Options in Compensation: Promises and Pitfalls*

Christian Riis Flor[a]
University of Southern Denmark

Hans Frimor[b]
Aarhus University

Claus Munk[c]
Aarhus University

September 27, 2011

Abstract. We derive the optimal compensation contract in a principal-agent setting where outcome is used to provide incentives for both effort and risky investments. Optimal compensation entails rewards for good as well as bad outcomes, to motivate investment, and is increasing at the mean outcome to motivate effort. If rewarding bad outcomes is infeasible, option-based compensation is a near-efficient means of overcoming the manager's induced aversion against undertaking risky investments, whereas stock-based compensation cannot. However, option-based compensation may induce excessively risky investments and capping pay can be important in curbing such behavior.

Keywords. Induced moral hazard, incentives for investment, option-based compensation, excessive risk taking, regulation.

---

*We appreciate comments from Anil Arya, Peter Ove Christensen, Joel Demski, and John Fellingham.

[a]Department of Business and Economics, University of Southern Denmark, Campusvej 55, DK–5230 Odense M, Denmark. E-mail: crf@sam.sdu.dk.

[b]Department of Economics and Business, Aarhus University, Bartholin's Alle 10, DK–8000 Aarhus C, Denmark. E-mail: hfrimor@econ.au.dk.

[c]Department of Economics and Business & Department of Mathematics, Aarhus University, Bartholin's Alle 10, DK–8000 Aarhus C, Denmark. E-mail: cmunk@econ.au.dk.

Options in Compensation: Promises and Pitfalls

Abstract. We derive the optimal compensation contract in a principal-agent setting where outcome is used to provide incentives for both effort and risky investments. Optimal compensation entails rewards for good as well as bad outcomes, to motivate investment, and is increasing at the mean outcome to motivate effort. If rewarding bad outcomes is infeasible, option-based compensation is a near-efficient means of overcoming the manager's induced aversion against undertaking risky investments, whereas stock-based compensation is not. However, option-based compensation may induce excessively risky investments and capping pay can be important in curbing such behavior.

Keywords. Induced moral hazard, incentives for investment, option-based compensation, excessive risk taking, regulation.

An important element in Corporate Governance is the set of mechanisms which serve to influence management decisions when ownership and control are separate. One important mechanism is the compensation of managers. The growing proliferation of equity-based compensation indicates market prices increasingly serve as the informational basis for performance appraisal.[1] Public outrage over the generosity of compensation is commonplace and especially so in times of crisis where compensation can seem excessive when compared to current performance. Compensation practices have the potential to seriously impact the economy and are thus of importance not only to managers and shareholders but to society at large. The current crisis and its depth has been blamed on common pay practises, which allegedly have lead to greediness, excessive risk taking and short-sightedness in the financial sector.[2] The political systems in the US and the EU have responded by passing legislation regulating executive compensation and, in particular, the use of option based compensation.

Regulating the use of options is problematic if option compensation is efficient. In an agency setting with induced moral hazard we demonstrate that options closely resemble the optimal contracts designed to overcome a manager's aversion against undertaking risky investments. We demonstrate that often the expected payoff generated under option compensation is very close to the payoff generated under the optimal contract. However, we also show that option compensation can be highly inefficient. The inefficiency arises from the fact option compensation can give the agent an incentive to take on excessive risk, and an obvious question is how to curb such incentives. We show that restricting

[1]Hall and Murphy (2002) report that 94% of the S&P 500 companies issued options to management in 1999. The grant (Black-Scholes) value of options awarded constituted 47% of annual top management compensation and for the industrial companies in S&P 500 the median market value of top management equity instrument holdings was $31 million. Fahlenbrach and Stulz (2011) report that the end of 2006 mean market value of bank CEOs holdings of stock and stock options amounted to $87.5 million.

