Pay schemes, bargaining, and competition for talent

Assar Lindbeck* and Jörgen Weibull†

Detailed version‡

December 21, 2015. Revised July 8, 2016.

Abstract. We present a novel model for analysis of contracts between a principal and an agent when the agent's task is to make a well-informed investment decision on behalf of the principal, with managers in large corporations as the most prominent example. The principal and agent initially bargain over the contract, and the bargaining outcome depends on both parties' outside options and relative bargaining power. Having signed a contract, the agent chooses how much effort to make in order to acquire information about the investment project at hand. This effort is non-contractable and the information obtained is private. The agent then decides whether or not to invest in the project in the light of his obtained information. We study the nature of such bargained contracts, in particular the role of bonuses and/or penalties. We show how all agreed-upon contracts balance the incentives for information gathering with those for the investment decision, a balance that may be represented in terms of what we call the contract's "carrot-stick ratio" and the contract's "power". In particular, no contract that consists of a fixed salary plus stocks in the project or company is ever agreed upon. Such contract provide insufficient incentive for the agent's information acquistion. By contrast, contracts that contain a balanced mix of fixed salary, stocks, and options are potential equilibrium contracts. We also briefly analyze the effects of behavioral biases such as overconfidence and optimism.

Keywords: Principal-agent, investment, information acquisition, 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, IAST Institute for Advanced Study in Toulouse, and KTH Royal Institute of Technology, Stockholm.

‡The authors thank Ingela Alger, Bo Becker, Jun Chen, Sven-Olof Fridolfson, Jungsuk Han, Bengt Holmström, Augustin Landier, Paul Milgrom, Francois Salanié, Francesco Sangiorgi, Per Strömberg, Thomas Tangerås and Peter Wikman for helpful comments to an earlier version of this paper. We thank Atahan Afsar for excellent research assistance and Kristaps Dzonsons for great 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 Foundatio, and the ANR-Labex IAST 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; Bakhija 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.[1] 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 these 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 these corporations during the last three decades of the 20th century, 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

[1] See Kaplan and Strömberg (2002) for an empirical study of contracts between venture capitalists and entrepreneurs.

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 hypothesis 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.[2] 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) and Edmans and Gabaix (2009).

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 an analytical framework, or "work horse", that can be used, and further developed, for studies of 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 combining elements from the standard principal-agent literature with elements from other literatures, in particular concerning economic agent's information acquisition.

The focus of our study is on the interaction between a *principal*, who has money, and an *agent*, who has talent for investment decisions. The principal may for example be the owner(s) of a corporation and the agent its current CEO or a candidate for this job. The principal faces a binary investment decision, or a sequence of such decisions, such as whether or not to undertake a risky project—say, purchase an asset, buy up another company, or undertake R&D for a certain production technology or product.

[2] In an empirical study, Böhm et al.(2015) find that the remuneration of employees in the financial sector has increased dramatically in recent decades, as compared to other sectors. At the same time they find hardly any evidence that the selection of talent into finance has improved. This suggests that the development cannot be explained as a result of increased competition for talent.

A project's return is random and the principal and the agent have the same prior beliefs about its probability distribution. The *ex post* return from the project is verifiable.

The two parties first enter a negotiation about the terms of a potential contract between them. The principal knows how talented the agent is. 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: successful investment, failed investment, and no investment. We require contracts to meet the "limited liability" constraint that the net pay from the principal to the agent is always nonnegative. 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.[3] 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 guides the agent's investment decision, once his information has been obtained. As will be seen, there is, in general, a tension between these two goals. A contract that is well designed for the first purpose (incentivize information acquisition) may induce the agent to be too "trigger happy", that is, to invest even when his information about the project is not sufficiently favorable or precise for the principal to wish investment to be made. This tension exists even when both parties are risk neutral and fully rational.

In our model, contracts consisting of a fixed salary plus some stocks (in the company or project at hand) are in general dominated by other contracts (available in the bargaining), since they provide too weak incentives for the agent's information acquisition. Contracts consisting of a fix salary plus an option to buy some future stocks at their current price are always dominated, since they provide no incentives for information acquisition and makes the agent invest irrespective of his private information. However, contract that consist of a fixed salary, some stocks and some options, in a correct balance, may be undominated and hence potential equilibrium outcomes of the bargaining.

To be more specific, we apply the generalized Nash-bargaining solution (Nash, 1950; Roth, 1979) to the initial contract negotiation stage between the principal and agent. Once a contract has been signed, 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 at or above her outside option,

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

and in expected utility to the agent at or above his outside option.[4] 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.[5] All other feasible contracts could be renegotiated to the advantage of at least one party. The generalized Nash bargaining solution selects an undominated contract and this selection depends both on the parties' outside options and on their individual bargaining power. The agent's outside option represents the intensity of competition for talent in the job market. The stiffer competition there is, the better is the agent's outside option.

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 dictate the conditions 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 over an outside candidate for the CEO position, while 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. In addition, there may also be considerable reputational cost for a company of firing its CEO.

This kind of principal-agent model is, to the best of our knowledge, novel. Moreover, it delivers new insights. Since the contract not only influences the agent's investment decision but also how well-informed the agent will choose to be, this is a model with endogenous asymmetric information, or "rational inattention". As shown in Weibull, Mattsson and Voorneveld (2007), such endogenous uncertainty in decision-making canonically leads to a non-convexity with ensuing potential discontinuity in the decision-maker's behavior with respect to the parameters defining the decision problem.[6] In the present model, such non-convexity induces "jumps" in the agent's information-acquisition behavior as reactions to gradual changes in the contract, project or investment environment. Intermediate levels of effort for information acquisition are typically not worthwhile for the decision maker in question.

Our model is schematically 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 $\boldsymbol{w} = (w_B, w_N, w_G)$, which specifies the agent's remuneration in each eventuality. The principal's net return (after having paid the agent) is denoted π_G, π_B and

[4] The outside option may be better for more talented agents, see Section 4.3.

[5] However, all feasible contracts are Pareto inefficient except when the agent has all bargaining power.

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

π_N, respectively. In the box right after a contract has been signed, the agent 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, 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.

By *bonus* we mean the additional pay after a successful investment, as compared with the remuneration when no investment was made (that is, the difference $w_G - w_N$). Likewise, by *penalty* we mean the reduction in the agent's remuneration after a failed investment, again as compared with his payment when no investment was made (that is, the difference $w_N - w_B$).[7]

In the special case when the agent is risk neutral and has a fixed and given signal precision, a whole range of positive bonuses and penalties are possible in equilibrium. The only conditions are the two participation constraints and the requirement that the bonus and penalty are balanced in a precise way. The ratio between these has to match the ratio between the underlying project's "upside" and "downside", defined below in terms of what we call the *carrot-stick ratio* in a contract. However, also "flat pay" contracts, i.e. contracts with neither bonus nor penalty, are possible in equilibrium in this case. Indeed, if the agent is even the slightest risk averse, only a flat salary is consistent with equilibrium. The principal, being risk neutral, will then absorb all risk—in line with standard results. The special case of exogenous signal precision thus lends no support for bonuses or penalties.

By contrast, when the agent's signal precision is fully flexible and chosen by the agent, as it is in our model, all equilibrium contracts include both a bonus and a penalty.[8] 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

[7] An 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.

[8] An canoncial example are stocks whose value depend on the agent's performance.

would be better off without the agent. Indeed, the agent's switch from active information gathering (positive effort bounded away from zero) to no information gathering at all (zero effort) may be discontinuous in his bonus and/or penalty and in the prevailing market conditions. Moreover, the limited-liability constraint (the lower bound on pay to the agent), is binding in all equilibrium contracts when the agent is risk neutral. The agent will then be paid nothing after a failed investment. In other words, his penalty will be maximal (loss of all income). This is a familiar property other principal-agent models but turns out to be non-trivial in the present model since it allows for the possibility that the agent to have bargaining power.

Another feature of our model 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, not even when the agent is risk-neutral. The reason is that such contracts may induce the agent, if hired, to gather no information at all. It may thus be in both parties' interest to pay the agent more. By way of numerical simulations, we show how large the equilibrium bonus rate is and how this depends on the primitives of the model.

We find that the effect of increased competition for talent upon the negotiated contract is qualitatively the same as the effect of an agent's increased bargaining power (for given outside options).[9] Moreover, while these factors influence the size of the bonus in the expected direction, the bonus rate is virtually unaffected. We also analyze conventional stock- or option-based contracts. Neither purely stock-based contracts nor purely option-based contracts are compatible with equilibrium. Stock contracts do not give the agent sufficient incentive to acquire information, and option-based contracts do not have any down-side for the agent and thus induce the agent to invest recklessly (and hence not acquire any information since this is of no use for him). However, combinations of a fixed salary with some stocks and some options, with the right balance, are potential equilibrium contracts.

The organization of the paper is as follows. Section 2 presents the model, which is then analyzed in Sections 3 and 4. In section 5 we discuss stock-based and options-based incentive contracts. Section 6 analyzes the effect of such behavioral biases as overconfidence and optimism. Section 7 discusses related literature and Section 8 concludes. Mathematical proofs are given in an appendix at the end of the paper.

2. Model

A risk-neutral principal has the opportunity to undertake a risky project that requires a lump-sum investment $I > 0$. The project yields a risky return, represented by a

[9] By contrast, the expected gains from trade are in general affected by bargaining power, even when both are risk neutral.

random variable Y. The project's net return is thus $Y - I$ and its *net return rate* is $X = (Y - I)/I$. The principal's outside option, in case she does not invest in the project and does not hire the agent, has expected net return rate $r_0 \geq 0$. In this study, we focus on the case of a binary return distribution; the project either yields a certain positive return, above the principal's return from her outside option, or it yields a negative net return. Formally, either $Y = R$ or $Y = I - L$, where $(R - I)/I > r_0$ and $L > 0$. Writing r_G for $(R - I)/I$ and r_B for $-L/I$ we thus have that either $X = r_G > r_0 \geq 0$ or $X = r_B < 0$. Let μ denote the probability for the first event; $\mu = \Pr(X = r_G)$ (and thus $\Pr(X = r_B) = 1 - \mu$). We will call μ the project's *prior* success probability and assume that it lies strictly between zero and one and is known by the principal. We will sometimes write θ for the success probability *ratio* $\mu/(1 - \mu)$.

We will here focus on the case when the *ex ante* expected return from the project falls short of the return from the principal's outside option,[10]

$$\mathbb{E}(X) = \mu r_G + (1 - \mu) r_B < r_0. \tag{1}$$

This is where the agent comes into the picture. At the outset the agent has the same information about the project as the principal, and thus shares the principal's prior belief μ. If employed by the principal, the agent is able to acquire relevant information about the project at hand. However, this requires effort on behalf of the agent. Both his effort and acquired information are unverifiable. Formally, the agent, if employed, will make some non-negative effort, where zero effort results in no more information than the one initially shared with the principal, but positive effort results in a noisy *private signal* S to the agent about the project,

$$S = X + \varepsilon, \tag{2}$$

for some *noise* ε that is statistically independent from X.[11] The noise term ε has a positive and continuous probability density ϕ_τ on $\mathbb{R}$, with cumulative distribution Φ_τ. The parameter τ, a function of the agent's effort to acquire information, belongs to some nonempty and closed set $T \subseteq \mathbb{R}_+$, and will be called the *precision* of the agent's signal. Canonical cases in the literature on asymmetric information are (i) $T = \{\tau\}$ for some $\tau > 0$ (a private signal of exogenous positive precision), (ii)

[10] The case when the ex ante value of the project exceeds the outside option can be analyzed in the same way. One then has to replace r_0 by $\mathbb{E}(X)$ in the subsequent analysis (or, more generally, replace r_0 by $\max\{r_0, \mathbb{E}(X)\}$). The general case can easily be analyzed but at the expense of more cumbersome notation.

[11] We thus assume that the agent receives noisy information about the return rate, rather than about the return, or about the state of nature. All these specifications are mathematically equivalent for a given project. The current specification is, arguably, the most natural when making comparisons across projects.

$T = \{\tau_L, \tau_H\}$ for $0 \leq \tau_L < \tau_H$ (an agent choosing between a given low and a given high signal precision), and (iii) $T = \mathbb{R}_+$ (an agent who can choose any non-negative signal precision). The main part of our analysis will focus on case (iii), since this allows the agent to continuously adjust his information-acquistion effort and allows us to highlight his marginal trade-offs. Unless otherwise stated, case (iii) is presumed.

We assume that the following version of the monotone likelihood-ratio property holds for the agent's signal:[12]

> [**MLRP**] *For any given* $\tau > 0$, *the density ratio* $\psi_\tau(s) = \phi_\tau(s - r_G)/\phi_\tau(s - r_B)$ *is strictly increasing in* $s \in \mathbb{R}$.

As will be seen presently, this property ensures that the high return rate, r_G, is more likely the higher is the observed signal value, and this is true at any given precision level of the signal. The agent's cost or disutility of exerting effort to acquire information is assumed to be increasing in the resulting signal precision and decreasing in the agent's talent or ability. More specifically, we write $g(\tau, a) \geq 0$ for this cost, where $a > 0$ is the agent's innate ability. Zero signal precision is costless; $g(0, a) = 0$ for all $a > 0$. We will refer to $\tau = 0$ as "no information acquisition" or "no signal".[13]

The agent's cost of obtaining information is increasing and convex in the precision of the information obtained, and decreasing in his ability. To be more specific, we will assume that g is twice differentiable with $\partial g(\tau, a)/\partial\tau \geq 0$, with strict inequality whenever $a, \tau > 0$, $\partial g(\tau, a)/\partial a < 0$, $\partial^2 g(\tau, a)/\partial\tau^2 \geq 0$, and $\partial^2 g(\tau, a)/\partial a\partial\tau < 0$.

Suppose that the agent has chosen a positive signal precision $\tau > 0$ and received his signal $s \in \mathbb{R}$. The agent's *posterior* probability for the positive return is then

$$\Pr(X = r_G \mid S = s) = \frac{\mu\phi_\tau(s - r_G)}{\mu\phi_\tau(s - r_G) + (1 - \mu)\phi_\tau(s - r_B)} = \frac{\theta\psi_\tau(s)}{\theta\psi_\tau(s) + 1}. \quad (3)$$

In force of the MLRP condition, this posterior is strictly increasing in the observed signal value s, for any given positive signal precision τ. For notational convenience we make one more assumption about the noise distribution, namely, that its density, ϕ_τ, has "thin tails" in a precise sense:[14]

> [**TT**] *For every positive signal precision* $\tau \in T$, *the density ratio* $\psi_\tau(s)$ *tends to zero as* $s \to -\infty$ *and to plus infinity as* $s \to +\infty$.

