



UvA-DARE (Digital Academic Repository)

Inducing good behavior

van der Veen, A.

[Link to publication](#)

Citation for published version (APA):
van der Veen, A. (2012). *Inducing good behavior*.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

4. How to Prevent Workers from Shirking: the Use and Effectiveness of Rewards and Punishments in the Inspection Game¹

4.1. Introduction

In the labor market, employers usually want workers to perform in a way that, left to themselves, they would not do. In many situations, workers will only deliver the desired performance level if there is a serious possibility that their work is inspected by the employer. Monitoring a worker is costly to the employer, though, and the employer would prefer not to do so if he were sufficiently sure that the worker would work hard. The essence of the interaction in such situations is described in the inspection game. In this game, the employer chooses to inspect or not, and the worker chooses to provide low or high effort. In every situation one of the players prefers to have chosen a different action. Basically, the inspection game is an asymmetric matching pennies game and the unique equilibrium is in mixed strategies.

To further encourage good behavior, after inspection the employer may consider punishing a worker who was found providing low effort or rewarding a worker who was found providing high effort. In this chapter, we investigate experimentally whether employers use rewards or punishments to incentivize their workers, and we compare the effectiveness of the two possibilities. Whether rewards for good behavior or punishments for bad behavior are more effective in preventing shirking is still an open question. Folk wisdom suggests that rewards may be more effective. As Benjamin Franklin (1744), one of America's founding fathers, put it: "... a spoonful of honey will catch more flies than (a) Gallon of vinegar". This folk wisdom is backed up by a strand of literature in psychology started by Skinner (1965). From his studies on animals, he concluded that rewards dominate punishments as punishments lose their effectiveness in the long term. In agreement with this conclusion, psychologists have reported that supervisors rewarding good behavior are more successful in encouraging subordinates to work hard than supervisors punishing bad behavior (Sims, 1980; Podsakoff, Bommer, Podsakoff, and MacKenzie, 2006; George, 1995).

¹This chapter is based on the identically titled paper joint with Daniele Nosenzo, Theo Offerman, and Martin Sefton. We are grateful to CREED programmer Jos Theelen for programming the experiment.

Typically, these studies draw their conclusions on the basis of questionnaires for employers and employees. This complicates the interpretation of the results because it is a priori not clear that rewards and punishments cause worker’s behavior or vice versa.

Controlled laboratory experiments investigating the strength of positive and negative reciprocity have been run, but not in the context of the inspection game. Previous studies consistently found relatively strong evidence for negative reciprocity and weak (or no) evidence for positive reciprocity (Abbink, Irlenbusch, and Renner, 2000; Brandts and Sola, 2001; Charness and Rabin, 2002; Offerman, 2002; Brandts and Charness, 2004; Falk, Fehr, and Fischbacher, 2003; Charness, 2004; Al-Ubaydli and Lee, 2009). The weak evidence for positive reciprocity casts doubt on the effectiveness of rewards in employer/worker relations. Ex ante it is hard to say what should be inferred from the stronger evidence for negative reciprocity for the case of the inspection game. On the one hand, employers using punishments may trigger a negative spiral of ongoing shirking and punishments, so that punishments may even have a counterproductive effect. On the other hand, workers may fear the possibility of punishment and work hard simply to avoid them. This would happen if the findings in the ultimatum game generalize to the inspection game. In the ultimatum game, proposers tend to behave well and propose fair offers to avoid the rejection (punishment) by responders (for a meta-study of ultimatum game experiments, see Oosterbeek, Sloof, and van de Kuilen, 2004). So evidence collected in controlled laboratory experiments in different environments is also rather inconclusive.²

We collect controlled evidence on the use and effectiveness of rewards and punishments in the inspection game in a $(1 + 3 \times 2)$ design. In all treatments, pairs are formed that consist of a worker and an employer interacting repeatedly for an indeterminate length of time. In the baseline treatment, subjects do not have the possibility to reward or punish, and they only interact through the inspection game. In the other treatments, two treatment variables are introduced. The first one is the tool to incentivize workers, which takes the form of (i) reward only, (ii) punish only, or (iii) reward and punish. The second treatment variable concerns the effectiveness of the tool itself, which is either low or high.³ With the low ratio, each reward or

²Our study also contributes to investigations of rewards and punishments in other applications. Andreoni, Harbaugh, and Vesterlund (2003) study the effects of rewards and punishments in a bargaining game where the proposer chooses an amount to transfer to the responder and the responder can then either punish or reward the proposer. They find that proposers’ transfers are particularly responsive to the threat of punishment, although rewards have a positive effect. Sefton, Shupp, and Walker (2007) examine the effect of rewards and punishments on contributions in a repeated public good game and find that punishments help sustaining higher cooperation levels in comparison to a baseline without reward/punishment opportunities, whereas the possibility of rewards has only a transient effect.

³In other settings, the effectiveness of rewards and punishments appears to depend on the rewarding/punishing technology. Sutter, Haigner, and Kocher (2010) obtain the result that when the benefit/cost of receiving reward/punishment is three times the cost of delivering it (i.e. with a 3:1 technology), both mechanisms are effective in encouraging contributions. Likewise, Rand, Dreber, Ellingsen, Fudenberg, and Nowak (2009) find that rewards are equally effective as punishments in sustaining cooperation in a repeated public good game with unknown time horizon and with a 3:1 reward/punishment technology. Gürer, Irlenbusch, and Rockenbach (2006) study a public good game with a 1:1 rewarding mechanism and a 3:1 punishment mechanism technology and find that only the latter affect contributions. Gürer, Irlenbusch, and Rockenbach (2009) study a public good game where one group member (the ‘leader’) can reward or punish the other contributors. Although both rewarding and punishment mechanisms employ a 3:1 technology, they find that punishments are more effective.

punishment point assigned by the employer yields or costs the worker one point and with the high ratio, each assigned reward or punishment point yields or costs the worker three points.

We obtain the following results. Like in public good games, the possibility to reward and/or punish has rather small effects on the interaction between employers and workers with the low ratio. With the high ratio, the following pattern emerges in our data. When employers can either only punish or only reward, workers shirk substantially less often than in the baseline game. The reduction in shirking behavior is approximately equally large with the two tools. With punishments, it is achieved with fewer inspections than with rewards. Therefore, employers are better off with punishments than with rewards. However, when employers have the possibility to use the two tools simultaneously, subjects still tend to employ the reward tool more often. This surprising result can be explained in the following way. When employers can use both tools simultaneously, punishments seem to be relatively less effective than in the case where only punishments are allowed, while rewards do not lose their effectiveness. Results from a questionnaire suggest that our subjects find rewards the more appropriate tool to incentivize workers. Thus, when both tools are available, employers can no longer hide behind the excuse that punishments provided the only way to get the workers to work hard. So there may be two factors contributing to the effect. On the one hand, workers seem to resist punishments when both rewards and punishments are possible, and on the other hand, employers prefer to make use of rewards instead of punishments. As a result, employers do not prefer the use of punishments when both tools are allowed.

