CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The University of Nottingham

Discussion Paper No. 2010-21

Daniele Nosenzo,
Theo Offerman,
Martin Sefton
and
Ailko van der Veen
December 2010

Inducing Good Behavior: Bonuses versus Fines in Inspection Games

CeDEx Discussion Paper Series

ISSN 1749 - 3293

CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The Centre for Decision Research and Experimental Economics was founded in 2000, and is based in the School of Economics at the University of Nottingham.

The focus for the Centre is research into individual and strategic decision-making using a combination of theoretical and experimental methods. On the theory side, members of the Centre investigate individual choice under uncertainty, cooperative and non-cooperative game theory, as well as theories of psychology, bounded rationality and evolutionary game theory. Members of the Centre have applied experimental methods in the fields of public economics, individual choice under risk and uncertainty, strategic interaction, and the performance of auctions, markets and other economic institutions. Much of the Centre's research involves collaborative projects with researchers from other departments in the UK and overseas.

Please visit http://www.nottingham.ac.uk/economics/cedex/ for more information about the Centre or contact

Sue Berry
Centre for Decision Research and Experimental Economics
School of Economics
University of Nottingham
University Park
Nottingham
NG7 2RD
Tel: +44 (0)115 95 15469
Fax: +44 (0) 115 95 14159
sue.berry@nottingham.ac.uk

The full list of CeDEx Discussion Papers is available at

http://www.nottingham.ac.uk/economics/cedex/papers/index.html

Inducing Good Behavior: Bonuses versus Fines in Inspection Games

by

Daniele Nosenzo[*], Theo Offerman[**], Martin Sefton[*] and Ailko van der Veen[**]

10 December 2010

Abstract

We examine the effectiveness of bonuses and fines in an ‘inspection game’ where an employer can learn the effort of a worker through costly inspection. Standard game theoretic analysis predicts that fines discourage shirking, whereas bonuses *encourage* shirking. In contrast, own-payoff effects suggest that both fines and bonuses discourage shirking. In an experiment we find that fines are more effective than bonuses in reducing shirking. However, we do not find that bonuses encourage shirking. Behavioral theories based on Impulse Balance Equilibrium or Quantal Response Equilibrium provide a good account of deviations from Nash equilibrium predictions.

Keywords: Inspection Games; Costly Monitoring; Rewards and Punishments; Bonuses and Fines; Quantal Response Equilibrium; Impulse Balance Equilibrium; Experiment.

JEL Classification Numbers: C70, C72, C92

Acknowledgement: 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 Foundation (ECF/2010/0636).

[*] School of Economics, University of Nottingham, Nottingham, NG7 2RD, United Kingdom

[**] CREED, Department of Economics, University of Amsterdam, 1018 WB, Netherlands

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 misbehaving.

To further encourage compliance, the authority may consider rewarding an individual who was inspected and found complying. For example, 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). Alternatively, the authority may consider increasing the sanctions on individuals who, upon inspection, are found not complying. For example, the Dutch government decided to increase the fine for undeclared savings from 100% to 300% in May 2009.[1] In this paper we study which of these two mechanisms is most successful in promoting good behavior.

The essence of such situations is captured by the 'inspection game', which we describe in Section 2. In this game an authority chooses to inspect or not, and an individual chooses to comply or not, and the unique Nash equilibrium is in mixed strategies, with positive probabilities of inspection and non-compliance. Perhaps unsurprisingly, fines for non-compliant behavior increase the equilibrium probability of compliance. On the other hand, and perhaps paradoxically, bonuses for compliant behavior reduce the equilibrium probability of compliant behavior. Thus, according to standard game theoretical reasoning, fines, and not bonuses, should be used to encourage compliance in such settings.

Previous experiments have revealed limited success of the Nash equilibrium for predicting behavior in games where the equilibrium is in mixed strategies (Ochs, 1995; Potters and van Winden, 1996; Goeree and Holt, 2001; Goeree et al., 2003). One of the reasons why the

[1] See https://zoek.officielebekendmakingen.nl/kst-31301-16.html.

