

UNIVERSITY OF CALIFORNIA
Los Angeles

**Essays on Self-Perception and Information
Revelation**

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Economics

by

Seda Ertac

2006

© Copyright by

Seda Ertac

2006

The dissertation of Seda Ertac is approved.

Michael S. Y. Chwe

Ichiro Obara

Leeat Yariv

David K. Levine, Committee Chair

University of California, Los Angeles

2006

To my family.

TABLE OF CONTENTS

1 Social Comparisons and Optimal Information Revelation: Theory and Experiments	1
1.1 Introduction	1
1.2 Model	10
1.2.1 Setup	10
1.2.2 Principal's Payoff	13
1.2.3 Information Policy	14
1.3 Results	15
1.3.1 Optimal Information Policy with Fixed, Independent Contracts	15
1.3.2 Dependent Contracts	28
1.3.3 Optimal Information Policy with Endogenous Contracts	34
1.3.4 Numerical Illustrations	37
1.3.5 Discussion	42
1.4 Extensions	44
1.4.1 Different Priors and Overconfidence	44
1.4.2 Partial or Selective Information Revelation	46
1.4.3 Informed Principal	47

1.5	An Experimental Test of the Theory	48
1.5.1	Experimental Design	48
1.5.2	Hypotheses	53
1.5.3	Procedures	55
1.5.4	Results	56
1.6	Conclusion	67
1.7	APPENDIX	71
	References	91
2	An Experimental Study of Ego-Relevance in Information Pro-	
	cessing	94
2.1	Introduction	94
2.2	Experimental Design	98
2.2.1	Overview	98
2.2.2	Procedures	102
2.2.3	Discussion	104
2.3	Results	107
2.3.1	Performance Rounds	107
2.3.2	Non-Performance Rounds	113
2.3.3	Direct Comparison of Performance and Non-Performance Rounds	116

2.3.4 Discussion and Extensions	117
2.4 Conclusions	120
2.5 APPENDIX	123
References	134

ACKNOWLEDGMENTS

There are many people who have had an impact on this dissertation; some with their academic guidance, some with their personal support, and many with both. First and foremost, I feel truly grateful to have had David Levine as an advisor. I cannot thank him enough for always giving me the most concrete, prompt and right advice in every matter, for teaching me how to be efficient as an academic, and most importantly, for his unwavering belief in and never-ending patience with me. Leeat Yariv has been an immediate role model and a constant source of support during my PhD, academically and otherwise, with her broad vision and empathy. I am extremely indebted to Bill Zame for his invaluable guidance with both my academic work and my career choices, and for being there to give me crucial advice and support especially throughout the job search process. I would like to thank Ichiro Obara for always being willing to spend time to discuss my research, and for his supportiveness and insight.

The research in this dissertation benefited greatly from discussions with Michael Chwe, Christian Hellwig, Hongbin Cai, Jozsef Molnar, John Riley and Jean-Laurent Rosenthal, and from the comments of the participants of the UCLA theory seminar, as well as the various conferences that the work was presented at. I would like to thank Raj Advani for designing the software used in the experiments conducted for this dissertation, and Walter Yuan for helping me with its implementation. The experiments in chapters 1 and 2 were generously funded by the Russell Sage Foundation (Small Grants in Behavioral Economics program) and the National Science Foundation (Doctoral Dissertation Grants program). Support from the UCLA Graduate Division (Dissertation Year Fellowship) is also gratefully acknowledged.

I am indebted to many friends for valuable discussions about my work, especially David, Joon, Matias and Yossi, and I would like to thank Alper, Emre, Eren, Burcu, Gonzalo, and Latika for making my time at UCLA a much more pleasant experience. I would also like to thank Lora Clarke and Estela Hopenhayn for making administrative matters as painless as can be.

Inci Gumus has been with me every step of this journey, and I feel lucky to have found in her a colleague, a roommate, and a lifelong friend.

Most importantly, I would like to thank my family for believing in me, and for always being there, no matter the distance.

VITA

January 1, 1979 Born, Izmir, Turkey

2000 B.A., Economics
Bilkent University, Ankara, Turkey

2000–2004 Graduate Fellowship
Department of Economics
University of California, Los Angeles

2001–2005 Teaching Assistant, Associate, and Fellow
Department of Economics
University of California, Los Angeles

2002–2003 Outstanding Teaching Assistant Award
Department of Economics
University of California, Los Angeles

2003 M.A., Economics
University of California, Los Angeles

2003 C.Phil., Economics
University of California, Los Angeles

2004–2005 Dissertation Year Fellowship
Graduate Division
University of California, Los Angeles

2005 Russell Sage Foundation Grant for Research in Behavioral Economics

- 2006 Russell Sage Foundation Grant for Research in Behavioral Economics
- 2006 National Science Foundation Doctoral Dissertation Grant

PUBLICATIONS AND PRESENTATIONS

Ertac, Seda (June 2005). “The Value and Acquisition of Ego-Relevant Information”. Presented at the 2005 International Meeting of the Economic Science Association, Montreal, Canada.

Ertac, Seda (September 2005). “The Value and Acquisition of Ego-Relevant Information”. Presented at the 2005 North American Meeting of the Economic Science Association, Tucson, AZ.

Ozcan, Kivilcim Metin, Asli Gunay, and Seda Ertac (2003), “Determinants of Private Savings Behavior in Turkey,” *Applied Economics*, 35(12), 1405-1416.

ABSTRACT OF THE DISSERTATION

Essays on Self-Perception and Information Revelation

by

Seda Ertac

Doctor of Philosophy in Economics

University of California, Los Angeles, 2006

Professor David K. Levine, Chair

In many economic settings, individuals do not have perfect information about their own traits, such as ability, and make decisions based on their beliefs about themselves. This dissertation, in broad terms, focuses on the theoretical and experimental analysis of the effects of “imperfect self-knowledge” on different types of behavior, and its implications for organizational design.

The first chapter focuses on social comparison information. When performance is affected by shocks that are common to everyone, performance comparisons with others affect perceived ability, and hence the motivation to exert effort. Given this effect on beliefs and behavior, I explore the organizational design problem where the amount of interim information disclosed to agents about each other’s performances is a choice variable of the principal. I find that the optimal information disclosure policy is fundamentally related to the degree of substitutability of different agents, and the contract structure. With exogenous contracts, if agents’ performances are sufficiently complementary from the perspective of the principal, withholding social comparison information may be optimal. However, when the wage scheme can be manipulated, full information revelation, coupled with

a “cooperative” incentive scheme, is the optimal policy. Findings from a laboratory experiment confirm many of the theoretical predictions, and the results are potentially applicable to many real-world situations, ranging from the revelation of grade distributions in classrooms to interim performance evaluations and team formation policies in firms.

In many settings, individuals are confronted with decision problems that involve information relevant to their self-image. The second chapter experimentally explores whether individuals process information differently depending on the information being self-relevant or not. The experiment implements two information processing tasks that are identical from a theoretical perspective but only differ in the type of information given to subjects (performance feedback versus information within the context of a purely statistical updating problem). Interestingly, the experiment yields no evidence for self-serving use of performance-related information. On the contrary, subjects tend to overweigh unfavorable performance information, and end up with overly pessimistic beliefs. This is in contrast to their updating behavior in the corresponding ego-irrelevant setup, where the feedback does not have any psychological connotation. In this case, subjects display an entirely different type of departure from Bayesian behavior and underweigh new information, consistent with conservatism in information processing.

CHAPTER 1

Social Comparisons and Optimal Information

Revelation: Theory and Experiments

1.1 Introduction

Self-perception is an important determinant of behavior in many settings that are characterized by imperfect self-knowledge. In the workplace, especially in the early phases of employment, individuals rarely have a precise idea of how suited they are for a particular task or career, and their prospects of advancement in it. In educational settings, students are often unsure about their propensity to succeed in a particular course or area of study, and tend to develop an academic self-concept over time. In the presence of imperfect self-knowledge, individuals tend to use previous successes and failures to learn about their unknown traits such as ability. Perceived ability, in turn, determines the return to taking a particular course of action, and affects crucial decisions such as whether or not to undertake a task, how much effort to exert, and whether to persevere or drop out in response to a failure. Given the effects on behavior, how much and what kind of interim performance information to give agents in order to influence their beliefs and maximize the potential for future success is an important organizational question. The issue of performance feedback, in fact, has long been a central issue in the management literature (e.g. Barr and Conlon (1994), Gibbs (1991),

Ilgen, Fisher and Taylor (1979), Morrison and Cummings (1992)).

Social comparison information is a particularly important type of information that affects self-perception. In performance settings, individuals often compare their performance or progress with others doing the same task. Failing in a task, for instance, usually has different implications on what we think about ourselves, depending on whether everyone else succeeded or not. Such dependence is especially well-documented in educational settings: gifted students in special programs for the gifted have been found to have lower levels of perceived ability than gifted students in normal programs (e.g. Zeidner and Schleyer (1999)), and academic self-concept is known to depend crucially on one's peer group, which has been termed the "big-fish-little-pond effect" (Marsh (1984), Marsh and Parker (2000)). Social comparisons also have important effects on behavior. An individual's motivation to exert further effort or to persevere in a task, her decision of whether to choose a particular career path etc. are likely to be influenced by how well she has done relative to others facing the same situation. The effects of comparisons on behavior suggest that manipulating the availability of social comparison information may potentially be an important tool in organizational design, for a principal who cares about the performance of multiple agents doing the same task. This paper presents a new theoretical model in which social comparisons affect effort through their effect on beliefs and self-perception, and analyzes, using a multi-agent framework, whether and when it would be optimal for a principal to release social comparison information to enhance future performance by her agents. The main predictions of the theoretical model are then tested in a laboratory experiment.

The effects of social comparisons on self-perception and their implications for organizational design have not been studied much in economics, although they

constitute an important part of the relevant literature in psychology and management¹. In many of the economic studies, social comparisons are taken to refer to relative income, wage, or status differences (e.g. Clark and Oswald (1996, 1998), Falk and Knell (2004), Hopkins and Kornienko (2004), Ok and Kockesen (2000)), and are usually built on the assumption that individuals are exogenously motivated by an intrinsic utility from rank, i.e., they get disutility from being worse-off than others, or utility from being ahead. Rather than assuming such an external, exogenous payoff dependence, this paper models social comparisons as a source of information. This is motivated by the idea that especially when the comparison concerns task performance, the impact of the comparison on behavior usually hinges on what the comparison has to say about one's aptitude in the task. For instance, it is usually not the same to be outperformed in an exam for which one has not studied at all while peers have been studying very hard and vice versa, or to be outperformed by someone who has had more previous training in the task. Such examples suggest that the informativeness of the comparison about ability matters. This informational aspect, which has thus far been neglected in economic models of social comparisons, is likely to be crucial in settings such as the workplace or the classroom, where individuals do not know their abilities perfectly and performance comparisons are commonplace.

The paper consists of two main parts: the first part presents the model and the theoretical results, and the second part reports on the design and results of an experiment conducted to test the theory. The theoretical model formalizes the idea that the effects of social comparisons are linked fundamentally to the ability inference that they lead to, and presents a new theory of social comparisons where comparisons affect behavior through their effect on an individual's per-

¹Suls and Wheeler (2000) provide a good selection from the vast psychology literature on the topic.

ceived *absolute* ability. I take a setting where agents do not know their abilities in a task, but use previous performances in the task to make inferences. Others' performances are valuable information for an individual who does not know her own ability, because of determinants of performance that are common to everyone performing the same task. A good example of such common shocks, taken from the realm of education, is the difficulty of the test or the lenience of the grader, which is unknown. When performance depends both on the common shock and the individual's ability (which is an idiosyncratic shock), learning how well others did is useful for separating the effects of the two factors in the inference process. If everyone succeeds in a test, it is more likely that the test was easy. Therefore, a failure when others succeed more strongly implicates low ability than in the case where one learns that others also have failed, or in the case where she does not know about others' performances at all. In this sense, the presence of correlated shocks, coupled with imperfect self-knowledge, leads to the result that "self-concept" is relatively formed². One's self-concept, in turn, affects her perceived return to exerting effort or investing in the task and thereby her effort level and future performance.

Given the potential effects of social comparisons on behavior, an important question that arises is whether a principal could use comparisons to generate higher effort and performance by her agents. There are several different ways in which this can be achieved. The first one is direct revelation or withholding of information, such as making the results of interim performance evaluations (e.g. pre-promotion or mid-career reviews) public or private in firms, or revealing the whole grade distribution as opposed to each student's own test score only in

²I use the term "self-concept" interchangeably with "perceived ability" or "self-confidence" throughout the paper. All definitions refer to the distribution of an individual's beliefs over her ability.

a classroom.³ A more indirect way to achieve the same effect is to form work groups or teams consisting of agents with particular characteristics (manipulating the type and content of the information that agents will likely receive, rather than directly withholding the information). Another indirect way is to use uniform (as opposed to differentiated) wage schemes, or to make performance bonuses public or keep them private, the idea being that agents can make an indirect inference about their abilities through the wage offered to themselves and others. In the current paper, I focus on direct performance feedback policies and analyze, using a two-period, two-agent model, whether it would be beneficial for a principal to create an environment where agents can observe the interim performances of each other, as opposed to one where they can only observe their own performance, or one where they cannot observe any interim information at all.

I start the analysis with the case where contracts are exogenous and only the information policy can be chosen, and then allow the principal to choose the compensation scheme in addition to the information policy. The main result is that the ranking of different information policies in terms of expected payoffs to the principal is fundamentally linked to (1) the degree of complementarity of the agents' performances in the principal's payoff, (2) the amount of discretion that the principal has in manipulating contracts. Complementarity in this context means that the principal gets extra payoff when both agents succeed at the same time, i.e. her payoff when both of the agents perform well is more than twice as high as her payoff when only one of the agents performs well. With continuous performance levels, this would correspond to a preference for a moderate performance by both agents instead of a very good performance by one agent and a

³For instance, the online grade entry/viewing software of UCLA, "Gradebook", gives the instructor the option of making the whole distribution of scores available to the students as opposed to their own scores only.

very bad one by the other. The importance of complementarities is likely to vary across organizations, depending on organizational goals and the production function. The funding that schools obtain from the government, for instance, may depend on the number of students that pass a certain threshold, rather than the achievement level of the best students, in which case complementarities would be important. Another setting characterized by strong complementarities is industries with “weakest-link”-type production functions, where the lowest effort or output matters immensely. With exogenously given contracts, I find that withholding social comparison information and revealing only one’s own performance can be better than revealing social comparison information in addition to own information, if the complementarities between the agents in the principal’s payoff function are strong enough. This exogenous contract setting captures situations where the compensation scheme is determined by another authority or institution, but a division manager in a firm, a coach or a teacher, who has no discretion in setting the wage levels and does not take into account the monetary costs of inducing effort, aims to maximize the motivation and performance of her agents using the amount of interim information they receive.⁴ When contracts can be fully manipulated along with the information policy, however, it becomes optimal to allow social comparisons even for strong levels of complementarity, since the shape of the contract can now be adjusted to mitigate any negative effects of social comparisons on motivation.

The second part of the paper reports results from a laboratory experiment designed to test the validity of the “informational theory” of social comparisons, i.e. to find out whether individuals use social comparison information in the correct way (if at all) when updating their beliefs, and to study how decisions are affected

⁴Indeed, attempts at confidence-management are commonly observed on the part of people who have less control over material incentives—parents, coaches, friends etc.

by the information received. An experimental analysis is very useful for assessing the validity of the theoretical model, since it allows beliefs and decisions to be observed, making it possible to pit them against optimal benchmarks, as well as analyze how beliefs and decisions respond to information. I closely replicate the structure of the theoretical model in the experimental design, and use a within-subject framework where participants are asked to submit their beliefs and make a decision once before and once after seeing others' outcomes. I find that the direction of the belief updating in response to social comparisons is generally in line with the main theoretical prediction, i.e. higher outcomes by the others lead to lower posteriors keeping own outcome fixed, although the absolute level of beliefs are sometimes different from their Bayesian counterparts.

The current paper fits into the broader research agenda of using economic models to study self-perception and its implications for organizations. It is therefore related to several branches of the literature in terms of its content and method of modeling. Because of the focus on self-perception, it is linked to the recently growing literature that studies imperfect self-knowledge in the context of organizations (e.g. Koszegi (2005), Gervais and Goldstein (2004), and especially to studies on "confidence management" using policies such as the reward structure or grouping (e.g. Benabou and Tirole (2003)). Since it analyzes incentives and performance in multi-agent settings, it is related to Prat (2000) who analyzes optimal team formation, and to Che and Yoo (2001), who study optimal incentives for teams.

Although the interpretation is based on social comparisons and ability perception, the type of model used in the paper can be applied to other situations in industrial organization, such as the decision of a small business owner of whether or not to persist operating in an industry in response to receiving information about

other firms' profits. In this sense it is linked, on an abstract level, to studies that focus on information revelation policies in different settings such as auctions and oligopoly (e.g. Mares and Harstad (2003), Molnar and Virag (2004)). Direct studies of the effects of interim performance appraisals in economics are few, although research in the subject seems to have gained momentum in the very recent years (e.g. Lizzeri, Meyer, and Persico (2003) and Fang and Moscarini (2005)).

Among the set of related papers mentioned above, the ones closest in motivation to the current paper are Prat (2000), Lizzeri, Meyer, and Persico(2003), and Fang and Moscarini (2005). Prat studies a team formation problem, in which the question is whether to form homogeneous or diverse work teams, i.e. whether to put agents with different or similar information structures into the same team. He shows that if the actions of the agents are complements in the team payoff, it is better to form teams of agents with similar information structures, and if they are substitutes, it is better to employ a diverse workforce. The reason is that agents with the same information structure commit positively correlated errors, which is good in the presence of complementarity. The current paper is based on a similar intuition in the sense that complementarity is key in determining the optimal information policy, but the main difference is that in my paper the correlation between the agents' outcomes arises endogeneously due to social learning and the information and compensation policies in effect. Lizzeri, Meyer and Persico, as I do, consider the effects of conducting interim performance evaluations on incentives in a two-period model. However, in their work the information is about one's own previous performance and not that of others, and the effect of the performance evaluation operates through a different channel, namely through the adjustment of first and second period effort to marginal incentives that are related

to the reward scheme. Their paper does not include agent learning, whereas in my model the effect of interim information on effort is tied to the agents' learning about their unknown ability and the marginal productivity of effort. Fang and Moscarini, on the other hand, look into how differential wage policies reveal information about an agent's own ability, thereby affecting morale and motivation, when agents are overconfident. The current paper shares a similar intuition, in that learning affects motivation and subsequent effort, but it differs in two important ways: In Fang and Moscarini, the information revelation is mediated through the principal's contract, rather than through direct feedback, and agents are assumed to be overconfident, a necessary condition for information withholding to occur in their model. In contrast, the current paper analyzes the effects of direct performance feedback rather than focusing on the signaling aspect of wage contracts, and does not assume that agents are overconfident, although this case is considered as an extension.

Although there are theoretical and experimental studies that deal with wage comparisons in the workplace where agents with relative payoff-dependent utility functions compare their levels of pay with each other or the principal (e.g. Englmaier and Wambach (2004), Gächter and Fehr (2001), Grund and Sliwka (2005), Itoh (2004)), the effects of social comparison information regarding task performance on ability perception, motivation and subsequent behavior remain mostly unexplored. To my knowledge, this paper is the first to analyze the effects of social comparison information on subsequent performance and its implications for optimal feedback policies, without assuming any external payoff dependence among the agents, either in terms of the utility function or through the reward scheme⁵.

⁵In my model, social comparisons can affect behavior even in the absence of any external payoff dependence. However, I also have a section in which I analyze the optimal information

The organization of the paper is as follows: in Section 2, I present the general model. In Section 3, I analyze the optimal information policy with different assumptions on the principal’s payoff function and the endogeneity of contracts. Section 4 considers some extensions to the theoretical model, Section 5 discusses the design and reports the results of the experiment, and Section 6 concludes.

1.2 Model

1.2.1 Setup

I consider a simple setting where there is one principal (she), and two risk-neutral agents (he) who engage in a task for two periods. The ability of agent i is denoted by $a_i > 0$, and η denotes a common shock which affects all agents’ performances in the same way. Throughout the paper, I will interpret this common shock as the difficulty of the task. a_i ’s are independent draws from the same distribution, and are independent of task difficulty. Abilities and task difficulty are unknown to all parties at all times, but distributions are common knowledge. The first period is a learning stage, in which performance depends on ability and the difficulty of the task. Therefore, performance in this period potentially provides some information about ability and difficulty. I henceforth refer to agent i ’s performance in the first period as his “signal”, denoted by s_i . These signals are assumed to be payoff-irrelevant for both the principal and the agents. After signals are realized, agents update their beliefs about ability using the signals they observe (if any), and based on these beliefs decide how much effort to exert in the second period. The effort exerted then leads to a distribution over final outcomes and therefore

policy when there *is* external payoff dependence, through positively or negatively interdependent contracts.

payoffs.

1.2.1.1 Timing

Period 0: a_i and η drawn

Period 1: s_i realized and possibly observed; efforts chosen

Period 2: Final performance and payoffs realized

1.2.1.2 The Signal Structure in the First Period

Performance in the first period depends on ability and a common component that affects all agents in the same way. Throughout the paper, I use two different signal structures, continuous and discrete, to maintain analytical tractability in different contractual settings, and for implementability in an experimental framework. The use of either signal structure does not change the main qualitative results of the paper. I start below with a continuous, multivariate normal model to illustrate the optimal information policy with independent contracts, and then move on to a discrete model for the analysis of dependent contracts. The discrete model is also used in the experimental design.

The Multivariate Normal Model I assume that the first period performance signals are realized according to $s_i = a_i + \eta$, where a_i is the ability level of agent i and η is a common shock to performance.⁶ I also assume that a_i and η are normally distributed, a_i , s_1 and s_2 are jointly normally distributed, and abilities

⁶Notice that η should be interpreted as the “easiness” rather than “difficulty” of the task, since it has a positive effect on performance.

are independent of each other and the common shock. All distributions are common knowledge. Formally:

$$a_i \sim N(\bar{a}_i, \sigma^2), \quad \bar{a}_i > 0, \quad i = 1, 2.$$

$$\eta \sim N(0, \psi^2)$$

$$\text{cov}(a_i, \eta) = 0, \quad i = 1, 2.$$

$$\text{cov}(a_1, a_2) = 0$$

and therefore

$$\begin{pmatrix} a_i \\ s_i \\ s_{-i} \end{pmatrix} \sim N \left(\begin{pmatrix} \bar{a}_i \\ \bar{a}_i \\ \bar{a}_{-i} \end{pmatrix}, \begin{bmatrix} \sigma^2 & \sigma^2 & 0 \\ \sigma^2 & \sigma^2 + \psi^2 & \psi^2 \\ 0 & \psi^2 & \sigma^2 + \psi^2 \end{bmatrix} \right)$$

where $\bar{a}_i = \bar{a}_{-i}$, since agents are assumed to be identical.

Notice that because of the common uncertainty factor affecting performance, the signals of the two agents are positively correlated, but are independent conditional on difficulty, due to the independence of ability and difficulty.