This chapter is organized in the following way. Section 4.2 describes the game and provides the standard theoretical benchmark based on selfish rational players. Section 4.3 presents the experimental design. Section 4.4 presents the experimental results and Section 4.5 concludes.

4.2. Inspection Game and Theoretical Benchmark

The inspection game involves two players and simultaneous moves. The employer chooses between inspect and not inspect, and the worker shirks or works. In the standard version of the game (see, e.g., Fudenberg and Tirole, 1992, p. 17), the employer incurs a cost of h from inspecting. If the worker provides high effort, the worker incurs a cost of c and the employer receives a revenue of v . If the employer does not inspect, the worker always receives a wage of w . If the employer inspects, the worker receives nothing when she shirks and she receives the wage when she works. The resulting payoffs are shown in the left panel of Figure 4.1 on the next page. We assume that all variables are positive and $v > c$, $w > h$, $w > c$. Note that joint payoffs are maximized when the worker supplies high effort and the employer does not inspect. The right panel presents the payoffs that we used in the experiment.⁴

⁴This means that in the experiment, we used the parameters $v = 40$, $w = 20$, $c = 15$ and $h = 15$. We added 15 to each of the worker's potential payoffs and 25 to each of the employer's possible payoffs because we wanted to prevent negative outcomes (which are problematic to implement in an experiment) and because we wanted the expected earnings in equilibrium not to differ too much between the two types of players.

Figure 4.1.: Inspection Game

		Canonical Game				Game used in Experiment					
		Work	Shirk			Work	Shirk				
Inspect	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="border: 1px solid black; padding: 5px;">$v - w - h$</td> <td style="border: 1px solid black; padding: 5px;">$-h$</td> </tr> <tr> <td style="border: 1px solid black; padding: 5px;">$w - c$</td> <td style="border: 1px solid black; padding: 5px;">0</td> </tr> </table>	$v - w - h$	$-h$	$w - c$	0	Inspect	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="border: 1px solid black; padding: 5px;">30</td> <td style="border: 1px solid black; padding: 5px;">10</td> </tr> <tr> <td style="border: 1px solid black; padding: 5px;">20</td> <td style="border: 1px solid black; padding: 5px;">15</td> </tr> </table>	30	10	20	15
$v - w - h$	$-h$										
$w - c$	0										
30	10										
20	15										
Not inspect	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="border: 1px solid black; padding: 5px;">$v - w$</td> <td style="border: 1px solid black; padding: 5px;">$-w$</td> </tr> <tr> <td style="border: 1px solid black; padding: 5px;">$w - c$</td> <td style="border: 1px solid black; padding: 5px;">w</td> </tr> </table>	$v - w$	$-w$	$w - c$	w	Not inspect	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="border: 1px solid black; padding: 5px;">45</td> <td style="border: 1px solid black; padding: 5px;">5</td> </tr> <tr> <td style="border: 1px solid black; padding: 5px;">20</td> <td style="border: 1px solid black; padding: 5px;">35</td> </tr> </table>	45	5	20	35
$v - w$	$-w$										
$w - c$	w										
45	5										
20	35										

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

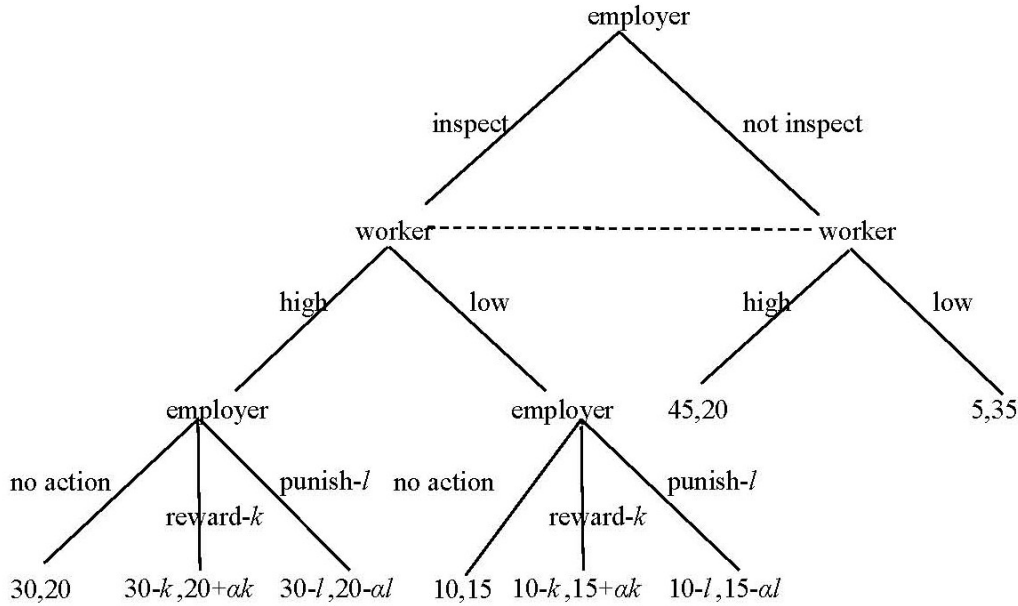
Let p denote the probability of inspection and q denote the probability of shirking. In the unique Nash equilibrium, the probabilities p and q are determined endogenously and must leave the players indifferent between actions. Thus, in equilibrium the employer inspects with probability $p_c = c/w$ and the worker chooses to shirk with probability $q_c = h/w$. The employer receives an expected payoff of $\pi_c^{employer} = v - w - hv/w$, the worker receives an expected payoff of $\pi_c^{worker} = w - c$, and joint payoffs are $\pi_c = v - c - hv/w$. In the version of the game used in the experiment, the employer inspects with probability $p = 3/4$ and the worker shirks with probability $q = 3/4$, and the employer's expected payoff equals 15 while the worker's expected payoff equals 20. The inspection game is the stage game in the baseline treatment.

In the games where we allow for punishments and rewards, the stage game of the baseline treatment is augmented in the following way. If the employer inspects, he observes the worker's choice to shirk or work, and then chooses between 'No action', 'Punish' and 'Reward'. If he chooses No action, then the payoffs are simply determined by the payoffs of the Inspection game. If he chooses Reward, he must assign the reward level k from the set 0, 1, 2, 3, 4, 5 and the employer's payoff from the inspection game is diminished by k while the worker's payoff is increased by αk . If he chooses Punish, he sets the punishment level l from the same set 0, 1, 2, 3, 4, 5 and the employer's payoff from the inspection game is diminished by l while the worker's payoff is decreased by αl . With the low ratio $\alpha = 1$ and with the high ratio $\alpha = 3$. Figure 4.2 on the facing page presents the augmented game graphically. In the games where we allow for reward only, the punishment option is chopped off from the game in Figure 4.2 and in the games where we allow for punishment only, the reward option is eliminated.