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

[13] And we will verify that the model is continuous at $\tau = 0$.

[14] Condition [TT] is not essential. The analysis can be carried out without this condition at the cost of heavier notation.

As is well-known, the normal distribution satisfies both MLRP and TT. Indeed, our canonical example will be normally distributed noise with mean zero and variance $\sigma^2 = 1/\tau > 0$. Then

$$\phi_\tau(x) = \sqrt{\frac{\tau}{2\pi}} \cdot e^{-\tau x^2/2} \quad \text{and} \quad \psi_\tau(x) = e^{(r_G - r_B)\left(x - \frac{r_G + r_B}{2}\right)\tau} \qquad \forall x \in \mathbb{R}. \tag{4}$$

As a basis for the subsequent analysis and discussion, we first specify a few key ingredients of the model.

2.1. Contracts. A *contract* between the principal and the agent specifies a remuneration for the agent, to be given when the project's return has been realized. This remuneration may depend on whether or not the agent invested, and, if investment was made, on the project's realized net return. We thus treat the (binary) investment decision and the realized net return from investment as verifiable and legally enforceable in court.[15] Under this assumption, any legally enforceable contract can be summarized as a vector $\boldsymbol{w} = (w_B, w_N, w_G) \in \mathbb{R}^3$ that specifies a (positive or negative) remuneration w_B to the agent if he invests and the state of nature turns out to be bad, a (positive or negative) remuneration w_N if the agent does not invest in the project at hand, and a (positive or negative) remuneration w_G if the agent invests and the state of nature turns out to be good. 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:[16]

$$W = \left\{\boldsymbol{w} \in \mathbb{R}^3 : w_B \leq w_N \leq w_G\right\}.$$

We call a contract $\boldsymbol{w} \in W$ *strictly monotonic* if both inequalities are strict ($w_B < w_N < w_G$). We will refer to the payoff difference $w_G - w_N$ as the *bonus* and call $w_N - w_B$ the *penalty*, the reduction of the agent's pay after a failed investment, as compared with the pay when no investment is made. By the *bonus rate* in a contract we will mean its ratio between the bonus and the pay when no investments is undertaken (granted, of course, the latter is positive), and we write $b = (w_G - w_N)/w_N$. Likewise, the *penalty rate* is defined as $p = (w_N - w_B)/w_N$.

Arguably, it is realistic in many situations to exclude contracts that in some state (or states) of the world prescribe a negative remuneration to the agent (or, more generally, a remuneration below some pre-specified level), if not for legal reasons perhaps because the agent does not have enough capital to cover a substantial loss or

[15] By contrast, the would-be return if the investment is not made, is not taken to be verifiable in court.

[16] The assumption about monotonicity will not be binding in equilibrium but faciliates the discussion and analysis.

may have means to hide or shelter his private capital from the principal by transferring ownership of assets to others. We refer to this non-negativity constraint as *limited liability*. Formally, such contracts constitute the subset W_0 of contracts $\boldsymbol{w} \in W$ with $w_B \geq 0$.

The vector $\boldsymbol{r} = (r_B, 0, r_G)$ will be called the project's *return-rate vector*, and will for notational convenience sometimes be written in the form $\boldsymbol{r} = (r_B, r_N, r_G)$, with $r_N = 0$. Any positive multiple λ of this vector is a strictly monotonic contract; $\lambda \boldsymbol{r} \in W$ and $\lambda r_B < \lambda r_N < \lambda r_G$. However, no such contract $\boldsymbol{w} = \lambda \boldsymbol{r}$ meets the limited-liability constraint since $r_B < 0$; $\lambda \boldsymbol{r} \notin W_0$.

2.2. Utilities and profits. The quality of the agent's information depends both on his innate talent or *ability* and on the *effort* he makes to acquire information. An agent who has signed a contract thus faces a two-stage decision, where the first stage is to choose the precision of his private signal and the second is to make the investment decision after he has received his signal (of precision determined by his first choice). Let $u(w_B)$, $u(w_N)$ and $u(w_G)$, respectively, denote the utility the agent obtains from his monetary remunerations in the three possible outcomes under signed contract $\boldsymbol{w} \in W$. The (Bernoulli or von Neumann-Morgenstern) function u is twice differentiable and unbounded with $u' > 0$ and $u'' \leq 0$. The agent's outside option, in case he does not sign a contract with the principal, has expected utility $\bar{u} = u(\bar{w}) \geq 0$, where $\bar{w}$ is his certainty-equivalent reservation wage.

Remark 1. *In the short manuscript version we write* $\ln f$ *for* u. *Hence, then* $f' = f \cdot u'$ *etc.*

Just as the agent's utility is random, so is the profit to the principal. If the agent is hired under contract $\boldsymbol{w} \in W$, then the principal's profit is $\pi_B = I \cdot r_B - w_B$ if investment is made in the bad state, $\pi_N = -w_N$ if no investment is made, and $\pi_G = I \cdot r_G - w_G$ if investment is made in the good state. Let $\bar{\pi}$ denote the principal's expected profit from her outside option, that is, $\bar{\pi} = I \cdot r_0$.

The agent's total utility is defined as the difference between his (random) remuneration utility and (deterministic) disutility (or cost) of effort. Hence, if he chooses signal precision $\tau > 0$, then his utility is

$$V(\boldsymbol{w}, \tau, a) = U(\boldsymbol{w}, \tau) - g(\tau, a),$$

where the random variable $U(\boldsymbol{w}, \tau)$ takes the values $u(w_B)$, $u(w_N)$ and $u(w_G)$, respectively, in the three possible outcomes (investment in the bad state, no investment, investment in the good state), the probabilities of which depend on the signal precision τ and the agent's investment strategy when he obtains his signal (his private information).

The associated random profit to the principal, having hired the agent under contract $\boldsymbol{w} \in W$, and supposing the agent has chosen signal precision $\tau > 0$, is the random variable $\Pi(\boldsymbol{w}, \tau)$ that takes one of the values π_B, π_N, or π_G, with probabilities depending on the agent's signal precision τ and the agent's investment strategy.

Finally, if the agent is hired under a contract $\boldsymbol{w} \in W$ and chooses signal precision $\tau = 0$ (i.e. makes no effort at all to acquire information), then his random utility, $V(\boldsymbol{w}, 0, a)$, is $u(w_N)$ with probability one if he does not invest. If he does invest, his utility is either $u(w_G)$, with probability μ, or $u(w_B)$, with probability $1 - \mu$. The principal's profit, $\Pi(\boldsymbol{w}, 0)$, is π_N if the agent does not invest, while if he invests it is either π_G, with probability μ, or π_B, with probability $1 - \mu$. For any such contract that meets the limited-liability constrain, the principal's expected profit is lower than from her outside option. To see this, note that if the agent does not invest, then she has to pay $w_N \geq 0$ but receive zero return, while if the agent does invest, then, by (**??**),

$$\mathbb{E}[\Pi(\boldsymbol{w}, 0)] < I \cdot \mathbb{E}[X] < Ir_0 = \bar{\pi}.$$

Hence, as one would expect, it is never in the principal's interest to sign a contract that gives the agent no incentive to acquire information.

2.3. Work plans for the agent. Suppose that the principal and agent have signed a contract $\boldsymbol{w} \in W$. A *work plan* for the agent is pair $\xi = (\tau, \vartheta)$, where $\tau \in T$ is a signal precision and $\vartheta : T \times \mathbb{R} \to \{0, 1\}$ an *investment strategy*, that is, a rule that to each signal precision $\tau \in T$ and every possible signal value $s \in \mathbb{R}$ assigns an investment decision $\vartheta(\tau, s) \in \{0, 1\}$, where $\vartheta(\tau, s) = 1$ represents "invest" and $\vartheta(\tau, s) = 0$ "do not invest". An investment strategy ϑ is *optimal* for the agent if it prescribes an investment decision that maximizes the agent's conditionally expected utility, given any observed signal value s with any given precision τ. A work plan $\xi = (\tau, \vartheta)$ is *sequentially rational* for the agent if the investment strategy ϑ is optimal and the signal precision τ maximizes the agent's (unconditional) expected utility (in anticipation of using ϑ).

We proceed to specify more precisely conditional and unconditional expected values for the agent, along with the agent's behavior. First, an investment strategy ϑ^* is optimal for the agent if, for every signal precision $\tau \in T$, it prescribes investment (no investment) whenever the agent's conditionally expected remuneration utility (given his signal) from investment exceeds (falls short of) his remuneration utility $u(w_N)$ from not investing. The cost of information acquisition here is irrelevant, since it is already sunk at the moment of investment and the agent's total utility is additively separable in remuneration utility and disutility of information acquisition.

The agent's associated *interim total expected utility*, after he has chosen his signal precision, but before he receives his private signal and applies an optimal investment strategy ϑ^*, is obtained by taking the expectation over all possible signal values and subtracting his cost for information acquisition:

$$\mathbb{E}_{\vartheta^*}\left[V\left(\boldsymbol{w},\tau,a\right)\right]=\mathbb{E}_{\vartheta^*}\left[U\left(\boldsymbol{w},\tau\right)\right]-g\left(\tau,a\right).$$

Under any agreed-upon contract, the agent will choose his signal precision $\tau \in T$ (where T is either a singleton, a two-point set, or $T = \mathbb{R}_+$), so as to maximize this interim expected total utility. The set of sequentially rational signal precisions for the agent is thus

$$T^*\left(\boldsymbol{w},a\right)=\arg\max_{\tau\in T}\mathbb{E}_{\vartheta^*}\left[V\left(\boldsymbol{w},\tau,a\right)\right]. \tag{5}$$

Let $V^*\left(\boldsymbol{w},a\right)$ denote the associated obtained total utility:

$$\mathbb{E}\left[V^*\left(\boldsymbol{w},a\right)\right]=\max_{\tau\in T}\mathbb{E}_{\vartheta^*}\left[V\left(\boldsymbol{w},\tau,a\right)\right], \tag{6}$$

Turning to the principal: if the agent has multiple optimal signal precisions—the non-generic case when $T^*\left(\boldsymbol{w},a\right)$ is not a singleton set—then we assume that the agent will choose a signal precision in $T^*\left(\boldsymbol{w},a\right)$ that maximizes the principal's expected profit (among all signal precisions in $T^*\left(\boldsymbol{w},a\right)$). In other words, the agent has a lexicographic loyalty (or sympathy) with the principal. Let $\Pi^*\left(\boldsymbol{w},a\right)$ denote the associated equilibrium random profit to the principal, that is, her net return after the state of nature has materialized and the agent has been paid his agreed-upon remuneration and acted sequentially optimally upon it.

2.4. Feasible and undominated contracts. A contract $\boldsymbol{w} \in W$ will be called feasible if it meets the limited liability constraint and both parties' participation constraints under some sequentially rational work plan for the agent:

Definition 1. *A contract $\boldsymbol{w}$ is **feasible** if $\boldsymbol{w} \in W_0$ and $\boldsymbol{w}$ admits a sequentially rational work plan for the agent such that*

$$\mathbb{E}\left[\Pi^*\left(\boldsymbol{w},a\right)\right]\geq\bar{\pi}\quad\textit{and}\quad\mathbb{E}\left[V^*\left(\boldsymbol{w},a\right)\right]\geq\bar{u}.$$

Let $W_F \subset W_0$ denote the (potentially empty) subset of feasible contracts. A feasible contract will be called undominated if there is no other feasible contract under which no party is worse off and at least one party is better off:

Definition 2. *A contract $\boldsymbol{w} \in W_F$ is **undominated** if there exists no contract $\boldsymbol{w}' \in W_F$ such that*

$$\mathbb{E}\left[\Pi^*\left(\boldsymbol{w}',a\right)\right]\geq\mathbb{E}\left[\Pi^*\left(\boldsymbol{w},a\right)\right]\quad\textit{and}\quad\mathbb{E}\left[V^*\left(\boldsymbol{w}',a\right)\right]\geq\mathbb{E}\left[V^*\left(\boldsymbol{w},a\right)\right]$$

with at least one inequality strict (for some sequentially rational work plan under $\boldsymbol{w}$ and some sequentially rational work plan under $\boldsymbol{w}'$).

Let $W_{UD} \subset W_F$ be the (potentially empty) subset of undominated contracts.

Remark 2. *Under all limited-liability contracts that give some of the potential gains of trade to the principal ($\mathbb{E}[\Pi^*(\boldsymbol{w}, a)] > \bar{\pi}$), the agent will reap only part of the fruits of his information acquisition effort, for which he bears all the cost. Hence, unless the agent has all the bargaining power, he will make socially suboptimal efforts to acquire information. For this reason, even undominated contracts are Pareto inefficient, unless they give all gains of trade to the agent.*

2.5. Bargaining. As mentioned in the introduction, the contract is initially 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 agent's outside option reflects the intensity of competition for 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 intensity 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.[17] 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 $W^{NBS}(\beta)$ be the set of generalized Nash-bargaining solutions when the principal's bargaining power is β and that of the agent is $1 - \beta$:

$$W^{NBS}(\beta) = \arg\max_{\boldsymbol{w} \in W_F} \left(\mathbb{E}[\Pi^*(\boldsymbol{w}, a)] - \bar{\pi}\right)^{\beta} \cdot \left(\mathbb{E}[V^*(\boldsymbol{w}, a)] - \bar{u}\right)^{1-\beta}. \qquad (7)$$

Clearly all Nash bargaining solutions are undominated; $W^{NBS}(\beta) \subseteq W_{UD}$.

2.6. Equilibrium. For a given specification of the model, including the bargaining power β, we will call a contract $\boldsymbol{w}$ an *equilibrium under bargaining power* β if it is a Nash bargaining solution, $\boldsymbol{w} \in W^{NBS}(\beta)$, and at least one party is better off than in autarky, $\mathbb{E}[\Pi^*(\boldsymbol{w}, a)] \geq \bar{\pi}$ and $\mathbb{E}[V^*(\boldsymbol{w}, a)] \geq \bar{u}$ with at least one strict inequality. Hence, equilibria do not exist for not very good projects, not very able agents, and/or parties with very good outside options.

