

Lindbeck, Assar; Weibull, Jörgen W.

Working Paper

Pay schemes, bargaining, and competition for talent

IFN Working Paper, No. 1100

Provided in Cooperation with:

Research Institute of Industrial Economics (IFN), Stockholm

Suggested Citation: Lindbeck, Assar; Weibull, Jörgen W. (2015) : Pay schemes, bargaining, and competition for talent, IFN Working Paper, No. 1100, Research Institute of Industrial Economics (IFN), Stockholm

This Version is available at:

<https://hdl.handle.net/10419/175045>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

IFN Working Paper No. 1100, 2015

Pay Schemes, Bargaining, and Competition for Talent

Assar Lindbeck and Jörgen Weibull

Pay schemes, bargaining, and competition for talent

ASSAR LINDBECK* AND JÖRGEN WEIBULL^{†‡}

December 21, 2015.

ABSTRACT. The paper provides a framework for analysis of remuneration to agents whose task is to make well-informed decisions on behalf of a principal, with managers in large corporations as the most prominent example. The principal and agent initially bargain over the pay scheme to the latter. The bargaining outcome depends both on competition for agents and on the relative bargaining power of the two parties, given their outside options, thus allowing for the possibility that the agent may be the current CEO who may have considerable power. Having signed a contract, the agent chooses how much effort to make to acquire information about the project at hand. This information is private and the agent uses it in his subsequent decision whether or not to invest in a given project. In *model A* the agent's effort to acquire information is exogenous, whereas in *model E* it is endogenous. Model A lends no support for other payment schemes than flat salaries is weak. Model E contains a double moral hazard problem; how much information to acquire and what investment decision to make. As a consequence, the equilibrium contracts in model E involve both bonuses and penalties. We identify lower and upper bounds on these, and study how the bonus and bonus rate depend on competition and bargaining power. We also analyze the nature of contracts when the agent is overconfident.

Keywords: Principal-agent, investment, endogenous uncertainty, contract, bonus, penalty.

JEL codes: D01, D82, D86, G11, G23, G30.

*Institute for International Economic Studies, Stockholm University, and Research Institute of Industrial Economics, Stockholm.

[†]Stockholm School of Economics, KTH Royal Institute of Technology, and Institute for Advanced Study in Toulouse.

[‡]The authors thank Bo Becker, Sven-Olof Fridolfson, Jungsuk Han, Augustin Landier, Paul Milgrom, Francois Salanié, Francesco Sangiorgi, Per Strömberg, and Thomas Tangerås for helpful discussions. We thank Atahan Afsar for research assistance and Kristaps Dzonsons for computer simulations. Assar Lindbeck thanks the Jan Wallander and Tom Hedelius Foundation, and the Marianne and Marcus Wallenberg Foundation for financial support. Jörgen Weibull thanks the Knut and Alice Wallenberg Research Foundation for financial support.

1. INTRODUCTION

The heated policy discussion in recent years of the high, and during certain periods rapidly increasing, pay to CEOs in large companies suggests that it is important to clarify the mechanisms behind this phenomenon. There are two dominating explanations in the literature. One is the stiff, and possibly gradually stiffer, competition for talents in open labor markets. The other explanation is the strong bargaining power of incumbent managers over their own pay scheme in negotiations with company boards and shareholders, resulting in top managers in large corporations being rewarded out of proportion to their contributions to the firms where they work. The empirical basis for each one of these hypotheses consists mainly of US data.

Advocates of the competition hypothesis point out that high, and during certain periods rapidly rising, earnings is not unique to CEOs in large and hierarchically organized corporations. Talented individuals in other professions are asserted to enjoy similar earnings, which suggest common explanations (Kaplan and Rauh, 2010; Bakhiya et al. 2012). For instance, the average individual in the highest 0.1 percent of the general distribution of earnings has enjoyed about equally high earnings as CEOs in large corporations at least from the mid-1930s and until about the mid-1980s (Kaplan 2013). This group of individuals includes, for instance, managers in private firms, hedge funds, private-equity firms and venture capital investors, as well as particularly successful lawyers, sportsmen, artists etc. Supporters of the competitive hypothesis have also emphasized that the compensation for CEOs in large corporations, in fact, has broadly developed in proportion to the development of the market value of large corporations during the last two decades; see, for instance, Kaplan (2013 Figure 16). Supporters of the competition hypothesis also point out that the firing of managers is related to poor performance of the stock values of the firms in which they work (Kaplan and Minton, 2012; Jenter and Lewellen, 2014). Their interpretation of both these observations is that managers of large corporations tend to be rewarded in relation to the performance of the firms where they work. They also point out that the pay policy for executives was approved by the boards in over 98 percent of S&P 500 and Russel 3000 companies in 2011.

By contrast, the supporters of the bargaining-power hypothesis instead point out that the earnings of CEOs in large corporations have increased much more rapidly than the market value of large corporations from the 1960s, although the developments have been quite proportional from the late 1990s (Mishel and Davis, 2014, Figure A). Advocates of the bargaining-power hypothesis also refer to the development of earnings during the last two decades. They report that the earnings of CEOs in large corporations, relative to average earnings for the 0.1 highest group in this period have been considerably above the previous long-term trend (Mishel and Davis, 2014; Bebchuck and Fried, 2003, 2004). Adherents to the bargaining-power hypothe-

sis argue that this situation reflects the ability of CEOs to exploit superior information when bargaining about their own pay. This superiority holds both when assessing the actual contribution of the CEO and when trying to understand the functioning of the highly incentive-oriented remuneration systems that emerged in the 1980s and 1990s – mainly in the form of bonuses, stocks and options. Other observers argue that some CEOs even use their information advantage to deliberately mislead boards and share-owners about the actual functioning of complex remuneration schemes — the so called “skimming hypothesis”, see, for instance, Bertrand and Mullainathan (2001).

In a comprehensive study of executive compensation, Murphy (2012) reaches the conclusion that both the competitive hypothesis and the bargaining-power hypothesis find support in available data. He suggests, however, that political factors also have played an important part, in particular policy actions such as disclosure requirements, tax policies, accounting rules and the general political climate. Our conclusion from the empirical literature in this field is that both the competition for talents in open markets and the bargaining power of top managers within firms are important factors behind the remuneration of CEOs. These two factors are also likely to interact with each other, since both the outside options of CEOs, due to competition for talent, and their administrative power within their firms, do influence the bargaining outcomes between company boards or owners and their CEOs. Indeed, policies such as those emphasized by Kevin Murphy, are likely to influence the remuneration of CEOs via both mechanisms.

The main ambition behind this paper is to provide a canonical analytical framework, or “work horse”, that can easily be used and further developed to apply to a wide range of important issues concerning the remuneration of agents whose tasks are to make well-informed decisions on behalf of some principal. We obtain a range of new inferences by way of combining elements from the principal-agent literature with elements from other literatures, in particular concerning production and distribution of information, competition for talent, strength of bargaining, over confidence and optimism etc. The focus of our study is on the interaction between a *principal*, who has money, and an *agent*, who has talent. The principal may for example be the owner of a firm and the agent its CEO or a candidate for this job. The principal faces a decision under uncertainty, or a sequence of such decisions, such as whether or not to undertake a risky project—say, make an investment or purchase an asset—or a sequence of such risky projects. The return from a project is a random variable and the principal and the agent have the same prior beliefs about its probability distribution. The two parties first enter a negotiation about the terms of a potential contract between them. If they fail to agree, each party picks up his or her outside option. If they **agree** on a contract, then this contract delegates the decision power to the agent whether or not to invest in the risky project at hand. The contract also

specifies a payment to the agent under each of three possible (verifiable) scenarios: investment in the project and “success”, investment and “failure”, and no investment. We require contracts to meet the “limited liability” constraint that the net pay from the principal to the agent is never negative. Having signed such a contract, the agent decides how much, if any, effort to make in order to acquire more information about the project at hand. The more talented the agent, the less effort he needs to spend in order to acquire a given level of precision in his information. The agent’s effort and acquired information are both modelled as variables on a continuum scale and they are his private information.¹ By contrast, the *ex post* return from the project are verifiable in court. The terms of the contract thus influence the agent in two distinct ways: it motivates the agent to acquire information about the project, and it also guides the agent’s investment decision, once his information has been obtained. As will be seen, there is a tension between these two goals; a contract may be well designed for the first purpose but may at the same time induce the agent to be too “trigger happy”, that is, to invest even when his information about the project is not sufficiently favorable for the principal. This tension exists even when both parties are risk neutral and fully rational.

More specifically, we apply the generalized Nash-bargaining solution (Nash, 1950; Roth, 1979) to the contract negotiation between the principal and agent. Once a contract has been signed by both parties, the agent is assumed to act sequentially rationally in his own self-interest, given the terms of the contract. We call a contract *feasible* if it results in expected profits to the principal above her outside option and in expected utility to the agent above his outside option. We call a feasible contract *undominated* if there exists no other feasible contract that would be at least as good for both parties and strictly better for at least one party. We require from an agreed-upon contract that it be immune against renegotiation in precisely the sense of being undominated. The generalized Nash bargaining solution selects an undominated contract from among the feasible contracts, and this selection depends both on the parties’ outside options and on their individual bargaining power. The parties’ outside options represent the intensity of competition for talent in the job market. The stiffer competition there is, the better is the agent’s outside option and the worse is that of the principal. We analyze the full spectrum of bargaining power relations between the two parties, ranging from a principal who can make a take-it-or leave it offer to the opposite extreme case of an agent who can make a take-it-or leave it offer to the principal. The intermediate case of equal bargaining power corresponds to Nash’s (1950) classical bargaining solution. The president of a company board may have a lot of bargaining power, given both parties’ outside

¹In another research project, one of the authors analyze such endogenous information acquisition in a model of credit, banking and securitization, see Axelson and Weibull (2014).

options, over an outside candidate for the CEO position. The opposite may hold when the agent is the company's current CEO. A sitting CEO usually has valuable connections, access to administrative resources, and inside information, which gives him an advantage in discussions with the company board and provides a potential threat in case he would leave. There may also be considerable reputational cost for a company of firing its CEO.

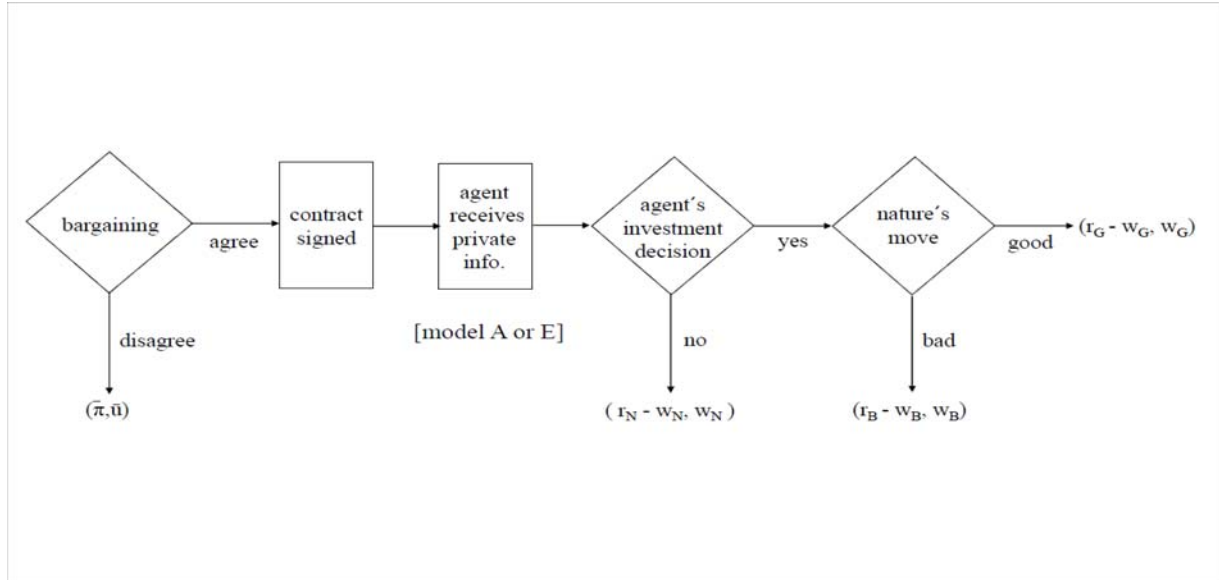
We first develop a base-line model, *model A*, in which the precision of the agent's information is fixed. The actual information obtained by the agent is still a random variable, though, so the agent has to choose an "investment strategy", a rule for when to invest, conditional upon his (future and random) information. The more talented the agent, the more precise will his information about the project be. The agent's investment strategy will depend on the terms of the contract. The properties of model version A are shown to be consistent with the principal-agent literature, and we use this simple model version as a bench-mark for our analysis of the double moral hazard and endogenous informational asymmetry that arises when the agent's information-gathering and investment decision both are endogenous.

This brings us to the extension of model A that we call *model E*. This model is, to the best of our knowledge novel, and it delivers new insights. Against the background of model A, the moral hazard problem in connection with endogenous information acquisition stands out clearly. In model E, the contract not only influences the agent's investment decision but also how well-informed the agent will choose to be when making his investment decision. Model E is thus a model of endogenous asymmetric information. As shown in Weibull, Mattsson and Voorneveld (2007), endogenous uncertainty in decision-making may introduce a major non-convexity in the decision problem.² In the present model, such non-convexity takes the following form. The agent will either make no effort at all or else he will make a significant effort, which depends continuously on the terms of the contract. He will never choose intermediate effort levels, given his talent and the terms of the contract. The reason for such a discontinuous "jump" in behavior is that a range of intermediate effort levels will not sufficiently enhance (diminish) the probability for the bonus (penalty) in the contract at hand; it is better to either make no effort at all or a relatively big effort. The agent's expected utility from remuneration, according to a given contract, is increasing with the precision of his information. However, the marginal expected utility from this remuneration is non-monotonic. It is always positive, but it tends to zero both when the agent has very imprecise information and when he has very precise information. In each case, the marginal benefit of more precision is largest when his signal precision is intermediate; then it matters most for his subsequent investment decision. He should

²Non-convexities in relation to information were first noted by Radner and Stiglitz (1984), see also Chade and Schlee (2002).

thus either make no effort at all or an effort high enough to affect his investment decision. This non-convexity suggests that even minute changes in the environment or in the contract may result in drastic deterioration of investment decisions.

Models A and E are summarized in the flow chart below. First, the two parties bargain over a potential contract between them. If they do not reach an agreement, they receive their outside options, the reservation profit, $\bar{\pi}$, to the principal and the reservation utility, \bar{u} , to the agent. Otherwise they sign a contract, denoted $w = (w_B, w_N, w_G)$, which specifies the agent's remuneration in each eventuality. The net return from investment in the good (bad) state of nature is denoted r_G (r_B), and r_N is the net return when not investing in the project at hand.³ Models A and E differ only in the box right after a contract has been signed. In model A, the agent there receives private information of exogenous precision. In model E, the agent there first decides how much information-gathering effort to make and then receives his private information with **a** precision determined by his effort. Once the agent has received his private information, in either model, he decides whether or not to invest in the project. If he invests, the project will succeed if the “state of nature” is good and fail if the state of nature is bad. Finally, the two parties receive their payoffs according to the signed contract. While model A thus includes only one moral hazard problem, in the sense of hidden action, namely the agent's investment strategy (for what signal values to invest), model E includes an additional moral hazard problem, namely, the agent's (for him costly) choice of signal precision, the fruits of which he has to share with the principal (unless the agent has all the bargaining power).



³This return may, but need not, be identical with the principal's outside option during the initial contract negotiation.

Our main results are as follows, with the term *bonus* referring to the additional pay after a successful investment as compared with the remuneration when no investment was made ($w_G - w_N$), and the term *penalty* referring to the reduction in the agent's remuneration after a failed investment, again as compared with his payment when no investment was made ($w_N - w_B$).⁴ In model A, a whole range of positive bonuses and penalties are compatible with equilibrium when the agent is risk neutral, granted the bonus and penalty are balanced in a precise way. The ratio between the bonus and penalty has to match the ratio between the project's "upside" and "downside" (defined below in a precise way). However, also "flat pay" contracts, i.e. contracts with no bonus and no penalty, are possible in equilibrium. The agent will then be indifferent between investing or not, irrespective of his private information, and may thus invest in a way that is in the best interest of the principal, given the agent's private information.⁵ However, if the agent is (even the slightest) risk averse in model A, only a flat salary is consistent with equilibrium. The principal, being risk neutral, will in equilibrium necessarily absorb all risk—in line with standard results. Model A thus lends no support to the argument that competition for talent requires bonuses. Competition may require high pay, but not necessarily in the form of bonuses and/or penalties. In this basic model version (with no tax distortions or other real-life complications), the only rationale for bonuses and penalties are to screen for more able agents when the principal has incomplete information about job candidates' talents.

By contrast, in model E, with its double moral-hazard problem, all equilibrium contracts include both a bonus and a penalty. Moreover, there are upper and lower bounds on these. If the bonus and/or penalty are set too low, the agent will make no information-gathering effort at all, and will thus be of no use for the principal, who would be better off without the agent. Indeed, the agent's switch from "diligent" information gathering (bounded away from zero) to "shirking" (zero effort) may be discontinuous. Even the slightest reduction of his bonus may induce the agent to make no information-gathering effort at all. Moreover, the limited-liability constraint (the lower bound on pay to the agent), is binding in all equilibrium contracts in model E; the agent will be paid nothing if investment fails. Typically, the equilibrium contracts in model E involves relatively high bonuses, but there is a uniform upper bound on them; the "bonus rate", defined as the ratio between the remuneration after a successful investment and after no investment cannot exceed the ex ante odds

⁴An alternative terminology—that avoids the term "penalty"—would be to call the extra pay after no investment, $w_N - w_B$, and after a succesful investment, $w_G - w_B$ (both as compared with the pay after a failed investment) "bonuses." However, we find that a less natural terminology.

⁵Hence, if the agent would have the slightest loyalty with the principal (which we do not here assume), he would strictly prefer to behave in such a way under such a contract.

for the bad state of nature.⁶ Hence, the more likely the good state is, ex ante, the lower is the upper bound on the bonus rate. In particular, under ex ante Laplacian uncertainty, the bonus rate never exceeds 100% (in line with a recommendation by the European Parliament).⁷

Another feature of model E, which makes it qualitatively different from standard principal-agent models is that even if the principal has all the bargaining power, and hence can give a take-it-or-leave-it offer to the agent, the offered contract need not force the agent down to his participation constraint. The reason is that such contracts may induce the agent, if hired, to make no effort at all to gather information. A certain minimal threshold is required to inspire the agent to work. By way of numerical simulations, we also show how large the equilibrium bonus rate is and how this depend on the primitives of the model. In particular, the effect of increased competition for talent is qualitatively the same as the effect of an agent's increased bargaining power. Moreover, while these factors influence the size of the bonus in the expected direction, the bonus rate is virtually unaffected. We also discuss conventional stock- or option-based contracts, as well as such behavioral biases as overconfidence and optimism. In particular, while stock-based contracts may arise in equilibrium, option-based contracts cannot. Such contracts do not have a down-side and would thus induce agents to invest also when available information is unfavorable. In order for an option-based contract to be compatible with equilibrium between rational and risk-neutral parties, the contract also has to contain some penalty after failed investments.

The organization of the paper is as follows. Section 2 provides notation, basic assumptions and a specification of model A, which is analyzed in Section 3. In section 4 we develop model E, and analyze this model in Section 5. Section 6 discusses stock-based and options-based incentive contracts. Section 7 analyzes such behavioral biases as overconfidence and optimism. Section 8 discusses related literature and Section 9 concludes. Mathematical proofs are provided in an appendix at the end of the paper.

2. MODEL A

A principal has the opportunity to undertake a risky project, which we represent as a random variable X that specifies the net return from the project. The principal's outside option has expected net return $\bar{\pi} \geq 0$. There are two states of nature, one

⁶The limited-liability constraint imposes a lower bound of 100% on the likewise defined "penalty rate".