Nash equilibrium does not provide an accurate description of behavior in these types of games is that it fails to capture ‘own-payoff effects’: players do change their behavior in response to changes in their own payoff, whereas the mixed strategy Nash equilibrium predicts that they will not. In the case of the inspection game, the own-payoff effect of introducing fines reinforces the theoretically expected effect: fines make non-compliance less attractive to the individual, and so the own-payoff effect points toward more compliance. However, the own-payoff effect of introducing bonuses for compliant behavior *reduces* the probability of non-compliance. Thus, Nash equilibrium and own-payoff effects point in different directions in this case, and so it is unclear whether the theoretical prediction that fines outperform bonuses in encouraging compliance will be supported in practice.

We describe our experiment for comparing the effectiveness of bonuses and fines in Section 3. Our inspection game is framed as an employer-worker scenario where an employer can either inspect or not and a worker can either supply high or low effort. 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 supplying low effort 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 supplied high effort. 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 supplied low effort. Finally, in the second part of the CONTROL treatment, subjects continued playing the same game as in the first part. This design allows us to examine whether bonuses or fines are more effective in encouraging working/discouraging shirking. 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 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. In general, standard comparative static predictions work well when own-payoff effects point in the same direction, but not otherwise. We show that observed deviations from Nash equilibrium predictions can be explained quite well by behavioral theories that incorporate loss aversion and can accommodate own payoff effects: Impulse Balance Equilibrium (Selten and Chmura, 2008) and an augmented version of Quantal Response Equilibrium (McKelvey and Palfrey, 1995). In Section 5 we discuss these results in relation to the existing literature and conclude.

2. INSPECTION GAMES

We study a simple simultaneous move inspection game. An employer can either inspect (I) or not inspect (N), and a worker can supply either high (H) or low (L) effort. In the canonical version of the game (see, e.g., Fudenberg and Tirole, 1992, p. 17) the employer incurs a cost of h from inspecting, and high effort 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 supplies low effort 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 supplies high effort and the employer does not inspect.

Figure 1. Inspection Games*

Canonical Inspection Game

	H	L
I	$v - w - h$ $w - c$	$-h$ 0
N	$v - w$ $w - c$	$-w$ w

Inspection Game with Fines

	H	L
I	$v - w - h$ $w - c$	$\mathbf{f} - h$ $\mathbf{-f}$
N	$v - w$ $w - c$	$-w$ w

Inspection Game with Bonuses

	H	L
I	$v - w \mathbf{- b} - h$ $w \mathbf{+ b} - c$	$-h$ 0
N	$v - w$ $w - c$	$-w$ 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.

Letting p denote the probability of inspection and q denote the probability of supplying low effort ('shirking'), the employer's expected net benefit from inspecting, that is the expected payoff from inspect minus the expected payoff from not inspect, is $qw - h$. This is the expected saving from not having to pay the wage to a shirker less the inspection cost. The worker's expected net benefit from shirking is $c - pw$; this is the saved effort cost less the expected loss in wages if the employer inspects. The probabilities p and q are determined endogenously and in equilibrium must leave the players indifferent between actions, i.e. must result in expected net benefits of zero. Thus, in the unique Nash equilibrium of the game the employer inspects with probability $p_c = c/w$ and the worker chooses low with probability $q_c = h/w$. The employer receives an expected payoff of $\pi_c^{employer} = v - w - hv/w$, the worker receives an expected payoff of $\pi_c^{worker} = w - c$, and joint payoffs are $\pi_c = v - c - hv/w$.

We now compare two possibilities for encouraging high effort relative to the canonical version of the game: imposing an additional fine on workers caught supplying low effort, versus paying a bonus to workers who are inspected and found supplying high effort.

