

Discretionary Sanctions and Rewards in the Repeated Inspection Game

by

Daniele Nosenzo^{*}, Theo Offerman^{**}, Martin Sefton^{*} and Ailko van der Veen^{***}

1 August 2012

Abstract:

We experimentally investigate a repeated “inspection game” where, in the stage game, an employee can either work or shirk and an employer simultaneously chooses to inspect or not inspect. Combined payoffs are maximized when the employee works and the employer does not inspect. However, the unique equilibrium of the stage game is in mixed strategies with positive probabilities of shirking/inspecting. We examine the effects of allowing the employer to sanction or reward the employee after she has inspected the employee. We find that rewards or sanctions can both discourage shirking, and have similar effects on joint earnings. In games allowing sanctions a reduction in shirking is accomplished with a lower inspection rate and the efficiency gains accrue to employers. In games allowing rewards employers actively reward employees for working and the efficiency gains are shared more equitably. A treatment where employers can combine sanctions and rewards leads to efficiencies similar to the single-instrument treatments, and outcomes more closely resemble those of the reward treatment in that the efficiency gains are shared.

Keywords: Inspection Game; Costly Monitoring; Discretionary Incentives; Rewards; Punishment; Experiment.

JEL Classification Numbers: C70, C72, C92

Acknowledgements: We thank participants at the 2012 International ESA Conference in New York, the 2011 CESS/CREED conference on experimental economics at NYU and the 2011 CREED-CeDEx-CBESS Meeting at the University of Nottingham for useful comments. We thank CREED programmer Jos Theelen for programming the experiment. Daniele Nosenzo acknowledges support from the Leverhulme Trust (ECF/2010/0636) and 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.

^{*} CeDEx, School of Economics, University of Nottingham, Nottingham, NG7 2RD, United Kingdom

^{**} CREED, Department of Economics, University of Amsterdam, 1018 WB, Netherlands

^{***} CBESS, School of Economics, University of East Anglia, Norwich, NR4 7TJ, United Kingdom

Nosenzo: email: daniele.nosenzo[at]nottingham.ac.uk, tel. +44 115 846 7492

Offerman: email: t.j.s.offerman[at]uva.nl, tel. +31 20 525 4294

Sefton: email: martin.sefton[at]nottingham.ac.uk, tel. +44 115 846 6130

van der Veen: email: ailko[at]xs4all.nl, tel. +44 1603 59 1685

1. INTRODUCTION

In this paper we compare the effectiveness of positive and negative incentives in an inspection game. This game is often used to represent strategic settings characterized by an imperfect alignment of interests between players (e.g., interactions between employers and employees, tax authorities and taxpayers, regulators and firms, law enforcement agencies and citizens, etc.).¹ Note that these settings typically have a hierarchical structure: an authority wishes to induce compliance from subordinates. A standard approach to encourage compliance is to use explicit contracts that specify automatic and fixed penalties in response to observed non-compliance. For example, labor contracts may specify penalties for employees who are found to underperform or violate the company's conduct policy. In addition to automatic incentives, authorities may also use *discretionary incentives* to align subordinates' interests with their own. For example, in the labor context, the nature and severity of the sanctions relating to underperformance may vary from verbal and written warnings to dismissal, and employers often have discretion over the disciplinary actions to be taken against employees. Moreover, in many settings authorities complement the use of sanctions for poor performance with the use of automatic and/or discretionary rewards. For example, again in the labor context, employers use a variety of bonus schemes to reward good performance.

Our experiment incorporates discretionary rewards and sanctions in a version of the standard inspection game discussed in Fudenberg and Tirole (1991). They discuss an interaction between an employee who chooses whether to work or shirk and an employer who simultaneously chooses whether or not to inspect the employee. Working is costly to the employee and generates revenue for the employer. Inspections are costly to the employer but reveal whether the employee works. The employee receives a wage from the employer unless she is caught shirking: in this case the employee is automatically sanctioned by the employer, who withholds her wage. Joint payoffs are maximized when the employee works and the employer does not inspect, but in the unique mixed-strategy Nash equilibrium of the one-shot game inspections and shirking occur with positive probability. To allow for discretionary incentives, we modify the inspection game described above by allowing employers to sanction or reward employees after an inspection. Both sanctions and rewards are costly for the employer, and sanctions reduce the employees' payoff while rewards increase it.

¹ See Avenhaus et al. (2002) for a review of the theory and discussion of applications of inspection games.

Because both discretionary sanctions and rewards are costly to implement, they will not be chosen by a selfish and rational employer in a subgame perfect equilibrium of the one-shot game, and equilibrium outcomes are thus unaffected relative to the standard game. However, as we discuss in Section 2, if the game is indefinitely repeated a broader set of outcomes can be sustained in equilibrium. Moreover, which outcomes can be sustained depends on which instrument is available. In particular, when discretionary sanctions are available the threat of sanctions may induce employees to work without the need for inspections, and so it is possible to attain the joint-payoff maximizing outcome in equilibrium. Note that since shirking is discouraged without incurring the cost of inspecting this is also the employer's preferred outcome. Thus, the employer may be made better off by the availability of discretionary sanctions relative to the stage game equilibrium, whereas the employee's payoff is unchanged. On the other hand, while the availability of discretionary rewards cannot sustain the joint-payoff maximizing outcome, they can improve efficiency relative to the stage game equilibrium by sustaining outcomes where both players are made better off.

We describe our laboratory experiment in detail in Section 3. In the experiment employer-employee pairs interact repeatedly for an indeterminate length of time. In a baseline treatment we use a standard inspection game with automatic sanctions and no discretionary incentives. In four other treatments we introduce the possibility for employers to use discretionary incentives to discipline employees. In these treatments we implement a 2x2 design where we vary the instrument employers can use to incentivize employees (either "punishment tokens" or "reward tokens"), and the power of the incentives (either low or high). In all treatments employers pay 1 point for each punishment/reward token assigned to the employee. Each punishment/reward token assigned decreases/increases the employee's payoff by 1 point in the "low-power" incentive treatments and by 3 points in the "high-power" treatments.

We report the results of the experiments in Section 4. In the baseline treatment with no discretionary incentives the proportion of inspections is 70% and employees shirk about 46% of the time. Low-power sanctions or rewards do not affect the proportions of shirking or inspecting relative to this baseline level. On the other hand, high-power incentives are effective: either the availability of sanctions or rewards reduces the proportion of shirking relative to the baseline treatment (to 29% in both cases). In the treatment with sanctions the reduction of shirking is achieved with a lower inspection rate than in the treatment with

rewards. An implication of this is that, while both instruments increase efficiency relative to the baseline treatment, the distribution of efficiency gains depends on the type of discretionary incentive available to the employer. Efficiency gains accrue solely to the employer when sanctions are available, whereas the efficiency gains are shared in the case of discretionary rewards.

In order to compare the effectiveness of discretionary rewards and sanctions our experiment varies the availability of the instruments across treatments, and employers have available at most one of the instruments. In natural workplaces both instruments are often available to employers. In Section 5 we report two additional treatments where employers can combine discretionary sanctions and rewards: after an inspection, employers choose whether to sanction, reward or take no action against an employee. Consistent with our earlier results we find that low-power incentives do not increase efficiency while high-power incentives reduce shirking and increase efficiency. In the case of high-power incentives we find that, as in the case where only discretionary rewards are available, the efficiency gains are shared by the employer and employee. Interestingly, when employers have access to both instruments they are no better off than when they only have access to sanctions. In fact, their use of discretionary rewards lowers their earnings slightly relative to the treatment where only sanctions are available.

Several related literatures compare the effectiveness of sanctions and rewards as incentive schemes, though in different settings and under different conditions than those studied here. One related literature focuses on social dilemma settings (e.g., Gülerk et al., 2006; Sefton et al., 2007; Rand et al., 2009; Sutter et al., 2010; Drouvelis and Jamison, 2012). There are several differences between the typical setup studied in this literature and our inspection game. A key difference between the settings is that in the inspection game players are asymmetric in terms of their ability to assign or receive punishments or rewards, whereas in the typical social dilemma situation players can mutually punish/reward each other. Thus, our setup seems better suited to study the effectiveness of positive and negative incentives in hierarchical interactions. In this sense, our study is also related to the literature on the use of bonuses and fines in principal-agent games (e.g., Fehr et al., 2007; Fehr and Schmidt, 2007). However, the focus of this literature is on the comparison between automatic (enforceable) incentives and discretionary incentives that cannot be enforced by a third party. In contrast, in this paper we focus on two different forms of discretionary incentives (rewards and sanctions)

and compare their effectiveness in disciplining shirking.² Most closely related is Nosenzo et al. (2010), who also examine the effectiveness of sanctions and rewards in an inspection game.³ However, differently from the present paper, Nosenzo et al. (2010) focus on automatic (non-discretionary) incentives that are pure transfers between players, and examine their effectiveness in one-shot inspection games. We discuss these related literatures and how their results relate to ours in more detail in Section 6.

In Section 7 we offer some concluding comments. Overall, our findings suggest that positive and negative discretionary incentives can be equally effective in disciplining the behavior of subordinates and increasing efficiency. However, the availability of either instrument has different implications for how the efficiency gains are distributed between players. Authorities can rely on the threatening power of sanctions to discipline subordinates, and can thus increase compliance without the need for a high inspection rate. An implication of this is that, by using sanctions, authorities can reap the most of the efficiency gains generated by the incentive tool. On the other hand, the effectiveness of rewards hinges upon the active use of the incentive tool: authorities need to constantly engage in costly inspections and reward the subordinate to induce compliance. An implication of this is that efficiency gains are distributed more equitably between players when rewards are available.

2. THE INSPECTION GAME: THEORETICAL BENCHMARKS

The inspection game involves two players and simultaneous moves. The employer chooses between inspect and not inspect, and the employee shirks or works. In the standard version of the game, the employer incurs a cost of h from inspecting. If the employee provides high effort, the employee incurs a cost of c and the employer receives a revenue of v . If the employer does not inspect, the employee always receives a wage of w . If the employer inspects, the employee receives the wage only if she works. The resulting payoffs are shown Figure 1(a). We assume that all variables are positive and $v > c$, $w > h$, $w > c$. Note that joint

² There are also related studies that compare economically equivalent contracts that are framed either as bonuses or fines (e.g., Hannan et al., 2005; Bigoni et al., 2011; Armantier and Boly, 2012; Hossain and List, *forthcoming*). In contrast to these studies, in our setting the difference between reward and sanctions is not simply a matter of framing, and the two instruments provide different incentives to the players.