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

By an equilibrium we mean a contract that is an equilibrium under some bargaining power $\beta \in [0,1]$. Let W^* denote the set of *equilibrium contracts*.

3. Analysis

3.1. First best. Before embarking on the main analysis, let us briefly consider a class of first-best contracts in the special case of a risk-neutral agent with no limited-liability constraint. So imagine, for the moment, that the principal and a risk neutral agent can sign and enforce any monotonic contract they like (but the agent will still decide how much information to gather and for what signals to invest). A natural and simple class of contracts then springs to mind, namely, that the principal sells the project to the agent. Then the agent, as both owner and expert, will fully internalize the consequences of his information-gathering effort and his investment decision. This first-best solution is not realistic, for at least two reasons, but is of theoretical interest, since it places an upper bound on what is achievable in equilibrium. One reason why it is not realistic is that in practice a potential agent does not have enough financial resources. Another reason is that this first-best solution presumes that the agent can acquire information, beyond the prior that he shares with the principal, only after he has bought the project.

To be more precise, let P be any price for the project, to be paid ex ante, that is before the agent can acquire more information, and independent of the state of nature. Such a contract can be formally represented in the present model by a strictly monotonic contract $\boldsymbol{w} = I\boldsymbol{r} - \boldsymbol{P} \in W$, defined as

$$\boldsymbol{w} = (Ir_B - P,\ -P,\ Ir_G - P) = I\boldsymbol{r} - (P, P, P) . \tag{8}$$

The last vector in this vector equation simply specifies that the agent has to pay the principal the price P, irrespective of the agent's subsequent information acquisition, investment decision and outcome of this investment, if made. Such a contract clearly violates the limited-liability constraint, since the agent's remuneration will be negative if he invests in the bad state. However, if there were no limited-liability constraint, then it would be feasible if and only if both parties' participation constraints are met,

$$\bar{\pi} \leq P \leq \mathbb{E}\left[V^*\left(I\boldsymbol{r}, a\right)\right] - \bar{u}.$$

Such a price P exists if and only if

$$\bar{\pi} + \bar{u} \leq \mathbb{E}\left[V^*\left(I\boldsymbol{r}, a\right)\right] . \tag{9}$$

In other words, the project has to be sufficiently attractive for the agent (which in part depends on the agent's ability and outside option). By contrast, if (**??**) does not hold, then there will, a fortiori, exist no feasible contract.

The intuition for the claimed social optimality of contracts of the form (**??**) is simple. Because if both parties are risk-neutral, then it is socially optimal to place all risk on the agent, who will (a) exert the socially optimal amount of information-gathering effort (since he will reap all the benefits of having precise information), and (b) he will invest exactly for those signal values that maximizes the expected social return from the project (which he "owns"). Both claims follow from the subsequent analysis. In particular, we will see that the agent's incentives, once hired, are the same under any contract $I\boldsymbol{r} - \boldsymbol{P}$ as under contract $I\boldsymbol{r}$, both when gathering information and when making the investment decision.

We turn to the more realistic case of limited-liability contracts.

3.2. The agent's investment decision. Suppose that the two parties have signed a strictly monotonic contract $\boldsymbol{w} \in W$, that the agent has chosen his signal precision $\tau \in T$ and observed his private signal, $S = s$. Should he invest or not? Being selfish, he will base this decision on his expected remuneration utility from each of the two options. By Bayes' rule, the agent's conditionally expected remuneration utility from investing (formally, choosing $\vartheta(s) = 1$) is

$$\begin{aligned}\mathbb{E}\left[U(\boldsymbol{w},\tau) \mid s\right] &= \frac{\mu\phi_\tau(s-r_G)\,u(w_G) + (1-\mu)\,\phi_\tau(s-r_B)\,u(w_B)}{\mu\phi_\tau(s-r_G) + (1-\mu)\,\phi_\tau(s-r_B)} \\ &= \frac{\theta\psi_\tau(s)\,u(w_G) + u(w_B)}{\theta\psi_\tau(s) + 1}.\end{aligned}$$

while from not investing (choosing $\vartheta(s) = 0$), it is $u(w_N)$. Hence, it is rational for the agent to invest if and only if

$$\mathbb{E}\left[U(\boldsymbol{w},\tau) \mid s\right] \geq u(w_N). \tag{10}$$

The agent's outside option and cost of information acquisition are bygones at this point in time, and are hence irrelevant for his investment decision. In force of the MLRP and TT conditions, (**??**) defines a simple optimal investment strategy, namely, to invest if and only if the signal exceeds a certain threshold value.

Proposition 1. *For any strictly monotonic contract $\boldsymbol{w} \in W$ and positive signal precision τ, it is optimal for the agent, who may be risk neutral or risk averse, to invest if and only if $s \geq s^*(\boldsymbol{w},\tau)$, where*

$$s^*(\boldsymbol{w},\tau) = \psi_\tau^{-1}\left[\frac{1}{\theta}\cdot\frac{u(w_N) - u(w_B)}{u(w_G) - u(w_N)}\right]. \tag{11}$$

In other words, an optimal investment strategy for the agent is the step function $\vartheta^* : T \times \mathbb{R} \to \{0,1\}$ defined by $\vartheta^*(\tau, s) = 1$ if and only if $s \geq s^*(\boldsymbol{w}, \tau)$, for all $\tau > 0$.[18] We see in (**??**) that the agent's critical signal threshold for investment is higher—the agent is more cautious—when the probability for the bad state is higher and/or the the ratio between his utility loss from the penalty and utility gain from the bonus is higher. The effect of the contract upon an agent's investment decision, given any signal precision that he may have chosen, can be summarized in a single number,

$$\rho(\boldsymbol{w}) = \theta \cdot \frac{u(w_G) - u(w_N)}{u(w_N) - u(w_B)}, \tag{12}$$

a number we will henceforth call the *carrot-stick ratio.* This is the probability-weighted ratio for the agent's utility gain from "always doing the right thing", namely, investing in the good state and abstaining from investment in the bad state. The higher the carrot-stick ratio, the lower is the agent's signal threshold, that is, the wider is the range of signal values for which he will invest. The carrot-stick ratio is independent of the agent's signal precision. For a risk neutral agent, the carrot-stick ratio of any contract is proportional to its ratio between the bonus and the penalty; $\rho(\boldsymbol{w}) = \theta b/p$. In particular, it then does not depend on the payment *levels*, $\rho(\lambda \boldsymbol{w}) = \rho(\boldsymbol{w})$ for all $\lambda > 0$.

In the same vein as the carrot-stick ratio, we define a contract's *power* as the square root of the probability-scaled product of the utility gain from making the right investment decision in each state of nature:

$$\kappa(\boldsymbol{w}) = \sqrt{\mu \cdot [u(w_G) - u(w_N)] \cdot (1-\mu) \cdot [u(w_N) - u(w_B)]}. \tag{13}$$

With this definition, all strictly monotonic contracts have positive power, and the bigger is the agent's utility gain from investing in the good state and his utility loss from investing in the bad, the higher is the power of his contract.

By proposition **??** it is only a contract's carrot-stick ratio, not its power, that matters for the agent's investment decision once he has obtained his private information. However, as will be seen, both the carrot-stick ratio and the power of the contract matter for the agent's effort to acquire information (and these are the only channels from the contract to the agent's behavior).

Suppose, temporarily, that the agent's signal precision is fixed and given; $T = \{\tau\}$ for some $\tau > 0$. If the agent is risk-neutral, then a contract is undominated if and only if the ratio between the bonus and penalty rates equals the project's return-rate ratio. The reason is that this is necessary and sufficient to align the incentives of the agent with those of the principal at the moment of the investment decision.

[18] Since the signal by assumption has a continuous distribution, this investment strategy is not unique. The agent may, without affecting his expected utility or the principal's expected profit, make arbitrary investment decisions on any signal subset that has Lebesgue measure zero.

Proposition 2. *Suppose that the agent is risk neutral and has a fixed signal precision $\tau > 0$. A strictly monotonic contract $\boldsymbol{w}$ is then undominated if and only if $\rho(\boldsymbol{w}) = \rho(\boldsymbol{r})$, or, equivalently, if and only if*

$$\frac{b}{p} = \frac{r_G}{|r_B|}. \tag{14}$$

We note that this result is independent of the agent's signal precision (as long as it is positive). However, as we will see below, the endogeneity of the agent's signal precision will in general introduce a trade-off between the desire to align the two parties' incentives both at the (early) moment of information acquisition and at the (later) moment of investment.

3.3. Interim expected utility and profit. We now move one step back in time, to the moment when the agent has chosen a signal precision $\tau > 0$ but not yet received his private signal. His expected utility from his future remuneration, as specified by the contract, and when using his sequentially rational investment strategy, ϑ^*, once he obtains his signal, is

$$\mathbb{E}[U(\boldsymbol{w},\tau)] = p_G(\boldsymbol{w},\tau)\cdot u(w_G) + p_N(\boldsymbol{w},\tau)\cdot u(w_N) + p_B(\boldsymbol{w},\tau)\cdot u(w_B), \tag{15}$$

where $p_G(\boldsymbol{w},\tau)$ is the probability that he will invest and the state turns out to be good, $p_B(\boldsymbol{w},\tau)$ the probability that he will invest and the state turns out to be bad, and $p_N(\boldsymbol{w},\tau) = 1 - p_G(\boldsymbol{w},\tau) - p_B(\boldsymbol{w},\tau)$ is the residual probability that he will not invest. In force of Proposition **??**, the above probabilities can, for any strictly monotonic contract and positive signal precision, be expressed in terms of the primitives of the model as

$$p_G(\boldsymbol{w},\tau) = \mu\cdot\int_{s^*(\boldsymbol{w},\tau)}^{+\infty}\phi_\tau(s - r_G)\,ds \tag{16}$$

and

$$p_B(\boldsymbol{w},\tau) = (1-\mu)\cdot\int_{s^*(\boldsymbol{w},\tau)}^{+\infty}\phi_\tau(s - r_B)\,ds. \tag{17}$$

The associated *expected profit* to the principal, once the contract has been signed and the agent has chosen his signal precision, is accordingly

$$\mathbb{E}[\Pi(\boldsymbol{w},\tau)] = \pi_G\cdot p_G(\boldsymbol{w},\tau) + \pi_N\cdot p_N(\boldsymbol{w},\tau) + \pi_B\cdot p_B(\boldsymbol{w},\tau). \tag{18}$$

3.4. The agent's information acquisition. We are now in a position to analyze the agent's choice of signal precision. As shown in Weibull, Mattsson and Voorneveld (2007), decision-making under endogenous uncertainty typically exhibits a non-convexity. This is true also in the present model and is illustrated in the diagram below, which shows the agent's total expected utility (expected utility from remuneration, net of his cost for information acquisition) as a function of his signal precision, under a given contract, for five different ability levels. The more able the agent, the higher his utility curve.

Figure 1: The agent's total expected utility as a function of his signal precision.

The thin horizontal line represents the agent's utility, $\mathbb{E}\left[U\left(\boldsymbol{w},0\right)\right]$, when not making any effort to gather information and not investing. The most unable agent's total expected utility (the lowest curve) is monotonically decreasing in signal precision. Such an agent thus makes no effort under the given contract. For slightly more able agents (the second lowest curve) there exists an local maximum at positive signal precision, but this is still not better than not making any effort at all and not investing. Hence, also such agents will choose to exert no effort. Highly able agents (the two highest curves) also have a local maximum at a positive signal precision, but this time this is also the agent's global maximum. Such an agent thus chooses a positive effort, and we see that this is higher for more able agents. The dashed curve shows the knife-edge case of an agent with intermediate ability. This agent is indifferent between (a) not gather any information and not invest, and (b) gathering information and then making an investment decision conditional upon his received information about the project. The positive signal precision that makes this agent indifferent is approximately $\tau = 0.5$ in this numerical example (the point where the dashed curve touches the horizontal line). Hence, under the given contract, no agent will ever choose a signal precision between zero and (approximately) 0.5. Any slight change of the contract may tip this agent to suddenly switch between zero signal precision and signal precision near 0.5, or in the market conditions—as represented by the prior μ probability for the good state of nature.

Does there, despite the just illustrated potential non-convexity, exist an optimal signal precision for every agent under every contract? The following result follows from standard arguments.

Proposition 3. *The subset $T^*(\boldsymbol{w}, a) \subseteq T$ is non-empty and compact under any contract $\boldsymbol{w} \in W$ and for any agent ability $a > 0$. If $\tau \in T^*(\boldsymbol{w}, a)$ and $\tau > 0$, then*

$$\frac{\partial}{\partial \tau}\mathbb{E}\left[U(\boldsymbol{w}, \tau)\right] = \frac{\partial}{\partial \tau} g(\tau, a). \tag{19}$$

In other words, the agent always has at least one optimal signal precision level, and if an optimal precision is positive then it meets the necessary first-order condition that the agent's marginal *expected remuneration utility* should equal his marginal information cost.

Remark 3. *The lack of convexity noted above is not an effect of the agency problem; it would remain even if the principal would herself acquire information and make the investment decision. For sufficiently high costs for information acquisition and/or low potential returns from the project, she would choose not to acquire any information and would not invest, while for sufficiently low costs and high potential returns, she would make significant effort to acquire information and then decide whether or not to invest, conditional upon the information obtained. She would never make small or intermediate efforts to acquire information.*

4. EQUILIBRIUM

We now proceed to analyze equilibrium under the assumption that the noise term ε in the agent's signal is normally distributed. We established above that under any strictly monotonic contract the agent will then either choose zero signal precision, or else he will choose a positive signal precision τ that meets the first-order condition (**??**), which for the normal distribution writes

$$\frac{\Delta \cdot \kappa(\boldsymbol{w})}{\sqrt{2\pi\tau}} \cdot \exp\left(-\frac{\Delta^2 \tau}{2} - \left(\frac{\ln \rho(\boldsymbol{w})}{2\Delta}\right)^2 \frac{1}{2\tau}\right) = \frac{\partial g(\tau, a)}{\partial \tau}, \tag{20}$$

where $\Delta = (r_G - r_B)/2$. The left-hand side of this equation represents the agent's marginal expected remuneration utility with respect to signal precision. We see that this marginal remuneration utility is higher the more high-powered is his contract (the larger $\kappa(\boldsymbol{w})$ is), for any signal precision $\tau > 0$ that the agent may choose. This monotonicity is not surprising: a contract that provides a higher bonus after a successful investment and incurs a higher penalty after a failed investment will make it more worthwhile for the agent to acquire information. One also sees that the agent's marginal remuneration utility is non-monotonic in the carrot-stick ratio of the contract, and that it is maximal, at all signal precision levels, when this ratio is one. This is precisely when the agent's signal threshold for investment is at the midpoint between the low and high return rates (r_B and r_G), a point at which the agent is most uncertain whether the return will be r_B or r_G. This is where his (bimodal) signal density is at its local minimum between the low and high return rates.

