Expert Advice, Control, and Heterogeneous Beliefs

Leonidas Enrique de la Rosa*
delaRosa@econ.au.dk

November 29, 2006

Abstract

This paper studies the effects of overconfidence in an investment-decision setting. A risk-averse agent privately observes information relevant to an investment decision, that he can report to a principal. In a standard common-priors setting, the optimal contract provides full insurance to the agent: the principal pays a fixed wage to the agent, asks him to reveal his information, and implements the efficient investment rule. When the agent overestimates the expected revenue of the project following investment, however, he is willing to "wager" on success against the (relatively pessimistic) principal, and hence bear some project risk in equilibrium. In addition, because what the principal considers to be the optimal investment rule is too conservative according to the agent's beliefs and the agent holds some stake in the choice of investment rule, he will accept a lower fixed payment in exchange for a more liberal investment rule. This can be interpreted as giving more control to the agent. It is somewhat counterintuitive that, the principal will surrender more control to an agent with whom she disagrees more sharply.

Keywords: overconfidence, heterogeneous beliefs, expert advice.

JEL Classification: A12, D81, D82.

---

*School of Economics and Management, University of Aarhus. I am particularly grateful to Benjamin Hermalin, Terrance Odean, Matthew Rabin, and especially Botond Kőszegi for very helpful discussions. All errors and omissions are my own.

1 Introduction

An important function of a firm's manager is to guide investment decisions; an experienced manager is likely to be better informed than the firm's owner regarding the profitability of potential investment projects. If the manager's attitudes towards risk differ from the owner's, however, full delegation of investment decisions may not be in the owner's best interest.

Consider, in particular, a case in which a risk-averse manager is hired by a risk-neutral owner. Imagine that the manager can share his private information regarding the project's expected profits with the owner. After receiving the manager's report, the owner then decides whether or not to invest in the project. In this principal-agent setting, under the assumption that manager and owner agree about the distribution of profits conditional on the manager's information, the manager's risk aversion does not pose an incentive problem. The owner can offer full insurance to the manager (i.e. offer him a fixed payment), and simply ask him to act with the owner's best interest in mind.

The common-priors assumption implicit in this logic is a strong one, however—in particular when the manager has some stake in the outcome of the investment decision following his recommendation. Studies in psychology suggest that individuals tend to be overly optimistic when evaluating their own skills and ability, and the likelihood of favorable outcomes. This bias seems to be stronger when evaluating outcomes that affect one's personal well-being. See Taylor and Brown (1988) for a discussion of the findings in psychology research regarding this bias. In our example, the manager's utility can be indirectly affected by the outcome of the investment decision (e.g. through ego, reputation). If the manager overestimates the returns to investment relative to the owner's beliefs, and his payment can be made contingent on the project's outcome following investment, then he will not receive full insurance in equilibrium.[1] In this case, the manager will in fact "wager" to some extent against the owner, so that he receives outcome-contingent pay in equilibrium. Because full insurance is no longer the equilibrium contract, the manager's well being is affected by the project's outcome following investment, and therefore by the owner's decision of whether or not to invest following the manager's recommendation. Disagreement about the expected returns to investment implies, in most instances, disagreement about the optimal course of action. As a result, the manager may be willing to "pay" for some degree of control over the investment decision. In contrast to the common-priors setting, in the presence of heterogeneous beliefs the investment rule implemented in equilibrium will almost surely not be the efficient investment rule according to the owner's beliefs. I briefly discuss the implications of this result in an insider/outsider (e.g. manager/consultant) setting in the concluding section.

[1] Gervais, Heaton and Odean (2003), and Van den Steen (2005) also study the effects of overconfidence/optimism in investment-decision settings. Some of the results in this paper are consistent with their findings.

Section 2 introduces a simple model that allows us to explore the effects of agent overconfidence in an investment-decision setting. I follow the basic structure employed by Holmstrom and Ricart i Costa (1986). The most notable characteristic of the model is that the agent privately observes some signal that is correlated with the expected profits of a potential investment project and can share that information with his employer (the principal). Principals compete to hire the agent by making simultaneous contract offers, which results in zero expected profits for the principal in equilibrium. A contract offer specifies both the payment scheme and the investment rule that will be followed (whether or not investment is undertaken following a given report by the agent). The expected revenue of the project depends stochastically on both the signal, which is private information to the agent, and on the agent's ability, which is unknown to both parties. The agent can choose whether to report the observed signal to the principal, but cannot misrepresent information.[2] Then, the principal invests according to the agreed-upon investment rule. If the agent chooses not to report the signal, no investment will be made: the agent has "veto power" after observing the signal, but the principal does not.[3] Once outcomes are realized, payments are distributed according to the terms of the contract, and the agency relationship ends.