1.2.1.3 Performance Technology and Compensation

After the learning stage, each agent chooses an effort level, which determines a probability distribution over final outcomes and payoffs. The performance (production) technology is given by $q_i = a_i e_i$, where $e_i > 0$ denotes agent i 's effort level and a_i his ability⁷. Notice that ability and effort are complements in

⁷In the discrete outcome model which I will consider later in the paper, this technology will correspond to $\pi_i = a_i e_i$, where π_i is the probability of getting a good outcome. The two frameworks are equivalent.

performance ($\frac{\partial^2 q_i}{\partial a_i \partial e_i} > 0$). The cost of effort is given by the function $c(e_i)$, with $c'(e_i) > 0$, $c''(e_i) > 0$ and $c'(0) = c(0) = 0$.⁸

1.2.2 Principal's Payoff

The principal's payoff depends on the final (second period) performances of the two agents, and is given by:

$$\Pi(q_1, q_2) = E[kq_1q_2 + V(q_1 + q_2)] = E[ka_1e_1a_2e_2 + V(a_1e_1 + a_2e_2)] \quad (1.2.1)$$

This particular form for the principal's payoff is chosen because it is the direct continuous equivalent of the payoff function in the discrete framework that will be used in the dependent contract case and the experimental analysis. Notice that the multiplicative term in the payoff function captures a complementarity between the agents' performances in the principal's payoff. Complementarity in this context means that the principal obtains a higher payoff when both agents have an average performance rather than one doing very well and the other doing very badly ($k > 0$). The case of $k = 0$ would correspond to perfect substitution, in which case the principal's payoff is just the expected sum of the agents' performances times a constant:

$$\Pi(q_1, q_2) = E[V(q_1 + q_2)]$$

Complementarities will turn out to play an important role in the optimal policy of the principal.

⁸Notice that I abstract from task difficulty in the second period. This is consistent with an interpretation in which the common shock is independent across periods, such as the difficulty of a particular exam, or the lenience of a particular grader.

1.2.3 Information Policy

The main focus of the paper is the manipulation of the agents' information set by the principal. Specifically, I analyze whether or not the principal would be better off making others' and/or own first-period signals observable to the agents, before they make the effort decision. I assume that the principal commits to an informational policy ex-ante (before signals are realized). Also, interim performances are assumed to be non-contractible.⁹ Let S^i denote the set of signals observable to agent i before he makes the effort decision. Then, the principal's options for agent i 's interim information set are $S^i = \{s_1, s_2\}$ (information about others revealed), $S^i = \{s_i\}$ (information about others withheld), or $S^i = \emptyset$ (all information withheld).

The optimal information policy will naturally depend on the objective function of the principal. In what follows, I analyze the effects of the principal's payoff function and especially the level of complementarity of the two agents' performances on the optimality of information revelation. In order to do this, I first consider the information revelation problem of a principal for an exogenously given compensation scheme, and leave aside the issue of choosing a compensation scheme. I then proceed to analyze the case where the principal can choose a general menu of contracts in addition to the information policy. The cases analyzed are given below:

- Exogenous, Independent Contracts (perfect substitutability vs. complementarities)
- Endogenous, Potentially Dependent Contracts (perfect substitutability vs. complementarities)

⁹I discuss this point in more detail in Section 4, which talks about extensions to the model.

1.3 Results

1.3.1 Optimal Information Policy with Fixed, Independent Contracts

In this section, I analyze the information revelation policy of a principal whose payoff depends only on a function of the outcomes of the agents, but not on the wages paid to them. I assume that the principal has no control over the compensation scheme, which is exogenously given, and does not bear the cost of paying agents, although she has incentives to maximize a function of their performance. As mentioned before, this case reflects situations where the principal is a division manager within a firm, a teacher, or a coach, who can employ only non-monetary policies to influence the motivation and performance of her agents, in a case where monetary incentives are determined by another authority. I also assume that the exogenously given compensation scheme is “independent”, i.e. one agent’s outcome has no effect on the other agent’s payoff.

1.3.1.1 Beliefs and Effort Choice

To be able to analyze the optimal information policy of the principal, I start by studying the agent side of the model. In the presence of imperfect information about ability, Bayesian agents update their beliefs about ability using the first period signal(s). Beliefs, in turn, are crucial in determining subsequent effort. The common uncertainty that affects the realization of the signals gives rise to the following lemma:

Lemma 1.1. (*Relativity of Self-Concept*) *An agent’s posterior beliefs about ability increase with his own signal. In the case where the other agent’s signal s_{-i} is observed, the ability posterior also depends on s_{-i} , and decreases with it.*

Proof—Using the properties of the multivariate normal distribution (see Appendix for a derivation), it is possible to show that agent i 's expectation of ability conditional on observing signals s_i and s_{-i} will be

$$E(a_i | s_i, s_{-i}) = \bar{a}_i + \frac{(\sigma^2 + \psi^2)(s_i - \bar{a}_i) - \psi^2(s_{-i} - \bar{a}_{-i})}{\sigma^2 + 2\psi^2} \quad (1.3.1)$$

If each agent observes his own signal only, then the conditional expectation of ability is given by:

$$E(a_i | s_i) = \bar{a}_i + \frac{\sigma^2(s_i - \bar{a}_i)}{\sigma^2 + \psi^2} \quad (1.3.2)$$

The claims follow immediately from the above equations. \square

The agents in this model basically face a signal extraction problem that requires them to filter the effects of difficulty and ability on the signal(s) observed. The intuition for the result that observing a good signal by the other agent lowers beliefs about own ability comes from the fact that the observation of a good signal by the other agent increases the likelihood that the task was easy, making it statistically less likely that ability was high. Therefore, keeping own performance constant, the decrease in difficulty perception induced by a better performance by the other agent leads to a downward revision in beliefs about own ability. This result is in accordance with evidence from psychology, which suggests that self-confidence decreases with unfavorable social comparisons (e.g. Alicke (2000), Brickman and Bulman (1977)).

After observing the first period signal(s) and updating their beliefs, agents decide on how much effort to exert in the second period. Optimal effort will naturally be influenced by the incentive scheme faced by the agent. For the purposes of this section, I assume a linear, independent compensation scheme, where an agent is

paid $w_i = \alpha + \beta q_i$, with $\alpha \geq 0$ and $\beta > 0$.¹⁰¹¹ Effort, along with true ability, determines the ultimate performance of the agent, according to the production technology $q_i = a_i e_i$.

Lemma 1.2. *Keeping one's own signal constant, the optimal effort level decreases with the other agent's signal.*

Proof—With independent, linear contracts, agent i's problem is to choose effort $e_i > 0$ to maximize:

$$\max_{e_i} E[\alpha + \beta q_i - c(e_i) | S^i]$$

From the first order condition (which is also sufficient), the optimal effort of agent i, e_i^* solves:

$$E(a_i | S^i) \beta - c'(e_i) = 0 \tag{1.3.3}$$

Now consider an increase in the conditional expectation $E(a_i | S^i)$. Notice that because of the convexity of the cost function, the second term in equation 4 increases with an increase in effort ($c''(e_i) > 0$), whereas the first term is independent of the effort level. This implies that the optimal effort has to increase when posterior beliefs about ability increase. That is, higher self-confidence leads to an increase in effort. The result that observing a higher signal by the other agent decreases effort follows directly from Lemma 1. □

As can be seen from equation (4), the optimal effort of an agent is increasing in his perceived ability, which I interpret as “self-confidence”, or “self-concept”. In the presence of imperfect knowledge about ability, self-confidence affects the agent's perceived prospects from exerting more effort, and hence his effort decision. The result stems from the fact that ability and effort are complements in

¹⁰Within the class of independent contracts, restricting attention to linear contracts is without loss of generality by standard arguments.

¹¹By the symmetry of the agents, wages will depend on performances only and not identities. Therefore, α and β will not be agent-specific.

the production function, and implies that an agent who is exposed to unfavorable comparisons will lose motivation and decrease his effort.¹²

Given the effect of comparisons on behavior, I now analyze the principal's optimal information policy in this independent, exogenous contracts case. From now on, I assume that the effort costs are quadratic, given by $c(e_i) = e_i^2/2$. With this specification, it is easy to verify that the optimal effort level will be equal to the conditional expectation of ability times the piece-rate, i.e. $e_i^* = E(a_i | S^i) \beta$. Since the principal does not bear the costs of inducing effort, her payoff function depends only indirectly on the contract, through the agents' optimal effort levels.

Recall that the principal's expected payoff, which is a function of the performances of the two agents, is given by:

$$\begin{aligned} \Pi &= E_{a_1, a_2, s_1, s_2} [k(q_1 q_2) + V(q_1 + q_2)] \\ &= E_{a_1, a_2, s_1, s_2} [k(a_1 e_1^* a_2 e_2^*) + (a_1 e_1^* + a_2 e_2^*)V] \end{aligned} \tag{1.3.4}$$

Notice that the expectation is taken over the joint distribution of abilities and signals. This is because the principal commits to an information revelation scheme ex-ante, before signals are realized. In what follows, I drop the subscripts in expectations to avoid notational clutter, unless necessary. Upon substitution of the equilibrium effort levels e_i^* , the principal's expected payoff becomes

$$\Pi = k\beta^2 E[a_1 E(a_1 | S^1) a_2 E(a_2 | S^2)] + V\beta E[a_1 E(a_1 | S^1) + a_2 E(a_2 | S^2)]$$

¹²However, it should be noted that the informational model of comparisons outlined above is also able to generate the opposite prediction (higher effort in response to an unfavorable comparison), with a different specification of the production technology or the reward scheme. If there is a "pass-fail" scheme in place, for instance, where a performance threshold must be met for a fixed reward to be obtained, a very confident agent may well choose to work less than a less confident one, and an unfavorable comparison that pulls beliefs downward may actually improve effort if initial confidence is too high to start with.

Define Π^{sc} , Π^{own} , and Π^{no} to be the expected payoff to the principal when she commits to revealing social comparison information, own performance information only, and no information at all, respectively¹³.

Proposition 1.3. *When $k=0$ (perfect substitutability), revealing all useful information to the agents is optimal for the principal, i.e. $\Pi^{sc} > \Pi^{own} > \Pi^{no}$.*

Proof—See Appendix. □

To gain some intuition for this result, it is useful to note that the principal’s expected payoff can be decomposed as

$$\begin{aligned}\Pi &= (E[a_1 E(a_1 | S^1)] + E[a_2 E(a_2 | S^2)])V\beta \\ &= (E(a_1) + cov(a_1, E(a_1 | S^1)) + E(a_2) + cov(a_2, E(a_2 | S^2)))V\beta\end{aligned}\tag{1.3.5}$$

The result that more information is better follows from two observations. First, as can be seen from Equation 6, the principal’s expected payoff is increasing in the covariance between an agent’s true ability and his ability posterior. Observing the other agent’s signal provides better information about own ability for each agent, and the principal (ex-ante) prefers this because her payoff function is such that the marginal return to higher effort by an agent is increasing in the agent’s ability, due to the complementarity of effort and ability in the production technology. In this sense, on an individual level, the principal would like to match higher beliefs with higher actions, and choose the information structure that makes self-confidence (effort) more tightly linked to true ability. Second, the separability of the principal’s payoff in the two agents’ performances and the agents being ex-ante identical ensures that the “more information is better” result will hold with two agents also. With substitutability, social comparison

¹³Throughout the paper, the term “revealing social comparison information” will always mean revealing others’ performances *in addition to* own performance.

information only serves to provide better information from the perspective of the principal. Therefore, the principal prefers giving agents information about themselves *and* others to giving them information only about themselves; and giving information only about themselves to not giving any information at all. This “ability-effort alignment effect” is the first main effect that plays a role in determining the principal’s optimal information policy.

1.3.1.2 The Case of Complementarities

The assumption that agents’ performances are perfect substitutes for the principal is likely to be unrealistic in many settings. A coach or a teacher, for instance, may prefer all her players or students having a reasonable performance to some performing extremely well and some extremely badly. Likewise, in a team production setting, success may require that everybody put some effort rather than some working very hard and others slacking. Such cases can be modeled through complementarities among the agents’ performances in the principal’s payoff function. As noted before, complementarity in this setting is captured by $k > 0$. The following proposition shows that withholding social comparison information can be better for the principal, if complementarities are strong enough.

Proposition 1.4. *a) For all k , $\Pi^{own} > \Pi^{no}$. b) There exist parameter configurations for which there is a threshold level of complementarity, \bar{k}_{scown} , such that for $k > \bar{k}_{scown}$, $\Pi^{own} > \Pi^{sc}$. c) There exist parameter configurations for which there is a threshold level of complementarity \bar{k}_{scno} such that for $k > \bar{k}_{scno}$, $\Pi^{no} > \Pi^{sc}$, and $\bar{k}_{scown} < \bar{k}_{scno}$ holds in the relevant region.*

Therefore, with appropriate parameter restrictions, the ranking of the policies is as follows:

$$\begin{array}{c}
0 \quad \underbrace{\hspace{1.5cm}} \quad [\bar{k}_{scown}] \quad \underbrace{\hspace{1.5cm}} \quad [\bar{k}_{scno}] \quad \underbrace{\hspace{1.5cm}} \quad k \\
\text{no} < \text{own} < \text{sc} & \text{no} < \text{sc} < \text{own} & \text{sc} < \text{no} < \text{own}
\end{array}$$

The intuition for this result is the following: In the presence of complementarities, the principal's expected payoff depends on both the marginal distribution of each agent's performance, and the association between the performances of the two agents. When agents can observe their own performances only, from an ex-ante perspective their self-confidence, effort, and performance levels will co-vary positively, due to the common uncertainty that affects the performances. Specifically, an agent is more likely to observe a high signal when his peer has observed a high signal as well, leading to a positive correlation in efforts. When agents can observe each other's signals and update their beliefs accordingly, however, the correlation between their posterior beliefs becomes negative, since one's perceived ability is negatively related to others' signals, and loosely speaking, one person's success is bad news for the other in terms of beliefs. This translates into a negative covariance between the efforts of the two agents, which the principal does not like if the agents' outcomes are complements in her payoff function. Therefore, with complementarities, there is a new effect that pushes the principal in the direction of withholding social comparison information. However, the previously mentioned ability-effort alignment effect is still present at the individual level (the expected performance of a single agent is higher with more valuable information). For certain parameter restrictions, the negative effect of social comparisons start to dominate at high enough levels of complementarities. The reason why the negative correlation does not *always* imply withholding of social comparison information for high enough complementarities is because the complementarity affects the importance of both the marginal distributions and the association between outcomes in the principal's expected payoff. Social com-

parison information increases the expected performance of a single agent ex-ante, although it decreases the co-movement of performances, and these two factors have conflicting effects on the principal's payoff. When the effect of a particular information system on the marginal distributions of and the association between outcomes go in the same direction, as is the case with own performance information, the effect on the principal's payoff is unambiguous, giving the result that providing agents with own performance information is always better than giving them no information. With social comparisons, however, the net effect depends on the trade-off between the within-person and across-person effects on the principal's payoff.

The comparison between giving social comparison information and no information at all also builds on the same type of intuition. That is, social comparisons can be dominated by giving no information at all, at high levels of complementarity. The threshold level of complementarity needed for that, however, is higher than that needed for own information to dominate, since withholding all information generates lower expected payoff in terms of the marginal distributions than giving own information only, in addition to a zero covariance between performances.

Example:

The below graph plots $\Pi^{sc} - \Pi^{own}$ and $\Pi^{sc} - \Pi^{no}$ against k , the level of complementarity, for $\beta = 0.5$, $\sigma^2 = 0.275$, $\psi^2 = 1$, $\bar{a} = 2$ and $V = 1$. As can be seen in the figure, social comparisons dominate "own information only" and "no information" policies for k small enough, and the ranking reverses after a threshold level of complementarity in both cases.

Figure 1.1: Expected Payoff Differences

Analyzing the threshold level of complementarity, k_{scown} , yields the following insights about the parameter restrictions required for the above results to hold:

Proposition 1.5. *There exists a positive threshold level of complementarity, k_{scown} , after which withholding social comparison information is optimal, if the mean of the ability distribution, \bar{a} is high enough. This threshold \bar{a} is increasing in the variance of the common shock, ψ^2 and the variance of the ability distribution, σ^2 .*

Proof—See Appendix. □

The above proposition reflects the trade-off between the standard value of information and its effects on motivation. It shows that withholding social comparisons may be optimal when the initial self-confidence of the agents is high enough, in which case they stand to lose more from an unfavorable social comparison. Also, increasing the variance of the common shock makes comparisons more informative, and this works to increase the value of information revelation.

The results of this section show that with exogenous, independent contracts and when the principal does not bear the cost of incentives, it may be optimal to

withhold social comparison information for strong enough complementarities. A question that comes to mind at this point is whether the information withholding result would hold if the principal took into account the cost of providing incentives. The following proposition shows that if the principal bears the cost of incentives without the power to manipulate them, she would still want to withhold social comparison information if complementarities are strong enough. The result is driven by the fact that agents' expected utilities will always be higher with more information in this model, giving the principal reason to withhold comparison information.

Proposition 1.6. *With exogenous contracts, if it is optimal to withhold social comparison information when the principal does not bear the cost of incentives, it will also be optimal to withhold information when payments to the agents are accounted for in the principal's payoff function.*

Proof—See Appendix. □

The above analyses show that social comparisons can affect effort and subsequent performance even in the absence of any external payoff dependence, through a purely informational channel. In many contexts, however, the reward scheme in place is not completely independent across agents: in tournaments or relative reward schemes, or team-based or cooperative incentive schemes, an agent's payment depends not only on his own performance or outcome, but also those of others. I now analyze how such dependence in the reward scheme might affect the results. To be able to analyze equilibrium beliefs and effort in this case, which now involve expectations of non-linear functions of random variables, I use a discrete model in the rest of the paper. This model closely replicates the multivariate normal model presented above, and all the results stated up to now continue to hold. The next section presents this discrete model, briefly reiterates

the previous results as they apply to this signal structure, and analyzes the case of dependent contracts.

1.3.1.3 The Discrete Model

Assume that the ability of agent i takes on one of two values, high and low, i.e. $a_i \in \{a_H, a_L\}$, with $0 < a_L < a_H \leq 1$. The difficulty of the task can be either high or low as well, and is denoted by $\eta \in \{\eta_L, \eta_H\}$, with $Pr(\eta = \eta_L) = \lambda$, where $0 < \lambda < 1$. a_i 's are independent draws from the same distribution, with $Pr(a_i = a_H) = \rho$ and $0 < \rho < 1$, and abilities are independent of task difficulty.

Fixing task difficulty, a higher ability level leads to a stochastically higher signal distribution for the agent, and vice versa. Specifically, I assume that high ability agents observe a good signal in an easy task for sure, low ability agents observe a bad signal in the difficult task for sure, and the probability of observing a good signal is $\mu > 0$ if the agent is of low ability and the task is easy, or the agent is of high ability and the task is difficult. Notice that this uncertainty reflects an additional idiosyncratic risk, which is assumed to be independent across agents.¹⁴

$$Pr(s_i = 1) = \begin{cases} \mu & \text{if } \eta = \eta_H \text{ and } a_i = a_H \\ 0 & \text{if } \eta = \eta_H \text{ and } a_i = a_L \\ 1 & \text{if } \eta = \eta_L \text{ and } a_i = a_H \\ \mu & \text{if } \eta = \eta_L \text{ and } a_i = a_L \end{cases}$$

After the learning stage, each agent chooses an effort level, which determines a probability distribution over final outcomes and payoffs. I assume that the second period outcome for each agent can be either a “success” (S) or a “failure”

¹⁴This extra idiosyncratic risk is not necessary for the results of the model to hold, but turns out to be more convenient for the experimental design.

(F), and the probability of success is given by $\pi_i^s = a_i e_i$, where $e_i \in (0, 1]$ denotes the effort level of the agent and a_i his ability. The cost of effort is given by the function $c(e_i)$, with $c'(e_i) > 0$, $c''(e_i) < 0$ and $c'(0) = c(0) = 0$.

The principal's payoff depends on the final (second-period) outcomes of the two agents, as summarized by the following table:

		<i>Agent 2</i>	
		<i>Success</i>	<i>Failure</i>
<i>Agent 1</i>	<i>Success</i>	V_{ss}	V_{sf}
	<i>Failure</i>	V_{fs}	V_{ff}

where $V_{sf} = V_{fs}$ by the symmetry of the agents, and I assume $V_{ss} > V_{sf} > V_{ff} \geq 0$. V_{ff} is normalized to be zero without loss of generality. The complementarity of the outcomes for the principal is fundamentally linked to the supermodularity of the principal's payoff, i.e. whether $(V_{ss} + V_{ff}) - (V_{sf} + V_{fs}) > 0$ or not. Intuitively, complementarity in this context means that the principal's payoff when both agents succeed at the same time is more than twice as high as her payoff when only one of the agents succeeds. Notice that if the principal does not obtain extra utility from both agents succeeding (or extra disutility from both failing), her payoff function would be linear and $(V_{ss} + V_{ff}) - (V_{sf} + V_{fs}) = 0$ would hold. Let $V_{sf} = V$, and define $k = V_{ss} - 2V$. In this case, $k \geq 0$ captures the degree of complementarity of the two agents in the principal's payoff, and $k = 0$ corresponds to the case of perfect substitutability. A general compensation scheme in this context specifies a wage corresponding to each possible outcome in the set $\{(S, S), (S, F), (F, S), (F, F)\}$, where the first element in $(., .)$ denotes the first agent's outcome. I denote these wages by $w_{SSi}, w_{SF_i}, w_{FS_i}, w_{FF_i}$, where w_{SF_i} is the payment to agent i when he succeeds and the other agent fails. By the

symmetry of the agents, wages will depend on outcomes only, and the subscript for identity is dropped henceforth.

It is straightforward to show that Lemma 1 applies in this discrete case as well: keeping own outcome constant, a higher outcome by the other agent decreases ability perception. Moreover, observing a good outcome by the other agent leads to a reduction in perceived ability as compared to the case where one observes only her own performance:

Lemma 1.7. *a) An agent's posterior beliefs about ability increase with his own signal. In the case where the other agent's signal s_{-i} is observed, the ability posterior also depends on s_{-i} , and decreases with it. Specifically,*

$$E(a_i | s_i = 0, s_j = 1) < E(a_i | s_i = 0, s_j = 0) < E(a_i | s_i = 1, s_j = 1) < E(a_i | s_i = 1, s_j = 0)$$

b) Observing that the other agent had a good (bad) signal lowers (increases) one's beliefs from their level when only one's own signal is observed. The complete ranking of beliefs is as follows:

$$\begin{aligned} E(a_i | s_i = 0, s_j = 1) &< E(a_i | s_i = 0) < E(a_i | s_i = 0, s_j = 0) \\ &< E(a_i | s_i = 1, s_j = 1) < E(a_i | s_i = 1) < E(a_i | s_i = 1, s_j = 0) \end{aligned}$$

Proof—See Appendix. □

With this discrete-outcome framework, an “independent” wage scheme corresponds to one where one agent's outcome has no effect on the other agent's payoff. Notice that this implies $w_{SS} = w_{SF} = w_S$ and $w_{FS} = w_{FF} = w_F$. As before, it is possible to show that lower beliefs will lead to lower effort with such wage schemes. Finally, the result about information withholding being optimal for high enough complementarities goes through in this case as well (see Appendix).

1.3.2 Dependent Contracts

This section maintains the assumption that the reward scheme is exogenously given and that the principal does not internalize payments to agents, and only drops the independence assumption in the contract. Studying this case first proves to be useful for building intuition for the endogenous contract case.