The subgame perfect equilibrium outcome of the augmented game is identified by backward induction. After inspection, a selfish and rational employer will either choose No action or choose free punishment ($k = 0$) or free reward ($l = 0$). This behavior is anticipated by the worker and the employer, and as a result, play in the phase preceding the final phase remains unaffected. Thus, in the subgame perfect equilibrium outcome subjects mix between their actions Inspect and Not inspect and actions Work and Shirk in precisely the same way as in the baseline treatment, i.e., $p = 3/4$ and $q = 3/4$.⁵

⁵The stage game does not have Nash equilibria where the employer uses positive reward or punishment levels. The employer can only use incredible punishments $l > 0$ if he never has to carry out the incredible threat.

Figure 4.2.: Inspection Game and the Possibility to Reward and Punish



In the actual labor market as well as in our experiment, employers and workers are engaged in a repeated interaction. Here, we consider the case where in each stage the game described above is played and where players' earnings are simply the sum of the earnings in 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. In such a setup, it is well-known that a continuum of outcomes can be supported in equilibrium when the continuation probability is sufficiently large. In particular, the cooperative outcome (Not inspect, Work) can be supported in equilibrium by threatening to set the other player on her minimax payoff if she ever deviates from the equilibrium path.

Instead of pursuing a full analysis of the repeated game (which is impossible because the number of possibilities explodes), we provide an intuitive argument for why it is easier to support cooperation in the versions of the game where punishments are allowed. In Figure 4.3 on page 47, we display in gray the pairs of (p, q) that correspond to equilibria where the players play according to a 'normal stationary stage game strategy' in each stage game, unless one of them deviates, in which case the deviating player is set on her minimax payoff forever. We assume that in the normal stationary stage game strategy, subjects mix with constant probabilities (p, q) , and after inspection employers punish a worker maximally if they find the worker shirking and if

This can only be accomplished if (i) he never inspects or (ii) he inspects with positive probability and the worker always works. In (i), the worker will want to shirk with $q = 1$, in which case the employer's strategy ceases to be a best response. In (ii), the employer prefers to deviate and never inspect. Likewise, it is easy to see that the employer cannot employ positive rewards $k > 0$ in any Nash equilibrium.

they are allowed to punish, and employers reward a worker maximally if they find the worker working and if they are allowed to reward. In the games that allow the employer to punish a deviating worker, cooperation can effectively be pursued. The expected future losses due to the unforgiving punishment outweigh the temptation to shirk. Without the possibility of punishment, full cooperation cannot be sustained in equilibrium. The promise that good behavior is rewarded may seduce the worker to work hard for a while, but if the employer never inspects and the reward therefore never materializes, the worker will be tempted to shirk. So from this perspective, games in which the employer can punish workers who are found shirking are expected to be more successful in generating actual cooperation.

4.3. Experimental Design and Procedures

The computerized experiment was carried out at the University of Nottingham. Subjects were recruited from a campus-wide distribution list. In total, 250 subjects participated in 21 sessions. Each session contained either five or six pairs of participants. Each subject participated in one session only. During a session no communication between subjects was allowed. Of each of the seven treatments, we carried out three sessions.

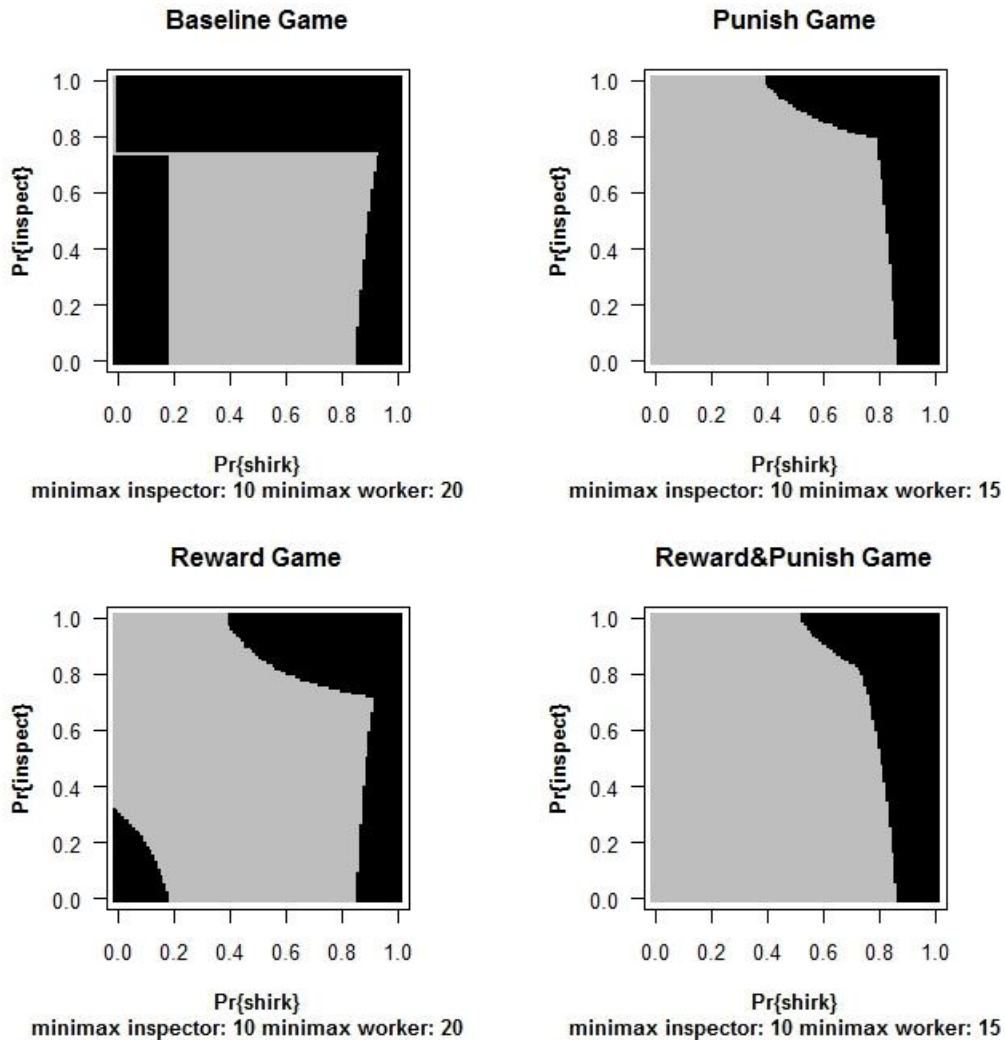
At the end of the session, subjects were paid in cash according to their accumulated point earnings from all rounds using an exchange rate of £0.007 per point. Sessions took about 40 minutes on average and earnings ranged between £5.6 and £23.0, averaging £12.1 (approximately US\$19.1 at the time of the experiment). Sessions started with a random assignment of subjects to computer terminals. Subjects received the instructions on paper, so that they could read along while an experimenter read the instructions 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.

At the start of the experiment, subjects were assigned to pairs and roles. Within each pair, one subject received the role of 'Employer' and the other the role of 'Worker'. Subjects knew that they would stay in the same role and in the same pair during the whole experiment. They were informed that each session consisted of at least 70 rounds, from round 70 on each round could be the last one with probability $1/5$. For comparability we kept the (computerized) random stopping draws constant across treatments: each treatment consisted therefore of three sessions with 71, 73 and 83 rounds, respectively.