An immediate and general observation is that in order for a contract to be an equilibrium it has to contain both a bonus and a penalty. This is seen in equation (**??**), where the left-hand side vanishes for every contract with zero power. Under such a contract, the agent would acquire no information, and thus the principal would either pay the agent nothing and gain nothing or make a loss compared with her outside option. Formally,

Proposition 4. *Assume normally distributed noise and an agent who may be risk neutral or risk averse. Then*

$$\boldsymbol{w} \in W^* \quad \Rightarrow \quad w_B < w_N < w_G.$$

*Moreover, the agent's signal precision τ is positive and satisfies (**??**) under all equilibrium contracts $\boldsymbol{w}$.*

We next turn to the special case of risk-neutral agents with arbitrary information costs, then consider agents with linear information costs but arbitrary risk attitudes, then consider risk neutral agents with linear information costs, and finally consider a class of risk-averse agents with non-linear information costs.

4.1. Risk-neutral agents. If the agent is risk neutral, then every contract with a positive pay to the agent after a failed investment is dominated. In other words, the limited liability constraint is binding for risk neutral agents. Formally,

Proposition 5. *Assume normally distributed noise and a risk neutral agent. Then*

$$\boldsymbol{w} \in W^* \quad \Rightarrow \quad w_B = 0. \tag{21}$$

However, while familiar in other principal-agent models, this result may first come as a surprise in the present model, since a transfer from the principal to the agent would neither affect the agent's information acquisition nor his subsequent investment decision. More precisely, if $\boldsymbol{w}' = \boldsymbol{w} + (\delta, \delta, \delta)$ for some $\delta > 0$, then $\rho(\boldsymbol{w}') = \rho(\boldsymbol{w})$ and $\kappa(\boldsymbol{w}') = \kappa(\boldsymbol{w})$. Hence, the agent's incentives are identical under both contracts. One might therefore be led to believe that if $\boldsymbol{w}$ is undominated, then so is $\boldsymbol{w}'$, which would contradict the claim that $w'_B = 0$ is necessary for $\boldsymbol{w}'$ to be undominated. However, such a contract $\boldsymbol{w}'$ is not undominated. The reason is that it is an interior contract in W_0 and hence allows us to vary all its three components in any direction. In particular, there will exist a *third* contract $\boldsymbol{w}''$, under which the agent is paid somewhat less after a failed investment, $0 < w''_B < w'_B$, and slightly more after a successful investment, $w''_G > w'_G$, inducing him to make more effort to acquire information, hence obtain a higher signal precision. Contract $\boldsymbol{w}''$ would thus result in a "bigger cake", in terms of expected remuneration utility and profit than $\boldsymbol{w}'$. The first-order effect on his disutility of effort is nil, by the envelope theorem, so the total cake under $\boldsymbol{w}''$ is bigger and each party can be given a slightly larger slice than under $\boldsymbol{w}'$.

The result in Proposition **??** is illustrated in the diagram below.

Figure 2: The feasible set of contracts with the principal's expected profit on the horizontal axis and the agent's expected utility on the vertical.

On the $x-$axis is the principal's expected profit and on the y-axis the agent's total expected utility. The indicated oblong areas (partly covering each other) are the sets of profit-utility pairs that are obtainable by feasible contracts with a pre-specified payment $w_B \geq 0$. These areas have been obtained by numerical simulations based on millions of feasible contracts. The smallest (light grey) oblong set was obtained for the highest such pay ($w_B = 0.01$), the intermediate (partly covered and dark gray) oblong set was obtained for the intermediate pay $w_B = 0.001$, and the largest (partly covered and black) oblong set to zero pay, $w_B = 0$. The set of *undominated* contract are thus those that result in profit-utility pairs on the north-eastern boundary of the biggest (black) oblong set. All other feasible contracts are dominated. The diagram thus confirms numerically the claim in Proposition **??**. We also see that, unlike most economics models, the Pareto frontier is non-monotonic. For feasible contracts with low expected utility, the maximum feasible expected profit is increasing, not decreasing, in the expected utility.

4.2. Linear information costs. We will here examine some qualitative properties of our model in the special case when the cost of information acquisition is linear. By contrast, we here do not now require the agent to be risk neutral. More precisely, we now assume that the agent has a cost, or disutility, of information acquisition that is proportional to the obtained signal precision and inversely proportional to the agent's ability:

[LIN] *$g(\tau, a) = c\tau/a$ for some $c > 0$.*

Condition LIN can be rephrased explicitly in terms of the agent's information-gathering *effort* $z \geq 0$ as follows. Instead of [LIN], assume linear disutility from effort, $g(z) = cz$, and let the resulting signal precision be given by $\tau = az$. That is, effort and ability (or talent) are *complements* in the production of signal precision. Then $g(z) = c\tau/a$, and we obtain [LIN]. Under this specification, and using (**??**), the necessary first-order condition (**??**) for a positive signal precision τ to be optimal for the agent (irrespective of whether he is risk-neutral or risk-averse) boils down to

$$\frac{\Delta \cdot \kappa(\boldsymbol{w})}{\sqrt{2\pi\tau}} \cdot \exp\left(-\frac{\Delta^2\tau}{2} - \left(\frac{\ln \rho(\boldsymbol{w})}{2\Delta}\right)^2 \frac{1}{2\tau}\right) = \frac{c}{a}. \tag{22}$$

This equation is illustrated in the diagram below, for a given contract $\boldsymbol{w}$. 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 given contract $\boldsymbol{w}$ with $\rho(\boldsymbol{w}) \neq 1$. The horizontal lines represent different values for the right-hand side, the (constant) marginal cost c/a to the agent of signal precision. For agents of high ability (given the marginal information cost c), the diagram shows that the necessary first-order condition (**??**) 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.

Figure 3: The agent's marginal expected remuneration-utility with respect to his signal precision, and alternative levels of his marginal cost of signal precision.

The dashed curve is the agent's marginal remuneration utility under a contract $\boldsymbol{w}'$ with unit carrot-stick ratio, $\rho(\boldsymbol{w}') = 1$. In this knife-edge case, the left-hand side in the necessary first-order condition (**??**) is strictly decreasing from plus infinity towards zero, and hence has a unique solution. We also see that the signal precision that is the unique local maximum (the one at the right-most intersection when $\rho(\boldsymbol{w}) \neq 1$) 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 (**??**) have exactly two solutions? And can a lower bound be given on the precision level that constitutes the local maximum?

Lemma 1. *Assume normally distributed noise and linear disutility of information acquisition. The agent may be risk neutral or risk averse. For any contract* $\boldsymbol{w} \in W$ *and agent ability* $a > 0$, *the agent's expected utility,* $\mathbb{E}[V(\boldsymbol{w}, \tau, a)]$, *has a local* *maximum at a positive signal-precision level if and only if*

$$\frac{\Delta^2 \cdot \kappa(\boldsymbol{w})}{\sqrt{\pi\left(\sqrt{(\ln\rho(\boldsymbol{w}))^2+1}-1\right)\exp\left(\sqrt{(\ln\rho(\boldsymbol{w}))^2+1}\right)}} > \frac{c}{a}. \tag{23}$$

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

$$\tau_{\min}(\boldsymbol{w}) = \frac{1}{2\Delta^2}\cdot\left(\sqrt{(\ln\rho(\boldsymbol{w}))^2+1}-1\right). \tag{24}$$

We note that the left-hand side in (**??**) blows up to plus infinity for contracts with unit carrot-stick ratio (and only then). Hence, this inequality is automatically satisfied for all such contracts, and the lower bound, $\tau_{\min}(\boldsymbol{w})$, is then zero. (See the dashed curve in Figure 3.) For all other contracts, those with $\rho(\boldsymbol{w}) \neq 1$, the left-hand side in (**??**) is a well-defined positive real number, and so is the lower bound $\tau_{\min}(\boldsymbol{w})$. For any such contract and any given marginal information cost $c > 0$, inequality (**??**) holds for all abilities a above a certain critical positive value. Agents with lower ability will choose signal precision zero, that is, acquire no information at all. Write $\tau^o(\boldsymbol{w}, a)$ for the unique solution to equation (**??**) that exceeds $\tau_{\min}(\boldsymbol{w})$, for any contract $\boldsymbol{w} \in W$ satisfying (**??**). This is the unique interior local maximum to the agent's decision problem concerning information acquisition. The following inequality is necessary and sufficient for an agent to weakly prefer working and acquiring information at precision $\tau^o(\boldsymbol{w}, a) > \tau_{\min}(\boldsymbol{w})$ over shirking (choosing $\tau = 0$):

$$\mathbb{E}\left[U\left(\boldsymbol{w}, \tau^o(\boldsymbol{w}, a)\right)\right] - \frac{c}{a}\tau^o(\boldsymbol{w}, a) \geq \mathbb{E}\left[U(\boldsymbol{w}, 0)\right]. \tag{25}$$

Proposition 6. *Assume normally distributed noise and linear disutility of information acquisition. The agent may be risk neutral or risk averse. The agent's signal precision τ^* is positive if (**??**) and (**??**) hold, and then $\tau^* = \tau^o(\boldsymbol{w}, a)$, the unique solution τ to (**??**) that exceeds $\tau_{\min}(\boldsymbol{w})$. Otherwise, $\tau^* = 0$ and the contract is not feasible. If (**??**) and (**??**) hold, then*

(i) $\tau^o(\boldsymbol{w}, a)$ is decreasing in the marginal cost of information, c, and increasing in the agent's ability, a.

(ii) $\tau^o(\boldsymbol{w}, a)$ is increasing in the contract's power when the contract's carrot-stick ratio is held fixed.

(iii) $\tau^o(\boldsymbol{w}, a)$ is non-monotonic in the contract's carrot-stick ratio when the contract's power is held fixed, and, given the contract's power, is maximal when the contract's carrot-stick ratio is unity.

(iv) Among contracts with unit carrot-stick ratio, $\tau^o(\boldsymbol{w}, a)$ is non-monotonic in the prior μ, and it is maximal when $\mu = 1/2$.

All claims except the last follow from observations made above. The last claim is intuitive, stating that (at least under certain contracts) the agent has maximal incentive to gather information when he is most uncertain.[19]

Remark 4. *Suppose that all payments in a contract $\boldsymbol{w}$ were multiplied by the same factor $\lambda > 0$. How would this affect the agent's behavior?*

If the agent is risk-neutral, then the new contract will have the same carrot-stick ratio for him as the original one, but the power will change by a factor λ^2. Hence, the agent will not change his investment decision, once he has his information, but he will acquire more (less) information under the new contract if $\lambda > 1$ ($\lambda < 1$), and his expected utility will go up (down).

By contrast, if the agent is risk averse and has a logarithmic Bernoulli function, then he will behave exactly in the same way under the new contract as under the original. The reason is that for such an agent, both the carrot-stick ratio and the power of the two contracts is the same: $\rho(\lambda\boldsymbol{w}) = \rho(\boldsymbol{w})$ and $\kappa(\lambda\boldsymbol{w}) = \kappa(\boldsymbol{w})$ for all $\lambda > 0$.[20]

4.3. Risk neutral agent with linear information cost. When the agent is risk-neutral, $w_B = 0$ in all undominated contracts. This allows graphical illustrations of the set of feasible contracts in two dimensions. The diagram below shows all contracts $\boldsymbol{w} \in W$ that have $w_B = 0$, with w_N on the horizontal axis and w_G on the vertical.[21]

By Corollary **??** and Proposition **??**, all undominated contracts for risk neutral agents are points (w_N, w_G) above the 45^o-diagonal, and each such point uniquely defines a feasible contract if and only if it lies inside the region where agent obtains at least his reservation utility and the principal her outside option profit. The oblong sets are upper contour sets for the principal's expected profit. The dashed curves are the agent's indifference curves (in terms of his expected total utility). The associated upper-contour set for the agent consists of all points above the indifference curve in question. Hence, the agent's upper contour sets are non-convex, although the agent is risk neutral. We note, in particular, the horizontal (when w_N is small) and vertical (when w_N is close to w_G) segments of the agent's indifference curves. They represent

[19] The claim is stated for the special case when $\rho(\boldsymbol{w}) = 1$ and then the observation is immediate. By continuity, the claim holds for all contracts with carrot-stick ratio close enough to unity.

[20] If $u(x) \equiv \ln x$, and $w_B > 0$, then

$$\rho(w) = \frac{\mu}{1-\mu} \cdot \frac{\ln(w_G/w_N)}{\ln(w_N/w_B)} \quad \text{and} \quad \kappa(\boldsymbol{w}) = \mu \cdot \ln(w_G/w_N) \cdot (1-\mu) \cdot \ln(w_N/w_B)$$

[21] The simulations behind these diagrams are based on the following numerical specification: $\mu = 0.25$, $I = 1 = r_G = -r_B$, $a = 3$, and $c = 0.001$.

contracts under which the agent makes no effort to acquire information (zero signal precision). The tangency points between the two parties' upper contour sets are indicated by small circles. These are undominated contracts if the resulting expected utility to the agent is above his reservation value and the principal's expected profit is above her reservation value (or, to be more precise, both participation constraints hold at least weakly and at least one holds strictly). The collection of all possible such tangency points defines a curve, the "contract curve" in the bargaining between the two parties.

Figure 4: Contour maps for the expected profit to the principal (solid lines) and for the expected utility to the agent (dashed lines).

Some observations are called for. First, we note in the diagram that there is an interior feasible contract that is optimal for the principal; the contract $\boldsymbol{w}^{**} \approx (0, 0.05, 0.13)$, located somewhere centrally in the smallest upper-contour set for the principal's expected profit. Suppose, for instance, that the agent's reservation utility is given by the lowest (dashed) utility isoquant in the diagram. Then the contract $\boldsymbol{w}^{**}$ results in expected utility to the agent above his reservation level. Indeed, this contract would be proposed by the principal even if she had all the bargaining power (that is, could make a take-it-or-leave it offer). The principal realizes that the agent would make too little effort to acquire information if hired at less well-paid contract (even if it would have the same carrot-stick ratio). This feature of our model—that the principal may not always wish to press down the agent to the agent's participation constraint—is a reflection of the non-monotonicity of the Pareto frontier in Figure 3.