[2]See, e.g., Blinder (2009) and Bebchuck and Spamann (2009) as well as the memorandum "Interagency Guidance on Sound Incentive Compensation Policies" by the US Department of the Treasury and the European Union Commission proposal COM(2009)0362 available, respectively, at `http://www.ots.treas.gov/_files/25354.pdf` and `http://www.ipex.eu/ipex/cms/home/Documents/doc_COM20090362FIN`. Further see, e.g., Cassidy (2002) and Madrick (2003) regarding the alleged role of stock-based compensation in earlier corporate scandals.

compensation to consist of fixed salary plus restricted stock, i.e., a linear contract, can be a very costly solution. We show that an alternative and much more efficient solution is to use option compensation where total remuneration is capped, that is, by removing both the up- and the downside associated with restricted stock.

The main theoretical framework for studying optimal executive compensation is the principal-agent model. The risk-neutral principal (representing the shareholders of the company) cannot observe the actions of the risk- and effort-averse agent (the manager) so that a compensation scheme depending on some observable performance measure (e.g., earnings or stock price) is needed to motivate the agent. In most applications of that model, the optimal compensation scheme is concave in outcome (at least for reasonable parameters) and, hence, compensation in the form of stocks and, in particular, stock options is inefficient. The key assumption behind this result is that the manager can only affect the first moment of the outcome distribution. While this may be reasonable for most day-to-day decisions, top managers also make strategic decisions with more pronounced impact on the distribution of future outcomes. We extend the principal-agent setting by assuming that, in addition to an unobservable effort decision affecting the mean, the agent decides how to allocate capital between a risky productive investment opportunity and a riskless investment (say, in bonds). The investment decision affects both the mean and the variance of the outcome. The investment decision is assumed unobservable to the principal who, as a consequence, must use outcome to induce both effort and investment. The agent has no direct preferences over investments but, as compensation is used to induce preferences over effort, the investment decision is subject to induced moral hazard.

We show that the optimal unrestricted compensation scheme in this setting has a "short butterfly" shape so that the agent is handsomely rewarded for both very low and very high outcomes as "extreme" outcomes signal the agent undertook the desired risky investment.[3] Rewarding failure is potentially problematic: first, if the agent can increase variance without bounds it is not a viable strategy to reward low outcomes equally handsomely as high outcomes. Second, rewarding low outcomes is not a viable strategy if the manager can underreport or destroy outcome. The optimal non-decreasing compensation scheme consists of a fixed payment up to some outcome level above which compensation is increasing in outcome and eventually flattens out for very high outcomes. In particular, the optimal compensation is (locally) convex and resembles a combination of a fixed salary and a number of call options. Our numerical examples show that often very little is lost by implementing the best "fixed plus options" contract instead of the more sophisticated optimal non-decreasing contract. Given very little is lost by restricting the contract to be piecewise linear, one may suspect very little would be lost by further restricting the contract to be linear (restricted stock). We demonstrate this conjecture is erroneous. The loss from using linear contracts may be substantial due to the agent's induced aversion against valuable but risky investments. Consequently, our model provides theoretical support for the widespread use of stock options in executive compensation.

Incentive effects of stock-based compensation are often approximated by the sensitivity of the agent's certainty equivalent or market value of the contract to changes in current stock price or stock price volatility. Such measures have subsequently been employed in optimizing incentives within firms and in comparing incentives across firms. We demonstrate that both uses may be problematic. First, an agent may chose actions to decrease current stock price if exposed to an alternate contract which is more sensitive to current stock price – even when stock-price is measured gross of compensation. Second, two contracts with virtually zero sensitivity to stock price volatility may lead to substantially different investment decisions and thus resulting volatility. Third, an agent's reaction to an incentive scheme or changes in an incentive scheme is closely related to the production technology operated by the agent making a comparison of incentives across industries or

[3] Ross (2004) similarly suggests that risk-taking behavior might optimally be induced through granting of put as well as call options.

even between individual firms within the same industry difficult. These effects are driven by a surprisingly large effect of the tail ends of the outcome distribution. One might expect that granting a number of options issued far in the money would be essentially equivalent to granting the same number of restricted stocks. However, these alternatives will often not lead to comparable decisions by the agent, even though we assume outcome is normally distributed – a distribution characterized by thin tails.

