similarly, Nash equilibrium predicts that fines reduce the inspection rate by about 37% relative to CONTROL, whereas inspection rates actually fall by about 19%. QRE without loss-aversion predicts a decrease in inspecting by 35%, whereas the predicted magnitude of the decrease is smaller in IBE and QRE with loss-aversion (respectively, 7% and 22%).

If the data were in accordance with an equilibrium, there would not be any relation between the history that subjects observe and how they play. Given that we observe differences between the actual data and any of the equilibrium concepts, it seems likely that the way subjects play is influenced by what they observe in their own session. An interesting question is then

whether they are more affected by the result of their own interaction in the previous round or by the social history information that they received about play in the previous 20 rounds of their session. To examine this we estimate regressions of the choice probabilities on past information from previous rounds. Specifically, for employers we examine how the probability of inspecting depends on the behavior of the worker they were matched with in the previous round, and on the average shirking rate in their session in the previous 20 rounds. Similarly, we estimate how the probability of shirking depends on inspecting behavior of the employer they were matched with in the previous round, and on the average inspection rate in the session.

Table 3 reports the results of logit regressions based on the last 60 rounds of Part Two (we drop the first 20 rounds of Part Two because then the proper social history was not yet complete). The probability of inspecting is systematically influenced by workers' shirking behavior: the more workers shirk, the more employers inspect. Moreover, while employers are significantly affected by the aggregate levels of shirking in the session, they are not influenced by the behavior of the workers they are matched with in the previous round. For workers the picture is similar: shirking rates are lower the higher is the probability of inspecting. Workers are significantly affected both by the aggregate levels of inspecting in the session and the behavior of the employers they are matched with in the previous round.

Table 3. History Dependence of Choices

	Inspect	Shirk
Shirk-1	0.16 (0.08)	
Shirk-20	5.33 (0.75)***	
Inspect-1		-0.40 (0.09)***
Inspect-20		-4.51 (0.86)***
Cons	-2.50 (0.32)***	1.94 (0.38)***
N.	9000	9000

Logit regression with individual fixed effects; robust standard errors using session level clusters in parentheses; shirk-1 is a dummy indicating whether the opponent shirked in the previous round; shirk-20 is the percentage of shirking in the last 20 rounds for all 5 workers in the session; inspect-1 is a dummy indicating whether the opponent inspected in the previous round; inspect-20 is the percentage of inspections in the last 20 rounds for all 5 employers in the session. * p<0.05, ** p<0.01, *** p<0.001

5. DISCUSSION OF THE RELATED LITERATURES

We compare the effectiveness of bonuses and fines as instruments for encouraging compliance in inspection games. In these games the unique Nash equilibrium is in mixed strategies with

positive probabilities of inspection and compliance. We find that bonuses targeted at those inspected and found complying are not effective in encouraging compliance: in fact, subjects in our experiment comply slightly less when bonuses are present, although the effect is not statistically significant. On the other hand, we find that introducing harsher fines for non-compliance is an effective tool for encouraging compliance.

Our paper contributes to an extensive literature on deterrence. The seminal papers of Becker (1968) and Stigler (1970) started an empirical and theoretical literature on how criminal behavior responds to incentives. The theoretical literature showed how crimes and deterrence can be studied in a market for offenses. In the approach of Ehrlich and Becker (1972) and Ehrlich (1996), the aggregate supply of offenses is based on factors affecting individuals' decision to pursue an illegal activity, such as the payoff per offense, the direct costs incurred by the offender in pursuing the payoff, the alternative wage rate in a legitimate activity, the probability of apprehension and conviction, and the prospective penalty if convicted. The supply of offenses increases with the overall net return of the illegal activity. The demand side of the market is shaped by potential victims' activities to prevent crime by employing safety measures such as burglar alarm systems, paying higher rents for living in safer neighborhoods, or hiring private guards. In the Ehrlich-Becker approach potential victims' demand-for-protection schedule increases with the crime rate. The 'demanded' (tolerated) crime by potential victims is then inversely related to the demand-for-protection schedule. Thus, the 'demanded' crime decreases with the crime rate. The equilibrium market crime rate is located at the intersection of the supply and demand for criminal activities.

In Ehrlich's framework public measures such as an increase in the probability of apprehension or an increase in the sanction of the criminal if convicted effectively diminish the equilibrium rate of crime. Likewise, a bonus for good behavior would induce fewer potential offenders to pursue a given criminal activity, and diminish the equilibrium rate of crime.

One theoretical difference between the Ehrlich-Becker approach and the game theoretic approach is that the former assumes a large multitude of individuals in the market for offenses, both on the supply and demand side, whereas in the latter there are effectively only two parties in the interaction. Another important difference is that in the market model, the introduction of bonuses or fines leads to changes in compliance because they change the relative returns from criminal versus legitimate activities. In the game theoretic approach monitoring and compliance are endogenous, and the introduction of bonuses and fines affects both sides of the interaction.