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. The employer's expected net benefit from inspecting is now $q(w + f) - h$ and the worker's expected net benefit from shirking is $c - p(w + f)$. The unique Nash equilibrium of the game has the employer inspect with probability $p_f = c/(w + f)$ and the worker shirk with probability $q_f = h/(w + f)$. The imposition of the fine increases the employer's expected payoff to $\pi_f^{employer} = v - w - hv/(w + f)$, but does not change the worker's expected payoff, who receives $\pi_f^{worker} = w - c$ as in the canonical version of the game. Thus, joint payoffs are increased to $\pi_f = v - c - hv/(w + f)$. Fines enhance efficiency because joint payoffs are reduced by low effort and/or inspection, 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 chosen high effort. The payoff matrix for this game is shown in the rightmost panel of Figure 1. The employer's expected net benefit from inspecting is now $qw - (1 - q)b - h$ and the worker's expected net benefit from shirking is $c - p(w + b)$. Now in equilibrium the employer inspects with probability $p_b = c/(w + b)$ and thus bonuses reduce the probability of inspection. However, in contrast to fines which reduce the probability of shirking, bonuses *raise* the probability of shirking to $q_b = (h + b)/(w + b)$. This results in counter-acting effects on joint payoffs. Relative to the equilibrium payoffs from the canonical version of the game, the worker's expected payoff is increased to $\pi_b^{worker} = w - c + cb/(w + b)$ and the inspector's expected payoff is reduced to $\pi_b^{employer} = v - w - v\,(h + b)/(w + b)$. Joint payoffs are

$\pi_b = (v - c - hv/w)\ w/(w + b) < \pi_c$. Thus, bonuses reduce efficiency because the beneficial effect of less frequent inspection is outweighed by the detrimental effect of increased shirking.

The comparative static predictions made by standard game theory follow from the counterintuitive property of the mixed-strategy Nash equilibrium whereby players' *own* decision probabilities are determined by the payoff differences of the *other* player, and not by *own* payoff differences. Figure 2 illustrates this by showing how the expected net benefit of inspecting/shirking varies across games.

Figure 2. Incentives to Shirk and Inspect*

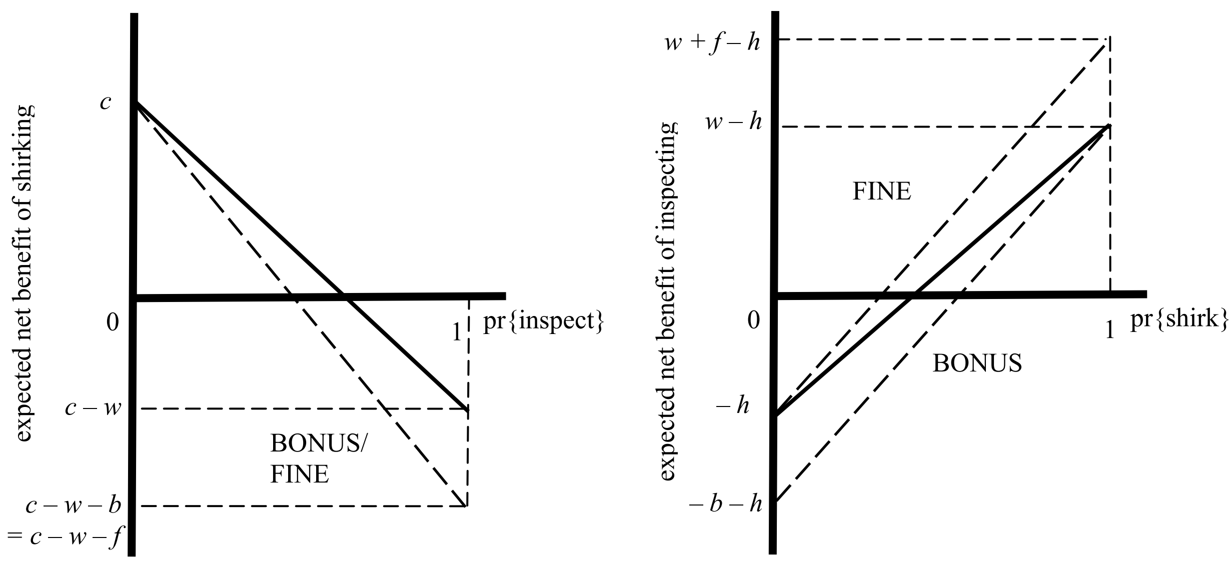

*Expected net benefit for canonical game shown as solid line, for other games shown as dashed line (assuming $b = f$).

