Essay 2

Biased Information and Effort

2.1 Introduction

The employer has to update her belief on the agent after receiving the result of the evaluation. This problem has been studied in the previous chapter. Therefore, evaluations allow the employer to increase her knowledge of the employee's ability as shown in Lazear (1995). The employer may know the employee's ability better than the employee. This information asymmetry is developed for example in Bénabou and Tirole (2003). In a setting where the principal is better informed of the cost of the agent's effort, they model the impact of the principal's choice of the incentive structure, such as rewards, encouragement, or criticism. They show that encouraging the agent may decrease his motivation, whereas criticism may not discourage him.

These asymmetric, imperfect self knowledge situations especially arise when the employee faces a new task or starts a new job. The employer has more experience in evaluating the employee's ability in the task. An example where a principal is better informed on the agent's type than the agent can be found in teacher/student relationships. A teacher learns the ability of her student by means of mid-term exams. The teacher is able to evaluate her student's ability whereas the student does not perfectly know his ability in the class before receiving his mark. Let us note that in all these situations, the principal cannot perfectly observe ability; her knowledge is still imperfect. Evaluations

allow the principal to gain some information about the agent but these evaluations cannot fully reveal the agent's ability.

Thus, during the evaluation interview, the employer can transmit some information to the employee about his performance or ability in the job. A teacher reveals to the student some information about his competence through the marks scored in the exams. This transmission of information is referred to as "feedback". Our aim is to study how this feedback influences the agent's motivation to exert effort. The impact of feedback on the agent's motivation to exert effort is unclear. In some situations, the feedback may encourage the agent to exert a higher effort. For example, a student who for the first time achieves a mark above the average, i.e. the feedback on his work, may exert more effort to continue to progress. However, the feedback may also discourage the agent and therefore it may reduce the agent's effort. For example, a student who makes effort and receives a bad mark, may be discouraged and decide to exert no effort the next time. Providing information about performance can lead to an increase in the effort level through encouragement and motivation enhancement as in Lizzeri, Meyer, and Persico (2002) and Ederer (2010). However, it can also have adverse effects such as discouraging employees to exert effort as in Ertac (2005).

Gibbs (1991) introduced the idea that the employer can use feedback to manipulate the employee's belief about his performance or ability. Following Gibbs (1991), we study how an employee revises his belief about his ability by receiving information from the employer who is better informed on the employee's ability. Employers may be tempted to strategically manipulate the information to improve the performance of the employee. Longenecker, Sims, and Gioia (1987) found evidence that executives manipulate appraisals in an intentional and systematic manner. Similarly, a teacher may be tempted to overestimate the student's performance to avoid discouragement or she may be tempted to underestimate the student's performance to limit overconfidence.

We examine these questions in a laboratory setting by studying a two-stage game with two players (a principal and an agent) under various information schemes. In

our benchmark treatment, the agent can have three different levels of ability. Neither the principal nor the agent is informed on the agent's ability. The principal gets a signal about the agent's ability. The principal sends a truthful message to the agent corresponding to the signal received. The agent observes the message and chooses an effort level. The agent's performance depends on his ability and effort choice. The performance determines both the principal's and the agent's payoffs.

We implement two other treatments. In our Bias treatment, the message can be manipulated by the principal who can overestimate or underestimate the signal received on the agent's ability. The information is cheap talk. The principal has an incentive to tell the agent that he has a medium ability. Therefore, we can test whether or not the principal will manipulate the information and whether or not the manipulation will have an impact on the effort decision of the agent. Our third treatment, called Cost treatment, is similar to our Bias treatment, except that here, the principal bears a small cost for manipulating the message. The principal still has an incentive to tell the agent that he has medium ability. The Cost treatment tests if the principal's decision to manipulate is affected by an added cost. In other words, it will assess whether or not the principal will only manipulate the information because it is free or if he supports a moral cost for lying (Mazar, Amir, and Ariely (2008)).

This study is the first to analyze the impact of a feedback on the agent's ability that may have been manipulated by the principal and where the manipulation is restricted to a small bias, on the agent's effort.

As predicted, we find that principals strategically manipulate the information to increase the agents' effort. They send information such that the agents revise their belief on their ability, and choose a higher level of effort. Employees do revise their belief in accordance with the message received. Effort is higher in the Bias treatment than in the benchmark treatment. As a whole, the possibility to bias information improves the agent's performance and the principal's payoffs but does not improve efficiency given as the sum of the employee's payoff and the employer's payoff. We also find that overestimation occurs far more often than underestimation, in contrast to our predictions.

43

It seems that the principals manipulate the signal in order to increase the agent's belief on his ability whereas they are more reluctant to use manipulation in order to decrease the agent's belief on his ability. Introducing a cost for biasing the message reduces this effect but it still persists. People who bear a moral cost of lying are more reluctant to use manipulation but this cost should be significant for it to cancel all bias.

This chapter is organized as follows. Section 2.2 describes the recent literature on feedback and lies. Section 2.3 describes the experimental design and delivers predictions. The data are analyzed in Section 2.4. In Section 2.5, we discuss our results and conclude.

2.2   Related Literature

The principal obtains information about her agent from the evaluation process. During the evaluation interview, she can transmit some of this information to the agent. This information transmission is studied in the economy as feedback. Varying the nature of feedbacks, recent theoretical works show that it can be better for the principal in some cases to conceal information about the agent's performance or ability. Their conclusion opposes those of the standard principal-agent theory where there is a benefit to use all the available information as in Holmstrom (1979). Lizzeri, Meyer, and Persico (2002) study a dynamic principal-agent model where the agent cannot fully observe his performance. They conclude that the agent works harder when information is revealed but the principal is better off if feedback is not provided. It is too costly to provide feedbacks. In a tournament setting, Ederer (2010) study the effort choices of two agents when they cannot observe their performance due to a random noise and an incomplete information of their ability levels. The principal privately observes the performance difference of the two agents. Then, she chooses to reveal no information to the agents or to reveal truthfully the performance difference to the two agents. Ederer (2010) shows that it can be better to conceal information on the agent's relative performance depending on the convexity of the marginal cost of effort and on the complementarity of the agent's ability and effort.

Some empirical studies have also analyzed these issues. Some studies confirm the idea that providing information is not always optimal. Ertac (2005) runs an experiment where the subjects have to choose a level of investment. The return of the investment depends on both an individual factor, i.e. an ability level,that can be equally either low or high, and a common factor, i.e. a signal, that can be also equally low or high. The probability of a good outcome is equal to 1 if both the common factor and the individual factor are high. It is equal to 0.5 if one of the two factor is high and the other is low. Finally, the probability of a good outcome is equal to 0 if both factors are low. The subjects are randomly match in 5-person groups with a particular common factor. The subject are asked their belief about the individual factor and the common factor once before and once after observing others' outcomes. She shows that when reward schemes are exogenous and independent and performances are sufficiently complementary, withholding information on social comparison (information on others) may be optimal because social comparison information induces a negative correlation among agents with regard to self-confidence, effort and outcome. Eriksson, Poulsen, and Villeval (2009) in a real-effort experiment conclude that mid-term or continuous feedback about relative performance does not improve performance regardless of the pay scheme used (piece rate or winner take all pay condition). Moreover, the continuous feedback on the relative performance decreases the quality of work because the number of mistakes is significantly higher than under the no feedback condition. However, other studies show that receiving information increases the performance. For example, Azmat and Iriberri (2010a) study, in a natural experiment, the effect of providing relative performance feedback on the performance. High School students receives the information of their absolute performance. During one year, they also receive the class average performance. The results show that the student's performance increases by 5% with the new information. The authors confirm this result in a laboratory setting (see Azmat and Iriberri (2010b)). Bandiera, Larcinese, and Rasul (2010) measure the causal effect of interim feedback on individual's performance using a natural experiment involving a UK university where different departments have historically different rules on the provision of feedback to

their student. Some departments provide students with their period one scores before stratng the period two, i.e. feedback regime, while other departments reveals the period one score at the end of the academic year, i.e. no-feedback regime. Their results show that the provision of feedback has a positive effect on the period two scores. This effect is more important for the more able students but it doesn't discourage less able students. Thus provision of feedback has litigate impact on the agent's performance.