In each treatment, a round started with a stage where at the same time the worker chose between 'high' (shirk) and 'low' (work) and the employer between 'inspect' and 'not inspect' which led to the payoffs presented in the right panel of Figure 4.1 on page 44. In the Baseline treatment, these were the only choices made in the round and subjects were immediately informed about the choices and payoff consequences for each one of them. At any time, subjects were informed of all choices and earnings of the own pair in previous rounds.

The other 6 treatments varied from the Baseline treatment in the tool that employers received

Figure 4.3.: Equilibria in the Repeated Game (continuation probability 0.8)



Notes: the pairs (p, q) in gray present the pairs that can be supported in this particular class of equilibria, while the pairs (p, q) in black cannot be supported in this class. In the ‘normal phase’, subjects mix with constant probabilities (p, q) in every stage game, and after inspection employers punish a worker maximally if they find the worker shirking and if they are allowed to punish, and employers reward a worker maximally if they find the worker working and if they are allowed to reward. The punish/reward games are based on the low ratio (1:1 technology). If a player deviates from the normal phase, she is set on her minimax payoff forever by the other player. In the Punish and Reward&Punish games, the minimax payoff of the worker decreases by 5 (because of the availability of a punishment of 5). In the games that allow punishments and rewards, the players may ignore the reward/punishment possibility, in which case the analysis coincides with the one for the baseline game. In this way, these graphs present additional equilibria offered by the relevant tool. We assume that a deviation of (p, q) is always immediately noticed, even with interior values of p and q . In reality, the normal phase should be carried out in “cycles” and players can only start punishing deviating players after a deviation from a cycle is observed. Therefore, in a “more realistic analysis”, the area of equilibrium pairs would diminish in each game, but the main qualitative features of the graphs would be preserved. For the more effective 1:3 technology, the pictures look very similar.

Table 4.1.: Experimental Design

Treatment	Reward	Punishment	Technology	Number of pairs
BL	no	no	–	17
R1:1	yes	no	1:1	18
P1:1	no	yes	1:1	18
R&P1:1	yes	yes	1:1	18
R1:3	yes	no	1:3	18
P1:3	no	yes	1:3	18
R&P1:3	yes	yes	1:3	18

to incentivize workers (Reward, Punish or Reward & Punish) and the effectiveness of the tool (Low or High). In each of these other treatments, the round was extended with an extra stage if the employer had chosen to inspect. In the extra stage, only the employer had to make a choice after receiving information of the worker’s choice between shirk and work. In the ‘R1:1’ and ‘R1:3’ treatments, the employer chose between ‘no action’ and ‘reward’, in the ‘P1:1’ and ‘P1:3’ treatments, between ‘no action’ and ‘punish’ and in the ‘R&P1:1’ and ‘R&P1:3’ treatments between ‘no action’, ‘reward’, and ‘punish’. If reward [punish] was chosen in the second stage, the employer chose the number of reward [punishment] tokens, a number from the set 0, 1, 2, 3, 4, 5. The employer paid a cost of 1 point per token. In the ‘1:1’ treatments the effectiveness ratio of the reward/punishment technology was low, meaning that each token increased (in case of reward) or decreased (in case of punishment) the payoff of the worker by one point. In the ‘1:3’ treatments, we employed a more effective 1:3 reward/punishment technology, in which case the worker’s payoff increased or decreased by three points for each token. Finally, both players in the pair were informed of the results in the pair (all choices and payoffs). Table 4.1 summarizes the experimental design.

4.4. Results

We present the experimental results in two parts. In Section 4.4.1, we present an overview of the aggregate results. This part provides the main answers to our research questions. In Section 4.4.2, we delve deeper into the data. There, we present the dynamics in the data and we provide an explanation of the main findings.

4.4.1. Overview

Figure 4.4 on page 50 displays how the inspect decisions of the employers and the shirk decisions of the workers developed over time. The two upper panels compare the Baseline treatment with the treatments with the low ratio. In all these treatments, there is a moderate upward trend in the frequency of inspection. In the second half of the experiment, the inspection probabilities are quite close to the stage game Nash benchmark of 75%. With the low ratio, inspection

probabilities do not differ much between treatments, although employers inspect to a somewhat lesser extent in P1:1 than in R1:1, R&P1:1 and BL. In contrast, the frequencies of shirking remain pretty constant across time in the low ratio treatments, at a substantially lower level than the stage game Nash benchmark. The treatments that allow for rewards and or punishments trigger somewhat less shirking than the Baseline treatment, but differences are modest.

The two lower panels provide the picture for the treatments with the high ratio. Here, the differences with the Baseline treatment are more pronounced. In R1:3 and Baseline, inspection frequencies are similar at the start and eventually grow to approximately the same level in the final rounds. In contrast, the inspection levels in R&P1:3 and P1:3 stay approximately constant, at lower levels than in the other two treatments. The right lower panel shows that subjects shirk substantially less in the treatments with the possibility of rewards and/or punishments than in the Baseline treatment. There are hardly any differences in the three treatments where employers have the possibility to incentivize workers through rewards and/or punishments. Thus, the decrease in inspection level in R&P1:3 and the even bigger decrease in inspection level in P1:3 do not come at the cost of higher shirking.

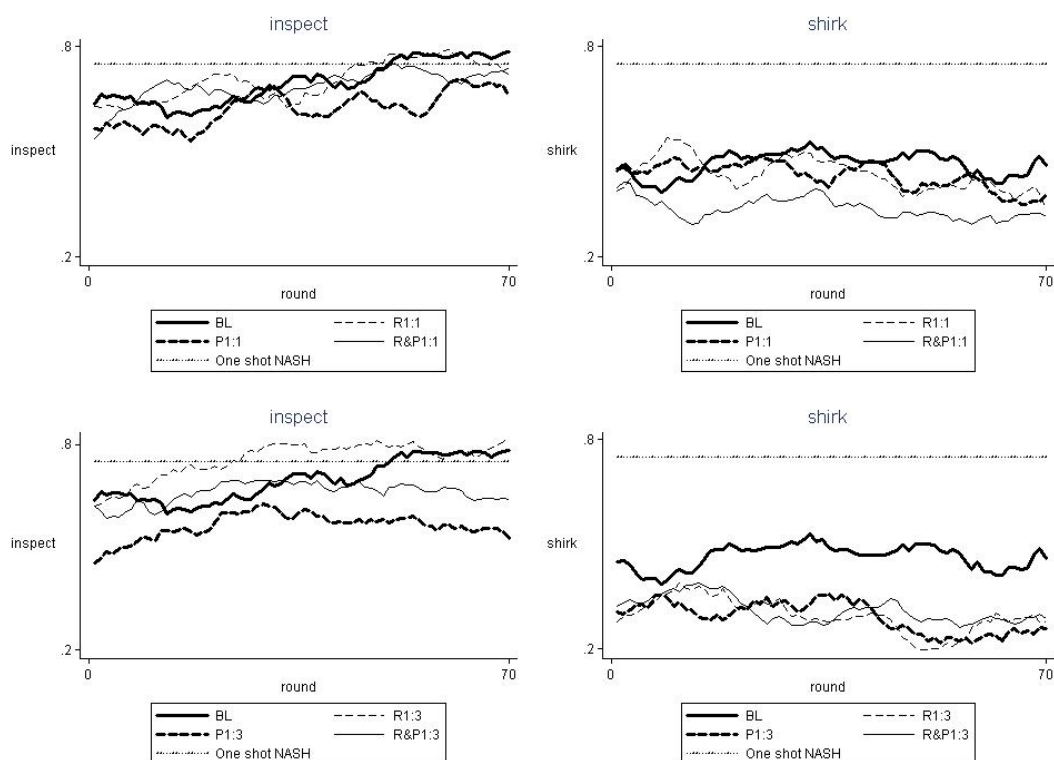
Because we are mainly interested in the comparison of the treatments after subjects have become familiar with the experiment, we focus on the second part of the experiment in the remainder of this chapter (unless we explicitly mention otherwise). Table 4.2 on page 51 presents the raw averages of inspections and shirking together with test results of hypotheses comparing the levels across treatments. Throughout this chapter, we employ a prudent test procedure with independent average statistics per pair of subjects. So each pair of subjects yields one data-point. We report the results of two-sided non-parametric ranksum tests.