The left panel shows the worker's expected net benefit from shirking. Bonuses and fines have similar effects on the worker's payoff: for any given inspection probability the expected benefit of shirking is decreased. The right panel shows the employer' expected net benefit from inspecting. Here, for any given shirking probability fines increase the expected benefit of inspecting while bonuses decrease it. While the left panel shows that fines and bonuses decrease the expected benefit of shirking, the effect on the equilibrium probability of shirking is seen in the *right* panel as the probability of shirking that gives the employers a zero expected net benefit from inspecting. Fines are predicted to decrease shirking relative to the canonical version of the game because of the impact they have on the employer's payoff: since the possibility of imposing a fine on shirkers makes inspections more attractive for the employer, in equilibrium the worker has to shirk less often relative to the canonical game in order to keep the employer indifferent between inspecting and not inspecting. Similarly, the perverse effect of bonuses on shirking follows from the decreased attractiveness of inspections in the game with bonuses relative to the canonical game: since the employer has now to pay a bonus to a worker who is found supplying high effort, a worker has to shirk *more* often to keep the employer indifferent between her two options. In the same way, the changes in the frequency of inspections in the inspection games with fines and bonuses relative to the canonical version of the game are a consequence of the impact of fines and bonuses on the worker's payoff.

Previous experimental work (e.g., Ochs, 1995; Goeree and Holt, 2001; Goeree et al., 2003) shows that these 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 bonuses into the inspection game generates an own-payoff effect that pulls workers' behavior in the opposite direction to Nash equilibrium predictions: shirking is less attractive when bonuses are paid to workers who supply high effort, and hence the supply of high effort is encouraged. On the other hand, introducing fines generates an own-payoff effect that pulls workers' behavior in the same direction as Nash equilibrium: imposing a fine on shirkers makes shirking less attractive and hence encourages the supply of high effort. 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.

In summary, given the evidence on the importance of own-payoff effects in previous experiments, it is not clear that experimental evidence will support the standard game theoretical analysis outlined above. In particular, the own-payoff effects arising when bonuses are paid to workers who are inspected and found supplying high effort may make them a more effective tool for encouraging effort than suggested by standard theory.

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.[2] 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. 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 canonical inspection game shown in the rightmost panel of Figure 3.[3] 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.

Figure 3. Parameterization of the Inspection Games Used in the Experiment*

Canonical Inspection Game

	H	L
I	52 25	12 20
N	60 25	0 40

Inspection Game with Fines

	H	L
I	52 25	**32** **0**
N	60 25	0 40

Inspection Game with Bonuses

	H	L
I	**32** **45**	12 20
N	60 25	0 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] Subjects were recruited through the online recruitment system ORSEE (Greiner, 2004). Instructions are available in Appendix A.

[3] Point earnings were derived from the game described in the previous section (see 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.

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 used the same earnings table as in Part One. In our five FINE sessions the earnings table was as in Part One except that the worker would pay a fine of 20 points to the employer if the worker chose low effort and the employer chose to inspect. Thus in Part Two of the experiment subjects in the FINE sessions played the inspection game shown in the middle panel of Figure 3. In our five BONUS sessions the earnings table was as in Part One except that the employer would pay a bonus of 20 points to the worker if the worker chose high effort and the employer chose to inspect (rightmost panel of Figure 3).

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 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

4.1 Inspection and Shirking Probabilities

Figure 4 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 1 reports 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.

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 (Kruskal-Wallis test p-values are 0.37 for shirk and 0.78 for inspect).[4] Averaged across all sessions the observed proportion of shirking decisions is 45% and the observed proportion of inspecting decisions is 78%: both statistics compare favorably with predictions made by Nash equilibrium (40% and 75%, respectively).

---

[4] 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.