³ As far as we are aware there have only been two other experimental studies of inspection games. Glimcher et al. (2005) discuss inspection games with different parameterizations of the inspection cost, while Rauhut (2009) studies the impact of the severity of automatic sanctions. Neither study compares sanctions with rewards.

payoffs are maximized when the employee works and the employer does not inspect. Figure 1(b) presents the payoffs that we used in the experiment.⁴

Let p denote the probability of inspection and q denote the probability of shirking. In the unique Nash equilibrium the employer inspects with probability $p = c/w$ and the employee chooses to shirk with probability $q = h/w$. The employer receives an expected payoff of $\pi^{employer} = v - w - hv/w$, the employee receives an expected payoff of $\pi^{employee} = w - c$, and joint payoffs are $\pi^{employer} + \pi^{employee} = v - c - hv/w$. In the version of the game used in the experiment the employer's equilibrium inspection probability is $p = 3/4$ and the employee's equilibrium shirking probability is $q = 3/4$, giving expected payoffs of 15 for the employer and 20 for the employee. This inspection game is the stage game in our baseline treatment.

Figure 1: Inspection game

	Work	Shirk		Work	Shirk
Inspect	$v - w - h$ $w - c$	$- h$ 0	Inspect	30 20	10 15
Not inspect	$v - w$ $w - c$	$- w$ w	Not inspect	45 20	5 35

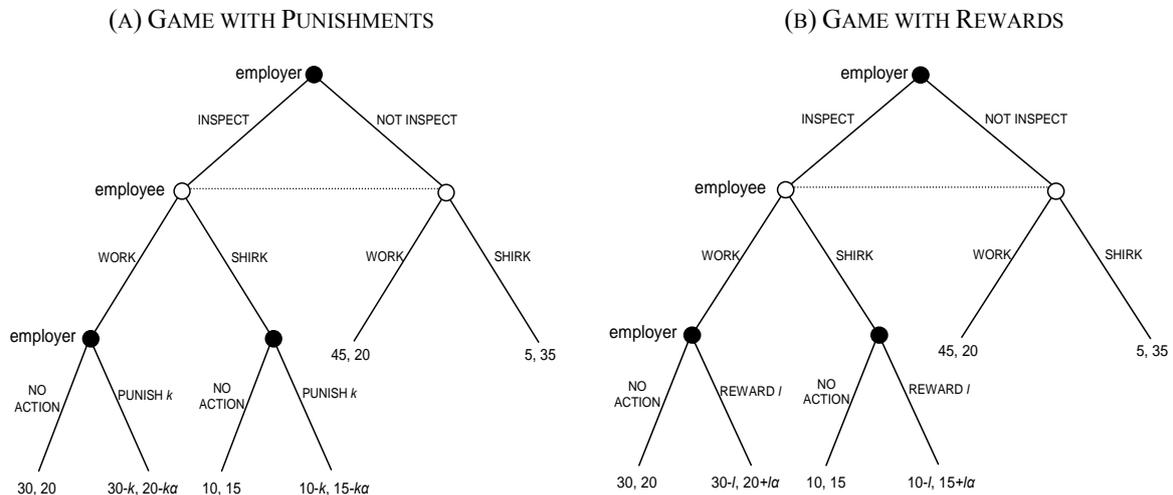
Notes: Employer is ROW player, Employee is COLUMN player. Within each cell, the Employer's payoff is shown at the top and the Employee's payoff at the bottom.

In the games where we allow for punishments the stage game of the baseline treatment is augmented in the following way. If the employer inspects, he observes the employee's choice and then chooses between "No action" and "Punish". If he chooses No action, then the payoffs are simply determined by the payoffs of the Inspection game. If he chooses Punish he must assign a punishment level k from the set $\{0, 1, 2, 3, 4, 5\}$ and the employer's payoff from the inspection game is then decreased by k while the employee's payoff is decreased by αk . Thus, these discretionary punishments are costly for both parties and have a negative

⁴ For the experiment we used the parameters $v = 40$, $w = 20$, $c = 15$ and $h = 15$. We then added 15 to each of the employee's potential payoffs and 20 to each of the employer's possible payoffs because we wanted to prevent negative outcomes (which are problematic to implement in an experiment). We also wanted to avoid zeros in the payoff matrix presented to subjects, and so added a further 5 points to each of the employer's possible payoffs.

direct impact on combined earnings.⁵ Figure 2(a) presents this augmented game graphically. The parameter α measures the impact of the instrument on the employee's payoff: in our experiment we use low-power ($\alpha = 1$) and high-power ($\alpha = 3$) punishments.

Figure 2: Augmented inspection games



Similarly, in the games where we allow for rewards the employer can choose between “No action” and “Reward” after an inspection. If he chooses Reward he then chooses the reward level l from the set $\{0, 1, 2, 3, 4, 5\}$ and the employer's payoff from the inspection game is then decreased by l while the employee's payoff is increased by αl . In our experiment we use both low-power ($\alpha = 1$) and high-power ($\alpha = 3$) rewards. Low-power rewards correspond to the case where rewards are pure transfers from the employer to employee, and thus have no direct impact on combined earnings. High-power rewards imply that rewards can increase combined earnings; a maximal reward costs the employer 5 points and benefits the employee 15 points, giving a net benefit of 10 points. Note, however, that rewards can only be given following an inspection, and the inspection cost (15 points) exceeds the net benefit from maximal rewards. Thus, combined earnings are still maximized when the employee works and the employer does not inspect. The augmented game with reward possibilities is shown in Figure 2(b).

Subgame perfect equilibria of the augmented games can be identified by backward induction. After inspection, a selfish and rational employer will either choose No Action or zero punishment ($k = 0$) or reward ($l = 0$). This behavior is anticipated by the players, and, as

⁵ Except, of course, in the case where the employer assigns zero punishment. We decided to include this in the set of available punishments as it may be useful for signaling purposes in settings where the game is played repeatedly, e.g. an employer might assign zero punishment tokens as a warning.

a result, play in the phase preceding the punishment/reward phase remains unaffected. Thus, in the subgame perfect equilibrium players mix between their actions in precisely the same way as in the baseline treatment, i.e., $p = \frac{3}{4}$ and $q = \frac{3}{4}$.

In naturally occurring workplace settings, and in our experiment, employers and employees are usually engaged in a repeated interaction. Here, we consider the case where in each stage the game described above is played and where a player's earnings are simply the sum of his earnings over all stage games. After each stage game, there will be a new stage game with independent probability δ and this process continues until it is terminated by chance. As is well known, repetition of the stage game equilibrium constitutes a subgame perfect equilibrium of the indefinitely repeated game, but other outcomes can be sustained as equilibria as well. What can be sustained in equilibrium depends on three factors: whether punishments or rewards are available, the parameter α , and the continuation probability.

First, when punishment is not possible repetition of the joint payoff maximizing outcome cannot be sustained in equilibrium. To see this note that for any pair of strategies yielding the outcome (Not Inspect, Work) in every stage, and hence a payoff of 20 for the employee in every stage, the employee can deviate to a strategy that specifies shirking in the first stage and working in all subsequent stages. This deviation is profitable since it yields 35 in the first stage and 20 in all subsequent stages.

In contrast, when punishments are available, and if the continuation probability is sufficiently high, it is possible to attain repetition of the (Not Inspect, Work) outcome in a subgame perfect equilibrium. This is possible because punishment allows the employer to reduce the employee's stage payoff below 20, and this can then serve as a threat that induces the employee to work. The critical discount factor needed to sustain the joint-profit maximizing outcome in a subgame perfect equilibrium depends on the parameter α . The key point is that with a higher α the threatened punishment is more severe, and so the employee can be discouraged from shirking for a larger set of discount factors. In this sense, in the game augmented by punishment the joint payoff maximizing outcome is easier to sustain when δ is high, or α is high (see Appendix A for technical details).

One of the features of the joint payoff maximizing outcome is that, relative to the stage game equilibrium, it offers an improved payoff *only to the employer*. When the inspection game is augmented by the possibility of rewards the stage game has pure strategy outcomes (where the employee works and the employer inspects and then rewards) with

payoffs that exceed the stage game equilibrium for *both* players. For example, when the employee works and the employer inspects and then maximally rewards the employee the employer receives a payoff of 25 and the employee receives a payoff of $20 + 5\alpha$. This outcome can also be sustained in a subgame perfect equilibrium for a range of (α, δ) . Again, it is easier to sustain this cooperative outcome for higher δ or higher α (see Appendix A).

In summary, when the game is indefinitely repeated and the continuation probability is sufficiently high a large set of outcomes can be sustained in a subgame perfect equilibrium. Which outcomes can be sustained depends on whether rewards or punishments are available, and the costs and consequences of using these. Combined payoffs can be maximized when punishment is available, essentially because it reduces the payoff the employee can guarantee herself and this introduces scope for threats that can keep the employee working without the need for inspections. An implication of this is that all of the efficiency gains accrue to the employer. These outcomes cannot be attained in the game augmented by rewards. On the other hand rewards can be used to achieve outcomes where both players are made better off.

3. EXPERIMENTAL DESIGN AND PROCEDURES

The computerized experiments were carried out at the University of Nottingham with subjects recruited from a campus-wide distribution list.⁶ In the initial study 178 subjects participated and no subject participated in more than one session. Three sessions were conducted for each of five treatments, with either five or six pairs of participants in a session. Sessions consisted of a number of rounds and at the end of a session subjects were paid in cash according to their accumulated point earnings from all rounds. Sessions took about 40 minutes on average and earnings ranged between £5.65 and £23.00, averaging £12.05 (by comparison, the adult minimum wage was £5.93 an hour at the time of the experiment).

At the beginning of a session subjects were randomly assigned to computer terminals and given paper copies of instructions, which an experimenter then read out loud. The instructions concluded with a series of questions testing subjects' understanding of the instructions. Answers were checked by the experimenters, who dealt privately with any remaining questions. During a session no communication between subjects was allowed.

⁶ Subjects were recruited through the online recruitment system ORSEE (Greiner, 2004). The instructions to subjects that were used in the experiment are reproduced in Appendix B.

After the instructional phase subjects were assigned to pairs and roles. Within each pair, one subject received the role of Employer and the other the role of Employee.⁷ Subjects knew that they would stay in the same role and in the same pair during the whole experiment. They were informed that the session consisted of at least 70 rounds. From round 70 onwards, each round could be the last one with probability 1/5.⁸

In each treatment, at the beginning of a round the Employee chose between “high effort” (work) and “low effort” (shirk) and, at the same time, the Employer chose between “inspect” and “not inspect”. Choices led to point earnings as presented in the right panel of Figure 1. In the **Baseline** treatment these were the only choices made in the round, and subjects were immediately informed about the choices and point earnings within their pair.

The other treatments varied from the Baseline treatment in the instruments available to employers for incentivizing employees (punishments or rewards), and the effectiveness of the tool ($\alpha = 1$ or $\alpha = 3$). In these treatments, after being informed of whether the Employee chose work or shirk, if the Employer had chosen to inspect he had to make an additional choice. In the “**P1:1**” and “**P1:3**” treatments the Employer chose between “no action” and “punish”, and if “punish” was chosen the Employer then chose the number of punishment tokens from the set $\{0, 1, 2, 3, 4, 5\}$. The Employer paid a cost of 1 point per token. In the “1:1” treatment each token reduced the Employee’s earnings by one point, while in the “1:3” treatment each token reduced the Employee’s earnings by three points. In the “**R1:1**” and “**R1:3**” treatments the Employer chose between “no action” and “reward”, and if “reward” was chosen he then had to choose the number of reward tokens from the set $\{0, 1, 2, 3, 4, 5\}$. Each token cost the Employer one point and increased the Employee’s earnings by either one (“1:1” treatment) or three (“1:3” treatment) points. Finally, both players in the pair were informed of the results in the pair (all choices and earnings). Table 1 summarizes the experimental design.