⁷The European Parliament approved a proposal for capital-requirement regulation (CRD4) in April 2013, a proposal that also suggested a cap on bankers' bonuses, namely, that the variable pay should not exceed 100% of the fixed pay, with some exceptions subject to shareholder approval (*Financial Times*, February 28, 2013)

“good” and one “bad.” The random variable X takes the value $r_G > \bar{\pi}$ in the good state and $r_B < \bar{\pi}$ in the bad. The *ex ante* expected net return from the project is thus

$$\mathbb{E}(X) = \mu r_G + (1 - \mu) r_B, \quad (1)$$

where $\mu \in (0, 1)$ is the probability for the good state to obtain. This is known by the principal and the agent. We focus on the case when this expectation falls short of the principal’s outside option, $\mathbb{E}(X) < \bar{\pi}$. Hence, without expert advice, the principal will not invest in the project. However, there is also a potential agent, an expert the principal may hire for the purpose of making a well-informed investment decision. The two parties first negotiate a contract between them. If they reach an agreement, the agent will obtain relevant information about the project before giving advice or making the investment decision on behalf of the principal. Formally, the agent will receive a *private signal* about the project, a random variable

$$S = X + \varepsilon, \quad (2)$$

where the *noise* ε is statistically independent from X . The noise term ε has a positive and continuous probability density ϕ_t on \mathbb{R} , with cumulative distribution Φ_t . Here $t > 0$ is a parameter that is exogenous in the present model version, but that will become endogenous in model E (Sections 4 and 5). In both model versions, t will be referred to as the *precision* of the agent’s signal, and will later on be specified as the inverse of the variance of the noise term ε . In the present model version, we call t the agent’s *talent*; the more talented he is, the better informed he will be once hired by the principal. In model E, the parameter t will sometimes be called the agent’s *competence*; there t will depend both on his *innate ability* and on his self-selected *work effort*.

2.1. The agent’s information. We assume throughout this study that the monotone likelihood-ratio property holds at any given signal precision t . In other words, the higher the observed signal value, the higher is the conditional probability for the high return (r_G). This is a statistical property shared by many probability distributions and it holds if noise is normally distributed. Formally, our assumption is⁸

[MLRP] For any signal precision $t > 0$, the density ratio $\psi_t(s) = \phi_t(s - r_G) / \phi_t(s - r_B)$ is strictly increasing in the signal s .

Indeed, under this condition, the *posterior* probability for the good state, given any signal realization s ,

$$\Pr(X = r_G \mid S = s) = \frac{\mu \phi_t(s - r_G)}{\mu \phi_t(s - r_G) + (1 - \mu) \phi_t(s - r_B)}, \quad (3)$$

⁸For its original and general formulations, see Karlin and Rubin (1956) and Milgrom (1981).

is strictly increasing in the observed signal value s .

We will impose one more condition on the noise distribution, also met by the normal distribution, namely, that the density ϕ_t has "thin tails" in a precise sense:

[TT] *For any signal precision $t > 0$, the density ratio $\psi_t(s)$ tends to zero as $s \rightarrow -\infty$ and to plus infinity as $s \rightarrow +\infty$.*

Together with condition MLRP, condition TT simplifies the subsequent analysis. Unlike the MLRP condition, which is of fundamental importance for the analysis, TT is merely a condition for analytical convenience, without consequence for the qualitative features of the results, see Remark 1 below.

2.2. Contracts. A *contract* between the principal and the agent specifies a remuneration for the agent, to be made *ex post*, once the project's return has been realized. The remuneration may depend on whether or not the agent invested (or recommended investment), and, if so, on the project's realized net return. We thus treat the investment decision and the realized net return from investment as verifiable and legally enforceable in court.⁹ Under this assumption, any legally enforceable contract can be summarized as a vector $w = (w_B, w_N, w_G) \in \mathbb{R}^3$ that specifies a (positive or negative) remuneration w_B to the agent if he invests in the bad state of nature, a remuneration w_G if the agent invests in the good state of nature, and a remuneration w_N if the agent does not invest in the project at hand. We focus on contracts that are *monotonic* in the sense that they pay most after an investment in the good state and least after an investment in the bad state. Let W denote this class of contracts; $W = \{w \in \mathbb{R}^3 : w_B \leq w_N \leq w_G\}$.¹⁰ We call a contract $w \in W$ *strictly monotonic* if both inequalities are strict ($w_B < w_N < w_G$).

The agent is supposed to be selfish and maximize the expected value of his random utility $U(w, t)$, which takes the values $u(w_B)$, $u(w_N)$ and $u(w_G)$, respectively, in these three events. We assume that his Bernoulli function $u : \mathbb{R} \rightarrow \mathbb{R}$ is continuous, strictly increasing, concave and unbounded. The agent's outside option, in case he does not sign a contract with the principal, has expected utility $\bar{u} = u(\bar{w})$, where $\bar{w} > 0$ is his certainty-equivalent reservation wage.

Arguably, it is realistic in many situations to exclude contracts under which the remuneration to the agent is negative in some states of the world, if not for legal reasons perhaps because the agent does not have enough capital to cover a loss or may have means to hide or shelter his private capital from the principal by transferring

⁹By contrast, the would-be return if the investment is not made is taken to be non-verifiable in court and can thus not be written as a condition for remuneration in the contract.

¹⁰The assumption about monotonicity will not be binding in equilibrium but facilitates the discussion and analysis.

ownership of assets to others. We refer to this non-negativity constraint as *limited liability*. Such contracts thus constitute the subset W_0 of contracts $w \in W$ with $w_B \geq 0$.¹¹

Just as the agent's utility $U(w, t)$ is a random variable, so is the *profit* $\Pi(w, t)$ to the principal; it is $r_B - w_B$ if investment is made in the bad state of nature, $r_N - w_N$ if no investment is made and $r_G - w_G$ if investment is made in the good state. The *non-investment* return r_N to the principal is assumed to be no better than the principal's initial outside option but better than the *ex ante* expected return from the project at hand,

$$\mathbb{E}[X] < r_N \leq \bar{\pi}. \quad (4)$$

This assumption is arguably realistic, since the non-investment option would typically also be available at the outset of a contract negotiation (in which case $\bar{\pi} < r_N$ cannot hold). We will call $r = (r_B, r_N, r_G)$ the project's *return vector*. We note that this vector, technically speaking, is an example of a strictly monotonic contract. Under this particular contract it is as if the principal hands over all (positive and negative) returns from the project to the agent. This particular contract does not meet the limited-liability constraint if any one of its three components is negative. Typically, the net return in the bad state of nature, r_B , is negative, in which case $r \notin W_0$.

When contemplating whether or not to accept any contract on the table, both parties anticipate that if the contract is signed, the agent will subsequently obtain his private information and will (not) invest if his *ex post* expected utility, conditional upon the terms of the contract and his then available private information, is thereby (not) maximized. Every strictly monotonic contract w has a well-defined associated *ex ante* expected utility $\mathbb{E}[U(w, t)]$ to the agent and *ex ante* expected profit $\mathbb{E}[\Pi(w, t)]$ to the principal (to be determined in the next section). By contrast, under a weakly monotonic contract $w \in W$ the agent may, with positive probability, be indifferent between investing and not investing. In order to make the principal's expected profit determined in such situations of indifference we assume that the agent will then choose the alternative that is more favorable to the principal.

The set of relevant contracts for bargaining between the two parties are those limited-liability contracts that result in outcome distributions that are not worse, in terms of their *ex ante* expected values, than the parties' outside options, and that, moreover, are not Pareto dominated by any other such contract—otherwise the parties would presumably renegotiate. Formally:

Definition 1. A contract $w \in W$ is **feasible** if (i) $w \in W_0$, (ii) $\mathbb{E}[\Pi(w, t)] \geq \bar{\pi}$ and (iii) $\mathbb{E}[U(w, t)] \geq \bar{u}$.

¹¹The same analysis goes through, mutatis mutandis, if the lower bound on payment is positive.

Definition 2. A contract w is **undominated** if it is feasible and there exists no feasible contract $w' \in W$ such that $\mathbb{E}[\Pi(w', t)] \geq \mathbb{E}[\Pi(w, t)]$ and $\mathbb{E}[U(w', t)] \geq \mathbb{E}[U(w, t)]$ with at least one inequality strict.

Let $F \subset W$ denote the set of feasible contracts and let $F^* \subset F$ be the subset of undominated contracts. Since we do not consider potential effects on third parties (employees, competitors or consumers), and since the contract space is rich enough to induce the agent to choose an optimal investment strategy, a contract is undominated in model A if and only if it is Pareto efficient.

2.3. Nash bargaining. As mentioned in the introduction, the contract is negotiated between the principal and the agent. In one extreme case, the principal makes a take-it-or-leave-it offer. In the opposite extreme case it is the agent who makes a take-it-or-leave-it offer. In practice, negotiations are arguably somewhere in between these two extremes. The parties' outside options reflect the intensity of competition for talented agents. Given the parties' outside options, the result of the negotiation may also depend on other aspects, here summarized under the rubric "bargaining power." For instance, if the agent is the current CEO of a company and the principal is the chairperson of its board, the agent may have considerable administrative resources, inside information and network contacts at his disposal when negotiating with the principal. It may also be costly for the company's reputation to fire a sitting CEO.

We include both parties' outside options (reflecting the degree of competition) and bargaining power in the analysis, and show how each aspect influences the bargained contract. For this purpose, we apply the generalized Nash-bargaining solution.¹² In the present model, a feasible contract is a *generalized Nash bargaining solution* if it maximizes the product of the two parties' profit/utility gains over their outside options, each gain raised to that party's "bargaining power". More precisely, for any $\beta \in [0, 1]$, let $F^{NBS}(\beta) \subseteq F$ be the set of generalized Nash-bargaining solutions when the principal's bargaining power is β and that of the agent is $1 - \beta$:

$$F^{NBS}(\beta) = \arg \max_{w \in F} (\mathbb{E}[\Pi(w, t)] - \bar{\pi})^\beta \cdot (\mathbb{E}[U(w, t)] - \bar{u})^{1-\beta}. \quad (5)$$

Clearly all such contracts are undominated; $F^{NBS}(\beta) \subseteq F^*$.

In sum, then, our solution concept is an application of the Nash bargaining solution to the contract negotiation phase, premised on the (selfish) agent's sequential rationality after a contract has been signed and after he has obtained his private information.

¹²In the pioneering paper (Nash, 1950) the two parties have equal bargaining power. The generalized, or asymmetric, bargaining solution appears in Roth (1979) and in subsequent work.

3. ANALYSIS OF MODEL A

Let w be any monotonic contract that the two parties have signed. In this post-bargaining situation, the agent's outside option is bygone and hence irrelevant for his investment decision.

3.1. The agent's investment strategy. Suppose the agent now has observed his private signal. Applying Bayes' rule, one readily obtains a closed-form expression for the agent's *conditionally expected* utility from investing, given any signal value $s \in \mathbb{R}$ that he may have observed,

$$\mathbb{E}[U(w, t) \mid s] = \frac{\mu \phi_t(s - r_G) u(w_G) + (1 - \mu) \phi_t(s - r_B) u(w_B)}{\mu \phi_t(s - r_G) + (1 - \mu) \phi_t(s - r_B)}. \quad (6)$$

In force of the MLRP condition and the (strict) monotonicity of the agent's Bernoulli function, this expected utility is, under any strictly monotonic contract, a continuous and strictly increasing function of the received signal s . Moreover, by condition TT, this conditionally expected utility ranges from $u(w_B)$ to $u(w_G)$ as the signal s runs from minus infinity to plus infinity. Hence, under any strictly monotonic contract the agent has a unique optimal *investment strategy*, namely, to invest if and only if his private signal surpasses a certain critical *threshold*, or *cut-off*, defined by the indifference condition that his conditionally expected utility from investing equals his expected utility from not investing;

$$\mathbb{E}[U(w, t) \mid s] = u(w_N). \quad (7)$$

Using this indifference condition, the agent's signal threshold is readily characterized in terms of the primitives of the model.

Proposition 1. *For any strictly monotonic contract $w \in W$, signal precision $t > 0$, and signal realization $s \in \mathbb{R}$, it is optimal for the agent to invest if and only if $s \geq s^*(w, t)$, where*

$$s^*(w, t) = \psi_t^{-1} \left[\frac{1 - \mu}{\mu} \cdot \frac{u(w_N) - u(w_B)}{u(w_G) - u(w_N)} \right]. \quad (8)$$

We see in (8) that the agent's signal threshold for investment, $s^*(w, t)$, is higher—the agent is more cautious—the higher is the odds for the bad state and the bigger is the ratio between his utility loss from the penalty and utility gain from the bonus. This is very natural and should be expected from a model such as this. The effect of

the contract upon an agent's investment decision, given his signal precision, can thus be summarized in a single number,

$$\rho(w) = \frac{\mu}{1-\mu} \cdot \frac{u(w_G) - u(w_N)}{u(w_N) - u(w_B)}, \quad (9)$$

(the inverse of the argument on the right-hand side of (8)), a number we will call the *carrot-stick ratio*. This ratio turns out to be useful in the subsequent analysis of models A and E. It is the probability-weighted ratio between the agent's *ex ante* expected utility gain from "always doing the right thing", namely, to invest in the good state and not in the bad. The higher the carrot-stick ratio, the *lower* is the agent's signal threshold for investing, that is, the wider is the range of signal values for which he will invest. The carrot-stick ratio, in turn, is higher the more likely is the good state, the better he is paid after a successful investment and after a failed investment, and the less he is paid when no investment is made. It also depends on his risk aversion since it depends on the ratio between two utility gains. We also note that the carrot-stick ratio is independent of the agent's signal precision (a fact that is useful in model E).

If the contract's carrot-stick ratio is below unity, then the agent would not invest in the absence of his private signal, while if the carrot-stick ratio is above unity he would invest even then:

$$\rho(w) \begin{matrix} \leq \\ \geq \end{matrix} 1 \quad \Leftrightarrow \quad \mu u(w_G) + (1-\mu)u(w_B) \begin{matrix} \leq \\ \geq \end{matrix} u(w_N). \quad (10)$$

Since by assumption uninformed investment is non-profitable, see (4), the principal certainly does not want the agent to invest if he would be uninformed.

Remark 1. *If condition [TT] were dropped, then it would be possible, for certain noise distributions and contracts, that equation (7) would hold for no signal $s \in \mathbb{R}$. The left-hand side of that equation would then either always be smaller than the right-hand side or it would always be larger. In the first (second) case, the agent will invest for no (all) signal values. For such noise distributions, feasibility of a contract thus requires that w_N be such that (7) holds for some signal $s \in \mathbb{R}$.*

3.2. Expected utility and profit. Knowing that he will receive an informative signal ($t > 0$), the agent's *ex ante* expected utility from accepting a contract is determined by his anticipated future information and subsequent investment decision. Formally, this *ex ante* expected utility can be expressed as

$$\mathbb{E}[U(w, t)] = p_G(w, t) \cdot u(w_G) + p_N(w, t) \cdot u(w_N) + p_B(w, t) \cdot u(w_B), \quad (11)$$

where $p_G(w, t)$ is the probability that the state is good *and* the agent invests, $p_B(w, t)$ the probability that the state is bad *and* (yet) the agent invests, and $p_N(w, t) =$

$1 - p_G(w, t) - p_B(w, t)$ is the residual probability that the agent does not invest. In force of Proposition 1 these probabilities can, for any strictly monotonic contract, be expressed in terms of the primitives of the model as

$$p_G(w, t) = \mu \int_{s^*(w, t)}^{+\infty} \phi_t(s - r_G) ds \quad (12)$$

and

$$p_B(w, t) = (1 - \mu) \int_{s^*(w, t)}^{+\infty} \phi_t(s - r_B) ds. \quad (13)$$

The associated profit to the principal in each of the three possible scenarios is the net return from the project, net of the payment to the agent, so the principal's ex ante expected profit is

$$\mathbb{E}[\Pi(w, t)] = (r_G - w_G) \cdot p_G(w, t) + (r_N - w_N) \cdot p_N(w, t) + (r_B - w_B) \cdot p_B(w, t). \quad (14)$$

3.3. Characterizing undominated contracts and Nash bargaining solutions. Having pinned down these ex ante expected values we are in a position to analyze the set of feasible contracts and its subset of undominated contracts. We first study the case of a risk-neutral agent and then turn to the case of risk averse agent.

3.3.1. Risk neutral agent. A feasible and strictly monotonic contract for a risk-neutral agent is undominated if and only if it induces the agent to use the same signal threshold as the principal would have used, had the principal had free access to the agent's private information. Moreover, if a feasible contract is not strictly monotonic, then it has to give the agent the same remuneration in all three outcomes. Formally:

Proposition 2. *If the agent is risk neutral, then a feasible and strictly monotonic contract $w \in W$ is undominated if and only if $\rho(w) = \rho(r)$, or, equivalently,*

$$\frac{w_G - w_N}{w_N - w_B} = \frac{r_G - r_N}{r_N - r_B}. \quad (15)$$

A feasible but not strictly monotonic contract $w \in W$ is undominated if and only if $w_G = w_N = w_B$.

In other words, a strictly monotonic and feasible contract is undominated if and only if its carrot-stick ratio (for a risk neutral agent) is the same as the carrot-stick ratio of the asset (also for a risk neutral agent). Equivalently, the ratio between its *bonus*, defined as $w_G - w_N$, and *penalty*, defined as $w_N - w_B$, has to match the corresponding ratio between the project's "potential gain", $r_G - r_N$, and "potential

loss", $r_N - r_B$. Under any such contract (there may be infinitely many), the agent will choose his signal threshold so as to maximize the sum of his and the principal's net gains—hence Pareto efficiency. We note that condition (15) is independent of the agent's signal precision (or talent). Moreover, it is independent of the project's "size"; condition (15) is unchanged if all three project returns are multiplied by one and the same positive scalar.

It follows from this result that the probabilities in (11) and (14), which in general depend on the contract at hand, are the same under all strictly increasing and undominated contracts.¹³ Moreover, by assumption, the *ex ante* expected return from the project, $\mathbb{E}[X]$, is lower than the no-investment return, r_N . This assumption is equivalent with $\rho(r) < 1$. Hence, Proposition 2 implies that the carrot-stick ratio is below unity for all undominated and strictly monotonic contracts, and this holds for any project in the class studied, and any signal precision:

Corollary 1. *If the agent is risk neutral, then any undominated and strictly monotonic contract $w \in W$ has a carrot-stick ratio below one, $\rho(w) < 1$.*

The second claim in the above proposition concerns contracts that are not strictly monotonic. It establishes that if such a contract is undominated then it is not possible that any one of the three payments differs from another. In other words, the agent then receives a flat salary. The reason is simple: if only one of the three payments would differ from the two others, then the agent will either always invest or never invest, and do so irrespective of his signal. Clearly, this is not in the interest of the principal who would then rather not hire the agent in the first place. Hence, all three payments have to either be distinct or be the same. In the latter case the agent's remuneration is the same whether or not he invests, irrespective of his private signal. In particular, he may then use the signal threshold $s^*(r, t)$ that is the one preferred by the principal, thus resulting in an equilibrium.

Another consequence of Proposition 2 is that the bonus and penalty may be high or low. All that matters is their *ratio*, $(w_G - w_N) / (w_N - w_B)$. This can also be given a geometric interpretation, namely, that all undominated contracts w for risk-neutral agents satisfy the linear equation

$$w_N = \frac{r_N - r_B}{r_G - r_B} w_G + \frac{r_G - r_N}{r_G - r_B} w_B. \quad (16)$$

Hence, when considering undominated contracts $w \in W$ for a risk-neutral agent, one may without loss of generality focus on any two of its components, say w_B and w_G , and neglect the third, here w_N , which then is given by (16). As noted above, such contracts represent the Pareto-efficient allocations.

¹³Then $p_G(w, t) = p_G(r, t)$, $p_N(w, t) = p_N(r, t)$ and $p_B(w, t) = p_B(r, t)$.