When the punishment/reward technology is relatively ineffective (1:1), the modest differences between the treatments appear not to be significant, with as only exception the comparison of the inspection level between P1:1 and Baseline, which is weakly significant at $p=0.10$. P1:1 is the only 1:1 treatment where the inspection level is (weakly) significantly less than the stage game Nash benchmark of 75% ($p = 0.06$). In the 1:1 treatments as well as the Baseline treatment, the shirking levels are significantly below the stage game Nash benchmark of 75%.

With a highly effective punishment/reward technology (1:3), the picture for inspections is qualitatively similar, but some differences are statistically more pronounced. The comparisons of the inspection levels remain insignificant with two exceptions: in P1:3 lower inspection levels are observed than in R1:3 ($p = 0.06$) and P1:3 is the only 1:3 treatment where the inspection level is significantly below the stage game Nash benchmark. In contrast, the shirking levels in R1:3, P1:3 and R&P1:3 are all substantially and significantly below the Baseline treatment. With regard to shirking, the 1:3 treatments are statistically indistinguishable from each other.

In the comparison of the 1:1 treatments and the 1:3 treatments, the differences in shirking are significant in the Reward treatments ($p=0.06$) and the Punish treatments ($p=0.01$).

Figure 4.4: Timeseries Inspect and Shirk



Notes: for each round, the average of the proportions in the interval $[\text{round} - 5, \text{round} + 5]$ is displayed.

Table 4.3 on page 52 shows how often employers chose no action, reward and punish after they inspected the worker and observed her decision to work or shirk. In total, employers rewarded workers more often than that they punished them. In R&P 1:1, after inspection employers rewarded workers in 53% of the cases and punished them in only 7% of the cases. In R&P1:3, rewards were assigned in 47% of the cases and punishments in 20% of the cases. Further insight is obtained if these numbers are broken down for whether the worker behaved well or shirked. Unsurprisingly, after the employer observed the worker shirking, he hardly rewarded her and after he observed the worker working he hardly punished her. In R1:1, the employer rewards working in 55% of the cases and in P1:1 the employer punishes shirking in 51% of the cases. Likewise, in R1:3 the employer rewards working in 64% of the cases and in P1:3 the employer punishes shirking in 52% of the cases. So conditional on the tool being appropriate for the action taken, it is used with an approximately equal frequency. In R&P1:1 a remarkable shift in the relative frequencies is observed: here, working is rewarded in 76% of the cases while shirking is only punished in 22% of the cases. So with the low ratio, employers favor rewards over punishments when either tool is allowed. A similar shift is not observed in R&P1:3, though. There, working

Table 4.2.: Actions in Stage 1

		Inspect					Shirk				
Treatment	N	Mean	p-values (ranksum)				Mean	p-values (ranksum)			
			R1:1	P1:1	R&P1:1	=75%		R1:1	P1:1	R&P1:1	=75%
BL	17	74%	0.72	0.10	0.87	0.52	47%	0.51	0.48	0.15	0.00
R1:1	18	75%		0.14	0.90	0.81	40%		0.96	0.35	0.00
P1:1	18	65%			0.34	0.06	42%			0.23	0.00
R&P1:1	18	72%				0.88	33%				0.00
Treatment	N	Mean	R1:3	P1:3	R&P1:3	=75%	Mean	R1:3	P1:3	R&P1:3	=75%
BL	17	74%	0.54	0.21	0.30	0.52	47%	0.01	0.02	0.03	0.00
R1:3	18	79%		0.06	0.12	0.42	27%		0.81	0.62	0.00
P1:3	18	57%			0.48	0.05	27%			0.70	0.00
R&P1:3	18	67%				0.20	29%				0.00
R1:1 vs R1:3		p=0.37					p=0.06				
P1:1 vs P1:3		p=0.54					p=0.01				
R&P1:1vs R&P1:3		p=0.26					p=0.79				

Notes: in the columns mean the average of the means of all pairs is displayed; the p-values are the results of the rank-sum tests between treatments within technologies; =75% gives the result of comparing inspect and shirk with the one shot mixed Nash equilibrium benchmark (75%, 75%); bottom 3 rows present the outcomes of ranksum tests between technologies within treatments. Rounds 36-70 only.

is rewarded in 61% of the cases while shirking is punished in 62% of the cases.

In the Baseline treatment, we observe an approximately equal number of inspect/work outcomes as inspect/shirk outcomes. In contrast, Table 4.3 on the following page shows that when employers chose to inspect, they encountered working much more often than shirking in the treatments where punishments and/or rewards are allowed. Thus, even though conditional on the appropriate action employers used each tool about equally frequently, we observe much more reward decisions than punishment decisions because inspect/work occurred substantially more often than inspect/shirk.

Table 4.4 on page 53 provides an overview of the number of tokens assigned by the employer, conditional on choosing a reward or a punishment. The Table shows that in all treatments the expected punishment of shirking behavior is approximately equally large, in the range of 3.34 to 3.90. In contrast, there is more variation in the extent to which employers reward working. In the 1:1 treatments, the expected rewards of working behavior (4.15 in R1:1 and 4.15 in R&P1:1) are higher than the expected punishments of shirking behavior, while in the 1:3 treatments the expected rewards of working behavior (3.21 in R1:3 and 2.74 in R&P1:1) are lower than the expected punishments of shirking behavior. Thus, the level of the reward depends on the technology, and subjects reward less when the ratio is high. Possibly this result is due to inequality aversion considerations.