Section 3 isolates the issue of the equilibrium payment scheme from the issue of the equilibrium investment rule (which we incorporate in Section 4 for tractability). In a common-priors setting, the equilibrium payment scheme is very simple: the principal offers full insurance to the agent, so that the agent is paid his expected productivity, and the agent always reveals his private information in equilibrium. If, however, the agent overestimates his ability, and thus the expected revenue of the project following investment relative to the principal's beliefs, then the equilibrium payment scheme will be outcome-dependent. We allow principal and agent to hold heterogeneous beliefs, in particular, about the distribution of revenue conditional on each signal realization. In equilibrium, the agent receives a higher payment following those realizations of revenue that he deems more likely than the principal does, and in return accepts a lower payment following those realizations that he believes are relatively unlikely. Even though the agent is risk averse, some degree of "wagering" of this kind will always be observed in equilibrium. A corollary of this result is that a more-overconfident agent will tend to receive a higher payment in the event that no investment is

[2] This is an important assumption in Holmstrom and Ricart i Costa (1986), that I maintain for tractability. The possibility of information misrepresentation is an interesting and relevant problem. Stocken (2000), for example, studies the credibility of information disclosure in a repeated cheap-talk game.

[3] Note that a risk-averse agent might prefer that no investment be made if his remuneration depends on project outcome. Allowing the agent to veto an investment project forces the principal to internalize this issue when designing the contract. If there are enough projects with large negative expected profits, it will be optimal for the principal to not invest following "no report" by the agent.

made, as long as it does not destroy the agent's incentives to fully reveal his private information. An overconfident agent overestimates his expected payment and his expected utility following the decision to invest; because he disagrees with the principal about his (unknown) ability but not about the ex-ante distribution of the signal (which both principal and agent know *ex-post* in equilibrium), he receives as much insurance as possible in terms of signal variability. Wishing to smooth his consumption across "invest" and "do not invest" events, a more overconfident agent will prefer a higher payment conditional on no investment, consistent with his expectation of higher payment if investment is undertaken.

Section 4 includes the investment rule as part of the optimization problem implicit in the equilibrium contract offer. Given that the agent receives full insurance in a common-priors setting, he is indifferent as to the implemented investment rule in that case. The principal thus implements the rule that she considers optimal, which, given her risk neutrality, is in fact the efficient rule. As a consequence of heterogeneous beliefs, the agent's payment is contingent on the outcome of investment whenever it is undertaken. Two sources of distortion (relative to the efficient rule according to the principal's beliefs) affect the equilibrium investment rule. First, the principal's objective function will almost surely no longer be aligned with efficiency. For the risk-averse agent to be willing to reveal his private information when the investment rule dictates investment (i.e. for him not to veto the project), it must be the case that he evaluates his *expected payment* conditional on investment as higher than what he would receive if no investment was made—which is a riskless payment. If the agent's expected payment conditional on investment is higher (or lower) according to the principal's beliefs as well, then the principal is relinquishing some of the marginal benefits of investment to the agent. As a consequence, the principal's objective function will, almost surely, not be aligned with efficiency. Second, if the agent believes that his *expected utility* will be higher conditional on investment than when no investment is made, he will be willing to forego some of his remuneration in exchange for a more liberal investment rule. The agent is "willing to pay for control" in that case, and the implemented investment rule will in fact be skewed away from what the principal would optimally choose if her decision was unilateral and non-contractible, and towards what the agent believes to be the optimal investment rule.

Section 5 provides an example of the equilibrium contract when the agent has an exponential utility function, and the productivity shocks are normally distributed. This purposefully simple example further illustrates the intuition behind the results of the model.

Section 6 concludes, and discusses a testable implication of the presence of overconfidence in an expert-advice setting, given that firms tend to rely on both outsiders' and insiders' advice to guide their investment decisions.