When contracts are dependent, we will no longer have $w_{SS} = w_{SF} = w_S$ and $w_{FS} = w_{FF} = w_F$, since an agent's payment will now depend on the other agent's outcome also. Recall that with independent contracts that are monotonically increasing in output (or a good outcome), a change in beliefs directly translates into a change in effort in the same direction: effort goes down when self-confidence goes down, and up when self-confidence goes up. Therefore, unfavorable social comparisons directly imply a decline in effort. With dependent contracts, the effect of comparisons on beliefs will still be the same: self-confidence will decline after an unfavorable comparison and increase after a favorable one. However, it is no longer clear whether a decline in beliefs will necessarily lead to a decline in effort. With independent contracts, the effort decision of an agent is influenced by his expectations about his own probability of success only, and hence only by his posteriors about his own ability, and this is why effort changes in the same direction in response to a change in beliefs. With dependent contracts, on the other hand, each agent's optimal effort will depend on his beliefs about the likelihood of all of the four possible states (success-success, success-failure, failure-success, failure-failure), and thereby his beliefs about the other agent's ability in addition to his own, given the signals he observes.

Formally, the equilibrium effort of agent i given his information set S^i solves the

following problem:

$$e_i^*(S^i) \in \underset{0 < e_i < 1}{\operatorname{argmax}} \left\{ E \left[\begin{array}{l} (a_i e_i(S^i) a_j e_j^*(S^j) w_{SS} + (a_i e_i(S^i) (1 - a_j e_j^*(S^j)) w_{SF} \\ + a_j e_j^*(S^j) (1 - a_i e_i(S^i)) w_{FS} \\ + (1 - a_i e_i^*(S^i)) (1 - a_j e_j^*(S^j)) w_{FF} - c(e_i) \end{array} \middle| S^i \right] \right\}$$

Notice that an agent's effort level also depends on his perception of what the other agent's effort and ability are likely to be. When social comparison information is revealed, rational expectations imply that given the signals observed, the effort level of each agent will be known by the other in equilibrium. When own information is given only, the equilibrium effort of an agent will depend on his expectations about the signal that the other agent might have received, conditional on his own signal, and therefore on a probability distribution on the other agent's equilibrium effort.

Lemma 1.8. *The equilibrium effort level of agent i will be given by:*

$$e_i^*(s_1, s_2) = \frac{(w_{SF} - w_{FF})(E(a_i | s_1, s_2) + E(a_{-i} | s_1, s_2) E(a_1 a_2 | s_1, s_2) (w_{SS} + w_{FF} - (w_{SF} + w_{FS})))}{1 - E(a_1 a_2 | s_1, s_2)^2 (w_{SS} + w_{FF} - (w_{SF} + w_{FS}))^2}$$

when social comparison information is observed, and

$$e_i^* = \frac{E[a_i](w_{SF} - w_{FF})}{1 - E[a_1]E[a_2](w_{SS} + w_{FF} - (w_{SF} + w_{FS}))^2}$$

when no information is observed.

With own information only, the equilibrium effort levels given observation of $s_i \in \{0, 1\}$, denoted by $e^*(0)$ and $e^*(1)$ solve:

$$\begin{aligned}
& (Pr(s_{-i} = 1 | s_i) e^*(1) E[a_1 a_2 | s_i, s_{-i} = 1] \\
& + Pr(s_{-i} = 0 | s_i) e^*(0) E[a_1 a_2 | s_i, s_{-i} = 0]) ((w_{SS} + w_{FF} - (w_{SF} + w_{FS})) \\
& + E(a_i | s_i) (w_{SF} - w_{FF}) - e_i(s_i) = 0.
\end{aligned}$$

for $i=1,2$.

Proof—From the direct solution of the two agents’ first order conditions for each informational specification, and application of the properties of expectations. See the Appendix for explicit expressions for the equilibrium efforts in the case of own information only. \square

With dependent reward schemes, an agent’s payment can depend on the other agent’s outcome positively or negatively. Assuming that $w_{FS} = w_{FF} = 0$ ¹⁵, the nature of the dependence boils down to the comparison of w_{SS} and w_{SF} . If $w_{SS} > w_{SF}$, an agent is rewarded more for his success when the other agent has succeeded than when he has failed, which I will call a “cooperative” incentive scheme. If $w_{SS} < w_{SF}$, on the other hand, success is rewarded more if it comes at a time when the other agent has failed, which means that the incentive scheme is “competitive”.

Recall that in the case of independent contracts, it was shown that the effort level of each agent is an increasing function of his own self-confidence only, and therefore, the ranking of beliefs (the fact that an agent thinks less highly of himself when others have received a good signal) translated directly into effort choices, i.e. $e^*(0,1) < e^*(0) < e^*(0,0) < e^*(1,1) < e^*(1) < e^*(1,0)$. With dependent contracts, however, this may no longer be the case, since an agent’s

¹⁵I will show later, in the endogenous contracts section, that this is indeed optimal when contracts can be chosen.

effort level is not only affected by his beliefs about his own probability of success but the probability of both agents succeeding. When the reward scheme is highly cooperative, learning that the other agent has observed a good signal, although it still decreases the agent's own ability perception, can lead to higher effort because it increases the marginal return to effort, which now depends on *both agents'* perceived productivities. The following example illustrates a potential reversal in the belief ranking given in Lemma 3, with dependent contracts.

Example:

Suppose that $w_{FS} = 0, w_{FF} = 0$, and let $w_{SF} = 0.5$. Also suppose that $\rho = 0.25, a_H = 0.8, a_L = 0.6, \mu = 0.6$, and $\lambda = 0.8$. The graphs below plot $e^*(0, 1) - e^*(0, 0)$ and $e^*(1, 1) - e^*(1, 0)$ against the level of "cooperativeness" of the wage scheme, as measured by $w_{SS} - w_{SF}$. When $w_{SS} - w_{SF}$ is high enough, observing that the other agent succeeded when you have failed can increase effort, as compared to the case of observing that he has failed also. Likewise, the other agent failing can decrease the motivation to exert effort.

Figure 1.2: Effort Difference

Figure 1.3: Effort Difference

Since contracts are exogenous, the principal's expected payoff, as before, does not contain wage payments, and is given by

$$\Pi = E(a_1 e_1 a_2 e_2)k + \sum_{i=1}^2 E(a_i e_i) V$$

or

$$\Pi = \sum_{i=1}^2 (E(a_i e_i))V + (E(a_1 e_1)E(a_2 e_2) + cov(a_1 e_1, a_2 e_2)) k$$

The comparison of informational policies again boils down to an assessment of the marginal distributions and covariances, with and without information:

Lemma 1.9. *The expected performance of a single agent, $E(a_1 e_1)$ is higher with social comparison information, for all w_{SS} and w_{SF} . Specifically, $E(a_1 e_1)_{sc} > E(a_1 e_1)_{own} > E(a_1 e_1)_{no}$.*

Proof—See Appendix. □

Lemma 1.10. *a) With own information, the covariance between effort levels (and outcomes) is positive, for all w_{SS} and w_{SF} . b) The covariance of effort levels when social comparison information is given increases in the “cooperativeness” of the reward scheme, $w_{SS} - w_{SF}$. There is a threshold level of $w_{SS} - w_{SF}$, call \bar{w} , such that if $w_{SS} - w_{SF} > (<)\bar{w}$, the covariance of effort levels (and outcomes) is positive (negative).*

Proof—See Appendix. □

Figures 1.4 and 1.5 plot the covariance between efforts against the cooperativeness of the reward scheme.

Figure 1.4: Covariance between efforts with SC information

Figure 1.5: Covariance between efforts with Own information

As can be seen, dependent (cooperative) contracts can create an external effect on the covariance of outcomes that counteracts the negative effect that social comparison information generates.

1.3.3 Optimal Information Policy with Endogenous Contracts

The above analyses assumed that the principal's payoff depends only on a function of the agents' outcomes, to reflect a setting where the compensation scheme is determined by a separate authority and the principal can choose only the information policy. In many contexts, however, the principal can manipulate not only the amount of information available to the agents, but also the incentive scheme. Therefore, I attempt next to answer the question of whether the optimal informational policy would be different when the principal can also choose the contract, along with the information policy. When contracts are variable, the principal chooses, along with the information policy, payments to be made to each agent in the four possible states, denoted by w_{SS} , w_{SF} , w_{FS} and w_{FF} . Recall

that here, w_{SF} denotes the payment to be made to the agent that succeeds, when the other agent fails. Due to the symmetry of the agents, wages will depend only on the outcome vector, and not on identities. We assume that the principal maximizes her ex-ante expected payoff subject to incentive compatibility, individual rationality and limited liability constraints. Formally, the principal's objective function is given by:

$$\begin{aligned} \Pi = & E[(a_1e_1a_2e_2)(V_{SS} - 2w_{SS}) + (a_1e_1(1 - a_2e_2) \\ & + a_2e_2(1 - a_1e_1))(V_{Sf} - (w_{SF} + w_{FS})) + \\ & (1 - a_1e_1)(1 - a_2e_2)(V_{ff} - 2w_{FF})] \end{aligned}$$

Therefore, the principal's problem can be written as:

$$\begin{aligned} & \max_{0 < w_{SS}, w_{SF}, w_{FS}, w_{FF} < 1} \{E(a_1e_1a_2e_2)(k - 2(w_{SS} + w_{FF} - (w_{SF} + w_{FS}))) \\ & + (E[a_1e_1] + E[a_2e_2])(V - w_{SF} - w_{FS} - 2w_{FF}) + V_{ff} - 2w_{FF}\} \end{aligned}$$

s.t.

$$e_i^* \in \underset{0 < e_i \leq 1}{argmax} \{E \left[\begin{array}{l} (a_i e_i a_j e_j^*)(w_{SS} + w_{FF} - (w_{SF} + w_{FS})) \\ + (a_i e_i)(w_{SF} - w_{FF}) \\ + a_2 e_2 (w_{FS} - w_{FF}) + w_{FF} - e_i^2/2 \end{array} \right] S^i \} \quad (IC_i)$$

$$\begin{aligned} & E[(a_i e_i^* a_j e_j^*)w_{SS} + (a_i e_i^*(1 - a_j e_j^*))w_{SF} + a_j e_j^*(1 - a_i e_i^*)w_{FS} \\ & + (1 - a_i e_i^*)(1 - a_j e_j^*)w_{FF} - c(e_i^*)] \geq 0 \quad (IR) \end{aligned}$$

$$w_{SS} \geq 0, w_{SF} \geq 0, w_{FS} \geq 0, w_{FF} \geq 0 \quad (LL)$$

Notice that $(w_{SS} + w_{FF} - (w_{SF} + w_{FS}))$ measures the supermodularity of the reward scheme, i.e. the reward scheme is supermodular if this expression is nonnegative. This property will play an important role in the analysis.

Lemma 1.11. *For any informational specification, the principal's optimal compensation policy will involve setting $w_{FS} = w_{FF} = 0$.*

Proof—See Appendix. □

Given $w_{FS} = w_{FF} = 0$, the principal's expected payoff reduces to

$$\Pi = E[a_1 e_1 a_2 e_2 (k - 2(w_{SS} - w_{SF}))] + (E[a_1 e_1] + E[a_2 e_2])(V - w_{SF})$$

and the supermodularity of the incentive scheme boils down to the comparison of w_{SF} and w_{SS} .¹⁶The following proposition gives more insight into the workings of the optimal contract in this case.

Proposition 1.12. *When the informational policy in place is to give no interim information at all, the principal's maximum payoff can be attained by independent, cooperative, or competitive contracts.*

Proof—See Appendix. □

Proposition 1.13. *When the outcomes of the agents are perfectly substitutable for the principal ($k=0$), then the maximum payoff for the principal can be attained by independent, cooperative, or competitive contracts, regardless of the information policy.*

Proof—See Appendix. □

The above proposition says that the shape of the optimal contract does not have an interaction with the information policy when there are no complementarities. Therefore, it is without loss of generality in this case to restrict attention

¹⁶Cooperative contracts correspond to supermodular incentive schemes, competitive contracts to submodular incentive schemes, and independent contracts to linear incentive schemes.

to independent compensation schemes, in which case our previous results on the optimality of full information revelation holds, since in the absence of complementarities between the agents, only the within-person effort-ability alignment effect is at work, making full information revelation optimal with general compensation schemes.

When there are complementarities ($k > 0$) and *some* information revelation, however, the shape of the contract may no longer be irrelevant. In this case, it is not possible to analytically solve for the optimal contract of the principal by inserting optimal efforts into the principal's payoff function and maximizing with respect to w_{SS} and w_{SF} . I therefore use numerical methods to illustrate the best information policy in such a case with general contracts and various levels of complementarity.

1.3.4 Numerical Illustrations

In this section, I numerically solve the wage-setting problem of the principal for all three informational policies (social comparison information *and* own information, own information only, and no information at all), and compare the maximum payoff attained under the different informational policies, given that wages are set at their respective optimal levels under each informational policy. I start by writing out the objective function of the principal for each informational policy, using the optimal efforts $e_i^*(S^i)$ given in Proposition 3. Plugging in the relevant $e_i^*(S^i)$, I then numerically solve the constrained optimization problem of the principal, which amounts to selecting w_{SS}, w_{SF} subject to the limited liability constraints and restrictions on the efforts.

1.3.4.1 The Optimal Wage Structure

Before analyzing the information policy, I first analyze the shape of the optimal contract under different informational structures.

Result 1 Cooperative incentive schemes, i.e. schemes that involve $w_{SS} > w_{SF}$ are optimal under all informational structures. While the same maximum payoff can also be achieved by competitive or independent wages when $k = 0$ or when the information policy is to give no information, with positive complementarities and an information policy that gives any kind of interim information (own or social comparison), cooperative wages do strictly better. Figure 1.6 plots the difference between the maximized payoffs for different levels of k , under the best competitive and cooperative contracts when social comparison information is given, with the following parameter settings: $V = 0.3, \mu = 0.7, \lambda = 0.4, \rho = 0.4, a_L = 0.4, a_H = 0.7$. Trials with different configurations of parameters do not change the qualitative nature of the result.

Figure 1.6: Payoff Difference

1.3.4.2 The Optimal Information Policy

Result 2: The principal's payoff is always higher with more information when general wage schemes can be chosen, i.e. $\Pi^{sc} \geq \Pi^{own} \geq \Pi^{no}$, with strict inequality for $k \neq 0$.

Figure 1.7 illustrates the principal's payoff under the three informational choices, as k varies. As can be seen, giving social comparison information along with own information is superior to giving own performance information only, which in turn is better than giving no interim information at all.

Figure 1.7: Expected Payoffs With Different Information Policies

The intuition for the result is that cooperative wage schemes increase the covariance between the efforts of the two agents, thereby mitigating the negative effects of social comparison information and restoring the optimality of full information revelation. Recall that when social comparison information is given, it leads to an

Figure 1.8: Expected Payoffs With Different Information Policies

increase in one agent's self-confidence and a decrease in that of the other, which leads to a decrease in the covariance of the efforts. Cooperative wages serve to mitigate this effect, since they induce an incentive for the agents to increase effort when the other agent does so. Figure 1.8 plots the covariance between the efforts at the best cooperative wage scheme (which is the optimal scheme), and the best competitive scheme (which actually is the same as the best independent scheme, because the constraint that $w_{SF} > w_{SS}$ will bind at the optimal solution). As can be seen from the figure, the covariance of efforts is higher with the cooperative wage scheme, and the difference increases with the level of complementarity.

1.3.4.3 Some Comparative Statics Results

I now look at how a change in the parameters of the model affects the difference between the maximized payoffs with social comparison information and own information. The following graphs illustrate the effects of a change in λ , $a_H - a_L$ and μ on the payoff difference.

Figure 1.9: Comparative Statics

As seen from the first graph in Figure 1.9, the payoff difference becomes higher as $a_H - a_L$ increases, because the variance of ability increases with the spread of the ability levels, making information more necessary and desired by the agent. Likewise, the payoff difference is at its highest when $\lambda=0.5$, which corresponds to the highest variance for the common factor, making social comparison information more valuable.

Looking at the third graph, which plots the maximized payoff difference against μ (the probability of the idiosyncratic risk being high), we see that the difference is at its lowest when μ is equal to 0.5, the point where the variance of the idiosyncratic shock is at its highest. The intuition for this result is that social comparisons are less informative when the variance of the idiosyncratic shock is high, blurring the informational value of the comparison.

1.3.5 Discussion

The above analysis illustrates that the principal's choice of whether to withhold or reveal social comparison information is based fundamentally on three main factors: (1) the complementarity between effort and ability on an individual level, (2) the complementarity of performances across agents, (3) the degree of control the principal has on monetary incentives. Recall that the within-agent complementarity pushes the principal toward revealing as much useful information as possible (including social comparison information), whereas the across-agent complementarity might work against revealing social comparison information under certain parameter restrictions. The complementarity at the individual level is fixed by our specification of the production (or probability of success) function. When the across-agent complementarity increases, assuming that the parameter restrictions are met, revealing social comparison information, which induces a

negative correlation in efforts and outcomes, becomes worse than revealing own performances only.

Suppose, for a moment, that the only options available to the principal were to give own performance information or no information at all. In this case, the within-person effort-ability alignment effect and the presence of complementarities would both push the principal toward the revelation of own information, since the correlation of outcomes when only own information is given is positive, because of the common uncertainty. Likewise, if a finer individualistic information structure were available (say, each agent observing two signals of own performance rather than one), this would always be preferred to observing only one signal, due to the better ability-effort alignment it induces. Although social comparison information, like any other kind of valuable information, is also beneficial for better alignment of effort and ability, letting agents observe each other's performances is fundamentally different than letting them observe another performance signal of their own, in the presence of complementarities, because of the across-agent outcome alignment issue. That is why we have a region where the optimal policy is to withhold social comparison information but reveal own performance information, if complementarities are strong enough.

In the above analysis, I have made several modeling choices and assumptions that deserve further discussion. I assumed, for instance, that the first period signal realizations depend on ability and difficulty, but not effort. This assumption is made mainly to be able to focus on the signal extraction problem caused by common uncertainty, and involves no loss of generality for the main result on belief updating. Given that agents start with identical priors and therefore exert the same effort, this equilibrium effort would be accounted for in the inference process and would not distort the belief updating in response to observing social com-

parison information. However, the first period effort levels would be influenced by a learning effect, whereby agents have an incentive to make their signals more precise by exerting more effort, depending on what kind of information policy is in effect, which in turn would be another factor that influences the optimal information policy. Such learning effects have been studied in Meyer and Vickers (1997) and elsewhere, so I choose to focus on the subsequent motivational effects of information by bringing effort into the picture only in the second period. Another assumption worthy of mentioning is that I abstract from task difficulty in the analysis of the second stage. This captures situations where the common uncertainty is uncorrelated across periods (e.g. an easy exam today does not make it more or less likely that the next exam will be easy). If the common shock is correlated across periods, however, agents' belief updating would now affect their optimal effort not only through ability perception but also through the inference of the common shock. In that case, observing good outcomes by other agents could generate a positive impact on effort through a decrease in difficulty perception, working against the negative impact generated by the decline in ability perception.

1.4 Extensions

1.4.1 Different Priors and Overconfidence

The current model assumes that the distribution of the agents' abilities are the same, and that this is common knowledge. This assumption can be relaxed in two main ways. First, in contexts where social comparisons generate information about abilities, it may be important to consider the possibility that agents' perceptions of themselves can be different than their perception of others—for

instance, agents may be overconfident (e.g. Santos-Pinto and Sobel (2003)). Second, all agents may share the belief that one agent is likely to be more able than the other (the expected abilities may be different), or simply less may be known about one agent than the other (the variance of the ability distribution may be different across agents). To capture such effects, the model can be generalized to account for differences in beliefs. In this section, we study one of the above possibilities: the effects of overconfidence on the information disclosure policy.

Suppose the principal's prior beliefs about the agents' abilities differ from those of the agents. Specifically, suppose the principal's prior is normal with mean \bar{a}^P and variance σ^2 , whereas each agent has beliefs about himself that are normal with mean \bar{a}^A and variance σ^2 , where $\bar{a}^P < \bar{a}^A$. Assume each agent's prior beliefs about the other agent's ability are the same as the principal's. In a sense, each agent has optimistic beliefs about himself, but "correct" beliefs about the others. I assume that the performances of the agents are perfect substitutes for the principal, and that the reward scheme is independent and exogenous. The below proposition shows that if overconfidence on the part of the agents is severe enough, it would be better for the principal to commit to withholding social comparison information, even when performances are perfectly substitutable.¹⁷

Proposition 1.14. *Suppose the principal's expected payoff is given by $\Pi = E[a_1e_1 + a_2e_2]V$ (perfect substitutability). If agents are sufficiently overconfident ($\bar{a}^A - \bar{a}^P$ is high enough), the expected payoff of the principal is higher if she makes the other agent's signal unobservable to each agent.*

Proof—See Appendix. □

¹⁷One might think of different ways to model overconfidence, such as assuming that all agents have the same priors about their own and others' abilities, which are more optimistic than the principal's. It is possible to show that the result of this section extends to that case. Moreover, the results directly extend to the discrete case as well.

By abstracting from the effect of complementarities, the perfect substitutes case enables us to identify the difference that overconfidence makes for the optimal feedback policy. Notice that, in the presence of overconfidence, we still have the "alignment effect" at the individual level, which tends to make more information better from the perspective of the principal. However, now there is also a new effect which makes the principal believe that information is likely to bring bad news for the agents, giving her an incentive to withhold signals (notice that given the principal's information set, she calculates expected confidence as $E^P(E_i^A(a_i | s_i, s_{-i})) < E^A(a_i)$). Notice also that social comparison information, in this context, is not different from any other type of "better information". For cases with complementarities, the intuition gleaned from the above analysis suggests that the threshold levels of substitutability required for the optimality of information revelation would go up in the presence of overconfidence, since overconfidence pushes the principal toward withholding information.

1.4.2 Partial or Selective Information Revelation

In this paper, I restrict attention to three types of information policies (no information at all, own performance information only, and social comparison information), and motivate this by the assumption that the principal will not have access to the interim performance signals herself, or that the interim signals are not contractible. Such a restriction is reasonable in some contexts and not in others. If we interpret information revelation as the decision of whether or not to create a setting where agents can observe each other (such as having two people work together on a similar task etc.), this restriction in focus is warranted. If, however, the principal has access to all the interim performance signals and decides on the disclosure policy after observing them, possibilities such as partial

disclosure (e.g. disclose all information whenever both agents get a good signal but not otherwise) or selective disclosure (e.g. disclose only to the agents that have observed a good signal) may arise, although they may not be practical in real-life feedback policies. An obvious extension to the current model, therefore, would be to pursue the extension of generalizing the information policy.

1.4.3 Informed Principal

The model in this paper assumes symmetric information between the principal and agents at the time the information policy is chosen, i.e. neither the principal nor the agents know abilities or the common shock at that point. An alternative would be for the principal to have private information about the agents' abilities before designing the information disclosure policy. This is especially realistic in educational settings, where the principal (teacher) accumulates superior information about the agents (students) through her expertise in judging performance. In this case, posterior beliefs would not only depend on the signals observed, but also on the disclosure policy itself.¹⁸ When the information policy is chosen without knowledge of abilities, as in the current model, observing a bad performance by another agent along with his own bad performance helps the agent preserve favorable beliefs about himself. If the principal's equilibrium policy, however, is such that only groups of agents with sufficiently low ability are shown each other's performances, the ability inference given others' bad outcomes may not be so favorable, since now the agent also updates his beliefs about the ability composition in the group. Ability grouping policies in education, for instance, may generate a conflicting impact on self-concept. On the one hand, for a struggling

¹⁸Ertac, Molnar and Virag (2005) study a related model with imperfect self-knowledge, where the agents use the principal's wage offer to update their beliefs about ability.

student, it may be morale-boosting to observe other students' bad performances in the same task. However, the very fact that the student has been grouped with low-performers might tell him something negative about his ability, if the grouping by the principal was made with knowledge of abilities. It might be interesting to contrast the equilibrium information disclosure policy in such a case with the current version where the principal does not have an informational advantage.