This literature always analyzes the impact of a truthful information. We are however interested in a situation where the principal can manipulate feedback. Therefore our work is also related to the literature on cheap talk, where a sender can send a free message to the receiver. This literature has revealed that when players' preferences are not aligned, a free message cannot theoretically convey any useful information. Then, the only equilibrium is a "babbling" equilibrium, in which the message should be disregarded by its recipient as in Farrell and Rabin (1996). Some experimental evidences however show that senders tend to send truthful messages more often than predicted and that recipients tend to trust the messages although they may not be credible theoretically. If players send more truthful messages than theoretically predicted, this might be due to lie aversion. Recent experimental studies on communication more precisely analyze lying behavior. Most subjects have a preference for truth-telling and feel some lie-aversion as in Gneezy (2005) and Sanchez-Pages and Vorsatz (2007). Gneezy (2005) study, in a sender-receiver game with conflict of interest, the lying behavior. It exists two payoffs distribution, A and B. The sender observes the payoffs of each players for the two distributions. Then he can communicate to the sender which is the distribution that maximizes the receiver's payoff. After observing the message, the receiver chooses A or B. The payoffs are implemented according to the receiver's choice. The author test the role of monetary incentives in the lying decision by varying the consequences of lies. The results show that 36% of the subjects choose to lie for a gain of $1 and a loss of $1 for the receiver. This proportion increases to 52% when the gain is $10 for the sender and the receiver's loss is also $10. People are sensitive to their profit when deciding to lie. More surprinsingly they also care about the other side payoff: when their gain for

lying is \$1 but the loss for the receiver is \$10, the proportion of lie decreases to 17%. This last result indicates that people are also sensitive to the negative consequences of their lie. Controlling with a dictator game, the author concludes that it is not only the care of others that motivate behavior but also lying aversion.

Until recently, these two strands of literature (that on feedbacks and that on lies) seem to ignore each other. Mohnen and Pokorny (2006) study the honesty of feedback in employer-employee relationships. Their results show that principals are influenced by the actual ability of their agent, and are less likely to give a positive feedback if they face a low ability agent. Moreover, principals prefer to give no feedback to a low ability agent in order to avoid telling lies. Finally, Mohnen and Pokorny (2006) find that agents adjust their effort according to the received feedback even so it may be a lie. However this result does not occur anymore in a long term relationship. The main difference with our experiment states in the type of feedback. In Mohnen and Pokorny (2006), principals choose whether to give feedback and whether to send the truth. We allow manipulated feedback such as they still convey information for the agent.

Ederer and Fehr (2007) compare in a tournament setting, the impact of no feedback, true feedback, and manipulated feedback on agent's performance. They find that most principals exhibit lie-aversion and send truthful feedback even when they can manipulate the feedback. The agents respond to feedback but decrease their effort when facing possibly manipulated feedback as compared to certain truthful feedback. Our experiment shows some similarities with Ederer and Fehr (2007) but it also differs in important respects. Their experiment is a two-stage game in a tournament setting. The output of the first period is the sum of the agent's effort and a random variable. The principal observes this output and sends a message on performance. Therefore the message contains information on the random variable. The principal is free to report any information. Like them, we run an experiment that compares truthful and biased messages. In our study, truthful message is about the agent's ability and not on his performance. The principal receives a signal about the agent's ability and sends a message to the agent. The agent chooses

an effort and the payoffs are determined. Our experiment matches one principal with one agent and the agent is paid a wage plus a bonus depending on performance. We constraint manipulation so that only a slight change in the truthful message is possible: the principal can only marginally overestimate or underestimate the truth. Thus, the biased information is always informative to the agent to some extent.

2.3   Experimental Design

2.3.1   The game

We study a two-stage game with two players (an employer and an employee). The agent can have three different levels of ability. Neither the principal nor the agent are informed on the agent's ability. The principal gets a signal about the agent's ability. He sends a message to the agent. The agent observes the message and chooses an effort level. The agent's performance depends on his ability and effort choice. The performance determines both the principal's and the agent's payoffs.

In our benchmark treatment, each employee is randomly matched with an employer. The employee's ability level a is randomly determined at each period. It can take any of three values (low, middle, or high) with equal probability: $a \in \{1, 3, 5\}$. This distribution is common information. The ability levels are i.i.d.. Neither the employee nor the employer are informed on the employee's ability. The employer receives a signal about the employee's ability, s, but the employee receives no information. As such, the employer is better informed than the employee. This signal is imperfect information on the employee's actual ability. With the same probability, each signal can underestimate, overestimate, or accurately describe the employee's actual ability. In expected terms, it is equal to the actual ability. The distribution of the signals is common information.

In the first stage, the employer sends a message, m, to the employee. In the benchmark treatment, the message always corresponds to the signal received ($m = s$). The employer is passive.

In the second stage, after observing the message, the employee chooses an effort level. This effort, e, can be either low, medium, or high, $e \in \{e_L, e_M, e_H\}$ [1]. The employee's and employer's payoffs depend on the employee's actual ability and on the level of effort. The employer's payoff depends positively on the effort chosen by the employee, e, and on his actual ability, a. The employer is always better off when the employee chooses a high effort regardless of his ability. The employer's payoff is given by

$$\Pi(q(a, e_L), f, B) < \Pi(q(a, e_M), f, B) < \Pi(q(a, e_H), f, B),$$

where f is the fixed wage, B is the bonus, q denotes performance, and z is standard. The employee's payoff is based on a flat wage minus the cost of effort. The employee can also earn a bonus if his performance reaches a certain threshold. The employee's payoff is given by

$$W = \begin{cases} f - c(e) & if \quad q(a, e) < z \\ f - c(e) + B & if \quad q(a, e) \geq z. \end{cases}$$

Table 2.1 displays the various possible payoffs

Ability	Effort level (chosen by the employee)	Employee's payoff	Employer's payoff
	low	9	5
1	medium	8	7
	high	4	9
	low	9	7
3	medium	8	9
	high	11	11
	low	9	9
5	medium	15	11
	high	11	13

Table 2.1: Payoffs

This bonus scheme provides strong incentive to a medium-ability employee to exert the highest level of effort. The bonus is not high enough to compensate a low-ability employee's cost to exert the highest level of effort[2]. Indeed, the employee who receives a message indicating that he has low ability is certain that he will never get a bonus. He

---

[1] To study the manipulation effect, three levels of effort are enough.

[2] It would have been too costly for the employer to incentivize the low-ability employee to exert the highest effort.

chooses a low effort level that minimizes his cost. In contrast, a high-ability employee is assured to get a bonus, and for him a medium effort level is sufficient as exerting the highest level of effort is unnecessarily costly. Therefore, the employees with the greatest incentives to work hard are the employees with a medium ability who must exert a high level of effort to reach the threshold and receive the bonus.