Table 1: Design of the initial study

Treatment	Punishments	Rewards	Technology	Number of pairs
Baseline	No	No	---	17
P1:1	Yes	No	1:1	18
P1:3	Yes	No	1:3	18
R1:1	No	Yes	1:1	18
R1:3	No	Yes	1:3	18

⁷ The actual labels used in the experiment were “Employer” and “Worker”.

⁸ In fact the last round was randomly determined according to these rules prior to the Baseline sessions and this resulted in three sessions with 71, 73 and 83 rounds, respectively. We then used these durations for the other treatments as well.

4. RESULTS

4.1 The Impact of Incentives on Inspecting and Shirking

Figure 3 displays the proportion of inspecting (top panels) and shirking (bottom panels) across rounds disaggregated by treatment.⁹ First, consider the left panels containing data from the Baseline, R1:1 and P1:1 treatments. In all three treatments the inspection rate (top-left panel) increases across rounds and approaches the Nash stage game equilibrium level (75%) in the last third of the experiment. In contrast, the proportion of shirking (bottom-left panel) is quite stable across rounds and much lower than the Nash stage game equilibrium level (75%). Averaging across rounds, the inspection rate is similar in the three treatments (62% in P1:1, 70% in R1:1, 70% in Baseline), and differences between treatments are not statistically significant ($p > 0.178$ in all bilateral comparisons).¹⁰ Similarly, the rate of shirking is similar across treatments (43% in P1:1, 43% in R1:1, 46% in Baseline), and none of the relevant bilateral comparisons are statistically significant (all $p > 0.338$).¹¹ Thus, the availability of low-power punishments and rewards (P1:1 and R1:1 treatments) does not affect the proportion of shirking or inspecting relative to a Baseline treatment where discretionary punishments and rewards are unavailable. In all cases, employees shirk about 40% of the times whereas the relative frequency of inspections is about 70%.

Treatment differences are more evident in the right panels of Figure 3, which present data from the 1:3 and Baseline treatments. First, the inspection rate in the P1:3 treatment (56%, averaged across rounds) is noticeably lower than in the R1:3 (76%) and Baseline (70%) treatments, although only the difference between P1:3 and R1:3 is significant ($p = 0.046$).¹² Second, the rate of shirking (bottom-right panel of Figure 3) is noticeably lower in P1:3 (29%) and R1:3 (29%) than in Baseline (46%), and these differences are statistically significant ($p = 0.027$ for P1:3 vs. Baseline; $p = 0.030$ for R1:3 vs. Baseline). It is worth noting that the rate of shirking is not significantly different in P1:3 and R1:3 ($p = 0.962$), and so the lower inspection

⁹ We restrict the analysis to the first 70 rounds where we have data from all 89 pairs who took part in the experiment. Analysis using all data is complicated by the fact that some pairs interacted for more rounds than others and we have very small sample sizes in the later rounds. However, all results reported here also hold in the full sample.

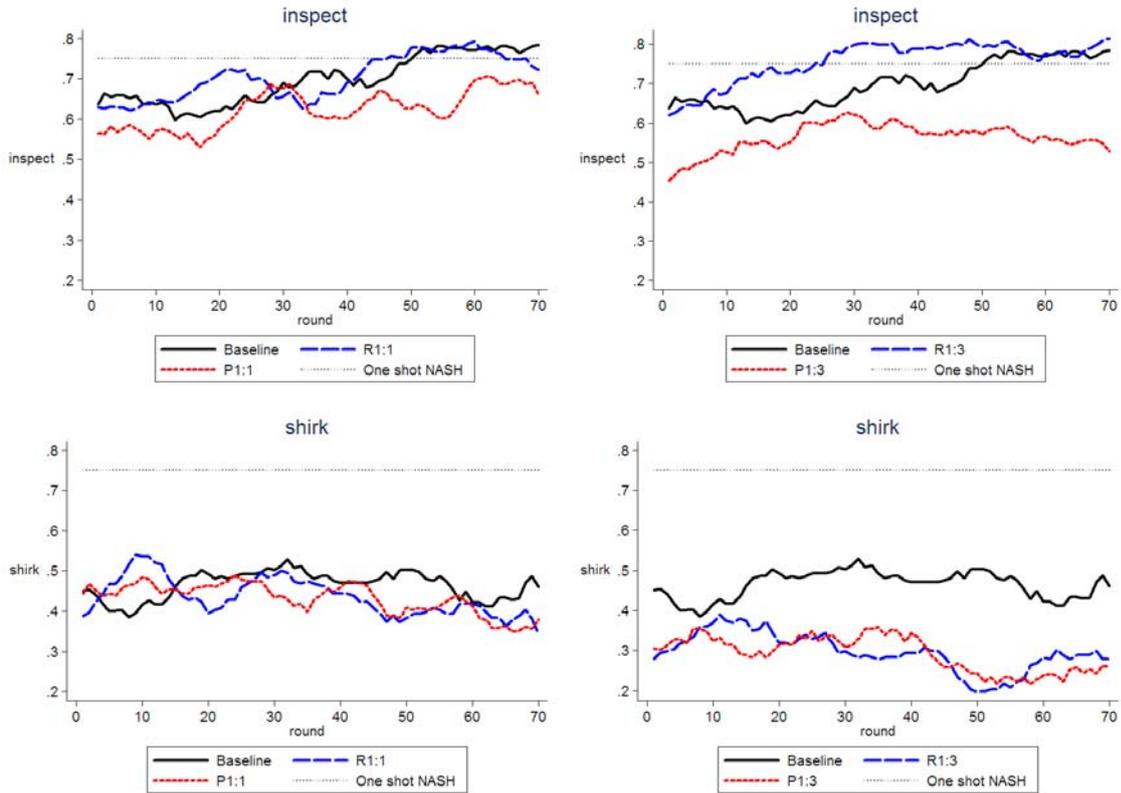
¹⁰ Our non-parametric analysis is based on tests applied to 18 independent observations per treatment (17 in Baseline). We consider data from each pair as one independent observation. Tests are applied to averages based on the first 70 rounds of the experiment. Unless otherwise specified, the reported p-values are based on two-tailed Wilcoxon rank-sum tests.

¹¹ Note that the rate of shirking in Baseline (46%) is considerably lower than the Nash equilibrium prediction for the stage game (75%). This suggests that repetition of the stage game in itself can be effective in disciplining employees.

¹² The differences between Baseline and R1:3 or P1:3 are not significant ($p > 0.209$).

frequency in P1:3, relative to R1:3, is not associated with higher shirking. Thus, when high-power incentives are used (P1:3 and R1:3 treatments) both incentive tools are equally effective in reducing shirking, but inspection rates are lower when the punishment tool is available.

Figure 3: Proportion of inspections (top panel) and shirking (bottom panel) across rounds



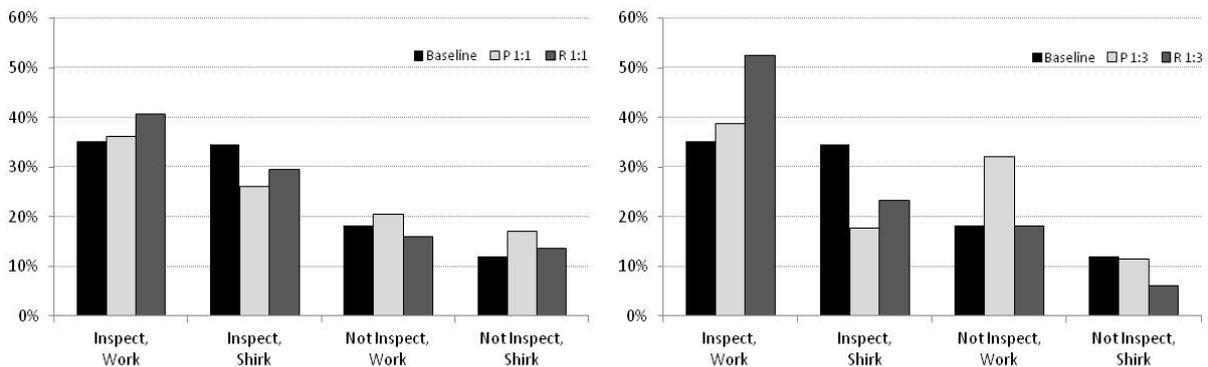
Notes: the Figure is based on data from 18 games per treatment (17 in Baseline).

Figure 4 shows the relative frequencies of the four possible outcomes (Inspect, Work; Not Inspect, Work; Inspect, Shirk; Not Inspect, Shirk). The 1:1 treatments are in the left panel and the 1:3 treatments in the right panel, with Baseline reproduced in both panels to facilitate comparisons. The distributions of outcomes in the 1:1 treatments are not very different from that observed in Baseline (see the left panel). In fact, the distributions of outcomes are not significantly different in any bilateral comparison of treatments (randomization tests: all $p > 0.551$).¹³ However, with higher-power incentives the distribution of outcomes in Baseline is significantly different from those in R1:3 and P1:3 (randomization tests: $p = 0.046$ and $p =$

¹³ To measure the difference in distributions across two treatments we computed the sum of the squared differences in the average relative frequencies of each outcome. P-values are based on two-sided randomization tests using Monte Carlo simulations (tsrtest command in Stata, see Kaiser and Lacy, 2009).

0.086, respectively).¹⁴ Inspection of Figure 4 reveals that the main difference between Baseline and R1:3 is that (Inspect, Work) is observed more frequently and (Inspect, Shirk) is observed less frequently in the treatment with rewards. The availability of punishment also reduces the frequency of the outcome (Inspect, Shirk) relative to Baseline. However, differently from the R1:3 treatment, this is combined with a marked increase in the frequency of the outcome (Not Inspect, Work), whereas the frequency of (Inspect, Work) is only slightly higher in P1:3 than in Baseline.

Figure 4: Distribution of outcomes by treatment



Notes: based on data from 1260 games per treatment (1190 games in Baseline)

The observed pattern of effects in the 1:3 treatments is consistent with the repeated game analysis presented earlier: the outcome that maximizes combined payoffs (Not Inspect, Work) is attained most easily when punishments are available (whereas the availability of rewards does not make much difference relative to the Baseline treatment), while the availability of rewards increases the frequency of the (Inspect, Work) outcome. The finding that differences between Baseline and the other treatments are only observed with high-power incentives is also consistent with the repeated game considerations discussed in Section 2, in that an increase in the impact of a punishment/reward token on the employee's payoff generally weakens the conditions required for punishments or rewards to sustain an outcome as part of an equilibrium.

