

General Introduction

Evaluation interviews are a usual practice in firms. Addison and Belfield (2008) show that 61% of British establishments in the 2004 Working Employment Relations Survey claim that all their workers are involved in performance appraisals. Brown and Heywood (2005) make a similar finding with Australian data. Diaye, Greenan, and Urdanivia (2007) find that 52% of employees in French manufacturing firms with over 50 employees have an annual individual evaluation interview. Formal appraisals allow employers to select, pay, promote, fire and motivate their employees. Evaluations are therefore an important management tool for the entire organization. Hence, understanding how evaluation impact on both employer's and employee's behavior is an important issue for the human resources management.

The purpose of evaluations is to measure ability, skill and effort that determine the employee's performance. However, evaluations are difficult and costly. There are many reasons why it would be difficult to make an accurate evaluation. For example, team work may give rise to some free-riding problems, performance may be partly due to luck and the worker's contributions can often be quite difficult to quantify. An inaccurate evaluation produces negative outcomes such as raising the employee's payoffs to compensate this risk, reducing the worker's motivation, or generating strong incentives to do the wrong things.

Lazear and Gibbs (2009) list two measures of performance: the quantitative method and the subjective evaluation. The quantitative method based on numeric measures of performance makes the relationship with compensation easier and is perceived as being

more objective. However, the latter advantage is not as obvious. Quantitative measures of performance may be manipulated by the employee by choosing observable actions that do not improve firm value. Moreover, performance may not be precisely measured. The firm therefore has to choose the scope of the measure when choosing to use the quantitative performance evaluation. A broader measure of performance takes into account most aspects of performance in the evaluation. However, it also includes aspects of the job that are uncontrollable by the employee making the incentives schemes riskier for the employee. The performance is quantified more easily with a precise measure of performance since it excludes most of the uncontrollable part of the employee's job reducing the risk of an inaccurate evaluation. However, it also increases the employee's incentive to manipulate the evaluation by putting his effort only in the observed tasks.

The other method to measure performance is a subjective evaluation. Adding subjectivity in evaluation enriches the quantitative performance measures by taking into account some dimensions of the job hard to quantify. It can also decrease the employee's manipulation of the evaluation. An employer who learn that an employee manipulates the quantitative measure, will adjust the evaluation in consequences. Thus, the manipulation has at least less impact on the evaluation. Moreover, subjective evaluations may improve communication since subjectivity imposes a requirement that the supervisor has to evaluate his employee more often and thus, give him advice more frequently. Communication is clearer and the employees will be more likely to trust subjective evaluations, making them more effective. However, this method leads to evaluation distortions such as favoritism, leniency bias (reluctance to give poor performance scores) or centrality bias (compression of performance scores). These biases cause inefficiencies for the firm. An employee who feels discriminated may quit, resulting in a loss of human capital. Biases may reduce the motivation and thus the effort of employees. It can even create wrong motivation. For example, an employee may spend much of his time trying to flatter his boss instead of working. Subjective evaluation can also create a yes man phenomenon (Prendergast and Topel (1993)) that is to say an incentive for workers to conform to the opinion of managers.

Economic literature (see Lazear (1995) or Prendergast (1999)) develops different solutions to solve the problem of potentially biased evaluations. First, the compensation of many employees may be less related to performance at least in the short term. One reason to adopt a fixed wage scheme is that it avoids the agent shirking and putting his efforts only in the visible tasks on which the payoff depends (see Holmstrom and Milgrom (1991)). Secondly, the firm may implement bureaucratic rules. For example, firms may use forced ranking in their performance appraisal systems. The managers have to categorize workers into a certain distribution (see Prendergast and Topel (1993)). Another solution, generally used in the real world, is the implementation of a long term relationship and promotion based reward. However, Gibbs (1991) considers a different approach and shows that in some cases, the manager should manipulate the information sent to his employee in order to increase effort.

This dissertation aims to study the evaluation in a context of imperfect information. A manager wants, first, to measure the employee's performance. The result of the performance appraisal, i.e. which may be more or less accurate but never perfect, provides the employer with a new set of information to assess the employee's productivity. The employer updates his initial belief with this new information. The first chapter studies how people update their beliefs with imperfect signals. The manager will then want to transmit some of his informational advantage to the employee in order to motivate him to exert a higher level of effort. This question is raised in the second chapter. Based on Gibbs (1991), we study how manipulated feedback impacts the choice of effort. The third chapter will then study supervisors' choices of different types of lies when they evaluate workers' performance.

The first essay analyses Bayesian updating behaviors when people receive a signal that restricts the number of possible states of the world. The employer by observing the employee working in the job has an initial belief about the employee's productivity. Moreover, the result of the performance appraisal is an imperfect signal of the produc-

tivity. Suppose that the signal indicates that the productivity is above a certain standard. This information reduces the possible levels of productivity without giving the employer any information on how far the employee's productivity is from the standard. Can an employer who imperfectly observes the employee's productivity be helped by a signal indicating whether the productivity reaches a certain threshold?