The bias and the cost treatments differ from the benchmark treatment only with respect to the messages that can be sent by the employer. Indeed, the employer is able to manipulate the signal received. She can send a message that overestimates or underestimates the signal or that corresponds honestly to the signal, $m \in \{s-1, s, s+1\}$. In the Cost treatment, the employer bears a cost, C, when she manipulates the message. The cost for manipulation equals 1 point per bias. Table 2.2 describes for each ability level (column 1), the different signals potentially received by the employer (column 2) and the messages that she sends to her employee in the benchmark treatment (column 3). The fourth column describes the various messages an employer can send to her employee in the bias and cost treatments conditional on the different signals she received.

Employee's ability (unobservable)	Signal (observed by the employer)	Message (sent by the employer)	
		Benchmark treatment	Bias and cost treatments
	0	0	-1, 0, or 1
1	1	1	0, 1, or 2
	2	2	1, 2, or 3
	2	2	1, 2, or 3
3	3	3	2, 3, or 4
	4	4	3, 4, or 5
	4	4	3, 4, or 5
5	5	5	4, 5, or 6
	6	6	5, 6, or 7

Table 2.2: Signals and messages by treatment

The comparison between the benchmark and bias treatments allows us to study how employers use a costless manipulation in order to increase their payoff and how employees condition their effort choices on the message they receive. The introduction of the Cost treatment allows us to test the impact of a costly manipulation on both the employers' and employees' decisions.

2.3.2 Procedures

The experiments have been conducted at the GATE laboratory, Lyon, France using the Regate software (Zeiliger (2000)). Via the ORSEE software (Greiner (2004)), we recruited 112 undergraduate students from local business and engineering schools. One session with the benchmark treatment, three sessions with the Bias treatment, and two sessions with the Cost treatment were organized. Each subject participated in only one session. We used a between-subject design.

At the beginning of the session, each subject was randomly assigned to a computer. The subjects first had to make lottery choices to control for their risk attitude. the instructions (a translation of all the instructions is available in the appendix 2.6) were read aloud. The subjects had to fill out a questionnaire to check their understanding of the game. The experimentalist checked the answers and re-explained the game in private. We used non-neutral framing in the instructions to facilitate the subjects' comprehension. Following this, the subjects played the game. At the end of the session, they filled out a demographic questionnaire. A session lasted 80 minutes on average.

Risk aversion can affect the employee's choice of effort level when they are uncertain about their ability and thus have to form beliefs about it. We used the lottery procedure of Holt and Laury (2002) to elicit the subjects' risk aversion. The subjects filled out a 10-decision questionnaire. Each decision consists of a choice between two lotteries, option A and option B. The payoffs for option A (the safer lottery) are either €2 or €1.60, whereas option B pays either €3.85 or €0.10 (the riskier lottery). In the first decision, the probability of the high payoff for both options is 1/10. In the second decision, the probability increases to 2/10, and so on. The high payoff's probability for each decision increases as the decision number increases. When the probability of the higher payoff is large enough (1/2), subjects should switch from option A to option B. Risk neutrality corresponds to a cross at the fifth decision: risk-loving subjects are expected to move earlier and risk-averse subjects, at the sixth decision or after.

Then, the subjects were divided equally into two groups — the employers' and the employees'groups — and pairs were formed randomly. There were 20 periods

and we use a stranger matching protocol. In each period of the main game, we asked employers to predict the effort level choice of their employee after they sent the message. We elicited these beliefs in order to know the employer's expectations regarding the employee's reaction to the message. To incentivize beliefs, the subjects were paid an extra €1, if their predictions were true at least 50% of the time[3].

Subjects were paid in cash in a separate room by someone who was not aware of the content of the experiment (this was made common information in 2.6). Subjects received a show-up fee of €3. In the Cost treatment, the cost for manipulation equals €0.05 per bias. One of the ten decisions of the Holt and Laury's lottery was drawn randomly at the end of the session, and in the payment room, subjects were paid according to their choice in this decision. On average, subjects earned €15.4 (min= €10.7; max= €18.5).

2.3.3 Theoretical Predictions

In this section, we describe the theoretical predictions of the experimental game. All the predictions are determined by assuming the optimizing behavior of self-interested and risk-neutral agents. We will discuss at the end of this section the consequences of introducing risk aversion. Consider the employee's prior belief on his ability before receiving a message from the employer. An employee knows that he can be low-, medium-, or high-ability employee with equal probability ($p = 1/3$). With no information, he will maximize his expected payoff by choosing a medium effort level ($E(W) = 9$ with e_L). For her part, an employer maximizes her payoff when the employee chooses a high effort level. Providing information may increase the employee's effort level by changing his belief about his ability level. The employee revises his belief on his ability by using the information he learns from the message sent by the employer (Bayesian inference). Which message will be optimal for the employer to send, knowing that she prefers the employee to choose a high effort level? In the following, we study each type of information (exact, imperfect, and manipulated) and discuss its impact on the employee's belief

---

[3]The payment is rather small to avoid strategic behaviors in this part of the game that may decrease the incentive for information manipulation.

about his ability level and effort choice by solving for a perfect Bayesian equilibrium. We study the efficiency in all these cases.

2.3.3.1 Game with perfect information

If the employee could perfectly know his ability, a low-ability employee ($a = 1$) would choose a low effort level (e_L); a medium-ability employee ($a = 3$) would choose a high effort level (e_H), and a high-ability employee ($a = 5$) would choose a medium effort level (e_M). The employees with the highest incentives to exert the highest effort level are the medium-ability employees. The employer is indifferent between a medium-ability employee ($a = 3$) who chooses a high effort level and a high-ability worker ($a = 5$) who will choose a medium effort level. However, the employer's payoff is higher when the employee chooses a high effort level (e_H) regardless of his ability.

2.3.3.2 Game with imperfect information (benchmark treatment)

In this game, the employee receives a message from the employer. The message has to be equal to the signal the employer has received, $m = s$. This signal is imperfect information about the employee's actual ability. For each ability level, the three possible signals have the same probability of being observed by the employer (2.2). Thus, the probability of each ability, conditional on the signal, is as follows:

- $P(a = 1 | s = 0 \cup s = 1) = 1$

- $P(a = 1 | s = 2) = P(a = 3 | s = 2) = 1/2$

- $P(a = 3 | s = 3) = 1$

- $P(a = 3 | s = 4) = P(a = 5 | s = 4) = 1/2$

- $P(a = 5 | s = 5 \cup s = 6) = 1$.

When the employee receives a message $s \in \{0, 1, 3, 5, 6\}$, he can perfectly learn his ability. As in the perfect information game, he maximizes his payoff, W

Max $W(e, s = 0 \cup s = 1) = 9$, by choosing e_L

Max $W(e, s = 3) = 11$, by choosing e_H

Max $W(e, s = 5 \cup s = 6) = 15$, by choosing e_M.

However, when he receives a message $s \in \{2, 4\}$, he is still uncertain about his ability level. In these cases, he maximizes his expected payoffs

Max $W(e|s = 2) = 9$, by choosing e_L

Max $W(e|s = 4) = 11.5$, by choosing e_M.

As in the perfect information game, the employer's payoff is higher when the employee chooses a high effort level regardless of his ability. The employer is better off when she can send the message $m = 3$, i.e., when she receives the signal $s = 3$.

2.3.3.3 Game with manipulated information (bias and cost treatments)

The employer can choose to send a truthful message ($m = s$) or manipulate the signal. The information is cheap talk in the Bias treatment. The employer can manipulate the information by overestimating or underestimating the signal, $m \in \{s - 1, s + 1\}$. Remember that the signal the employer received is imperfect information about the employee's actual ability. Table 2.2 describes for each ability level, the different signals that can be received by the employer (the same as in the benchmark treatment) and the various messages the employer can send to the employee.