1.5 An Experimental Test of the Theory

In this section, I report the main results from an experiment designed to test the informational theory of social comparisons. Assessing the validity of the theoretical model in the paper is crucial for evaluating the reliability of its policy implications. Controlled experiments are very useful for this purpose, since they provide a setup where the value of the comparison information can be clearly assessed and optimal decisions calculated. Moreover, laboratory experiments make measurement of individuals' beliefs possible, which, given the importance of "self-perception" in the issue of social comparisons, is of utmost importance, and especially hard to achieve in field settings. The experiment presented here tests the agent side of the theoretical model, which is a prerequisite for the validity of the implications regarding the principal's optimal policy.

1.5.1 Experimental Design

The experiment is designed to test the informational theory of social comparisons by analyzing whether individuals use valuable social comparison information correctly in (1) forming judgments and (2) making decisions. Subjects are faced with a decision problem in which the optimal solution should depend on the perceived

value of an individual-specific random variable, which I call “individual factor” for neutrality in wording. This variable corresponds to ability in the theoretical model. The individual factor is randomly assigned to be either high or low for each subject each round, with equal chance. Individual factors are independently distributed across subjects, and over the rounds.

Part of the payoffs come from an investment decision, where investing more has higher return if the individual factor is high. Before the investment decision is made, there is a learning stage where each individual observes an interim outcome (a good or a bad signal), which is determined by the interplay of the individual factor and a “common factor” (an unknown random variable that affects the outcome of everyone in the same group in the same way). The common factor, like the individual factor, can be either “high” or “low”, with equal chance, and is drawn independently of the individual factor.¹⁹ The computer randomly matches subjects in 5-person groups each round, with each group being assigned a particular common factor in that round. Observing a good interim outcome is more likely when one’s individual factor is high *or* when the common shock is favorable. Specifically, the probability of a good outcome is represented by the following matrix:²⁰:

	Common Factor High	Common Factor Low
Ind. Factor High	1	0.5
Ind. Factor Low	0.5	0

¹⁹The “common factor” is meant to capture “difficulty” in the theoretical model, except for the slight change in interpretation that a “higher common factor” refers to an easy task (a good shock).

²⁰Notice that this specification corresponds to the parameterization $\rho = 0.5$, $\lambda = 0.5$, $\mu = 0.5$ in the theoretical model.

Social comparison information in this context refers to information on how many individuals have observed a good interim outcome in a given round, coupled with the information on an individual's own outcome. As shown in the theoretical section, keeping one's own outcome constant, Bayesian updating implies a posterior belief about ability that decreases in the number of good outcomes in the group. Therefore, letting s_i denote subject i 's interim outcome, IF_i her individual factor, CF the common factor and ng the number of good outcomes in the rest of the group, we have the following:

(i) $Pr_i(IF_i = H | s_i, ng)$ is decreasing in ng .

(ii) $Pr_i(CF = H | s_i, ng)$ is increasing in ng .

The following figures show the Bayesian posterior beliefs about the individual and the common factor, after different levels of social comparison information, ranging from 0 good outcomes to 4 good outcomes.

Figure 1.10: Bayesian posteriors

Figure 1.11: Bayesian posteriors

I use a within-person design in order to analyze how individuals update their beliefs and decisions when they see their own outcome and when they see others' outcomes. The stages of a given round is as follows:

- Individual and common factors are randomly chosen by the computer (true values never observed by the subjects).
- Own outcome is observed (good or bad).
- Beliefs about individual factor and the common factor are submitted.
- An investment decision is made.
- The number of good outcomes by the other people in one's group is observed.
- Beliefs about individual factor and common factor are revised.
- Another investment decision is made.

- One of the two belief-decision pairs is randomly picked to “count” for payoffs in that round.

In what follows, I explain the procedures in more detail.

Belief Elicitation: To obtain insight into how beliefs are updated when social comparison information is observed, subjects’ beliefs about the individual factor and the common factor were elicited using a quadratic scoring rule, which rewards the individual for more accurate beliefs (see Selten (1998) and Sonnemans and Offerman (2001) for a discussion of such rules in experiments). The elicitation procedure amounts to subjects submitting probabilities for the individual factor or the common factor being high and low, and is incentive-compatible under the assumptions of risk-neutrality and expected payoff maximization. The belief elicitation procedure is implemented twice in a given round: once after observing own outcome, and a second time after observing the number of good outcomes in the group. In each belief elicitation stage, subjects submit their beliefs about both the individual factor and the common factor.

The Decision Problem: After submitting beliefs, subjects are asked to choose a level of investment. The investment decision is structured such that an individual with a high individual factor obtains a higher marginal return from investing than one with a low individual factor. Specifically, every unit invested has a return of 12 points if the individual factor is high and 3 if the individual factor is low. Investment costs are increasing, and convex in the level of investment, given by $C(I) = I^2/2$, where I is the level of investment. Therefore, if individual factors were known, the optimal level of investment by a subject with a high-individual factor would be higher than that of one with a low individual factor.

When individual factors are unknown, more positive beliefs about the individual factor should translate into a higher optimal investment choice. Maximizing the expected payoffs from investment yields

$$I^* = 3 + 9(Pr(IF = H))$$

as the optimal investment choice given beliefs about the individual factor being high. Rather than the costs and benefits of investment, subjects were given a table that directly presents the net payoff from investment depending on ability being high or low. As with belief elicitation, the investment decision was made twice in a given round: pre- and post-social comparison information.

Payoffs: The payoffs in the experiment come from three sources: payoffs from beliefs, payoffs from decisions, and a fixed payoff to be paid in every round. The computer randomly picks either the pre-social comparison or post-social comparison sets of beliefs and decisions to use in the calculation of the actual payoffs in a given round²¹. The specific quadratic scoring rule formula used in the calculation of payoffs from estimation is given in the online appendix, along with the payoffs to different levels of investment.

1.5.2 Hypotheses

The hypotheses generated by the experimental design are as follows:

Hypotheses About Beliefs:

Hypothesis 1: (Actual versus Bayesian posteriors) Posterior beliefs (both pre- and post-social comparison information) are equal to their Bayesian counterparts.

²¹This is done in order to give incentives for the subjects to pay as much attention to pre-information choices as the post-information ones.

Hypothesis 2: (Within-subject ranking of posteriors) Each individual's posterior beliefs about the individual factor should be increasing in her own outcome and decreasing in the number of good outcomes by the others. Posterior beliefs about the common factor should be increasing in own outcome as well as the number of good outcomes by the others.

Hypothesis 3: (Direction of the change in beliefs with social comparison information)

a) The probability assigned to having a high individual factor should increase, compared to its pre-social comparison information level, if less than two of the four outcomes of the remaining people in the group are good. If more than two good outcomes are observed, the probability should decrease, and it should stay the same if exactly two good outcomes are observed.

b) The probability assigned to the common factor being high should increase (decrease) compared to its pre-social comparison information level, if at least (less than) two of the four outcomes of the remaining people in the group are good.

Hypotheses About Investment:

Hypothesis 4: (Optimality of Investment Given Beliefs) Subjects' investment levels are given by $I = 3 + 9(Pr(IF = H))$, where $Pr(IF = H)$ is the submitted belief about having a high individual factor.

Hypothesis 5: (Actual versus "Truly Optimal" Investment) Subjects' investment decisions reflect the optimal investment level given their outcome (and given the Bayesian posterior given that outcome).

Hypothesis 6: (Within-subject ranking of investment levels) Each individual's

investment levels should be increasing in her own outcome and decreasing in the number of good outcomes by the others.

Hypothesis 7: (Direction of the change in investment levels with social comparison information) Investment should increase (decrease), compared to its pre-social comparison information level, if less (more) than two of the four outcomes of the remaining people in the group are good. Investment should stay the same if exactly two good outcomes are observed.

This within-person design enables us to test the predictions of the informational theory of social comparisons by generating data on (1) whether individuals use social comparison information “correctly” in forming their perceptions, and (2) whether their investment decisions are affected by their beliefs (and indirectly the social comparison information they observe) in the way that the theory predicts.

1.5.3 Procedures

The experiments were conducted at the California Social Science Experimental Laboratory (CASSEL) at UCLA, using a Java-based computer program designed for this experiment. Seventy undergraduate students participated in the experiments²², and typically the sessions lasted close to one and a half hours. Twenty rounds were run for pay, and subjects participated in 3 practice rounds before the “real rounds”. After the practice rounds, they were given a computerized quiz, which they needed to complete correctly to be able to proceed. Earnings in the experiment were denominated in “points”, with an exchange rate of 100 points=\$0.70. Average earnings in the experiment were around \$20. Subjects were given two handouts at the start of the experiment: one describing the po-

²²Apart from the sessions whose results are reported here, two pilot sessions were also run.

tential payoffs from estimation (derived from the quadratic scoring rule), and the other the payoffs to different levels of investment, depending on whether the individual factor was high or not. At the end of each round, the subjects were given information about which of their two investment choices was picked, the payoff they obtained, and their total payoff up to that point in the experiment. At the end of the experiment, a survey was given, which collected information about their perceived level of understanding of the decision problem, gender, and major. Below, I present the results from the analysis of the 20 rounds of data from 67 of the 70 subjects that participated.²³ The instructions that were read to the subjects are given in the Appendix.

1.5.4 Results

Consistent with the hypotheses stated above, there are several dimensions in which the data from the experiment can be analyzed. First of all, we can pit the actual beliefs and investment choices of the subjects against the optimal benchmarks given by the theory, and analyze how close to the theoretical model actual behavior is, in terms of magnitude. Second, we can focus on submitted beliefs alone, to see whether the *ranking* of submitted beliefs after different kinds of social comparison information is observed is consistent with the theoretical ranking. Third, since we have within-subject data on how beliefs change upon receiving social comparison information, we can test whether the direction and magnitude of belief-updating is consistent with the theory. In what follows, I analyze first the beliefs and decisions after observing own outcomes only, and

²³Three subjects were taken out of the analysis because they either submitted extremely high investment choices that left them with a hugely negative payoff over the rounds, or their behavior in the experiment as well as their own response to the survey question of how well they understood the setup indicated poor understanding of the instructions.

then move on to the analysis of social comparisons.

1.5.4.1 Pre-Social Comparison Information Results:

I first start by analyzing the submitted beliefs and decisions after observing own outcomes. The following tables show some descriptive statistics regarding the mean and variances of the submitted beliefs and investment choices in the first stage, before social comparison information is observed. The first number in the optimal investment cell gives the optimal investment level given the subject's submitted beliefs, and the second number gives the optimal investment level given the posteriors that *should have been* submitted if subjects were perfectly Bayesian.

Table 1.1: Pre-Information Beliefs and Decisions, Own Outcome=Good			
	<i>Mean(Submitted)</i>	<i>Optimal</i>	<i>Std.Dev.(Submitted)</i>
$Pr(IF = high)$	0.7035	0.75	0.1625
$Pr(CF = high)$	0.6528	0.75	0.1680
<i>Investment.</i>	7.785	9.331/9.75	2.841

Table 1.2: Pre-Information Beliefs and Decisions, Own Outcome=Bad			
	<i>Mean(Submitted)</i>	<i>Optimal</i>	<i>Std.Dev.(Submitted)</i>
$Pr(IF = high)$	0.3519	0.25	0.1919
$Pr(CF = high)$	0.399	0.25	0.1787
<i>Investment</i>	5.013	6.175/5.25	2.594

Actual versus Bayesian posteriors:

The first result to be noted about the first-stage beliefs is that subjects attribute a lower probability than they should to having a high individual factor when they see a good outcome, and a lower probability than they should to having a low individual factor when they see a bad outcome. The same pattern is observed with the common factor: subjects underattribute to having a high factor when they see a good outcome, and underattribute to having a low factor when they see a bad outcome. The Wilcoxon sign-rank test²⁴ shows that when a good outcome is observed, submitted beliefs for having a high individual factor are significantly lower than Bayesian beliefs ($z=-2.920$, $p=0.0035$). When a bad outcome is observed, on the other hand, we have that the probability assigned to having a high individual factor is higher than the Bayesian benchmark ($z=4.881$, $p\text{-value}=0.0000$). Likewise for the common factor, the probability assigned to the high state is lower than it should be when a good outcome is observed ($z=-4.800$, $p=0.0000$), and higher than it should be ($z=5.843$, $p=0.0000$) when a bad outcome is observed. In general, subjects seem to have a strong tendency to assign lower probability than the Bayesian benchmark to the state that is more likely given the outcome (low state when bad outcome is observed, high state when good outcome is observed)²⁵.

Ranking of Beliefs:

Next, we look at the ranking of beliefs when a good outcome is observed versus a bad outcome. As expected, the probability assigned to the high state is higher when a good outcome is observed than when a bad outcome is observed, for both

²⁴To avoid the effects of dependence arising from potential subject-specific effects, the data is averaged across rounds for each subject in the nonparametric tests.

²⁵For all the reported results from non-parametric tests, parametric counterparts such as the t-test give the same result, with similar p-values.

the individual factor and the common factor. Mann-Whitney tests show that this difference is statistically very significant with a p-value of 0.0000 ($z = -9.513$).

Comparing the payoffs that subjects would get from estimating the individual factor and the common factor, we see that there is no statistically significant difference between the two ($p = 0.8628$).

Investment Levels:

The second question is how beliefs translate into decisions. There are two different optimality benchmarks, against which actual investment levels can be compared: optimality *given* submitted beliefs, and optimality against a full Bayesian benchmark (optimal decision given Bayesian posteriors rather than submitted beliefs). Analyzing the investment decisions in the first stage, subjects seem to have a significant tendency to underinvest given their beliefs, and this tendency is present both after observing a good outcome and a bad outcome (the Wilcoxon sign-rank test statistic is $z = -5.168$ when $\text{outcome} = \text{good}$, and $z = -4.712$ when $\text{outcome} = \text{bad}$, both significant at $p = 0.000$). If we compare the observed investment levels with the optimal investment levels with *optimal* beliefs (rather than the actual beliefs submitted by the individual), we get the interesting result that subjects significantly underinvest when they see a good outcome, but the hypothesis that investment levels are equal to the optimal levels when a bad outcome is observed cannot be rejected. The reason for this difference is that subjects tend to underinvest *given* their beliefs, but this effect is counteracted by their tendency to overestimate the probability of having a high individual factor when they see a bad outcome, as mentioned above, which makes observed investment levels closer to the optimal ones with Bayesian beliefs.²⁶

²⁶There are possible explanations for the underinvestment behavior, such as risk-aversion, but we bypass those issues since the within-person design allows us to focus on the change in beliefs and decisions when comparison information is observed and analyze the effects for a

We now move on to the analysis of the effect of social comparison information on beliefs and decisions, to test the main premise of the theoretical model, which is that beliefs about ability (the individual factor) are decreasing in the number of good outcomes in the group.

1.5.4.2 Post-Social Comparison Information Results:

Beliefs: The first question, as before, is how close post-comparison information posteriors are to the Bayesian benchmarks. Figures 1.12 and 1.13 (please see Appendix) plot the beliefs (averaged across all subjects and rounds) about the individual and the common factor, given different levels of social comparisons, against the Bayesian posteriors.

As can be seen from figure 1.13, mean beliefs about the common factor being high increase with the number of outcomes, as the theory predicts, but the submitted belief profile seems to be more conservative than Bayesian posteriors (especially at the extremes), which may be a consequence of risk-aversion. Results from sign-rank tests that pit the observed beliefs against the optimal benchmarks for different types of comparison information sometimes indicate significant differences between the actual beliefs and the Bayesian posteriors, both for the individual factor and the common factor estimation. Mainly, the observation that subjects assign too low a probability to having a high individual factor when the outcome is good and too high a probability when the outcome is bad seems to be preserved. Another observation is that subjects update much less than they should when they see “extreme” types of social comparison information that are in the same direction as their own outcome (i.e. observing 4 good outcomes when one’s own outcome is good, or observing 4 bad outcomes when one’s own outcome

given risk-aversion parameter.

is bad)²⁷.

Comparing the actual beliefs with the Bayesian posteriors is not very informative about the way subjects process the social comparison information, however, because the absolute level of the submitted second-stage beliefs does not say much about the true effect of social comparison information, since any such effect will be confounded with the general updating errors and tendencies of the subjects. The more interesting analysis involves looking at the *change* in beliefs after social comparison information is observed, to see whether the direction of belief updating in response to the observation of others' outcomes is correct. Recall, from Hypothesis 4, that beliefs about the individual factor should go down (up) when more (less) than 2 good outcomes are observed, and beliefs about the common factor should go up when at least 2 good outcomes are observed. I use a random-effects model to test the direction of the learning. Specifically, I regress the change in beliefs after social comparison information is observed on a constant and outcome dummies, allowing for subject-specific effects²⁸.

The following tables illustrate the results of the regressions for the individual factor and the common factor²⁹. We are interested in the effects of the outcome dummies on the change in beliefs when social comparison information is observed. The dependent variable, therefore, is the second-stage belief minus the first stage belief.

²⁷It should be noted that the optimal change in beliefs is lower in these cases too. Suppose that the first stage outcome is good. In this case, observing 4 good outcomes will change beliefs less than observing 4 bad outcomes.

²⁸Session dummies used to account for potential session-specific effects turn out to be insignificant and are not included.

²⁹* denotes statistical significance at the 1% level.

Change in Beliefs About the Individual Factor

	Coeff.	Change	Std. Dev.	z
constant	0.1689	0.1689	0.025	6.75*
numG=1	-0.1104	0.0585	0.0294	-3.76*
numG=2	-0.1835	-0.0146	0.026	-7.06*
numG=3	-0.1963	-0.0274	0.248	-7.9*
numG=4	-0.2077	-0.0388	0.0269	-7.71*

	Coeff.	Change	Std. Dev.	z
constant	0.0361	0.0246	0.0187	1.32
numG=1	-0.0156	0.0205	0.2027	-0.77
numG=2	0.0004	0.0365	0.0249	0.02
numG=3	-0.01826	-0.1465	0.0238	-7.65*
numG=4	-0.2182	-0.1822	0.0279	-7.82*

Change in Beliefs About the Common Factor

	Coeff.	Change	Std. Dev.	z
constant	-0.2246	-0.2246	0.0239	-9.38*
numG=1	0.0238	-0.2007	0.0283	0.84
numG=2	0.0983	-0.1263	0.0306	3.21*
numG=3	0.2910	0.0664	0.0282	10.3*
numG=4	0.4159	0.1914	0.0275	15.12*
numG=5	0.4834	0.2589	0.0331	14.59*

The coefficients of the outcome dummies in the above tables illustrate the effect on beliefs compared to the benchmark case, which is the case where the number of good outcomes in the rest of the group is zero (captured by the constant in the regression). Therefore, in order to understand the effect of observing a particular outcome on the pre- and post-information belief difference, I add the coefficient of the constant to the other coefficients, as given in the “change” column. For instance, looking at the first table, we see that seeing one good outcome by the others increases beliefs by about 6%.

Recall that the prediction of the model is that subjects should increase their beliefs about the common factor being high if they see more than two good outcomes by the other subjects, and decrease them otherwise. Looking at the third table, we see that this is very much borne out in the data: the change in beliefs from the first round to the second round is negative up to three good outcomes, the direction of belief change with different numbers of good outcomes is always in the correct direction, and almost always statistically significant.

The prediction of the model is that first-stage beliefs should be reduced (increased) when more (less) than two outcomes by the remaining subjects are observed, and kept the same if exactly two outcomes by the remaining subjects are observed. The first and second tables above present the results from two separate regressions, depending on the subject’s own outcome. When own outcome is good, it is possible to see that the direction of belief updating between stages with different types of social comparison information is correct. That is, beliefs about the individual factor decrease with the observation of more good outcomes by others. Also, the revised beliefs are higher than first-stage beliefs up to two good outcomes, and lower after that, in line with the theory. One observation that is worth noting is that there is not a big change in beliefs when four good

outcomes are observed as opposed to three.

The behavior when own outcome is bad is somewhat more erratic in the sense that there is little updating in terms of magnitude when 0 to 2 good outcomes are observed by others, although it is mostly in the correct direction (the updating becomes statistically significant if we have fewer social outcome dummies, e.g. smaller than, equal to, and greater than 2 good outcomes). After 2 good outcomes, however, we see a marked and statistically significant decline in beliefs, which theoretically should be the case. Looking at both regressions (when outcome is good and bad), the general tendency seems to be that subjects update less strongly than they should when they see extreme observations of social comparison information that go along with their own outcomes: i.e. observing 4 good outcomes when one's own outcome is good or observing 4 bad outcomes when one's own outcome is bad. In terms of direction, however, the predictions of the informational theory of social comparisons are mainly upheld by the data.

Looking at the first table (where we analyze the effect of social comparisons when own outcome is good), we can see that beliefs increase up to two good outcomes, and decrease after that, consistent with the theoretical prediction, and that social comparison has a statistically significant effect on the change in beliefs. Looking at the second table, the direction of the belief change with a bad own outcome is such that when the good outcomes observed by the rest of the group is greater than or equal to 3, an decline in beliefs are observed, and an increase is observed otherwise, but the effect of observing 0 or 1 good outcomes is not statistically significant.

Another dimension to analyze the data along is to look at the ranking of the beliefs for different levels of social comparison information. Mann-Whitney tests that test the equality of beliefs when less than two or more than two good outcomes

by others is observed (keeping own outcome constant) shows that beliefs about the individual factor are significantly higher when less than two good outcomes by others is observed ($p=0.0001$, $p=0.0000$ for good and bad own outcomes, respectively). Likewise, beliefs about the common factor are significantly higher when more than two good outcomes by others are observed ($p=0.0000$ for both good and bad own outcomes), in line with the theoretical prediction.

One important question, of course, is how valuable social comparison information is for payoffs. A comparison of the pre- and post-information estimation payoffs gives the result that information is valuable for both estimating the individual factor and the common factor, although more so for the common factor ($z = 2.861$, $p=0.0042$ for the individual factor and $z = 10.115$, $p= 0.0000$ for the common factor). Related to this, looking at the payoff difference in the second-stage estimations of the individual and the common factors, it is possible to see that subjects have an easier time using the social comparison information in updating beliefs about the common factor. The payoffs from estimation of the common factor are significantly higher than those from estimation of the individual factor in the second stage (Sign-rank test, $z=-4.722$, $p=0.0000$), a difference which was absent in the first stage when only own information is observed.

Decisions: Figures 1.14 and 1.15 (Appendix) plot the actual investment level (averaged across individuals) against the optimal investment *given* submitted beliefs, and the optimal investment given Bayesian posteriors. It is possible to see that the tendency of the subjects to underinvest continues after observing social comparison information, with a notable exception in the extreme case of a bad own outcome and 4 good outcomes by others. The reason for this difference is twofold: First of all, subjects do have a tendency to overassign a probability to having a high individual factor when they observe a bad outcome, which tends to make investment levels higher than optimal. Second, subjects do not update as much as they should when they see extreme cases of social comparison information (in this case they do not update downward enough). These two effects make the actual investment level higher than optimal in those cases.

A Mann-Whitney test shows that there is a significantly more positive change in investment levels as compared to the first stage when less than two good outcomes by others are observed, for any given own outcome ($p=0.07$ for a good own outcome, and $p=0.0324$ for a bad own outcome). A comparison of absolute investment levels in the second stage, when less or more than two good outcomes by others are observed, shows that investment is higher in magnitude when less than two outcomes are observed, but this effect is not significant at the 5% level.