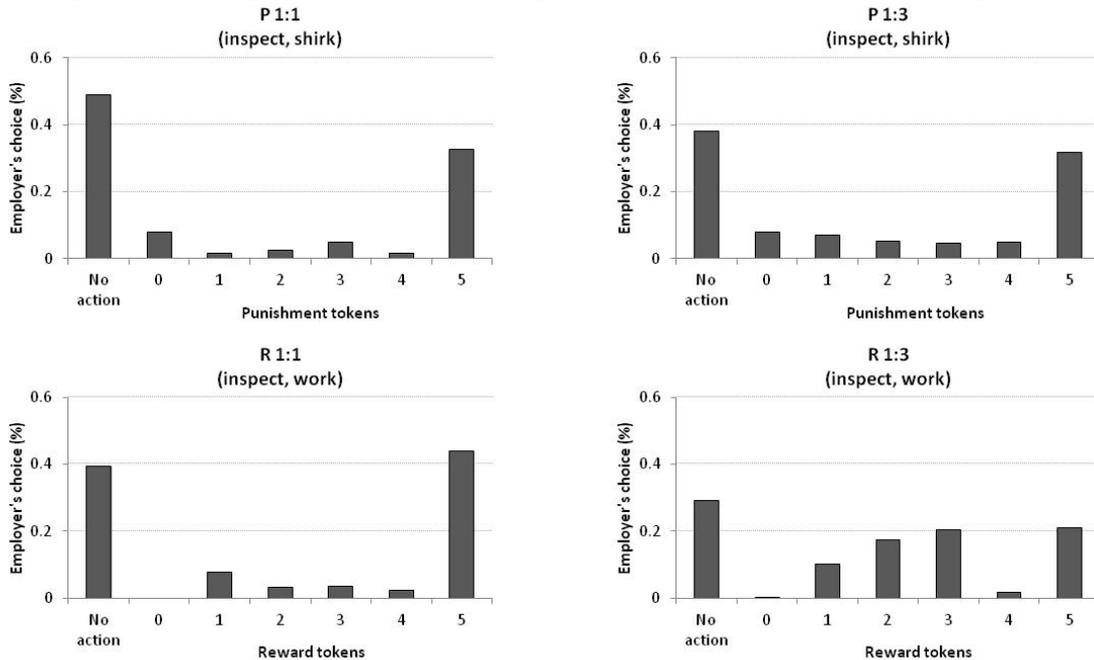
4.2 The Use of Punishments and Rewards

We next examine how employers used discretionary incentives in the treatments with punishments or rewards. Following an inspection, the employer learns the choice of the

¹⁴ The distributions of outcomes in R1:3 and P1:3 are also significantly different from each other (randomization test: $p = 0.054$).

employee, and then decides whether to take no action or to assign punishment or reward tokens (depending on the treatment). Figure 5 shows the proportion of “No Action” decisions and punishment/rewards tokens assignments disaggregated by treatment. For the P1:1 and P1:3 treatments we report punishment decisions following the outcome (Inspect, Shirk), whereas for the R1:1 and R1:3 treatments we report reward decisions following the outcome (Inspect, Work).¹⁵

Figure 5: Use of punishments after (Inspect, Shirk) and rewards after (Inspect, Work)



Notes: based on 328 games in P1:1, 223 games in P1:3, 513 games in R1:1 and 660 games in R1:3.

In the P1:1 (top-left) and P1:3 (top-right) treatments punishment happens more often than not when an employee is caught shirking (51% of the games in P1:1 and 62% in P1:3). In both treatments by far the most common use of the incentive tool is to assign maximal punishment to the employee (5 tokens).¹⁶ Overall, the expected number of punishment tokens assigned to an employee caught shirking is equal to 1.90 in P1:1 and 2.10 in P1:3.

¹⁵ In the experiment punishments are mainly targeted at shirkers and rewards are predominantly given to employees observed to have worked in that period. Thus, we observe very little use of punishment after (Inspect, Work) – this occurs in 50 out of 944 games - and very little use of rewards after (Inspect, Shirk) – in 55 out of 667 games.

¹⁶ None of the other levels of punishment are used particularly frequently, but it is interesting to note that the second most frequent level of punishment (in both treatments) is to assign zero tokens. These actions may be intended as (costless) signals from employers of their readiness to punish the employee. However, following an inspection zero tokens were only assigned in about 8% of games.

In the R1:1 treatment (bottom-left) employers reward employees found working in 61% of the games, and in the R1:3 treatment (bottom-right) in 71% of the games. Thus, employees observed to have worked are rewarded somewhat more frequently than shirkers are punished: the frequency of “reward” decisions conditional on finding an employee working is higher than the frequency of “punish” decisions conditional on finding an employee shirking (61% vs. 51% in the 1:1 treatments, 71% vs. 62% in the 1:3 treatments).

We also find that high-power rewards are used differently from low-power rewards. When low-power rewards are available, these are most frequently used to reward employees maximally (this occurs in 44% of the games in R1:1). However, when the benefit of the reward to the employee is three times larger than the cost to the employer, employers tend to use maximal rewards less often (21% of the time), and assigning 2 or 3 reward tokens become more frequent (used respectively 17% and 20% of the time each). As a consequence, the expected number of reward tokens assigned to an employee who is inspected and found working drops from 2.53 in R1:1 to 2.18 in R1:3. This change in the way rewards are used may reflect employers’ concerns with relative earnings. For example, following (Inspect, Work), differences in earnings are minimized by assigning 5 reward tokens in R1:1 and by assigning 2 or 3 reward tokens in R1:3.

4.3 Efficiency and Earnings

Table 2 shows average earnings per game across treatments. We focus on total earnings, i.e. the earnings that players received at the end of each round, including any cost or benefit following the use of rewards and punishments. The Table reports players’ individual earnings as well as combined earnings.

Table 2: Efficiency and earnings

	Baseline (n = 17)	R 1:1 (n = 18)	P 1:1 (n = 18)	R 1:3 (n = 18)	P 1:3 (n = 18)
Employer’s Earnings	22.82 (7.66)	22.02 (3.46)	23.06 (4.53)	25.30 (4.22)	27.97 (6.93)
Employee’s Earnings	20.09 (2.42)	21.62 (1.67)	20.74 (2.46)	23.37 (3.49)	19.54 (2.06)
Combined Earnings	42.91 (8.64)	43.64 (4.18)	43.80 (5.34)	48.67 (7.20)	47.51 (7.69)

Notes: “Combined Earnings” shows the sum of the earnings of the employer and the employee. Standard deviations based on group averages in parentheses.

In Baseline combined earnings can range from 25 points (when the employer Inspects and the Employee shirks) to 65 points (when the Employer does not inspect and the Employee works). The efficiency level predicted by the Nash stage game equilibrium is 35 points. In the experiment, combined earnings are 23% higher than this, averaging 42.91 points across rounds. Averaged over all pairs, the main recipient of this efficiency gain is the employer, who earns much more than predicted (22.82 vs. 15 points) whereas employees' earnings are close to the predicted level (20.09 vs. 20 points).

Earnings in the treatments with low-power incentives do not differ much from Baseline. With the sole exception of the comparison between employees' earnings in R1:1 and Baseline (21.62 vs. 20.09 points; $p = 0.032$), individual and combined earnings in the incentive treatments are within one point of those in Baseline, and the differences are not statistically significant (all $p > 0.552$).

In contrast, high-power incentives have a greater impact on earnings. Both punishment and rewards have a similar impact on efficiency, which increases from 42.91 in Baseline to 47.51 in P1:3 ($p = 0.089$) and 48.67 in R1:3 ($p = 0.024$). Although combined earnings are slightly higher in R1:3 than P1:3 the difference is not statistically significant ($p = 0.569$).¹⁷ The ability to punish is beneficial to employers: relative to Baseline employers' earnings increase from 22.82 to 27.97 points ($p = 0.017$) in P1:3. At the same time employee's earnings are reduced, from 20.09 to 19.54 points, but not significantly so ($p = 0.478$). In contrast, the efficiency gains from rewards are shared by employers and employees. Relative to Baseline employees' earning increase by 16% in R1:3 (to 23.37 points; $p = 0.008$), and employers' earnings increase by 11% (to 25.30 points, $p = 0.086$). In summary, both high power rewards and punishments enhance efficiency, with the effects of both instruments being significant and similar in magnitude, but the efficiency gains accrue to employers in the case of punishments and are shared in the case of rewards.

¹⁷ Of course, each time a reward token is assigned in R1:3 joint earnings increase by 2 points so it is perhaps not surprising that joint earnings are higher in R1:3 (however, note that given our parameterization of the game in all treatments joint earnings are maximized when the employee works and the employer does not inspect). We also calculated efficiencies net of the costs and benefits of reward/punishment tokens. In this case efficiency is still significantly higher in the incentive treatments than Baseline, and the difference between efficiencies in R1:3 and P1:3 is still insignificant.

5. A FURTHER STUDY: REWARDS AND PUNISHMENT

Our initial study shows that, relative to Baseline, the availability of either rewards or punishments increases joint earnings, though by different channels and with different implications for the distribution of earnings. Relative to Baseline, in P1:3 there is an increased frequency of the (Not Inspect, Work) outcome, while in R1:3 there is an increased frequency of the (Inspect, Work) outcome followed by the use of rewards. Thus the availability of either punishment or rewards reduces shirking, with similar effects in the R1:3 and P1:3 treatments, although this is achieved in P1:3 with a lower inspection rate. The main recipients of efficiency gains are the employer when punishments are available, whereas efficiency gains are shared more equitably when rewards are available. These results raise an obvious question as to what would follow from the availability of both rewards *and* punishment, since both instruments are available to employers in many naturally occurring settings.

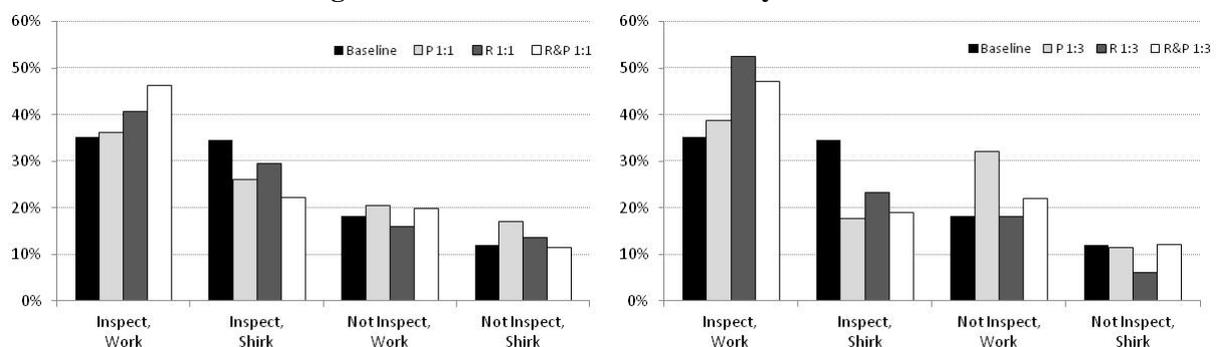
In this Section we report two additional treatments that examine this question. In these treatments employers could follow up an inspection with “no action”, “punish”, or “reward”, and, if “punish” or “reward” were chosen the employer could assign punishment or reward tokens. We ran treatments with low-power (**R&P1:1**) and high-power (**R&P1:3**) incentives. Apart from the expanded set of options following an inspection, the sessions were conducted in the same way as those of the initial study. In all we recruited 72 new subjects and ran three sessions of each treatment with twelve subjects per session. These sessions took about 40 minutes on average and earnings ranged between £7.10 and £23.00, averaging £12.33.