The employer maximizes her payoff[4] when the employee chooses a high effort level. Thus, the employer has an incentive to send the message $m = 3$ whenever possible. Table 2.2 indicates that she can send this message when she observes $s \in \{2, 3, 4\}$. Therefore, the employer should systematically overestimate signal $s = 2$ and underestimate signal $s = 4$, and send message $m = 3$. For all other signals, the manipulation is useless (for

---

[4]The manipulation cost is so small that the theoretical predictions for the bias and the cost treatments are the same in our game.

example manipulating $s = 1$ will never convince the employee that he has medium ability).

In the Bias treatment, we assume that the employer manipulates the information only when it is optimal for her to do so. We assume that the subjects have a preference for truth-telling when they cannot increase their payoff by manipulating signals. This assumption will be tested by comparing decisions in the bias and the cost treatments. We should observe a decrease in the unnecessary manipulations when the employers bear a cost of manipulation. Therefore, we expect that in both the bias and the cost treatments, when receiving signal $s \in \{0, 1, 3, 5, 6\}$, the employer sends a message equal to the signal.

What is the effort choice of the employee knowing that the message can indicate a manipulated signal? When the employee receives the message $m \in \{-1, 0, 6, 7\}$, he can perfectly infer the employer's signal and revise his beliefs on his ability level accordingly. As in the perfect information game, the employee maximizes his payoff

Max $W(e, m = -1 \cup m = 0) = 9$, by choosing e_L

Max $W(e, m = 6 \cup m = 7) = 15$, by choosing e_M

When he receives message $m \in \{1, 2, 3, 4, 5\}$, the employee remains uncertain about the signal and his ability level but can revise his prior belief by taking into account the fact that the employer may have biased the signal. His revised belief on his ability is as follows:

- $P(a = 1|m = 1) = 3/4$ and $P(a = 3|m = 1) = 1/4$

- $P(a = 1|m = 2) = P(a = 3|m = 2) = 1/2$

- $P(a = 1|m = 3) = P(a = 5|m = 3) = 1/5$ and $P(a = 3|m = 3) = 3/5$

- $P(a = 3|m = 4) = P(a = 5|m = 4) = 1/2$

- $P(a = 3|m = 5) = 1/4$ and $P(a = 5|m = 5) = 3/4$.

55

In these cases, he chooses to maximize his expected payoffs

$MaxW(e, m = 1 \cup m = 2)$, by choosing e_L

$MaxW(e, m = 3)$, by choosing e_H

$MaxW(e, m = 4 \cup m = 5)$, by choosing e_M

However, if the employee also assumes that when his employer is indifferent between possible messages, she has a preference for truth-telling or bears a manipulation cost, then only the message $m = 3$ is ambiguous. That does not change our predictions.

Table 2.3 summarizes the theoretical predictions for each treatment under the assumptions of selfishness and risk neutrality.

Signal observed by the employer	Benchmark treatment		Bias and cost treatments	
	Message sent by the employer	Effort level chosen by the employee	Message sent by the employer	Effort level chosen by the employee
0	0	low	0	low
1	1	low	1	low
2	2	low	3	high
3	3	high	3	high
4	4	medium	3	high
5	5	medium	5	medium
6	6	medium	6	medium

Table 2.3: Messages and predictions about the efforts in the three treatments

In the bias and the cost treatments, the employers should never send messages 2 and 4. If they do so, the employees will choose a low and a medium effort level respectively. Therefore, we can make the following predictions.

Prediction 1: In the benchmark treatment ($m = s$), the employee chooses a low effort level, e_L, conditional on observing $m \in \{0, 1, 2\}$. An employee observing $m = 3$ chooses a high effort level, e_H. An employee observing $m \in \{4, 5, 6\}$ chooses a medium effort level, e_M.

Prediction 2: In the bias and the cost treatments, the employer sends a biased message, $m = 3$, when she receives $s \in \{2, 4\}$. For the other signals, she sends a message equal to the signal received, $m = s$.

Prediction 3: In the bias and the cost treatments, the employee receiving $m \in \{-1, 0, 1, 2\}$ chooses a low effort level, e_L. An employee receiving $m = 3$ chooses a high effort level, e_H. An employee receiving $m \in \{4, 5, 6, 7\}$ chooses a medium effort level, e_M.

Prediction 4: A small manipulation cost should not interfere in the manipulation decision if the preference for truth-telling assumption is verified. However, if this is not the case, we can expect that the introduction of a manipulation cost decreases useless manipulation.

Prediction 5: Manipulation of information by the employer leads to the employee choosing a high effort level more often than truthful information.

If we release the assumption of risk neutrality on the employees' side, the predictions in the case of uncertain messages are affected. In the benchmark treatment, risk aversion changes predictions for the message $m = 4$. Remember that $P(a = 3|s = 4) = P(a = 5|s = 4) = 1/2$. Receiving the message $m = 4$, the employee has to compare his expected payoff with the different effort choices. A low effort level, e_L, gives him a certain payoff of $W = 9$ (i.e., whatever his ability). Respectively his expected payoffs are 11.5 and 11 with a medium effort level, e_M, and a high effort level, e_H. Depending on his degree of risk aversion, a risk averse employee can prefer the certain payoff, 11, and choose the high effort level, e_H.

In the manipulation treatments, the employee's risk aversion has no impact when the message is fully informative on the employer's signal ($m \in \{-1, 0, 6, 7\}$). Risk aversion has an impact only for employees receiving message $m \in \{3, 4, 5\}$. For example, an employee who receives message $m = 4$ ($P(a = 3|m = 4) = P(a = 5|m = 4) = 1/2$) compares his payoffs with three possible effort levels:

- Low effort level, e_L, gives him a certain payoff 9.

- Medium effort level, e_M, leads to an uncertain payoff with expected value $E(W = 11.5)$.

- High effort level, e_H, gives him a certain payoff 11.

A sufficiently risk averse employee can therefore switch from a medium effort level, e_M, to a high effort level, e_H, to guarantee a certain payoff.

<u>Prediction 6</u>: Releasing the assumption of employees' risk neutrality changes the predictions as follows:

- In the benchmark treatment, a risk averse employee observing message $m = 4$ chooses a high effort level, e_H.

- In the bias and the cost treatments, a risk averse employee observing message $m = 3$ chooses a low effort level, e_L, and when observing message $m \in \{4, 5\}$, he chooses a high effort level, e_H.

An employer who expects her employee to be risk averse will also adjust her behavior. She will send message $m \in \{4, 5\}$ as often as possible in the manipulation treatments. Therefore, she will overestimate signal $s = 3$ and underestimate signal $s \in \{5, 6\}$.

2.3.3.4 Efficiency

Does manipulation increase efficiency? In this subsection, we study the employee's payoff, the employer's payoff, and overall efficiency for the perfect information game, imperfect information game (benchmark treatment), and manipulated information games (bias and cost treatments).

In the perfect information game, the employee has the probability $P(a = 1) = P(a = 3) = P(a = 5) = 1/3$ of being a low-, medium-, or high-ability employee. His mean expected payoff is 11.66 when he chooses his effort level to maximize his payoff[5]. As a consequence, the employer's mean expected payoff is 9 and efficiency is 20.66.

In the imperfect information game (benchmark treatment), the employee also has the probability $P(a = 1) = P(a = 3) = P(a = 5) = 1/3$ of being low-, medium-, or

---

[5]A low-ability employee, $a = 1$, chooses a low effort level; a medium-ability employee, $a = 3$, chooses a high effort level; and a high-ability employee, $a = 5$, chooses a medium effort level.

high-ability employee. However, he has the probability $P(s=0)=P(s=1)=P(a=3)=P(a=5)=P(s=6)=1/9$ and the probability $P(s=2)=P(s=4)=2/9$ of observing each signal. His mean payoff is 11.11 when he chooses his effort level to maximize his payoff. The employer's mean expected payoff becomes 8.33. The imperfect information situation decreases both the employee's and the employer's expected payoffs. Efficiency is 19.44.