We investigate how subjects update their belief about the state of the nature after receiving this type of signal. We create an experiment on individual choice. The subject assesses an urn's composition. The urn is composed of 20 balls, either yellow or blue in color. Therefore, 21 possible urns may exist. The subject successively receives two types of information: first, a sample of observations from the urn and, second, a signal that restricts the set of the possible urns. The signal indicates if the urn is composed by more or less X yellow balls. The subject is asked twice what he estimates the composition of the urn to be, a first time observing the sample, i.e. his initial belief, and a second time after receiving the signal, i.e. his posterior belief. The subject's initial belief can include the new subset of possible states of nature. The objective of this paper is to analyze how the signals that reduce the number of possible states of nature are used by the subjects to update their belief about the state of the nature.

We find that the signal helps the subjects to assess the composition of the urn. Unsurprisingly, an invalidating signal significantly increases the probability of correct estimation. However, more surprisingly, a validating signal (somehow confirming the initial belief) reduces the frequency of correct estimations compared to no signal. In 49% of the cases, subjects change their estimation even if they receive a validating signal. A more precise analysis shows that this unpredicted updating behavior after receiving a validating signal occurs when the signal does not help the subjects enough and leaves the subjects with a high uncertainty level. This work shows that on average the employer is able to better assess the employee's productivity when observing the results of the performance appraisal which indicates whether the productivity is above a specified standard. However, this type of signal may also lead to some assessment mistakes when the new information leaves the employer highly uncertain, i.e. when the standard is

either reached too easily or with much difficulty.

Second, the evaluation interview following the performance appraisal, allows the employer to transmit information to the employee on his performance and thus increasing his motivation.

The literature on feedback shows that providing feedback increases the agent's performance. Azmat and Iriberry (2010b) study the effect of providing relative performance feedback on individual performance when agents are rewarded according to their absolute performance. They find that performance increases by 17% when relative performance feedback is provided. Bandiera, Larcinese, and Rasul (2010) measure the causal effect of interim feedback on an 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 students. Their results show that the provision of feedback has a positive effect on the period two scores. Some studies although show that providing feedback is not always optimal. In a tournament setting, Ederer (2010) studies the effort choices of two agents when they cannot observe their performance. The principal privately observes the difference in performance between the two agents. Then, the principal chooses whether or not to reveal this information to the agents. Ederer (2010) shows that it can be better to conceal information on the agent's relative performance. 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 and a common factor. The subjects are asked their belief about the individual factor and the common factor once before and once after observing others' outcomes. The result show that withholding information on social comparison (information of others) may be optimal because social comparison information induces a negative correlation among agents with regard to self-confidence, effort and outcome. Depending on the type of feedback, the impact of feedback is not clearly determined. The principal has to decide how and what kind of information to transmit during the performance appraisal. Evaluations allow the employer to increase her knowledge of the employee's ability

as shown previously. It creates an information asymmetry between the principal and the agent at the principal's advantage. The employer knows the employee's ability better than the employee does. The feedback given during the performance appraisal may be a strategic tool for the employer. Gibbs (1991)'s idea is that a principal may influence, through a manipulated feedback, the agent's behavior in order to increase his effort and motivation. In a race, the front runner has less incentive to run fast if he knows that he is far ahead. Identically, the underdogs have less incentive to run fast if winning is hopeless. The runners with the highest incentives are the runners in the middle. Therefore the coach may report a lower distance between the front runner (or the underdog) and his opponents in order to give him the incentive to continue to exert high effort. Farrell and Rabin (1996) introduce for the first time the concept of cheap talk, i.e. the signal that does not directly affect the payoffs and has an impact on global efficiency. Crawford and Sobel (1982) develop a model in which a better-informed sender sends a signal to a receiver who then takes an action. They find that a noisy signal may be an equilibrium.

In the second essay, we focus on the feedback an employer sends to the employee during the evaluation interview. It deals also with a situation of imperfect knowledge about the state of nature. Here, the state of nature corresponds to an agent's ability and is not known by anyone (neither the principal nor the agent). It can be either low, medium or high. A principal receives imperfect information about the agent's ability, i.e. a signal, and has to send a message to him. This signal may perfectly reveal the agent's ability or may leave the principal uncertain about the agent's ability. The principal then sends a message to the agent who then chooses a level of effort. Both effort and real ability affect both players' payoffs. In the Benchmark treatment, the principal has to send a message corresponding exactly to the signal she received. In the Biased treatment, she can without any cost, manipulate the message. The manipulation is restricted to a small bias compared to the signal (± 1 from the signal). In a third treatment, Cost treatment, the principal bears a small cost for manipulating the message. In the two treatments where manipulation is allowed, the principal should send a message corresponding