2 Model Setup

Consider a simple one-period expert-advice model, in which an agent gives information to a principal, who in turn includes the agent's "expert advice" in her decision-making process. We will follow the notation and basic model structure used by Holmstrom and Ricart i Costa (1986), applied to a one-period setting and allowing for the agent's payment to be contingent on same-period outcome.

The principal has the possibility of investing in a given project. At the beginning of the agency relationship, the agent observes a signal s that is correlated with the project's profitability. The profitability of the project also depends on the agent's ability. The project's revenue y, conditional on a decision to undertake investment, can be written as

$$y = s + \varepsilon,$$

, where s is privately observed by the agent and has cumulative distribution function N, and the distribution of ε depends on the agent's ability—which is unknown to both parties. Under the assumption of common prior beliefs (which we will relax below), principal and agent share the belief that ε has cumulative distribution function F with mean μ.[4] The agent can decide whether or not to report the true value of s (which can be interpreted as whether or not to "recommend" investment), but cannot misrepresent information in the report. The principal can always guarantee zero revenue by not investing. The principal is risk neutral, and his objective is thus to maximize expected profits:

$$\max \mathbb{E}\left[y - w\right]$$

where w stands for the agent's "wage"—a money transfer to the agent. The agent is risk averse, and his objective is to maximize his expected utility of consumption: $\max \mathbb{E}\left[u\left(c\right)\right]$. Assume that the agent's only income is his wage, so that the agent's objective can be expressed as

$$\max \mathbb{E}\left[u\left(w\right)\right].$$

Assume that several principals have access to this type of investment projects, and compete to contract with the agent.[5] The timing of the model is as follows: first, principals make simultaneous

[4] Holmstrom and Ricart i Costa (1986) assume that the agent can be either talented or not, and the prior probability of the agent being talented is p. In their model, $\varepsilon \sim G$ with mean μ_G if the agent is talented (G for "good"), and $\varepsilon \sim B$ with mean μ_B if the agent is not talented, where $\mu_G > \mu_B$. This specific form is one way to parametrize agent ability, particularly useful in a career-concerns setting; the generalization that $\varepsilon \sim F$ with mean μ allows for a more intuitive presentation of the results of this paper.

[5] The results in a setting where one principal makes a take-it-or-leave-it offer to an agent are nearly identical. See de la Rosa (2006a) for an extensive explanation. In short, these are dual problems, with the main difference being who retains the expected gains from trade.

contract offers to the agent. A contract offer consists of an investment rule and a payment schedule to the agent. The agent then chooses which offer (if any) to accept. When and if the agent chooses to accept a contract offer, he privately observes signal s, and chooses whether or not to report it to the principal. If the agent reports the signal, investment is undertaken or not according to the agreed-upon investment rule. If investment is undertaken, ε is then realized and correctly inferred by both parties (given that project profitability is publicly observed). There is no investment if the agent chooses not to report s. Payoffs are then distributed according to the agreed-upon payment schedule, and the agency relationship ends. If the agent chooses not to accept any contract, no project will be undertaken, and the players receive payoffs according to some outside option. I assume that the agent's outside option is low enough so that he always accepts some contract offer in equilibrium.

The equilibrium contract solves

$$\max \mathbb{E}\left[u\left(w\right)\right]$$
$$\text{s.t. } \mathbb{E}\left[y-w\right] \geq 0.$$

Note that the non-negative expected profits restriction must bind in equilibrium: if a principal made positive expected profits in equilibrium, another principal could outbid the "equilibrium" contract by offering δ more to the agent in every state of the world, attract the agent, and earn strictly positive expected profits.