In the Bias treatment, the employee has the probability $P(m=0)=P(m=1)=P(m=5)=P(m=6)=1/9$ and the probability $P(m=3)=5/9$ of receiving each message[6]. In this case, his mean payoff is 10.66 and the employer's mean payoff is 9.66. In our game, introducing information manipulation increases the employer's expected payoff and decreases the employee's expected payoff. However, the total payoff 20.32 is higher than in the benchmark treatment.

In the Cost treatment, the employee has the same mean payoff, 10.66, as in the Bias treatment. However, the employer's mean payoff decreases to 9.22. Introducing a cost of manipulation decreases efficiency, 19.88, as compared to in the Bias treatment; however, efficiency is still higher than in the benchmark treatment.

2.4 Results

First, we study the employer's decision. Then, we compare the effort level decision in the different treatments.

2.4.1 Summary statistics

Table 2.4 summarizes the main results.

---

[6]We suppose that the employer manipulates the information according to our predictions.

		Benchmark treatment		Bias treatment		Cost treatment	
				Frequency (%)			
	−1			12	(2.07%)	1	(0,28%)
	0	18	(10%)	22	(3.79%)	28	(7,78%)
	1	26	(14.44%)	56	(9.66%)	42	(11,67%)
Distribution	2	46	(25.56%)	66	(11.38%)	59	(16,39%)
of	3	18	(10%)	145	(25%)	79	(21,94%)
messages	4	26	(14.44%)	121	(20.86%)	76	(21,11%)
	5	19	(10.56%)	87	(15%)	42	(11,67%)
	6	27	(15%)	46	(7.93%)	33	(9,17%)
	7			25	(4.31%)	0	(0%)
Frequency	underestimation (-1)			109	(18.79%)	32	(8,89%)
of biased	no bias (0)			127	(21.90%)	217	(60,28%)
information	overestimation (+1)			344	(59.31%)	111	(30,83%)
Choice	low	71	(39.44%)	200	(34.48%)	149	(41,39%)
of the	medium	70	(38.89%)	211	(36.38%)	133	(36,94%)
effort level	high	39	(21.67%)	169	(29.14%)	78	(21,67%)
Prediction	low	76	(42.22%)	159	(27.41%)	106	(29,44%)
of the	medium	57	(31.67%)	199	(34.31%)	106	(29,44%)
effort level	high	47	(26.11%)	222	(38.28%)	148	(41,11%)
Total		180		580		360	

Table 2.4: Descriptive statistics

A two-sample Kolmogorov-Smirnov test[7] for equality of distribution functions indicates that the message distributions are significantly different ($p < 0.05$) across all our treatments. We observe that the messages are more concentrated around $m = 3$ when manipulation is allowed. The whole distribution of choices shows that only 21.90% of the subjects choose to send a non-biased message in the Bias treatment. However, this proportion increases in the Cost treatment to 60.28%. The employers use manipulation but we can also observe that overestimation is relatively more frequently used as compared to underestimation in both manipulation treatments. The effort decision of the employees differs significantly between the benchmark and the bias treatments ($p = 0.0676$*, Mann-Whitney rank-sum test) and between the bias and the cost treatments ($p = 0.0072$***, Mann-Whitney rank-sum test). Table 2.5 displays the distribution of effort levels in our treatments for each signal received by the employer.

---

[7]In all statistical tests reported in this paper, we use a conservative test based on the average decision of the individual across periods as one unit of observation.

Signal	Treatment	Low effort level	Medium effort level	High effort level	Total	No. of obs.
Signal 0	Benchmark	100.00%	0.00%	0.00%	100%	18
	Bias	96.92%	0.00%	3.08%	100%	65
	Cost	100.00%	0.00%	0.00%	100%	37
Signal 1	Benchmark	100.00%	0.00%	0.00%	100%	26
	Bias	83.08%	9.23%	7.69%	100%	65
	Cost	92.31%	5.77%	1.92%	100%	52
Signal 2	Benchmark	60.00%	15.56%	24.44%	100%	46
	Bias	46.83%	21.43%	31.75%	100%	126
	Cost	54.44%	21.11%	24.44%	100%	90
Signal 3	Benchmark	0.00%	0.00%	100.00%	100%	18
	Bias	16.07%	21.43%	62.50%	100%	56
	Cost	25.81%	19.35%	54.84%	100%	31
Signal 4	Benchmark	0.00%	61.54%	38.46%	100%	26
	Bias	10.14%	47.10%	42.75%	100%	138
	Cost	8.43%	56.63%	34.94%	100%	83
Signal 5	Benchmark	0.00%	100.00%	0.00%	100%	19
	Bias	1.41%	64.79%	33.80%	100%	71
	Cost	.00%	73.33%	26.67%	100%	30
Signal 6	Benchmark	0.00%	100.00%	0.00%	100%	27
	Bias	0.00%	93.22%	6.78%	100%	59
	Cost	0.00%	97.30%	2.70%	100%	37

Table 2.5: Distribution of effort level conditional on the signals in the benchmark and the bias treatments

Table 2.5 shows that the proportion of high effort level is higher for all possible signals except $s = 3$.

We can however notice that the employers predict different effort level across the benchmark treatment and both manipulation treatments ($p = 0.001 * **$ and $p = 0.003 * **$, Mann-Whitney rank-sum test).

These descriptive results show significant differences in the subjects' behavior between the treatments. Next, we study the employers' decisions and then analyze the employees' effort level choices.

2.4.2 Decision to bias information

Employers should send the message $m = 3$ as often as possible. The employers should therefore bias signals $s = 2$ and $s = 4$, and send $m = 3$.

The message distribution differs significantly between the benchmark treatment and the manipulation treatments ($p < 0.0001$) whereas the signal distribution does not differ between treatments. The employers manipulate information when allowed to do so. Only 21.90% of the messages correspond to the signal in the Bias treatment. On the

whole, we observe that in the Bias treatment, more biases (78.10% of the biased signals) than predicted (45.52% of the signals should have been biased) and 56% of manipulations are useless. The introduction of a manipulation cost inverses the manipulation effect: we observe 39.72% of the signals, whereas the predictions state that 48.06% of the signals should have been biased. Introducing a cost reduces the relative frequency of non-strategic biases (only 23% of the manipulations are useless).