As in the initial study low-power incentives have little impact on behavior. In the R&P1:1 treatment neither the rate of shirking (34%) nor inspection rate (69%) is significantly different from Baseline or any of the other 1:1 incentive treatments ($p \geq 0.117$). In contrast, also as in our initial study, high-power incentives have a greater impact on behavior: the rate of shirking in R&P1:3 (31%) is similar to those in R1:3 and P1:3 (29%), and is significantly lower than in Baseline (31% vs. 46%, $p = 0.049$). An important result from the initial study is that the reduction of shirking is achieved with a lower inspection rate when the punishment tool is available than when rewards are available. When both punishments and rewards are simultaneously available the frequency of inspections is 66%, higher than when only punishments are available (P1:3, 56%), but lower than when only rewards are available (R1:3,

76%). The inspection rate in R&P1:3 is not significantly different from Baseline or any of the other two incentive treatments ($p \geq 0.199$).

Figure 6 shows the effect on outcomes by adding the data from the R&P treatments to the data previously shown in Figure 4. The Figure illustrates how the distribution of outcomes in the R&P treatments shares features of both the R and P treatments: in both panels there appears to be a reduction in the frequency of (Inspect, Shirk) in the R&P treatments relative to Baseline. This is combined partly with an increase in the frequency of (Inspect, Work), as in the treatments allowing rewards only, and partly with an increase of (Not Inspect, Work), as in the treatments allowing punishment only.¹⁸

Figure 6: Distribution of outcomes by treatment



Notes: based on data from 1260 games per treatment (1190 games in Baseline)

We emphasize two main findings from the R&P treatments. First, employees' behavior is not disciplined more effectively when employers can combine punishments and rewards, compared to the case where only one incentive tool is available. In fact, although the simultaneous availability of high-power punishments and rewards leads to higher efficiencies than in Baseline (47.15 vs. 42.91 points, $p = 0.069$), combined earnings are slightly lower than in R1:3 (48.67) or P1:3 (47.51).¹⁹

Second, the efficiency gains from combining punishment and reward instruments are shared by the employer and employee. Compared to Baseline, in R&P1:3 employees' earnings increase by about 10% (22.12 vs. 20.09 points, $p = 0.021$) and employers' earnings

¹⁸ For the case of R&P1:1 the shift is smaller and the distribution of outcomes is not significantly different from any of the other treatments (randomization tests: all $p \geq 0.277$). The effects are stronger in the R&P1:3 treatment where the distribution of outcomes is different from Baseline (randomization test: $p = 0.066$), but not from either P1:3 or R1:3 (randomization tests: $p \geq 0.265$).

¹⁹ However, efficiency in R&P1:3 is not significantly different from R1:3 ($p = 0.429$) or P1:3 ($p = 0.912$). Efficiencies in R&P1:1 (45.85 points on average) are also not significantly different from Baseline or the other 1:1 treatments.

also increase by about 10% (albeit insignificantly so; 25.03 vs. 22.82 points, $p = 0.121$).²⁰

This is different from the P1:3 treatment where efficiency gains accrue only to the employer, and more closely resembles the pattern in the R1:3 treatment.

It is interesting to note that employers' earnings are lower (although insignificantly so) when both instruments are available than when they can only punish (25.03 vs. 27.97, $p = 0.184$). This is counterintuitive since an employer offered an additional instrument could always choose not to use it. To examine this further we looked at the success of rewarding and punishing strategies in the R&P1:3 treatment versus the single-instrument treatments.²¹ We classified employees as "workers" or "shirkers", and employers as "punishers", "rewarders", "punishers and rewarders" or "non-punishers/non-rewarders" based on their behavior in the first 10 rounds of the experiment.²² We then looked at the relation between these behaviors and average employers' earnings in rounds 11-70. In the P1:3 treatment punishment seems to be an effective device for disciplining shirkers: "punishers" earn 29% more than "non-punishers" when matched with a "shirker". Also, in R1:3 rewards seem to be an effective device for encouraging workers: "rewarders" earn 16% more than "non-rewarders" when they are matched with a "worker".²³ In R&P1:3, however, punishment loses much of its bite: when matched with a "shirker", "punishers" earn just 9% more than "non-punishers".²⁴ In contrast, rewarding workers who behave well seems to become more remunerative in R&P1:3: "rewarders" earn 30% more than "non-rewarders" when matched

²⁰ Employers' earnings in R&P1:1 (23.99 points) are not significantly different from Baseline or the other 1:1 treatments. Employee's earnings in R&P1:1 (21.86 points) are significantly higher than in Baseline ($p = 0.016$): this is because shirking and inspection rates are similar to R1:1 and, as in that treatment, when employers inspect and observe the employee worked they usually reward, and often reward maximally.

²¹ The use of punishment and rewards in the R&P treatments is similar to that in the other incentive treatments: rewards are mainly assigned to employees who are found working and punishments to employees caught shirking. Moreover, rewards tend to be used more frequently than punishments: employees who are observed working are rewarded more frequently than shirkers are punished (74% vs. 26% in R&P1:1; 64% vs. 60% in R&P1:3). Figure A.1 in Appendix C shows the proportion of punishment/rewards tokens assignments in the R&P treatments.

²² Specifically, employers are classified on the basis of the average assigned reward tokens (r) and punishment tokens (p) over the first 10 rounds: if $\max(r, p) < 0.5$ then the employer is classified as "non-punisher/non-rewarder"; if $\max(r, p) \geq 0.5$ and $|r - p| < 0.25$ then the employer is classified as "punisher and rewarder"; if $\max(r, p) \geq 0.5$ and $(r - p) \geq 0.25$ then the employer is classified as "rewarder"; if $\max(r, p) \geq 0.5$ and $(p - r) \geq 0.25$ then the employer is classified as "punisher". In each 1:3 treatment, employees are classified on the basis of how often they shirked in the first 10 rounds: the 9 employees shirking least often in a session are classified as "workers", the other 9 employees as "shirkers".

²³ In contrast, punishment seems counter-productive in P1:3 when employers are matched with a "worker" – "punishers" earn 33% less than "non-punishers". In R1:3 rewards have much less impact when employers are matched with a "shirker" – "rewarders" earn just 4% more than "non-rewarders".

²⁴ Employers who are classified as "punishers and rewarders" do worse than this when matched with a shirker: they earn just 7% more than "non-punishers/non-rewarders". Similarly, rewards are not effective in encouraging shirkers: "rewarders" earn just 0.3% more than "non-rewarders" when matched with a "shirker".

with a “worker”.²⁵ This suggests that the effectiveness of punishments from the employer’s perspective is reduced when rewards are also available, and this may resolve the puzzle that employers offered both rewards and punishments seem to do worse than employers who are only offered the punishment instrument. In the next section we discuss how this result relates to findings from related literatures.

6. DISCUSSION OF MAIN FINDINGS IN RELATION TO THE LITERATURE

In our setting we find that both discretionary rewards and punishments can reduce shirking and enhance efficiency. Crucial to their effectiveness is the ratio of the cost of delivering reward/punishment to the benefit/cost of receiving it. High-power incentives (where the ratio is 1:3) are effective in reducing shirking and enhancing efficiency, while low-power incentives (1:1 ratio) have little impact on behavior or earnings. This finding is echoed in the literature on the effectiveness of discretionary rewards and punishments in public goods experiments. While high-power rewards and punishments are typically found to be effective in encouraging contributions (e.g., Rand et al., 2009; Sutter et al., 2010), low-power incentives are usually much less effective (e.g., Walker and Halloran, 2004; Sefton et al., 2007).²⁶

Also as in the public goods experiments literature (e.g., Rand et al., 2009; Sutter et al., 2010), we find that (high-power) rewards and punishments are equally effective in discouraging shirking. Relative to our baseline treatment, allowing (high-power) rewards results in a 16% decrease in shirking while allowing (high-power) punishments reduces shirking by 17%.

One finding that contrasts to the public goods experiments literature is that high-power rewards and punishments have similar effects on joint earnings in our inspection games (rewards increase combined earnings by 11%, and punishments increase combined earnings by 13% compared to the Baseline). In contrast, high-power rewards are typically found to have an efficiency-advantage in the public goods literature. This reflects an important difference between set-ups. In our set-up rewards can only be used following a costly inspection, and the cost of inspecting outweighs the efficiency gains associated with

²⁵ Combining punishments and rewards further increases employer earnings: “punishers and rewarders” earn 33% more than “non-punishers/non-rewarders” when matched with a “worker”.

²⁶ See Milinski and Rockenbach (2012) and Nosenzo and Sefton (2012) for a review of the literature comparing rewards and punishments in public goods experiments. See Balliet et al. (2011) for further discussion and for a meta-analysis of the effectiveness of rewards and punishments in social dilemmas.

the act of rewarding. Thus, inspecting and rewarding a worker has a negative impact on combined payoffs, whereas in the public goods literature the mere act of rewarding raises joint payoffs.

Our set-up presents several other important differences from the settings used in the public goods literature. One key difference is that in our setting the availability of the incentive tools is restricted to one player (the employer), whereas in the public goods literature typically players can mutually punish or reward each other.²⁷ In this sense, our setup may be better suited to examine the role of discretionary punishments and rewards in environments where there are hierarchical relations between agents.

Another difference is that in our setting there are asymmetries in the action spaces and payoff functions of the players, whereas players are usually symmetric in public goods games. A consequence of this is that there is little emphasis in the public goods literature on how the efficiency gains from using punishments and/or rewards are distributed across players. In contrast, in our setting, the type of instrument available has important implications for how the reduction in shirking is achieved and how the efficiency gains are distributed between the employer and the employee. When only punishment is available, employers rely more on the threat rather than the use of punishments, and the reduction in shirking is associated with a 15% decrease in the inspection rate relative to the Baseline treatment. An implication of this is that the employer is the one who gains from the increase in efficiency while employees are hardly affected: employer's earnings increase by 22% while employee's earnings decrease by 2%. In contrast, the disciplining effect of rewards is achieved by inspecting and rewarding employees who are observed working. An implication of this is that both players share the efficiency gains: employer's earnings increase by 11% and employee's earnings by 16%.