Figure 4. Proportions of Shirking (left panel) and Inspecting (right panel) across Treatments*

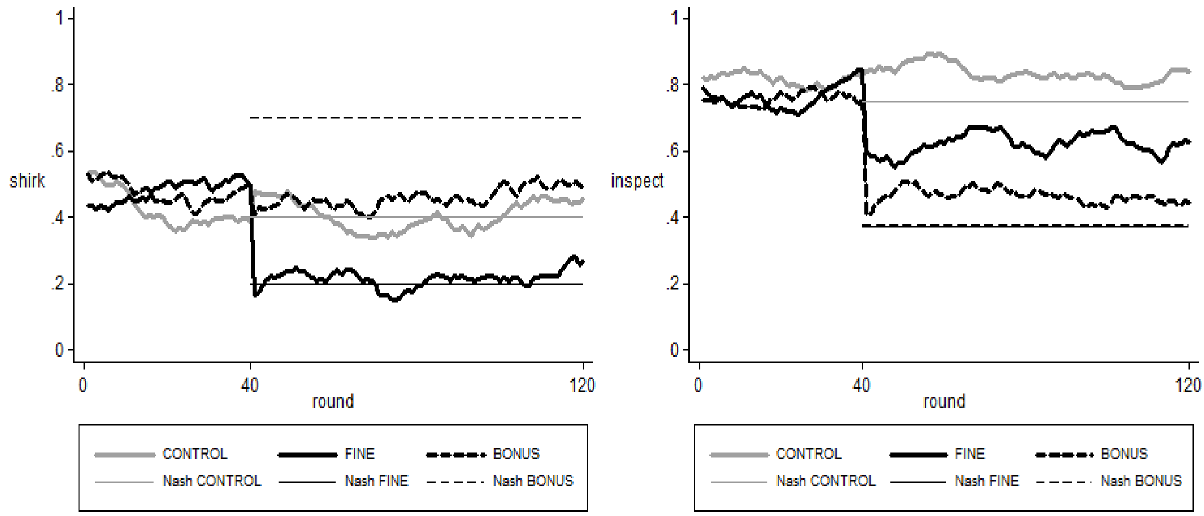

*In each round, the average is displayed of the proportions of (max) 5 previous rounds, the current round and (max) 5 future rounds.

Table 1. Choice Proportions, 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
Nash	*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
Nash	*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.

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% 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 inspection probabilities, these are significantly lower in FINE than CONTROL (p = 0.01) and BONUS than CONTROL (p = 0.01). We also note, however, that the inspection probability 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 and fines have the same effect on inspection probabilities, we find a statistically significant difference in the proportions of inspections between FINE and BONUS (p = 0.01).

4.2 Earnings

Table 2 reports average earnings per game across treatments in the last 20 rounds of Part Two of the experiment. Nash equilibrium predictions are also reported for comparison.

Table 2. Earnings in Part Two, Average by Treatment*

	Part Two		
	CONTROL	FINE	BONUS
Joint Earnings	58.7 (5.75)	69.6 (2.64)	58.9 (2.40)
Nash	*61.0*	*73.0*	*50.5*
Worker Earnings	24.2 (1.08)	22.5 (1.38)	32.7 (1.01)
Nash	*25.0*	*25.0*	*32.5*
Employer Earnings	34.5 (5.11)	47.1 (1.35)	26.1 (2.30)
Nash	*36.0*	*48.0*	*18.0*

* Table shows average point earnings per game (last 20 rounds only). Standard deviations based on session averages in parentheses.

In principle, joint earnings can range from 32 points (when the employer inspects and the worker shirks) to 85 (when the employer does not inspect and the worker works). Theory predicts that joint earnings are equal to 61 points in the game used in CONTROL. In the experiment, earnings in our CONTROL sessions are close to this, averaging 58.7 points across the last 20 rounds of Part Two. Theory also predicts that fines are beneficial and bonuses are detrimental for efficiency. Using Mann-Whitney rank-sum tests, we find that, consistent with these predictions, joint earnings in FINE are higher than in CONTROL, and the difference in the distributions is statistically significant (p = 0.01). On the contrary, we find no evidence that bonuses hamper efficiency: in fact, introducing bonuses slightly *increases* on average joint earnings relative to CONTROL, although the effect is not statistically significant (p = 0.85).