Moreover, for risk-neutral agents there are infinitely many Nash bargaining solutions associated with each bargaining-power parameter $\beta \in [0, 1]$. For any given β , all such solutions result in the same expected profit to the principal and the same expected utility to the agent. The indeterminacy arises because under risk-neutrality both parties are indifferent as to who takes on more risk, and there are a continuum of ways to share the risk while keeping the expected values constant. To put this more formally, let

$$\Delta = r_G \cdot p_G(r, t) + r_N \cdot p_N(r, t) + r_B \cdot p_B(r, t) - \bar{\pi} - \bar{w}. \quad (17)$$

This is the *potential gain of trade* between the two parties. It is the maximal expected total return from the investment opportunity, given the information available to the agent, net of both parties' outside options. The Nash bargaining solution prescribes that this potential gain of trade be divided in proportion to the parties' bargaining powers:

Proposition 3. *Suppose that the agent is risk neutral and that $F \neq \emptyset$. For any $\beta \in [0, 1]$, the principal's expected profit and the agent's expected utility are both constant on the set $F^{NBS}(\beta) \subset F$ of Nash bargaining solutions, with $\mathbb{E}[\Pi(w, t)] = \bar{\pi} + \beta\Delta$ and $\mathbb{E}[U(w, t)] = \bar{w} + (1 - \beta)\Delta$ for all $w \in F^{NBS}(\beta)$.*

We illustrate the above results by way of a numerical example. See Figure 1 below, which is drawn for $\mu = 0.5$, $r_G = 0.5$, $r_N = 0$ and $r_B = -1$. In other words, the good and bad states are equally likely, the ex ante expected net return from the project is negative, and the no-investment option yields zero net return. We consider a risk-neutral agent with normally distributed noise that has mean value zero and variance $\sigma^2 = 1/t$. The two parties' outside options are $\bar{\pi} = 0$ and $\bar{u} = \bar{w} = 0.1$. The diagram has w_B on the horizontal axis, w_G on the vertical, and w_N is set according to (16), that is, $w_N = 2w_G/3 + w_B/3$. The upper solid line represents the principal's participation constraint and the lower solid line the agent's participation constraint, both drawn for signal precision $t = 5$. The striped areas, one with positively sloped thin lines (for the principal) the other with negatively sloped lines (for the agent), represent those limited-liability contracts that meet each party's participation constraint, that is, lie below (above) the principal's (agent's) participation constraint. Their intersection, the hatched area, thus represents the set of undominated contracts. We see that this subset is defined by linear constraints and forms a convex subset of the set W_0 of limited-liability contracts; the area between the vertical axis (where $w_B = 0$) and the 45-degree line (where $w_G = w_B$). The dashed straight lines represent the case of a less talented agent ($t = 2$) with the same outside option as the more talented agent ($t = 5$). The participation constraints are still linear but define a smaller feasible set of contracts in comparison with the more talented agent, for obvious reasons. We also

note the slightly different slope of the new participation constraints, to be discussed in see Section 3.4.

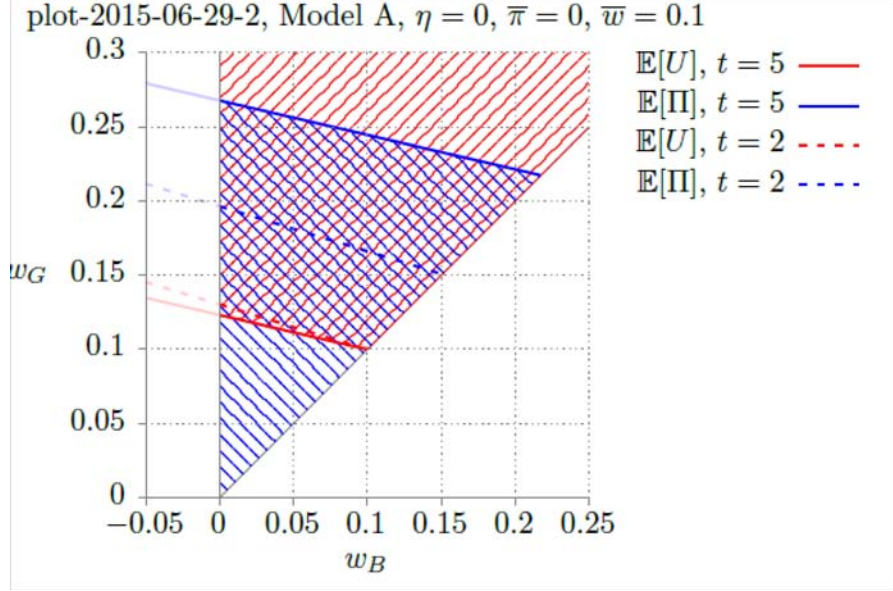


Figure 1: The set of undominated contracts in model A when the agent is risk neutral.

The Nash bargaining solutions are not indicated in the diagram but would make up straight lines that are parallel with the two parties' participation constraints. The higher β is, the closer will the NBS line be to the agent's participation constraint.

Remark 2. In the literature on CEO compensation, regularities have been found with respect to the amount of remuneration and the size of the company, see e.g. Gabaix and Landier (2008) and Kaplan and Rauh (2010). In the present model, this aspect can be illuminated by the following thought experiment. Suppose that all three net returns, r_B , r_N and r_G , along with the principal's outside option, $\bar{\pi}$, were scaled up by a factor $\lambda > 1$. This would represent scaling up the size of the company in question. The agent's signal precision t , being the inverse of the variance in his noise, would then have to be scaled down accordingly, by the factor $1/\lambda^2$. Holding his outside option constant, and assuming that the agent is risk neutral, his expected remuneration would, by (17), increase by a factor slightly below λ . More exactly, it would be λ times his former expected remuneration, minus $\beta\bar{w}$, that is, his outside option scaled by the principal's bargaining power. Hence, pay to a CEO of a larger company will be higher, and will increase a bit less than proportionately to the size of the company.

3.3.2. Risk averse agent. We now turn to the case of a risk averse agent, one with a strictly concave Bernoulli function u . Since there is no moral hazard involved in this model A, a feasible contract is undominated if and only if all risk is absorbed by the principal. For otherwise there would exist another feasible contract with the same expected utility and profit but with less risk placed on the agent. Hence, a risk-averse agent will receive a flat salary in all undominated contracts. Under such contracts, he is indifferent between investing or not, irrespective of his signal. In particular, it will be optimal for him to use the signal cut-off $s^*(r, t)$ that maximizes the asset's total expected return (to be shared between the two parties). This way, the agent has "traded away" all risk to the principal and the cake to be divided has grown (in expected utility terms).

Proposition 4. *For a risk averse agent, a feasible contract w is undominated if and only if $w_G = w_N = w_B$.*

In sum, then, model A does not lend any support to pay schemes involving bonuses when the agent is (even the slightest) risk averse. The next diagram shows indifference curves for both parties when the agent has CARA Bernoulli function $u(y) = y^{1-\eta}/(1-\eta)$ for $\eta = 0.9$, with all other parameters the same as in Figure 1 (and $\bar{u} = u(\bar{w}) = (0.1)^{0.1}/0.1$).

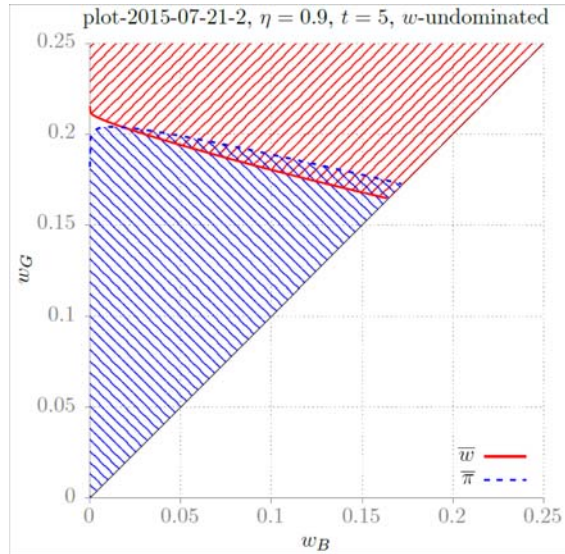


Figure 2: The set of feasible contracts when the agent is risk-averse.

Like in Figure 1, w_N is hidden and set according to (16). This is only for the sake of comparison. For risk-averse agents, (16) is not necessary for a contract to be

undominated.¹⁴ The subset of feasible contracts shown in Figure 2 is still convex, but no longer given by linear constraints. Moreover, unlike in the case of a risk-neutral agent, the two parties indifference curves intersect in such a way that, for every contract w in the interior of the cone W_0 of limited-liability contracts, there exist a wedge-shaped area of feasible contracts that are better for both parties, and those "better" contracts lie closer to the diagonal, at which the agent obtains a flat salary. This illustrates the logic behind the theoretical claim in Proposition 4, that all undominated contracts (and hence also all Nash bargaining solutions) take the form of a flat salary.

3.4. Agent with unknown talent. We here briefly consider the case when agents know their own signal precision (talent) but the principal does not know which agent has what signal precision (talent). The principal may then post a contract and use bonuses and penalties to deter less talented agents and attract more talented ones. A full analysis falls outside the scope of this paper, but we show how such an extension of model A can be made.

The dashed lines in Figure 1 above indicate the participation constraint also when the agent is less talented ($t = 2$ instead of $t = 5$) but has the same outside option. Such an agent is less valuable to the principal, so the principal's participation constraint is represented by a (much) lower line. More importantly, we also note that the dashed lines ($t = 2$) are slightly steeper than the solid ones ($t = 5$). This is because any increase in the penalty has to be compensated by a larger increase in the bonus in order to keep the less able agent indifferent. As a consequence, if the principal would offer a contract in the wedge between the solid and dashed lower lines (the participation constraints for agents of type $t = 2$ and $t = 5$), then the less talented agent would reject the contract but the more talented agent would accept it. A contract with such a (relatively high) bonus and penalty can thus be used by the principal as a screening device. However, as suggested by the diagram, the effectiveness of such screening depends on the talent distribution among agents and also on their outside options. In particular, if less talented agents have worse outside options—which may well be the case—then also such agents may "bite" on contracts with high bonuses and penalties, in which case screening does not work.

The following example illustrates a simple case. Suppose that there is a pool of risk-neutral agents of differing talent (signal precision) and that the principal has the possibility to make a take-it-or-leave-it offer to any one or all of these. Let $t_H > 0$ be the highest talent in the agent pool, that is, at least one agent in the pool has this

¹⁴For risk-averse agents, equation (16) is but one of infinitely many ways to specify w_N . Contracts of the form $w_N = \lambda w_G + (1 - \lambda) w_B$ for some $\lambda \in (0, 1)$ define a subspace that contains the straight line defined in Proposition 4. Diagrams drawn for a varitey of λ -subspaces show the same features as seen in Figure 2.

talent, and assume also that there is a positive probability that some other agent in the pool has a talent just below t_H .¹⁵ Suppose that each agent in the pool knows his own talent and that the principal knows the talent distribution but does not know which agent in the pool has what signal precision. If all agents have the same outside option, \bar{u} , then the principal will offer the contract $\hat{w} = (0, y, y + by)$ to all agents in the pool, where the *bonus rate* b , defined as $(w_G - w_N) / w_N$, equals $(r_G - r_N) / (r_N - r_B)$, and the remuneration y after non-investment is uniquely determined by the most talented agent's participation constraint, $\mathbb{E}[U(\hat{w}, t_H)] = \bar{u}$. Only agents with talent t_H will accept this contract, and the principal will hire one of these.

4. MODEL E

In model A, the agent's only choice was whether or not to invest in the risky project, a decision he based on his private signal, with due regard to the contract's terms and his exogenous signal precision. We now extend the model by letting the agent's signal precision be endogenous, to depend not only on his innate ability but also on how much effort the agent makes in order to acquire information. Hence, an agent who has signed a contract now faces a two-stage decision process, where the first stage is to choose signal precision and the second is to observe his private signal with the chosen precision and make the investment decision. The agent still obtains utility $u(w_B)$, $u(w_N)$ and $u(w_G)$, respectively, from his monetary remunerations in the three outcomes under any contract $w \in W$, but in addition he also experiences disutility (or cost) $C(t, a)$ from his effort to obtain signal precision $t \geq 0$, given his ability $a > 0$. Hence, his total utility is now

$$\tilde{U}(w, t, a) = U(w, t) - C(t, a),$$

where the random variable $U(w, t)$ is defined in model A. Zero precision, $t = 0$, represents a completely uninformative signal, or effectively “no signal.” This is the case when the agent does not gather any information at all about the project and is just as ill-informed as the principal. The agent's disutility or cost of obtaining zero signal precision is zero; $C(0, a) = 0$ irrespective of $a > 0$. At positive signal precisions t , the cost $C(t, a)$ is twice differentiable with $\partial C(t, a) / \partial t > 0$, $\partial C(t, a) / \partial a < 0$ and $\partial^2 C(t, a) / \partial t^2 \geq 0$. In other words, the agent's cost of obtaining information is increasing in the precision of his signal and decreasing in his's ability, and the marginal cost of obtaining information is non-decreasing. As a consequence, the agent's total expected utility becomes negative at sufficiently high signal-precision levels. This follows from the fact that his expected utility from remuneration is bounded from

¹⁵More precisely, for every $\varepsilon > 0$ there is a positive probability that at least one agent in the pool has talent in the interval $(t_H - \varepsilon, t_H)$

above (by $\mu u(w_G)$) while his disutility from effort is strictly increasing and convex and hence unbounded from above.¹⁶

For any contract $w \in W$, agent ability $a > 0$ and signal precision $t > 0$, the agent's conditionally expected total utility, given any received signal $s \in \mathbb{R}$, is

$$\mathbb{E} \left[\tilde{U}(w, t, a) \mid s \right] = \mathbb{E} [U(w, t) \mid s] - C(t, a), \quad (18)$$

where $\mathbb{E} [U(w, t) \mid s]$ is given in (6). Proposition 1 applies verbatim also to this extended model, since when the agent decides whether to invest his signal precision is already fixed and his disutility of effort sunk.

The agent's *interim* expected utility, that is, after he has chosen signal precision but before he observes his signal, is $\mathbb{E} [\tilde{U}(w, t, a)] = \mathbb{E} [U(w, t)] - C(t, a)$, obtained by integration of (18) over all signal values. We again use the analysis of model A, where $\mathbb{E} [U(w, t)]$ is given in equation (11) and the associated probabilities are given in equations (12) and (13). Write $T^*(w, a)$ for the set of signal precisions that maximize the agent's interim expected utility under a contract $w \in W$,

$$T^*(w, a) = \arg \max_{t \geq 0} \mathbb{E} [\tilde{U}(w, t, a)]. \quad (19)$$

As will be shown below, the maximand in (19) is in general not quasi-concave. Hence, under some contracts the agent will be indifferent between two (or more) quite distinct signal precisions. However, by standard arguments it can be verified that the set $T^*(w, a)$ is always nonempty and compact.

Let $\tilde{U}(w, a)$ denote the (random) utility to the agent under contract w when his ability is $a > 0$ and he acts sequentially rationally, that is, first chooses his signal precision t optimally (for him) and then makes his investment decision optimally (for him), given the signal he will receive. Formally,

$$\mathbb{E} [\tilde{U}(w, a)] = \max_{t \geq 0} \mathbb{E} [\tilde{U}(w, t, a)].$$

Likewise, let $\tilde{\Pi}(w, a)$ denote the associated (also random) profit to the principal, that is, her net return after the agent has chosen his signal precision, made his investment decision and been paid his agreed-upon remuneration. For the principal, it is the agent's signal precision and the agent's investment decision that matters, not how costly or strenuous it is for the agent to obtain information. Hence, for any precision $t > 0$ that the agent may choose, the principal's expected profit is the same as in model A, $\mathbb{E} [\Pi(w, t)]$, given in equation (14). If there are multiple signal

¹⁶In fact, it will never be optimal for the agent to choose signal precision above the precision level t where $C(t, a) = \mu u(w_G)$.

precisions that are optimal for the agent under the given contract, that is, if $T^*(w, a)$ is not a singleton set, then we will assume that he will select the signal-precision t in $T^*(w, a)$ that is best for the principal. (In other words, the agent has a lexicographic form of loyalty with the principal.) Then

$$\mathbb{E} \left[\tilde{\Pi}(w, a) \right] = \max_{t \in T^*(w, a)} \mathbb{E} [\Pi(w, t)].$$

This tie-breaking rule determines the principal's expected utility under all contracts.

Feasibility and dominance of contracts can now be defined as in model A:

Definition 3. A contract $w \in W$ is **feasible** in model E if

$$w \in W_0, \quad \mathbb{E} \left[\tilde{\Pi}(w, a) \right] \geq \bar{\pi} \quad \text{and} \quad \mathbb{E} \left[\tilde{U}(w, a) \right] \geq \bar{u}.$$

Definition 4. A contract w is **undominated** if is feasible and there exists no feasible contract w' such that

$$\mathbb{E} \left[\tilde{\Pi}(w', a) \right] \geq \mathbb{E} \left[\tilde{\Pi}(w, a) \right] \quad \text{and} \quad \mathbb{E} \left[\tilde{U}(w', a) \right] \geq \mathbb{E} \left[\tilde{U}(w, a) \right]$$

with at least one inequality strict.

Let $\tilde{F} \subset W$ denote the set of feasible contracts in model E and let $\tilde{F}^* \subset \tilde{F}$ be the subset of undominated contracts. Let $\tilde{F}^{NBS}(\beta) \subseteq \tilde{F}$ be the subset of feasible contracts that are generalized Nash-bargaining solutions when the principal's bargaining power is β ,

$$\tilde{F}^{NBS}(\beta) = \arg \max_{w \in \tilde{F}} \left(\mathbb{E} \left[\tilde{\Pi}(w, a) \right] - \bar{\pi} \right)^\beta \cdot \left(\mathbb{E} \left[\tilde{U}(w, a) \right] - \bar{u} \right)^{1-\beta}. \quad (20)$$

Because of the moral hazard involved in the agent's choice of signal precision (from which he has to share the fruits with the principal except when $\beta = 0$), undominated contracts are, unlike in model A, in general not Pareto efficient in model E.

We proceed to analyze this extended model version in the special case of normally distributed noise.¹⁷ More precisely, when the agent chooses positive precision t , the noise term ε in the agent's signal (2) is henceforth taken to be normally distributed with mean value zero and variance $\sigma^2 = 1/t$,

$$\phi_t(x) = \sqrt{\frac{t}{2\pi}} \cdot e^{-tx^2/2} \quad \forall x \in \mathbb{R}. \quad (21)$$

¹⁷We conjecture that subsequent results hold qualitatively for all noise distributions satisfying conditions MLRP and TT.

It is easily verified that this distribution meets both MLRP and TT.¹⁸ If the agent opts for zero signal precision, then he receives no signal.¹⁹

5. ANALYSIS OF MODEL E

We proceed by analyzing the agent's decision-making backwards in time.

5.1. The agent's investment decision. Suppose thus that the agent is at his final decision stage, that is, after he has signed a contract and decided upon his signal precision. From the analysis of model A we readily obtain closed-form expressions for the agent's optimal signal threshold under any strictly monotonic contract (see Proposition 1, (13), (12) and the appendix):

Lemma 1. *Assume normally distributed noise. The agent's optimal signal threshold under any strictly monotonic contract w , and given any signal precision $t > 0$, is*

$$s^*(w, t) = \frac{r_G + r_B}{2} - \frac{\ln \rho(w)}{(r_G - r_B)t}, \quad (22)$$

where $\rho(w)$, the carrot-stick ratio, is defined in (9). Moreover,

$$p_G(w, t) = \mu \Phi_1 \left(\frac{\ln \rho(w)}{(r_G - r_B)\sqrt{t}} + \frac{r_G - r_B}{2}\sqrt{t} \right) \quad (23)$$

and

$$p_B(w, t) = (1 - \mu) \Phi_1 \left(\frac{\ln \rho(w)}{(r_G - r_B)\sqrt{t}} - \frac{r_G - r_B}{2}\sqrt{t} \right). \quad (24)$$

It follows from these expressions that if the agent's signal precision would be extremely high—a virtually "noise-free" signal—then he would almost surely always "do the right thing"; invest if and only if the state of nature is good ($p_G(w, t) \rightarrow \mu$ and $p_B(w, t) \rightarrow 0$ as $t \rightarrow \infty$). By contrast, if his signal precision would be extremely low or zero—a virtually uninformative signal—then the agent would either (almost) never or (almost) always invest, depending on whether the carrot-stick ratio is below or above unity. If this ratio is below one, the agent's signal threshold will be very

¹⁸Its density ratio which is exponentially increasing in the signal s , for any signal precision $t > 0$:

$$\psi_t(s) = \exp \left[(r_G - r_B) \left(s - \frac{r_G + r_B}{2} \right) t \right].$$

¹⁹It is easily verified that the model is continuous at signal precision $t = 0$, that is, results for $t > 0$ converge as $t \rightarrow 0$ to the results for $t = 0$.

high when his signal is very imprecise; he is then willing to invest only when the signal is very positive. By contrast, if this ratio is above one, then the agent's signal threshold will be very low when his signal is imprecise; he is then willing to invest unless the signal is very negative. The reason for his apparently counter-intuitive behavior in the second case is that the carrot-stick ratio is above one precisely when his remuneration utility when not investing is lower than his expected remuneration utility from investing "blindly" (see 10). The last situation is highly undesirable for the principal; the agent then gambles with the principal's money without being better informed than the principal.