Second, the diagram suggests that the effect of changes in bargaining power upon the equilibrium contract 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 "contract curve". Third, we note that 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 outside option value, the bonus *rate* is virtually unaffected.[22] Fourth, the above diagram shows that it is not always in the best 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 his ability, but (iii) the principal believes it is higher, and (iv) they both know that the agent's outside option has a very low value. Instead of offering the above-mentioned contract $\boldsymbol{w}^{**}$, the principal would then offer another contract, which may well have lower expected utility for the agent.[23]

We conclude by illustrating how the principal's equilibrium profit and agent's utility depend on the agent's ability, see diagram below. The diagram is drawn for the classical case in the principal-agent literature when the principal has all the bargaining power, $\beta = 1$. The principal will then give a take-it-or-leave-it contract (as indicated by the small circles in Figure 4). The increasing and concave curve is the principal's expected profit at the globally optimal contract for her, $\boldsymbol{w}^{**}$, adapted to the agent's ability, granted the agent would accept the contract. The decreasing and convex curve is the agent's expected (total) utility if he would be employed according to the principal's globally optimal contract, $\boldsymbol{w}^{**}$. The dashed horizontal line represents the principal's outside option, $\bar{\pi} = r_0 I$, and the dashed upward-sloping straight line the agent's outside option, $\bar{u}$.

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

[23] In the numerical specifiction behind the diagram, $a = 3$, while for $a = 5$, the principal would propose $\boldsymbol{w}^o \approx (0, 0.035, 0.115)$, a worse contract for the agent. This 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.

Figure 5: The principal's equilibrium profit and agent's equilibrium utility as functions of the agent's ablity (or talent).

As seen in Figure 5, the principal will not hire an agent with ability below $a_{\min} = 1$. Such *low-ability agents* are not worthwhile to hire for the principal since they need to be paid a very high bonus in order to work hard enough to obtain sufficient signal precision for the investment decision to be worthwhile. Agents with ability below this critical level, $a < a_{\min}$, will thus not be hired and both parties will instead take their outside options. In the knife-edge case of an agent with exactly the critical ability, $a = a_{\min}$, the principal will be indifferent between hiring and not hiring. Such an agent will be very happy to be employed. His expected utility would take a big jump up from his reservation utility. The reason for this jump is that the principal needs to pay him a lot in order to make him work enough to obtain a precise enough signal. We also see in the diagram that the agent's total expected utility is decreasing in his ability when his ability is in a certain range above the minimum level, $a_{\min}$. The reason for this, perhaps surprising, result is that the agent's participation constraint is not yet binding (since he is not that good). For such agents of *intermediate ability*, the principal will offer her globally optimal contract, $\boldsymbol{w}^{**}$, as adapted to the agent's ability. The more able the agent is, the less does the principal have to pay in order for the agent to acquire information, and the principal's expected profit is then increasing in the agent's ability, where we note that the marginal profit value of increased ability is declining. The reason is that the maximal return from investment is bounded and making an already precise signal even more precise is not so valuable to the principal. This holds for a range of intermediate ability levels.

However, as the agent's ability is gradually increased over this interval, sooner or later his expected utility from working under the principal's contract, $\boldsymbol{w}^{**}$ (as adapted to his ability) will reach his reservation utility, $\bar{u}$. In the diagram, this happens at the second critical ability level $a_{med} \approx 2.83$, namely, where the agent's downward-sloping expected-utility curve intersects his dashed and upward-sloping reservation utility curve. For agents with higher ability, $a > a_{med}$, the principal's contract will need to match the agent's outside option. Hence, the agent's expected utility curve has a kink at the critical ability level a_{med}, and for higher ability levels his expected utility from working for the principal tracks his reservation utility. Such *high-ability agents* are thus more and more expensive for the principal, the more able they are, so her expected profit no longer increases according to her globally optimal contract. It may first increase, but at a lower rate, and then decrease, as indicated in the diagram. Sooner or later, as we increase the agent's ability, the principal's expected profit will reach the profit from her outside option. In the diagram, this happens at the third critical ability level, $a_{\max} \approx 4.75$, the point where the solid and downward-sloping profit curve intersects the dashed and horizontal outside-option profit line. Agents with higher ability, the "*super stars*", are not worthwhile to hire for our principal, since their outside options are too good. The principal does better by taking her outside option.

In sum, for a given project (or company), the agents can be divided into four categories, (I) those with low ability, who are not worthwhile to hire, (II) those of intermediate ability, who are worth hiring, the more so the better they are, (III) agents with high ability, who are worth hiring but less so the better they are, and (IV) super stars, agents with very high ability and very attractive outside options, and these are not worth hiring either. The principal's expected profit is a continuous but non-monotonic function of the agent's ability, and is maximal when the agent's ability is just above a_{med}, that is, for agents with high, but not very high, ability. The expected utility for an agent of given low ability is thus first increasing with his outside option (in category I), jumping up discontinuously and then continuously diminishing with ability when intermediate (category II), then again increasing with the outside option for high-ability and super-star agents (categories III and IV). The optimal ability, from the principal's viewpoint, in general depends on the other parameters in the model, in particular on the size I of the project or firm.

Double-or-nothing projects. We finally give a detailed specification of our model in the special case of a project that either returns double or nothing from the investment, $r_G = -r_B = 1$, still for a risk neutral agent with linear information costs. Then $\Delta = 1$ and $\rho(\boldsymbol{r}) = \theta$. We will here focus on equilibrium contracts $\boldsymbol{w} \in W^*$. Hence $0 = w_B < w_N < w_G$. Writing y for the "salary" w_N and b for the bonus rate $(w_G - w_N)/w_N$, we then have, $\rho(\boldsymbol{w}) = \theta b$ and $\kappa(\boldsymbol{w}) = y\sqrt{\mu(1-\mu)b}$. The agent's signal threshold for investment is now

$$s^*(\boldsymbol{w}, \tau) = -\frac{1}{2\tau}(\ln\theta + \ln b)$$

(see Appendix A). The socially optimal signal-threshold can be shown to be

$$\hat{s}(\boldsymbol{r}, \tau) = -(2\tau)^{-1}\ln\theta$$

(see Appendix A). Hence, the agent's signal threshold lies above (below) the socially optimal signal threshold if the bonus rate is below (above) unity. This is not surprising, since the bonus rate is one if and only if the agent's contract also gives double or nothing, in which case the two parties' incentives are perfectly aligned at the investment decision stage.

Second, under any undominated contract the agent's equilibrium signal precision, τ^*, is positive and satisfies the first-order condition (**??**), which now boils down to

$$y \cdot \sqrt{\frac{\mu(1-\mu)b}{2\pi\tau}} \cdot \exp\left[-\frac{\tau}{2} - \frac{(\ln\theta + \ln b)^2}{8\tau}\right] = \frac{c}{a}. \tag{26}$$

The left-hand side is linearly increasing in the contract's salary, y, at all positive signal precisions and bonus rates. Hence, the higher the agent's salary, at a fixed bonus rate, the more effort will he make to acquire information about the project. By contrast, the left-hand side is not monotonic in the contract's bonus rate when this is very high. An increase in the bonus rate may then induce the agent to make less, rather than more effort to acquire information. However, for all bonus rates $b \leq 1/\theta$ the left-hand side is strictly increasing in b.

Third, the agent's expected total utility is $\mathbb{E}[V(\boldsymbol{w}, \tau^*)]$, where, for any signal precision $\tau > 0$,

$$\begin{aligned}\mathbb{E}[V(\boldsymbol{w}, \tau)] &= \mu b y \cdot \Phi_1\left(\frac{\ln\theta + \ln b}{2\sqrt{\tau}} + \sqrt{\tau}\right) \\ &\quad - (1-\mu)\, y \cdot \Phi_1\left(\frac{\ln\theta + \ln b}{2\sqrt{\tau}} - \sqrt{\tau}\right) + y - \frac{c\tau}{a}\end{aligned}$$

(see (**??**)). Moreover, since the agent's signal precision is endogenous, $\partial\mathbb{E}[V(\boldsymbol{w}, \tau)]/\partial\tau = 0$ when $\tau = \tau^*$, which is precisely equation (**??**).[24]

[24] Hence, by the envelope theorem,

$$\frac{d}{db}\mathbb{E}[V(\boldsymbol{w}, \tau^*(b))] = \frac{\partial}{\partial b}\mathbb{E}[V(\boldsymbol{w}, \tau^*)]$$

Fourth, the associated expected profit for the principal is $\mathbb{E}[\Pi(\boldsymbol{w},\tau^*)]$, where, for any signal precision $\tau>0$,

$$\begin{aligned}\mathbb{E}[\Pi(\boldsymbol{w},\tau)] &= \mu(I-by)\cdot\Phi_1\left(\frac{\ln\theta+\ln b}{2\sqrt{\tau}}+\sqrt{\tau}\right)\\ &\quad -(1-\mu)(I-y)\cdot\Phi_1\left(\frac{\ln\theta+\ln b}{2\sqrt{\tau}}-\sqrt{\tau}\right)-y. \qquad (27)\end{aligned}$$

As a comparative-statics thought experiment, suppose that the project's size is increased. Then an increase in the agent's salary by the same factor (while keeping the bonus rate constant) will increase the principal's expected profit more than proportionately. To see this, let $\lambda>1$ and consider a project that differs from the original project only in its required investment, $I'=\lambda I$. Suppose that the new contract for the new project differs from the original contract only by offering the agent a proportionally higher salary: $y'=\lambda y$ and $b'=b$. If the original contract was feasible, so will the new contract be if it gives the principal at least the same expected profit as from the original project under the original contract. Moreover, as was noted above, the agent will choose a higher signal precision under the new contract (since his salary has increased). Consequently, the probability for a successful (failed) investment, if investment will be made, will be higher (lower). In view of (**??**), the principal's expected profit from the new project under the new contract will thus be at least $\lambda\mathbb{E}[\Pi(\boldsymbol{w},\tau)]$, since, for any given $\tau>0$, the right hand side of (**??**) for the new project and new contract will be multiplied by the factor λ. These comparative statics properties appear to be broadly in line with the findings in Edmans, Gabaix, and Landier (2009).

We finally note that the total surplus from the contract is

$$\begin{aligned}G(\boldsymbol{w},\tau^*) &= \mathbb{E}[\Pi(\boldsymbol{w},\tau^*)]+\mathbb{E}[V(\boldsymbol{w},\tau^*)]=\mu I\cdot\Phi_1\left(\frac{\ln\theta+\ln b}{2\sqrt{\tau^*}}+\sqrt{\tau^*}\right)\\ &\quad -(1-\mu)I\cdot\Phi_1\left(\frac{\ln\theta+\ln b}{2\sqrt{\tau^*}}-\sqrt{\tau^*}\right)-\frac{c}{a}\tau^*.\end{aligned}$$

The total surplus thus depends on the agent's salary y only indirectly, by way of influencing his choice of signal precision. As noted above, the higher his pay y is, the higher will his signal precision be. Numerical simulations suggest that if the principal has all the bargaining power and the agent's reservation utility is low (so that the principal will offer the agent the principal's optimal contract $\boldsymbol{w}^{**}$, discussed above), then the contract will pay the agent more the bigger the investment, but mostly by way of increasing the salary, not the bonus rate. For example, in one such simulation, a 25% increase in the project size I resulted in approximately 12% higher salary y^{**} to the agent but no numerically visible change of the bonus rate b^{**}.

We summarize results from numerical simulations for double-or-nothing projects in two tables. In Table A the principal has all the bargaining power while in Table B the two parties have equal bargaining power. Hence, one may think of Table A as representing the case of an outside CEO candidate and Table B as the case of a current CEO who renegotiates his contract with the board. In both tables the principal's outside option has a 1% return rate, $r_0 = 0.01$, and the agent's ability is the same, $a = 3$. The table shows equilibrium outcomes for a few combinations of I, μ, and $\overline{u}$, representing the "size" of the project or the firm (I), the "investment climate" (μ), and the "intensity of competition" for talent , respectively. The associated equilibrium outcomes are described in terms of the agent's salary y^*, bonus rate, b^*, his signal precision τ^*, the principal's expected profit $\mathbb{E}[\Pi^*]$ and the agent's expected utility $\mathbb{E}[V^*]$.

Not surprisingly, increased bargaining power to the agent reduces the principal's expected profit and raises the agent's expected utility. Less evident is that the agent will work harder, and thus be better informed about the project at hand, when he has more bargaining power. His salary then goes up but his bonus rate may go up or down. Bigger projects (larger corporations or funds), by and large, give rise to higher salaries and higher bonus rates, but not always. In better investment climates the agent works less hard and thus has less precise information. We also see that increased bargaining power to the agent results in higher gains of trade; the sum of the last two columns is higher in every row in Table B than in the same row in Table A. These are but numerical examples and should therefore be taken with a grain of salt. However, they show that the model is operational, can be brought to data, illuminate causal links, and predict how equilibrium contracts change when conditions change.

TABLE A: $\beta = 1$

I	μ	$\overline{u}$	y^*	b^*	τ^*	$\mathbb{E}[\Pi^*]$	$\mathbb{E}[V^*]$	G^*
1	0.2	0	0.015	2.33	3.51	0.153	0.020	0.173
1	0.2	0.05	0.035	2.57	4.97	0.136	0.051	0.187
1	0.4	0	0.015	1.67	3.58	0.346	0.023	0.369
1	0.4	0.05	0.035	1.29	4.74	0.333	0.051	0.384
0.5	0.2	0	0.010	2.50	2.95	0.067	0.013	0.080
0.5	0.2	0.05	0.035	2.57	4.97	0.042	0.051	0.093
0.5	0.4	0	0.010	1.50	2.88	0.162	0.015	0.177
0.5	0.4	0.05	0.035	1.29	4.74	0.140	0.051	0.191

TABLE B: $\beta = 0.5$

I	μ	$\overline{u}$	y^*	b^*	τ^*	$\mathbb{E}[\Pi^*]$	$\mathbb{E}[V^*]$	G^*
1	0.2	0	0.070	1.93	5.89	0.097	0.094	0.191
1	0.2	0.05	0.085	2.06	6.28	0.075	0.117	0.192
1	0.4	0	0.125	1.44	7.01	0.200	0.194	0.394
1	0.4	0.05	0.140	1.46	7.22	0.175	0.219	0.394
0.5	0.2	0	0.035	2.00	4.75	0.046	0.047	0.093
0.5	0.2	0.05	0.050	2.20	5.43	0.024	0.070	0.094
0.5	0.4	0	0.060	1.58	5.82	0.099	0.095	0.194
0.5	0.4	0.05	0.075	1.60	6.21	0.074	0.120	0.194

5. Contracts with stocks and/or options

The model 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. 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 has a current 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. If the company board and CEO candidate agree on some contract $\boldsymbol{w} \in W$, then the company's stock value instead becomes $\Pi_0 + \mathbb{E}\left[\Pi^*\left(\boldsymbol{w}, \tau\right)\right]$. For ease of exposition, we here assume that the stock value of the company is always positive (under all contracts to be considered). Real-life contracts often involve a salary and some company stock and/or an option to buy such stock at some future time at some preset price, which typically is the current stock price. We here briefly consider such payment schemes.