Furthermore, in the 1:1 treatments the mode of the distribution is to assign 5 tokens in all cases. That is, given than an employer chose to reward or punish, he tended to assign the maximum number of tokens. Again, the picture looks differently for rewards in the 1:3 treatments; there the

Table 4.3.: Actions in Stage 2

Treatment	after	N	no action	reward	punish
R1:1	work	286	45%	55%	
	shirk	184	97%	3%	
	all	470	66%	34%	
P1:1	work	243	98%		2%
	shirk	164	49%		51%
	all	407	78%		22%
R&P1:1	work	313	24%	76%	0%
	shirk	143	76%	3%	22%
	all	456	40%	53%	7%
R1:3	work	357	36%	64%	
	shirk	139	88%	12%	
	all	496	51%	49%	
P1:3	work	256	94%		6%
	shirk	102	48%		52%
	all	358	81%		19%
R&P1:3	work	310	34%	61%	5%
	shirk	110	28%	10%	62%
	all	420	33%	47%	20%

Notes: results conditional on inspecting in stage 1. Rounds 36-70 only.

mode of the distribution shifts to cheaper rewards of 2 or 3 tokens. It is also worth mentioning that employers sometimes used free punishments of 0 tokens if the worker shirked, while they almost never used free rewards of 0 points to reward if the worker worked. Possibly, employers regard a punishment of 0 tokens as a useful warning while they fear that a free reward backfires.

Table 4.5 on page 54 presents the efficiency levels of the firms on the left hand side and employer's and worker's total earnings on the right hand side. We define efficiency as the sum of the worker's and employer's earnings in stage 1. Arguably, this is the statistic that would be most interesting to the owners of the firm because it deals with the primary money streams in the firm (in actual firms rewards and punishments are not necessarily expressed in monetary terms).

When the technology is relatively ineffective (1:1), efficiency is only marginally and usually insignificantly enhanced by the possibility to reward and/or punish. Treatment P1:1 provides the exception, where the efficiency level is weakly significantly increased compared to the BL treatment. This is due to the fact that the same level of shirking is accomplished with fewer inspections in P1:1. Interestingly, in the 1:1 treatments the employer does not benefit from the possibility to reward and/or punish, while the worker is better off when rewards are allowed (both in R1:1 and R&P1:1, workers earn significantly more than in BL).

The picture is different in the 1:3 treatments where rewards and punishments are more effective. There, the efficiency levels are significantly enhanced when rewards and/or punishments are

Table 4.4.: Assignment of Tokens

Treatment	Actions		N	Tokens						Exp. Value
	stage II	stage I		0	1	2	3	4	5	
R1:1	reward	Work	157	0.00	0.14	0.06	0.04	0.03	0.73	4.15
		Shirk	5	0.20	0.40	0.40	0.00	0.00	0.00	1.20
		All	162	0.01	0.15	0.07	0.04	0.02	0.71	4.06
P1:1	punish	Work	4	0.75	0.00	0.25	0.00	0.00	0.00	0.50
		Shirk	84	0.17	0.02	0.02	0.02	0.05	0.71	3.90
		All	88	0.19	0.02	0.03	0.02	0.05	0.68	3.75
R&P1:1	reward	Work	238	0.05	0.08	0.08	0.03	0.01	0.76	4.15
		Shirk	4	0.00	0.00	0.00	0.50	0.25	0.25	3.75
		All	242	0.05	0.08	0.08	0.03	0.02	0.75	4.14
	punish	Work	0	-	-	-	-	-	-	-
		Shirk	31	0.19	0.00	0.10	0.10	0.03	0.58	3.52
		All	31	0.19	0.00	0.10	0.10	0.03	0.58	3.52
R1:3	reward	Work	229	0.00	0.06	0.28	0.34	0.03	0.29	3.21
		Shirk	16	0.06	0.31	0.19	0.38	0.00	0.06	2.13
		All	245	0.01	0.08	0.27	0.34	0.03	0.28	3.13
P1:3	punish	Work	16	0.38	0.19	0.06	0.06	0.06	0.25	2.00
		Shirk	53	0.19	0.08	0.06	0.09	0.06	0.53	3.34
		All	69	0.23	0.10	0.06	0.09	0.06	0.46	3.03
R&P1:3	reward	Work	188	0.01	0.27	0.32	0.06	0.03	0.31	2.74
		Shirk	11	0.00	0.18	0.45	0.27	0.00	0.09	2.36
		All	199	0.01	0.27	0.33	0.07	0.03	0.30	2.72
	punish	Work	16	0.00	0.13	0.00	0.13	0.25	0.50	4.00
		Shirk	68	0.00	0.03	0.19	0.21	0.03	0.54	3.87
		All	84	0.00	0.05	0.15	0.19	0.07	0.54	3.89

Notes: conditional on a reward or punishment decision, the average relative frequency of the number of tokens assigned in a treatment for the worker's decision is listed. The expected value is calculated as the sum of the products of the tokens and the relative frequencies; rounds 36-70 only.

allowed and employers are better off compared to the BL treatment. Remarkably, although the employers are the ones who decide whether they want to punish or reward, and therefore could ignore the possibility to reward if both tools are allowed, employers earned less in P1:3 than in R&P1:3. The difference is (weakly) significant at $p = 0.09$. In Section 4.4.2, we come back to this surprising result. The workers also benefit significantly from employers' ability to incentivize them, except in the treatment P1:3 where only punishments are allowed, in which case they earned approximately the same as in the BL.

4.4.2. Dynamics and Explanation

The previous section dealt with the aggregate static outcomes of the experiment. In this section, we present the behavioral dynamics and we provide an explanation of the main results. Table 4.6 on page 55 presents how often combinations of employer and worker decisions occurred in the different treatments. In addition, it displays transitions by listing the frequencies of outcomes in a new round conditional on the outcomes in the previous round.

In the columns 'freq', the relative frequencies of employer/worker decisions are listed. In BL, the most common combinations are inspect/work and inspect/shirk, which occur approximately

Table 4.6.: Played Combinations and Transitions

Treatment	t=t	freq.	t=t+1				Treatment	t=t	freq.	t=t+1			
			ni/w	ni/s	in/w	in/s				ni/w	ni/s	in/w	in/s
BL	ni/w	16%	47%	19%	16%	17%	BL	ni/w	16%	47%	19%	16%	17%
	ni/s	9%	22%	24%	22%	33%		ni/s	9%	22%	24%	22%	33%
	in/w	37%	14%	4%	60%	23%		in/w	37%	14%	4%	60%	23%
	in/s	37%	4%	6%	29%	62%		in/s	37%	4%	6%	29%	62%
R1:1	ni/w	14%	20%	14%	36%	30%	R1:3	ni/w	17%	14%	7%	61%	18%
	ni/s	11%	27%	16%	25%	31%		ni/s	4%	21%	11%	50%	18%
	in/w	45%	15%	6%	63%	15%		in/w	57%	20%	4%	65%	12%
	in/s	29%	7%	15%	30%	49%		in/s	22%	9%	4%	33%	54%
P1:1	ni/w	20%	29%	29%	26%	16%	P1:3	ni/w	32%	60%	22%	13%	7%
	ni/s	16%	27%	28%	25%	19%		ni/s	11%	42%	8%	31%	19%
	in/w	39%	21%	9%	53%	18%		in/w	41%	18%	2%	68%	12%
	in/s	26%	8%	8%	36%	48%		in/s	16%	8%	12%	38%	41%
R&P1:1	ni/w	18%	46%	12%	28%	14%	R&P1:3	ni/w	21%	24%	17%	43%	15%
	ni/s	10%	23%	25%	25%	27%		ni/s	12%	36%	22%	32%	9%
	in/w	50%	12%	5%	71%	11%		in/w	49%	21%	6%	58%	15%
	in/s	23%	6%	11%	30%	54%		in/s	17%	8%	15%	46%	31%

Notes: freq. gives the frequencies of all combinations employer/worker decisions in rounds 36-70; t=t presents the frequency in the current round and t=t+1 presents the outcomes in the subsequent round conditional on the combination of the current round; ni=not inspect, in=inspect, w=work, s=shirk.

inspects and the worker shirks, in which case subjects often stubbornly repeat their previous choices.