The perverse effect of bonuses follows because the potential offender anticipates that inspection has become less attractive to the authority since it may involve the payment of the bonus.

Clearly, some applications are better described by the market approach while others are more closely approximated by the game theoretic approach. For instance, the case where home owners protect their property with burglar alarms while petty thieves consider their chances to earn a quick buck is better described in the market approach, while the interaction between a tax authority and citizens deciding whether or not to declare their savings is better described by the game theoretic approach.¹³

The empirical literature on deterrence is surveyed in Freeman (1983), Cameron (1988) and Ehrlich (1996). Ehrlich concludes that, taken as a whole, these studies offer massive evidence for the hypothesis that incentives have a deterrent effect on crime, even though a minority of the studies fails to identify an effect. The empirical evidence of whether negative or positive incentives are more effective in encouraging compliance is less conclusive. Freeman (1983) concludes that the deterrent effects of positive measures, such as higher legitimate earnings possibilities, rehabilitation measures and a lower disparity in the in the distribution of income in society, are weaker than the negative deterrent effects of both the probability and the severity of sanctions. Ehrlich (1996) is less convinced by the decisiveness of the evidence and points to some econometric problems with existing studies that prevent a definite conclusion. For instance, in the empirical studies positive and negative incentives are usually correlated, the causal direction between incentives and crime rates remains uncertain, and there are difficulties in measuring relevant variables such as the occurrence and returns from criminal activity.

Experimental methods sidestep some of these measurement difficulties, and so present a useful framework for studying the effectiveness of incentives. Several recent experimental studies have examined whether monetary incentives deter criminal behavior in a variety of settings (e.g., Abbink et al., 2002; Abbink, 2006; Cameron et al., 2009; Schildberg-Hörisch and Strassmair, *forthcoming*; Serra, *forthcoming*). This literature typically focuses on the impact of sanctions on criminal behavior, whereas we compare the effectiveness of rewards and sanctions.

One strand of research that has investigated both positive and negative incentives focuses on how individuals respond to different reward/penalty schemes in a tax reporting setting (see,

¹³ The fact that the tax authority is involved in many separate inspection games with multiple citizens does not change the game theoretic nature of the game.

for example, Alm et al., 1992; Bazart and Pickhardt, 2011; Kastlunger et al., 2011). Similarly, a literature on public goods provision has studied how players respond to rewards/punishments that are automatically assigned to them depending on how their contributions compare with others (e.g., Falkinger et al., 2000; Dickinson, 2001; Andreoni and Gee, 2011a; 2011b).

In this research, like ours, the reward/penalty schemes are fixed by the experimenter and interest is directed at how individuals respond to these. Our research differs from this research on tax compliance and voluntary contributions to public goods in that the authority is an *active player* in our setup. This has the implication that an exogenous change in reward/penalty schemes affects monitoring behavior. It is the interaction of effects on monitoring/compliance behavior that determines the ultimate effect of a change in penalty/reward scheme.

Another strand of literature has used two-stage games where in the second stage, after having observed choices made in the first stage, players can incur costs to punish or reward other players. Players are not predicted to use costly rewards or punishments if they are solely concerned about own earnings, but they might if they have preferences for reciprocity. In fact, a large experimental literature documents the willingness of some people to eschew private interests and react positively toward those that treat them well (positive reciprocity) or negatively toward those that treat them poorly (negative reciprocity). In particular, early studies of games that allow for both positive and negative reciprocity found that the latter has a particularly strong impact (Abbink et al., 2000; Offerman, 2002; Charness and Rabin, 2002). These findings are echoed in Andreoni et al. (2003) who investigate the effects of rewards and punishments in a proposer-responder game where the proposer chooses an amount to transfer to the responder and the responder can then either punish or reward the proposer. They find that proposers' transfers are particularly sensitive to the threat of punishment, although rewards have also positive effects. Similarly, Nosenzo and Sefton (2012) review the literature on rewards and punishments in public good games, and conclude that punishments are at least as effective as rewards.¹⁴

¹⁴ The relative effectiveness of the instruments depends on the rewarding/punishing technology. Both high-power rewards and punishments (where the benefit/cost of receiving reward/punishment is three times larger than the cost of delivering it) are found to be effective (Rand et al., 2009; Sutter et al., 2010). The evidence is more mixed with low-power incentives (where the benefit/cost of receiving reward/punishment is equal to the cost of delivering it), but punishments are generally found to be more effective than rewards (Walker and Halloran, 2004; Sefton et al., 2007; Sutter et al., 2010; Nosenzo and Sefton, 2012). Gülerk et al. (2006) study a public good game with low-power rewards and high-power punishments. They find that only the latter have an impact on contributions. Gülerk et al. (2009) use a public goods game where one group member (the 'leader') can use high-power rewards or punishments to encourage contributions, and find that contributions are higher when punishments are used.