The effectiveness of sanctions and rewards in inspection games has also been studied in Nosenzo et al. (2010), although there are several important differences between the two settings. They study a one-shot inspection game, whereas we study a repeated game. Their punishments and rewards are pure monetary transfers with no direct effect on joint payoffs, whereas our punishments are costly to both parties and our rewards may increase joint payoffs. Their punishments and rewards are triggered automatically in response to specific combinations of actions chosen by the players (inspect-shirk for punishments, inspect-work

²⁷ Exceptions are Gürer et al. (2009), Heijden van der et al. (2009), O'Gorman et al. (2009), Nosenzo and Sefton (2012) and Carpenter et al. (*forthcoming*) who study settings where the ability to punish/reward group members is restricted to one player.

for rewards), whereas we study discretionary instruments. Despite these differences, both studies find that punishments effectively discourage shirking, reduce inspection rates, and lead to higher efficiencies. In contrast, differently from the present study, Nosenzo et al. (2010) find that rewards are ineffective in reducing shirking or raising efficiency.

Thus the disciplining power of punishment is robust across the two contexts whereas the effectiveness of rewards seems to be more sensitive to details of the environment. A theoretical analysis of how rewards and punishment affect behavior can be used to reconcile these findings. In Nosenzo et al. (2010) the fact that punishment is more effective than reward for discouraging shirking is consistent with the equilibrium predictions of their one-shot game. In our setup the subgame perfect equilibrium of the stage game is unaffected by the possibility of using discretionary rewards or punishments because they are costly and so should not be used by a profit-maximizing employer. Nevertheless, either punishments or rewards can discourage shirking in a repeated game. Thus, as also noted by Rand et al. (2009) in a public goods context, rewards may be more effective in repeated game environments. An interesting avenue for further research would be to examine more systematically the factors required to facilitate the effectiveness of positive incentives.

A further difference of our study from Nosenzo et al. (2010) is that we also study two additional treatments where we allow employers to use both sanctions and rewards. In these treatments we find again that only high-power incentives are effective. Somewhat differently from some of the previous studies that also examined the joint availability of sanctions and rewards (e.g. Sefton et al., 2007; Andreoni et al., 2003), we do not find that combining the instruments enhances efficiency relative to settings where only one instrument is available. Allowing (high-power) rewards and punishments reduces shirking by 15% (vs. 16% in the rewards-only treatment and 17% in the punishments-only treatment) and increases combined earnings by 10% (vs. 13% and 11% in rewards-only and punishments-only treatments). We also find that, as in the rewards-only treatments, players share the efficiency gains when both instruments are available: employer's and employee's earnings both increase by 10%. Interestingly, among the incentive treatments, we find that employers earn most when they can only punish, and least when they can use both punishments and rewards.

The finding that allowing both instruments does not allow further improvement on the availability of a single instrument is in line with findings from the principal-agent literature (Fehr and Schmidt, 2007), where contracts combining bonuses and penalties do not induce

significantly more effort than contracts that only specify bonuses. Moreover, the finding that employers earn more when they can only punish than when they can use both punishments and rewards is also similar to the finding in Fehr and Schmidt (2007) that the joint availability of bonuses and penalties reduces the principal's earnings.

In our experiment, the negative effects of making punishments and rewards jointly available seem to stem from the reduced effectiveness of the punishment instrument relative to the case when only punishments are available. This echoes the findings on the reduced effectiveness of sanctions when these are perceived as unkind or hostile by the recipient of the punishment (e.g., Fehr and Gächter, 2002; Fehr and Rockenbach, 2003; Fehr and List, 2004; Dickinson and Villeval, 2008; Houser et al., 2008). Relying on the use of punishment when rewards are also available may be perceived as particularly unkind, and may thus trigger more negative responses from employees than when sanctions are the only disciplinary instrument available to the employer.

7. CONCLUSIONS

Taken together, the findings from our paper and the related literatures suggest that both discretionary sanctions and rewards can be equally effective in encouraging compliance and influencing behavior in the direction of more socially efficient outcomes. The effectiveness of sanctions appears to be less sensitive to the details of the social and economic environment, whereas rewards can be more effective in some environments than others (e.g. in environments where players interact repeatedly). Moreover, the power of sanctions relies on the threat of punishment rather than on its use, whereas the effectiveness of rewards requires the incentive tool to be actively used. An implication of this is that the use of rewards results in a re-distribution of wealth between authorities and subordinates, whereas sanctions can be used by authorities to reap most of the benefits generated by the incentive tool.

REFERENCES

- Andreoni, J., W. Harbaugh, and L. Vesterlund. 2003. The carrot or the stick: Rewards, punishments, and cooperation. *American Economic Review* 93(3), 893–902.
- Armantier, O., and A. Boly. 2012. Framing of Incentives and Effort Provision. Unpublished manuscript, available at <https://sites.google.com/site/olivierarmantier/>.
- 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.
- Balliet, D., L.B. Mulder, and P.A.M. Van Lange. 2011. Reward, punishment, and cooperation: A meta-analysis. *Psychological Bulletin* 137(4), 594–615.
- Bigoni, M., G. Spagnolo, and P. Valbonesi. 2011. Sticks and Carrots in Procurement: an experimental exploration. Unpublished manuscript, available at http://works.bepress.com/paola_valbonesi/19.
- Carpenter, J., S. Kariv, and A. Schotter. Network architecture, cooperation and punishment in public good experiments. *Review of Economic Design*, forthcoming, 1–26.
- Dickinson, D., and M.C. Villeval. 2008. Does monitoring decrease work effort? The complementarity between agency and crowding-out theories. *Games and Economic Behavior* 63(1), 56–76.
- Drouvelis, M., and J. Jamison. 2012. Selecting Public Goods Institutions: Who Likes to Punish and Reward? Federal Reserve Bank of Boston, WP No. 12-5.
- Fehr, E., and S. Gächter. 2002. Do Incentive Contracts Undermine Voluntary Cooperation? University of Zurich Working Paper No. 34.
- Fehr, E., A. Klein, and K.M. Schmidt. 2007. Fairness and contract design. *Econometrica* 75(1), 121–154.
- Fehr, E., and J. List. 2004. The Hidden Costs and Returns of Incentives - Trust and Trustworthiness Among CEOs. *Journal of the European Economic Association* 2(5), 743–771.
- 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.
- 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.
- 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.
- Güererk, Ö., B. Irlenbusch, and B. Rockenbach. 2006. The competitive advantage of sanctioning institutions. *Science* 312(108), 108–111.
- Güererk, 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.
- Hannan, L.R., V. Hoffman, and D. Moser. 2005. Bonus versus Penalty: Does Contract Frame Affect Employee Effort? In *Experimental Business Research: Economic and*

- Managerial Perspectives*, ed. Amnon Rapoport and Rami Zwick, 2:151–169. Holland: Springer.
- Heijden, van der E., J. Potters, and M. Sefton. 2009. Hierarchy and opportunism in teams. *Journal of Economic Behavior & Organization* 69(1), 39–50.
- Hossain, T., and J. List. The Behavioralist Visits the Factory: Increasing Productivity Using Simple Framing Manipulations. *Management Science*, forthcoming.
- Houser, D., E. Xiao, K. McCabe, and V. Smith. 2008. When punishment fails: Research on sanctions, intentions and non-cooperation. *Games and Economic Behavior* 62(2), 509–532.
- Kaiser, J., and M.G. Lacy. 2009. A general-purpose method for two-group randomization tests. *Stata Journal* 9(1), 70–85.
- Milinski, M., and B. Rockenbach. 2012. On the interaction of the stick and the carrot in social dilemmas. *Journal of Theoretical Biology* 299, 139–143.
- Nosenzo, D., T. Offerman, M. Sefton, and A. van der Veen. 2010. Inducing Good Behavior: Bonuses versus Fines in Inspection Games. CeDEx Discussion Paper 2010-21, University of Nottingham.
- Nosenzo, D., and M. Sefton. 2012. Promoting Cooperation: The distribution of reward and punishment power. CeDEx Discussion Paper 2012.
- O’Gorman, R., J. Henrich, and M. van Vugt. 2009. Constraining free riding in public goods games: designated solitary punishers can sustain human cooperation. *Proceedings of the Royal Society B-Biological Sciences* 276(1655), 323–329.
- 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.
- 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.
- 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.

APPENDICES

Appendix A

A1: Sustaining (Not Inspect, Work) in the game augmented by punishments.

Simple Nash-reversion strategies cannot sustain (Not Inspect, Work) in a subgame perfect equilibrium of the indefinitely repeated game. Note, in fact, that the Employee receives the same payoff ($w-c$) in the mixed-strategy stage-game Nash equilibrium and in the (Not Inspect, Work) outcome. Thus, the Employee will not be deterred from shirking by the threat of reversion to the stage-game Nash equilibrium outcome. In order to deter the Employee from shirking, the Employer needs to be able to (credibly) threaten the Employee with a payoff lower than $w - c$. This may be achieved when the Employer can punish the Employee. If punishments are available, then for some (α, δ) parameters values the joint profit maximizing outcome can be sustained in a subgame perfect equilibrium using the type of “stick and carrot” strategies introduced in Abreu (1986; “Extremal equilibria of oligopolistic supergames”, *Journal of Economic Theory* 39, 191-225). A player’s strategy specifies actions for each stage depending on whether the game is in a “cooperative phase” or “disciplinary phase”. The game begins in a cooperative phase and remains in the cooperative phase unless a player deviates from the cooperative action. A deviation from the cooperative action triggers the disciplinary phase. After the employer has carried out T consecutive stages of his disciplinary actions the game returns to the cooperative phase.

The stick and carrot strategies are:

- Cooperative phase: Employee chooses Work
Employer chooses Not Inspect.
- Disciplinary phase: Employee chooses Shirk
Employer chooses Inspect, assigns 5 punishment points if Work,
assigns $5 - 5/\alpha$ punishment points if Shirk.

Before we check the conditions under which these form a subgame perfect equilibrium we note several features of these strategies.

i) On the equilibrium path (Not Inspect, Work) is chosen and stage payoffs are (45, 20) in every stage.

ii) In the disciplinary stage if the employer follows the stick and carrot strategy the employee receives a stage payoff of $20 - 5\alpha$ regardless of whether she works or shirks, and the employer receives a stage payoff of 25 if the employee works and $5 + 5/\alpha$ if the employee shirks.

iii) the employee shirks in the disciplinary phase; if the employee works in the disciplinary phase the employer never has an incentive to inspect and the threatened punishment is not credible.

iv) if the transition from disciplinary to cooperative phase were to require both players to use their disciplinary strategies for T consecutive periods it would not be credible that an employer would punish an employee who chooses work in the disciplinary phase. By choosing work the employee would delay reversion to the cooperative phase regardless of the employer's response and so the employer would have no incentive to punish.

v) Because the transition from disciplinary to cooperative phase is governed by the employer's use of the disciplinary actions the employee will only find it in her interest to shirk in the disciplinary phase if her stage payoff from shirking is as high as her stage payoff from working. Thus the most the employer can punish the employee when she shirks is the amount that reduces the employee's stage payoff to the same level that she would get from choosing work. This means that the employee's stage payoff in the disciplinary phase cannot be reduced to less than $20 - 5\alpha$. This is the level attained when the employee chooses Work and receives 5 punishment points (i.e. is maximally punished).