Figures 2.1 and 2.2 display the distribution of the biases for each possible signal in the bias and in the cost treatments, respectively.
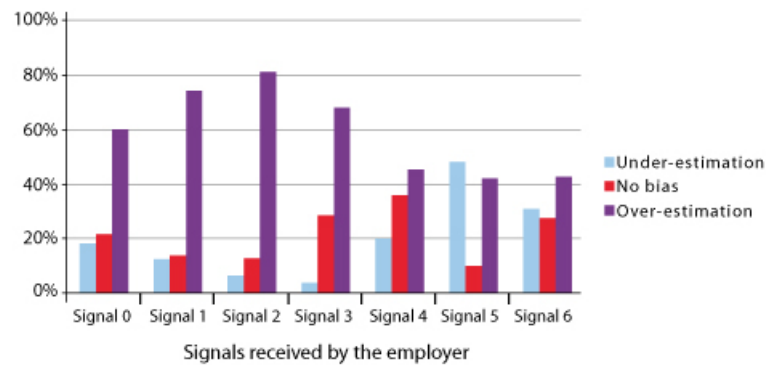

Figure 2.1: Distribution of messages for each signal in the Bias treatment
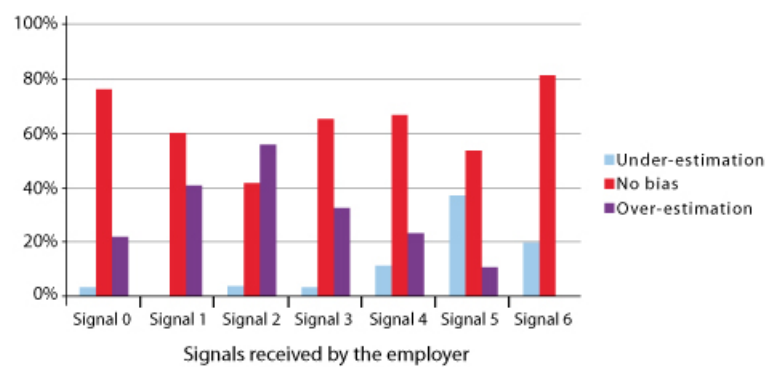

Figure 2.2: Distribution of messages for each signal in the Cost treatment

Figures 2.1 and 2.2 show different types of behavior depending on the signal received. Unsurprisingly, we observe that the employer overestimates the signals that indicate low-ability, $s \in \{0, 1, 2\}$, more frequently than the other signals. Employers underestimate the signals indicating that the employee has high ability, $s \in \{4, 5, 6\}$.

Table 2.6 displays the distribution of the predicted effort level for the signals observed by the employer and for the messages sent, by treatment.

Signal	Message	Benchmark treatment Predicted effort level			Bias treatment Predicted effort level			Cost treatment Predicted effort level		
		Low	Medium	High	Low	Medium	High	Low	Medium	High
	-1				100.00%	0.00%	0.00%	100.00%	0.00%	0.00%
0	0	100.00%	0.00%	0.00%	92.86%	0.00%	7.14%	100.00%	0.00%	0.00%
	1				87.18%	10.26%	2.56%	50.00%	0.00%	50.00%
	0				100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
1	1	100.00%	0.00%	0.00%	100.00%	0.00%	0.00%	80.65%	19.35%	0.00%
	2				75.00%	12.50%	12.50%	52.38%	4.76%	42.86%
	1				75.00%	12.50%	12.50%	100.00%	0.00%	0.00%
2	2	67.39%	13.04%	19.57%	68.75%	25.00%	6.25%	67.57%	21.62%	10.81%
	3				23.53%	24.51%	51.96%	14.00%	10.00%	76.00%
	2				100.00%	0.00%	0.00%	100.00%	0.00%	0.00%
3	3	0.00%	5.56%	94.44%	0.00%	6.25%	93.75%	0.00%	20.00%	80.00%
	4				2.63%	10.53%	86.84%	0.00%	20.00%	80.00%
	3				7.41%	3.70%	88.89%	11.11%	0.00%	88.89%
4	4	3.85%	19.23%	76.92%	0.00%	18.37%	81.63%	0.00%	29.09%	70.91%
	5				0.00%	79.03%	20.97%	0.00%	84.21%	15.79%
	4				2.94%	23.53%	73.53%	0.00%	0.00%	100.00%
5	5	0.00%	100.00%	0.00%	0.00%	42.86%	57.14%	0.00%	93.75%	6.25%
	6				0.00%	96.67%	3.33%	0.00%	33.33%	66.67%
	5				0.00%	83.33%	16.67%	0.00%	57.14%	42.86%
6	6	0.00%	96.30%	3.70%	0.00%	100.00%	0.00%	0.00%	93.33%	6.67%
	7				0.00%	96.00%	4.00%	0.00%	0.00%	0.00%

Table 2.6: Predicted effort levels

Table 2.6 shows that in the Bias treatment, when the employer observes the signal $s = 0$ and sends an unbiased message ($m = s = 0$), the low effort level is predicted 92.86% of the time. This proportion decreases when the employer sends a biased message ($m = s + 1 = 1$) to 87.18%. The employers send messages such that a low-ability employee thinks he is a medium-ability employee by overestimating the signals that suggest low-ability. The employers also try to change the high-ability employees' beliefs such that they think that they have medium-ability by sending underestimated messages. The employers expect that overestimated signals indicating low-ability and underestimation of high-ability signals induce higher effort. The employers understand

the effect of the manipulation on the effort.

However, figures 2.1 and 2.2 show that the employers overestimate the signal (proportion of overestimation = 59.31% in the Bias treatment and 30.83% in the Cost treatment) more often than they underestimate it (18.79% and 8.89%, respectively). However, they do not underestimate messages as often as optimally required. It seems that manipulation by overestimation or by underestimation is not perceived similarly by the employers. The results on the expected effort level cannot explain why underestimation is so rarely chosen.

In order to better explain information manipulation, we estimate a multinomial logit regression model[8] in which the explained variable is the decision to bias information (bias+1 vs bias−1, bias= 0 is the reference category). We estimate on the pooled data of the bias and cost treatments. The signals are introduced as dummy variables with signal 3 as the reference category. The period may be a proxy explaining the learning effect, that is, the employer's understanding of the employee's role and how the employee can be manipulated by the information sent. We also include a dummy variable for the Cost treatment to test whether the employers manipulate only because it is free. Table 2.7 displays the results of this estimation.

---

[8]As the subjects' decisions are repeated, we adjust the standard errors for the intragroup correlation, that is, cluster the standard errors on the individuals.

Dependent variable:Bias decision	Bias -1		Bias +1	
Signal 0 received by the employer	1.0579	(.7420)	-.4236	(.4051)
Signal 1 received by the employer	.9914	(.8328)	.4700	(.4594)
Signal 2 received by the employer	.8141	(.7990)	1.1470***	(.4363)
Signal 3 received by the employer	ref.		ref.	
Signal 4 received by the employer	1.0409	(.8019)	-.7212	(.4679)
Signal 5 received by the employer	3.1089***	(.7767)	.0368	(.4697)
Signal 6 received by the employer	1.7083**	(.7601)	-.8897*	(.4572)
Period	.0214	(.0333)	-.0010	(.0250)
Cost treatment	-1.7998***	(.3765)	-1.9601***	(.3224)
Constant	-1.8323**	(.8913)	1.0435**	(.5295)
Observations		705		
Wald Chi^2		167.44		
Prob > Chi^2		.0000		
Pseudo R^2		.1847		

Multinomial logit regression model on pooled data of Bias and Cost treatments. Standard errors in parentheses. Levels of significance: * 10%; ** 5%; *** 1%.

Table 2.7: Determinants of the bias decisions

Table 2.7 shows that the employers overestimate the signal $s = 2$ more often than the signal $s = 3$. The employers bias the signal in order to manipulate her employee's belief such that the employee with low-ability thinks he has medium-ability. However, the employers do not manipulate the signal $s = 4$ but bias the signals $s = 5$ and $s = 6$ in order to underestimate the employee's ability. These results are explained by the effort level predicted by the employer. When observing the signal $s = 4$ and sending the message $m = 4$, the employers predicts a high effort level in 81.63% of the cases, a medium effort level in 18.37% of the cases, and 0% of low effort level in the Bias treatment (respectively 70.91%, 29.09%, and 0% in the Cost treatment). When observing the signal $s = 4$ and sending the message $m = 3$, the employers predicts a high effort level in 88.89% of the cases, a medium effort level in 3.70% of the cases, and 7.41% of low effort level in the Bias treatment (respectively 88.89%, 0%, and 11.11% in the Cost treatment). The employer may be averse to send $m = 3$ when observing $s = 4$ because they predict a higher proportion of low effort level compared to sending $m = 4$. The

predicted effort level explains also that the signals $s = 5$ and $s = 6$. The proportion of high effort level predicted increases significantly when choosing to underestimate these signals compared to sending an unbiased message. For example, the proportion of high effort level predicted in the cias treatment when observing the signal $s = 5$ and sending the message $m = 5$ is 6.25% compared to 100% when sending the message $m = 4$. The Cost treatment has a strong negative impact on the bias decision. Time has no impact on both positive or negative manipulation, which indicates a small learning effect in this game.