Our research differs from these studies in that we do not study *discretionary*, or informal, rewards and punishments, but we rather focus on formal bonuses and fines that are *automatically* triggered after specific combinations of actions chosen by the players.¹⁵ Moreover, we study bonuses and fines that are pure transfers from one party to another, and so have no direct efficiency implications. Thus, bonuses or fines can only enhance performance to the extent that they succeed in inducing behavior that is more aligned with the group interest. Finally, unlike previous research on the effect of rewards/punishments in social dilemmas, in our game standard theory predicts that bonuses and fines will affect performance.

A striking result of our study is that monetary incentives do not always have the desired effect. This is somewhat reminiscent of findings from a literature studying the potential pitfalls, or ‘crowding-out’ effects, of monetary incentives (see Bowles and Polania-Reyes, 2012 for a review). For example, Gneezy and Rustichini (2000) find that the introduction of a small fine for parents being late when picking up their children at day-care centers results in *more* late arrivals. Bohnet et al. (2001) examine a contractual relationship in which a first mover decides whether or not to enter a relation with a second mover under various levels of contract enforceability. In a dynamic model of preference adaptation, intrinsic trustworthiness is crowded out with ‘medium’ enforcement probabilities. They report experimental support for their model. In an experiment, Fehr and Rockenbach (2003) find that sanctions revealing selfish intentions destroy altruistic cooperation whereas sanctions perceived as fair leave altruism alive. In a situation with asymmetric information, Benabou and Tirole (2003) show how an informed principal (e.g., a parent or teacher) may discourage an uninformed agent (e.g., a child or student) who infers from the use of the incentive scheme negative news about her own ability or the importance of the task at hand. In a model where agents are motivated by intrinsic, extrinsic and reputational concerns, Benabou and Tirole (2006) show that the presence of an extrinsic incentive may spoil an individual’s reputational motivation to perform a good deed to convey a positive self-image.

¹⁵ There is a small literature that investigates the attractiveness of contracts that make use of *automatic* probabilistic fines in comparison to contracts that make use of *discretionary* bonuses. For instance, Fehr et al. (2007) show that principals prefer the use of a bonus contract that offers the agent a voluntary and unenforceable bonus for satisfactory performance over an incentive contract that offers an automatic probabilistic punishment in case of poor performance as well as a trust contract that pays a generous wage up front. The principals’ choices make sense because the bonus contract is superior to the other types of contracts. In a follow-up experiment, Fehr and Schmidt (2007) investigate the possibility that a combined contract that uses both a voluntary bonus and a fine dominates the bonus contract. In the experiment, the combined contract does not induce better performance by the agents, and two-thirds of the principals prefer the pure bonus contract.

Differently from this literature the perverse effect of bonuses in our setting does not stem from the interaction of pecuniary and non-pecuniary motivations, but is directly implied by pecuniary incentives. Moreover, the fact that bonuses and fines are automatically triggered rather than result from an intentional and discretionary choice of a player may limit any motivational crowding-out effects in our inspection games.

As far as we are aware there have only been two experimental studies of inspection games. Dorris and Glimcher (2004) observe the behavior of human and monkey subjects in inspection games with different parameterizations of the inspection cost.¹⁶ In some experiments they had humans playing against humans, whereas in others they had humans or monkeys in the role of Worker playing against a computer in the role of Inspector. They find that (human and monkey) Workers' behavior is close to Nash equilibrium predictions only for high inspection costs. Dorris and Glimcher (2004) do not study the impact of bonus or fines in their setup. Rauhut (2009) studies the impact of the severity of the punishment in an inspection game. His set up differs from ours in that the punishment hurts the inspectee but does not affect the payoff of the inspector in any way. A consequence is that an increase in the punishment decreases the probability of inspection but leaves the probability of shirking unaffected in the Nash equilibrium. Nevertheless, he finds that inspectees shirk less often when the punishment is increased, in agreement with the own-payoff effect.¹⁷ Our paper differs from his also in that we study reward as well as punishment.