A second aspect of our data is worth discussing: while according to Nash equilibrium the introduction of fines is Pareto improving, as it is predicted to leave the workers' earnings unchanged relative to CONTROL and to increase the employer's payoff, we find that fines are in fact detrimental for workers. In FINE, workers earn about 1.5 points per game less than in CONTROL, and the difference is (weakly) statistically significant (p = 0.06). Fines are instead beneficial for the employer as predicted (p = 0.01). Thus, the introduction of fines has distributive consequences that are not accounted for by standard theory: employers are better off when fines are introduced, but this occurs at the expenses of workers who are worse off relative to CONTROL. The introduction of bonuses has instead the predicted distributive consequences: it significantly increases the worker's payoff and decreases the employer's payoff (p = 0.01 and p = 0.02 respectively).

4.3 Explaining Observed Behavior

Whereas Nash equilibrium predictions seem to capture well the comparative static effects of fines on shirking behavior and bonuses on inspecting behavior, they do not capture observed effects of fines on inspections or bonuses on effort. 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 equilibrium effects.

Table 3 contains predicted choice probabilities made by two alternative concepts: Quantal Response Equilibrium (QRE) and Impulse Balance Equilibrium (IBE).[5] The predictions are for our Part Two data. Part One data is used to estimate the QRE precision parameter λ.[6] For the estimated value of λ QRE predictions are generally close to Nash equilibrium predictions.

Table 3. Predicted Choice Probabilities

	Probability of Shirking			*Probability of Inspecting*		
	CONTROL	FINE	BONUS	CONTROL	FINE	BONUS
Nash	0.40	0.20	0.70	0.75	0.375	0.375
QRE (λ=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
$\text{Nash}^{\text{with loss-aversion}}$	0.25	0.11	0.54	0.60	0.23	0.33
$\text{QRE}^{\text{with loss-aversion}}$ (λ=0.289)	0.42	0.10	0.46	0.69	0.47	0.36

IBE is based on the idea that players look at forgone payoffs when they adjust their decision probabilities: choosing an option that yields a lower payoff than the alternative option generates an 'impulse' in the direction of the non-chosen option. Impulses generated by foregone

[5] Appendix B contains details on the procedures used to derive the equilibrium predictions for IBE and QRE.
[6] When $\lambda = 0$ players choose actions equi-probably and in the limit as λ approaches ∞ players always choose their best-response. As in Selten and Chmura (2008) and Brunner et al. (*forthcoming*), 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. Unlike these authors, who use all data to estimate λ, we estimate the parameter using data from Part One of the experiment and then use this estimated value of λ to predict behavior in Part Two of the experiment.

payoffs that represent a 'loss' relative to a player's security payoff level (her pure strategy maximin value) weigh twice as much as foregone 'gains'. In equilibrium, players choose the decision probabilities such that the impulses of foregone payoffs are equal across options. IBE predictions differ markedly from Nash equilibrium when own payoff and Nash equilibrium effects are in conflict: the IBE predicted probability of shirking in BONUS is 43% (versus the 70% Nash prediction) and the predicted probability of inspecting in FINE is 61% (versus 37.5%).

The fact that Nash equilibrium and QRE are not augmented by loss-aversion while IBE is has generated a recent debate about whether the incorporation of loss-aversion is what drives the observed differences in performance across these equilibrium concepts (see Selten and Chmura, 2008; Brunner et al., *forthcoming*; Goerg et al., *forthcoming*). To examine this possibility, Table 3 also reports predictions made by Nash equilibrium and QRE when these concepts are augmented with loss-aversion.[7] Incorporating loss-aversion into the concepts generally improves the performance of QRE, but not the performance of Nash equilibrium.

Overall, 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. This is summarized in Figure 5. The Figure shows how the introduction of bonuses and fines affect the probability of shirking and inspecting relative to CONTROL according to the three solution concepts, as well as in the data for the last 20 rounds of Part Two.