Let $\alpha\left(s\right) \in \{0,1\}$ denote the investment rule implemented following announcement s from the agent. As we will see below, in an common-priors setting, the information asymmetry does not pose an incentive problem, so that the efficient investment rule $\alpha^*\left(s\right)$ (which is optimal from the principal's perspective) is implemented in equilibrium. Investment is efficient if and only if the project yields positive expected profits conditional on the realization of the signal. That is, $\alpha^*\left(s\right)=1$ if and only if

$$\mathbb{E}\left[y \mid s\right] \geq 0,$$

which implies

$$s+\int \varepsilon \, dF\left(\varepsilon\right) \geq 0,$$

and thus

$$s+\mu \geq 0.$$

The efficient investment rule can also be expressed as a "hurdle rate" $s^*\left(\mu\right)$, so that the project is undertaken if and only if $s \geq s^*\left(\mu\right)$. $s^*\left(\mu\right)$ satisfies

$$s^*\left(\mu\right)+\mu=0.$$

It follows that the equilibrium contract solves

$$\max_{s^*(\mu),w} \mathbb{E}\left[u\left(w\right)\right]$$
$$\text{s.t. } \mathbb{E}\left[w\right] = \left(\mathbb{E}\left[s \mid s \geq s^*\left(\mu\right)\right] + \mu\right) \cdot \Pr\left(s \geq s^*\left(\mu\right)\right).$$

When principal and agent agree on their beliefs regarding the distribution of ε, a fixed-wage contract is the equilibrium contract. The agent receives the full expected productivity of his private information, and, given that he is fully insured, there is never an incentive for him to withhold information (i.e. to "veto" an investment decision). In fact, failing to report s over a non-measure zero subset of signals when the expected profits from investment conditional on s are strictly positive would reduce the agent's payment in equilibrium. In what follows, I assume that an agent who is indifferent between reporting and not reporting will choose to report.

We will subsequently allow for the possibility that principal and agent hold heterogeneous beliefs about the distribution of ε. In particular, consider the case in which the agent overestimates his expected productivity. Let a tilde denote the agent's beliefs: the agent holds the prior belief that ε has a cumulative distribution function $\tilde{F}$ with mean $\tilde{\mu}$. We will say that the agent is *overconfident* in this setting if

$$\tilde{\mu} > \mu,$$

which means that an overconfident agent overestimates the expected return of an investment project, conditional on any given signal, relative to the principals' beliefs.[6]

3 Expert Advice

To isolate the consequences of heterogeneous beliefs on the contracted payment schedule from the issue of control (the negotiated investment rule), assume for now that the principal always implements the investment rule which is efficient according to her beliefs. The equilibrium investment

[6] There are other dimensions over which the beliefs of principal and agent could differ. One could think about an agent who is "overconfident about the precision of ε." The effects of this type of overconfidence will become apparent when we characterize the equilibrium contract allowing for agent and principal to evaluate the likelihood of a given realization of ε according to different distribution functions in Section 3. The agent could also be overconfident about the distribution of s. In this case the agent's payment would depend on s in equilibrium. He would be willing to trade a higher payment under those realizations of s that he believes to be more likely than the principal does, in return for lower payment under those realizations of s that he believes to be relatively less likely. Absent heterogeneous beliefs about the distribution of ε, full insurance with respect to the realization of ε will still be optimal. I abstract from the possibility of disagreement about the distribution of s because it would necessarily distort the principal's objective function away from efficiency when choosing $\alpha\left(s\right)$, adding one more source of distortion to the two we study in Section 4.

rule $\alpha^*(s)$ and respective hurdle rate $s^*(\mu)$ thus remain unchanged from the common-priors setting discussed in Section 2 above. Note that this investment rule seems too conservative from the agent's point of view: for some realizations of s, the principal chooses not to invest even though the project is profitable according to the agent's beliefs. If the agent's remuneration depends on the profitability of projects that are undertaken in equilibrium, it is possible that he would be willing to "pay" the principal in exchange for some decision power. We will address this question in Section 4 next.

Recall that the equilibrium contract solves

$$\max \tilde{\mathbb{E}}\left[u\left(w\right)\right]$$
$$\text{s.t. } \mathbb{E}\left[y-w\right]\geq 0,$$

where the tilde over the first expectations operator points to the fact that the agent is interested in maximizing his perceived expected utility, which depends on his beliefs rather than the principal's. As noted before, the non-negative expected profits constraint binds, so we can write

$$\mathbb{E}\left[w\right]=\left(\mathbb{E}\left[s\mid s\geq s^*\left(\mu\right)\right]+\mu\right)\cdot\Pr\left(s\geq s^*\left(\mu\right)\right).$$