Finally, our study also contributes to a recent literature evaluating different solution concepts for predicting behavior in games with mixed strategy equilibria (e.g., Selten and Chmura, 2008; Brunner et al., 2011; Selten, et al., 2011). Standard game theoretical analysis applied to the game used in our experiment yields the perhaps paradoxical result that introducing bonuses increases considerably the probability that the employee will shirk. While in our experiment we do observe a slight increase in shirking in the presence of bonuses, this effect is much smaller than predicted by Nash equilibrium and is not statistically significant. This is more in line with the predictions made by alternative concepts such as Impulse Balance Equilibrium and Quantal Response Equilibrium. Interestingly, for our data, the latter concept performs better than Nash equilibrium only if it incorporates loss aversion, like Impulse Balance Equilibrium

¹⁶ See also Glimcher et al. (2005).

¹⁷ In fact, Rauhut studies a game where two inspectors interact with two inspectees who are involved in a prisoners' dilemma. Under some assumptions, this expanded game has the same characteristics as an inspection game.

does. Our conjecture is that loss aversion is important in our setting because the amount a player can unilaterally guarantee offers a salient security payoff, and in the equilibrium of the games both the employer and the worker face the danger of earning less than this.¹⁸ More generally, our results show that when Nash equilibrium and alternative predictions diverge we find more support for the latter than for the former.

6. CONCLUSIONS

In this study we have focused on the case where rewards and punishments are simple transfers between the interacting parties (e.g. monetary fines for noncompliance or bonuses for compliance). This seems to be a useful starting point as the connections between incentives, behavior, and earnings are straightforward to interpret: bonuses and fines have no direct efficiency consequences unless they induce a change in behavior. We find that fines, but not bonuses, enhance efficiency. An interesting extension would be to settings where the costs and benefits of rewarding/being rewarded are asymmetric (e.g., when bonuses consist of medals and prizes that may have more value for the person receiving them than for the person awarding them). If the bonus remains equally costly to the inspector while it becomes more beneficial to the inspectee, our results suggest that the inspectee will shirk less often because of the enhanced own-payoff effect of working. Thus, in such a setup bonuses may have a positive effect on compliance. Also, in this study we examine the performance of exogenously imposed mechanisms. In our experiment, workers chose whether to work or shirk and employers chose whether to inspect or not inspect. Fines and bonuses were then triggered automatically in response to the actions chosen by the players. Another interesting avenue for further research would be to explore the effect of discretionary fines and bonuses and the endogenous choice of punishing and rewarding mechanisms.¹⁹

¹⁸ Note that our game was parameterized so that subjects could not make real losses. The possibility of making real losses may even accentuate the importance of loss aversion. For example, in the domain of gift exchange games, Dittrich and Zieglmeyer (2006) and Gose and Sadrieh (2011) show that principals offer lower wages when losses are possible. Notwithstanding this, principals' behavior depends on such factors as how much they can lose and how large the efficiency gains are even when losses are covered out of an initial endowment - see, for instance, Engelmann and Ortmann, 2009.

¹⁹ In recent follow-up work we examine a related setting where the use of incentives is discretionary and the employer can choose whether to punish or reward an employee after an inspection. See Nosenzo et al. (2012) for details. See also Gürer et al. (2006) and Sutter et al. (2010) for recent studies of the endogenous choice of punishing and rewarding mechanisms for encouraging voluntary contributions to public goods.