Using the expressions in Lemma 1 we obtain a closed-form expression for the agent's interim expected remuneration utility (that is, after he has chosen signal precision but before he receives his signal) under any contract $w \in W$ and any signal precision $t > 0$:

$$\begin{aligned} \mathbb{E}[U(w, t)] = & u(w_N) + \mu[u(w_G) - u(w_N)] \cdot \Phi_1\left(\frac{\ln \rho(w)}{(r_G - r_B)\sqrt{t}} + \frac{r_G - r_B}{2}\sqrt{t}\right) \\ & - (1 - \mu)[u(w_N) - u(w_B)] \cdot \Phi_1\left(\frac{\ln \rho(w)}{(r_G - r_B)\sqrt{t}} - \frac{r_G - r_B}{2}\sqrt{t}\right) \end{aligned} \quad (25)$$

General features of how this depends on the agent's signal precision are readily established. First, under any strictly increasing contract w this expected utility is, not surprisingly, strictly increasing in the agent's signal precision. Second, if the agent's signal precision is very high (he will then receive an almost perfectly accurate signal), then his expected utility will be close to the a priori expected utility from "doing the right thing in each state of nature",

$$\lim_{t \rightarrow \infty} \mathbb{E}[U(w, t)] = \mu u(w_G) + (1 - \mu) u(w_N). \quad (26)$$

By contrast, if the agent's signal precision is zero (or very low), then the agent's expected utility will be (close to) his a priori expected utility when acting sequentially rationally under the contract,

$$\mathbb{E}[U(w, 0)] = \max\{u(w_N), \mu u(w_G) + (1 - \mu) u(w_B)\}. \quad (27)$$

The first element of the set on the right-hand side exceeds the second if and only if the carrot-stick ratio $\rho(w)$ is below unity.

5.2. The agent's information acquisition. Having considered the agent's final decision stage, we now turn to his decision before that stage, his choice of signal precision. Of particular interest for understanding the agent's incentives when choosing his signal precision is his marginal expected remuneration utility with respect

to his chosen signal precision (that is, not considering the disutility of effort). Not surprisingly, this marginal utility is positive under any strictly monotonic contract. However, it is non-monotonic. Indeed, it is close to zero both when the signal is very noisy and when it is very precise. In other words, then a small increase in the agent's signal precision has little effect on his expected remuneration utility. More precisely, let

$$\kappa(w) = \mu(1 - \mu) \cdot [u(w_G) - u(w_N)] \cdot [u(w_N) - u(w_B)]. \quad (28)$$

We will refer to $\kappa(w)$ as the *power* of the contract w , the probability-weighted product of the agent's utility gains when "doing the right thing" in each state of nature. The bigger $\kappa(w)$ is, the more is at stake for the agent.

Lemma 2. *Assume normally distributed noise. For any strictly monotonic contract $w \in W$, the agent's expected utility from remuneration, $\mathbb{E}[U(w, t)]$, is differentiable in his signal precision $t > 0$. The associated marginal utility is everywhere positive but tends to zero as $t \rightarrow 0$ and as $t \rightarrow +\infty$. Moreover, at any precision level $t > 0$ this marginal utility satisfies*

$$\frac{\partial}{\partial t} \mathbb{E}[U(w, t)] = \frac{r_G - r_B}{2} \sqrt{\frac{\kappa(w)}{2\pi t}} \cdot \exp \left[-\frac{1}{2t} \left(\frac{\ln \rho(w)}{r_G - r_B} \right)^2 - \frac{t}{2} \left(\frac{r_G - r_B}{2} \right)^2 \right]. \quad (29)$$

We note that the marginal expected utility from remuneration is higher the higher is the power of the contract and the closer its carrot-stick ratio is to unity. The reason for the latter is that the carrot-stick ratio is unity precisely when the agent's signal threshold is at the midpoint between the low and high returns (r_B and r_G), a point at which the signal density is at its minimum (in the interval between the low and high returns) and hence its precision is most valuable for the agent's investment decision.

For any contract $w \in W$, signal precision $t > 0$ and agent ability $a > 0$, the agent's expected total utility, $\mathbb{E}[\tilde{U}(w, t, a)]$, is thus differentiable in t and can be decomposed into two terms,

$$\frac{\partial}{\partial t} \mathbb{E}[\tilde{U}(w, t, a)] = \frac{\partial}{\partial t} \mathbb{E}[U(w, t)] - \frac{\partial}{\partial t} C(t, a),$$

where the first term is given in (29) and the second simply is his marginal cost or disutility of obtaining signal precision t , given his ability $a > 0$. From this equation it follows that if the agent finds it optimal to choose a positive signal precision, then this optimal precision has to meet the first-order condition that its marginal benefit to the agent, in terms of his expected remuneration utility, equals its marginal disutility or cost to the agent from making the effort necessary to obtain information of that precision. For some contracts the agent may well instead prefer to exert no effort and obtain signal precision zero.

Recalling the definition of the set $T^*(w, a)$ of optimal signal precisions in (19), we have obtained that if $t \in T^*(w, a)$, then either $t = 0$, or else $t > 0$ and

$$\frac{\partial}{\partial t} \mathbb{E}[U(w, t)] = \frac{\partial}{\partial t} C(t, a), \quad (30)$$

that is, such that the expected marginal utility of better information equals the marginal costs of acquiring it.

As shown in Weibull, Mattsson and Voorneveld (2007), endogenous information acquisition in decision-making under uncertainty exhibit non-convexities. This is true also in the present model and is illustrated in the diagram below, which shows the agent's total expected utility as a function of his signal precision t , under a given contract, for five different ability levels ($a = 0.6, 0.8, 0.93, 1.2, 2$).²⁰

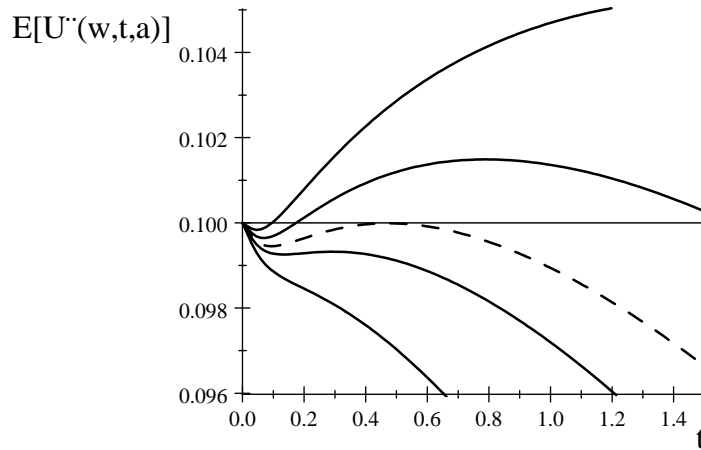


Figure 3: The agent's expected utility as a function of his signal precision.

The more able the agent, the higher his utility curve. The thin horizontal line represents the agent's utility $u(w_N)$ when not making any effort to gather information and not investing. We see that if his ability is very low, then his total expected utility is monotonically decreasing in signal precision (the lowest curve, $a = 0.6$). Such an agent thus makes no effort under the given contract. If the agent's ability is a bit, but not much, higher (say $a = 0.8$), then there exists an interior local maximum to his expected utility, but still the agent will choose to exert no effort. A highly able agents

²⁰The cost of information acquisition is here proportional to the precision. This is why zero precision may be a solution. If the cost of information acquisition is instead quadratic in the precision, the local maximum at zero is replaced by a local maximum near zero.

(say, $a = 1.2$ or $a = 2$) also has a local maximum to his expected utility, but this time this is also his global maximum. Such an agent thus strictly prefers to make a certain positive effort, and we see that this is higher the more able he is. The dashed curve shows the knife-edge case of an agent with intermediate ability ($a \approx 0.9$). This agent is indifferent between (i) making no effort and no investment (ii) and exerting a certain effort resulting in positive signal precision, $t \approx 0.5$, the point where the dashed curve touches the horizontal line. A slightly higher bonus would make him exert a positive (slightly higher) effort, while a slightly lower bonus would make his effort discontinuously drop to zero. Under the given contract no agent will choose a signal precision between zero and (approximately) 0.5, the latter thus being a point of discontinuity in behavior with respect to ability, under this particular contract.

Also the expected profit to the principal, defined in (14), can be expressed in terms of the primitives of the model for any contract $w \in W$ and signal precision $t > 0$:

$$\begin{aligned} \mathbb{E}[\Pi(w, t)] &= r_N - w_N + \\ &+ \mu [(r_G - w_G) - (r_N - w_N)] \cdot \Phi_1 \left(\frac{\ln \rho(w)}{(r_G - r_B) \sqrt{t}} + \frac{r_G - r_B}{2} \sqrt{t} \right) \\ &+ (1 - \mu) [(r_B - w_B) - (r_N - w_N)] \cdot \Phi_1 \left(\frac{\ln \rho(w)}{(r_G - r_B) \sqrt{t}} - \frac{r_G - r_B}{2} \sqrt{t} \right). \end{aligned} \quad (31)$$

Some general qualitative properties of this profit function will be noted. First, if the agent chooses a very high signal precision, then the ex ante expected profit to the principal will be close to the probability-weighted average profit after a successful investment and the profit when no investment is made,

$$\lim_{t \rightarrow \infty} \mathbb{E}[\Pi(w, t)] = \mu (r_G - w_G) + (1 - \mu) (r_N - w_N).$$

Second, if the agent chooses zero signal precision (no private information) and the carrot-stick ratio is below one, then he will not invest and hence the principal obtains the profit associated with non-investment,

$$\rho(w) < 1 \quad \Rightarrow \quad \mathbb{E}[\Pi(w, 0)] = r_N - w_N.$$

Third, and finally, if the carrot-stick ratio is above one, then an agent who chooses zero precision will instead invest and the principal's expected profit will thus be the probability-weighted average profit after a successful and failed investment,

$$\rho(w) > 1 \quad \Rightarrow \quad \mathbb{E}[\Pi(w, 0)] = \mu (r_G - w_G) + (1 - \mu) (r_B - w_B).$$

This expected profit is lower than the principal's outside option, $\mathbb{E}[\Pi(w, 0)] \leq \mathbb{E}[X] < \bar{\pi}$.

To push the analysis further without losing ourselves in involved calculations, we will often refer to the special case when the agent's disutility of information gathering is linearly increasing in the signal precision and inversely proportional to the agent's ability:

$$[\text{LIN}] \quad C(t, a) = ct/a \text{ for some } c > 0.$$

We will refer to this condition as "linear disutility of information acquisition". In other words, the agent's disutility of information acquisition, given his ability $a > 0$, is a linear function of the obtained signal precision t . Expressed instead in terms of the variance $\sigma^2 = 1/t$ of the noise term in his signal, his disutility of obtaining any given variance $\sigma^2 > 0$ is inversely proportional to the variance, $c/(a\sigma^2)$. In particular, the marginal disutility from variance-reduction is higher the lower the variance is and tends to plus infinity as the variance goes to zero (no noise), and it tends to zero as the variance goes to infinity (pure noise). The more able he is, the better the agent will be informed, at any given effort level.²¹ The parameter c can be interpreted as an information cost parameter. Arguably, the simple specification in [LIN] is a natural bench-mark case. We believe the qualitative conclusions under this specification hold qualitatively under more general specifications.²²

Under this specification, and using (29), the necessary first-order condition (30) for a *positive* signal precision t to be optimal for the agent boils down to

$$\frac{r_G - r_B}{2} \sqrt{\frac{\kappa(w)}{2\pi t}} \cdot \exp \left[-\frac{1}{2t} \left(\frac{\ln \rho(w)}{r_G - r_B} \right)^2 - \frac{t}{2} \left(\frac{r_G - r_B}{2} \right)^2 \right] = \frac{c}{a}. \quad (32)$$

This equation is illustrated in the diagram below, for a given contract w . The solid curve is the left-hand side of the equation, showing how the agent's marginal remuneration utility depends on his signal precision under a contract w with $\rho(w) \neq 1$. The horizontal lines represent different values for the right-hand side, the (constant) marginal cost to the agent of signal precision. The lowest of the three horizontal lines represents an agent with relatively high ability ($a = 2$), and the highest an agent with low ability ($a = 0.6$).

²¹Condition LIN is equivalent with assuming that the agent's signal precision t is bilinear in his information-gathering *effort* $z \geq 0$ and ability a , let $t = az$ and let his disutility from effort be $\psi(z) = cz$.

²²For example, if the cost instead is quadratic in the signal precision, $C(t, a) = ct^2/a$, then the agent will never choose zero signal precision. Instead, there will exist a locally optimal precision level, close to zero, and that positive but low precision level will play the same qualitative role as zero precision does in the current specification.

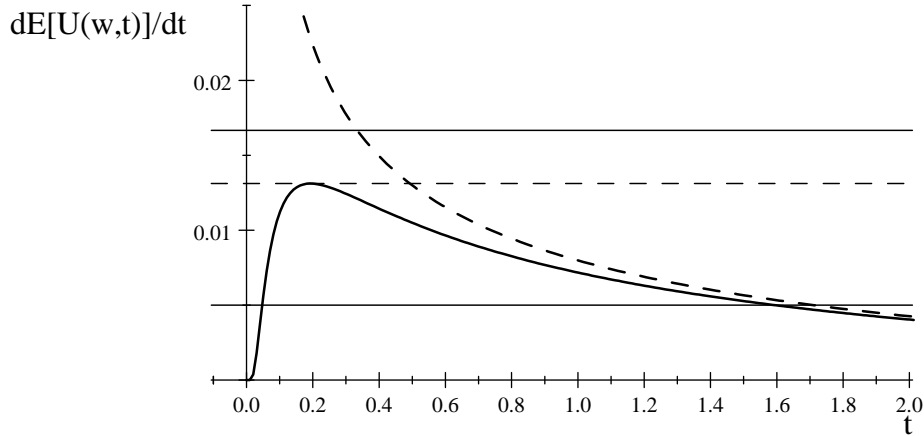


Figure 4: The agent's marginal expected utility from remuneration as a function of his signal precision.

For agents of high ability, the diagram shows that the necessary first-order condition (32) has two solutions (two intersections between the curve and the associated horizontal line). By contrast, for agents with low ability, the equation has no solution at all (no intersection); such an agent will exert effort zero since his marginal disutility of effort always exceeds his marginal utility from remuneration. At a critical intermediate ability level, indicated by the dashed horizontal line, the equation will have a unique solution. For any ability level that is high enough to warrant two solutions, the left-most intersection between the curve and the associated horizontal line is a local minimum of the agent's expected utility and the right-most intersection a local maximum. The dashed curve is the agent's marginal remuneration utility under a contract w' with unit carrot-stick ratio, $\rho(w') = 1$. In this knife-edge case, the left-hand side in the necessary first-order condition (32) is strictly decreasing from plus infinity towards zero, and hence has a unique solution.

In the same diagram we also see that the signal precision that is the unique local maximum (the one at the right-most intersection when $\rho(w) \neq 1$) is strictly increasing in the agent's ability and decreasing in the information cost parameter c . This is not surprising; more able agents will choose higher signal precisions under any given information cost, and agents with given ability will choose higher signal precisions if information costs fall, *ceteris paribus*.

For what contracts and agent types does then equation (32) have exactly two solutions? And can a lower bound be given on the precision level that constitutes the local maximum?

Lemma 3. *Assume normally distributed noise and linear disutility of information acquisition. For any contract $w \in W$ and agent ability $a > 0$, the agent's expected*

utility, $\mathbb{E} [\tilde{U}(w, t, a)]$, has a local maximum at a positive signal-precision level if and only if

$$\frac{1}{\sqrt{\pi}} \left(\frac{r_G - r_B}{2} \right)^2 \sqrt{\frac{\kappa(w)}{\left(\sqrt{[\ln \rho(w)]^2 + 1} - 1 \right) \exp \left(\sqrt{[\ln \rho(w)]^2 + 1} \right)}} > \frac{c}{a}, \quad (33)$$

Moreover, there exists at most one positive locally optimal signal precision, and this precision is the unique solution t to (32) that exceeds

$$\tau(w) = \frac{2}{(r_G - r_B)^2} \left(\sqrt{(\ln \rho(w))^2 + 1} - 1 \right). \quad (34)$$

We note that the left-hand side of condition (33) is plus infinity if and only if $\rho(w) = 1$. Hence, the condition is met by all contracts with unit carrot-stick ratio. We also note that the lower bound, $\tau(w)$, on positive equilibrium signal precisions, is zero if and only if the carrot-stick ratio is 1, while for all other carrot-stick ratios it is positive.

We are now in a position to characterize the signal-precisions chosen by any agent under any monotonic contract. To state this result, write $t^o(w, a)$ for the unique solution to equation (32) that exceeds $\tau(w)$, for any contract $w \in W$ and agent ability $a > 0$ satisfying (33). This is the unique interior local maximum to the agent's decision problem concerning information acquisition. Recall from (27) that the agent's investment decision at zero signal precision depends on the contract's carrot-stick ratio. The following inequality is necessary and sufficient for the agent to weakly prefer working (acquiring information) over shirking:

$$\mathbb{E}[U(w, t^o(w, a))] - C(t, a) \geq \mathbb{E}[U(w, 0)]. \quad (35)$$

Proposition 5. *Assume normally distributed noise and linear disutility of information acquisition. Then the agent will choose $t = t^o(w, a) > 0$ if (33) and (35) hold. Otherwise he will choose $t = 0$.*

In force of Proposition 5, we immediately obtain expressions for the agent's expected utility and the principal's expected profit from any contract. To be more specific: for any contract $w \in W$ satisfying (33), the agent's expected total utility is

$$\mathbb{E} [\tilde{U}(w, a)] = \max \{ \mathbb{E}[U(w, t^o(w, a))] - C(t^o(w, a), a), \mathbb{E}[U(w, 0)] \} \quad (36)$$

Otherwise $\mathbb{E}[\tilde{U}(w, a)] = \mathbb{E}[U(w, 0)]$. The principal's expected profit is $\mathbb{E}[\tilde{\Pi}(w, a)] = \mathbb{E}[\Pi(w, t^o(w, a))]$ if (33) and (35) hold. Otherwise $\mathbb{E}[\tilde{\Pi}(w, a)] = r_N - w_N$.

5.3. Feasible contracts, undominated contracts, and bargaining. We are finally in a position to analyze the set of feasible contracts, its subset of undominated contracts, and the Nash bargaining solutions. For this purpose, we first note that (33) and (35) are necessary conditions for a contract to be feasible. For otherwise the agent will not acquire any information (choose $t = 0$), in which case the principal would be better off without the agent—who the principal has to pay but is of no help. In particular, feasibility requires that there both be a bonus and a penalty. Mathematically, this follows from Lemma 2, which shows that if $\kappa(w) = 0$ then the agent’s marginal expected remuneration utility, not surprisingly, is zero. Hence:

Corollary 2. *Assume normally distributed noise and linear disutility of information acquisition. A necessary condition for a contract $w \in W$ to be feasible is that it is strictly increasing ($w_B < w_N < w_G$).*

In other words, both the bonus, $w_G - w_N$, and the penalty, $w_N - w_B$, must be positive. In order to be feasible in model E, a contract has to reward successful investments and punish failed investments, both in comparison with the agent’s pay when not undertaking the project.

Moreover, it can be established that if the agent is risk neutral, then the limited-liability constraint is binding for all undominated contracts; all feasible contracts with positive pay after a failed investment are dominated. The result is valid under any information costs that satisfy our general requirements.²³

Proposition 6. *Assume normally distributed noise. If the agent is risk neutral, then every undominated contract satisfies the limited-liability constraint with equality,*

$$w \in \tilde{F} \quad \Rightarrow \quad w_B = 0. \quad (37)$$

This result may come as a surprise, since a shift from a feasible contract $w = (w_B, w_N, w_G)$ to $w' = (w_B + \delta, w_N + \delta, w_G + \delta)$, for any $\delta \in \mathbb{R}$ such that also w' is feasible, changes neither the agent’s choice of signal precision nor his subsequent investment decision, since such a shift neither affects the carrot-stick ratio nor the power of the contract; $\rho(w') = \rho(w)$ and $\kappa(w') = \kappa(w)$. Hence, the agent’s incentives are identical under both contracts. One might therefore be led to believe that if w is undominated, then so is w' , contradicting the claimed result that $w_B = 0$ is necessary for a contract to be undominated. The explanation is that if $w_B = 0$ and $\delta > 0$, then w' would be dominated by a third contract w'' in which the agent would be paid zero after a failed investment and choose a higher signal precision than in w' , and

²³More precisely, we here assume that the cost $C(t, a)$ is twice differentiable with $\partial C(t, a) / \partial t > 0$ and $\partial^2 C(t, a) / \partial t^2 \geq 0$.

thus also higher than in the original contract w . Contract w'' would thus result in a "bigger cake" (in terms of expectations) to be divided between the two parties. The moral-hazard problem that arises when the agent chooses his information precision is due to the limited-liability constraint and the fact that the return in the bad state of nature is negative. In the absence of the limited liability constraint, one can show (see appendix) that the two parties would write a contract in which *all* the risk would be on the agent, though both parties are risk neutral. The principal would then keep the same residual sum of money in all three outcomes and fully incentivize the agent's information gathering, which would be the same as if the agent would own the whole project/asset himself.