In the Punish treatment P1:3, the efficient outcome not inspect/work is repeated in a clear majority of the cases where it occurs. Likewise, inspect/work and inspect/shirk are also often repeated, both in P1:1 and P1:3. In contrast, in P1:3, the outcome not inspect/shirk is almost always abandoned, most often in favor of the outcome where the worker gives in (not inspect/work). In this treatment, the fear of punishment seems to loom large. In the reward and punish treatment R&P1:3, the dynamics are similar as in the reward treatment to the extent that the combination of inspect and work absorbs many previous outcomes. In R&P1:3 the outcome of inspect and work is repeated even more often once it is reached, but here it does not absorb behavior from the other cells. Here, the outcomes not inspect/work and inspect/shirk tend to be repeated, while after no inspect/shirk any outcome may occur.

A striking feature shared by all treatments is that both the employer and the worker tended to stubbornly repeat their choices when the bad outcome was reached where the employer inspects and the worker shirks. Table 4.7 on the following page zooms in on the question how likely such ‘battles of the will’ were, how long they lasted and how they tended to be resolved. In the 1:1 treatments, runs occurred approximately equally frequently in R1:1 and P1:1 as in BL, but they occurred to a lesser extent in R&P1:1. In the treatments where punishments and/or rewards were possible, the average lengths of these runs were smaller than in the baseline treatment. In contrast, in all effective technology 1:3 treatments, runs occurred much less frequently than in the baseline treatment, and if they occurred, they lasted shorter, except for R1:3. In all cases, it was the worker who was more likely to give in after a battle of the wills by changing her behavior

Table 4.7.: Battle of the Wills: Who Gives in?

Treatment	#runs	length (sd)	behavior changed by			Treatment	#runs	length (sd)	behavior changed by		
			work.	empl.	both				work.	empl.	both
BL	18	4.83 (2.75)	67%	22%	11%	BL	18	4.83 (2.75)	67%	22%	11%
R1:1	19	3.84 (0.76)	42%	42%	16%	R1:3	10	5.60 (3.86)	100%	0%	0%
P1:1	15	4.27 (2.12)	93%	7%	0%	P1:3	11	3.64 (1.80)	64%	27%	9%
R&P1:1	10	4.20 (2.49)	60%	30%	10%	R&P1:3	9	3.67 (0.71)	78%	22%	0%

Notes: a run is a series of consecutive rounds where the worker shirks and the employer inspects; runs shorter than 3 are discarded; we only consider runs that had their first round and their last round between 36 and 69.

to working.

In Section 4.4.1, we reported the remarkable result that even though employers made more money when they used punishments to incentivize workers in P1:3 than when they used rewards to encourage workers in R1:3, they did not shift toward using punishments when both tools were allowed in R&P1:3. Ideally, to investigate the success of rewarding versus punishing, one would like to classify employers as ‘punishers’, ‘rewarders’, ‘punishers and rewarders’ and ‘no-punishers and no-rewarders’ and the workers as ‘shirkers’ or ‘workers’ on the basis of an external measure. Then we could compare the occurrence of either type of employers across treatments, and we could compare their performance when matched with shirkers, and when matched with workers. We do not have such independent measures in our experiment, and therefore use behavior in the first 10 rounds as a proxy for the measure, and we use the rounds 11-70 to determine the success of various strategies. Table 4.8 on the next page presents employers’ earnings as a function of their own type and the type of worker they were matched with.

For completeness, the Table presents the results for the 1:1 treatments as well as the 1:3 treatments. Here, we focus on the 1:3 treatments because in those treatments we observed real differences between the treatments. In the treatment where employers are restricted to using rewards R1:3, employers classified as rewarder make clearly more money when they are matched with a worker who is not a shirker than employers who do not make use of the possibility to reward. If rewarders are matched with shirkers they make approximately the same amount as money as employers who do not use the reward tool. In the treatment where employers can make use of punishments but not rewards P1:3, when matched with a shirker employers make substantially more money when they are punishers than when they are not. In contrast, when matched with workers who work, the punishment strategy is counter productive and punishers earn less than the employers who refrain from punishing. Remarkably, when matched with workers who work, employers who refrain from punishing in P1:3 earn substantially more than employers who refrain from rewarding in R1:3. Possibly, the latent threat of (not used) punishments encouraged workers to behave well in P1:3.

Table 4.8.: Employers' Strategies and Earnings

Treatment	Worker	employer							
		punisher		no punisher/ no rewarder		rewarder		punisher/ rewarder	
		N	Mean (sd)	N	mean (sd)	N	mean (sd)	N	mean (sd)
R1:1	Worker			3	20.21 (1.57)	6	23.78 (4.93)		
	Shirker			6	21.27 (3.01)	3	22.27 (3.55)		
P1:1	Worker	3	25.34 (4.16)	6	25.39 (6.06)				
	Shirker	5	20.39 (1.70)	4	20.74 (3.38)				
R&P1:1	Worker	1	13.25	2	32.25 (18.03)	5	27.96 (3.35)	1	18.60
	Shirker	2	20.21 (2.89)	3	20.90 (2.07)	4	23.47 (1.19)		
R1:3	Worker			2	24.57 (1.08)	7	28.49 (3.37)		
	Shirker			4	22.36 (2.58)	5	23.23 (4.92)		
P1:3	Worker	3	25.48 (3.79)	6	33.85 (7.88)				
	Shirker	5	27.76 (4.42)	4	21.55 (3.50)				
R&P1:3	Worker			3	22.99 (2.94)	5	29.90 (3.70)	1	30.60
	Shirker	2	23.53 (3.65)	2	21.53 (4.78)	3	21.61 (1.83)	2	23.08 (4.86)

Notes: workers and employers are classified on the basis of their behavior in the first 10 rounds; employers' average earnings are based on rounds 11-70 (stage 1 and 2 earnings added); workers are classified on the basis of how often they shirked in the first 10 round, the 9 workers shirking fewest are classified as "workers", the other 9 as "shirkers"; employers are classified on the basis of the average assigned reward tokens (x_1) and the average punish tokens (x_2) over the first 10 rounds: if $\max(x_1, x_2) < 0.5$ then the employer is classified as "no punisher /no rewarder", if $\max(x_1, x_2) \geq 0.5$ and $|x_1 - x_2| < 0.25$ then the employer is classified as "punisher /rewarder", if $\max(x_1, x_2) \geq 0.5$ and $x_1 - x_2 \geq 0.25$ then the employer is classified as "rewarder", if $\max(x_1, x_2) \geq 0.5$ and $x_2 - x_1 \geq 0.25$ then the employer is classified as "punisher".