REFERENCES

- Abbink, K. 2006. Laboratory experiments on corruption. In *International Handbook on the Economics of Corruption*, ed. Susan Rose-Ackerman. Cheltenham, UK: Edward Elgar.
- Abbink, K., B. Irlenbusch, and E. Renner. 2000. The moonlighting game - An experimental study on reciprocity and retribution. *Journal of Economic Behavior & Organization* 42(2), 265–277.
- Abbink, K., B. Irlenbusch, and E. Renner. 2002. An Experimental Bribery Game. *Journal of Law, Economics, and Organization* 18(2), 428–454.
- Alm, J., B.R. Jackson, and M. McKee. 1992. Deterrence and Beyond: Toward a Kinder, Gentler IRS. In *Why People pay Taxes*, ed. Joel Slemrod, 311–29. Ann Arbor: University of Michigan Press.
- Andreoni, J., and L. Gee. 2011a. The Hired Gun Mechanism. NBER Working Paper No. 17032.
- Andreoni, J., and L. Gee. 2011b. Gun For Hire: Does Delegated Enforcement Crowd out Peer Punishment in Giving to Public Goods? NBER Working Paper No. 17033.
- Andreoni, J., W. Harbaugh, and L. Vesterlund. 2003. The carrot or the stick: Rewards, punishments, and cooperation. *American Economic Review* 93(3), 893–902.
- Andreozzi, L. 2004. Rewarding Policemen Increases Crime. Another Surprising Result from the Inspection Game. *Public Choice* 121(1-2), 69–82.
- Avenhaus, R., B. Von Stengel, and S. Zamir. 2002. Inspection Games. In *Handbook of Game Theory*, ed. R. Aumann and S. Hart, 3:1947–87. Amsterdam: North-Holland.
- Bazart, C., and M. Pickhardt. 2011. Fighting Income Tax Evasion with Positive Rewards. *Public Finance Review* 39(1), 124–149.
- Becker, G.S. 1968. Crime and Punishment - An Economic Approach. *Journal of Political Economy* 76(2), 169–217.
- Benabou, R., and J. Tirole. 2003. Intrinsic and extrinsic motivation. *Review of Economic Studies* 70(3), 489–520.
- Benabou, R., and J. Tirole. 2006. Incentives and prosocial behavior. *American Economic Review* 96(5), 1652–1678.
- Bianco, W.T., P.C. Ordeshook, and G. Tsebelis. 1990. Crime and Punishment: Are One-Shot, Two-Person Games Enough? *The American Political Science Review* 84(2), 569–586.
- Blackwell, C. 2007. A Meta-Analysis of Tax Compliance Experiments. International Studies Program Working Paper N. 0724.
- Bohnet, I., B. Frey, and S. Huck. 2001. More Order with Less Law: On Contract Enforcement, Trust, and Crowding. *The American Political Science Review* 95(1), 131–144.
- Bowles, S., and S. Polania-Reyes. 2012. Economic Incentives and Social Preferences: Substitutes or Complements? *Journal of Economic Literature* 50(2), 368–425.
- Brunner, C., C. Camerer, and J.K. Goeree. 2011. A Correction and Re-Examination of ‘Stationary Concepts for Experimental 2 x 2 Games’. *American Economic Review* 101(2), 1029–40.
- Camerer, C.F.F., T.H. Ho, and J.K. Chong. 2004. A cognitive hierarchy model of games. *Quarterly Journal of Economics* 119(3), 861–898.
- Cameron, L., A. Chaudhuri, N. Erkal, and L. Gangadharan. 2009. Propensities to engage in and punish corrupt behavior: Experimental evidence from Australia, India, Indonesia and Singapore. *Journal of Public Economics* 93(7-8), 843–851.
- Cameron, S. 1988. The Economics of Crime Deterrence: A Survey of Theory and Evidence. *Kyklos* 41(2), 301–323.

- Charness, G., and M. Rabin. 2002. Understanding social preferences with simple tests. *Quarterly Journal of Economics* 117(3), 817–69.
- Cox, G.W. 1994. A note on crime and punishment. *Public Choice* 78, 115–124.
- Dickinson, D.L. 2001. The carrot vs. the stick in work team motivation. *Experimental Economics* 4(1), 107–124.
- Dittrich, D.A.V., and A. Ziegelmeyer. 2006. Laboratory bilateral gift exchange: The impact of loss aversion. Papers on strategic interaction, Max Planck Institute of Economics, Jena.
- Dorris, M.C., and P.W. Glimcher. 2004. Activity in Posterior Parietal Cortex Is Correlated with the Relative Subjective Desirability of Action. *Neuron* 44(2), 365–378.
- Ehrlich, I. 1996. Crime, Punishment, and the Market for Offenses. *The Journal of Economic Perspectives* 10(1), 43–67.
- Ehrlich, I., and G.S. Becker. 1972. Market Insurance, Self-Insurance, and Self-Protection. *Journal of Political Economy* 80(4), 623–648.
- Engelmann, D., and A. Ortmann. 2009. The Robustness of Laboratory Gift Exchange: A Reconsideration.” Mimeo. Working Paper available at <http://engelmann.vwl.uni-mannheim.de/2570.0.html>.
- Falkinger, J., E. Fehr, S. Gächter, and R. Winter-Ebmer. 2000. A simple mechanism for the efficient provision of public goods: Experimental evidence. *American Economic Review* 90(1), 247–264.
- Fehr, E., A. Klein, and K.M. Schmidt. 2007. Fairness and contract design. *Econometrica* 75(1), 121–154.
- Fehr, E., and B. Rockenbach. 2003. Detrimental effects of sanctions on human altruism. *Nature* 422(6928) (March 13), 137–140.
- Fehr, E., and K.M. Schmidt. 2007. Adding a Stick to the Carrot? The Interaction of Bonuses and Fines. *The American Economic Review* 97(2), 177–181.
- Freeman, R. 1983. Crime and Unemployment. In *Crime and Public Policy*, ed. J.Q. Wilson, 89–106. San Francisco: Institute for Contemporary Studies.
- Fudenberg, D., and J. Tirole. 1991. *Game Theory*. Cambridge, MA: MIT Press.
- Glimcher, P.W., M.C. Dorris, and H.M. Bayer. 2005. Physiological utility theory and the neuroeconomics of choice. *Games and Economic Behavior* 52(2), 213–256.
- Gneezy, U., and A. Rustichini. 2000. A fine is a price. *Journal of Legal Studies* 29(1), 1–17.
- Goeree, J.K., and C.A. Holt. 2001. Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review* 91(5), 1402–1422.
- Goeree, J.K., C.A. Holt, and T.R. Palfrey. 2003. Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior* 45(1), 97–113.
- Gose, K., and A. Sadrieh. 2011. Fair Wages When Employers Face the Risk of Losing Money. FEMM Working Paper 2011_09.
- Greenhouse, S. 2005. Day Laborer Battle Runs Outside Home Depot. *The New York Times*, October 10.
- Greiner, B. 2004. An Online Recruitment System for Economic Experiments. In *Forschung und wissenschaftliches Rechnen GWDG Bericht 63*, 79–93. Göttingen: Gesellschaft für Wissenschaftliche Datenverarbeitung.
- Gürerk, Ö., B. Irlenbusch, and B. Rockenbach. 2006. The competitive advantage of sanctioning institutions. *Science* 312(108), 108–111.
- Gürerk, Ö., B. Irlenbusch, and B. Rockenbach. 2009. Motivating teammates: The leader’s choice between positive and negative incentives. *Journal of Economic Psychology* 30(4), 591–607.