The optimal non-decreasing contract trades off incentives for effort and for risk taking. In order to provide the manager incentives for exerting effort, the compensation must be increasing preferably around the mean outcome, which suggests the options should be issued in the money. On the other hand, that might conflict with the proper incentives for investments so that at-the-money or out-of-the-money options may emerge.[4]

It is well known that contracts derived employing the standard first-order approach to solving the principal-agent problem, may not induce the projected actions by the agent.[5] For some parameterizations of our model we find that when the agent is compensated with the seemingly optimal "fixed plus options" contract derived using the first-order approach his expected utility is only locally concave in the risky investment and starts to grow without bounds as the risky investment is increased above some level. This is most likely to happen when the production technology allows for investment to increase variance without too much harm to expected outcome. Such a production technology – combined with the agent's ability to obtain additional funding without the principal's knowledge – will lead to excessive risk taking via dramatic overinvestment in the risky project. The non-decreasing contract where compensation flattens out for very high outcomes is less susceptible to this problem, which suggests the problem with the "fixed plus options" contract can be resolved by appropriately capping total compensation to the

[4]Taxes are absent from our model but may clearly affect the design of compensation contracts in general and the moneyness of granted options in particular.

[5]Under the first-order approach, the agent's choice problem is represented by the first-order conditions from maximizing his expected utility of compensation with respect to the choice variables, in our case the effort level and the risky investment.

manager. We show this is the case and that capping compensation has the added benefit of making the contract more robust to misspecifications of the production function.

Our observations have implications for regulation. Following the financial crisis of 2008-2009, both the US and EU passed legislation regulating executive pay with the purpose of curbing future excessive risk taking and short-sightedness.[6] Under the US regulation, firms receiving assistance under the Troubled Asset Relief Program (TARP) are only allowed to use restricted stock for salary supplements and bonuses, whereas options are prohibited. The stocks so granted cannot vest while the company has obligations related to assistance under TARP. Golden parachute payments are prohibited; this also encompasses payments made in connection with a change in control of the company. As commanded by the subsequent Dodd-Frank Act, the federal bank regulators issued the "Interagency Guidance on Sound Incentive Compensation Policies" on June 21st, 2010. The guidance does not explicitly prohibit any form of compensation but outlines four methods to make compensation more sensitive to risk: (1) risk-adjusting compensation, (2) deferring payment of rewards so that the payments may be adjusted as risks are realized or become better known, (3) using longer performance periods, and (4) reducing the sensitivity of rewards to measures of short-term performance. The guideline further stresses that (i) banking organizations must ensure golden parachute arrangements do not encourage imprudent risk taking, and (ii) compensation should be reduced if an employee exposes the organization to material risk – as measured ex-ante.

In Europe, individual countries similarly regulated compensation as a condition for receiving national support and subsequently all financial institutions have been encompassed by EU regulation – of which parts may be extended to all listed companies. The EU regulation requires that at least 50% of variable remuneration must be composed of "shares or share-linked instruments" and contingent capital whose value reflects the credit quality of the institution. At least 40 % of the variable remuneration component

[6] Core and Guay (2010) and Murphy (2010, 2011) review the regulation and its consequences.

is deferred not less than three years and variable remuneration, including the deferred portion, is paid or vests only if it is sustainable according to the financial situation of the credit institution as a whole. Also, employees subject to the regulation are not allowed to hedge or otherwise insure compensation risks. Furthermore, financial institutions cannot reward failure, e.g., in the form of golden parachutes.

From the US and EU regulation it appears the hope is excessive risk taking can be prevented by making decision makers share "equally" in success and failure through, e.g., restricted stock, or share more in the total value of the enterprise through deferred compensation linked to corporate debt. Our analysis suggests remuneration consisting of fixed pay and restricted stock – as under TARP – may be excessively costly, as the decision maker will refrain from undertaking valuable risky investments. However, we also demonstrate compensation consisting of fixed pay and a number of options may incentivize the agent to undertake excessively risky investments and such compensation is not precluded under the general US or EU regulation. This constitutes a dilemma: if forcing executives to share equally in success and failure – as under the US TARP regulation – is too strong while the US/EU general regulation is too lax, how should remuneration then be regulated? According to our analysis the solution is to prevent decision makers from sharing equally in success and failure. That is, compensation should be bounded both below and above and should thus not entail harsh punishments or lavish rewards. More specifically our analysis suggests compensation consisting of options but where total compensation is capped will strike the right balance between allowing incentives for risk taking while effectively curbing excessive risk taking.