Overall, it is possible to say that the predictions of the informational theory of social comparisons is borne out by the experimental data in that subjects seem to update their beliefs about the individual and the common factor in the right direction most of the time. The magnitude of the change, however, is not always in line with what the theory predicts. In particular, subjects seem to be less responsive than they should be to extreme outcomes by the group that are in the same direction as their own outcome. In other words, they do not

decrease their beliefs enough when they have observed a good outcome and see that everyone else has also observed a good outcome, and they do not increase their beliefs enough when they have observed a bad outcome and see that everyone else has also observed a bad outcome. As for investment levels, a major tendency that is observed in the data is for the subjects to invest less than the optimal amount given the beliefs they submit, but the underinvestment tendency is in some cases mitigated by the inaccuracy of the beliefs they submit (assigning a higher probability to the individual factor being high), making investment levels closer to the true optimal level (with Bayesian beliefs). Investment levels also seem to be less sensitive to social comparison information than beliefs are.

1.6 Conclusion

This paper puts forward a novel framework for analyzing social comparisons; modeling comparisons as a source of information for ability inference when individuals have imperfect self-knowledge. The model captures, using a standard economic framework and without the need for any external payoff dependence, the fundamental insight that performing worse than others makes an individual feel less competent. It also explains many related findings from the psychology literature. This framework is likely to prove useful for thinking about the effects of comparisons in several different settings, and for predicting when the same comparison is likely to have an effect on behavior and when not (for instance, the model predicts that only “informative” comparisons, e.g. with similar agents, would have an influence on beliefs and hence behavior). Moreover, the results suggest policy recommendations for improving overall performance in multi-agent settings such as the workplace or the classroom, through the manipulation of the

availability and content of information, and hence agents' beliefs.

Restricting attention to exogenous, independent reward schemes, I find that the comparison of different information policies boils down to two main effects: an effort-ability alignment effect at the individual level, and a performance alignment effect across agents. Variants of the first effect has been discussed in other information-economic models in the literature (e.g. Athey and Levin (1998)). The second effect, which may lead the principal to withhold useful social comparison information from the agents is new. In the presence of complementarities across agents in the principal's payoff function, the principal does not like the fact that social comparison information induces a negative correlation between self-confidences, efforts, and outcomes. Consequently, I find that there exist parameter configurations for which the principal would prefer to disclose to agents own performance information but withhold social comparison information, when complementarities are strong enough. When a general incentive scheme with potentially dependent payments across agents can be chosen along with the informational policy, however, the principal now has another tool for manipulating the correlation of agents' efforts and outcomes, and chooses to use "cooperative contracts", which mitigate the negative effects of social comparison information on the association of performances between agents, and makes full revelation optimal.

The model generates several testable implications, and the results relate to organizational design in a fundamental way. The first set of results, as noted above, are based on the nature of the principal's objective function. In settings where the substitutability of performance is high (such as when only the "best" performance matters, as in R&D contexts, or where agents do not possess individual-specific skills that make them irreplaceable in production), the model suggests that so-

cial comparisons should be allowed to the maximum extent, whereas in settings where each agent's effort and performance is not perfectly substitutable for the principal, it may be best for the principal to suppress social comparisons. It is worth noting again, that the policy recommendations from the model need not be exclusively interpreted in terms of direct performance feedback revelation. The same type of effect on motivation may be achieved through a reduction in the comparability of performances across agents (e.g. by making agents who will necessarily observe each other work on different tasks, or forming work groups or teams of agents with divergent backgrounds), instead of the principal directly revealing or hiding the performances of agents from each other. The main results of the paper would be applicable with the appropriate interpretation in such settings as well.

Another important area where results can be applied is educational settings. The self-concept that students develop in school is often crucial for the effort they put into classes, their drop-out decisions, aspirations and, career choices; and an important source of information that shapes the academic self-concept is comparisons with peers. The effects of social comparisons on self-concept and behavior suggest that manipulating the availability and content of comparisons may potentially lead to an improvement in overall educational performance. One way that this can be achieved, where the paper's results would directly apply, is to decide whether to make grade distributions publicly available, or to reveal to students their own scores only. Another, less direct policy tool is ability grouping, since by manipulating the comparison group of students, it might be possible to influence their self-perception and therefore effort. As mentioned before, this has been found to be important in the case of gifted students, and is likely to be important in general in cases where "confidence-management" is an important

objective of policymakers.

The second set of results links the optimal policy to the amount of discretion the principal has in choosing the compensation scheme. In some settings, feedback policies and compensation policies are determined by separate authorities (as in the case of a division manager who has no control over the compensation scheme), and in others, the principal has full control of any policy; monetary or informational. The results indicate that in the presence of complementarities, it would be optimal to use cooperative wage schemes along with full information revelation. This is because cooperative wage schemes where each agent's payment depends positively on the performance of the other agent mitigate the negative effect of social comparisons on the correlation between the two agents' outcomes, thereby restoring the optimality of full information revelation. In terms of testable implications of the model, this suggests that in team-based settings with endogenous cooperative monetary rewards, we should expect to see more frequent interim performance evaluations, and that cooperative wages can be used by employers who are concerned about the effects of relative performance evaluations on morale.

1.7 APPENDIX

FIGURES:

Figure 1.12: Actual vs. Bayesian posteriors, Individual

Figure 1.13: Actual vs. Bayesian posteriors, Common

Figure 1.14: Actual vs. Optimal Investment, Good Outcome

Figure 1.15: Actual vs. Optimal Investment, Bad Outcome

PROOFS:

Derivation of Posterior Beliefs for Lemma 1— (e.g. Ruud (2000)) Let $\mathbf{y} \sim N(\boldsymbol{\mu}, \boldsymbol{\Omega})$. If we partition $\mathbf{y} = \begin{bmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \end{bmatrix}$, $\boldsymbol{\mu} = \begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix}$ and $\boldsymbol{\Omega} = \begin{bmatrix} \boldsymbol{\Omega}_{11} & \boldsymbol{\Omega}_{12} \\ \boldsymbol{\Omega}_{21} & \boldsymbol{\Omega}_{22} \end{bmatrix}$, then

$$\mathbf{y}_1 | \mathbf{y}_2 \sim N[\mu_1 + \boldsymbol{\Omega}_{12} \boldsymbol{\Omega}_{22}^{-1} (\mathbf{y}_2 - \mu_2), \boldsymbol{\Omega}_{11} - \boldsymbol{\Omega}_{12} \boldsymbol{\Omega}_{22}^{-1} \boldsymbol{\Omega}'_{12}]$$

Using the distributions of the random variables in the model,

$$E(a_i | s_i, s_{-i}) = \bar{a}_i + \begin{bmatrix} \sigma^2 & 0 \end{bmatrix} \begin{bmatrix} \sigma^2 + \psi^2 & \psi^2 \\ \psi^2 & \sigma^2 + \psi^2 \end{bmatrix}^{-1} \begin{bmatrix} s_i - \bar{a}_i \\ s_{-i} - \bar{a}_{-i} \end{bmatrix}$$

Hence,

$$E(a_i | s_i, s_{-i}) = \bar{a}_i + \frac{(\sigma^2 + \psi^2)(s_i - \bar{a}_i) - \psi^2(s_{-i} - \bar{a}_{-i})}{\sigma^2 + 2\psi^2}$$

□

Proof of Proposition 1— When $k=0$, the principal's payoff function reduces to

$$\Pi = (E[a_1 e_1^*] + E[a_2 e_2^*])V$$

Inserting the agents' equilibrium efforts yields:

$$\Pi = \beta(E_s[a_1 E(a_1 | S^1)] + E_s[a_2 E(a_2 | S^2)])V$$

By the law of iterated expectations, this is equal to:

$$\begin{aligned} \Pi &= \beta E_s [E [a_1 E(a_1 | S^1) | S^1] + E[a_2 E(a_2 | S^2) | S^2]]V \\ &= \beta E[E(a_1 | S^1)^2 + E(a_2 | S^2)^2]V \\ &= \beta(E(a_1)^2 + E(a_2)^2 + Var(E(a_1 | S^1)) + Var(E(a_2 | S^2)))V \end{aligned}$$

We are interested in the difference in expected profits in the two informational scenarios, $S^i = \{s_i, s_{-i}\}$ (information about others revealed) and $S^i = \{s_i\}$ (information about others withheld), for $i = 1, 2$, which will be given by:

$$\sum_{i=1}^2 Var(E(a_i | s_i, s_{-i})) - Var(E(a_i | s_i))$$

Since agents are ex-ante identical, the above difference is greater than zero if and only if $Var(E(a_i | s_i, s_{-i})) > Var(E(a_i | s_i))$. Since the variance of posterior beliefs will always be higher with more information, more information always increases the principal's expected payoff. In fact, with this particular signal structure, it is possible to show that

$$Var(E(a_i | s_i, s_{-i})) = \frac{\sigma^2(\sigma^2 + \psi^2)}{(\sigma^2 + 2\psi^2)}$$

$$Var(E(a_i | s_i)) = \frac{\sigma^4}{(\sigma^2 + \psi^2)}$$

and

$$Var(E(a_i)) = 0$$

Therefore, $Var(E(a_i | s_i, s_{-i})) > Var(E(a_i | s_i)) > Var(E(a_i))$, which implies $\Pi^{sc} > \Pi^{own} > \Pi^{no}$. \square

Proof of Proposition 2— The principal's expected payoff is given by:

$$\begin{aligned}\Pi &= E(a_1 e_1 a_2 e_2) k + \sum_{i=1}^2 E(a_i e_i)(V) \\ &= \beta^2 E[a_1 E(a_1|S) a_2 E(a_2|S)] k + \beta \sum_{i=1}^2 E(a_i E(a_i|S)) V\end{aligned}$$

Using the law of iterated expectations, the above equation becomes:

$$\begin{aligned}\Pi &= \beta^2 ((E(a_1)^2 + cov(a_1, E(a_1|S)))(E(a_2)^2 + cov(a_2, E(a_2|S))) \\ &\quad + cov(a_1 e_1, a_2 e_2)) k + \beta \left(\sum_{i=1}^2 (E(a_i)^2 + cov(a_i, E(a_i|S))) \right) V\end{aligned}$$

or alternatively,

$$\begin{aligned}\Pi &= (E_S[E(a_1|S^1)^2] E_S[E(a_2|S^2)^2] \\ &\quad + cov(a_1 e_1, a_2 e_2)) k + \left(\sum_{i=1}^2 E_S[E(a_i|S^i)^2] \right) V\end{aligned}$$

where S^i denotes the set of signals available to agent i before he makes his effort decision.

a) The optimal effort level will be equal to $e_i^* = E(a_i)$ when no interim information is given. By Jensen's inequality and the law of iterated expectations, $E_S[E(a_i|s_i)^2] > E(a_i)^2$. The term $cov(a_1 e_1, a_2 e_2)$ in this context captures the correlation in the performances of the agents. If the performances were independent, notice that we would have $E(a_1 e_1 a_2 e_2) = E(a_1 e_1) E(a_2 e_2)$. When no interim information is given, performances are independent ($cov(a_1 e_1, a_2 e_2) = 0$), since effort is equal to the unconditional expectation of ability, and abilities are independently drawn. When own information is given, we will have $cov(a_1 e_1, a_2 e_2) > 0$, because of the common uncertainty assumption, and therefore $E(a_1 e_1 a_2 e_2) > E(a_1 e_1) E(a_2 e_2)$. The specific expression for the covariance with the current distributional assumptions is:

$$cov(a_1 e_1, a_2 e_2)_{own} = \beta^2 \frac{\sigma^4}{(\sigma^2 + \psi^2)}.$$

The result that $\Pi^{own} > \Pi^{no}$ follows.

b)

$$\begin{aligned}
& \Pi^{sc} - \Pi^{own} = \\
& \underbrace{\beta^2(E(a_1)^2 + cov(a_1, E(a_1|s_1, s_2)))(E(a_2)^2 + cov(a_2, E(a_2|s_1, s_2)))}_{>0} \\
& \quad - \underbrace{((E(a_1)^2 + cov(a_1, E(a_1|s_1)))(E(a_2)^2 + cov(a_2, E(a_2|s_2))))}_{>0} \\
& + \beta^2 \underbrace{cov(a_1 E(a_1|s_1, s_2), a_2 E(a_2|s_1, s_2)) - cov(a_1 E(a_1|s_1), a_2 E(a_2|s_2))}_{<0} k \\
& + \beta \underbrace{\left(\sum_{i=1}^2 cov(a_i, E(a_i|s_1, s_2)) - cov(a_i, E(a_i|s_i)) \right) V}_{>0}
\end{aligned}$$

It is possible to show that

$$cov(a_i, E(a_i|s_1, s_2)) > cov(a_i, E(a_i|s_i))$$

because of better alignment between the true state and the posteriors with more information. However,

$$cov(E(a_1|s_1, s_2), E(a_2|s_1, s_2)) < 0 < cov(E(a_1|s_1), E(a_2|s_2))$$

and therefore the second set of terms in the principal's payoff is negative. After deriving the above expressions with our assumptions on the signal structure, we can show that there exist parameters for which revealing social comparison information is dominated by revealing own information only, for k high enough. Since $cov(a_1, e_1) = cov(a_1, E(a_1)) = 0$ and $cov(a_1 e_1, a_2 e_2) = cov(a_1 E(a_1), a_2 E(a_2)) = 0$ with no information, the expected performance of a single agent will be lower with no information than with social comparison information, but the association between the performances will be higher. Therefore, revealing no information at all can dominate revealing social comparison information also, if k is high enough.

Specific expressions for the covariances are given by:

$$\begin{aligned} cov(a_i, E(a_i | s_i, s_{-i})) &= \frac{\sigma^2(\sigma^2 + \psi^2)}{(\sigma^2 + 2\psi^2)} \\ cov(a_i, E(a_i | s_i)) &= \frac{\sigma^4}{(\sigma^2 + \psi^2)} \\ cov(a_1e_1, a_2e_2)_{own} &= \beta^2 \frac{\bar{a}^2 \psi^2 \sigma^4}{(\sigma^2 + \psi^2)^2} \\ cov(a_1e_1, a_2e_2)_{sc} &= \beta^2 \frac{\psi^2 \sigma^2 (\psi^2 \sigma^2 - 3 \bar{a}^2 (\sigma^2 + 2\psi^2))}{(\sigma^2 + 2\psi^2)^2} \end{aligned}$$

Solving for the level of complementarity that sets the payoff differences equal to zero, we obtain

$$\bar{k}_{scown} = \frac{2\psi^2(\sigma^2 + \psi^2)(\sigma^2 + 2\psi^2)}{\beta(-\psi^2\sigma^2(2\psi^4 + 6\psi^2\sigma^2 + 3\sigma^4) + \bar{a}^2(\sigma^2 + 2\psi^2)(4\sigma^4 + 6\psi^2\sigma^2 + \psi^4))} \quad (1.7.1)$$

and

$$\bar{k}_{scno} = \frac{2(\sigma^2 + \psi^2)(\sigma^2 + 2\psi^2)}{\beta(-\sigma^2(2\psi^4 + 2\psi^2\sigma^2 + \sigma^4) + \bar{a}^2(-2\sigma^2 + \psi^2)(\sigma^2 + 2\psi^2))} \quad (1.7.2)$$

The derivative of the payoff difference between social comparison and own information with respect to k is given by:

$$\frac{\partial(\Pi^{sc} - \Pi^{own})}{\partial k} = \frac{\beta\psi^2\sigma^2(\psi^2\sigma^2(3\sigma^4 + 2\psi^4 + 6\psi^2\sigma^2) - \bar{a}^2(\sigma^2 + 2\psi^2)(\psi^4 + 6\psi^2\sigma^2 + 4\sigma^4))}{(\sigma^2 + 2\psi^2)^2(\sigma^2 + \psi^2)^2} \quad (1.7.3)$$

Using (8) and (10), it is possible to see that the payoff difference will be decreasing in k if and only if $\bar{k}_{scown} > 0$. Therefore, when $k > \bar{k}_{scown} > 0$, withholding information is optimal. Comparing the expressions for \bar{k}_{scown} and \bar{k}_{scno} , it is possible to show that if $\bar{k}_{scno} > 0$, then $\bar{k}_{scno} > \bar{k}_{scown} > 0$. \square

Proof of Proposition 3— Since $\bar{k}_{scown} > (<)0$ whenever the payoff difference is decreasing (increasing) in k , we look for conditions under which $\bar{k}_{scown} > 0$. Solving $\bar{k}_{scown} = 0$ for \bar{a} yields:

$$\bar{a}^* = \sqrt{\frac{\psi^2 \sigma^2 (2\psi^4 + 6\sigma^2 \psi^2 + 3\sigma^4)}{(2\psi^2 + \sigma^2)(\psi^4 + 6\sigma^2 \psi^2 + 4\sigma^4)}} \quad (1.7.4)$$

Noting that

$$\frac{\partial \bar{k}_{scown}}{\partial \bar{a}} = \frac{-4\bar{a}\psi^2(\sigma^2 + \psi^2)(\sigma^2 + 2\psi^2)^2(\psi^4 + 6\psi^2\sigma^2 + 4\sigma^4)}{\beta(\psi^2\sigma^2(2\psi^4 + 6\sigma^2\psi^2 + 3\sigma^4) - \bar{a}^2(2\psi^2 + \sigma^2)(\psi^4 + 6\sigma^2\psi^2 + 4\sigma^4))^2} < 0$$

and evaluating the derivatives of the right-hand side of Equation 11 with respect to ψ^2 and σ^2 , which yields $\frac{\partial \bar{a}^*}{\partial \psi^2} > 0$ and $\frac{\partial \bar{a}^*}{\partial \sigma^2} > 0$ establishes the result. \square

Proof of Proposition 4— When the principal bears the cost of incentives, her payoff function is given by

$$\Pi = E(a_1 e_1 a_2 e_2)(k) + (E[a_1 e_1] + E[a_2 e_2])(V - \beta) - 2\alpha$$

Agent i 's expected utility, when he exerts effort e_i will be given by

$$U_i = E[\alpha + a_i e_i \beta - e_i^2/2]$$

The principal's payoff, therefore, can be written as

$$\Pi = E(a_1 e_1 a_2 e_2)(k) + (E[a_1 e_1] + E[a_2 e_2])V - \sum_{i=1}^2 U_i - \sum_{i=1}^2 E(c(e_i)).$$

The cost of effort function, $c(e_i) = e_i^2/2$, is a convex function of the conditional expectation and therefore of the posterior probabilities. Likewise, the expected utility of the agent is also convex in beliefs. By Blackwell's theorem, better information (more signals in our context) raises the expectation of any convex function of posterior beliefs (see Kihlstrom (1984) for a proof). So, $\sum_{i=1}^2 U_i + \sum_{i=1}^2 E[c(e_i)]$ will be higher with social comparison information than it is with own performance information only. Therefore, keeping wages constant, it will be better for the principal to withhold information when payments are made by

the principal, in all cases where withholding information would be optimal when wages do not enter directly into the principal's payoff. \square

Proof of Lemma 4— To reduce notational clutter, let $Pr(s_{-i} = 1 | s_i = 0) = Pr(1|0)$, $Pr(s_{-i} = 1 | s_i = 1) = Pr(1|1)$, $Pr(s_{-i} = 0 | s_i = 1) = Pr(0|1)$ and $Pr(s_{-i} = 0 | s_i = 0) = Pr(0|0)$. Also let $d = w_{SS} + w_{FF} - (w_{SF} + w_{FS})$.

Solving

$$\begin{aligned} & (Pr(s_{-i} = 1 | s_i) e^*(1) E[a_1 a_2 | s_{-i} = 1, s_i] \\ & + Pr(s_{-i} = 0 | s_i) e^*(0) E[a_1 a_2 | s_{-i} = 0, s_i]) ((w_{SS} + w_{FF} - (w_{SF} + w_{FS})) \\ & + E(a_i | s_i) (w_{SF} - w_{FF}) - e_i(s_i)) = 0. \end{aligned}$$

for $i=1,2$ for $e^*(1)$ and $e^*(0)$ yields the following equilibrium effort levels in response to own signals:

$$\begin{aligned} e_i^*(1) &= \frac{((w_{SF} - w_{FF})(E[a_i | 1] - E[a_i | 1] E[a_1 a_2 | 0, 0] Pr(0|0)d + \\ & E[a_i | 0] E[a_1 a_2 | 0, 1] Pr(0|1)d))}{(1 - \alpha(E[a_1 a_2 | 1, 1] Pr(1|1) + E[a_1 a_2 | 0, 1]^2 Pr(0|1) Pr(1|0)d) + \\ & E[a_1 a_2 | 0, 0] Pr(0|0)d(-1 + E[a_1 a_2 | 1, 1] Pr(1|1)d))} \\ e_i^*(0) &= \frac{((w_{SF} - w_{FF})(E[a_i | 0] - E[a_i | 1] E[a_1 a_2 | 0, 1] Pr(1|0)d + \\ & E[a_i | 0] E[a_1 a_2 | 1, 1] Pr(1|1)d))}{(1 - d(E[a_1 a_2 | 1, 1] Pr(1|1) + E[a_1 a_2 | 0, 1]^2 Pr(0|1) Pr(1|0)d) + \\ & E[a_1 a_2 | 0, 0](Pr(0|0)d(-1 + E[a_1 a_2 | 1, 1] Pr(1|1)d)))} \end{aligned}$$

\square

Proof of Lemmas 5 and 6— Calculation and simplification of the covariances, which involve quite messy algebra, are given in the online appendix to this paper (<http://www.bol.ucla.edu/~sertac/research/appendix.html>). \square

Proof of Lemma 7— In the social comparison information case, recall that the effort level was given by:

$$e_i^*(s_1, s_2) = \frac{(w_{SF} - w_{FF})(E(a_i | s_1, s_2) + E(a_{-i} | s_1, s_2)E(a_1 a_2 | s_1, s_2)(w_{SS} + w_{FF} - (w_{SF} + w_{FS})))}{1 - E(a_1 a_2 | s_1, s_2)^2(w_{SS} + w_{FF} - (w_{SF} + w_{FS}))^2}$$

It is possible to see, looking at the above equation, that reducing w_{FF} by an amount x and increasing w_{SS} by the same amount would lead to a higher effort. Since the total cost of providing incentives would be the same in that case, and the principal is risk-neutral, the principal's payoff increases. This means that $w_{FF} > 0$ cannot be optimal for the principal. i.e, for any $w_{FF} > 0$, the principal could increase effort by lowering w_{FF} and increasing w_{SS} , keeping the total wage bill the same. Therefore, $w_{FF}^* = 0$. Setting $w_{FF} = 0$, the same type of argument shows that w_{FS} should be set equal to zero as well. The same logic is applicable in the other cases (own information and no information). \square

Proof of Proposition 5— Define Δ to be the cooperativeness of the wage scheme, i.e. $\Delta = w_{SS} - w_{SF}$. When no interim information is given, agents do not update their beliefs, so their effort levels are known for certain ex-ante. The principal's payoff is given by

$$\Pi = E[a_1 e_1 a_2 e_2 (k - 2(w_{SS} - w_{SF}))] + (E[a_1 e_1] + E[a_2 e_2])(V - w_{SF})$$