- Hanson, J., and M. Bell. 2007. Harvest Trails in Australia: Patterns of Seasonal Migration in the Fruit and Vegetable Industry. *Journal of Rural Studies* 23(1), 101–117.
- Kastlunger, B., S. Muehlbacher, E. Kirchler, and L. Mittone. 2011. What Goes Around Comes Around? Experimental Evidence of the Effect of Rewards on Tax Compliance. *Public Finance Review* 39(1), 150–167.
- Nagel, R. 1995. Unraveling in the guessing games: An experimental study. *American Economic Review* 85(5), 1313–1326.
- Nosenzo, D., T. Offerman, M. Sefton, and A. van der Veen. 2012. Discretionary Sanctions and Rewards in the Repeated Inspection Game. CeDEx Discussion Paper 2012-10, University of Nottingham.
- Nosenzo, D., and M. Sefton. 2012. Promoting Cooperation: The distribution of reward and punishment power. CeDEx Discussion Paper 2012-08, University of Nottingham.
- Ochs, J. 1995. Games with Unique, Mixed Strategy Equilibria: An Experimental Study. *Games and Economic Behavior* 10, 202–217.
- Offerman, T. 2002. Hurting hurts more than helping helps. *European Economic Review* 46(8), 1423–1437.
- Rand, D.G., A. Dreber, T. Ellingsen, D. Fudenberg, and M.A. Nowak. 2009. Positive Interactions Promote Public Cooperation. *Science* 325(5945), 1272–1275.
- Rauhut, H. 2009. Higher Punishment, Less Control?: Experimental Evidence On the Inspection Game. *Rationality and Society* 21(3), 359–392.
- Schildberg-Hörisch, H., and C. Strassmair. An Experimental Test of the Deterrence Hypothesis. *forthcoming in Journal of Law, Economics, and Organization DOI:10.1093/jleo/ewq015*.
- Sefton, M., R. Shupp, and J.M. Walker. 2007. The effect of rewards and sanctions in provision of public goods. *Economic Inquiry* 45(4), 671–690.
- Selten, R., and T. Chmura. 2008. Stationary Concepts for Experimental 2x2-Games. *American Economic Review* 98(3), 938–66.
- Selten, R., T. Chmura, and S. Goerg. 2011. Correction and Re-Examination of Stationary Concepts for Experimental 2x2 Games: A Reply. *American Economic Review* 101(2), 1041–44.
- Serra, D. Combining Top-Down and Bottom-Up Accountability: Evidence from a Bribery Experiment. *forthcoming in Journal of Law, Economics, and Organization DOI: 10.1093/jleo/ewr010*.
- Stahl, D.O., and P.W. Wilson. 1995. On Players' Models of Other Players: Theory and Experimental Evidence. *Games and Economic Behavior* 10(1), 218–254.
- Stigler, G.J. 1970. The Optimum Enforcement of Laws. *Journal of Political Economy* 78(3), 526–536.
- Sutter, M., S. Haigner, and M.G. Kocher. 2010. Choosing the Carrot or the Stick? Endogenous Institutional Choice in Social Dilemma Situations. *Review of Economic Studies* 77(4), 1540–1566.
- Tsebelis, G. 1989. The Abuse of Probability In Political Analysis: The Robinson Crusoe Fallacy. *The American Political Science Review* 83(1), 77–91.
- Walker, J.M., and M.A. Halloran. 2004. Rewards and sanctions and the provision of public goods in one-shot settings. *Experimental Economics* 7(3), 235–247.
- Wan, J. 2010. The Incentive to Declare Taxes and Tax Revenue: The Lottery Receipt Experiment in China. *Review of Development Economics* 14(3), 611–624.