5.1. Stock-based contracts. Suppose that a CEO's contract consists of a fixed pay y and a positive share λ of the company's future stock value (after the realization of the project at hand). Such a *stock-based* contract is formally equivalent with a contract $\boldsymbol{w} \in W$ of the form $w_j = y + \lambda \cdot (\Pi_0 + Ir_j - w_j)$ for $j = B$, N and G, where $r_N = 0$. Consequently,

$$w_j = \frac{y + \lambda \cdot (\Pi_0 + Ir_j)}{1 + \lambda} \quad \text{for } j = B, N, G, \tag{28}$$

The linear dependence of the agent's payments on the project's returns implies that

$$w_G - w_N = \frac{\lambda I}{1+\lambda} \cdot r_G \quad \text{and} \quad w_N - w_B = \frac{\lambda I}{1+\lambda} \cdot |r_B| \,. \tag{29}$$

For a risk neutral agent, this implies that $\rho(\boldsymbol{w}) = \rho(\boldsymbol{r})$. Hence, such contracts align the agent's and principal's incentives perfectly at the moment when the agent takes the investment decision, see Proposition **??**. By Proposition **??**, all undominated contracts with a risk-neutral agent satisfy $w_B = 0$.[25] Hence, if a stock-based contract with a risk neutral agent is to be undominated (in the space of all contracts, stock-based or not) we necessarily have

$$w_B = 0, \quad w_N = \frac{\lambda I}{1+\lambda} \cdot |r_B| \quad \text{and} \quad w_G = \frac{\lambda I}{1+\lambda} \cdot (r_G + |r_B|) \,. \tag{30}$$

As seen in Figure 5 above, the previous section, there are situations in which all undominated contracts have carrot-stick ratios above $\rho(\boldsymbol{r})$ (and this may be so in general). As a consequence, stock-based contracts for risk-neutral agents may be dominated by other, non-stock based contracts. An intuition for the potential suboptimality of stock-based contracts for risk neutral agents is that although they make the agent's incentives perfectly aligned with those of the principal when it comes to the investment decision, at any given signal precision, they may provide too little power to incentivize the agent to make sufficient effort to acquire information. Unless the agent has all the bargaining power ($\beta = 0$), he knows he will have to share the fruits of his information-gathering efforts with the principal, and (being selfish by hypothesis) will exert less effort and thus obtain a less precise signal than what would be in both parties' best interest.

Maximization of the principal's expected profits under the constraint $\rho(\boldsymbol{w}) = \rho(\boldsymbol{r})$, that is, to perfectly align the two parties at the investment stage, does not reach the principal's optimal contract, at w^{**}, since it does not trade off the benefit of this alignment against the benefit to the principal of inducing the agent to acquire much information. This lack of incentive in stock-based contracts, in the diagram demonstrated for a risk-neutral agent, would be even more pronounced for a risk-averse agent, since such an agent will need a higher bonus and lesser penalty in order to set his signal threshold for investment at the principal's preferred level (or else there will be an interval of signals in which the principal would invest, if he had the information, but when the risk averse agent will not invest).

5.2. Option-based contracts. By hypothesis, the company's future stock value, $\Pi_0 + \Pi^*(\boldsymbol{w}, \tau)$, can exceed its status quo value, $\Pi_0 + \bar{\pi}$, only after investment in the

[25]This is obtained by setting $y = -\lambda(\Pi_0 + r_B)$.

good state of nature.[26] Hence, if the contract contains an option for the new CEO to buy a pre-specified amount of the company's future stocks, after the project has been realized but at the current stock price—that is, before the current CEO has been replaced—then the new CEO will exercise his option only after a successful investment. Hence, if such a standard option-based contract $\boldsymbol{w}$ consists of a fixed pay y and the option to buy a positive share λ of the company's future stocks, then $w_B = w_N = y$ and

$$w_G = y + \lambda (Ir_G - w_G) - \lambda\bar{\pi},$$

or, equivalently,

$$w_G = \frac{y + \lambda (Ir_G - \bar{\pi})}{1 + \lambda}, \tag{31}$$

Since there is no penalty in such a contract ($p = w_N - w_B = 0$), the CEO will always invest. Moreover, he thus has no incentive to acquire any information. Hence, such a standard option-based contracts are not used in any undominated contract.

Proposition 7. *All standard option-based contracts—a fixed pay plus an option to buy a pre-specified share of the company's future stock at its current price—are dominated both for risk-neutral and risk-averse agents.*

An alternative to a standard option—the right to buy future stocks at their current price—is to offer the agent the option to buy future stocks at a price below the current price. As an extreme case, the price could be zero, that is, the principal gives the option for free. As pointed out to us by Bengt Holmström, such an option is equivalent with giving the agent stocks. As we will see in the next subsection, equilibrium contracts may actually contain such free options alongside some standard options—a result that may appear surprising at first sight.

5.3. Mixed contracts. We saw above that all contracts that combine a fixed salary with some stocks are dominated, the reason being that, unless the agent is given all the stocks, he will under-invest in information acquisition. We also saw that all contracts that combine a fixed salary and some option to buy future stocks are dominated. There, the reason was that this leads the agent to acquire no information and to always invest. Expressed in terms of our model, the first class of contracts have a too low carrot-stick ratio while the second class have an infinite carrot-stick ratio. This suggests that perhaps contracts that contain both stocks and options could in fact be undominated.

In order to clarify this possibility, consider contracts that consist of a fixed salary y (to be paid in all circumstances), a share λ_1 of the company's future stock, and the

[26] If no investment is made, $\Pi (w, t) = -w_N \leq 0$. If investment is made in the bad state, $\Pi (w, t) = Ir_B - w_B \leq 0$.

option to buy an additional share λ_2 of the company's future stocks (both after the project has been realized). Such a mixed stock-option contract $\boldsymbol{w}$ has to satisfy

$$\begin{cases} w_B = (y + \lambda_1 \Pi_0) / (1 + \lambda_1) - \lambda_1 I |r_B| / (1 + \lambda_1) \\ w_N = (y + \lambda_1 \Pi_0) / (1 + \lambda_1) \\ w_G = (y + \lambda_1 \Pi_0) / (1 + \lambda_1 + \lambda_2) + (\lambda_1 + \lambda_2) I r_G / (1 + \lambda_1 + \lambda_2) - \lambda_2 \bar{\pi} \end{cases}$$

By varying the fixed pay y, the stock share $\lambda_1 > 0$ and the options share $\lambda_2 \geq 0$ in such a mixed contract $\boldsymbol{w}$, any carrot-stick ratio $\rho(\boldsymbol{w})$ from $\rho(\boldsymbol{r})$ and up, and any power $\kappa(\boldsymbol{w}) > 0$ can be obtained. Hence, without loss of generality, one may restrict attention to mixed contracts.

5.4. From binary to continuum return distributions. Our analysis has been focused on the case of a binary distribution for a project's return rate. This is, of course, a heroic simplification. In particular, it neglects the possibility of rare but disastrous outcomes. We here briefly sketch how the analysis can be adapted to the more realistic case of a continuum distribution. Suppose, thus that the noise in the agent's signal S is still additive, $S = X + \varepsilon$, and that the noise term ε is still statistically independent of the return rate X, but now X has a continuous distribution on the real line. A strictly monotonic contract is now an increasing and right-continuous function, $\boldsymbol{w} : \mathbb{R} \to \mathbb{R}$. Write W for the class of all such contracts. Under the MLRP (as restated for the continuum case), the agent's investment decision, given his signal precision τ, can still be represented in the form of a signal threshold such that it is optimal for the agent to invest if and only if his signal exceeds (or equals) this threshold.

To be more specific, given any contract $\boldsymbol{w} \in W$ and given any signal precision τ that the agent has chosen, it is optimal for the agent to invest if and only if his signal observation exceeds $s^*(\boldsymbol{w}, \tau)$, where the latter is either $-\infty$, meaning that he should invest for all signals, or it is $+\infty$, meaning that he should invest for no signals, or it satisfies the indifference condition

$$\mathbb{E}[\boldsymbol{w}(X) \mid S = s] = u(w_N).$$

Under MLRP, the left-hand side is a continuous and increasing function of the signal value s.

In the case of binary return-rate distributions, we defined a contract's carrot-stick ratio and could characterize the agent's signal threshold in terms of this ratio. This is no longer possible. However, a similar analysis as the one carried out above for the binary case seems feasible at least for some special cases. For example, if the return rate X is normally distributed with mean value μ and variance σ^2, and the noise term ε in the agent's signal is normally distributed with zero mean and variance $1/\tau$, then the agent's signal S will be normally distributed with mean-value μ and variance $\sigma^2 + 1/\tau$, permitting operational expressions for a risk neutral agent's expected utility and the principal's expected profit. In particular if contracts are restricted to be composed of a fixed salary, some stock share and some options, as in the mixed contracts just discussed, then it would seem possible to analyze equilibrium compositions of these three components.

6. 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 optimism. By the latter we mean biases in one's information gathering, such as attaching more importance to positive than negative news about a pet project.[27] Such overconfidence and/or optimism among managers has been documented in the literature, see Kahneman (2011, pp. 254-256), Daniel and Hirshleifer (2015), and Van den Steen (2011). Kindelberger (2005) has given a lively account of overconfidence among actors in financial markets. We here study the effects of overconfidence and optimism in the present model. For the sake of analytical tractability and brevity, we then focus entirely on the investment decision, and take the signal precision as exogenous and fixed. Hence, we will not study the effect of overconfidence on the agent's choice of effort in information acquisition. His talent is by assumption a complementary input to his effort in production of signal precision, and we saw before that an agent's signal precision is increasing in his (even if only self assessed) ability, see Proposition **??**. However, whether more able agents exert more or less effort is a priori an open question, and the answer depends on the income- and substitution effects, just as in models of labor supply. We focus throughout this section on the case of a risk neutral agent and a principal who has all the bargaining power.

6.1. Overconfidence. Suppose that the agent's signal precision is fixed at some level τ, and that the principal knows this precision level while the agent believes his signal precision to be $\tau' = (\gamma + 1)\,\tau$ for some $\gamma > 0$. In other words, the agent is biased, but the principal knows this. We call γ the agent's *degree of overconfidence* (a negative γ would represent under-confidence). The agent's overconfidence will benefit the (insightful) principal, since the agent will 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, and the principal can "tilt" the contract in such a way as to make the agent choose the principal's preferred signal threshold for the investment decision. What is the effect of such a behavioral bias on the contract and outcome?

[27] Clearly overconfident can arise out of such optimism concerning signals about one's own ability.

Since the agent by hypothesis is risk neutral, but overestimates his signal precision, the principal will offer a contract that just meets the agent's (subjective) participation constraint (that is, in the eyes of the agent). The principal will then set the bonus and penalty so that the agent will end up using the signal-threshold he would have used had he known his own signal precision. The set of feasible contracts is now defined by the principal's participation constraint and the agent's participation constraint, which now takes the form $\mathbb{E}\left[U\left(\boldsymbol{w},(\gamma+1)\,\tau\right)\right] \geq \bar{u}$. These constraints reflect the principal's knowledge of the agent's true signal precision and the agent's degree of overconfidence.

Proposition 8. *Consider an agent with fixed signal precision $\tau > 0$, known by the principal who also knows the agent's degree of overconfidence $\gamma > 0$. 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 $\hat{\boldsymbol{w}} = (0, y, y + by) \in W_0$, where*

$$b = \frac{r_G}{|r_B|} \cdot e^{\gamma \cdot \ln \rho(\boldsymbol{r})}, \tag{32}$$

and where $y > 0$ is uniquely determined by the agent's subjective participation constraint, met with equality, $\mathbb{E}\left[U\left(\hat{\boldsymbol{w}},(\gamma+1)\,\tau\right)\right] = \bar{u}$.

This result calls for some remarks. First, we note that although the agent has a biased estimate of his ability, his equilibrium signal threshold under the equilibrium contract $\hat{\boldsymbol{w}}$ is unaffected by this bias. The principal sets the carrot-stick ratio such that it induces the agent to use the signal threshold for investment that the principal desires. Second, we see in (**??**) that the more overconfident the agent, the lower will be his bonus (recall that $\rho(\boldsymbol{r}) < 1$ by hypothesis). Under the hypothesis of Proposition **??**, the equilibrium bonus rate b is thus exponentially decreasing in the agent's degree of overconfidence, and this rate is given by $\ln \rho(\boldsymbol{r}) < 0$.

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 overinvestment 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.

6.2. Optimism. Overconfidence is arguably a form of optimism concerning one's own ability. Another form of optimism (or pessimism) is when a decision-maker holds an unrealistically positive (negative) belief about the future success probability of a given project. In the present model, 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.[28] 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 non-zero. A positive (negative) mean value, η, represents an agent who has a tendency to over-sample favorable (disfavoring) information about his project. We will call η the agent's *degree of optimism*, with $\eta = 0$ being the bench-mark case of an unbiased agent.

Suppose that the principal knows the agent's degree of optimism. Formally, this can be inserted in our model by letting the noise term ε in the agent's private signal be distributed $N(\eta, 1/\tau)$ for some $\tau > 0$ and $\eta \in \mathbb{R}$ while the agent believes its distribution is $N(0, 1/\tau)$ (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 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?