In sum, our paper makes the following main contributions. First, we extend the standard principal-agent framework to the more realistic case where the agent makes a strategic investment decision in addition to the effort decision, and the agent can hence affect both the mean and the variance of the outcome. Second, we show the optimal unrestricted compensation contract has a short butterfly form to motivate risky investments.

Third, we show that among the non-decreasing contracts, the optimal compensation scheme is locally convex. Fourth, we provide evidence little is lost by using a "fixed plus options" contract instead of the more complicated optimal compensation scheme. Fifth, we demonstrate that standard measures of the incentives provided by a compensation scheme are flawed when the manager controls investments. Related to this, we show the agent may have an incentive to take on excessive risk and that one way of curbing this excessive risk-taking behavior is by capping the option payoff.

The paper proceeds as follows. Section 1 links our analysis to the literature. Section 2 describes the principal-agent problem when the agent can affect outcome via both effort and investment choice. Section 3 discusses various outcome-dependent compensation contracts and reformulates the setting to stock-based compensation to facilitate comparison with the literature. Section 4 allows the agent to invest without bounds and discusses caps on options payoffs to curb incentives to overinvest. Section 5 concludes. All proofs are in the Appendix.

1 Ties to the literature

Within the standard principal-agent setting where the agent exerts only effort, the shape of the optimal contract is determined by the information content in the variable used in assessing the agent's performance (such as outcome or stock price) and the agent's utility function. If higher outcome is a signal the agent chose the desired action, then optimal compensation is increasing and often concave. For many distributions where effort affects only the mean outcome, higher outcome is a favorable signal, and then local convexity can arise if compensation is bounded below by a limited liability constraint (Lambert and Larcker 2004) or if the agent's preferences exhibit loss aversion (Dittmann, Maug, and Spalt 2010). Hemmer, Kim, and Verrechia (1999) show that convexity in compensation can arise endogenously if the variance of outcome directly depends on the agent's effort and the agent has power utility, but only for levels of risk aversion not empirically supported. Feltham and Wu (2001), restricting the choice of contracts to stock and at-the-money options, find that stock can be outperformed by at-the-money options when the variance of outcome directly depends on the agent's effort and the agent has mean-variance preferences. In our model, convexity is caused by induced moral hazard regarding an investment decision, the effects of which can be varied through our choice of technology. Further, in our model convexity does not rely on limited liability or the level of risk aversion.

It is well-known that the provision of incentives for one task (effort) might lead to an induced moral hazard problem regarding other tasks (investments), see, e.g., Feltham and Xie (1994). In Demski and Sappington (1987) and Lambert (1986) an agent privately acquires information and subsequently implements investments based on the information acquired in the planning stage. As the outcome of implementation is informative about the agent's planning activities, the principal is forced to use the outcome in influencing the agent. Induced moral hazard is present if the agent's incentives to affect the signal about his planning activities, i.e., implementation incentives, cannot be ignored. Our model is similar in the sense that effort and investment decisions interact, when investment affects the productivity of effort or when investments affects the variance of outcome.[7]

Some papers evaluate specific compensation structures without explicit modeling of the underlying incentive problems. A typical approach is to measure incentives provided by a contract as the sensitivity of the manager's certainty equivalent to current share price, see, e.g., Hall and Murphy (2002). However, when the agent also controls investments, this measure of incentives can be misleading. For example, we show that shifting the contract to one with higher incentives according to this measure may lead to both lower effort and lower investments with detrimental consequences for the principal.[8] Dittmann

[7]Lambert (1986) shows that incentive problems concerning the personally costly identification and analysis of investments and the subsequent investment decision can lead to convexity in the agent's compensation. In Lambert (1986) the agent has two investment alternatives. Our model allows independently variable unbounded effort and investment opportunity sets.