We do not yet have a complete proof for the following claim, but we do have a few theoretical and numerical observations in favor of it:

$$w \in \tilde{F} \quad \Rightarrow \quad \rho(w) \leq 1. \quad (38)$$

Before providing these observations we note that it may in fact be surprising that there exists a uniform upper bound on the bonus rate—one that holds across agents of all abilities and for all cost functions for information acquisition. The reason is that a higher bonus-rate induces the agent to make more effort to gather information, an activity that is beneficial for the principal. The reason why there, nevertheless, may exist an upper bound is the distortion that a high bonus rate causes in the agent's subsequent investment decision, once he has obtained his private information. For if his bonus rate is high, then the agent will become "trigger happy", that is, invest even when the signal is too weak to result in a conditionally expected investment profit above her outside option. In fact, in all equilibrium contracts in model E, there is a trade-off between these two forces—and this holds even when the agent is risk neutral.²⁴

The diagram below, based on one of many computer simulations, illustrates the forces at work. It shows the set of contracts $w \in W$ that have $w_B = 0$, with w_N on the horizontal axis and w_G on the vertical. By Proposition 6, all undominated contracts for risk neutral agents are points (w_N, w_G) above the diagonal in this diagram, and each such point uniquely defines an undominated contract. The oblong sets are upper contour sets for the principal. Each such set thus represents all feasible contracts (with $w_B = 0$) that result in expected profits above a given level. The dashed curves indicate the lower boundaries of (the non-convex) upper contour sets for the agent. Each such set thus represents all feasible contracts (with $w_B = 0$) that result in an expected utility to the agent (including his cost for information acquisition) above a given level. The tangency points between the two parties' upper contour sets are thus undominated contracts, granted both parties' outside options give them less profit

²⁴We conjecture that (38) holds also for a risk-averse agent.

and utility, respectively, than in the upper contour sets in question. The collection of all such tangency points defines a curve, the "meta contract-curve" (in line with the notion of a contract curve in exchange economies).

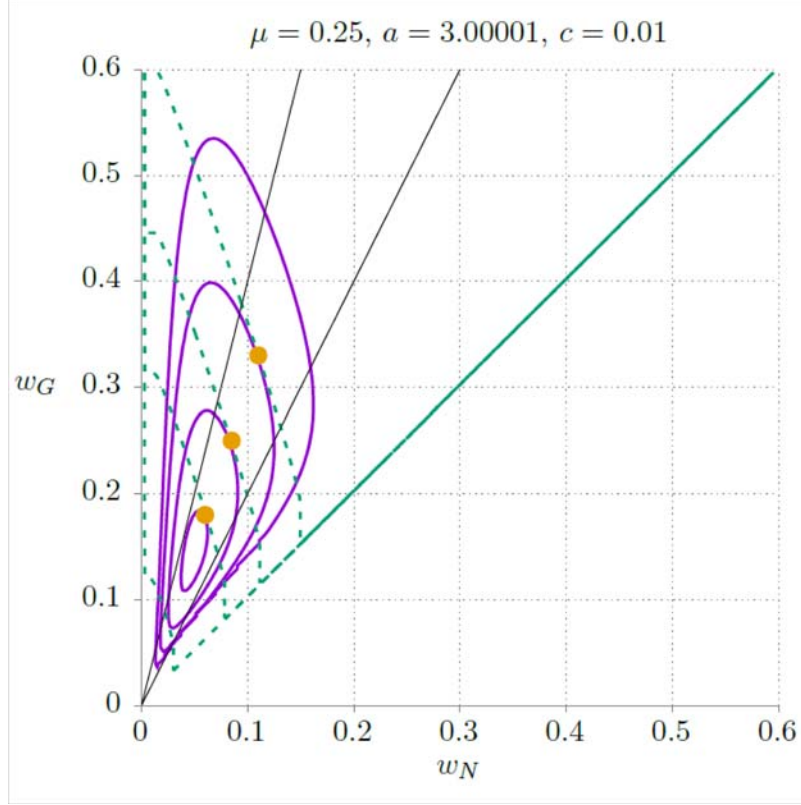


Figure 5: A contour map for the expected profit to the principal and the expected utility to the agent.

Some observations are called for. First, we note that all undominated contracts w in the diagram satisfy the double inequality $\rho(r) < \rho(w) < 1$, or, equivalently in this numerical example, $2 < w_G/w_N < 4$, the lower and upper bounds indicated in the diagram by straight lines.

Second, in the diagram there is an interior and unique globally optimal contract from the viewpoint of the principal, \hat{w} , in the diagram $\hat{w} \approx (0, 0.05, 0.13)$. This observation contrasts sharply with standard principal-agent models, where the principal typically drives the agent down to the agent's participation constraint. As the diagram shows, this is not always the case in the present model. For if the lowest utility isoquant in the diagram represents the agent's outside option, then the contract \hat{w}

would result in expected utility above that level. Indeed, this would be the equilibrium contract if the principal had all the bargaining power ($\beta = 1$), that is, could make a take-it-or-leave-it offer.

Third, the diagram suggests that the effect of changes in bargaining power are qualitatively indistinguishable from the effects of changes in the parties' outside options—for instance due to changed competition for talented agents. The effect of such changes is simply a movement along the meta contract-curve.

Fourth, the contract curve in the diagram has an approximately constant slope. Hence, while the size of the bonus increases with the agent's increased bargaining power, or increased competition for talent, the bonus rate is virtually unaffected; it stays close to 3.²⁵

Fifth and last, Figure 5 and similar numerical simulations shows that it may not be in the interest of the agent to make the principal believe that he is more able than he in fact is. For suppose that (i) the principal has all the bargaining power, (ii) the agent knows that his ability is $a = 3$ (as in the diagram), but (iii) the principal believes it is $a = 5$, and (iv) they both know that the agent's outside option has a very low value (say zero). Instead of offering the contract $\hat{w} \approx (0, 0.05, 0.13)$, the principal will then offer the contract $w^o \approx (0, 0.035, 0.115)$, with evidently lower expected utility for the agent.²⁶

One of our theoretical results that points towards (38), is that, for every contract with carrot-stick ratio above unity, there exists a contract with carrot-stick ratio below unity under which the sum of the principal's expected profit and the agent's (total) expected utility is higher. Hence, at least one of the two parties is better off in this alternative contract. This result holds for all risk neutral agents, of arbitrary ability and with arbitrary information-cost functions. Formally:

Proposition 7. *Consider a risk neutral agent and normally distributed noise. Suppose that $w = (0, w_N, w_G)$ is a feasible contract with $\rho(w) > 1$. There exists a contract $w' = (0, w'_N, w'_G)$ with $\rho(w') < 1$ such that*

$$\mathbb{E} \left[\tilde{\Pi}(w', a) \right] + \mathbb{E} \left[\tilde{U}(w', a) \right] > \mathbb{E} \left[\tilde{\Pi}(w, a) \right] + \mathbb{E} \left[\tilde{U}(w, a) \right]. \quad (39)$$

Turning back to the claim (38), we note that for risk neutral agents, the inequality in (38) defines an upper bound on the bonus rate in all undominated contracts.²⁷

²⁵Geometrically, the bonus in a contract is the vertical distance from the contract to the 45-degree line.

²⁶This points to a tension in the agent's incentives for truth-telling. On the one hand he may want to exaggerate his talent before being called to a job interview. On the other hand, during the negotiation he may have an incentive to down-play his talent, or to exaggerate the difficulties of the work task in question.

²⁷The bonus rate being defined, as before, as the ratio between the pay after a successful investment and the pay when no investment is made.

To see this, note that since $w_B = 0$ in all such contracts, and since then $\rho(w) = \mu(b - 1)/(1 - \mu)$:

$$\rho(w) \leq 1 \quad \Leftrightarrow \quad b \leq \frac{1 - \mu}{\mu}. \quad (40)$$

In other words, the higher the odds for the bad state of nature, the higher is the upper bound on the bonus rate. In particular, the upper bound is unity under Laplacian uncertainty, that is, when both states of nature are deemed equally likely. (If one defines a contract's *penalty rate* p in the same vein as its bonus rate b , then the limited-liability constraint evidently imposes unity as an upper bound; $p = (w_N - w_B)/w_N \leq 1$.)

We finally note that the European Parliament approved a proposal for capital-requirement regulation (CRD4) in April 2013, a proposal that also suggested a cap on bankers' bonuses, namely, that the variable pay (bonus) should not be above 100% of the fixed pay ("salary"), with some exceptions subject to shareholder approval (Financial Times, February 28, 2013). In terms of our model, this requirement amounts to $b \leq 1$.

6. CONTRACTS WITH STOCKS OR OPTIONS

Models A and E may be interpreted in terms of a company, represented by its board, here the principal, and a CEO candidate, here the agent. The latter may either be the current CEO or an outside candidate for the job. In both cases, the two parties enter a negotiation that may result in a contract being signed. If not, each party picks up his or her outside option. If the agent is the current CEO, the principal's outside option is to look for a new CEO and the agent's outside option is to take another job. If the agent is an outside candidate, the principal's outside option may be to continue with its current CEO and the agent's outside option may be to either continue on his current job or look for another job. We here focus on this latter case, that is, when the company already has a CEO and the agent is an outside candidate who may take over as CEO. In this case, let $\Pi_0 + \bar{\pi}$ be the company's stock value under its current CEO, where $\Pi_0 > 0$ is a constant that is independent of who is the CEO. The agent then receives utility \bar{u} . If the board and new CEO candidate agree on some contract $w \in W$, then the company's stock value instead becomes $\Pi_0 + \mathbb{E}[\Pi(w, t)]$, and the new CEO's expected utility is then $\mathbb{E}[U(w, t)]$. The company's future stock value, after the new CEO has made his investment decision and the state of nature has been realized, is the random variable $V(w, t) = \Pi_0 + \Pi(w, t)$. The project's profit, $\Pi(w, t)$, is the difference between its (state dependent) return and remuneration to the new CEO. Hence, $V(w, t) = \Pi_0 + r_\omega - w_\omega$ in each of the three possible outcomes, $\omega = B, N, G$. For the sake of brevity, we here assume that this future stock value is positive with probability one under all contracts w that will be considered. In other

words, the hiring of the new CEO and the implementation of the project at hand cannot result in bankruptcy of the company.

Real-life contracts often involve a salary and some share of the company's stock and/or options to buy such stock at some future time (but at their current price). Such payment schemes can be represented formally in the present model as follows.

6.1. Stock-based contracts. Suppose that the CEO's contract w consists of a fixed (positive or negative) pay y and a share $q > 0$ of the company's stock. Such a *stock-based* contract $w = (w_B, w_N, w_G)$ then has to satisfy the balance equation $w_\omega = y + q \cdot (\Pi_0 + r_\omega - w_\omega)$ for $\omega = B, N$ and G . Consequently, any such *stock-based* contract is of the form

$$w_\omega = \frac{y + q \cdot (\Pi_0 + r_\omega)}{1 + q} \quad \text{for } \omega = B, N, G.$$

Accordingly,

$$w_G - w_N = \frac{q}{1 + q} \cdot (r_G - r_N) \quad \text{and} \quad w_N - w_B = \frac{q}{1 + q} \cdot (r_N - r_B).$$

We note that for a risk neutral agent this implies that $\rho(w) = \rho(r)$. Consequently, in model A, the incentives of the company owners and a new and risk neutral CEO are perfectly aligned (by Proposition 2). Such a CEO will base his investment decision on the signal threshold preferred by the owners. We note that this holds for any stock-based contract. By contrast, a risk averse CEO should never be offered a stock-based contract in model A, since then all risk should be held by the company owners, see Proposition 4.

In model E, we know from Proposition 6 that every undominated contract for a risk-neutral agent has to satisfy $w_B = 0$.²⁸ This implies that

$$w_N = \frac{q}{1 + q} \cdot (r_N - r_B) \quad \text{and} \quad w_G = \frac{q}{1 + q} \cdot (r_G - r_B). \quad (41)$$

Hence, the only "free parameter" in undominated stock-based contracts in model E is the stock share, q , which is determined by the two parties' outside options and bargaining power.²⁹ We also note that all such contracts have the same bonus rate,

$$b = \frac{w_G}{w_N} - 1 = \frac{r_G - r_B}{r_N - r_B} - 1 = \frac{r_G - r_N}{r_N - r_B}.$$

In sum:

²⁸This is obtained by setting $y = -q(\Pi_0 + r_B)$.

²⁹The necessary condition for undominatedness, (40), is met since it is equivalent to our maintained hypothesis $\mathbb{E}[X] \leq r_N$.

Proposition 8. *Every stock-based contract—a fixed (positive or negative) pay plus a stock share—that meets the parties’ participation constraints (strictly) is undominated in model A if the agent is risk neutral. No stock-based contract is undominated in model A if the agent is risk averse. In model E, all undominated stock-based contracts for a risk neutral agents have the same bonus rate $b^* = (r_G - r_N) / (r_N - r_B)$.*

In model A, the bonus rate in an undominated stock-based contract is of the form

$$b = \frac{w_G}{w_N} - 1 = \frac{q \cdot (r_G - r_N)}{y + q \cdot (\Pi_0 + r_N)},$$

where y is the fixed (positive or negative) pay and $q > 0$ the stock share. We note that none of these bonus rates exceed the bonus rate b^* . In a sense, thus, the bonus rate in model E is driven up to the maximum bonus rate in model A. The intuition for this is that in model E there is an additional incentive effect, namely, to induce the agent to exert effort to obtain information.

6.2. Option-based contracts. The company’s future stock value, $V(w, t) = \Pi_0 + \Pi(w, t)$, can exceed its status quo value, $\Pi_0 + \bar{\pi}$, only after investment in the good state of nature.³⁰ Hence, if the contract contains an option for the new CEO to buy a pre-specified share $v > 0$ of the company’s future stocks, after the project has been realized but at the current stock price under status quo—that is, before the old CEO has been replaced—then the hired CEO will exercise his option only after a successful investment. Hence, if such an option-based contract $w = (w_B, w_N, w_G) \in W_0$ consists of a fixed pay $y \geq 0$ and the option to buy a share $v > 0$ of future stocks, then $w_B = w_N = y$ and

$$w_G = y + v \cdot [\Pi_0 + r_G - w_G - (\Pi_0 + \bar{\pi})],$$

or, equivalently,

$$w_G = \frac{1}{1+v} \cdot y + \frac{v}{1+v} \cdot (r_G - \bar{\pi}) \quad (42)$$

and hence $0 \leq w_B = w_N = y < w_G$.

It follows that, irrespective of the option share $v > 0$, equation (15) in Proposition 2 is violated. The reason is simple; there is no penalty in an option-based contract. Hence, with a risk-neutral new CEO, no agreed upon contract will be of this option-based kind in model A. Evidently, this also holds if the new CEO is risk averse—since with an option in his hand (and no penalties for bad investments), the CEO will always invest. In model E, all undominated contracts are strictly monotonic (by Corollary 2) and hence must pay less after a failed investment than after no investment, which is not the case with any option-based contract. In sum:

³⁰If no investment is made, $\Pi(w, t) = r_N - w_N \leq r_N \leq \bar{\pi}$. If investment is made in the bad state, $\Pi(w, t) = r_B - w_B \leq r_B < r_N \leq \bar{\pi}$.

Proposition 9. *All option-based contracts—a fixed pay plus options—are dominated in models A and E, for risk-neutral as well as risk-averse agents.*

This negative result concerning options suggests a class of *option-and-penalty* based contracts, namely, contracts $w \in W_0$ with $w_B = 0$, $w_N = y > 0$ and w_G as in (42). The necessary condition (40) for such a contract to be undominated in model E when the agent is risk neutral takes the form

$$v \leq \frac{(1 - \mu)y}{\mu(r_G - \bar{\pi}) - y} \quad (43)$$

In other words, in any undominated option-and-penalty contract, the pay y after non-investment and the option share v have to together satisfy this inequality (in addition to satisfying both parties' participation constraints). For any given pay $y \in (0, \mu(r_G - \bar{\pi}))$, this inequality provides an upper bound on the option share, v . This upper bound is increasing in the pay y and in the status quo profit of the company (under its current CEO), and is decreasing in the project's return in the good state of nature and in the probability for the good state of nature.

7. BEHAVIORAL BIASES

As is by now well documented in the economics literature, human decision-makers often exhibit behavioral biases, such as overconfidence in one's own ability and biases in one's information gathering. Indeed, such overconfidence among managers has been documented in the research literature, see Kahneman (2011, pp. 254-256) and Daniel and Hirshleifer (2015). Kindelberger (2005) has given a lively account of such overconfidence among actors on financial markets. Moreover, we may be unaware of these biases in ourselves. We here study the effects of some such possibilities on the resulting contracts and outcomes in models A and E in some special cases. We will first consider the case when both the principal and agent have the same exaggerated beliefs about the agent's ability or signal precision. The agent is overconfident but the principal does not realize this. We then turn to the case when the agent is overconfident but the principal understands this. We focus throughout on the case of a risk neutral agent.

7.1. Shared overconfidence. First, reconsider model A when the principal and agent both believe that the agent's signal precision (or talent) is t , while in fact it is $t' < t$. Hence, the agent is overconfident (about his talent or signal precision) and the principal shares this inflated belief. We will focus on the case of normally distributed noise. The analysis in section 3 clearly applies without any change. The effect of the shared overconfidence only appears after the fact, that is, in the probability

distribution over outcomes. The two parties will either agree on a strictly monotonic contract that induces the agent to choose an incorrect signal cut-off, adapted to the two parties exaggerated estimate of the agent's signal precision, or else they will agree on a flat salary in which case the principal will suggest the indifferent agent to anyhow use the same incorrect signal threshold, see Proposition 2. Whether the agent will be too trigger-happy or too cautious because of this shared misperception depends on the carrot-stick ratio, as seen in Proposition 1. In the arguably most interesting case of contracts with carrot-stick ratio below 1 (see Corollary 1), the agent will be too trigger-happy; his signal cut-off will induce him to invest in situations in which he shouldn't.

More precisely, assume a risk-neutral agent and normally distributed noise, and suppose that both parties believe the agent's signal precision is t but in fact it is $t' < t$. They will then agree on some contract w that satisfies (15). The agent will then invest if and only if his signal is at least $s^*(w, t)$ while for both parties the optimal signal threshold in fact is higher, $s^*(w, t')$. Investment will mistakenly be made for signals in the interval between these signal values, the interval

$$I = \left(\frac{r_G + r_B}{2} + \frac{|\ln \rho(w)|}{(r_G - r_B)t}, \frac{r_G + r_B}{2} + \frac{|\ln \rho(w)|}{(r_G - r_B)t'} \right). \quad (44)$$

We see that the less talented the agent is, the larger is this interval of signal values, in which the agent will invest although it would be in both parties' best interest not to do so, and the larger will the expected loss be to both parties be, given the agent's true signal precision (talent).

7.2. One-sided overconfidence. Second, again reconsider model A, but now assume that the agent believes his signal precision is $(\theta + 1)t$ while in fact it is t and the principal knows it is t . In other words, the agent misperceives his own talent but now the principal realizes this. We will call θ , when positive, the agent's *degree of overconfidence*. The agent is unaware of his overconfidence. Formally, we can modify model A by letting the noise term ε in the agent's private signal still be distributed $N(0, 1/t)$ while the agent believes its distribution is $N(0, 1/[(\theta + 1)t])$. Assuming that the agent's outside option is common knowledge between the principal and agent, the agent's overconfidence will benefit the (insightful) principal, since the agent will then overestimate his expected future remuneration from any contract that they may discuss, so the principal can offer a "leaner" contract than had the agent understood his true and lower signal precision. What is, more precisely the effect of such a behavioral bias upon the contract and outcome? This will, in part, depend on the two parties' bargaining power. We here focus on the special case when the principal

has all the bargaining power.³¹ Since the agent by hypothesis is risk neutral, but overestimates his signal precision, the principal will offer the contract that just meets the agent's (subjective) participation constraint (that is, in the eyes of the agent, given his exaggerated belief about his talent) and that is the most high-powered among such contracts (see Figure 1). The principal will set the bonus and penalty so that the agent will end up using the signal-threshold he should have used had he know his own signal precision. The set of feasible contracts is now defined by the two participation constraints, $\mathbb{E}[\Pi(w, t, \theta)] \geq \bar{\pi}$ and $\mathbb{E}[U(w, (\theta + 1)t)] \geq \bar{u}$ respectively, where the first inequality accounts for the principal's knowledge of the agent's true signal precision and the agent's degree of overconfidence, and the latter accounts for the agent's inflated belief about his signal precision.