APPENDIX A

In this appendix, we explain the procedure to derive the equilibrium predictions of Impulse Balance Equilibrium (IBE) and Quantal Response Equilibrium (QRE) in the context of the Baseline Inspection Game. We also discuss how we derived predictions for QRE and Nash equilibrium augmented with loss aversion.

In IBE players judge the payoffs according to how they relate to their security level. A player's security level s is determined by the player's pure maximin payoff, the maximum of the minimum payoffs corresponding to the player's actions. The left panel of Figure A.1 presents the baseline inspection game that we actually employed in the experiment, in which the inspector can secure a payoff of 12 and the worker a payoff of 25. The payoff matrix is then transformed to account for loss aversion in the following way. From each payoff exceeding a player's security level half the difference between the payoff and the security level is subtracted (the other payoffs remain unchanged). Or, each payoff x is replaced by $x - \max\{\frac{1}{2}(x - s), 0\}$. As a consequence, losses compared to the reference point weigh twice the amount that gains weigh. The middle panel of Figure A.1 presents the Transformed inspection game. From the Transformed game, the Impulse matrix is derived with the following procedure. Each set of two payoffs of a player corresponding to the same action of the other player is transformed such that the highest payoff becomes 0 and the lowest becomes the difference between the highest and the lowest. The resulting numbers represent the impulses to choose the other action given the action chosen by the other player. The impulse matrix is presented in the right panel of Figure A.1.

Figure A.1 Baseline game, Transformed game and Impulse matrix

		Baseline Inspection Game		Transformed Inspection Game		Impulse Matrix	
		W	S	W	S	W	S
I	I	52	12	32	12	4	0
	N	25	20	25	20	0	5
N	I	60	0	36	0	0	12
	N	25	40	25	32.5	7.5	0

In the IBE, a player's expected impulse from one action to the other equals the expected impulse from the other action to the one action. Let p represent the probability that the employer chooses I, and q the probability that the worker chooses S, then p and q follow from the solution of the impulse balance equations:

$$4p(1-q) = 12(1-p)q, \quad 7.5(1-p)(1-q) = 5pq.$$

In QRE, players maximize expected utility taking the actual response function of the other player into account, but make mistakes. Let $E_{\text{player}}[a]$ represent a player's expected utility from choosing action a , then:

$$p = \frac{e^{\lambda E_{\text{employer}}[I]}}{e^{\lambda E_{\text{employer}}[I]} + e^{\lambda E_{\text{employer}}[NI]}}, \quad q = \frac{e^{\lambda E_{\text{worker}}[S]}}{e^{\lambda E_{\text{worker}}[S]} + e^{\lambda E_{\text{worker}}[W]}}$$

where λ represents the player's rationality parameter that is estimated from the data.

The QRE prediction for the games without loss aversion is found using the ordinary payoffs matrices, e.g., for the Baseline Inspection Game the payoffs listed in the left panel of Figure A.1. For QRE with loss aversion, the payoffs of the Transformed games are used. For example, for the Baseline Inspection Game, p and q follow from the solution of:

$$p = \frac{e^{\lambda[32(1-q)+12q]}}{e^{\lambda[32(1-q)+12q]} + e^{\lambda[36(1-q)]}}, \quad q = \frac{e^{\lambda[25]}}{e^{\lambda[25]} + e^{\lambda[20p+32.5(1-p)]}}.$$

The predictions of Nash equilibrium with loss aversion are also calculated using the payoffs of the Transformed games.

APPENDIX B

Instructions

Introduction

This is an experiment about decision-making. In the room, there are ten people who are participating in this experiment. You must not communicate with any other participant in any way during the experiment. At the end of the experiment you will be paid in private and in cash. The amount of money you earn will depend on the decisions that you and the other participants make. The experiment consists of two parts, each part consisting of a number of rounds. In each round you can earn points. At the end of the experiment you will be paid according to the sum of your total point earnings from all rounds in both parts at a rate of 0.4 pence per point. You will receive the instructions for the second part after the first part is finished.

Part One

At the beginning of Part One five of the participants will get the role of "employers" and five will get the role of "workers". You will find out whether you are an employer or worker when the decision-making part of the experiment begins. If you are an employer you will remain an employer throughout the first part, and if you are a worker you will remain a worker throughout the first part.

Part One will consist of 40 rounds. In each round the employers will be paired with the workers. Thus, if you are an employer you will be paired with one of the workers, and if you are a worker you will be paired with one of the employers. The people you are paired with will change randomly from round to round.