Notice that in this case $E[a_i e_i] = E[a_i]E[e_i]$ and that $E[a_i e_i a_j e_j] = E[a_i]E[a_j]E[e_i]E[e_j]$ because of the independence of abilities, and the independence of abilities and effort levels. Inserting the optimal effort in the case of no information, we obtain

$$\Pi = \frac{E(a)^2 w_{SF} (-2(-1 + \Delta E(a)^2 V + (-2 + k E(a)^2) w_{SF})}{(-1 + \Delta E(a)^2)^2}$$

Maximizing this with respect to Δ yields

$$\Delta^* = \frac{V + (-2 + kE(a)^2)w_{SF}}{E(a)^2 V} \quad (1.7.5)$$

As can be seen, any wage policy (w_{SF}, w_{SS}) that satisfies 1.7.5 will be payoff-maximizing. It then follows that the set of contracts that satisfy 1.7.5 will include cooperative, competitive and independent wages, since Δ^* can be positive, negative, or zero. \square

Proof of Proposition 6— The proof relies on an envelope theorem. Again, let $\Delta = w_{SS} - w_{SF}$. Notice that the principal's equilibrium expected payoff is given by:

$$\begin{aligned} \Pi &= E(a_1 e_1^* a_2 e_2^*)(-2(w_{SS} - w_{SF})) + (E[a_1 e_1^*] + E[a_2 e_2^*])(V - w_{SF}) \\ &= (E[a_1 e_1^*] + E[a_2 e_2^*])V - ((E[a_1 e_1^*] + E[a_2 e_2^*])w_{SF} + 2 E(a_1 e_1^* a_2 e_2^*) \Delta) \end{aligned}$$

From agent i 's first-order-condition,

$$E(a_i | S)w_{SF} + E(a_1 a_2 e_j^* | S)\Delta = c'(e_i^*)$$

Multiplying by e_i^* and adding the two agents' first-order conditions gives:

$$[E(a_1 | S)e_1^* + E(a_2 | S)e_2^*]w_{SF} + [e_1^* E(a_1 a_2 e_2^* | S) + e_2^* E(a_1 a_2 e_1^* | S)]\Delta = e_1^* c'(e_1^*) + e_2^* c'(e_2^*)$$

and

$$E_S[E(a_1 | S)e_1^* + E(a_2 | S)e_2^*]w_{SF} + [e_1^* E(a_1 a_2 e_2^* | S) + e_2^* E(a_1 a_2 e_1^* | S)]\Delta = E_S[e_1^* c'(e_1^*) + e_2^* c'(e_2^*)]$$

Applying the law of iterated expectations, it is possible to rewrite the principal's equilibrium expected payoff as:

$$\Pi = (E[a_1 e_1^*] + E[a_2 e_2^*])V - E_S[e_1^* c'(e_1^*) + e_2^* c'(e_2^*)]$$

Therefore, the principal's expected payoff does not depend on the shape of the compensation scheme. \square

Proof of Proposition 7— Expected profits for the principal can be expressed as before; however, now the first expectation is taken over the principal's beliefs about ability and signals, which are different than the agents', and therefore we need superscripts P and A (for principal and agent, respectively) in the expectations. . Formally, we have:

$$\begin{aligned}\Pi &= \beta E^P[a_1 E^A((a_1 | S^1))] + E^P[a_2 E^A((a_2 | S^2))] \\ &= \beta \sum_{i=1}^2 E^P(a_i) E^P(E^A(a_i | S^i)) + cov(a_i, E(a_i | S^i))\end{aligned}$$

Using the above equation along with the formulas for posterior mean beliefs and the law of iterated expectations, we can express expected profits when both signals are observed as follows:

$$\begin{aligned}\Pi_{sc} &= \beta \sum_{i=1}^2 \left(\overline{a^P} E[\overline{a^A} + \frac{(\sigma^2 + \psi^2)}{(\sigma^2 + 2\psi^2)}(s_i - \overline{a^A})] + \frac{\sigma^2(\sigma^2 + \psi^2)}{(\sigma^2 + 2\psi^2)} \right) \\ &= \beta \sum_{i=1}^2 \left((\overline{a^P a^A} + \frac{(\sigma^2 + \psi^2)}{(\sigma^2 + 2\psi^2)} \overline{a^P}(\overline{a^P} - \overline{a^A})) + \frac{\sigma^2(\sigma^2 + \psi^2)}{(\sigma^2 + 2\psi^2)} \right)\end{aligned}$$

whereas the expected payoff in the case where only one's own signal is observed is given by

$$\begin{aligned}\Pi_{own} &= \beta \sum_{i=1}^2 E(a_i) E(E(a_i | s_i)) + cov(a_i, E(a_i | s_i)) \\ &= \beta \sum_{i=1}^2 \left(\overline{a^P} E[\overline{a^A} + \frac{\sigma^2}{(\sigma^2 + \psi^2)}(s_i - \overline{a^A})] + \frac{\sigma^4}{(\sigma^2 + \psi^2)} \right) \\ &= \beta \sum_{i=1}^2 \left((\overline{a^P a^A} + \frac{\sigma^2}{\sigma^2 + \psi^2} \overline{a^P}(\overline{a^P} - \overline{a^A})) + \frac{\sigma^4}{(\sigma^2 + \psi^2)} \right)\end{aligned}$$

With some algebra, it is possible to show that making only one's own signal observable yields higher expected profits than making both signals observable if and only if $\overline{a^P}(\overline{a^A} - \overline{a^P}) > \sigma^2$. \square

Analogues of Some of the First Set of Results in the Discrete Case:

Proof of Lemma 1— To cut back on notation, I will use the symmetry of the agents and denote $Pr(a_i = H | s_i = 0, s_j = 1)$ as $Pr(H | 0, 1)$. Therefore, $Pr(H | s_i, s_j)$ denotes the posterior belief of someone who has observed s_i as her own signal and s_j as her peer's signal of having a high ability. Using Bayes' rule, the posteriors are calculated in the following way:

$$Pr(H | 1) = \frac{Pr(H, 1)}{Pr(1)} = \frac{(\mu(1 - \lambda) + \lambda)\rho}{(\mu(1 - \lambda) + \lambda)\rho + (1 - \rho)\lambda\mu}$$

$$Pr(H | 0) = \frac{Pr(H, 0)}{Pr(0)} = \frac{(1 - \mu)(1 - \lambda)\rho}{(1 - \mu)(1 - \lambda)\rho + (1 - \lambda + (1 - \mu)\lambda)\rho}$$

$$\begin{aligned} Pr(H | 0, 0) &= \frac{Pr(H, 0, 0)}{Pr(0, 0)} \\ &= \frac{(1 - \lambda)\rho(1 - \mu)(\rho(1 - \mu) + (1 - \rho))}{\lambda((1 - \rho)^2(1 - \mu)^2) + (1 - \lambda)(\rho^2(1 - \mu)^2 + 2\rho(1 - \rho)(1 - \mu) + (1 - \rho)^2)} \end{aligned}$$

$$\begin{aligned} Pr(H | 1, 0) &= \frac{Pr(H, 1, 0)}{Pr(1, 0)} \\ &= \frac{\rho(1 - \mu)\lambda(1 - \rho) + \mu(1 - \lambda)(1 - \rho + (1 - \mu)\rho)}{\lambda(1 - \rho)(\mu^2(1 - \rho) + \mu\rho) + \rho(\mu^2(1 - \lambda)\rho + \lambda(\mu(1 - \rho) + \rho))} \end{aligned}$$

$$\begin{aligned}
Pr(H|1,1) &= \frac{Pr(H,1,1)}{Pr(1,1)} \\
&= \frac{\rho(\mu^2(1-\lambda)\rho + \lambda(\mu(1-\rho) + \rho))}{\lambda(1-\rho)(\mu^2(1-\rho) + \mu\rho) + \rho(\mu^2(1-\lambda)\rho + \lambda(\mu(1-\rho) + \rho))}
\end{aligned}$$

$$\begin{aligned}
Pr(H|0,1) &= \frac{Pr(H,0,1)}{Pr(0,1)} \\
&= \frac{\rho(\mu^2(1-\lambda)\rho + \lambda(\mu(1-\rho) + \rho))}{\lambda(1-\rho)(\mu^2(1-\rho) + \mu\rho) + \rho(\mu^2(1-\lambda)\rho + \lambda(\mu(1-\rho) + \rho))}
\end{aligned}$$

Calculating the difference between the beliefs and some algebra gives us the belief ranking given in Lemma 1. \square

Proof of Proposition 3— When the principal bears the cost of incentives, her payoff function is given by

$$\Pi = E(a_1e_1a_2e_2)(k) + (E[a_1e_1] + E[a_2e_2])(V - (w_S - w_F)) - 2w_F.$$

Agent i 's expected utility, when he exerts effort e_i will be given by

$$U_i = E[a_i e_i w_S + (1 - a_i e_i) w_F - e_i^2/2]$$

The principal's payoff, therefore, can be written as

$$\Pi = E(a_1e_1a_2e_2)(k) + (E[a_1e_1] + E[a_2e_2])V - \sum_{i=1}^2 U_i - \sum_{i=1}^2 E(c(e_i)).$$

Hence, the principal's payoff when she bears the cost of incentives is her payoff without contracts minus the total expected payoffs of the agents and the total cost of effort. Let ρ' denote the posterior probability of being of high ability. We can write the agent's second-period equilibrium payoff as

$$U_i = \rho' a_H(e_i^*) w_S + (1 - \rho') a_L(e_i^*) w_F - \frac{e_i^{*2}}{2}$$

Inserting $w_F = 0$, and $e_i^* = .E(a_i|S^i) w_S = (\rho' a_H + (1 - \rho')a_L)w_S$, we have:

$$\begin{aligned} U_i &= \rho' a_H e_i^* w_S - \frac{e_i^{*2}}{2} \\ &= \rho' a_H (\rho' a_H + (1 - \rho')a_L) w_S^2 + (1 - \rho')a_L (\rho' a_H + (1 - \rho')a_L) w_S^2 \\ &\quad - \frac{(\rho' a_H + (1 - \rho')a_L)^2 w_S^2}{2} \end{aligned}$$

which is convex in the posterior beliefs, since

$$\frac{\partial^2 U_i}{\partial \rho'^2} = (a_H - a_L)^2 w_S^2 > 0.$$

The cost of effort, $c(e_i) = e_i^{*2}/2$, is also a convex function of the conditional expectation and therefore of the posterior probabilities. By Blackwell's theorem, better information (more signals in our context) raises the expectation of any convex function of posterior beliefs (see Kihlstrom (1984) for a proof). So, $\sum_{i=1}^2 U_i + \sum_{i=1}^2 E[c(e_i)]$ will be higher with social comparison information than it is with own performance information only. Therefore, keeping wages constant, it will be better for the principal to withhold information when payments are made by the principal, in all cases where withholding information would be optimal when wages do not enter directly into the principal's payoff. \square

INSTRUCTIONS FOR THE EXPERIMENT ³⁰

Welcome to CASSEL. This is an experiment on decision-making. Your earnings will depend partly on your decisions, and partly on chance. Your earnings in this experiment will be denominated in “points”. At the end of the experiment, the total number of points that you have earned will be converted to dollars at an exchange rate of 0.0070, which means that every 100 points is worth 70 cents. In addition, a \$5 show-up fee will be added to your earnings, for your participation. You will be paid privately, in cash, at the end of the experiment. You are under no obligation to tell anyone how much you have earned in the experiment. Now, please write your name and computer number (starting with SSEL) on the record sheets that you have been given. This will be your identification number. I will now read the instructions for the experiment. Please listen to these instructions carefully, and do not hesitate to interrupt if anything is unclear.

There will be 20 rounds in this experiment. The computer will randomly assign you to a group with 4 other participants each round. Who you will be put in the same group with is completely randomly determined. In each round, the computer will assign you an “individual factor”, which can have either a high value or a low value, with 50% chance. Individual factors are independent across participants. This means that your having a high individual factor does not make it more or less likely that other participants have a high individual factor. In each round, the computer will also select a “common factor”, which also could be either high or low, with 50% chance. The common factor will take the SAME VALUE for everyone in your group. In other words, the individual factors can be different across people (everyone will randomly receive an individual factor that

³⁰The instructions were read to the subjects, and were supplemented by a powerpoint presentation that included bullet-point and graphical representations of the main points, and some screenshots from the experiment.

can be either high or low), but the common factor will be either high for everyone in your group or low for everyone in your group in a given round. Rounds are independent, meaning that your individual factor, whom you are in a group with, and the common factor selected for your group in a particular round does not at all affect your individual factor, whom you are in a group with, or the common factor in the next round.

[show Slide 1]

The individual factor and the common factor together determine your “outcome” each round, which can be good or bad.

[show Slide 2]

If the common factor for your group is high in a particular round, you get a good outcome FOR SURE if your individual factor is high, and you get a good outcome with 50% chance if your individual factor is low.

If the common factor is low, you get a bad outcome FOR SURE if your individual factor is low, and you get a good outcome with 50% chance if your individual factor is high.

[show Slide 3]

At the start of each round, you will be shown what your outcome was in that round. You will not know what your individual factor was, and you will not know what the common factor was.

Part of your payoffs in this experiment will come from how accurately you guess your individual factor and the common factor. Another part of your payoffs will come from a decision you will make. I will now explain each of these stages in detail.

Estimation Stage:

After you see your outcome, you will be asked to state your beliefs on how likely you think it is that you had a high individual factor in that round. In particular, you will see two statements on your computer screen:

Percent chance that my individual factor was high

Percent chance that my individual factor was low

[Show the belief submission stage screenshot—Slide 4]

And for each statement, you are going to enter the percent chance with which you think that statement is true. For example, if you think that there is a 60% chance that your individual factor was high, after seeing your outcome, you will put 60 in the upper text box and the program will automatically display 40 (100 minus 60) in the lower text box.

You will be compensated for the accuracy of your estimation. This means that the payoff you will get from the percent chances you submit depends on which of the two statements was actually true.

[show quadratic scoring rule payoff graphs—Slides 5-7]

Now, please look at Handout 1, which shows the payoffs corresponding to each possible set of beliefs (percent chances) that you may submit. You can see that the worst payoff you can possibly get is zero, and this happens if you assign a chance of 100% for a statement that is not true. You can also see that your payoff increases with the percent chance that you assign to the true statement. These payoffs have been generated according to a formula that is designed in such a way that you will be maximizing your expected payoff if you enter what you truly believe about each statement being correct.

[Show Slide with QSR example—Slide 8]

After you complete this estimation, the same procedures will be repeated for the common factor. That is, after submitting your beliefs about the individual factor, you will also be asked to submit your beliefs about the common factor in the same fashion.

[Show Slide with common factor screenshot-Slide 9]

Your total payoffs from estimation will be the sum of your payoffs from your estimation of your individual factor and your estimation of the common factor.

Up to this point in a round, you will have seen your own outcome and submitted your beliefs about whether you had a high or low individual factor and common factor in that particular round. The next stage of the experiment is the decision stage.

In the decision stage, you are going to be asked to make an investment decision. This involves a choice of how many units you would like to invest. The return to your investment in a round depends only on your individual factor in that round (it does not depend on the common factor). Your payoffs from investment will be higher if you have a high individual factor. Please look at the payoff table given in Handout 2. (show payoff table on-screen as well). You can see that for each investment level, there are two corresponding payoffs, depending on whether you have a high or low individual factor. The payoff to a certain level of investment when you have a high individual factor is higher than the case when you have a low individual factor.

[Show Slide 10—investment example]

After you complete the first three stages of a round (observe outcome, make

estimations about the individual and the common factor, make investment decision), you will be given information on the outcomes of the 4 other participants in your group. You will see, on your computer screen, how many of the four other participants observed a good outcome in that round.

[Show Screenshot of the group information stage]

After receiving this information, you will go through the stages of estimating your individual factor and the common factor again. You will also make a second investment decision. Therefore, at the end of the round, you will have made two estimates for the individual factor, two estimates for the common factor, and two investment decisions: once before observing the outcomes of others in your group, and once after. The computer program will randomly pick one of these sets of estimates and decisions with equal chance, and those will be used in the calculation of your payoffs. That is, there is a 50% chance that your initial decisions may “count” for your payoffs.

In total, your total payoffs in a round will be given by:

[Show payoff slide—Slide 12]

We are going to run three practice rounds on the computer for you to get familiar with the program and the procedures. Your decisions in these three rounds will not affect your earnings. After these rounds, you will be asked to complete a short quiz to test your understanding of the procedures. We will then start the real rounds.

Please click on the icon titled “MC” (red dog) on your desktop, and log in when prompted. Please also pull out your dividers at this point. If you have questions at any point in the experiment, please raise your hand and the experimenter will assist you.

REFERENCES

- [1] Alicke, M.D. (2000). "Evaluating Social Comparison Targets" in J. Suls and L. Wheeler (Eds.) *Handbook of Social Comparison: Theory and Research* (pp.271-293). New York: Kluwer Academic/Plenum Publishers.
- [2] Athey, S. & Levin, J. (1998). "The Value of Information in Monotone Decision Problems", MIT Working Paper, No. 98-24.
- [3] Barr, S.H., & Conlon, E. (1994). "Effects of Distribution of Feedback in Work Groups" *Academy of Management Journal*, 37(3), 641-655.
- [4] Benabou, R., & Tirole, J. (2002). "Self-Confidence and Personal Motivation", *Quarterly Journal of Economics*, 117(3), 871-915.
- [5] Benabou, R., & Tirole, J. (2003). "Intrinsic and Extrinsic Motivation", *Review of Economic Studies*, 70, 489-520.
- [6] Brickman, P., & Bulman, R. (1977). "Pleasure and pain in social comparison" in J. Suls & R. L. Miller (Eds.) *Social Comparison Processes: Theoretical and Empirical Perspectives* (pp. 149-186). Washington, DC: Hemisphere.
- [7] Che, Y.. & Yoo, S. (2001). "Optimal Incentives for Teams", *American Economic Review*, 91(3), pp.525-41.
- [8] Clark, A. E., & Oswald, A. J. (1996). "Satisfaction and Comparison Income", *Journal of Public Economics*, 61(3), 359-381.
- [9] Clark, A.E., & Oswald, A. J. (1998). "Comparison-Concave Utility and Following Behavior in Social and Economic Settings", *Journal of Public Economics*, 70, 133-155.
- [10] Englmaier, F., & Wambach, A. (2002). "Contracts and Inequity Aversion", IZA Working Paper, No. 1643.
- [11] Ertac, S., Molnar, J. & Virag, G. (2005). "The Signaling Role of Contracts with Imperfect Self-Knowledge and Type-Dependent Outside Options", in progress.
- [12] Falk, A. & Knell, M. (2004). "Choosing the Joneses: Endogenous Goals and Reference Standards", *Scandinavian Journal of Economics*, 106 (3), 417-435.

- [13] Fang, H. & Moscarini, G. (2005). "Morale Hazard", *Journal of Monetary Economics*, 52, 749-777.
- [14] Gervais, S. & Goldstein, I. (2004) "Overconfidence and Team Coordination" Working Paper, Duke University Fuqua School of Business.
- [15] Gibbs, M. J. (1994). "An Economic Approach to Process in Pay and Performance Appraisals", Working Paper.
- [16] Grund, C., & Sliwka, D. (2005), "Envy and Compassion in Tournaments" *Journal of Economics and Management Strategy*, 14 (1), 187-20.
- [17] Hopkins, E. & Kornienko, T. (2004) "Running to Keep in the Same Place: Consumer Choice as a Game of Status", *American Economic Review*, 94 (4), 1085-1107.
- [18] Ilgen, D.R., Fisher, C.D., & Taylor, S.M. (1979). "Consequences of Individual Feedback on Behavior in Organizations", *Journal of Applied Psychology*, 64, 349-371.
- [19] Itoh, H. (2004). "Moral Hazard and Other-Regarding Preferences", *Japanese Economic Review*, 55(1), 18-45.
- [20] Kihlstrom, R.E. (1984) "A Bayesian Exposition of Blackwell's Theorem on the Comparison of Experiments", in M. Boyer and R.E. Kihlstrom (eds.), *Bayesian Models of Economic Theory*, Elsevier.
- [21] Koszegi, B. (2005). "Ego-Utility, Overconfidence and Task Choice", forthcoming in the *Journal of the European Economic Association*.
- [22] Lizzeri, A., Meyer, M. & Persico, N. (2003), "The Incentive Effects of Interim Performance Evaluations", mimeo, New York University.
- [23] Mares, V. & R. Harstad (2003). "Private Information Revelation in Common-Value Auctions", *Journal of Economic Theory*, 109, pp. 264-282.
- [24] Marsh, H. (1987). "The Big-Fish-Little-Pond Effect on Academic Self-Concept", *Journal of Educational Psychology*, 79, 280-295.
- [25] Marsh, H. W., Kong, C. K., & Hau, K. T. (2000). "Longitudinal Multi-level Models of the Big-Fish-Little-Pond Effect on Academic Self-Concept: Counterbalancing Contrast and Reflected-Glory Effects in Hong Kong Schools," *Journal of Personality and Social Psychology* 78, 337-349.
- [26] Meyer, M. & Vickers, J. (1997). "Performance Comparisons and Dynamic Incentives", *Journal of Political Economy*, 105(3), 547-581.

- [27] Molnar, J. & Virag, G. (2004). "Revenue Maximizing Auctions with Externalities and Signaling", Working Paper, University of Rochester.
- [28] Morrison, E.W., & Cummings, L.L. (1992). "The Impact of Diagnosticity and Performance Expectations on Feedback Seeking Behavior", *Human Performance*, 5, 251-264.
- [29] Ok, E.A. & Kockesen, L. (2000). "Negatively Interdependent Preferences", *Social Choice and Welfare*, 17, 533-558.
- [30] Prat, A. (2002). "Should a Team Be Homogeneous?" *European Economic Review*, 46(7), 1187-1207.
- [31] Ruud, P.A. (2000). "An Introduction to Classical Econometric Theory", New York: Oxford University Press.
- [32] Santos-Pinto, L. & Sobel, J. (2003). "A Model of Positive Self-Image in Subjective Assessments." Department of Economics, University of California, San Diego.
- [33] Selten, R. (1998). "Axiomatic Characterization of the Quadratic Scoring Rule", *Experimental Economics*, 1, 43-62.
- [34] Sonnemans, J., & Offerman, T. (2001). "Is the Quadratic Scoring Rule Really Incentive Compatible?," Working paper, CREED, University of Amsterdam.
- [35] Suls, M.R. & L. Wheeler (Eds.) (2000). "Handbook of Social Comparison: Theory and Research", New York: Kluwer Academic/Plenum Publishers.
- [36] Zeidner, M. & Schleyer, E. (1999). "The Big-Fish-Little-Pond Effect for Academic Self-concept, Test Anxiety, and School Grades in Gifted Children", *Contemporary Educational Psychology*, 24, 305-329.