The set of feasible contracts is now defined in terms of the accordingly adapted participation constraints, $\mathbb{E}[\Pi(\boldsymbol{w}, \tau, \eta)] \geq \bar{\pi}$ and $\mathbb{E}[U(\boldsymbol{w}, \tau)] \geq \bar{u}$, 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.

Proposition 9. *Consider a principal who knows that the agent's signal is* $\varepsilon \sim N(\eta, 1/\tau)$ *while the agent believes it is* $N(0, 1/\tau)$. *Suppose that the agent is risk neutral and the principal can make a take-it-or-leave-it offer. Then she will offer the contract* $\hat{\boldsymbol{w}} = (0, y, y + by)$, *where*

$$b = \frac{r_G}{|r_B|} \cdot e^{-\tau \cdot (r_G - r_B) \cdot \eta} \tag{33}$$

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

[28] Alternatively, 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.

This result calls for some remarks. First, we note that in the special case of an unbiased agent ($\eta = 0$) the bonus rate is $b = r_G / |r_B|$, that is, the agent's incentives when making the investments decision (given his private signal) are perfectly aligned with the principal's interest. Second, we see in (**??**) that the equilibrium bonus rate b is decreasing in the agent's signal bias, η. In other words, the more optimistic the agent is, the lower will the principal set the 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 (**??**) and (**??**), we see that the effects differ. In particular, the bonus to an agent with degree of optimism η is lower than the bonus to an agent with degree of overconfidence γ if and only if

$$\eta \cdot (r_G - r_B)\,\tau \;+\; \gamma \cdot \ln \rho\,(\boldsymbol{r}) \;>\; 0.$$

This inequality 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 offered a balanced agent, one who is neither optimistic nor pessimistic and has a correct belief about his own signal precision.)

6.3. Career concerns. Third, 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, but may increase the agent's incentive 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 our model 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.

7. 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 Bebchuk and Fried (2003), and Edmans and Gabaix (2009), and, for an early study of performance pay, Jensen and Murphy (1990). 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 (Rosen, 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 (2013), 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. (2012) 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 (Garicano 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 5.1 and 5.2 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.[29] 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 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)

[29] In 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 mangers are portrayed in the press.

and the literature reported there.[30]

8. 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 the preceding section, 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 principal-agent relations where the task of the agent—say, a manager of a corporation—is to make well-informed decisions on behalf of the principal—say, the owner(s) of a corporation or pension fund.

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.

The agent's information is 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. If the agent's information would be exogenously supplied, then 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. However, when the agent's effort to acquire information is endogenous, then the model contains a moral hazard problem; how much information should the agent acquire? As a consequence, all equilibrium contracts 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

[30] Otto (2014) also assesses each CEO's optimism with two measures (although different ones than those in the preeding 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.

these issues it turns out that what we call the carrot-stick ratio is a useful measure of the incentive balance of a contract. 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, it turns out from our numerical simulations that the effect of these two shifts are qualitatively the same. Other examples are the consequences for the set of undominated contracts of the agent's ability, the costs of information acquisition, and the size of the project or firm. We here assumed that the agent's ability (or talent) for information acquisition is a complementary production factor with his effort to acquire information, but we only studied the simplest such case. An interesting avenue for future research is to study the nature of contracts under more realistic information-production functions, with information technology as one input, the agent's ability (or talent) as another, and the agent's work effort as a third. We also studied the role of stock- and option-based contracts, and found that all contracts consisting of a fixed salary and some stocks in the project or company are Pareto dominated by other feasible contracts. While inducing the agent to make the right investment decisions, given his signal precision and private information, such stock-based contracts provide the agent with too weak incentives for information acquisition, since he will not reap all the benefits of more precise information but bear all the cost. However, contracts that consist of a mix of fixed salary, some stocks and some options (to buy future stocks at their current price), if correctly balanced, are undominated.

Finally, we apply our model to an analysis of the consequences of behavior bias such as overconfidence and overoptimism. If the principal knows about the latter's overconfidence, and if the agent's outside option is common knowledge between the principal and the agent, then the latter's overconfidence will benefit the principal. The reason is that the agent overestimates 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.

9. APPENDIX A: NORMAL NOISE

We here derive some basic results for the special case when the noise in the agent's signal is normally distributed as specified in equation (**??**). The agent may be risk-neutral or risk averse. One then obtains a clean expression of the agent's optimal signal cut-off, the critical signal level above which he invests. This critical level is anchored at the arithmetic mid-point between the good and bad returns to the project, and adjusted upwards or downwards from there, depending on the terms of the contract. This formula for the agent's signal threshold permits closed-form expressions for the probabilities for successful and failed investments.

Lemma 2. *Assume $\varepsilon \sim N(0, 1/\tau)$. Then the agent's optimal signal threshold under any strictly monotonic contract $\boldsymbol{w} \in W$ is*

$$s^*(\boldsymbol{w}, \tau) = \frac{r_G + r_B}{2} - \frac{\ln \rho(\boldsymbol{w})}{(r_G - r_B)\tau}, \tag{34}$$

*where the carrot-stick ratio $\rho(\boldsymbol{w})$ is defined in (**??**). Moreover, when he uses this investment strategy, the probability that a successful investment is made is*

$$p_G(\boldsymbol{w}, \tau) = \mu \cdot \Phi_1\left(\frac{\ln \rho(\boldsymbol{w})}{(r_G - r_B)\sqrt{\tau}} + \frac{r_G - r_B}{2}\sqrt{\tau}\right), \tag{35}$$

and that a failed investment is made is

$$p_B(\boldsymbol{w}, \tau) = (1 - \mu) \cdot \Phi_1\left(\frac{\ln \rho(\boldsymbol{w})}{(r_G - r_B)\sqrt{\tau}} - \frac{r_G - r_B}{2}\sqrt{\tau}\right). \tag{36}$$

The probability that no investment is made is thus $p_N(\boldsymbol{w}, \tau) = 1 - p_G(\boldsymbol{w}, \tau) - p_B(\boldsymbol{w}, \tau)$.

Proof: These expressions follow from (**??**), (**??**) and (**??**). In the present special case, (**??**) can be written as

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

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

or

$$\mu e^{-\tau(s-r_G)^2/2} \left[u(w_G) - u(w_N)\right] \geq (1-\mu) e^{-\tau(s-r_B)^2/2} \left[u(w_N) - u(w_B)\right]$$

or

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

or:

$$\ln \rho(\boldsymbol{w}) \geq \frac{1}{2}\tau(r_G - r_B)(r_G + r_B) - \tau(r_G - r_B)s$$

or

$$\tau(r_G - r_B)s \geq -\ln \rho(\boldsymbol{w}) + \frac{1}{2}\tau(r_G - r_B)(r_B + r_G)$$

or

$$s \geq \frac{1}{2}(r_G + r_B) - \frac{\ln \rho(\boldsymbol{w})}{(r_G - r_B)\tau}$$

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

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

and

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

Q.E.D.

Using the expressions in Lemma **??** we obtain a closed-form expression for the agent's interim expected remuneration utility under any contract $\boldsymbol{w} \in W$ and for any signal precision $\tau > 0$:

$$\begin{aligned} \mathbb{E}[U(\boldsymbol{w},\tau)] &= u(w_N) + \mu[u(w_G) - u(w_N)]\cdot\Phi_1\left(\frac{\ln \rho(\boldsymbol{w})}{(r_G - r_B)\sqrt{\tau}} + \frac{r_G - r_B}{2}\sqrt{\tau}\right) \\ &\quad - (1-\mu)[u(w_N) - u(w_B)]\cdot\Phi_1\left(\frac{\ln \rho(\boldsymbol{w})}{(r_G - r_B)\sqrt{\tau}} - \frac{r_G - r_B}{2}\sqrt{\tau}\right) \end{aligned} \quad (37)$$

General features of how this depends on the agent's signal precision are readily established. First, under any strictly increasing contract $\boldsymbol{w}$ this expected utility is, not surprisingly, strictly increasing in the agent's signal precision. Second, if the agent's signal precision is very high, then he will almost surely invest in the good state and not in the bad, so

$$\lim_{\tau\to\infty} \mathbb{E}\left[U\left(\boldsymbol{w},\tau\right)\right] = \mu u\left(w_G\right) + \left(1-\mu\right) u\left(w_N\right). \tag{38}$$

By contrast, as the agent's signal precision tends to zero, his interim expected utility *either* tends to his utility from the pay he receives when not investing *or* to his expected utility from investing without any signal, whatever is the best for him:

$$\lim_{\tau\to 0} \mathbb{E}\left[U\left(\boldsymbol{w},\tau\right)\right] = \begin{cases} u\left(w_N\right) & \text{if } \rho\left(\boldsymbol{w}\right) < 1 \\ \mu u\left(w_G\right) + \left(1-\mu\right) u\left(w_B\right) & \text{if } \rho\left(\boldsymbol{w}\right) \geq 1 \end{cases}, \tag{39}$$

where $u\left(w_N\right) > \mu u\left(w_G\right) + \left(1-\mu\right) u\left(w_B\right)$ if and only if $\rho\left(\boldsymbol{w}\right) < 1$. Hence, an uninformed agent will not gamble with the principal's money if $\rho\left(\boldsymbol{w}\right) < 1$, while he will gamble if $\rho\left(\boldsymbol{w}\right) > 1$. The latter situation is clearly undesirable for the principal, since then the agent is not better informed than the principal (and by assumption $\mathbb{E}\left[X\right] < r_0$).

Also the expected profit to the principal, defined in (**??**), can be expressed in terms of the primitives of the model for any contract $\boldsymbol{w} \in W$ and signal precision $\tau > 0$:

$$\begin{aligned} \mathbb{E}\left[\Pi\left(\boldsymbol{w},\tau\right)\right] &= \pi_N + \\ &+ \mu \cdot \left(\pi_G - \pi_N\right) \cdot \Phi_1\left(\frac{\ln \rho\left(\boldsymbol{w}\right)}{\left(r_G - r_B\right)\sqrt{\tau}} + \frac{r_G - r_B}{2}\sqrt{\tau}\right) \\ &+ \left(1-\mu\right) \cdot \left(\pi_B - \pi_N\right) \cdot \Phi_1\left(\frac{\ln \rho\left(\boldsymbol{w}\right)}{\left(r_G - r_B\right)\sqrt{\tau}} - \frac{r_G - r_B}{2}\sqrt{\tau}\right). \end{aligned} \tag{40}$$

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,

$$\lim_{\tau\to\infty} \mathbb{E}\left[\Pi\left(\boldsymbol{w},\tau\right)\right] = \mu \pi_G + \left(1-\mu\right) \pi_N.$$

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,

$$\rho\left(\boldsymbol{w}\right) < 1 \qquad \Rightarrow \qquad \mathbb{E}\left[\Pi\left(\boldsymbol{w},0\right)\right] = \pi_N.$$

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,

$$\rho(\boldsymbol{w}) > 1 \qquad \Rightarrow \qquad \mathbb{E}[\Pi(\boldsymbol{w}, 0)] = \mu\pi_G + (1-\mu)\pi_B.$$

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

Next, we obtain a closed-form expression for the agent's marginal interim expected utility from remuneration, showing exactly how it depends on the project returns and on terms of the contract. We believe the obtained qualitative properties holds for a wide class of unimodal noise distributions.

Lemma 3. *Assume $\varepsilon \sim N(0, 1/\tau)$. For any strictly monotonic contract $\boldsymbol{w} \in W$, the agent's interim expected utility from remuneration, $\mathbb{E}[U(\boldsymbol{w}, \tau)]$, is differentiable in his signal precision τ, and the associated marginal utility is given by*

$$\frac{\partial}{\partial\tau}\mathbb{E}[U(\boldsymbol{w}, \tau)] = \frac{r_G - r_B}{2} \cdot \frac{\kappa(\boldsymbol{w})}{\sqrt{2\pi\tau}} \cdot \exp\left[-\frac{1}{2\tau}\left(\frac{\ln\rho(\boldsymbol{w})}{r_G - r_B}\right)^2 - \frac{\tau}{2}\left(\frac{r_G - r_B}{2}\right)^2\right]. \quad (41)$$

In particular, this marginal utility is positive for all positive signal precisions $\tau > 0$ and tends to zero both when the signal tends to become virtually uninformative, $\tau \to 0$, and when the signal tends to become virtually perfectly informative, $\tau \to +\infty$.

Proof: We first write

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

where

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

Clearly $\mathbb{E}[U(\boldsymbol{w}, \tau)]$ is differentiable in τ for any $\tau > 0$, and

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

where

$$\frac{\partial}{\partial\tau}\Phi_1\left(\frac{\ln\rho(\boldsymbol{w})}{(r_G-r_B)\sqrt{\tau}}+\frac{r_G-r_B}{2}\sqrt{\tau}\right)=$$

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

and

$$\frac{\partial}{\partial\tau}\Phi_1\left(\frac{\ln\rho(\boldsymbol{w})}{(r_G-r_B)\sqrt{\tau}}-\frac{r_G-r_B}{2}\sqrt{\tau}\right)=$$

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

Hence,

$$A\cdot\frac{\partial}{\partial\tau}\mathbb{E}\left[U(\boldsymbol{w},\tau)\right]=$$

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

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

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

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

or

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

This establishes (**??**). **Q.E.D.**

10. APPENDIX B: MATHEMATICAL PROOFS OF OTHER RESULTS

Throughout this appendix, we will write $\boldsymbol{r} = (r_B, r_N, r_G)$ with $r_N = 0$.