These findings yield results 1 and 2:

Result 1: Employers bias information even in non-strategic situations.

Result 2: Introducing a cost of manipulation decreases the proportion of non-strategic biases.

2.4.3 Effort level

We predict that the employees choose a high effort level only when they receive message $m = 3$. We expect that the manipulation treatments lead to more frequent high effort level choices than the benchmark treatment. Figure 2.3 displays the distribution of the effort level for each message in each treatment.
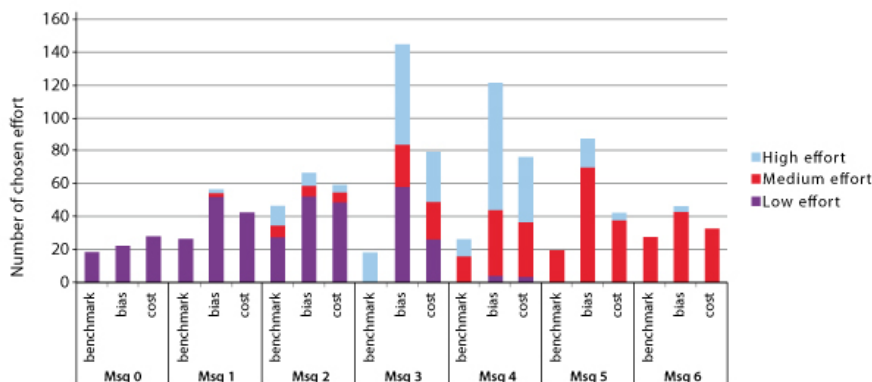
Figure 2.3: Distribution of effort levels conditional on the messages received in each treatment

In the benchmark treatment, the effort decisions are consistent with our predictions. 78.89% of the employees observing a message that reveals a high probability of being a low-ability employee ($m \in \{0, 1, 2\}$) choose a low effort level. Messages indicating high-ability ($m \in \{4, 5, 6\}$) lead to a medium effort level choice in 85.92% of the cases. moreover, the message $m = 3$ that reveals medium-ability induces the employee to choose a high effort level in 100% of the cases.

In the manipulation treatments, the effort decision conditional on the message $m \in \{0, 1, 2, 5, 6\}$ is equal to the predicted effort. $m \in \{0, 1, 2\}$ lead to a low effort level choice in 89.74% of the cases. 87.98% of the employees observing $m \in \{5, 6\}$ choose a medium effort level. However, the messages $m = 3$ and $m = 4$ lead to more variance in the effort decision. $m = 3$ leads to a high effort level choice (predicted effort) in 40.63% of the cases, low effort level choice in 37.50% of the cases, and medium effort level choice in 21.88% of the cases. $m = 4$ leads to medium effort level choice (predicted effort) in 37.56% of the cases, high effort level choice in 58.88% of the cases, and low effort level choice in 3.55% of the cases. This may be explained by the higher level of risk in these treatments for these messages. The message $m = 3$ in the benchmark treatment perfectly reveals medium-ability. However, in the manipulation treatments the message $m = 3$ corresponds to low-, medium-, or high-ability. The message $m = 4$ gives an uncertain information about the ability in all the treatments. However, the manipulation allowed in the manipulation treatments can make the employees believe that the message $m = 4$ is more uncertain in these treatments.

We estimate an ordered logit model to explain the effort level decision[9] in the pooled treatment data. We introduce in the regression two dummy variables indicating the Bias treatment and the Cost treatment. We include each message as a dummy variable with message 3 as the reference category. We expect the messages indicating low-ability $m = \{0, 1, 2\}$ or the messages indicating high-ability $m = \{4, 5, 6\}$ to decrease the effort level decision as compared to the message $m = 3$ indicating medium-ability. We exclude the message $m = 0$ because it leads to low effort level decision all the time. We also

---

[9]The effort level decision equals 0 if the effort is low, 1 if the effort is medium, and 2 if the effort is high.

exclude the messages $m = -1$ (that leads to a low effort level decision in all the cases) and $m = 7$ (that leads to a medium effort level decision in 92% of the cases) that appear only in the manipulation treatments. Risk aversion is given by the number of times the subjects choose option A (safe option) in the Holt and Laury test. Finally, we add two variables interacting the manipulation treatments and the risk aversion degree to control for the fact that the subjects are in a riskier situation in these treatments. A time trend is also included. Table 2.8 displays the results of this regression.

Dependent variable: effort level decision		
Message 1 received by the employee	-5.2761***	(.8440)
Message 2 received by the employee	-2.6439***	(.4396)
Message 3 received by the employee	ref.	
Message 4 received by the employee	1.2118***	(.4190)
Message 5 received by the employee	-.0018	(.3892)
Message 6 received by the employee	-0.2106	(.3619)
Bias treatment	1.2250	(.8320)
Cost treatment	.4221	(.6833)
Risk aversion	.1606**	(.0764)
Bias treatment*Risk aversion	-.2520**	(.1168)
Cost treatment*Risk aversion	-.1748	(.1201)
Period	-.0287	(.0192)
Observations	759	
LR *Chi*^{2}	164.30	
Prob > *Chi*^{2}	.0000	
Pseudo R^2	.2528	

Ordered logit regression model on the pooled data of our treatments. Standard errors are under parenthesis. Levels of significance: * 10%; ** 5%; *** 1%.

Table 2.8: Determinants of the effort level decisions

Table 2.8 indicates that being in the Bias treatment or in the Cost treatment has no impact on the effort decision. This result may have two reasons: the employees do not expect the strategic behavior of the employers or they cannot identify the truth from the lies.

The effort level choice increases in the employee's risk aversion. Moreover, the variable

interacting the Bias treatment and the risk aversion degree indicates that the aversion of risk has a less important impact in the Bias treatment than in the two other treatments. This result is lined with our predictions when releasinf the assumption of risk neutrality. All messages have the predicted impact except $m = 4$ that increases the effort level decision compared to the predicted effort level. A risk-averse employee chooses a high effort level when he receives the message $m = 4$. The data fits the theoretical predictions on effort level choices when the assumption of risk neutrality is relaxed.

These findings yield result 3:

Result 3: Controlling for risk attitudes, employees choose their effort level according to the message sent by their employer in the three treatments.

2.4.4 Efficiency

If the effort decision does not change across treatments, conditional on the message received by the employee, then biased information improves the employer's payoff. By manipulating the messages, employers increase the employee's effort and increases their payoff. In contrast, the employee's payoff decreases. Manipulation leads to some employees choosing a too costly effort level. However, does manipulation increase efficiency? We compare our treatments with an ordinary least squares model[10] to explain the employer's payoff, the employee's payoff, and efficiency[11] by the Bias treatment (Cost treatment) and the signal received by the employer.

As expected, the Bias treatment has a positive (negative) impact on the employer's payoff (employee's payoff) as compared to the benchmark treatment. However, on the whole it has no impact on total payoffs in contradiction with our predictions. Efficiency is not affected by manipulation. Distorting information mainly affects the surplus distribution. The introduction of the manipulation cost cancels this effect: the employer's payoff and efficiency decrease whereas no impact is seen on the employee's payoff as compared to the Bias treatment.