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

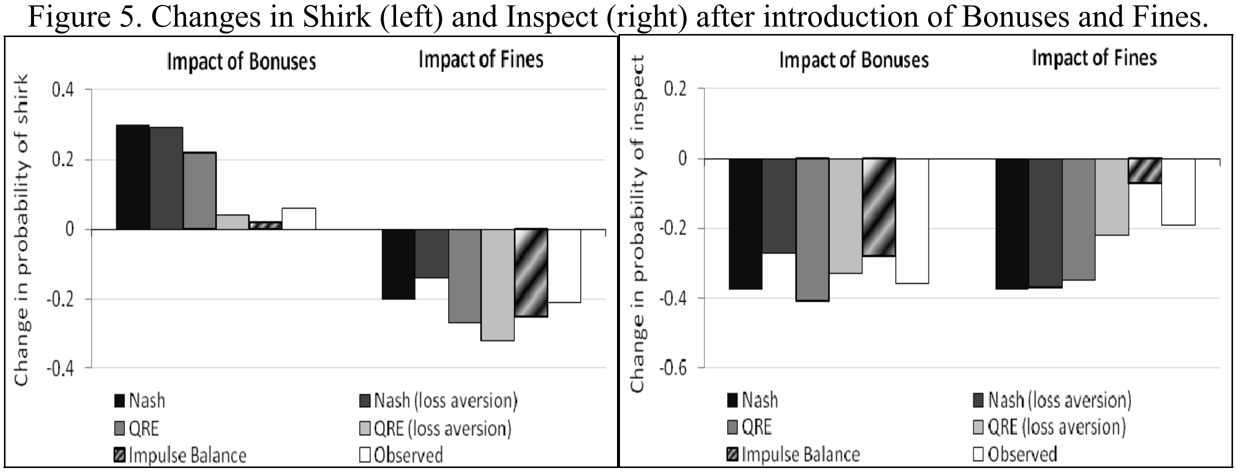

[7] As in Selten and Chmura (2008) we incorporate loss aversion by transforming payoffs above the security level as follows. If x is the payoff and m is the security level, any payoff $x > m$ is transformed into $x' = m + (x-m)/2$. The exact procedure is discussed in Appendix B.

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 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 (about 20% or less).

5. DISCUSSION AND CONCLUSIONS

We compare the effectiveness of bonuses and fines as instruments for encouraging compliance in inspection games. In our setting the incentive for a worker to work is given by the monitoring activity of an employer and the costs/benefits incurred by the worker when she is inspected and found to have worked or shirked. The unique Nash equilibrium of the game is in mixed strategies with positive probabilities of inspection and shirking. We find that bonuses targeted at those inspected and found working are not effective in encouraging working: in fact, subjects in our experiment shirk slightly more often when bonuses are present, although the effect is not statistically significant. On the other hand, we find that introducing harsher fines for shirkers is an effective tool for encouraging working.

The question of whether rewards or punishments are a better tool for inducing socially desirable behavior has been addressed in previous experimental work. Most of the 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, Sefton et al. (2007) examine the effect of rewards and punishments on contributions in a repeated public good game and find that punishments help subjects to sustain higher cooperation levels compared to a control game with no reward/punishment opportunities, whereas the possibility of rewards has only a transient effect.[8]

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.[9] 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.

Most closely related to our work is Rauhut (2009) who studies the impact of the severity of the punishment in an inspection game. His set up differs from ours in that the punishment

[8] More recent research has shown that the effectiveness of rewards and punishments in settings such as this depends on the rewarding/punishing technology. Sutter et al. (2010) find that when the benefit/cost of receiving reward/punishment is three times larger than the cost of delivering it (i.e. with a 3:1 technology), both mechanisms are effective in encouraging contributions. Similarly, Rand et al. (2009) find that rewards are as effective as punishments in sustaining cooperation in a repeated public good game experiment with unknown time horizon and with a 3:1 reward/punishment technology. Gürerk et al. (2006) study a public good game where the rewarding mechanism displays a 1:1 technology and a punishment mechanism displays a 3:1 technology. They find that only the latter have an impact on contributions. Gürerk et al. (2009) use a public goods game where one group member (the 'leader') can reward or punish the other contributors. Although both rewarding and punishment mechanisms display a 3:1 technology, they find that contributions are higher when punishments are used.