Proposition 10. *Consider a principal who knows that the agent's signal has precision $t > 0$ and degree of overconfidence $\theta > 0$. Suppose that the agent is risk neutral, that there exists a feasible contract, and that the principal can make a take-it-or-leave-it offer. She will then offer the contract $\hat{w} = (0, y, y + by) \in W_0$, where*

$$b = \left(\frac{\mu}{1 - \mu} \right)^\theta \cdot \left(\frac{r_G - r_N}{r_N - r_B} \right)^{1+\theta} = \rho(r)^\theta \cdot \left(\frac{r_G - r_N}{r_N - r_B} \right), \quad (45)$$

and where $y > 0$ is uniquely determined by the agent's subjective participation constraint met with equality, $\mathbb{E}[U(\hat{w}, (\theta + 1)t)] = \bar{u}$.

This result calls for some remarks. First, we note that the contract is of the same form as screening contracts for agents of unknown ability, see in Section 3.4 (then $b = (r_G - r_N) / (r_N - r_B)$, in accordance with (45) for $\theta = 0$). Second, although the agent has a biased estimate of his ability, his equilibrium signal threshold under the equilibrium contract \hat{w} is unaffected by this bias. The principal will set the bonus rate such that it induces the agent to use the signal threshold for investment that the principal desires. Hence, it depends only on the principal's estimate t of the agent's signal precision and degree of overconfidence θ . Third, we see in (45) how the bonus rate depends on the agent's degree of overconfidence. The more overconfident the agent, the lower will be his bonus (recall that $\rho(r) < 1$ by hypothesis).

We finally note that the bonus can be expressed in terms of the carrot-stick ratio for the project itself (and a risk neutral agent):

$$b = \frac{r_G - r_N}{r_N - r_B} \cdot [\rho(r)]^\theta = \frac{r_G - r_N}{r_N - r_B} \cdot e^{\theta \ln \rho(r)} \quad (46)$$

³¹Other cases are more difficult to analyze since the two parties will then not agree on the size of the pie.

As noted before, our maintained hypothesis, that the ex ante expected return from the project is lower than the no-investment return, is equivalent with the project's carrot-stick ratio $\rho(r)$ being a number strictly between zero and one. We thus have

Corollary 3. *Under the hypothesis of Proposition 10, the bonus b is exponentially decreasing in the agent's degree of overconfidence, at the rate $\ln \rho(r) < 0$, ceteris paribus.*

Malmendier and Tate (2005) and Malmendier and Taylor (2015) have also analyzed the consequences of overconfidence for various types of corporate decisions. However, by contrast to our analysis they abstract from agency problems and information asymmetries. They conclude that a CEO's overconfidence results in over-investment when the firm's investment decisions are sensitive to the availability of internal funds and other forms of capital that the CEO perceives to be relatively cheap. They also find empirical support for this prediction, as well as for a number of other predicted consequences of overconfidence for firms' investment decisions.

Daniel and Hirshleifer (2015) study the possibility that overconfidence makes actors in financial markets more aggressive, and that this generates excessive trading and volatility in asset prices. Thus, their paper studies asset markets rather than principal-agent problems, and they do not study issues about the remuneration of CEOs.

7.3. Optimism. Overconfidence is arguably a form of optimism, or over-optimism. Another form of optimism (or pessimism) is when a decision-maker holds an unrealistically positive (negative) prior belief about the future success probability of a given project. In the present models, this would amount to the principal and/or agent attaching a too high (low) probability μ (given the information available in society at large concerning the project at hand) for the successful outcome of the project. This could either be a common prior, "shared optimism/pessimism," or distinct priors.³² We here briefly study the case when the agent believes that his signal is normally distributed with mean value zero (as we have hitherto assumed), while in fact it is normally distributed with a mean value that may be positive, zero or negative. A positive (negative) mean value δ represents an agent who has a tendency to over-sample favorable (disfavoring) information about his project. We will call δ the agent's *degree of optimism*, with $\delta = 0$ being the bench-mark case of an agent who samples information in an unbiased way.

Suppose that the principal knows the agent's degree of optimism. Formally, this can be inserted in model A by letting the noise term ε in the agent's private signal

³²Alternatively, optimism or pessimism could also take the form of an exaggeration of the returns (r_G and r_B) in the good and bad states of nature, respectively.

be distributed $N(\delta, 1/t)$ for some $t > 0$ and $\delta \in \mathbb{R}$ while the agent believes its distribution is $N(0, 1/t)$ (evidently, this bias can be combined with overconfidence on behalf of the agent, in which case his self-estimated signal precision may differ from his true signal precision). In this case, when the agent believes his information is balanced but it is in fact biased, and the principal knows the bias, the latter, can then design a contract accordingly. If the agent is optimistic or pessimistic, the principal can compensate the bias by “tilting” the contract in such a way that the agent will, in fact, use the signal cut-off desired by the principal. What is then the effect of such behavioral biases upon the contract? This will, in part, depend on the two parties’ bargaining power. We here focus on the case when the principal has all the bargaining power and the agent is risk neutral. Even then, there are multiple equilibrium contracts (see Proposition 2 and Figure 1). For the sake of definiteness, we here focus on the unique equilibrium contract that has maximal screening power, see Section 3.4.

The set of feasible contracts is now defined in terms of the accordingly adapted participation constraints, $\mathbb{E}[\Pi(w, t, \delta)] \geq \bar{\pi}$ and $\mathbb{E}[U(w, t)] \geq \bar{u}$, respectively, where the first inequality accounts for the principal’s awareness of the agent’s bias and the latter accounts for the agent’s naïve (and incorrect) belief that his signal is unbiased (and is hence calculated as before). Let $\hat{w}(y)$ denote contracts of the form discussed in Section 3.4.

Proposition 11. *Consider a principal who knows that the agent’s signal is $\varepsilon \sim N(\delta, 1/t)$ while the agent believes it is $\varepsilon' \sim N(0, 1/t)$. Suppose that the agent is risk neutral and the principal can make a take-it-or-leave-it offer ($\beta = 1$). Then she will offer the contract with maximal screening power, namely, $\hat{w}(y) = (0, y, y + by)$, where*

$$b = \frac{r_G - r_N}{r_N - r_B} \cdot e^{-t \cdot (r_G - r_B) \cdot \delta} \quad (47)$$

and the pay $y > 0$ is uniquely determined by the indifference condition $\mathbb{E}[U(\hat{w}(y), t)] = \bar{u}$.

This result calls for some remarks. First, we note that in the special case of an unbiased agent ($\delta = 0$) the contract $\hat{w}(y')$, not surprisingly, is identical with the contract in Section 3.4. Then the bonus rate is $b = (r_G - r_N) / (r_N - r_B)$. Second, we see in (47) that the equilibrium bonus rate b is exponentially decreasing in the agent’s signal bias, δ . In other words, the more optimistic the agent is, the lower will the principal set his bonus rate. The equilibrium bonus rate for an optimistic (pessimistic) agent is lower (higher) than that for an unbiased agent.

It is also interesting to compare the effects of optimism with the effects of overconfidence. Comparing equations (47) and (45), we see that the effects differ. In

particular, the bonus to an agent with degree of optimism δ is lower than the bonus to an agent with degree of overconfidence θ if and only if

$$\delta \cdot (r_G - r_B) \cdot t + \theta \cdot \ln \rho(r) > 0.$$

We see from this inequality that it is more likely to hold the higher the agent's signal-precision is, that is, the more talented he is. More precisely, for any given positive degrees of overconfidence and optimism, there exists a critical signal precision such that optimism results in a lower bonus than overconfidence for all agents with signal precision above this critical level. (Recall that both bonuses are lower than the bonus to a balanced agent, one who is neither optimistic nor pessimistic and has a correct belief about his own signal precision.)

7.4. Loyalty. Third, consider an agent who feels some “substantial” loyalty towards the principal, more substantial than the lexicographic loyalty we already have assumed, such as a CEO who cares about the company (its owners and/or other employees). Such loyalty may be represented in the form of an extra term in the agent's utility function, a term that is strictly increasing in the principal's profit (or, say, proportional to it). Evidently, this may have an effect on the bargaining over the contract (weakening the agent's bargaining position). Moreover, if the agreed-upon contract is strictly increasing and we consider model A, then such loyalty will have no effect at all on the agent's investment decision, since the contract still necessarily meets condition (15). If the contract is a flat salary, then the agent will no longer be indifferent (as he is in the original version of model A) but will have an incentive to pick the right signal cut-off for his investment decision. Hence, loyalty so to speak reinforces such flat salary contracts, and therefore such contracts may be more common in situations like that in model A. In model E, the effect of loyalty is likely to induce the agent to make more effort to gather information. Hence, the principal will gain from employing such an agent and may not need to pay the full prize (since the agent is loyal and thus willing to accept a slightly worse contract).

7.5. Career concerns. Fourth and finally, we briefly consider an agent who cares about his reputation. If future employers are outsiders to the principal-agent relationship, it may be hard for them to judge the project, the parties' outside options, the agent's behavior etc. However, a successful or failed investment is usually noted also by outsiders, and such observations may play a role for the agent's future employment opportunities. Such career concerns will not change the agent's investment decision in model A, but may increase the agent's incentive in model E to gather information about investment opportunities, and hence the agent may obtain more precise information than had he only cared about his expected remuneration. Such

career concerns could easily be incorporated in model E by way of his Bernoulli function, for instance, by letting it contain an additional positive (negative) term after a successful (failed) investment, a term that represents the present value of his reputation for future employments. Arguably, this would lower the equilibrium bonus rates in all undominated contracts. In particular, the upper bound in (38) and (40) should still hold, *a fortiori*.

8. RELATED LITERATURE

Theoretical studies of the remuneration of CEOs in large firms are in most cases designed to deal with specific issues. Examples are the importance for the pay of firm size, conflicts between extrinsic (economic) and intrinsic work motivation, the relevance of incentives for innovative tasks, and the consequences of new communication and information technologies, increased volatility of the business environment, overconfidence or optimism among agents, etc. For a survey of the literature on CEO compensation as an agency problem, see Babchuk and Fried (2013). In this paper we have modeled the determination of pay for CEOs as the result of the interaction between two mechanisms that have dominated studies in this field for a long time: (i) keen competition for talent in the market for managers, and (ii) strong bargaining power of managers within firms.

Models relying on the first mechanisms (usually in the form of competitive sorting of heterogeneous managers to heterogeneous firms) are largely inspired by Sherwin Rosen's seminal paper (1981) on "the theory of superstars". However, different authors in this tradition have relied on different versions of this theory, depending on which precise question they address. In early versions of this theory, in which moral hazard was neglected, a main conclusion was that the most talented managers were assigned to the largest firms, since that is where CEOs have the strongest impact. However, Edmans and Gabaix (2011) have shown that this allocation is distorted in the presence of risk, risk aversion and moral hazard. They also show that the size of this distortion increases in the dispersion of managerial ability and decreases in the dispersion in firm size. Moreover, quantitative simulations by Gabaix and Landier (2008) suggest that even quite modest differences in talent may result in huge differences in remuneration – basically as a result of the (assumed complementary) interaction between talent and firm size. Moreover, it is rather generally believed among observers that generous bonus systems, combined with limited liability, contributes to excessive risk-taking. However, Malcolmson (2010) shows that this effect may be mitigated by properly designed contracts between principals and agents.

A number of papers also deal with the consequences of the interaction between economic incentives and intrinsic work ethic in a competitive environment, see Carlin and Gervais (2009), Kosfeld and von Siemens (2011), and Bénabou and Tirole

(2013). A result of these analyses is that tasks pursued on the basis of intrinsic work motivation – tasks that often provide positive externalities within firms – tend to be squeezed out by simpler tasks when high-powered economic incentives are applied. The mechanism is, of course, similar to that in the general literature on multitasking (Holmström and Milgrom (1991)). As pointed out by Bénabou and Tirole, a “bonus culture” also tends to develop within firms under these circumstances.

Some authors (e.g. Holmström, 2004) have suggested that agents’ short-sightedness may be mitigated by lengthening the vesting horizon in share or option contracts. Edmans et al. (2006) have worked out more dynamic contract arrangements with an optimal horizon of the incentive structure. The contract is imbedded in a life-cycle model for CEOs where they may adjust their private saving, as well as inflate their remuneration by temporary discretionary actions. The contracts are tied to what the authors call dynamic incentive accounts, which include both cash and the firm’s equity. The accounts include rebalancing arrangements to guarantee that the equity proportion all the time is sufficient to generate satisfactory incentives.

Axelson and Bond (2015) instead apply a dynamic (multi-period) perspective on tasks by managers of funds in the context of an efficiency-wage model. The difficulties to monitor agents with financial tasks create a moral hazard problem, which is only partly solved by equilibrating contracts that provide higher remuneration than for other tasks (and sectors). The authors also conclude that it is profitable for a firm to assign even young employees to tasks with strong temptations to moral hazard, and with relatively high remuneration, rather than postpone such contracts until the employee has succeeded on tasks with less moral hazard. A main conclusion of the paper is that moral-hazard problems endogenously worsen in good times, the reason being that the financial sector has a higher proportion of relatively inexperienced employees in such situations. By contrast to many other papers, the study of Axelson and Bond highlights not only the misallocation of talented individuals among tasks (and sectors in the economy) but also the optimal sequencing of different tasks for employees over their life cycle. In particular, the emphasis on endogenous career dynamics, and not just the distribution of talents on tasks, differentiate this paper from other in the same field (such as Terviö, 2009).

Axelson and Bond (2015) also suggest that the analysis may be deepened in the future by including employee costs of learning about the functioning of asset management. It should then be observed that this knowledge is acquired before the agents are confronted with concrete investment decisions, hence at a time when the effort costs of learning is already sunk. Indeed, to do just that is one of the attempted contributions of our paper.

Not only competitive-sorting models, but also models with strong bargaining power of managers appear in different versions. A common feature of these models is, however, that the remuneration of CEOs tend to be inflated since share owners

and board members may be unable to evaluate the consequences of complicated and generous remuneration schemes, such as stock options; see, for instance, Hall and Murphy (2003).

Gustavo Manso (2011) has instead studied the issue of designing optimum contracts for motivating managers to innovate. He concludes that optimal incentive schemes designed to stimulate innovation should be structured differently than contracts designed to avoid moral hazard in the connection with general work effort. In particular, optimal contracts should, according to Manso, exhibit considerable tolerance for early failures, combined with rewards for long-term success. More concretely, an optimum contract should, according to Manso, include a combination of stock options with long vesting periods and “golden parachutes”.

Different authors also provide different explanations for the rapid increase in the remuneration of CEOs in recent decades. Some authors – such as Gabaix and Landier (2008) and Kaplan and Rauh (2010) – refer to the gradually higher capital values of large firms. Other authors argue that new communication and information technologies have resulted in a gradually more complicated role for CEOs and that this has increased the demand for high-quality managers (Caricano and Rossi-Hansberg, 2006; and Giannetti, 2011). Several authors have also asserted that large firms today require general skill as a complement to specialists skills, and that this has increased the demand for talent; see, for instance, Custodio, Ferreira, and Matos (2013). Some authors also refer to increased volatility of the business environment in recent decades, which is also asserted to *increase* the demand for highly skilled manager (e.g., Dow and Raposo, 2005).

There is also a literature on how overconfidence influences the remuneration of CEOs. Following much of the literature in this field (as well as sections 8.1 and 8.2 of this paper), one then defines overconfidence as overestimation of one’s own talent. There is a well-documented tendency of individuals to consider themselves “above average” on positive characteristics; see, for instance, the brief survey in Malmendier and Tate, 2005.³³ People typically overestimate the precision in their private information, judgement and intuition, with substantial effect on their economic behavior, see Heller (2014) for an analysis of such phenomena and their evolutionary foundations. Arguably, such overconfidence tends to result in excessive investment and risk-taking. Indeed, such overconfidence is a widely held explanation of the recent international financial crisis, in particular after the so called “leveraged buyout revolution” in the 1980s, see e.g. Holmström and Kaplan (2001). Malmendier and Tate argue that in the case of overconfidence stock options are unhelpful, since there is no need to boost

³³In their empirical analysis, Malmendier and Tate use two alternative measures of overconfidence. The first is based on the manager’s own personal portfolio transactions. The second measure is based on how managers are portrayed in the press.

investment incentives in this case. Indeed, as pointed out by Gervais et al. (2011), and analyzed here, if principals have information about the overoptimism of managers, they may be able to mitigate excessively high risk-taking by avoiding contracts with strong incentives for risk-taking.

Finally, there is a literature on the consequences of overoptimism, rather than overconfidence, among managers – in the sense that managers’ expectations about the return on investment are unrealistically optimistic. However, as in the case of overconfidence, properly designed contracts may mitigate the tendencies among managers to take high risks. Indeed, empirical studies indicate that contracts in the real world, in fact, often provide weaker incentives, and result in smaller total remuneration, for highly optimistic managers than for others, see, for instance, Otto (2014) and the literature reported there.³⁴

9. CONCLUDING REMARKS

Discussions about the distribution of income in developed countries have in recent decades focused on the large, and in several countries increasing, share of national income received by managers in large corporations. This means that clarification of the mechanism behind the remuneration to managers is of great general interest, including policy interest. As pointed out in Section 8, a number of specific aspects of this issue have already been discussed in the scholarly literature in recent years. Generally speaking, the purpose of this paper is to provide a canonical analytical framework, or “work horse”, that can be applied to analyses of many different aspects of the remuneration of managers, or more generally agents in principal-agent relations where the task of an agent, such as a manager of a corporation, is to make well-informed (investment) decisions on behalf of a principal who owns capital, or somebody who represents capital owners such as owners of stocks in corporations.

The principal and the agent initially bargain over the pay scheme for the latter. The bargaining outcome then depends both on competition for agents in the open market and on the relative bargaining power of the two parties inside the organization for which the agent works, allowing for the possibility that the agent may be the current CEO with considerable bargaining power. Indeed, our analysis highlights the interaction between these two determinants of remuneration contracts.

We have presented two versions of our work-horse model – model A with exogenous effort by the agent to acquire information about investment opportunities, and model E with endogenous effort on this matter. In both models, the agent’s information is

³⁴Otto (2014) also assesses each CEO’s optimism with two measures (although different ones than those in the preceding footnote). One is based on how eager the CEO turns out to be to cash in his options. The other measure of optimism is based on voluntarily released earnings forecast by the management.

private and he uses it in a subsequent decision whether or not to invest in a certain project. This combination of forces may result in high total pay for the agent in both models. However, in model A the case for other payment schemes than flat salaries is weak when the agent is risk averse – although bonuses combined with penalties may help the principal to screen and select suitable agents when the principal is uncertain about the agent’s competence.

By contrast, model E contains a double moral hazard problem: how much information should the agent acquire and what investment decision should he make? As a consequence, all equilibrium contracts in model E involve both bonuses and penalties. Indeed, if the bonus and/or penalty are set sufficiently low, the agent will abstain altogether from gathering information, and his investment decision will be as uninformed as it would be for the principal herself. The principal would then be better off without the agent, since she then avoids remuneration costs. Moreover, in some situation the agent may switch discontinuously from solid information gathering to zero-gathering in response even to a slight reduction of the bonus. When analyzing these issues it turns out that what we call the carrot-stick ratio is a useful measure of the incentive power of a contract, both in models A and E. This measure is simply the probability-weighted ratio between the agents ex ante expected utility gain from investing in the good state and not in the bad.

We also study the consequences for the equilibrium contract of a number of changes in the environment (comparative statics). Examples are shifts in the degree of competition for agents or in the relative bargaining power of principals and agents. Indeed, in model E it turns out that the effect of these two shifts are rather similar. We have also applied our work-horse model to a number of other questions, such as the consequences for the remuneration contract of the size of firms. We have also studied the role of stock- and option-based contracts. For instance, while a fixed salary in the equilibrium contract may be combined with stocks, it is never combined with an option share.