At the beginning of a round all participants will make their decisions. Employers must choose either INSPECT or NOT INSPECT. Workers must choose either HIGH effort or LOW effort. At the end of the round, after everyone has made their decision, the computer will inform you of the choices made by you and the person you were paired with and your point earnings for the round.

The number of points you earn in a round will depend on the decisions made by you and the person you are paired with in that round, as described in the tables below:

Employer's point earnings		
	HIGH	LOW
INSPECT	52	12
NOT INSPECT	60	0

Worker's point earnings		
	HIGH	LOW
INSPECT	25	20
NOT INSPECT	25	40

For example, if the employer chooses NOT INSPECT and the worker chooses LOW the employer earns 0 points and the worker earns 40 points.

In addition, on your screen you will see your accumulated point earnings so far, and a table summarizing the decisions made by all participants in previous rounds. The table will be like the one shown below (although the data in the table has been chosen for illustrative purposes only: in the experiment the data will correspond to the actual decisions made by participants).

Results of last 20 rounds			
	HIGH	LOW	Total
INSPECT	10%	20%	30%
NOT INSPECT	30%	40%	70%
Total	40%	60%	100%

For example, the table tells you that the combination (INSPECT, HIGH) occurred in 10% of the cases, that the employers chose INSPECT in 30% of the cases, and the workers chose HIGH in 40% of the cases. The table is based on the results of the most recent 20 rounds only.

To make sure everyone understands the instructions so far, please complete the questions about Part One below. In a couple of minutes someone will come to your desk to check the answers.

1. Will you be matched with the same person from round to round? _____

2. How many points will you earn in a round if you are an employer, choose NOT INSPECT, and the worker you are matched with chooses HIGH? _____

3. How many points will you earn in a round if you are a worker, choose HIGH, and the employer you are matched with chooses NOT INSPECT? _____

4. Is the following statement true: the screen summarizing the history so far always contains information on all previous rounds _____

5. Is the following statement true: the screen summarizing the history so far contains information on the choices of all 10 participants in the room _____

Part Two

In Part Two you will keep the same role as you had in Part One. Again, you will be matched with a different person in the other role in each round. Part Two will consist of an additional 80 rounds, starting with round 41 and ending after round 120. Your decisions together with the decisions of the people that you will be matched with will determine your earnings that will be added to your total earnings in points from Part One. At the beginning of a round, employers must again choose either INSPECT or NOT INSPECT, while workers must choose either HIGH effort or LOW effort. At the end of the round, the computer will inform you of the outcome of the round for you and the person you are paired with.

[*CONTROL*: The point earnings that the employer and worker receive in each of the four cases (INSPECT, HIGH); (INSPECT, LOW); (NOT INSPECT, HIGH); (NOT INSPECT, LOW) will remain exactly the same as in Part One, as shown below.

Employer's point earnings		
	HIGH	LOW
INSPECT	52	12
NOT INSPECT	60	0

Worker's point earnings		
	HIGH	LOW
INSPECT	25	20
NOT INSPECT	25	40

]

[*FINE*: The only difference between Part One and Two will be that the worker will pay a fine of 20 points to the employer when the worker was inspected and chose low effort. So after INSPECT and LOW the employer's point earnings increase by 20 points and the worker's point earnings decrease by 20 points, as shown in the tables below:

Employer's point earnings		
	HIGH	LOW
INSPECT	52	32
NOT INSPECT	60	0

Worker's point earnings		
	HIGH	LOW
INSPECT	25	0
NOT INSPECT	25	40

Thus, if the employer chooses INSPECT and the worker chooses LOW the employer earns 32 points and the worker earns 0 points. In all other cases the payoffs remain the same as in Part One.]

[*BONUS*: The only difference between Part One and Two will be that the employer will give a reward of 20 points to the worker when he or she inspected the worker and found out that the worker chose high effort. So after INSPECT and HIGH the employer's point earnings decrease by 20 points and the worker's point earnings increase by 20 points, as shown in the new earnings tables below:

Employer's point earnings		
	HIGH	LOW
INSPECT	32	12
NOT INSPECT	60	0

Worker's point earnings		
	HIGH	LOW
INSPECT	45	20
NOT INSPECT	25	40

Thus, if the employer chooses INSPECT and the worker chooses HIGH the employer earns 32 points and the worker earns 45 points. In all other cases the payoffs remain the same as in Part One.]

As before, your screen will display your accumulated point earnings (including your earnings from Part One). You will also see a table summarizing the decisions made by all participants in previous rounds. At the start of period 41, this table will be empty. The table will again list the results of the most recent 20 rounds after round 41.

Ending the session

At the end of round 120 your total points from all rounds will be converted to cash at a rate of 0.4 pence per point and you will be paid this amount in private and in cash. Now please begin making your Part Two decisions.