CHAPTER 2

An Experimental Study of Ego-Relevance in Information Processing

2.1 Introduction

Given the ubiquity of situations where economic agents face uncertainty regarding a decision, individuals' attitudes to information and the way they process information once they receive it is a central issue in economics. In many decision-making contexts, individuals receive (usually ambiguous) feedback in the interim, and are confronted with the task of using this information to revise their beliefs and perceptions about the problem. The main assumption of standard economic theory about the processing of information is that agents are "Bayesian", in the sense that they use Bayes' rule in updating their prior beliefs when faced new information, which then affects the decisions that they make. Although standard economic theory makes no distinctions in these predictions according to the "nature" or the content of the information, in many settings the attitudes toward and the processing of information crucially depend on the type of information in question. Information received in work settings or academic settings (e.g. the outcome of a project or an investment decision, or performance in a task that requires skill) usually has a bearing on one's self-image and morale, and is "ego-relevant". On the other hand, some types of information collected or

received before making decisions can be completely irrelevant to one's self-image (e.g. evaluating the success of a project that has been designed by other people). It is natural to think that the reaction of individuals to information and the way they process it may be different when the information is ego-relevant, from both the case of information that is unrelated to the self, and the predictions of standard economic theory. For instance, when information is ambiguous, individuals might use the information in a self-serving way, e.g. downplay or ignore negative information or information that conflicts with their initial beliefs. Indeed, there is a considerable literature in psychology on motivated reasoning, ego-protection and ego-enhancement motives (see Baumeister (1998) for a review). It is also a well-known fact, from both psychology and economics literatures, that judgment biases such as overconfidence may come into play when the decision at hand involves a trait that is important to individuals' self-image, potentially affecting how beliefs are formed and information is processed. Since in reality, many of the decisions in economic settings involve such ego-relevant information (especially in organizational settings), addressing the issue of ego-relevance in self-assessments and information processing is crucial from an economic standpoint.

This paper experimentally tests the predictions of economic theory that individuals process information optimally (that is, using Bayes' rule), in the context of both information that is relevant to ego and information that is ego-irrelevant. Performance feedback from tasks that require ability and effort is used as the ego-relevant information, since ability, or "doing well" in tasks is usually the most important trait that affects self-image in organizational and educational settings. Behavior in this setting is contrasted with a case where individuals face an updating problem that does not have a personal aspect. By maintaining the exact statistical structure of the updating task, but changing the nature

(self-relevance) of the information, I study whether any observed departures in behavior from the theoretical predictions are due to the nature of the information, as opposed to general errors in Bayesian updating that are common to both types of information.

Given the importance of information processing in decision-making under uncertainty, the attitudes toward and the processing of information has been the focus of both theoretical and empirical work. Psychologists have long identified several heuristics and biases in judgment under certainty, such as “representativeness” (tendency of individuals to judge the likelihood of an event by the degree to which it resembles a stereotype, regardless of its base rate frequency—see, for example, Kahneman, Slovic & Tversky (1982)). Rabin (1998) provides a review of these biases as they relate to economic decision-making. As mentioned before, there is also a big literature in psychology on how “ego-protection” and “ego-enhancement” motives affect attitudes to information and attributions when ambiguous information is received.

Bayesian updating and more generally, information processing has been studied in several contexts in the economics literature as well. There are, for instance, several studies that test experimentally for confirmation bias, which refers to the tendency to interpret new evidence to fit existing beliefs (Jones (2003), Jones and Sugden (2001), Dave and Wolfe (2005)), and Rabin and Schrag (1999) have put forward a theoretical model of how this bias affects beliefs and decisions. Ertac (2005) studies Bayesian updating in a signal extraction problem in the context of social comparisons, where subjects update their beliefs about two unknown factors using their own outcome as well as the outcomes of others, and finds evidence of conservatism (insufficient use of information). Charness and Levin (2005) study the interaction between Bayesian updating and the reinforcement

heuristic. Belief updating has also been studied in the context of social learning and information cascades (e.g. Kraemer, Noth, and Weber (2001)).

In addition to the experimental studies, there are also several theoretical studies that focus on departures from Bayesian behavior in information processing (e.g. Epstein, Noor and Sandroni (2005)). Koszegi (2004, 2005) analyzes the effects of “ego utility” on information collection and decision-making, whereas Benabou and Tirole (2002) put forward a model with “selective memory” where individuals have a probability of not remembering unfavorable information.

Because it uses data from a real-effort performance task and analyzes self-knowledge, the current paper is also related to the economics literature on overconfidence and self-assessments (e.g. Hoelzl and Rustichini (2005), Camerer and Lovallo (1999)). The paper closest to the current paper in motivation and design is Clark and Friesen (2005). This paper also studies self-assessments in a performance context by eliciting beliefs, and analyzes how individuals make predictions about their future performance and revise their predictions after completing the task. An important difference of the current paper with this paper is that it elicits beliefs about performance in an already completed task rather than about future performance, and gives individuals direct feedback about that same past performance and looks at the revision in beliefs. This design allows us to explicitly calculate Bayesian probabilities, and compare the submitted data with the theoretical predictions. Moreover, the current study also implements a set of non-performance updating rounds, which helps isolate the role that the nature of the information plays in belief updating. A major contribution of this paper to the literature, in fact, is that the experimental design allows a direct comparison of ego-relevant and ego-irrelevant information, by keeping the statistical nature of the updating problem and hence the theoretical predictions exactly the same, and changing

only the nature of the information, which helps isolate the reason behind any departures from Bayesian behavior.

The primary goal of the paper is to identify the effects of information being “personal” on its use. Given that much of the information that people are confronted with in real life have a personal connotation, the answer to the question of how such information is actually processed is crucial for addressing issues such as how much performance feedback to give employees, which is an important question in organizational design.

The organization of the paper is as follows: Section 2 puts forward the experimental design, procedures, and hypotheses. Section 3 presents and discusses the results, and Section 4 concludes.

2.2 Experimental Design

2.2.1 Overview

To study the processing of ego-relevant information, I use an experimental setting where information is valuable from a material payoff perspective, but potentially affects self-image. This design is meant to capture the usual trade-off between the material benefits of using information optimally (improved decision-making), and the potential psychological costs/consequences. The experiment is composed of two main parts, which differ in the type of updating task and the information that the participants are confronted with. In the first part of the experiment, participants work on tasks that require ability and effort, and the information in question is their performance. In the second part of the experiment, there is no task, and the information that participants are asked to process is information

about a state of the world that has no relevance to self-image. The material benefit of holding accurate beliefs is induced by paying the subjects for accuracy of their estimation of the “true state”. In the ego-relevant rounds, the true state is the subject’s place in the performance distribution in the task that was just completed. Specifically, subjects work on a task that requires ability and effort in each round, and then are asked to state how likely they think they were in the top 20%, middle 60%, or bottom 20% of the performance distribution. The beliefs they submit for each of these three states (top, middle, and bottom) are compensated according to a quadratic scoring rule. In the ego-irrelevant rounds, on the other hand, there is no performance stage, and the true state is one of three outcomes (top, middle, bottom) that is randomly picked by the computer according to a pre-determined probability distribution. In these rounds, the states “top”, “middle” and “bottom” have no relation to performance, but are just possible outcomes that have different likelihoods of being chosen. As in the ego-relevant rounds, subjects are asked to assign probabilities to each state being picked, and are compensated for the accuracy of these beliefs.

2.2.1.1 Stages of the Performance Rounds

In the “performance rounds” of the experiment, subjects work on a task for a specified amount of time, and are compensated for each question they answer correctly¹. They are initially not given any information on their performance, and are asked to state their beliefs about their likelihood of being in the top 20%, middle 60%, or bottom 20% of the performance distribution. These estimates, as mentioned before, are compensated using a quadratic scoring rule².

¹More information about the tasks is given in Section 2.2.2.

²The specific formula used for the quadratic scoring rule is given in Section 2.2.2.

To study how ego-relevant information is processed, after the initial estimation of performance, participants are given (sometimes ambiguous) feedback about their performance. Specifically, subjects are told whether they were in the top of the distribution or not. Naturally, if they receive information that they were in the top 20%, there is no remaining uncertainty about performance, and if they learn that they were not in the top, they need to update their beliefs about how likely they are to be in the middle or the bottom. After getting this feedback, subjects are asked to submit their beliefs about their performance again, and this second estimation is also compensated in the same fashion as the first, using a quadratic scoring rule. The total earnings in a given performance round is therefore equal to the sum of the payoffs received from performance (which is compensated using a piece-rate for each question correctly solved), and the payoffs received from accurate estimation of performance, which includes both the pre-feedback and post-feedback estimation payoffs.

2.2.1.2 Stages of the Non-Performance Rounds

In the non-performance rounds of the experiment, there is no task performance stage. Instead of being linked to performance, the states “top”, “middle” and “bottom” are just outcomes that have certain prior probabilities of being picked by the computer. At the beginning of each round, subjects are given these prior probabilities, and then are asked to assign a likelihood to each of the three states being chosen in that round. As in the performance rounds, these beliefs are compensated using a quadratic scoring rule³.

After this initial estimation, they are then given information on whether the

³Notice that assuming risk-neutrality, it is optimal for subjects to just submit the objective prior probabilities in this first estimation.

chosen state was “top” or not. That is, they learn the true state if it was top, and if not, they know that it can either be middle or bottom. After the information stage, subjects are asked once again to submit probabilities for the three states.

The innovative point of the current design is that it replicates the exact updating problem from a statistical perspective, with and without a real task. The way this is achieved is as follows: the computer program records the probabilities that were submitted by each participant in the performance rounds, and uses exactly those probabilities as the prior probabilities of each state being picked in the non-task rounds. In other words, for every performance round, there is a corresponding non-task round where the subject faces the same prior probabilities for the three states, the difference being only that the states are not related to performance anymore. This allows us to compare behavior across the cases of ego-relevant and ego-irrelevant information, when the statistical nature of the problem is the same and economic theory makes the exact same predictions for behavior. Notice, however, that due to the random nature of the outcome in the non-performance rounds, it is not always possible to replicate the same feedback in the non-performance rounds as the corresponding performance round. For instance, although priors are the same across the two rounds, one performance and one not, the true state in the non-performance round can be different from that in the performance round, leading to a difference in the information received by the subject and hindering comparability.

The next section explains the procedures of the experiment in more detail.

2.2.2 Procedures

The experiments were conducted at the California Social Science Experimental Laboratory (CASSEL) at UCLA, using undergraduate UCLA students as subjects. The experiment was run on the computer, using a JAVA-based software⁴.

7 sessions with 10 participants each were conducted. Additionally, one pilot session was run to test subjects' understanding of the instructions and the workings of the computer program. Each session consisted of 4 practice rounds and 16 paid rounds, and lasted about an hour and fifteen minutes, with payments averaging around \$20. The exchange rate used in the experiment was 0.025. After the instructions were read, participants played 3 practice rounds, and then completed a short quiz for testing their understanding

The first 8 rounds of the experiment were "performance rounds" where subjects had to work on two different tasks that require ability and effort. After these rounds, the instructions for the second part of the experiment (the non-performance rounds) were read, and subjects participated in another practice round as well as a short quiz, before starting the real rounds. A survey was given at the end, which elicited information about how much value subjects attached to good performance in each task, which task they enjoyed more, and how they were affected by the feedback they received, along with other information about the clarity of the instructions etc.(see Appendix for the set of questions used in the survey).

Tasks and Performance The task that was used in the first four rounds of the performance part of the experiment was an algebra task. In these rounds,

⁴The program was developed by Raj Advani.

subjects were given 2 minutes to complete as many addition questions as they could (there were 16 questions in each round). The questions involved adding five two-digit numbers⁵. Earnings from performance were calculated according to a piece-rate, with each correct answer being worth 50 points. The computer then ranked everyone’s performance (number of questions correctly solved), and determined the “top”, “middle” and “bottom” groups. Since each session consisted of 10 subjects, the top 20% group corresponded to the top two performers, the bottom 20% to the worst two performers, and the middle 60% the six subjects in between. Any ties in the number of questions solved were broken using a timed “tie-breaker” question, given to the subjects after completing the normal 2-minute task⁶.

The second task used in the experiment was a verbal task, which consisted of working on 16 questions in 2.5 minutes. Subjects were given multiple-choice questions from GRE verbal tests, which involved analogies and antonyms⁷. Four rounds were run with the addition task, and four rounds with the verbal task.

Belief Elicitation To obtain insight into how subjects estimate their performance and how they update their beliefs when information is observed, beliefs were elicited using a quadratic scoring rule, which rewards the individual for more accurate beliefs (see Selten (1998) and Sonnemans and Offerman (2001) for a discussion of the use of this procedure in experiments). The elicitation procedure, in our design, amounts to the subjects submitting percent chances for being in

⁵This task has also been used in the performance experiments in Niederle and Vesterlund (2006).

⁶The subject who got the tie-breaker question correctly and in the least amount of time was ranked higher.

⁷The questions were taken from “GRE: Practicing to Take the General Test” (2001), and were used in the experiment with permission of Educational Testing Service, the copyright owner.

the top 20%, middle 60% or the bottom 20% of the performance distribution in the performance rounds, and for the state top, middle or bottom being picked in the non-performance rounds. The payoff from these beliefs depends on the actual true state, and is increasing in the probability assigned to it. The exact formula used in the experiment is given by:

$$\pi(p_i, p_{-i}) = 50 + 100(p_i) - 50(p_i^2 + \sum p_{-i}^2)$$

where p_i denotes the probability assigned to the true state, which can be either top, middle or bottom.

As evident from the above formula, the maximum payoff that can be obtained from a single estimation is 100, and the minimum is 0. This belief-elicitation procedure is incentive-compatible under the assumptions of risk-neutrality and expected payoff maximization, that is, it is always optimal for the subjects to submit their true beliefs.

The software used in the experiment was designed such that subjects could observe their potential payoffs from their estimation on their screen as they entered probabilities in the relevant boxes.

2.2.3 Discussion

The experimental design described above has numerous advantages. First of all, the elicitation procedure makes it possible to clearly observe the prior and posterior beliefs of subjects, and allows us to study both the updating behavior of the subjects by pitting observed posteriors against Bayesian benchmarks, and the correctness of their priors in the performance context. Moreover, it provides a clear and simple way of motivating the subjects for holding accurate beliefs,

without actually confronting them with a complicated decision problem. The data on the initial priors of the subjects after completing the task is useful by itself, since it gives information on the level of overconfidence (or underconfidence) of the subjects, which can be interesting to correlate with the way each subject responds to different types of feedback.

The most important feature of the design, however, is that it can directly contrast information processing in performance and non-performance contexts. It is quite likely that the updating behavior of individuals will be affected by some general errors, heuristics and biases, regardless of the nature of the information in question. For instance, subjects may have a general tendency to behave conservatively in information processing, i.e. not update their beliefs enough in response to new information. The replication of the exact same updating task from a statistical perspective, only with a different type of information, can help isolate the role (if any) that the performance context has on behavior. The use of the non-performance context also aids in addressing a potential issue with belief elicitation in this setup. One general problem with quadratic scoring rules is that the beliefs submitted may potentially depend on risk preferences. For instance, while risk-neutral subjects would always submit their true beliefs, risk-lovers might have an incentive to distort their beliefs toward the extremes, assigning a probability of 1 to the event that they think most likely, and betting on getting the highest monetary payoff if it turns out to be true. By giving subjects the objective probabilities of each state being chosen in the non-performance task and asking them to still submit beliefs about each state being chosen, compensated by the scoring rule, we obtain data that can reveal whether subjects have incentives to submit anything other than their true beliefs, when the scoring rule is used.

2.2.3.1 Hypotheses

Given the data collected from the experiment, there are several interesting questions that one can look at, and several different dimensions to analyze the data across. The three main issues are summarized by the following hypotheses:

1. Subjects hold correct initial beliefs.

Many psychology studies have documented “overconfidence”, in the sense that people usually label themselves “better-than-average” in ability (e.g. Alicke (1985)), or are unrealistically optimistic about their prospects in life. With the data on the initial priors submitted by the subjects, it is possible to see whether subjects are overconfident or have correct priors about their performance.

2. Subjects are Bayesian: Subjects use Bayes’ rule in computing the probabilities for each state after getting information. That is, the submitted posterior probabilities for the three states are equal to the Bayesian posteriors, which are given by:

$$p'_T = 0, \quad p'_M = \frac{p_M}{p_M + p_B}, \quad p'_B = \frac{p_B}{p_M + p_B}$$

if the information observed indicates that the true state is “not top”, where p_M and p_B denote the prior probabilities for the states “middle” and “bottom”, respectively, and

$$p'_T = 1, \quad p'_M = 0, \quad \text{and} \quad p'_B = 0$$

if the information observed indicates that the true state is “top”.

3. Given the same priors in the performance and non-performance treatments, the posterior beliefs submitted should be the same.

2.3 Results

It is useful to start the analysis with a description of the data obtained from the experiment. Since we have 8 performance rounds and 8 non-performance rounds in each session, and 7 sessions of 10 subjects each, we have 560 observations in performance rounds, and 560 in non-performance rounds. On average, subjects solve 3.73 questions in the addition task, and 4 questions in the verbal task, and the actual performance in the two tasks is not significantly different. One potential issue with the belief elicitation procedure and the task payoffs could be that the compensation of beliefs may give rise to “moral hazard”, in the sense of subjects purposefully not solving questions in the task part and assigning a 100% chance to being in the bottom of the distribution, to reap sure payoffs from estimation. In our design, the maximum payoff that can be received from a single estimation is 100 points, and the marginal payoff to solving one question is 50 points. Looking at the data, we see that there is no evidence to suggest moral hazard: in very few instances do subjects get all questions wrong (or less than 2 questions right), and those cases are usually not coupled with a tendency to submit 100% for bottom. In the next three sections, we analyze behavior in the performance and non-performance rounds of the experiment, and then compare the two.

2.3.1 Performance Rounds

2.3.1.1 Initial Self-Assessments

We first start analyzing the data from the performance rounds. It is useful to start with some observations about the initial beliefs in the 8 performance rounds.

A quick look at the data suggests that subjects understand the belief submission procedure quite well. In all 560 observations, the beliefs submitted are coherent in the sense that if the probability assigned to top is higher than that assigned to the bottom, then the probability assigned to middle is also higher than that assigned to the bottom, which makes sense since performance is being assessed. The quadratic scoring rule does not seem to induce individuals to submit extreme beliefs, since there are only 43 (out of 560) observations where the initial beliefs involve assigning 100% to one state. And out of these 43 observations, only 10 occur in the initial round of a task.

The first interesting question that arises while analyzing the prior belief data from the performance rounds is how accurate subjects' self-assessments are, and whether they are over- or underconfident. A conservative measure of overconfidence in the aggregate would be to look at the "state choices" of each subject, ignoring the actual probabilities submitted. In order to do this, we construct an "initial choice" variable by looking at which state each subject picked as being most likely in each round. For instance, if a subject submitted a probability of 40% for top, 30% for middle and 30% for bottom, her "choice" would be classified as "top". If there is no miscalibration in the aggregate, we would expect 2 people in each session to have chosen top, 2 bottom, and 6 middle.

The table below represents these initial choices separately for the addition task and the verbal task (the notation "state A ~ state B" is used to indicate that the subject has submitted the same probability for both states, which is higher than the third state):

Table 2.1: Initial Choices of States

	Top	Top~Middle	Middle	Middle~Bottom	Bottom
Addition	57(20%)	11(4%)	164(59%)	25(9%)	23(8%)
Verbal	30(10%)	17(6%)	139(50%)	47(17%)	47(17%)

As can be seen crudely from the above table, there seems to be a marked difference between aggregate miscalibration in the two tasks, addition and verbal, with the subjects being clearly underconfident as a group in the verbal task and somewhat overconfident in the addition task.

At the individual level, overconfidence can be measured by comparing the subjects' beliefs about their performance by their actual performance level. Again using the "choice" of each subject and pitting it against that subject's true performance, we get the following distributions for each task .

Table 2.2: Actual versus Perceived Performance in the Addition Task

<i>Choice</i>	<i>Actual Performance</i>		
	Bottom	Middle	Top
Bottom	8	15	0
Bottom~Middle	12	12	1
Middle	33	107	24
Middle~Top	2	5	4
Top	1	29	27
Total	56	168	56

Table 2.3: Actual versus Perceived Performance in the Verbal Task

<i>Choice</i>	<i>Actual Performance</i>		
	<i>Bottom</i>	Middle	<i>Top</i>
Bottom	22	23	2
Bottom~Middle	14	30	3
Middle	19	95	25
Middle~Top	0	8	9
Top	1	12	17
Total	56	168	56

As the above tables show, subjects can be classified as “correct” in 142 instances (51% of the time), “underconfident” in 56 instances (20% of the time), and “overconfident” in 82 instances (29% of the time). In the verbal task, people are correct in 134 instances (48% of the time), underconfident in 92 instances (33% of the time), and overconfident 54 times (19% of the time). A t-test that tests the hypotheses $H_0 : choice - actual\ performance = 0$ is easily rejected for the verbal task, with $t = -2.54$. For the addition task, the overestimation is borderline significant at the 5% level with $t = 1.98$.

For a more accurate measure of the strength of biases in self-assessments in the two tasks, it is possible to construct the following expected performance index, by assigning values 2, 1, and 0 to “top”, “middle”, and “bottom”, respectively:

$$E[p] = p_T (2) + p_M (1) + p_B (0)$$

which can be compared with the actual performance of the subject, by assigning

2 to the state of “top”, 1 to the state of “middle”, and 0 to the state of “bottom”. With this comparison, for example, a subject whose expected performance was greater than 1 and whose actual performance was “middle” would be classified as overconfident. And, it is possible to construct a measure of the strength of the misprediction by taking the difference between the expected and actual performance indices. The following table shows some summary statistics for this difference.

Table 2.4: Performance Assessments and Overconfidence: Summary Statistics

	Expected Performance Index	Overconfidence Index (Expected-Actual Performance)
Addition	1.044679 (0.462)	0.045 (0.596)
Verbal	0.85375 (0.45)	-0.146 (0.556)

Note: Numbers reported are means, with the standard deviations in parantheses.

As can be seen from the above tables, in terms of beliefs compared to actual performance, there seems to again be a difference between the two tasks, with subjects being more confident in the addition task. A two-sample t-test⁸ confirms that overconfidence, as defined above, is higher in the addition task than in the verbal task (p=0.05). This variation in the distribution of initial beliefs between the two tasks could potentially be helpful in answering the question of whether there is a link between overconfidence and the degree of Bayesianism. In terms of the level of “confidence” of the subjects in their estimates (measured by the variance of the expected performance index), there is no significant difference between the tasks (p=0.46), suggesting that subjects do not feel more comfortable guessing performance in one task than the other.

⁸The data was averaged for each subject across rounds, to control for dependence.

2.3.1.2 Reactions to Feedback

Given the focus of the paper, a very important question is how individuals react to feedback. The feedback structure used in this experiment is such that in each round, two of the subjects learn their true performance perfectly, and the others get ambiguous information. Not surprisingly, subjects who learn that they are in the top revise their beliefs as they should most of the time, assigning 100% to the state top. However, there are 14 instances (7 in the performance rounds) where subjects still attribute a positive probability to being in the top when the computer tells them they aren't. This can be due either to errors, misunderstanding of the instructions, or disbelief on the subjects' part that feedback is always truthful.

As expected, information decreases the variance of expected performance, if we define the variance as:

$$Var[p] = p_T (2 - E[p])^2 + p_M (1 - E[p])^2 + p_B (0 - E[p])^2$$

Notice that this measure will take the value of 0 if one state is assigned a probability of 1. The average of this across subjects and rounds is 0.264 before receiving the performance feedback, and 0.20 after receiving the feedback.