Finally, we apply our model to an analysis of the consequences of behavior bias such as overconfidence and overoptimism. In one version of overconfidence both the principal and the agent overestimate the agent’s signal precision (or talent) – what we have called “shared overconfidence”. A risk-neutral but overconfident agent may invest even if it would be in both parties’ interest if he did not. The situation is quite different if the principal, by contrast to the agent, is fully informed about the latter’s overconfidence – what we have called “one-sided overconfidence”. Assuming that the agent’s outside option is common knowledge between the principal and the agent, the latter’s overconfidence will actually benefit the (insightful) principal since the agent will then overestimate his expected future remuneration from any contract that they may agree about – a point made by Gervais et. al. (2011). The explanation is that the principal and the agent, as the result of the overconfidence, may agree

about only modestly incentivized contract. Overoptimism – in the sense that the agent has an unrealistically positive prior belief about the future success probability of a given project – has somewhat different implications than overconfidence. Indeed, we specify under what conditions the equilibrium bonus is lower for an agent with given level of overoptimism than for an agent with the same degree of overconfidence.

10. APPENDIX

Proposition 1: Clearly

$$\mathbb{E}[U(w, t) \mid s] = \frac{\mu\psi_t(s)}{\mu\psi_t(s) + 1 - \mu}u(w_G) + \frac{1 - \mu}{\mu\psi_t(s) + 1 - \mu}u(w_B), \quad (48)$$

where $u(w_G) > u(w_B)$ for every contract $w \in W$ and any strictly increasing Bernoulli function u . Hence, under MLRP and TT, $v(s) = \mathbb{E}[U(w, t) \mid s]$ defines a continuous and strictly increasing surjection from \mathbb{R} onto the open interval $(u(w_B), u(w_G))$, which for every contract $w \in W$ contains $u(w_N)$. Using (48) in (7), (8) is obtained.

Proposition 2: Assume that $w^* \in F$ is strictly monotonic and let

$$B(w^*) = \{w \in W_0 : \mathbb{E}[\Pi(w, t)] \geq \mathbb{E}[\Pi(w^*, t)] \quad \wedge \quad \mathbb{E}[U(w, t)] \geq \mathbb{E}[U(w^*, t)]\}.$$

Clearly $w^* \in B(w^*)$. Moreover, w^* is undominated if and only if

$$w^* \in \arg \max_{w \in B(w^*)} G(w, t)$$

where $G(w, t)$ is the expected value of the *sum of payments* to the two parties,

$$G(w, t) = \sum_{\omega=B, N, G} [(r_\omega - w_\omega) + w_\omega] \cdot p_\omega(w, t)$$

Thus $G(w, t)$ is the expected remuneration to an agent who has contract $r \in W$. From Proposition 1, we thus obtain

$$w^* \in \arg \max_{w \in W} G(w, t) \quad \Leftrightarrow \quad s^*(w^*, t) = s^*(r, t),$$

where (see (8))

$$s^*(w^*, t) = \psi_t^{-1} \left[\frac{1 - \mu}{\mu} \cdot \frac{w_N^* - w_B^*}{w_G^* - w_N^*} \right] \quad \text{and} \quad s^*(r, t) = \psi_t^{-1} \left[\frac{1 - \mu}{\mu} \cdot \frac{r_N - r_B}{r_G - r_N} \right].$$

Since ψ_t^{-1} is strictly monotonic,

$$w^* \in \arg \max_{w \in W} G(w, t) \Leftrightarrow \frac{w_N^* - w_B^*}{w_G^* - w_N^*} = \frac{r_N - r_B}{r_G - r_N} \quad (49)$$

We are now in a position to drive home the claim. First, since $B(w^*) \subset W$:

$$\frac{w_N^* - w_B^*}{w_G^* - w_N^*} = \frac{r_N - r_B}{r_G - r_N} \Rightarrow w^* \in \arg \max_{w \in B(w^*)} G(w, t).$$

In other words, $\rho(w^*) = \rho(r)$ is sufficient for a contract $w^* \in F$ to be undominated.

Second, in order to show that $\rho(w^*) = \rho(r)$ is also necessary, suppose that w^* does not satisfy this equation. Then $s^*(w^*, t) \neq s^*(r, t)$. Suppose that $s^*(w^*, t) < s^*(r, t)$. (A similar argument can be developed for the case $s^*(w^*, t) > s^*(r, t)$.) For any $\delta, \varepsilon \geq 0$ write $\hat{w}(\delta, \varepsilon)$ for the contract \hat{w} with $\hat{w}_B = w_B^* + \varepsilon$, $\hat{w}_N = w_N^* + \varepsilon$ and $\hat{w}_G = w_G^* + \varepsilon - \delta$. Then

$$\frac{w_N^* - w_B^*}{w_G^* - w_N^*} < \frac{\hat{w}_N - \hat{w}_B}{\hat{w}_G - \hat{w}_N} = \frac{w_N^* - w_B^*}{w_G^* - w_N^* - \delta}$$

By continuity, there exists a $\bar{\delta} > 0$ such that

$$\frac{w_N^* - w_B^*}{w_G^* - w_N^*} < \frac{\hat{w}_N - \hat{w}_B}{\hat{w}_G - \hat{w}_N} < \frac{r_N - r_B}{r_G - r_N}$$

for all $\delta \in (0, \bar{\delta})$. By Proposition 1, $s^*(w^*, t) < s^*(\hat{w}(\delta, \varepsilon), t) < s^*(r, t)$ for all $\delta \in (0, \bar{\delta})$. For all such δ , $G(\hat{w}(\delta, \varepsilon), t) > G(w^*, t)$, irrespective of $\varepsilon \geq 0$ (since the signal threshold of $\hat{w}(\delta, \varepsilon)$ is closer to the optimal signal threshold). In other words, the expected value of the sum of payments to the two parties is bigger under contract $\hat{w}(\delta, \varepsilon)$ than under contract w^* . For $\varepsilon = 0$ the principal's expected payment under $\hat{w}(\delta, \varepsilon)$ is higher than under contract w^* for all $\delta \in (0, \bar{\delta})$. Take any such δ and now let ε increase until the expected payment to the principal under contract $\hat{w}(\delta, \varepsilon)$ equals that under contract w^* . (This is always possible since the expected payment is continuous and strictly decreasing in ε). Then the expected payment to the agent is positive. By continuity, if ε is reduced slightly, both parties' expected payments under contract $\hat{w}(\delta, \varepsilon)$ is higher than under w^* . Clearly $\hat{w}(\delta, \varepsilon) \in B(w^*)$. Since $G(\hat{w}(\delta, \varepsilon), t) > G(w^*, t)$, $w^* \notin \arg \max_{w \in B(w^*)} G(w, t)$. Since $B(w^*) \subset W$, $w^* \notin \arg \max_{w \in W} G(w, t)$. Q.E.D.

Lemma 1: In the present special case, (7) can be written as

$$\mu e^{-t(s-r_G)^2/2} u(w_G) + (1 - \mu) e^{-t(s-r_B)^2/2} u(w_B) =$$

$$= \left[\mu e^{-t(s-r_G)^2/2} + (1-\mu) e^{-t(s-r_B)^2/2} \right] u(w_N)$$

or

$$\mu e^{-t(s-r_G)^2/2} [u(w_G) - u(w_N)] = (1-\mu) e^{-t(s-r_B)^2/2} [u(w_N) - u(w_B)]$$

or

$$\rho(w) = e^{-t(s-r_B)^2/2+t(s-r_G)^2/2} = e^{\frac{1}{2}t(r_G-r_B)(r_B-2s+r_G)}$$

or:

$$\ln \rho(w) = \frac{1}{2}t(r_G - r_B)(r_G + r_B) - t(r_G - r_B)s$$

or

$$t(r_G - r_B)s = -\ln \rho(w) + \frac{1}{2}t(r_G - r_B)(r_B + r_G)$$

or

$$s = \frac{1}{2}(r_G + r_B) - \frac{\ln \rho(w)}{(r_G - r_B)t}$$

which results in (22). Using this result, one obtains

$$\begin{aligned} p_G(w, t) &= \mu \int_{s^*(w, t)}^{+\infty} \phi_t(s - r_G) ds = \mu \int_{s^*(w, t) - r_G}^{+\infty} \phi_t(x) dx \\ &= \mu \sqrt{\frac{t}{2\pi}} \cdot \int_{s^*(w, t) - r_G}^{+\infty} e^{-tx^2/2} dx = \frac{\mu}{\sqrt{2\pi}} \cdot \int_{(s^*(w, t) - r_G)\sqrt{t}}^{+\infty} e^{-z^2/2} dz \\ &= \mu \left(1 - \Phi_1 \left[(s^*(w, t) - r_G) \sqrt{t} \right] \right) = \mu \Phi_1 \left[(r_G - s^*(w, t)) \sqrt{t} \right] \\ &= \mu \Phi_1 \left[\left(r_G - \frac{1}{2}(r_B + r_G) + \frac{\ln \rho(w)}{(r_G - r_B)t} \right) \sqrt{t} \right] \\ &= \mu \Phi_1 \left[\left(\frac{r_G - r_B}{2} + \frac{\ln \rho(w)}{(r_G - r_B)t} \right) \sqrt{t} \right] \end{aligned}$$

resulting in (23). Similarly,

$$\begin{aligned} p_B(w, t) &= (1-\mu) \int_{s^*(w, t)}^{+\infty} \phi_t(s - r_B) ds = (1-\mu) \Phi_1 \left[(r_B - s^*(w, t)) \sqrt{t} \right] \\ &= (1-\mu) \Phi_1 \left[\left(r_B - \frac{1}{2}(r_B + r_G) + \frac{\ln \rho(w)}{(r_G - r_B)t} \right) \sqrt{t} \right], \end{aligned}$$

resulting in (24). Q.E.D.

Lemma 2: We first write

$$\begin{aligned} A \cdot [\mathbb{E}[U(w, t)] - u(w_N)] &= \rho(w) \Phi_1 \left(\frac{\ln \rho(w)}{(r_G - r_B) \sqrt{t}} + \frac{r_G - r_B}{2} \sqrt{t} \right) \\ &\quad - \Phi_1 \left(\frac{\ln \rho(w)}{(r_G - r_B) \sqrt{t}} - \frac{r_G - r_B}{2} \sqrt{t} \right) \end{aligned}$$

for

$$A = \frac{1}{(1 - \mu) [u(w_N) - u(w_B)]}$$

It follows immediately that $\mathbb{E}[U(w, t)]$ is differentiable in t for any $t > 0$, and

$$\begin{aligned} A \cdot \frac{\partial}{\partial t} \mathbb{E}[U(w, t)] &= \rho(w) \cdot \frac{\partial}{\partial t} \Phi_1 \left(\frac{\ln \rho(w)}{(r_G - r_B) \sqrt{t}} + \frac{r_G - r_B}{2} \sqrt{t} \right) \\ &\quad - \frac{\partial}{\partial t} \Phi_1 \left(\frac{\ln \rho(w)}{(r_G - r_B) \sqrt{t}} - \frac{r_G - r_B}{2} \sqrt{t} \right) \end{aligned}$$

where

$$\begin{aligned} &\frac{\partial}{\partial t} \Phi_1 \left(\frac{\ln \rho(w)}{(r_G - r_B) \sqrt{t}} + \frac{r_G - r_B}{2} \sqrt{t} \right) = \\ &= \frac{\partial}{\partial t} \left[\frac{\ln \rho(w)}{(r_G - r_B) \sqrt{t}} + \frac{r_G - r_B}{2} \sqrt{t} \right] \cdot \phi_1 \left(\frac{\ln \rho(w)}{(r_G - r_B) \sqrt{t}} + \frac{r_G - r_B}{2} \sqrt{t} \right) \\ &= \left[\frac{r_G - r_B}{4\sqrt{t}} - \frac{\ln \rho(w)}{2(r_G - r_B) t \sqrt{t}} \right] \cdot \phi_1 \left(\frac{\ln \rho(w)}{(r_G - r_B) \sqrt{t}} + \frac{r_G - r_B}{2} \sqrt{t} \right) \end{aligned}$$

and

$$\begin{aligned} &\frac{\partial}{\partial t} \Phi_1 \left(\frac{\ln \rho(w)}{(r_G - r_B) \sqrt{t}} - \frac{r_G - r_B}{2} \sqrt{t} \right) = \\ &= \frac{\partial}{\partial t} \left[\frac{\ln \rho(w)}{(r_G - r_B) \sqrt{t}} - \frac{r_G - r_B}{2} \sqrt{t} \right] \cdot \phi_1 \left(\frac{\ln \rho(w)}{(r_G - r_B) \sqrt{t}} - \frac{r_G - r_B}{2} \sqrt{t} \right) \\ &= - \left[\frac{r_G - r_B}{4\sqrt{t}} + \frac{\ln \rho(w)}{2(r_G - r_B) t \sqrt{t}} \right] \cdot \phi_1 \left(\frac{\ln \rho(w)}{(r_G - r_B) \sqrt{t}} - \frac{r_G - r_B}{2} \sqrt{t} \right) \end{aligned}$$

Hence,

$$\begin{aligned} A \cdot \frac{\partial}{\partial t} \mathbb{E}[U(w, t)] &= \\ &= \rho(w) \cdot \left[\frac{r_G - r_B}{4\sqrt{t}} - \frac{\ln \rho(w)}{2(r_G - r_B) t \sqrt{t}} \right] \cdot \phi_1 \left(\frac{\ln \rho(w)}{(r_G - r_B) \sqrt{t}} + \frac{r_G - r_B}{2} \sqrt{t} \right) \\ &\quad + \left[\frac{r_G - r_B}{4\sqrt{t}} + \frac{\ln \rho(w)}{2(r_G - r_B) t \sqrt{t}} \right] \cdot \phi_1 \left(\frac{\ln \rho(w)}{(r_G - r_B) \sqrt{t}} - \frac{r_G - r_B}{2} \sqrt{t} \right), \end{aligned}$$

or (since $\exp[-(a+b)^2/2] = \exp[-(a^2+b^2)/2] \cdot \exp(-ab)$ and $\exp[-(a-b)^2/2] = \exp[-(a^2+b^2)/2] \cdot \exp(ab)$):

$$A\sqrt{2\pi} \cdot \frac{\partial}{\partial t} \mathbb{E}[U(w, t)] \cdot \exp \left[\frac{1}{2} \left(\frac{\ln \rho(w)}{(r_G - r_B) \sqrt{t}} \right)^2 + \frac{1}{2} \left(\frac{r_G - r_B}{2} \sqrt{t} \right)^2 \right] =$$

$$\begin{aligned}
 &= \rho(w) \cdot \left[\frac{r_G - r_B}{4\sqrt{t}} - \frac{\ln \rho(w)}{2(r_G - r_B)t\sqrt{t}} \right] \cdot \exp\left(-\frac{\ln \rho(w)}{2}\right) \\
 &\quad + \left[\frac{r_G - r_B}{4\sqrt{t}} + \frac{\ln \rho(w)}{2(r_G - r_B)t\sqrt{t}} \right] \cdot \exp\left(\frac{\ln \rho(w)}{2}\right) \\
 &= \rho(w) \cdot \left[\frac{r_G - r_B}{4\sqrt{t}} - \frac{\ln \rho(w)}{2(r_G - r_B)t\sqrt{t}} \right] \cdot \frac{1}{\sqrt{\rho(w)}} \\
 &\quad + \left[\frac{r_G - r_B}{4\sqrt{t}} + \frac{\ln \rho(w)}{2(r_G - r_B)t\sqrt{t}} \right] \cdot \sqrt{\rho(w)} \\
 &= \frac{r_G - r_B}{2} \cdot \sqrt{\frac{\rho(w)}{t}}.
 \end{aligned}$$

or

$$\begin{aligned}
 \frac{\partial}{\partial t} \mathbb{E}[U(w, t)] &= \frac{r_G - r_B}{2A} \cdot \sqrt{\frac{\rho(w)}{2\pi t}} \\
 &\quad \cdot \exp \left[-\frac{1}{2} \left(\frac{\ln \rho(w)}{(r_G - r_B)\sqrt{t}} \right)^2 - \frac{1}{2} \left(\frac{r_G - r_B}{2} \sqrt{t} \right)^2 \right]
 \end{aligned}$$

This establishes the claimed formula. Q.E.D.

Lemma 3: For any given contract $w \in W$, let

$$\begin{aligned}
 D(t) &= \frac{\partial}{\partial t} \mathbb{E}[U(w, t)] = \frac{r_G - r_B}{2} \sqrt{\frac{\kappa(w)}{2\pi t}} \\
 &\quad \cdot \exp \left[-\frac{1}{2t} \left(\frac{\ln \rho(w)}{r_G - r_B} \right)^2 - \frac{t}{2} \left(\frac{r_G - r_B}{2} \right)^2 \right]
 \end{aligned}$$

Thus $D(t)$ is the left-hand side of (32). We wish to characterize situations in which the maximum of $D(t)$ exceeds c/a . For notational convenience, first write

$$D(t) = \frac{\gamma}{\sqrt{t}} e^{-\alpha/t - \beta t}$$

for

$$\alpha = \frac{1}{2} \left(\frac{\ln \rho(w)}{r_G - r_B} \right)^2 > 0, \beta = \frac{1}{2} \left(\frac{r_G - r_B}{2} \right)^2 > 0 \text{ and } \gamma = \frac{r_G - r_B}{2} \sqrt{\frac{\kappa(w)}{2\pi}} > 0$$

Then $D'(t)$ has the same sign as

$$-\frac{\partial}{\partial t} \left[\sqrt{t} \cdot \exp \left(\frac{\alpha}{t} + \beta t \right) \right] = -\sqrt{t} \cdot \left(\frac{1}{2t} + \beta - \frac{\alpha}{t^2} \right) \cdot \exp \left(\frac{\alpha}{t} + \beta t \right)$$

For $t > 0$, $D'(t) = 0$ if $2\beta t^2 + t - 2\alpha = 0$, or equivalently (for positive t), $t = T$ where

$$T = \frac{1}{4\beta} \left(\sqrt{16\alpha\beta + 1} - 1 \right)$$

This is the t -value at which the right-hand side in (32) is maximal. Substituting back α and β :

$$T = \frac{2}{(r_G - r_B)^2} \left(\sqrt{[\ln \rho(w)]^2 + 1} - 1 \right).$$

Thus

$$\begin{aligned} D(T) &= \frac{r_G - r_B}{2} \sqrt{\frac{\kappa(w)}{2\pi T}} \\ &\quad \cdot \exp \left[-\frac{1}{2T} \left(\frac{\ln \rho(w)}{r_G - r_B} \right)^2 - \frac{T}{2} \left(\frac{r_G - r_B}{2} \right)^2 \right] \end{aligned}$$

The exponent can be simplified:

$$\begin{aligned} & -\frac{1}{2T} \left(\frac{\ln \rho(w)}{r_G - r_B} \right)^2 - \frac{T}{2} \left(\frac{r_G - r_B}{2} \right)^2 \\ &= -\frac{[\ln \rho(w)]^2}{4 \left(\sqrt{[\ln \rho(w)]^2 + 1} - 1 \right)} - \frac{\left(\sqrt{[\ln \rho(w)]^2 + 1} - 1 \right)}{4} \\ &= \frac{-1}{4 \left(\sqrt{[\ln \rho(w)]^2 + 1} - 1 \right)} \cdot \left([\ln \rho(w)]^2 + \left(\sqrt{[\ln \rho(w)]^2 + 1} - 1 \right)^2 \right) \\ &= \frac{-1}{4 \left(\sqrt{[\ln \rho(w)]^2 + 1} - 1 \right)} \cdot \left([\ln \rho(w)]^2 + [\ln \rho(w)]^2 + 1 - 2\sqrt{[\ln \rho(w)]^2 + 1} + 1 \right) \\ &= \frac{-1}{2 \left(\sqrt{[\ln \rho(w)]^2 + 1} - 1 \right)} \cdot \left([\ln \rho(w)]^2 + 1 - \sqrt{[\ln \rho(w)]^2 + 1} \right) \\ &= \frac{-\sqrt{[\ln \rho(w)]^2 + 1}}{2 \left(\sqrt{[\ln \rho(w)]^2 + 1} - 1 \right)} \cdot \left(\sqrt{[\ln \rho(w)]^2 + 1} - 1 \right) = -\frac{1}{2} \sqrt{[\ln \rho(w)]^2 + 1} \end{aligned}$$

Hence (33), and (32) has two solutions if and only if (33) holds. Moreover, the larger of these solution exceeds T , which gives (34). Q.E.D.