[9] See Dickinson (2001) for a public good game experiment where rewards/punishments points are automatically assigned to the highest/lowest contributor in the group.

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.[10] Our paper differs from his also in that we study reward as well as punishment. As far as we are aware ours is the first paper to experimentally compare positive and negative incentives in inspection games.

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., *forthcoming*; Goerg, et al., *forthcoming*). 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 (although, for our data, the latter concept performs better than Nash equilibrium only if it incorporates loss aversion). More generally, our results show that when Nash equilibrium and alternative predictions diverge we find more support for the latter than for the former.

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 misconduct or bonuses for good conduct). 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 one 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

[10] 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.

own-payoff effect of working. Thus, in such a setup bonuses may have a positive effect on inspectees' good behavior. 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 endogenous choice of punishing and rewarding mechanisms.

REFERENCES

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.

Andreoni, J., W. Harbaugh, and L. Vesterlund. 2003. The carrot or the stick: Rewards, punishments, and cooperation. *American Economic Review* 93(3), 893-902.

Brunner, C., C. Camerer, and J. Goeree. A Correction and Re-Examination of "Stationary Concepts for Experimental 2 x 2 Games". *American Economic Review, forthcoming*.

Charness, G., and M. Rabin. 2002. Understanding social preferences with simple tests. *Quarterly Journal of Economics* 117(3), 817-69.

Dickinson, D.L. 2001. The carrot vs. the stick in work team motivation. *Experimental Economics* 4(1), 107-124.

Fudenberg, D., and J. Tirole. 1992. *Game Theory*. Cambridge, MA: MIT Press.

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.

Goerg, S., T. Chmura, and R. Selten. Correction and Re-Examination of Stationary Concepts for Experimental 2x2 Games: A Reply. *American Economic Review, forthcoming*.

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.

McKelvey, R.D., and T.R. Palfrey. 1995. Quantal Response Equilibria for normal form games. *Games and Economic Behavior* 10(1), 6-38.

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) (September), 1423-1437.

Potters, J., and F. van Winden. 1996. Comparative statics of a signaling game: An experimental study. *International Journal of Game Theory* 25(3), 329-353.

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.

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.

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.

APPENDIX A (not for publication)

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.

APPENDIX B

In this appendix, we explain the procedure to derive the equilibrium predictions of IBE and QRE with loss aversion in the context of the canonical inspection game. Selten and Chmura (2008) provide a more general discussion for IBE and Brunner et al. (forthcoming) for QRE.

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 B.1 presents the canonical inspection game, 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 $x - \max\{½(x - s),0\}$. As a consequence, losses compared to the reference point weigh twice the amount that gains weigh. The middle panel of Figure B.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 B.1.

Figure B.1 Canonical game, Transformed game and Impulse matrix

Canonical Inspection Game

	H	L
I	52, 25	12, 20
N	60, 25	0, 40

Transformed Inspection Game

	H	L
I	32, 25	12, 20
N	36, 25	0, 32.5

Impulse Matrix

	H	L
I	4, 0	0, 5
N	0, 7.5	12, 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 L, then p and q follow from the solution of *the impulse balance equations*:

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

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:

$$p=\frac{e^{\lambda E_{employer}[I]}}{e^{\lambda E_{employer}[I]}+e^{\lambda E_{employer}[NI]}}, \qquad q=\frac{e^{\lambda E_{worker}[L]}}{e^{\lambda E_{worker}[L]}+e^{\lambda E_{worker}[H]}},$$

where λ represents the player's rationality parameter that is estimated from the data. For QRE with loss aversion, the payoffs of the Transformed inspection game are used. In this case, p and q follow from the solution of:

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

The QRE prediction for the game without loss aversion is similarly found using the ordinary payoffs listed in the left panel of Figure B.1.