What we are more interested in, of course, is the direction and strength of subjects' belief updating when they learn that they are not in the top. For this, we compare the revised beliefs of each subject with the corresponding Bayesian posteriors. Notice that observing that one is not in the top still leaves room for not being in the bottom of the distribution. If subjects would like to avoid thinking that they are in the bottom of the distribution, for ego-utility reasons,

it might make sense for them to engage in self-serving processing of information, which would lead to a higher-than-Bayesian probability being assigned to the state “middle” and lower-than-Bayesian to “bottom”. The following table shows the mean and standard deviation of the revised minus Bayesian posteriors, for each task type:

Table 2.5: Bias=Actual-Bayesian Posterior for Middle

Addition	-0.0623 (0.1515)
Verbal	-0.025 (0.1283)
Non-Performance Rounds	0.1022 (0.287)

Note: Standard deviations in parantheses

A Wilcoxon signed-rank test⁹ that tests the bias (defined as the actual revised belief minus the Bayesian probability) being equal to zero when feedback is ambiguous (“not top”) shows, interestingly, that the posterior probability assigned to “middle” is significantly less than the Bayesian posterior, which implies that the probability assigned to bottom is higher than it should be. This effect, however, is more significant in the addition task ($z = -4.169$, $p=0.0000$) than in the verbal task ($z=-2.552$, $p=0.01$).

2.3.2 Non-Performance Rounds

In the non-performance rounds, rather than revealing their priors themselves, subjects are given the objective probabilities with which each state is to be picked by the computer. As mentioned before, to allow direct comparison with the use of

⁹Observations were averaged across subjects to account for dependence. Parametric tests that tests the hypothesis that the bias=0 also generate similar results, with ($t=-4.45$ for the addition task, and $t=-2.34$ for the verbal task).

information in the performance rounds, the prior probabilities were taken directly from each subject's submitted probabilities in the performance rounds. Before making the comparison with the performance rounds, however, we analyze the non-performance rounds in isolation and study how close to Bayesian subjects' information processing is in this case.

The first estimation that subjects make after they are given information on the prior probabilities is useful for understanding how well the quadratic scoring rule works in eliciting true beliefs. Given the scoring rule and the objective probabilities, a subject should submit exactly the prior probabilities in the first estimation round, if she is a risk-neutral expected utility maximizer. In total, there 560 comparisons of objective priors with submitted probabilities are available. In the majority of cases, there is little to no difference between the objective and submitted probabilities (see Figures 2.1 and 2.2 in the Appendix for the distribution of the differences), suggesting that risk-aversion or risk-loving was not much of an issue and that the quadratic scoring rule worked quite well in inducing truthful revelation.

In order to analyze how information is processed, we again compare the revised beliefs with their Bayesian counterparts. The following table summarizes the "bias" when the subject sees ambiguous information, where bias is defined as the actual posterior minus the Bayesian posterior assigned to the state "middle". On average, subjects' posterior beliefs for the state middle are about 10% higher than they should be.

Table 2.6: Bias=Actual-Bayesian Posterior for Middle

	Mean	Std. Dev
Non-Performance Rounds	0.1022	0.287

N=398

A t-test reveals that the probability assigned to middle is significantly higher than the Bayesian posterior when the feedback is ambiguous (middle or bottom), with $t=2.5867$, and $p=0.01$. Notice that this is in contrast to the result that was observed in the performance rounds, which is that subjects attribute a lower probability than they should to being in the middle. In order to test whether the level of bias in the beliefs submitted for the state “middle” is different in the performance versus non-performance rounds, we use a non-parametric Mann-Whitney test, which confirms that the bias is significantly more positive in the non-performance rounds than the performance rounds ($z=4.171$, $p=0.0000$). A two-sample t-test also confirms this finding ($t=3.6034$, $p=0.0004$).

While the above tests clearly show that there is over-attribution of probability to “middle” in the non-performance rounds and under-attribution to it in the performance rounds, another important question is whether updating is more accurate in one setting than the other. To analyze this, we calculate the absolute value of the bias attributed to the state middle when feedback is ambiguous, and compare the performance and non-performance rounds. A Mann-Whitney test indicates that the absolute value of the bias, interestingly, is higher in the non-performance rounds than in the performance rounds, and this is significant at the 5% level.

2.3.3 Direct Comparison of Performance and Non-Performance Rounds

Since the prior probabilities used in the non-performance rounds were taken directly from the prior probabilities submitted in the performance rounds, our design allows for a stronger comparison between information processing in the cases of ego-relevant and ego-irrelevant information.

We take a conservative approach and focus only on the cases in which the submitted probabilities in the first stage estimation of the non-performance task are exactly the same as the probabilities submitted in the performance task. Another limitation on the comparability of the data from the performance and non-performance rounds is that even though the priors are the same, the feedback can be different in a matched pair of rounds, because of the randomness involved in the determination of the outcome in the non-performance rounds. Therefore, we restrict our attention to cases where the initial submitted probabilities as well as the feedback observed are the same across the performance and the corresponding non-performance rounds.

Our main question, of course, is how subjects use information in the two settings. We focus on non-trivial cases of updating, i.e. ignore the rounds in which the prior indicates that one state will be picked with probability 1. Figure 2.3 (Appendix) shows the distribution of the posterior probability difference.

Statistical tests indicate that the posterior probability difference in the non-performance versus performance rounds is significantly positive. A Wilcoxon sign-rank test that is performed on the within-subject non-task/task probability difference (again, averaged over rounds for each subject to control for dependence across rounds) gives us $z=3.011$ and $p=0.0026$, and a t-test also confirms this result with $t=2.8689$ and $p=0.0052$. That is, both parametric and non-parametric

tests confirm that the actual posterior for the state “middle” is significantly higher when the updating problem does not involve performance-related information. Since the Bayesian posterior is the same in the two settings, due to the statistical equivalence of the two updating problems, this means that the bias is more likely to be on the positive side (over-assignment of probability to the state “middle”) in the non-performance rounds.

Another question is whether subjects are better Bayesian updaters when performance information is involved than when the information in question has no self-relevance. In order to answer this question, we calculate the absolute value of the bias by taking the discrepancy between actual posteriors and the Bayesian posteriors for “middle”, for both the performance rounds and the corresponding non-performance rounds. We then take the difference of this absolute bias in the performance round and the corresponding non-performance round. A t-test reveals that the absolute bias difference is very significantly higher in the non-performance rounds than in the performance rounds ($t=-5.28$, $p=0.0000$), which is confirmed by non-parametric tests also. This result is also consistent with the results obtained using the across-subject comparison before.

2.3.4 Discussion and Extensions

The results on the initial beliefs of the subjects indicate that there is a difference between the two tasks, addition and verbal, in terms of subjects’ level of overconfidence. This difference is useful, since it gives us a chance to see whether there are any correlations between the type of initial beliefs and the way information is processed. Subjects seem to be significantly more pessimistic about their performance in the verbal task, and the incidence of overconfidence is certainly higher in the addition task. This finding is consistent with the studies

that find that people are more likely to be overconfident in domains that they perceive to be “easy” (e.g. Kruger (1999), Moore and Cain (2005)), since most of the subjects have revealed in the survey that they found the verbal task more difficult. It is also consistent with Clark and Friesen (2005), who find evidence of underconfidence in a verbal task.

In terms of Bayesian updating, all levels of analysis seem to yield the following main conclusion: When ambiguous negative information is received, subjects attribute a lower-than-Bayesian probability to the state “middle” when the information is self-relevant, and a higher-than-Bayesian probability when the updating task is irrelevant to the self. This result is quite unexpected because if subjects are affected by a self-serving bias in the interpretation of ego-relevant information, they would be expected to assign a higher probability to the more favorable state (middle) than Bayesian, after ambiguous negative information is received. The data indicate, on the contrary, that subjects over-use the information received in the self-relevant case. The reason for the observed behavior may be that subjects are prone to getting overly pessimistic once unfavorable information is received, or a psychological tendency to hedge themselves by receiving good monetary payoffs, in case they are indeed in the bottom. On the other hand, the observed behavior may also be because the initial priors in the performance task are “fuzzy”, and the subjects are not confident in their beliefs at all to start with, which may potentially skew their response to information. In order to see whether this explanation may be correct, it makes sense to inquire whether the tendency to over-assign probability to “bottom” is correlated with the initial variance of beliefs (how “confident” the subject was in her estimation). The correlation between these two is found to be only 0.04, so this explanation seems to be unlikely to be true.

Unlike the observations in the performance rounds, in the non-performance rounds, there is a clear tendency on the part of the subjects to assign too much probability to the middle state when “not top” is observed. This may be due to the fact that the probabilities given in the non-performance rounds are ultimately taken from the performance rounds, which makes it likely that middle has a higher probability than bottom, given the likelihood of the two states in the performance rounds. In this sense, the assignment of too much probability to middle is consistent with “conservatism” in information processing, as in Dave and Wolfe (2005) and Ertac (2005), who find evidence of insufficient use of information in similar, ego-irrelevant updating problems.

In terms of the question of whether subjects have ego about the tasks that were used in the experiment, it may be useful to consider the survey responses. Although these responses are not compensated and therefore are open to misrepresentation, they may still give us a glimpse into how subjects viewed the tasks. The histograms describing the frequencies of the answers to the question of “how important was it to you to do better than others in the addition (verbal) task?” are given in the Appendix (Figures 2.4 and 2.5), and show that in many cases, people have a preference as to where they are in the performance distribution.

The current experimental results raise a few important questions. One important point to consider is that reactions to ego-relevant information may be dependent on how favorable the information is. That is, it is a valid question whether there would be any differences in Bayesian updating if the feedback structure revealed “bottom” vs. “not bottom” rather than “top” vs. “not top”. To investigate this, further sessions should be run with this feedback structure, and results compared with those obtained with the “top/not top” type of feedback. Also, running a session where beliefs are not compensated may help understand whether the

non-existence of self-serving bias is due to the reward structure.

In terms of the current data, the next step would be to more thoroughly analyze the data at the individual level, e.g. look for correlations between a subject's initial bias level, the way she responds to information, and classify different types of behavior among the subject population. In a related sense, the data should also be more closely analyzed in the light of individual responses to the survey questions. Another interesting issue to explore would be to look at whether gender affects the overconfidence level of the subjects, and also how they respond to information.

2.4 Conclusions

Information processing, because of its crucial role in decision-making under uncertainty, is a fundamental issue in economics. Bayesian updating, which is one of the fundamental building blocks of economic theories of information, has been an interesting topic to explore for both economists and psychologists, and laboratory evidence has suggested that individuals' updating behavior may not always be Bayesian in nature.

In many economic settings, such as the workplace, the decisions that agents need to make or the information that they need to process is relevant to themselves, such as how they have performed in a certain task, whether they should continue investing time and effort into a certain course of action etc. From an organizational design point of view, it is important to know what agents' response to feedback about their performance would be, for the optimal design of interim performance feedback policies. It is well-accepted in the psychology literature that individuals tend to interpret evidence in ways that make it look favorable

to them, or try to preserve favorable beliefs, when the information in question has a bearing on their self-image. However, there is not much work that analyzes how self-relevant information is processed using rigorous methods, in a context with belief elicitation and financial motivation to hold accurate beliefs. This study attempts to fill this gap by constructing a setting where individuals are rewarded for holding accurate beliefs, but may have psychological preferences over the information they receive. The paper analyzes information processing with both self-relevant and self-irrelevant information, keeping the statistical nature of the updating problem intact. A major strength of the experimental design is the comparison of two settings where the theoretical predictions for information processing are exactly the same, but the nature of the information is different.

The results from the first set of experiments described in this paper are quite puzzling, in the sense that they go against the received hypotheses that individuals are likely interpret self-relevant feedback in a self-serving way, which, in the context of this experiment, would correspond to assigning a higher-than-Bayesian probability to the remaining favorable state (middle) when ambiguous negative information is received. Instead, the results show that subjects over-assign probability to the outcome that is not psychologically favorable (being in the bottom). Behavior in the non-performance rounds departs from the Bayesian hypothesis too, but is in stark contrast to the behavior in the corresponding rounds: subjects assign too much probability to the state “middle”, which is consistent with conservatism in information processing, i.e. not using the information sufficiently in belief-updating.

Overall, the results give insight into how individuals process performance information to form beliefs about themselves, and suggest that the way individuals respond to information can be quite different depending on the type of the in-

formation in question. While more experiments are needed to pinpoint the exact reason behind the observed differences between the two settings, the existing data suggest that individuals may not always be employing self-serving strategies in responding to performance feedback, but rather may end up with a more negative self-image than they should after receiving ambiguous unfavorable information, which is important for how feedback policies should be constructed in organizations where inducing accurate decisions while preserving self-confidence may be an important goal.

2.5 APPENDIX

INSTRUCTIONS¹⁰

Welcome to CASSEL. This is an experiment on decision-making. Your earnings will depend partly on your performance, partly on your decisions, and partly on chance. Earnings in this experiment will be denominated in “points”. At the end of the experiment, the total number of points that you have earned will be converted to dollars at an exchange rate of 0.0025, which means that every 100 points is worth 25 cents. In addition, a \$5 show-up fee will be added to your earnings, for your participation. You will be paid privately, in cash, at the end of the experiment. You are under no obligation to tell anyone how much you have earned in the experiment. Now, please write your name and computer number (starting with SSEL) on the record sheets that you have been given.

The experiment will consist of two parts. We will now explain the procedures of the first part of the experiment. You will also have a chance to play some practice rounds, to help you get familiar with the computer program. Please raise your hand if you have any questions throughout the experiment, and the experimenter will assist you.

Instructions for Part 1—Performance Rounds

This experiment will have two main parts, and each part will consist of several rounds. In the first part of the experiment, you will be working on tasks that require ability and effort. In each round, you will be presented with a set of questions to work on within a fixed amount of time. Part of your earnings will

¹⁰The instructions were read to the subjects, and were supplemented by a powerpoint presentation that included some graphs, bullet-point representation of the main points, and some screenshots from the experiment.

come from your performance in these questions. For each question that you answer correctly within this time frame, you will earn 50 points. [show Slide 1] Once this stage is completed, the computer will record and rank everyone's performance. A tie-breaker question will be used to resolve ties, if any. We will give you more information about the tie-breaking procedures later. Once performances are ranked, the computer will form three groups: the top 20% of the performance distribution, which corresponds to the top 2 performers of the 10 participants, the middle 60%, which corresponds to the next 6 performers of the 10 participants, and the bottom 20%, which corresponds to the worst 2 performers of the 10 participants. [show Slide 2]

You will then be asked to submit your beliefs about the likelihood of being in each performance group [show Screenshot 1]. You will see three statements on your computer screen (I was in the top 20% of the distribution, I was in the middle 60% of the distribution, I was in the bottom 20% of the distribution), with text boxes across from each statement. You will then be asked to submit the percent chances with which you think each statement is true. Naturally, the percent chances that you submit should add up to 100%, because your performance will be in one of these three groups.

The second part of earnings comes from the accuracy of your assessment of your performance. The payoff you are going to get from the beliefs you submit about your performance depends on what your performance actually was. When you submit your beliefs, you will see the potential payoffs that you may get on your screen, depending on what your true performance was. The payoff rule is such that the worst thing to do in terms of payoffs would be to assign 100% chance to an incorrect statement, and the best thing to do would be to assign 100% chance to the correct statement. Your potential payoffs will vary with the specific

probabilities you submit, and the correct state. The payments are determined such that you will be maximizing your ex-ante expected payoff if you are truthful about your assessment of your performance.

After you have made your initial assessment of your performance, you will receive feedback. This feedback will be in the form of revealing whether you were actually in the top 20% of the distribution or not. That is, if you were actually in the top, you will learn it. If you were not in the top, you will get a message that says: “you were not in the top, and may be in the middle or the bottom”. Remember that the feedback given by the computer will always be truthful. After receiving the feedback, you will be asked to submit your beliefs about your performance again. This estimation will be compensated using the same payoff rule as the first estimation. After you have made this second estimation, the round will end.

Your total earnings in a given round in this first part of the experiment will be the sum of the earnings you get from task performance (50 points*number of correct questions), and the payoffs you get from the initial assessment of your performance and the second assessment of your performance after the feedback.

We will now start the computer program and run three unpaid “practice rounds” for you to get more familiar with the stages of the experiment and the software. The “task” that you will face in these practice rounds will be different than the tasks that you will face in the real rounds¹¹. You will also face more questions and have more time in the actual paid rounds. These practice rounds are just for you to familiarize yourself with Now, please pull out your dividers, and double-click on the MC (red dog) icon on your desktop.

¹¹A number sequence task was used in the practice rounds, only for illustration purposes. That is, subjects were given questions involving the guessing of the logically following number in a sequence of numbers with a pattern. Subjects were given two questions, and only 30 seconds to work on the task in the practice rounds.

[Go over instructions briefly once again, step by step as they actually play the stages on the computer]

[Talk about the tie-breaker question as they see it on their screen: “After the normal time for the task is up, you will see a single tie-breaker question on your screen. This question will be used by the computer to resolve any ties in the ranking and determination of the top, middle and bottom groups. If one or more participants have the same number of correct questions, the computer ranks the person who got the tie-breaker question right and in a shorter amount of time higher”]

The task that you will be working on in the first four rounds of the experiment is “addition”. You will be given 15 addition questions at the start of each round, and will have 120 seconds to solve as many questions as you can.

[Show Slide with examples]¹²

Instructions for Part 2—Non-performance rounds

In this second part of the experiment, there will be no task performance stage. In these rounds, “top”, “middle” and “bottom” will have nothing to do with your performance, but will just be “states of the world” that have certain chances of being chosen by the computer. In other words, in each round, the computer will pick either top, or middle, or bottom. You will be given information on how likely each state is to be picked, before the round starts. For example: if the computer has a 50% chance of picking top, 25% chance of middle, and 25% chance of bottom, you will see this on your computer screen. You will then be asked to submit your beliefs about the likelihood of different states to be picked

¹²Before the 4 rounds with the verbal task, subjects were shown an example of the type of questions they are going to face.

in that round. You will be paid for the accuracy of these beliefs, in the same fashion as in the performance rounds.

After you make this initial estimation, you will be given some information about the state actually picked by the computer. The computer will tell you whether the chosen state was top, or not top. After receiving this information you will be asked to submit your beliefs again about the likelihood of the three states being the chosen state. This second assessment is also compensated in the same way as the first. We will run 8 rounds with this configuration.

[After the 8 rounds] The decision part of the experiment is now over. You will now see a survey on your screen. Please fill in your answers in the text boxes. Thank you very much for your participation.

SURVEY QUESTIONS:

1. Which of the two tasks (algebra and verbal) do you think is more reflective of a person's overall ability or intelligence?
2. Which of the two tasks did you find more difficult?
3. Which of the two tasks did you enjoy more?
4. If/when you received favorable feedback about your performance (the message "you were in the top"), did it affect your morale? Please answer on a scale of -10 to 10: -10=extremely negatively, 0=not at all, 10=extremely positively. If not applicable, please write N/A.
5. If/when you received favorable feedback about your performance in a round, did it affect your morale? Please answer on a scale of -10 to 10: -10=extremely negatively, 0=not at all, 10=extremely positively. If not applicable, please write N/A.
6. If/when you received unfavorable feedback about your performance (the message "you were not in the top"), did it affect your morale? Please answer on a scale of -10 to 10: -10=extremely negatively, 0=not at all, 10=extremely positively. If not applicable, please write N/A.
7. If/when you received unfavorable information about your performance in a round, do you think it affected your performance in the subsequent rounds? Please answer on a scale of -10 to 10: -10=extremely negatively, 0=not at all, 10=extremely positively. If not applicable, please write N/A.
8. How important was it for you to do better than others in the addition

task? Please answer on a scale of 0 to 10: 0=not important at all, 10=extremely important.

9. How important was it for you to do better than others in the verbal task? Please answer on a scale of 0 to 10: 0=not important at all, 10=extremely important.

10. Please rate your understanding of the experiment on a scale of 0 to 10. How clear were the instructions?

Some Summary Statistics and A Correlation Table:

Table 2.7: Frequency of Choices of Subjects in Questions 1-3

	Addition	Verbal
Q1-Link to General Intelligence	39 (56%)	31 (54%)
Q2-Perceived Difficulty	20 (29%)	50 (71%)
Q3-Enjoyment	47 (67%)	23 (33%)

Table 2.8: Correlation Table for Questions 1-3

	Relation to Ability	Difficulty	Enjoyment
Relation to Ability	1		
Difficulty	-0.7826	1	
Enjoyment	0.5526	-0.9385	1

FIGURES

Figure 2.1: Difference Between the Submitted and Objective Prior Probabilities
For the State “TOP”

Figure 2.2: Difference Between the Submitted and Objective Prior Probabilities
For the State “MIDDLE”

Figure 2.3: The Difference in the Posterior Probabilities for Middle in the Non-Performance and Performance Rounds

Figure 2.4: Histogram of Responses to Survey Question 8

Figure 2.5: Histogram of Responses to Survey Question 9

REFERENCES

- [1] Baumeister, R. F. (1998). "The Self." in D. Gilbert, S. Fiske, and G. Lindzey, (Eds.), *The Handbook of Social Psychology*, vol. 1 (pp.680-740). Boston: McGraw-Hill.
- [2] Benabou, R. & Tirole, J. (2002). "Self-Confidence and Personal Motivation" *Quarterly Journal of Economics*, 117(3), 871-915.
- [3] Camerer, C. & Lovallo, D. (1999). "Overconfidence and Excess Entry: An Experimental Approach" *American Economic Review*, 89(1), 306–318.
- [4] Charness, G. & Levin, D. (2005) "When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect", *American Economic Review*, 95 (4), 1300-1309.
- [5] Chetan, D.& Wolfe, W. W. (2004). "On Confirmation Bias and Deviations from Bayesian Updating" Working paper.
- [6] Clark, J. & Friesen, L. (2005). "Rational Expectations of Own Performance: An Experimental Study" Working paper.
- [7] Ertac, S.(2005). "Social Comparisons and Optimal Information Revelation: Theory and Experiments", Working paper.
- [8] Epstein, L.G., Noor, J. & Sandroni, A. (2005). "Non-Bayesian Updating: A Theoretical Framework" Working Paper, University of Rochester.
- [9] "GRE: Practicing to Take the General Test" 9th Edition, Educational Testing Service, 2001.
- [10] Hoelzl, E. & Rustichini, A. (2005). "Overconfident: Do You Put Your Money on It?", *Economic Journal* 115, 305-318.
- [11] Jones, M. (2003). "Positive Confirmation in Rational and Irrational Learning" Working Paper No. 150, Department of Economic Studies, University of Dundee.
- [12] Jones, M. & R. Sugden (2001). "Positive Confirmation Bias in the Acquisition of Information." *Theory and Decision*, 50, 59-99.
- [13] Koszegi, B. (2001). "Ego Utility and Information Acquisition.", Working Paper, University of California, Berkeley.

- [14] Kozsegi, B. (2005). "Ego-Utility, Overconfidence and Task Choice", forthcoming in the Journal of the European Economic Association.
- [15] Kraemer, C., M. Noth, and M. Weber (2001). "Information Aggregation with Costly Information and Random Ordering: Experimental Evidence." Working Paper, University of Mannheim.
- [16] Kruger, J. (1999). "Lake Wobegon Be Gone! The "Below-average Effect" and the Egocentric Nature of Comparative Ability Judgments" *Journal of Personality and Social Psychology*, 77, 221-232.
- [17] Moore, D. A. & Cain, D. M. (2005). "Overconfidence and Underconfidence" Working Paper, Carnegie Mellon Tepper School of Business.
- [18] Niederle, M. and Vesterlund, L. (2006). "Do Women Shy Away From Competition?" Working Paper, University of Pittsburgh.
- [19] Rabin, M. (1998). "Psychology and economics" *Journal of Economic Literature*, 36, 11-46.
- [20] Selten, R. (1998). "Axiomatic Characterization of the Quadratic Scoring Rule", *Experimental Economics*, 1, 43-62.
- [21] Sonnemans, J., & Offerman, T. (2001). "Is the Quadratic Scoring Rule Really Incentive Compatible?," Working paper, CREED, University of Amsterdam.