To check that these strategies form a subgame perfect equilibrium we must check that there is no incentive to deviate in any subgame. In each case considered below we assume that one player follows the stick and carrot strategy and we compare the other player's repeated game payoff from following the stick and carrot strategy with that from delaying following it for one stage. Due to the stationary nature of the strategies, if it does not pay to delay for one stage it will not pay to delay for more than one stage.

1. subgames beginning in the cooperative phase.

The employer clearly has no incentive to deviate from the stick and carrot strategy.

The employee's payoff from following the stick and carrot strategy is $20/(1 - \delta)$. By deviating she gets $35 + (20 - 5\alpha)(\delta + \dots + \delta^T) + \delta^{T+1}20/(1 - \delta)$. The employee has no incentive to deviate if $20/(1 - \delta) \geq 35 + (20 - 5\alpha)(\delta + \dots + \delta^T) + \delta^{T+1}20/(1 - \delta)$, or

$$\delta(1 - \delta^T)/(1 - \delta) \geq 3/\alpha \quad (1).$$

2. subgames beginning after (Inspect, Work) in a disciplinary phase.

Suppose we are in the first stage of a disciplinary phase. To comply with the stick and carrot strategy the employer inflicts maximal punishment in this stage, giving him a stage payoff of 25. In the next $T - 1$ stages the employer receives $5 + 5/\alpha$ per stage. The game then reverts to the cooperative phase. Thus the employer's payoff from complying is $25 + (5 + 5/\alpha)(\delta + \dots + \delta^{T-1}) + \delta^T 45/(1 - \delta)$. By not punishing the employer gets $30 + (5 + 5/\alpha)(\delta + \dots + \delta^T) + \delta^{T+1} 45/(1 - \delta)$. The employer has no incentive to deviate if

$$25 + (5 + 5/\alpha)(\delta + \dots + \delta^{T-1}) + \delta^T 45/(1 - \delta) \geq 30 + (5 + 5/\alpha)(\delta + \dots + \delta^T) + \delta^{T+1} 45/(1 - \delta).$$

If this condition holds the gain of 5 payoff units reaped by not punishing in the current stage is outweighed by the discounted sum of losses from delaying reversion to the cooperative phase by one stage. This condition simplifies to

$$24 - 3/\delta^T \geq 3/\alpha \quad (2)$$

Analogously, if we are in the k^{th} stage of a disciplinary phase the employer will have no incentive to deviate from the stick and carrot strategy if $24 - 3/\delta^{T+1-k} \geq 3/\alpha$. Clearly, if condition (2) is satisfied then so is this condition.

3. subgames beginning after (Inspect, Shirk) in a disciplinary phase.

Suppose we are in the first stage of a disciplinary phase. Complying with the stick and carrot strategy gives the employer a payoff of $(5 + 5/\alpha)(1 + \dots + \delta^{T-1}) + \delta^T 45/(1 - \delta)$. By not punishing the employer gets $10 + (5 + 5/\alpha)(\delta + \dots + \delta^T) + \delta^{T+1} 45/(1 - \delta)$. Here the gain from not punishing in the current period is $10 - 5 - 5/\alpha < 5$ and as long as condition (2) holds this is again outweighed by the sum of discounted losses from delaying reversion to the cooperative phase. Similarly, if we are in the k^{th} stage of a disciplinary phase then if condition (2) is met this implies that the employer has no incentive to deviate from the stick and carrot strategy.

4. subgames beginning at the start of a stage in the disciplinary phase

The employee's action does not affect her stage payoff (which is $20 - 5\alpha$ whether she works or shirks) and does not affect her future stage payoffs (since her action does not affect the duration of the disciplinary phase). Thus she has no incentive to deviate from the stick and carrot strategy.

Given that the employee shirks in the disciplinary stage the employer's stage payoff from inspecting and punishing is at least as high as that from not inspecting, and so the employer has no incentive to deviate from the stick and carrot strategy.

Thus the stick and carrot strategies form a subgame perfect equilibrium when conditions (1) and (2) are both met. Combining these conditions we see that the stick and carrot strategies form a subgame perfect equilibrium if $3/\alpha \leq \min\{24 - 3/\delta^T, \delta(1 - \delta^T)/(1 - \delta)\}$.

Note that for given δ the length of the disciplinary phase can be set to make the right hand side as large as possible. The condition then divides the α -parameter space into values where the efficient outcome can or cannot be sustained as a subgame perfect equilibrium. For a given δ sustaining the efficient outcome in a subgame perfect equilibrium requires a sufficiently high α . Similarly, since both arguments of the $\min\{.,.\}$ function are increasing in δ , for a given α a sufficiently high δ is required to sustain efficiency in a subgame perfect equilibrium.

A.2 Sustaining (Inspect, Work, maximal rewards) in the game augmented by rewards.

For some values of (α, δ) the following strategies form a subgame perfect equilibrium:

- Cooperative phase: Employee chooses Work,
Employer chooses Inspect, assign 5 reward points if Work,
assign 0 reward points if Shirk
- Disciplinary phase: Employee chooses Shirk,
Employer chooses Inspect, assigns 0 reward points

1. subgames beginning in the cooperative phase

Given that the employer follows the stick and carrot strategy the employee maximizes her stage game payoff by working, and so there is no incentive to deviate.

Given that the employee follows the stick and carrot strategy the employer must have no incentive to not reward after inspecting. For it to be optimal for the employer to reward an employee who works we require $25/(1 - \delta) \geq 30 + 10(\delta + \dots + \delta^T) + \delta^{T+1} 25/(1 - \delta)$. Also we require that the employer must have no incentive to not inspect. For it to be optimal to inspect we require $25/(1 - \delta) \geq 45 + 10(\delta + \dots + \delta^T) + \delta^{T+1} 25/(1 - \delta)$. Clearly the latter condition provides the sterner test. This condition simplifies to

$$\delta(1 - \delta^T)/(1 - \delta) \geq 4/3. \quad (3)$$

2. subgames beginning in the disciplinary phase

The employer maximizes her stage game payoff by complying with the stick and carrot strategy and so there is no incentive to deviate for the employer.

Suppose we are in the first stage of a disciplinary phase. If the employee complies with the stick and carrot strategy she receives $15(1 + \delta + \dots + \delta^{T-1}) + \delta^T(20 + 5\alpha)/(1 - \delta)$. If she works and delays compliance for one stage the employee can obtain $20 + 15(\delta + \dots + \delta^T) + \delta^{T+1}(20 + 5\alpha)/(1 - \delta)$. There is no incentive to delay if

$$\delta^T \geq 1/(1 + \alpha). \quad (4)$$

As long as $\delta > 4/7$ condition (3) can be met for sufficiently high T . Let $\delta > 4/7$ be given and suppose T^* is chosen to just satisfy (3). Then the cooperative outcome can be sustained if and only if $\delta^{T^*} \geq 1/(1 + \alpha)$, or $\alpha \geq (1 - \delta^{T^*})/\delta^{T^*}$. Thus for a given δ we require α to be sufficiently high. Similarly, for a given α conditions (3) and (4) divide the δ -parameter space into values that are too low to sustain and sufficiently high to sustain the cooperative outcome as a subgame perfect equilibrium.

Appendix B

Baseline treatment

Instructions

Introduction

This is an experiment about decision-making. There are other people in the room who are also participating in this experiment. You must not communicate with any other participant in any way during the experiment. If you have a question at any time, raise your hand and a monitor will come to your desk to answer it. The experiment consists of a number of rounds, in each of which you earn points. At the end of the experiment you will be paid, in private and in cash, according to the sum of your total point earnings from all rounds at a rate of 0.7 pence per point.

Structure of the experiment

At the beginning of the first round you will be randomly paired with another participant, and you will be paired with this same person in every round throughout the experiment. One of you will be randomly assigned the role of “Employer” and the other will be assigned the role of “Worker”. At the beginning of the first round you will be informed of your role, and you will keep this role in every round throughout the experiment.

The experiment will consist of at least 70 rounds. After round 70 the computer will randomly determine whether the experiment ends or continues. There will be a 20% chance that the experiment ends, and a 80% chance the experiment continues. If the experiment continues, then at the end of round 71 the computer will again randomly determine whether the experiment ends or continues. Again, there will be a 20% chance that the experiment ends, and a 80% chance the experiment continues. This process will repeat until the experiment ends. Thus no participant will know in advance how many rounds the experiment will consist of, or which round will be the last.

Structure of a round

At the beginning of a round the Employer chooses either INSPECT or NOT INSPECT. At the same time the Worker chooses either HIGH or LOW effort. Point earnings depend on choices as described in the table below:

	HIGH	LOW
INSPECT	Employer earns 30 Worker earns 20	Employer earns 10 Worker earns 15
NOT INSPECT	Employer earns 45 Worker earns 20	Employer earns 5 Worker earns 35

For example, if the Employer chooses NOT INSPECT and the Worker chooses LOW the Employer earns 5 points and the Worker earns 35 points. As another example, if the Employer chooses INSPECT and the Worker chooses HIGH the Employer earns 30 points and the Worker earns 20 points.

The computer will then inform you of the choices made by you and the person you are paired with, and your point earnings for the round.

On your screen you will also see in which round you are, your role, your total point earnings so far, and a table summarizing the decisions and earnings made in previous rounds by you and the person you are paired with.

Ending the session

After the last round your total points from all rounds will be converted to cash at a rate of 0.7 pence per point and you will be paid this amount in private and in cash.

Quiz

Before the decision making part of the experiment begins we want to be sure everyone understands the instructions. Please complete the questions below. In a couple of minutes someone will come to your desk to check the answers. (The decisions and earnings used for the questions below are simply for illustrative purposes. In the experiment decisions and earnings will depend on the actual choices of the participants.)

If you have any questions please raise your hand and a monitor will come to your desk to answer it.

1. Will you be matched with the same person during the whole experiment? _____
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. How many points will you earn in a round if you are an Employer, choose INSPECT, and the Worker you are matched with chooses LOW _____
5. How many points will you earn in a round if you are a Worker, choose LOW, and the Employer you are matched with chooses INSPECT _____
6. The experiment will last at least ____ rounds.
7. Suppose the experiment has reached round 83. How likely is it that the experiment will continue to round 84? a) Impossible; b) 20% chance; c) 80% chance; d) 100% chance;

R1:3 treatment

Instructions

Introduction

This is an experiment about decision-making. There are other people in the room who are also participating in this experiment. You must not communicate with any other participant in any way during the experiment. If you have a question at any time, raise your hand and a monitor will come to your desk to answer it. The experiment consists of a number of rounds, in each of which you earn points. At the end of the experiment you will be paid, in private and in cash, according to the sum of your total point earnings from all rounds at a rate of 0.7 pence per point.