---

[10]The estimations are not reported here but are available upon request.

[11]Efficiency is measured as the sum of the employee's payoff and the employer's payoff

These observations yield results 4 and 5:

Result 4: Efficiency is not affected by manipulation when manipulation is free. Only the distribution of payoffs is affected.

Result 5: Information manipulation improves the employee's effort level and thus the employer's payoff, but at a cost for the employee.

2.5 Conclusion

In this chapter, we study the impact of biased information on motivation and efficiency when employers are better informed than the employees on their ability. We show that it is possible to increase the employees' effort levels by manipulating information. Biased messages allow employers to manipulate the employees' beliefs about their ability.

We ran an experiment to investigate the use and impact of biased information in an employer-employee relationship. We compare three treatments. The first treatment is the truthful information treatment. In the second treatment, the employer can bias the information sent to the employee. In the third treatment, we introduce a cost of manipulating the information.

Our results are consistent with our theoretical predictions. The employer sends biased messages to the employee. The employers use the manipulation even if they seem relatively averse to underestimating the signal on the employee's ability. Introducing a cost decreases non-strategic biases but manipulation persists. The employee chooses his effort level according to the message received, though the message may be biased. We also see that risk aversion plays a role in the employee's choice especially in the manipulation treatments where the information is more vague. Biases lead to both low-ability employees and high-ability employees providing a higher effort level. Finally, manipulation increases the employers' payoff without improving efficiency.

Our results may have some implications in organization management. Formal appraisals are an important tool for employers to increase employees' motivation. Gibbs (1991) analyzed the employer/employee relationship as a game of information and perceptions.

If the employer can credibly manipulate the information, since she does not always have an incentive to lie and the employee cannot tell if the employer is lying or telling the truth, the employer can increase the employee's effort level. Our experiment supports this idea. Our study provides new evidence on the use and impact of information on the employee's ability.

For future, it would be interesting to more precisely study the possible asymmetric behavior in transmitting overestimated or underestimated signals. Moreover, the aversion to lying found in Gneezy (2005) may be different depending on the positive or negative aspects of the lie. Moreover, understanding the impact of the performance appraisal is an important tool for firms to become more competitive. It may also be interesting to study how performance appraisals are used in different types of firms (public/private firms, for example) and how they differently affect the employees.

2.6   Appendix: Instructions (original in French)

You are about to participate in an experiment in economics during which you can earn money.

Most of your earnings during this experiment are expressed in points, with the following conversion rate:

20 points = €1

This experiment will have two parts. The earnings of each part will be added at the end of the session. You will receive €3 more for participating in this experiment. Your earnings will be paid to you in cash in private and in a separate room by someone who doesn't know the experiment.

It is forbidden to communicate with other subjects during the experiment. If you have any question regarding these instructions, please raise your hand.

PART 1

We thank you for filling out this form that enables you to earn additional money. Your screen shows ten decisions. Each decision is a paired choice between "Option A" and "Option B". You will make ten choices, but only one of them will be used in the end to determine your additional earnings. Let us explain how these choices will affect your earnings. Here is a ten-sided die that will be used to determine this payoff. The faces are numbered from 1 to 10 (the "0" face of the die will serve as 10). After you have made all of your choices, and when you come to the other office to receive your payment, you will throw this die to determine what your payoff is for the option you chose, A or B, for the particular decision selected.

Even though we ask you to make ten decisions, only one of these will end up affecting your earnings. However, you will not know in advance which decision will be used. Obviously, each decision has an equal chance of being used in the end.

Look at Decision 1.

Option A pays €2 if the throw of the dice is 1, and it pays €1.6 if the throw is 2-10. Option

72

B yields €3.85 if the throw of the dice is 1 and it pays €0.1 if the throw is 2-10.

Look at Decision 2.

Option A pays €2 if the throw of the dice is 1 or 2, and it pays €1.6 if the throw is 3-10.
Option B yields €3.85 if the throw of the dice is 1 or 2 and it pays €0.1 if the throw is 3-10.

The other decisions are similar, except that as you move down the table, the chances of a higher payoff for each option increase. In fact, for Decision 10 in the bottom row, the dice will not be needed since each option pays the highest payoff for sure, so your choice here is between €2 and €3.85.

To summarize,

- you will make ten choices. For each decision row, you will have to choose between Option A and Option B. You may choose A for some decision rows and B for other rows. You may change your decisions and make them in any order.

- the software will randomly select one of the ten decision which will be used for paiements.

- Then, you will throw the die again to determine your money earnings for the Option you chose for that Decision.

Earnings (in Euros) for this choice will be added to your earnings, and you will be paid all earnings in cash. If you have any question, please raise your hand. Your questions will be answered in private. Please do not talk with anyone.

PART 2

There is two types of subjects (in equal number): employers and employees. Your type is randomly assigned to you and you keep it for all this part. Your type will be written on your screen at the start. The employee has to produce an effort. He doesn't know his ability. The employer has to send an information to his employee about his ability.

During this part, there are 20 periods.

Period description

At each period, paires are randomly formed with one employer and one employee. You will never know your co-participant identity.

1. The employee ability is randomly determined $1, 3, 5$. Neither the employee nor the employer know the ability level.

2. The employer receives a signal about his employee's ability. The following table shows you for each ability level, the possible signals in the 2nd row.
   For example, if the real ability is 1, the employer can receive the signal 0, 1 or 2. But he can also receive signal 2 when the real ability is 3.

3. After observing the signal, the employer sends a message to his employee about his ability. The employer can choose to send to his employee the signal he received or a different one. If he chooses to send a different signal, the signal can deviate only about 1 point (+1 or -1).
   For example, if the signal received by the employer is 1, he can decide to send message 0, 1 or 2. If he receives signal 4, he can send message 3, 4 or 5.
   The different possible messages are described in the following table.

Employee's ability (unknown)	Signal (observed by the employer)	Possible feedbacks (send by the employer)
1	0	-1, 0 or 1
	1	0, 1 or 2
	2	1, 2 or 3
3	2	1, 2 or 3
	3	2, 3 or 4
	4	3, 4 or 5
5	4	3, 4 or 5
	5	4, 5 or 6
	6	5, 6 or 7

4. After observing the message sent by his employer, the employee chooses an effort level: low, medium or high effort.

5. The earnings depend on the chosen effort and the real employee's ability. Earnings are described in the following table.

Ability	Effort (choosen by the employee)	Employee's wage	Employer's profit
	low	9	5
1	medium	8	7
	high	4	9
	low	9	7
3	medium	8	9
	high	11	11
	low	9	9
5	medium	15	11
	high	11	13

Example 1: if the ability level is 3 and the employee chooses a high effort, the employee earns 11 points and the employer earns 11 points.

Example 2: if the ability level is 1 and the employee chooses a medium effort, the employee earns 8 points and the employer earns 7 points.

You observe your earnings for this period. And an other period starts.

At each new period, new draws determine:

- your co-participant

- the employee's ability

- the signal received by the employer

Moreover, each employer has to predict employee's effort choice. 20 more points will be added to their earnings if they predict correctly at least 50/100 of the periods.

To summarize, at each period:

- An ability level is randomly assigned 1, 3, 5 at each employee. Neither the employee nor the employer know this ability level.

- The employer receives a signal about his employee's ability.

- The employer sends a message to his employee: the true signal, the signal -1, or the signal +1.

- The employee observes the message and chooses an effort low, medium, high. The employer predicts the chosen effort.

- You observe your earnings.