to a medium ability and the agent should respond in choosing a high effort. The cost supports by the principal to manipulate does not change these predictions. The cost treatment tests whether the principal will manipulate the information merely because it is free or whether she supports a moral cost for lying (Mazar, Amir, and Ariely (2008)). The result shows that principals do not hesitate to manipulate the feedback. This manipulation increases the agent's effort and thus the principal's profit. Moreover, we observe that there is more overestimation than underestimation in the choice of manipulation. This essay shows that employers use their informational advantage in order to manipulate the employee's beliefs about his ability and thus the employee increases his effort. Moreover, it seems that manipulation in evaluating someone leads to an asymmetric behavior regarding whether it increases the performance (good news) or decreases it (bad news). This manipulation may be interpreted as a lie. These results show that a fraction of people lie at the expense of someone else.

The economic literature on lies not only shows that some people lie but also that a significant proportion of subjects are averse to lying. For example, Gneezy (2005) studies, in a sender-receiver game with conflict of interest, the lying behavior. Two distribution payoffs exist, namely A and B. The sender observes the payoffs of each player in both distributions. She can then communicate to the receiver which distribution maximizes the receiver's payoff. After observing the message, the receiver chooses either A or B. The payoffs are implemented according to the receiver's choice. The author tests the role of monetary incentives in the lying decision by varying the consequences of lies. The results show that people are sensitive to their profit when deciding to lie. More surprisingly, people are also sensitive to the negative consequences of their lies. Controlling with a dictator game, the author concludes that it is not only the care of others that motivates behavior but also the aversion to lying.

The third essay focuses on lying behavior in evaluation. More precisely, we study by means of a real effort experiment, lying behavior and the nature of lies when a supervisor evaluates an agent. In our experiment, a supervisor has to report to the experimentalist

the output of an agent who is performing a real effort task. The supervisor should always report a performance level that corresponds to a certain threshold. Our design allows us to observe different lies, based on Erat and Gneezy (2009): by reporting a performance different from the actual performance, supervisors can either improve both players' payoffs (Pareto white lie), improve their own payoffs and reduce the agent's payoff (selfish black lie), or only improve the agent's payoff (altruistic lie). Our main behavioral hypotheses are: i) people tell the truth or lie but not as much as predicted if they are lie-averse or guilt-averse; ii) guilt aversion occurs if the correlation between second order beliefs and propensity to lie is significant and positive.

We designed two treatments. In the baseline, the agents and the supervisors are located in different rooms. Once the agents have performed the task, the answers sheets are transferred to the other room and each supervisor counts the number of correct answers of the agent. Then, each supervisor reports orally to the experimenter a number of correct answers and gets paid according to the verbal report. Meanwhile, the agents are asked to predict how many correct answers will be reported by their supervisor. Once all the supervisors have left the laboratory, each agent gets paid individually. The Second Order Belief (SOB) treatment is similar to the baseline except that the supervisors are asked their belief about their agent's prediction. We find that the majority of supervisors are willing to bias their report in order to earn more. Few supervisors tell spite black lies and altruistic white lies (according to Erat and Gneezy (2009)'s terminology), while both selfish black lies and Pareto white lies are frequent. In most situations, making the second order beliefs more salient affects neither the propensity of lying nor the nature of lies. The analysis of beliefs show that workers overestimate their actual performance and that supervisors anticipate a higher number of lies predicted by their workers than what they actually do. An econometric analysis shows that the decision to lie is positively correlated with the second order belief. The supervisor who believes that his worker expects him to tell the truth are less likely to lie than a supervisor who believes that his worker expects him to lie. Thus, supervisors seem more guilt-averse than lie-averse.

The evaluation of employees is essential in the incentives schemes in a firm. However, the performance measures are often imperfect and lead to biases in evaluation. Among the potential reasons for those biases, this dissertation studies more precisely two of them: problem in updating beliefs and voluntary manipulation of the evaluation. The experimental method is particularly relevant here since it is impossible to collect accurate data on this topic in the field. Of course, a potential drawback is related to external validity of laboratory experiments. With this reservation in mind, we can however suggest that uncertainty allows biases in evaluation that are not detrimental to the firm in the short run. A natural extension would be to check the long-term effects on trust and productivity of those biases.

First, updating beliefs after an imperfect signal, such as an inaccurate measure of the employee's performance, seems to lead to more evaluation errors if the accuracy of the signal is really low. The effects of voluntary manipulation may be positive. We show that in some cases a small manipulation may increase the effort choices of the agents. However, in our last experiment, supervisors are willing to disregard the actual performance in order to increase their own payoff.