When both tools become available in R&P1:3, the picture becomes different. Unlike in P1:3, employers who are matched with shirkers earn hardly more when they act as punisher than when they refrain from punishing and rewarding. So punishing loses much of its bite when both tools are available. In contrast, employers who are matched with workers who work earn much more when they pursue a rewarding strategy than when they refrain from using, and the difference is bigger than in R1:3. So rewarding workers who behave well seems to become more remunerative when both tools are allowed. Another striking feature is that employers who are matched with well-behaving workers and who refrain from punishing and rewarding in R&P1:3 earn much less than employers who are matched with well-behaving workers and who refrain from punishing in P1&3. This suggests that the unused threat of punishing loses much of its force when employers can use rewards as well as punishments.

Table 4.9.: Questionnaire

		enjoyment of employer by using		aim is to influence behavior by using		appropriateness	
Treatment	Type	q1 reward (sd)	q2 punishment (sd)	q3 reward (sd)	q4 punishment (sd)	q5 reward (sd)	q6 punishment (sd)
R1:1	employer	4.08 (2.23)		5.42 (2.15)		5.83 (1.47)	
	worker	3.92 (2.11)		3.58 (2.19)		6.67 (0.65)	
	employer vs worker MW	$p = 0.84$		$p = 0.04$		$p = 0.13$	
P1:1	employer		2.50 (1.68)		4.42 (2.68)		4.08 (1.88)
	worker		2.50 (2.11)		4.00 (2.59)		4.92 (1.44)
	employer vs worker MW		$p = 0.69$		$p = 0.70$		$p = 0.29$
R1:1 vs P1:1	employer MW	$p = 0.07$		$p = 0.30$		$p = 0.02$	
	worker MW	$p = 0.09$		$p = 0.70$		$p = 0.00$	
R&P1:1	employer	4.92 (1.98)	2.92 (1.88)	5.42 (2.19)	4.58 (2.50)	5.67 (1.78)	4.67 (2.27)
	worker	3.67 (2.39)	3.00 (2.26)	4.42 (2.91)	3.92 (2.39)	6.75 (0.45)	5.08 (1.93)
	employer vs worker MW	$p = 0.18$	$p = 0.93$	$p = 0.71$	$p = 0.58$	$p = 0.05$	$p = 0.70$
Wilcoxon		q1 vs q2		q3 vs q4		q5 vs q6	
R&P1:1	employer	$p = 0.02$		$p = 0.26$		$p = 0.12$	
	worker	$p = 0.09$		$p = 0.51$		$p = 0.02$	

Notes: the questionnaire was filled out by the subjects of the last 6 sessions equally divided over R1:1; P1:1 and R&P1:1; MW=Mann-Whitney test; 7[1] = completely [dis]agree; q1="After inspection, I enjoyed rewarding the worker if he or she provided high effort/ I think the employer enjoyed rewarding me after inspecting if I provided high"; q2="After inspection, I enjoyed punishing the worker if he or she provided low effort/ I think the employer enjoyed punishing me after inspecting if I provided low effort"; q3="I assigned reward points to reinforce the worker's behavior/ I think the employer assigned reward points to reinforce my behavior"; q4="I assigned punishment points to change the worker's behavior/reward points to reinforce the worker's behavior /I think the employer assigned punishment points to change my behavior"; q5="It is appropriate to reward a worker who provides high effort"; q6="It is appropriate to punish a worker who provides low effort".

The success of the different strategies lines up with their actual use. In P1:3 where punishments were effective, 56% (5 out of 9) of the employers who were matched with a shirker pursued a punishing strategy. In R&P1:3, the percentage of employers exclusively relying on punishments decreased to 22% (2 out of 9).

In the final 6 sessions, we administered a questionnaire to further explore the reasons for an asymmetry between rewards and punishments. In the questionnaire, we asked employers as well as workers whether they felt that the employer enjoyed punishing/rewarding, whether the employer's aim was to influence the worker's behavior and to what extent the uses of punishments and rewards were appropriate. Table 4.9 presents the results. Employers and workers tend

to agree that employers enjoy rewarding good behavior, while they do not enjoy punishing bad behavior. Employers as well as workers think that rewards and punishments are used to influence the worker's behavior. Interestingly, the employers agree more with these statements than workers do, although the differences are usually not significant. Most informative are the answers regarding the appropriateness of the uses of rewards and punishments. Both employers and workers agree very much with the statement that it is appropriate to reward a well-behaving worker, while they agree substantially and significantly less with the statement that punishments are appropriate when the worker shirks. The difference in feelings about the appropriateness of the two tools may explain why many employers primarily chose to reward and why punishments lost part of their effectiveness when both tools were available.

4.5. Conclusion

Employers who want to stimulate workers to work hard may consider using rewards and punishments to achieve their goal. The use and effectiveness of rewards and punishments by employers is often hotly debated. Many people have strong opinions on how workers should be encouraged. It is surprising that this important discussion has not yet been backed up by controlled laboratory evidence. In this chapter, we have contributed to filling this gap.

We have obtained the following results. When rewards and punishments are relatively ineffective, as in our 1:1 treatments, rewards and punishments have only modest effects that are often not significant. Instead, when we introduced effective rewards and punishments in our 1:3 treatments, we observed substantial and significant effects. In the treatments where employers could use only punishments or only rewards, as well as in the treatment where both tools were allowed, we observed a common substantial decrease in the rate of shirking compared to the baseline treatment. In the treatment where employers were restricted to punishments, this was accomplished with much fewer costly inspections than when employers were restricted to rewards. As a result, employers earned more when they could only use punishments than when they could only use rewards. A remarkable result was that when employers could use both rewards and punishments, they did not shift in the direction of using punishments. To the contrary, employers continued to reward more often than punish when both tools were allowed.

A closer analysis reveals that the punishment strategy loses much of its force when both rewards and punishments are allowed. Pursuing a punishment strategy is more remunerative when employers cannot reward than when they can. In addition, employers as well as workers report that they feel that rewarding a well-behaving worker is more appropriate than punishing a shirker. The bottom line is that when employers can use rewards and punishments, our results suggest that they will primarily incentivize their workers through rewards, and for good reasons because the effectiveness of punishments may be eroded when rewards are possible. From the firm's perspective, shirking behavior is most efficiently reduced when the manager does not have the possibility to reward good behavior of the workers. So if the government (or the owners of

the firm) limits the extent to which bonuses can be given, superior results for the firm may be obtained.