Structure of the experiment

At the beginning of the first round you will be randomly paired with another participant, and you will be paired with this same person in every round throughout the experiment. One of you will be randomly assigned the role of “Employer” and the other will be assigned the role of “Worker”. At the beginning of the first round you will be informed of your role, and you will keep this role in every round throughout the experiment.

The experiment will consist of at least 70 rounds. After round 70 the computer will randomly determine whether the experiment ends or continues. There will be a 20% chance that the experiment ends, and a 80% chance the experiment continues. If the experiment continues, then at the end of round 71 the computer will again randomly determine whether the experiment ends or continues. Again, there will be a 20% chance that the experiment ends, and a 80% chance the experiment continues. This process will repeat until the experiment ends. Thus no participant will know in advance how many rounds the experiment will consist of, or which round will be the last.

Structure of a round

At the beginning of a round the Employer chooses either INSPECT or NOT INSPECT. At the same time the Worker chooses either HIGH or LOW effort. Point earnings depend on choices as described in the table below:

	HIGH	LOW
INSPECT	Employer earns 30 Worker earns 20	Employer earns 10 Worker earns 15
NOT INSPECT	Employer earns 45 Worker earns 20	Employer earns 5 Worker earns 35

For example, if the Employer chooses NOT INSPECT and the Worker chooses LOW the Employer earns 5 points and the Worker earns 35 points. As another example, if the Employer chooses INSPECT and the Worker chooses HIGH the Employer earns 30 points and the Worker earns 20 points.

If the Employer chooses NOT INSPECT, the round will end immediately and the computer will inform you of the choices made by you and the person you are paired with, and your point earnings for the round.

If the Employer chooses INSPECT, then the round will have a second stage. In stage two the Employer is informed of the choice of the Worker (HIGH or LOW). Then the Employer chooses between NO ACTION and REWARD. If the Employer chooses NO ACTION earnings for the round are unchanged. If the Employer chooses REWARD, he or she has to decide how many tokens, from zero to five inclusive, to assign to the Worker. Each token assigned reduces the Employer's earnings by 1 point and increases the Worker's earnings by 3 points. At the end of stage two the computer will inform you of all the choices made by you and the person you are paired with, and your point earnings for the entire round.

On your screen you will also see in which round you are, your role, your total point earnings so far, and a table summarizing the decisions and earnings made in previous rounds by you and the person you are paired with.

Ending the session

After the last round your total points from all rounds will be converted to cash at a rate of 0.7 pence per point and you will be paid this amount in private and in cash.

Quiz

Before the decision making part of the experiment begins we want to be sure everyone understands the instructions. Please complete the questions below. In a couple of minutes someone will come to your desk to check the answers. (The decisions and earnings used for the questions below are simply for illustrative purposes. In the experiment decisions and earnings will depend on the actual choices of the participants.)

If you have any questions please raise your hand and a monitor will come to your desk to answer it.

1. Will you be matched with the same person during the whole experiment? _____
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. How many points will you earn in a round if you are an Employer, choose INSPECT, and the Worker you are matched with chooses LOW, and
 - A) In stage two you then choose NO ACTION? _____
 - B) In stage two you then choose to REWARD the Worker with zero tokens? _____
 - C) In stage two you then choose to REWARD the Worker with five tokens? _____
5. How many points will you earn in a round if you are a Worker, choose LOW, and the Employer you are matched with chooses INSPECT, and
 - A) In stage two the Employer then chooses NO ACTION? _____
 - B) In stage two the Employer then chooses to REWARD you with zero tokens? _____
 - C) In stage two the Employer then chooses to REWARD you with five tokens? _____

6. How many stages will each round consist of?
a) One; b) Two; c) Two if the Employer chooses INSPECT; d) Two if the Employer chooses NOT INSPECT;
7. The experiment will last at least ____ rounds.
8. Suppose the experiment has reached round 83. How likely is it that the experiment will continue to round 84? a) Impossible; b) 20% chance; c) 80% chance; d) 100% chance;

P1:3 treatment

Instructions

Introduction

This is an experiment about decision-making. There are other people in the room who are also participating in this experiment. You must not communicate with any other participant in any way during the experiment. If you have a question at any time, raise your hand and a monitor will come to your desk to answer it. The experiment consists of a number of rounds, in each of which you earn points. At the end of the experiment you will be paid, in private and in cash, according to the sum of your total point earnings from all rounds at a rate of 0.7 pence per point.

Structure of the experiment

At the beginning of the first round you will be randomly paired with another participant, and you will be paired with this same person in every round throughout the experiment. One of you will be randomly assigned the role of “Employer” and the other will be assigned the role of “Worker”. At the beginning of the first round you will be informed of your role, and you will keep this role in every round throughout the experiment.

The experiment will consist of at least 70 rounds. After round 70 the computer will randomly determine whether the experiment ends or continues. There will be a 20% chance that the experiment ends, and a 80% chance the experiment continues. If the experiment continues, then at the end of round 71 the computer will again randomly determine whether the experiment ends or continues. Again, there will be a 20% chance that the experiment ends, and a 80% chance the experiment continues. This process will repeat until the experiment ends. Thus no participant will know in advance how many rounds the experiment will consist of, or which round will be the last.

Structure of a round

At the beginning of a round the Employer chooses either INSPECT or NOT INSPECT. At the same time the Worker chooses either HIGH or LOW effort. Point earnings depend on choices as described in the table below:

	HIGH	LOW
INSPECT	Employer earns 30 Worker earns 20	Employer earns 10 Worker earns 15
NOT INSPECT	Employer earns 45 Worker earns 20	Employer earns 5 Worker earns 35

For example, if the Employer chooses NOT INSPECT and the Worker chooses LOW the Employer earns 5 points and the Worker earns 35 points. As another example, if the Employer chooses INSPECT and the Worker chooses HIGH the Employer earns 30 points and the Worker earns 20 points.

If the Employer chooses NOT INSPECT, the round will end immediately and the computer will inform you of the choices made by you and the person you are paired with, and your point earnings for the round.

If the Employer chooses INSPECT, then the round will have a second stage. In stage two the Employer is informed of the choice of the Worker (HIGH or LOW). Then the Employer chooses between NO ACTION and PUNISH. If the Employer chooses NO ACTION earnings for the round are unchanged. If the Employer chooses PUNISH, he or she has to decide how many tokens, from zero to five inclusive, to assign to the Worker. Each token assigned reduces the Employer's earnings by 1 point and reduces the Worker's earnings by 3 points. At the end of stage two the computer will inform you of all the choices made by you and the person you are paired with, and your point earnings for the entire round.

On your screen you will also see in which round you are, your role, your total point earnings so far, and a table summarizing the decisions and earnings made in previous rounds by you and the person you are paired with.

Ending the session

After the last round your total points from all rounds will be converted to cash at a rate of 0.7 pence per point and you will be paid this amount in private and in cash.

Quiz

Before the decision making part of the experiment begins we want to be sure everyone understands the instructions. Please complete the questions below. In a couple of minutes someone will come to your desk to check the answers. (The decisions and earnings used for the questions below are simply for illustrative purposes. In the experiment decisions and earnings will depend on the actual choices of the participants.)

If you have any questions please raise your hand and a monitor will come to your desk to answer it.

1. Will you be matched with the same person during the whole experiment? _____
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. How many points will you earn in a round if you are an Employer, choose INSPECT, and the Worker you are matched with chooses LOW, and
 - A) In stage two you then choose NO ACTION? _____
 - B) In stage two you then choose to PUNISH the Worker with zero tokens? _____
 - C) In stage two you then choose to PUNISH the Worker with five tokens? _____
5. How many points will you earn in a round if you are a Worker, choose LOW, and the Employer you are matched with chooses INSPECT, and

A) In stage two the Employer then chooses NO ACTION? _____

B) In stage two the Employer then chooses to PUNISH you with zero tokens? _____

C) In stage two the Employer then chooses to PUNISH you with five tokens? _____

6. How many stages will each round consist of?

a) One; b) Two; c) Two if the Employer chooses INSPECT; d) Two if the Employer chooses NOT INSPECT;

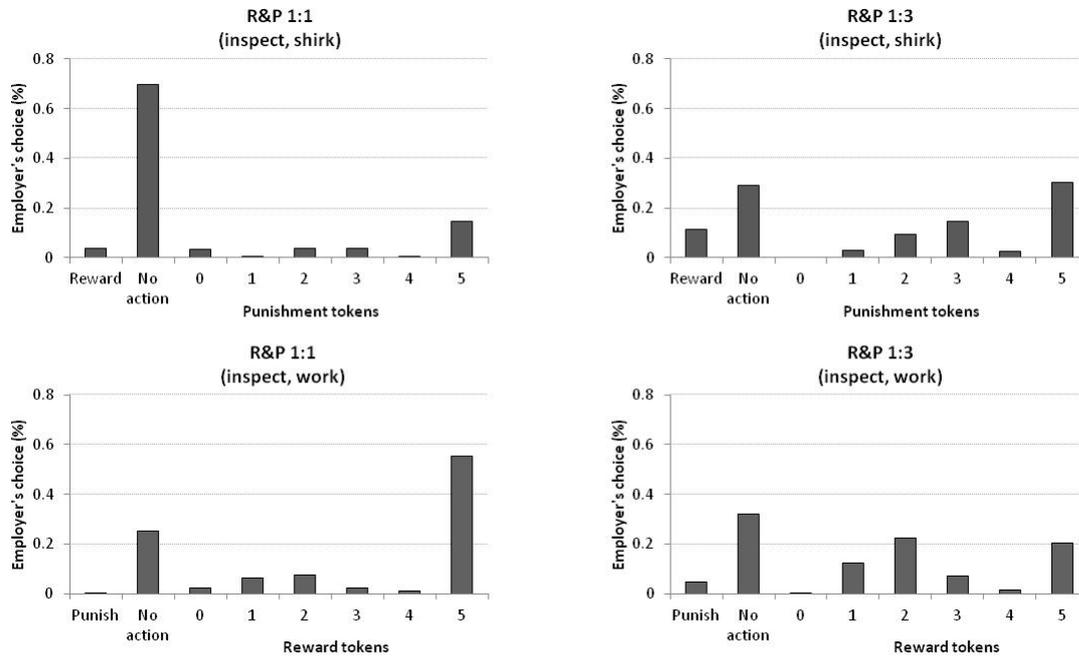
7. The experiment will last at least _____ rounds.

8. Suppose the experiment has reached round 83. How likely is it that the experiment will continue to round 84?

a) Impossible; b) 20% chance; c) 80% chance; d) 100% chance;

Appendix C

Figure A.1: Employer's decisions after (Inspect, Shirk) and (Inspect, Work) in the R&P treatments



Notes: based on 864 games in R&P1:1, 831 games in R&P1:3.