Proposition 6: In order to establish (37), assume that w is a feasible and undominated contract with $w_B > 0$. Being an interior point in W_0 , and using the observations in the proof of Proposition 2, w then maximizes $G(w, t^o(w, a)) - C(t^o(w, a), a)$. The signal threshold, given the precision $t = t^o(w, a)$ need to be optimal, which, by Proposition 2 requires $\rho(w) = \rho(r)$ for a risk neutral agent. Moreover, when the agent chooses his signal precision, given w , his choice $t = t^o(w^*, a)$ necessarily satisfies (32) for $w = r$. Since $\rho(w) = \rho(r)$, it is also necessary that $\kappa(w) = \kappa(r)$. In sum, w has to solve the following two equations:

$$(w_G - w_N)(r_N - r_B) = (w_N - w_B)(r_G - r_N) \quad (50)$$

and

$$(w_G - w_N)(w_N - w_B) = (r_G - r_N)(r_N - r_B) \quad (51)$$

Division of each side gives $(r_N - r_B)^2 = (w_N - w_B)^2$. Hence, $r_B - w_B = r_N - w_N$. Likewise, multiplication of each side gives $(w_G - w_N)^2 = (r_G - r_N)^2$ and hence $r_G - w_G = r_N - w_N$. In other words, the principal earns the same amount in each of the three outcomes. Since this amount cannot be negative, by the principal's participation constraint, we necessarily have $r_B > w_B$, which implies $w_B < 0$, thus violating the limited liability constraint. Hence, no contract w in the interior of W_0 is undominated. Hence, $w_B = 0$ is necessary. Q.E.D.

Proposition 7: First, we identify a contract w^* with such that $\rho(w^*) = 1/\rho(w)$ and $\kappa(w^*) = \kappa(w)$:

$$w_G^* = \frac{\mu}{1-\mu}(w_G - w_N) + \frac{1-\mu}{\mu}w_N, \quad w_N^* = \frac{\mu}{1-\mu}(w_G - w_N) \quad \text{and} \quad w_B^* = 0.$$

Then

$$\rho(w^*) = \frac{\mu}{1-\mu} \cdot \frac{w_G^* - w_N^*}{w_N^*} = \frac{1-\mu}{\mu} \cdot \frac{w_N}{w_G - w_N} = \frac{1}{\rho(w)}$$

and

$$\kappa(w^*) = \mu(1-\mu)(w_G^* - w_N^*)w_N^* = \kappa(w)$$

Hence, the necessary FOC (30) for the agent's choice of signal precision is the same under both contracts (irrespective of the agent's cost function for information acquisition) see (29). Hence, the agent will choose the same signal precision under both contracts, $t^o(w^*, a) = t^o(w, a)$ for all $a > 0$. Since the agent's information cost is the

same in both contracts, it remains to prove that the expected *total return* under w^* is higher than under w . For this purpose, we first note that

$$s^*(w, t) = \frac{r_G + r_B}{2} - \frac{\ln \rho(w)}{(r_G - r_B)t} < s^*(r, t) = \frac{r_G + r_B}{2} - \frac{\ln \rho(r)}{(r_G - r_B)t}$$

and

$$s^*(w^*, t) = \frac{r_G + r_B}{2} + \frac{\ln \rho(w)}{(r_G - r_B)t}.$$

In particular, $s^*(w, t) < s^*(w^*, t)$. There are two cases to consider.

Case I: $s^*(w^*, t) \leq s^*(r, t)$. Then the expected total return (the principal's and agent's added together) is higher under w^* than under w , since the threshold under w^* is still lower than first-best, but deviates less from first-best than w , and the expected total return is strictly increasing in the signal threshold when this is below its first-best optimum value. Hence, w^* results in a higher expected total return.

Case II: $s^*(w^*, t) > s^*(r, t)$. Then the expected total return is again higher under w^* than under w . The reason is that w^* induces an upward deviation from the first-best optimum that is smaller than the downward deviation induced by w :

$$s^*(r, t) - s^*(w, t) = \frac{\ln \rho(w) - \ln \rho(r)}{(r_G - r_B)t} > \frac{\ln \rho(w) + \ln \rho(r)}{(r_G - r_B)t} = s^*(w^*, t) - s^*(r, t).$$

Moreover, as is shown below, the expected total return, as a function of the signal threshold, is asymmetric around its optimum point, $s^*(r, t)$, decreasing faster to the left than to the right. Consequently, the expected total return under w^* is higher than under w .

It thus remains to analyze the expected total return from the asset when the agent uses an arbitrary signal threshold x (at a fixed signal precision t), which we write as

$$\begin{aligned} R(x) &= r_N + \mu(r_G - r_N) \Pr(s > x \mid G) - (1 - \mu)(r_N - r_B) \Pr(s > x \mid B) \\ &= r_N + \mu(r_G - r_N) \int_{x-r_G}^{+\infty} \phi_t(z) dz - (1 - \mu)(r_N - r_B) \int_{x-r_B}^{+\infty} \phi_t(z) dz \end{aligned}$$

We have

$$\begin{aligned} R'(x) &= -\mu(r_G - r_N) \phi_t(x - r_G) + (1 - \mu)(r_N - r_B) \phi_t(x - r_B) \\ &= (1 - \mu)(r_N - r_B) \cdot [\phi_t(x - r_B) - \rho(r) \phi_t(x - r_G)], \end{aligned}$$

and we know that $R'(x)$ vanishes at $x = s^*(r, t)$, is positive when $x < s^*(r, t)$ and negative when $x > s^*(r, t)$. Writing $z \in \mathbb{R}$ for the deviation from the optimal signal threshold,

$$x = z + s^*(r, t) = z + \frac{r_G + r_B}{2} - \frac{\ln \rho(r)}{(r_G - r_B)t}$$

we obtain

$$\begin{aligned} R'(z + s^*(r, t)) &\propto \exp \left[-\frac{t}{2} \left(z + \frac{r_G - r_B}{2} - \frac{\ln \rho(r)}{(r_G - r_B)t} \right)^2 \right] \\ &\quad - \rho(r) \exp \left[-\frac{t}{2} \left(z - \frac{r_G - r_B}{2} - \frac{\ln \rho(r)}{(r_G - r_B)t} \right)^2 \right] \\ &= \exp(A) - \exp(B) \end{aligned}$$

Where

$$\begin{aligned} A &= -\frac{t}{2} \left[z^2 + 2z \left(\frac{r_G - r_B}{2} - \frac{\ln \rho(r)}{(r_G - r_B)t} \right) + \left(\frac{r_G - r_B}{2} - \frac{\ln \rho(r)}{(r_G - r_B)t} \right)^2 \right] \\ &= -\frac{t}{2} z^2 - z \left(t \frac{r_G - r_B}{2} - \frac{\ln \rho(r)}{r_G - r_B} \right) - \frac{t}{2} \left(\frac{r_G - r_B}{2} \right)^2 + \frac{1}{2} \ln \rho(r) - \left(\frac{\ln \rho(r)}{r_G - r_B} \right)^2 \end{aligned}$$

and

$$\begin{aligned} B &= \ln \rho(r) - \frac{t}{2} \left[z^2 - 2z \left(\frac{r_G - r_B}{2} + \frac{\ln \rho(r)}{(r_G - r_B)t} \right) + \left(\frac{r_G - r_B}{2} + \frac{\ln \rho(r)}{(r_G - r_B)t} \right)^2 \right] \\ &= \ln \rho(r) - \frac{t}{2} z^2 + z \left(t \frac{r_G - r_B}{2} + \frac{\ln \rho(r)}{r_G - r_B} \right) - \frac{t}{2} \left(\frac{r_G - r_B}{2} \right)^2 - \frac{1}{2} \ln \rho(r) - \left(\frac{\ln \rho(r)}{r_G - r_B} \right)^2 \\ &= -\frac{t}{2} z^2 + z \left(t \frac{r_G - r_B}{2} + \frac{\ln \rho(r)}{r_G - r_B} \right) - \frac{t}{2} \left(\frac{r_G - r_B}{2} \right)^2 + \frac{1}{2} \ln \rho(r) - \left(\frac{\ln \rho(r)}{r_G - r_B} \right)^2 \end{aligned}$$

Hence,

$$\begin{aligned} \frac{R'(s^*(r, t) + z)}{R'(s^*(r, t) - z)} &= \frac{\exp \left[-z \left(t \frac{r_G - r_B}{2} - \frac{\ln \rho(r)}{r_G - r_B} \right) \right] - \exp \left[z \left(t \frac{r_G - r_B}{2} + \frac{\ln \rho(r)}{r_G - r_B} \right) \right]}{\exp \left[z \left(t \frac{r_G - r_B}{2} - \frac{\ln \rho(r)}{r_G - r_B} \right) \right] - \exp \left[-z \left(t \frac{r_G - r_B}{2} + \frac{\ln \rho(r)}{r_G - r_B} \right) \right]} \\ &= \frac{\exp \left[-zt \frac{r_G - r_B}{2} \right] - \exp \left[zt \frac{r_G - r_B}{2} \right]}{\exp \left[zt \frac{r_G - r_B}{2} \right] - \exp \left[-zt \frac{r_G - r_B}{2} \right]} \cdot \exp \left[\frac{2 \ln \rho(r)}{r_G - r_B} \right] \end{aligned}$$

For $z > 0$, the nominator is negative and denominator positive. Hence, for such z :

$$\frac{|R'(s^*(r, t) + z)|}{R'(s^*(r, t) - z)} = \exp \left[\frac{2 \ln \rho(r)}{r_G - r_B} \right]$$

By hypothesis, $\rho(r) < 1$, so

$$\frac{|R'(s^*(r, t) + z)|}{R'(s^*(r, t) - z)} < 1$$

for all $z > 0$. This establishes that the expected total return declines faster when the signal threshold moves to the left than when it moves to the right. Consequently, $R[s^*(w, t)] < R[s^*(w^*, t)]$. Q.E.D.

Proposition 10: We here prove a slightly more general result, allowing the agent to also have a bias in his information gathering. The principal will choose the contract that induces the agent to use the signal cut-off that maximizes the gross return to the asset, given the principal's belief that the agent's signal has noise $\varepsilon \sim N(\delta, 1/t)$ for some $\delta \in \mathbb{R}$. From Proposition 1 and Lemma 1 we deduce that the optimal signal-threshold then is

$$\hat{s} = \frac{r_G + r_B}{2} - \frac{\rho(r)}{t \cdot (r_G - r_B)} + \delta$$

By contrast, given the agent's (distorted) belief about his noise distribution, he will use the following signal threshold under any contract $w = (0, w_N, w_G)$, again using Proposition 1:

$$\tilde{s} = \frac{r_G + r_B}{2} - \frac{\rho(w)}{(1 + \theta)t \cdot (r_G - r_B)}$$

We find the bonus rate, $b = w_G/w_N - 1$, by setting $\tilde{s} = \hat{s}$:

$$\rho(w) = (1 + \theta) [\rho(r) - t(r_G - r_B)\delta]$$

Hence, the principal will set

$$b = \frac{w_G}{w_N} - 1 = \rho(r)^\theta \cdot e^{-(1+\theta)t \cdot (r_G - r_B) \cdot \delta}$$

The contract will thus be of the form $\hat{w} = (0, y, y + by)$ where y is such that the agent believes, *ex ante*, that his expected remuneration from the contract equates his outside option:

$$[1 - p_B(\hat{w}, (1 + \theta)t)] \cdot y + p_G(\hat{w}, (1 + \theta)t) \cdot b \cdot y = \bar{u}$$

or

$$y = \frac{\bar{u}}{1 - p_B(\hat{w}, (1 + \theta)t) + b \cdot p_G(\hat{w}, (1 + \theta)t)}$$

Under this contract \hat{w} , the risk-neutral agent's signal cut-off will be $\tilde{s} = s^*(r, t) + \delta$. Q.E.D.

REFERENCES

- [1] Axelson, U. and P. Bond (2015): “Wall Street occupations”, *Journal of Finance* 70, 1949-1996.
- [2] Axelson, U. and J. Weibull (2014): “Informed lending and credit crashes”, mimeo. London School of Economics and Stockholm School of Economics.
- [3] Bakija, J., A. Cole and B. Heim (2012): “Jobs and income growth of top earners and the causes of changing income inequality: Evidence from U.S. tax returns data”, Working paper, Indiana University.
- [4] Bebchuk, L. and J. Fried (2003): “Executive compensation as an agency problem”, *Journal of Economic Perspectives* 17, 71-92.
- [5] Bebchuk, L. and J. Fried (2004): *Pay without Performance: The Unfulfilled Promise of Executive Compensation*. Harvard University Press.
- [6] Bénabou, R., and J. Tirole (2013): “Bonus culture: Competitive pay, screening, and multitasking”, CEPR Discussion Paper 9419.
- [7] Bertrand, M. and S. Mullainathan (2001): “Are CEOs rewarded for luck? The ones without principals are”, *Quarterly Journal of Economics* 116, 901-932.
- [8] Carlin, B. and S. Gervais (2009) “Work ethic, employment contracts and firm value”, *Journal of Finance* 64, 785-821.
- [9] Chade H. and E. Schlee (2002): “Another look at the Radner-Stiglitz nonconcavity in the value of information”, *Journal of Economic Theory* 107, 421-452.
- [10] Cronqvist, H. and R. Fahlenbach (2013): “CEO contract design: How do strong principals do it?”, *Journal of Financial Economics* 108, 659-674.
- [11] Custodio, C., M. Ferreira, and P. Matos (2013): “Generalists versus specialists: Lifetime work experience and CEO pay”, *Journal of Financial Economics* 108, 471-492.
- [12] Daniel, K. and D. Hirshleifer (2015): “Overconfident investors, predictable returns, and excessive trading”, *Journal of Economic Perspectives* 29, 61-87.
- [13] DeMarzo, P., and M. Fishman (2007): “Agency and optimal investment dynamics”, *Review of Financial Studies* 20, 151-188.
- [14] DeMarzo, P., M. Fishman, Z. He, and N. Wang (2012): “Dynamic agency and the Q theory of investment”, *Journal of Finance* 67, 2295-2340.

- [15] Dow, J. and C. Raposo (2005): “CEO compensation, change and corporate strategy”, *Journal of Finance* 60, 2701-2727.
- [16] Edmans, A., X. Gabaix, T. Sadzik, and Y. Sannikov (2012): “Dynamic CEO Compensation”, *Journal of Finance* 67, 1603–1647
- [17] Edmans, A., Gabaix, X. and A. Landier (2009): “A multiplicative model of optimal CEO incentives in market equilibrium”, *Review of Financial Studies* 22, 4881-4917.
- [18] Edmans, A., and X. Gabaix (2009): “Is CEO pay really inefficient? A survey of new optimal contracting theories”, *European Financial Management* 15, 486-496.
- [19] Edmans, A. and X. Gabaix (2011): “The effects of risk on the CEO market”, *Review of Financial Studies* 24, 2822-2863.
- [20] Gabaix X, and A. Landier (2008): “Why has CEO pay increased so much?”, *Quarterly Journal of Economics* 12, 49-100.
- [21] Garicano, L. and E. Rossi-Hansberg (2006): “Organization and inequality in a knowledge economy”, *Quarterly Journal of Economics* 121, 1383-1435.
- [22] Gervais, S., J. B. Heaton, and T. Odean (2011): “Overconfidence, compensation, contracts and capital nudgiting”, *Journal of Finance* 66, 1735-1777.
- [23] Giannetti, M. (2011): “Serial CEO incentives and the structure of managerial contracts”, *Journal of Financial Intermediation* 20, 633-662.
- [24] Hall, B. and K. Murphy (2003): “The trouble with stock options”, *Journal of Economic Perspectives* 17, 49-70.
- [25] Heller, Y. (2014): “Overconfidence and diversification”, *American Economic Journal: Microeconomics* 6, 134-153.
- [26] Holmström, B. and S. Kaplan (2001): “Corporate governance and merger activity in the U.S.”, *Journal of Economic Perspectives* 15, 121-144.
- [27] Holmström, B. and S. Kaplan (2003): “The state of US corporate governance: What’s right and what’s wrong?”, *Journal of Applied Corporate Finance* 15, 8-20.
- [28] Holmström, B. (2004): “Pay without performance and the managerial power hypothesis: A comment”, *Journal of Corporation Law* 30, 703-715.

- [29] Holmström, B. and P. Milgrom (1991): “Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design”, *Journal of Law, Economics and Organization* 7, 24-52.
- [30] Jensen, M. and K. Murphy (1990): “Performance pay and top-management incentives”, *Journal of Political Economy* 98, 225-264.
- [31] Jenter, D., and K. Lewellen (2014): “Performance induced CEO turnover”, Research Paper 3054, Graduate School of Business, Stanford University.
- [32] Kahneman, D. (2011): *Thinking, Fast and Slow*. London: Allen Lane.
- [33] Kaplan, S. and P. Strömberg (2002): “Financial contracting theory meets the real world: An empirical analysis of venture capital contracts”, *Review of Economic Studies* 100, 1-35.
- [34] Kaplan, S. and P. Strömberg (2004): “Characteristics, contracts, and actions: Evidence from vernture capitalist analyses”, *Journal of Finance* 59, 2177-2210.
- [35] Kaplan, S. (2013): “CEO pay and corporate governance in the U.S.: Perceptions, facts and challenges”, *Journal of Corporate Finance* 25, 8-25
- [36] Kaplan, S., and B. Minton (2012): “How has CEO turnover changed?”, *International Review of Finance* 12, 57-87.
- [37] Kaplan, S., and J. Rauh (2010): “Wall Street and Main Street: What constitutes to the rise in the highest incomes?”, *Review of Financial Studies* 23, 1004-1050.
- [38] Karlin S., and H. Rubin (1956): “The theory of decision procedures for distributions with monotone likelihood ratio”, *Annals of Mathematical Statistics* 27, 272-299.
- [39] Kindleberger, D. (2005): *Manias, Panics and Crashes: History of Financial Crises*. New York: Basic Books.
- [40] Kosfeld, M. and F. von Siemens (2011): “Competition, cooperation and corporate culture”, *Rand Journal of Economics* 42, 23-43.
- [41] Malcolmson, J. (2010): “Do Mangers with limited liability take more risky decisions? An informatioan acquisition model”, CESifo Working Paper No. 2943, February.
- [42] Malmendier, U. and G. Tate (2005): “Does overconfidence affect corporate investment? CEO overconfidence measures revisited”, *European Financial Managemement* 11, 649-659.

- [43] Malmendier, U. and G. Taylor (2015): “Behavioral CEO:s. The role of managerial overconfidence”, *Journal of Economic Perspectives* 29, 37-60.
- [44] Manso G. (2011): “Motivating innovation”, *Journal of Finance* 66, 1823-1860.
- [45] Mishel, L., and A. Davis (2014): “CEO pay continues to rise as typical workers are paid less”. Issue Brief, Economic Policy Institute, no. 380.
- [46] Milgrom, P. (1981): “Good news and bad news: Some representation results”, *Bell Journal of Economics* 12, 380-391.
- [47] Morgan, J., and P. Stocken (2003): “An analysis of stock recommendations”, *RAND Journal of Economics* 34, 183-220.
- [48] Murphy, K., J., (2012): “Executive Compensation: where we are and how we got there”, in G. Constantinides, M. Harris and R. Stulz (eds.), *Handbook of the Economics of Finance*. Amsterdam: North-Holland.
- [49] Nash, J. (1950), "The bargaining problem", *Econometrica* 18, 155-162.
- [50] Otto C. A., (2014), “CEO optimism and incentive compensation”, *Journal of Financial Economics* 114, 366-404.
- [51] Prendergast, C. (1999): “The provision of incentives in firms”, *Journal of Economic Literature* 37, 7-63.
- [52] Radner R. and J. Stiglitz (1984): “A nonconcavity in the value of information”, pp. 33-52, in M. Boyer and R.E. Kihlstrom (eds.), *Bayesian Models in Economic Theory*. North-Holland.
- [53] Rosen, S. (1981): “The economics of superstars”, *American Economic Review* 71, 845-858.
- [54] Roth, A. (1979): *Axiomatic Models of Bargaining*, Lecture Notes in Economics and Mathematical Systems, 170. Springer Verlag.
- [55] Terviö, M. (2009): “Superstars and mediocrities: Market failure in the discovery of talent”, *Review of Economic Studies* 76, 829-850.
- [56] Weibull J., L.-G. Mattsson and M. Voorneveld (2007): “Better may be worse: some monotonicity results and paradoxes in discrete choice under uncertainty”, *Theory and Decision* 63, 121-141.