Recall that the agent can choose not to report his private information to the principal. In the case of common priors, full insurance for the agent implies that the agent will always reveal his information in equilibrium. If the agent is to fully reveal his information, it must be the case that

$$\tilde{\mathbb{E}}\left[u\left(w\right)\right]\mid_{\alpha(s)=1}\geq\tilde{\mathbb{E}}\left[u\left(w\right)\right]\mid_{\alpha(s)=0}$$

whenever the agent observes an s such that $\alpha(s)=1$. This follows from the assumption in Holmstrom and Ricart i Costa (1986) that no investment is made following no report (i.e. no recommendation to invest) from the agent. As noted by these authors, the principal will in fact decide not to invest following no report from the agent in equilibrium if there are sufficiently "bad" projects. The contract must be such that the agent does not veto a project whenever the equilibrium investment rule dictates investment. This *full-revelation constraint* may or may not bind in equilibrium.

As the following remark notes, in a heterogeneous-beliefs setting the agent will bear some risk following a decision to invest.

Remark 1 *If $\tilde{\mu}>\mu$, then full insurance to the agent is not an equilibrium.*

A formal proof is in the appendix. The intuition is that an overconfident agent is willing to wager against the principal on investment outcome. Consider two potential equilibrium contracts:

one that offers full insurance to the agent, and an alternative contract under which the agent receives higher payment following a high realization of revenue (e.g. profit-sharing, or an option on a share of the profits). Recall that equilibrium requires that the principal receive zero profits in expectation. Therefore, both contracts must yield the same expected profits—and the same expected payment to the agent—*according to the principal's beliefs*. Given that an overconfident agent overestimates his expected productivity, he believes that the risky contract gives him a higher expected payment. This constitutes a first-order gain in terms of his perceived expected utility, while the loss that the agent incurs from bearing risk is only of second order when evaluated at full insurance. It follows that an overconfident agent, even though risk averse, will always be exposed to some risk in equilibrium.

In general, a contract offer will be characterized by an investment rule $\alpha(s)$, and payment to the agent $w(s, \alpha(s), \varepsilon)$. We can restrict attention, without loss of generality, to a wage function of the form

$$w = \beta + \sigma(\varepsilon),$$

noting that ε will be realized only if the project is undertaken (if $\alpha(s) = 1$). Given that ε cannot be observed if no investment is made, $w(s, \alpha(s), \varepsilon) = w(s', \alpha(s), \varepsilon)$ whenever $\alpha(s) = \alpha(s') = 0$, so we can write $w = \beta$ whenever no investment is made. It is also true that $w(s, \alpha(s), \varepsilon) = w(s', \alpha(s), \varepsilon)$ for all ε (almost everywhere) whenever $\alpha(s) = \alpha(s') = 1$. This follows from the agent's risk aversion, independence between s and ε, and the fact that principal and agent agree about the distribution of s. Intuitively, it is costless for the principal to insure the agent against variability in s, and given that principal and agent disagree about the distribution of profits only insofar as they disagree about the distribution of ε, it is optimal for the risk-neutral principal to absorb all the risk from variability in s following investment.[7]

Taking this into account, we can rewrite the principal's zero-expected-profits condition as

$$\begin{aligned}\beta &= \int_{s^*(\mu)}^{\infty} \int (s + \varepsilon - \sigma(\varepsilon))\, dF(\varepsilon)\, dN(s) \\ &= \left(\mathbb{E}[s \mid s \geq s^*(\mu)] + \mu - \int \sigma(\varepsilon)\, dF(\varepsilon)\right) \cdot \Pr(s \geq s^*(\mu)).\end{aligned}$$

In equilibrium, $\sigma(\cdot)$ is Pareto optimal from the perspective of securities trading under uncer-

[7] To see why, consider the following example. For simplicity, imagine that s and ε are discretely distributed. Consider two realizations of the signal, s' and s'', such that $\Pr(s'), \Pr(s'') > 0$. Assume that $w(s'', 1, \varepsilon') > w(s', 1, \varepsilon')$ for some realization ε' of ε such that $\Pr(\varepsilon') > 0$ according to both the principal's and the agent's beliefs. If the principal offers more insurance to the agent in an actuarially fair manner (reducing $w(s'', 1, \varepsilon')$ and increasing $w(s', 1, \varepsilon')$ marginally so that expected profits to the principal remain constant), the risk-averse agent's perceived expected utility will increase, independent of what principal and agent believe to be the actual likelihood of ε'.