10.1. Proposition ??. Assume that $\boldsymbol{w}^* \in W_F$ is strictly monotonic and let

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

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

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

where $G(\boldsymbol{w},\tau)$ is the expected value of the *sum of payments* to the two parties,

$$G(\boldsymbol{w},\tau) = \sum_{j=B,N,G} (\pi_j + w_j) \cdot p_j(\boldsymbol{w},\tau) = \sum_{j=B,N,G} r_j \cdot p_j(\boldsymbol{w},\tau),$$

where the last equality follows from $\pi_j = r_j - w_j$ for $j = B, N, G$. From Proposition **??**, we obtain that for a contract $\boldsymbol{w}$ to maximize $G(\boldsymbol{w},\tau)$ in the set W it is necessary and sufficient that $\boldsymbol{w}$ induces the same signal cut-off as would a risk-neutral agent under contract $\boldsymbol{r}$:

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

where (see (**??**))

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

Since ψ_τ^{-1} is strictly monotonic,

$$\boldsymbol{w}^* \in \arg\max_{\boldsymbol{w}\in W} G(\boldsymbol{w},\tau) \quad \Leftrightarrow \quad \frac{w_N^*-w_B^*}{w_G^*-w_N^*} = \frac{r_N-r_B}{r_G-r_N} \tag{42}$$

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

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

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

Second, in order to show that $\rho(\boldsymbol{w}^*) = \rho(\boldsymbol{r})$ is also necessary, suppose that $\boldsymbol{w}^*$ does not satisfy this equation. Then $s^*(\boldsymbol{w}^*,\tau) \neq s^*(\boldsymbol{r},\tau)$. Suppose that $s^*(\boldsymbol{w}^*,\tau) < s^*(\boldsymbol{r},\tau)$. (A similar argument can be developed for the case $s^*(\boldsymbol{w}^*,\tau) > s^*(\boldsymbol{r},\tau)$.) For any $\delta,\varepsilon \geq 0$ write $\hat{\boldsymbol{w}}(\delta,\varepsilon)$ for the contract $\hat{\boldsymbol{w}}$ 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 **??**, $s^*(\boldsymbol{w}^*,\tau) < s^*(\hat{\boldsymbol{w}}(\delta,\varepsilon),\tau) < s^*(\boldsymbol{r},\tau)$ for all $\delta \in (0,\bar{\delta})$. For all such δ, $G(\hat{\boldsymbol{w}}(\delta,y),\tau) > G(\boldsymbol{w}^*,\tau)$, irrespective of $\varepsilon \geq 0$ (since the signal threshold of $\hat{\boldsymbol{w}}(\delta,\varepsilon)$ 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 $\hat{\boldsymbol{w}}(\delta,y)$ than under contract $\boldsymbol{w}^*$. For $\varepsilon = 0$ the principal's expected payment under $\hat{\boldsymbol{w}}(\delta,\varepsilon)$ is higher than under contract $\boldsymbol{w}^*$ 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 $\hat{\boldsymbol{w}}(\delta,\varepsilon)$ equals that under contract $\boldsymbol{w}^*$. (This is always possible since the expected payment is continuous and strictly decreasing in ε). Then the expected payment to the agent is greater than under $\boldsymbol{w}^*$. By continuity, if ε is reduces slightly, both parties' expected payments under contract $\hat{w}(\delta,\varepsilon)$ is higher than under $\boldsymbol{w}^*$. Clearly $\hat{\boldsymbol{w}}(\delta,\varepsilon) \in B(\boldsymbol{w}^*)$. Since $G(\hat{\boldsymbol{w}}(\delta, y), \tau) > G(\boldsymbol{w}^*, \tau)$, $\boldsymbol{w}^* \notin \arg\max_{\boldsymbol{w}\in B(\boldsymbol{w}^*)} G(\boldsymbol{w},\tau)$. Since since $B(\boldsymbol{w}^*) \subset W$, $\boldsymbol{w}^* \notin \arg\max_{\boldsymbol{w}\in W} G(\boldsymbol{w},\tau)$.

10.2. Proof of Proposition ??. The existence claim follows from Weierstrass' maximum theorem, since the maximand in (**??**) is clearly continuous and, without loss of generality, we may assume that the set T is compact, either it is a singleton, a two-point set or $T = \mathbb{R}_+$, and in the latter case we use the fact that $U(\boldsymbol{w},\cdot)$ is bounded from above by $u(w_G)$ while $g(\cdot, a)$ is increasing and convex, hence unbounded. Thus, in this latter case there exists a $\tau > 0$ large enough to make the maximand negative, while the maximand is zero at $\tau = 0$.

10.3. Proof of Proposition ??. In order to establish (**??**), assume that $\boldsymbol{w}$ is a feasible and undominated contract with $w_B > 0$. Then $\boldsymbol{w}$ is an interior point in the closed cone W_0, so (by the observations made in the proof of Proposition **??**) $\boldsymbol{w}$ then maximizes $H(\boldsymbol{w}) = G(\boldsymbol{w}, \tau^o(\boldsymbol{w}, a)) - g(\tau^o(\boldsymbol{w}, a), a)$ in the interior of W_0. In particular, the signal threshold used at the investment decision point and given the precision $\tau = \tau^o(\boldsymbol{w}, a)$, need to maximize $G(\boldsymbol{w},\tau)$, which, by Proposition **??** requires $\rho(\boldsymbol{w}) = \rho(\boldsymbol{r})$ for a risk neutral agent. Moreover, when the agent chooses his signal precision, given $\boldsymbol{w}$, his choice $\tau = \tau^o(\boldsymbol{w}, a)$ necessarily satisfies (**??**) for $\boldsymbol{w} = I\boldsymbol{r}$ (as if he owned the project). Since $\rho(\boldsymbol{w}) = \rho(I\boldsymbol{r}) = \rho(\boldsymbol{r})$, it is also necessary that $\kappa(\boldsymbol{w}) = \kappa(I\boldsymbol{r})$. In sum, $\boldsymbol{w}$ has to solve the following two equations:

$$(w_G - w_N)(r_N - r_B) = (w_N - w_B)(r_G - r_N) \tag{43}$$

and

$$(w_G - w_N)(w_N - w_B) = I^2 (r_G - r_N)(r_N - r_B) \tag{44}$$

(for $r_N = 0$). Division of each side gives $I^2 (r_N - r_B)^2 = (w_N - w_B)^2$, or, equivalently, $I(r_N - r_B) = w_N - w_B$ Hence, $Ir_N - w_N = Ir_B - w_B$. Likewise, multiplication of each side gives $(w_G - w_N)^2 = I^2 (r_G - r_N)^2$ and hence $Ir_G - w_G = Ir_N - w_N$. 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 $Ir_B > w_B$, which implies $w_B < 0$, thus violating the limited-liability constraint. Hence, no contract $\boldsymbol{w}$ in the interior of W_0 is undominated. Hence, $w_B = 0$ is necessary.

10.4. Proof of Lemma ??. For any given contract $\boldsymbol{w} \in W$, let

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

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

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

for

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

Then $D'(\tau)$ has the same sign as

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

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

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

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

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

Thus

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

The exponent can be simplified:

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

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

Hence (**??**), and (**??**) has two solutions if and only if (**??**) holds. Moreover, the larger of these solution exceeds T, which gives (**??**).

10.5. Proof of Proposition ??. 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 $\varepsilon \sim N(\eta, 1/\tau)$ for some $\eta \in \mathbb{R}$. From Proposition **??** and Lemma **??** we deduce that the optimal signal-threshold then is

$$\hat{s}=\frac{r_G+r_B}{2}-\frac{\rho(\boldsymbol{r})}{\tau\cdot(r_G-r_B)}+\eta$$

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

$$\tilde{s} = \frac{r_G + r_B}{2} - \frac{\rho(\boldsymbol{w})}{(1+\gamma)\,\tau \cdot (r_G - r_B)}$$

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

$$\rho(\boldsymbol{w}) = (1+\gamma)\left[\rho(\boldsymbol{r}) - \tau\,(r_G - r_B)\,\eta\right]$$

Hence, the principal will set

$$b = \frac{w_G}{w_N} - 1 = \rho(\boldsymbol{r})^{\gamma} \cdot e^{-(1+\gamma)\tau\cdot(r_G - r_B)\cdot\eta}$$

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:

$$\left[1 - p_B\left(\hat{\boldsymbol{w}}, (1+\gamma)\,\tau\right)\right] \cdot y \;+\; p_G\left(\hat{\boldsymbol{w}}, (1+\gamma)\,\tau\right) \cdot b \cdot y \;=\; \bar{u}$$

or

$$y = \frac{\bar{u}}{1 - p_B\left(\hat{\boldsymbol{w}}, (1+\gamma)\,\tau\right) + b \cdot p_G\left(\hat{\boldsymbol{w}}, (1+\gamma)\,\tau\right)}$$

Under this contract $\hat{\boldsymbol{w}}$, the risk-neutral agent's signal cut-off will be $\tilde{s} = s^*(\boldsymbol{r}, \tau) + \eta$.

References

Axelson, U., and P. Bond (2015): “Wall Street occupations”, *Journal of Finance* 70, 1949-1996.

Axelson, U., and J. Weibull (2014): “Informed lending and credit crashes", mimeo. London School of Economics and Stockholm School of Economics.

Bakhija, J., A. Cole and B. Heim (2012): “Jobs and income growth of top earners and the causes of changing income ineqauality: Evidence from U.S. tax returns data”, Working paper, Indiana University.

Bebchuk, L., and J. Fried (2003): “Executive compensation as an agency problem”, *Journal of Economic Perspectives 17, 71-92.*

Bebchuk, L., and J. Fried (2004): *Pay without Performance: The Unfulfilled Promise of Executive Compensation.* Harvard University Press.

Bénabou, R., and J. Tirole (2013): "Bonus culture: Competitive pay, screening, and multitasking", CEPR Discussion Paper 9419.

Bertrand, M., and S. Mullainathan (2001): "Are CEOs rewarded for luck? The ones without prinipals are", *Quarterly Journal of Economics* 116, 901-932.

Böhm, M., D. Metzger and P. Strömberg (2015): "Since you're so rich, you must be really smart: talent and the finance wage premium", Sveriges Riksbank Working Paper series No. 313.

Carlin, B., and S. Gervais (2009) "Work ethic, employment contracts and firm value", *Journal of Finance* 64, 785-821.

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.

Custodio, C., M. Ferreira, and P. Matos (2013): "Generalists versus specialists: Lifetime work experience and CEO pay", J*ournal of Financial Economics* 108, 471-492.

Daniel, K., and D. Hirshleifer (2015): "Overconfident investors, predictable returns, and excessive trading", *Journal of Economic Perspectives* 29, 61-87.

DeMarzo, P., M. Fishman, Z. He, and N. Wang (2012): "Dynamic agency and the Q theory of investment", *Journal of Finance* 67, 2295-2340.

Dow, J., and C. Raposo (2005): "CEO compensation, change and corporate strategy", *Journal of Finance* 60, 2701-2727.

Edmans, A., and X. Gabaix (2009): "Is CEO pay really inefficient? A survey of new optimal contracting theories", *European Financial Management* 15, 486-496.

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.

Edmans, A., and X. Gabaix (2011): "The effects of risk on the CEO market", *Review of Financial Studies* 24, 2822-2863.

Edmans, A., X. Gabaix, T. Sadzik, and Y. Sannikov (2012): "Dynamic CEO Compensation", *Journal of Finance* 67, 1603–1647.

Gabaix, X., and A. Landier (2008): "Why has CEO pay increased so much?", *Quarterly Journal of Economics* 12, 49-100.

Garicano, L., and E. Rossi-Hansberg (2006): "Organization and inequality in a knowledge economy", *Quarterly Journal of Economics* 121, 1383-1435.

Gervais, S., J. B. Heaton, and T. Odean (2011): "Overconfidence, compensation, contracts and capital nudgeting", *Journal of Finance* 66, 1735-1777.

Giannetti, M. (2011): "Serial CEO incentives and the structure of managerial contracts", *Journal of Financial Intermediation* 20, 633-662.

Hall, B., and K. Murphy (2003): "The trouble with stock options", *Journal of Economic Perspectives* 17, 49-70.

Heller, Y. (2014): "Overconfidence and diversification", *American Economic Journal: Microeconomics* 6, 134-153.

Holmström, B., and S. Kaplan (2001): "Corporate governance and merger activity in the U.S.", *Journal of Economic Perspectives* 15, 121-144.

Holmström, B. (2004): "Pay without performance and the managerial power hypothesis: A comment", *Journal of Corporation Law* 30, 703-715.

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.

Jensen, M., and K. Murphy (1990): "Performance pay and top-management incentives", *Journal of Political Economy* 98, 225-264.

Jenter, D., and K. Lewellen (2014): "Performance induced CEO turnover", Research Paper 3054, Graduate School of Business, Stanford University.

Kahneman, D. (2011): *Thinking, Fast and Slow.* London: Allen Lane.

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.

Kaplan, S. and P. Strömberg (2004): "Characteristics, contracts, and actions: Evidence from vernture capitalist analyses", *Journal of Finance* 59, 2177-2210.

Kaplan, S. (2013): "CEO pay and corporate governance in the U.S.: Perceptions, facts and challenges", *Journal of Corporate Finance* 25, 8-25

Kaplan, S., and B. Minton (2012): "How has CEO turnover changed?", *International Review of Finance* 12, 57-87.

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.

Karlin S., and H. Rubin (1956): “The theory of decision procedures for distributions with monotone likelihood ratio”, *Annals of Mathematical Statistics* 27, 272-299.

Kindelberger, D. (2005): *Manias, Panics and Crashes: History of Financial Crises.* New York: Basic Books.

Kosfeld, M., and F. von Siemens (2011): “Competition, cooperation and corporate culture”, *Rand Journal of Economics* 42, 23-43.

Malcolmson, J. (2010): “Do Mangers with limited liability take more risky decisions? An informatioan acquisition model”, CESifo Working Paper No. 2943, February.

Malmendier, U., and G. Tate (2005): “Does overconfidence affect corporate investment? CEO overconfidence measures revisited”, *European Financial Management* 11, 649-659.

Malmendier, U., and T. Taylor (2015): “Behavioral CEO:s. The role of managerial overconfidence”, *Journal of Economic Perspectives* 29, 37-60.

Manso G. (2011): “Motivating innovation”, *Journal of Finance* 66, 1823-1860.

Mishel, L., and A. Davis (2014): “CEO pay continues to rise as typical workers are paid less”. Issue Brief, Economic Policy Institute, no. 380.

Milgrom, P. (1981): “Good news and bad news: Some representation results”, *Bell Journal of Economics* 12, 380-391.

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.

Nash, J. (1950), "The bargaining problem", *Econometrica* 18, 155-162.

Otto, C. A. (2014), “CEO optimism and incentive compensation”, *Journal of Financial Economics* 114, 366-404.

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.

Rosen, S. (1981): “The economics of superstars”, *American Economic Review* 71, 845-858.

Roth, A. (1979): *Axiomatic Models of Bargaining*, Lecture Notes in Economics and Mathematical Systems, 170. Springer Verlag.

Van den Steen, E. (2011): "Overconfidence by Bayesian Rational Agents", *Management Science* 57, 884-896.

Terviö, M. (2009): "Superstars and mediocrities: Market failure in the discovery of talent", *Review of Economic Studies* 76, 829-850.

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.