[8]Lambert, Larcker, and Verrecchia (1991) demonstrate that a risk-averse manager holding a call option can have an incentive to reduce variance – which in their setting does not harm expected outcome

and Maug (2007) compare actual executive compensation data (often with plenty of stocks and options) to the combination of fixed pay, stocks, and stock options that induces a given sensitivity – and thus supposedly a given effort level of the agent – at the lowest cost to the principal. As stock options are not part of the compensation scheme that optimally motivates the agent to provide effort when effort is the only concern, it is hardly surprising that Dittmann and Maug conclude that most observed contracts are highly inefficient. Given our theoretical and numerical results, real-life compensation contracts appear much more efficient.

A related problem arises when compensation is compared across firms or industries. For example, Core and Guay (2010) conclude that the compensation and incentive structures of bank CEOs are similar to that of CEOs of non-banks; see DeYoung, Peng, and Yan (2010) and Fahlenbrach and Stulz (2011) for related studies. According to our analysis, seemingly negligible differences in contracts or the production technology can lead to significantly different behavior. Hence, it is difficult to compare incentives even after controlling for observable technological differences.

If our modeling captures fundamental issues in providing incentives, one would expect firms whose CEO receives option-based compensation to invest more in risky assets and (consequently) to have a higher stock price volatility than firms whose CEO holds common stock.[9] A number of empirical papers provide evidence in line with these predictions. Sanders and Hambrick (2007) document that CEOs with large option holdings tend to undertake investments which are larger and lead to more extreme performance. Coles, Naveen, and Naveen (2006) find that higher sensitivity of CEO wealth to stock price volatility leads the CEO to invest more in R&D and less in property plant and equipment, and the CEOs tend to fund these investments by borrowing. Smith and

– despite the fact that this reduces the market value of the option (and of the company).

[9]Admittedly, our model does not directly lead to this latter conjecture as volatility per se does not lead to the use of options. Instead it is the marginal effect of investments on volatility that leads to the use of options, hence, using stock price volatility as the dependent variable would ideally require a control for differences in production technologies across firms.

Swan (2008) report evidence that "CEO incentives with option-based asymmetric payoffs greatly increase the likelihood that a firm will increase risk by undertaking both major real investments and acquisitions. In contrast, equity-based incentives that induce upside and downside symmetric payoffs are associated with fewer major acquisitions and neither encourages nor discourages real investments." Guay (1999) finds that the sensitivity of CEOs' wealth to equity risk is positively related to firms' investment opportunities. Similar evidence is provided by Rajgopal and Shevlin (2002) for gas and oil producers and by Mehran and Rosenberg (2007) for banks. The evidence given in these papers suggests that the risk-taking incentives of executive stock options are neither negligible nor small and, thus, should not be ignored in empirical analysis.

Finally, we emphasize that our principal-agent setting captures the effort and productive investment choice of a manager (agent) employed by an owner (principal) who, if well-diversified, can reasonably be assumed risk neutral towards the outcome. This contrasts the case where one or several portfolio managers (agents) undertake financial investment – and maybe information acquisition – decisions on behalf of an investor (principal) who is presumably risk averse to the outcome of the investment. The theoretical literature on delegated portfolio management is abundant (see survey by Stracca 2006) and addresses issues such as the incentives provided by specific compensation contracts and the design of optimal contracts (e.g., Stoughton 1993, Admati and Pfleiderer 1997, Carpenter 2000, Ou-Yang 2003, van Binsbergen, Brandt, and Koijen 2008), as well as equilibrium implications (e.g., Berk and Green 2004, Cuoco and Kaniel 2011).

2 A principal-agent model with effort and investment choice

We consider a one-period principal-agent model in which an important variation on the usual story is introduced: the agent has available a certain amount of working capital, which the agent invests in two distinct technologies, one risky and one riskless technology. We initially assume the optimal investment in the risky technology is less than the available working capital and that the principal can observe any borrowing by the agent. That is, the agent cannot invest in excess of the capital already at his disposal (the consequences of relaxing this assumption are discussed in Section 4). The principal observes neither effort choice nor the allocation of capital between the technologies. And thus, the only contracting variable available to the principal is future outcome, payoff or stock price resulting from the investment and production decisions made by the agent. The game begins at time $t = 0$, where a contract – specifying effort, capital allocation, and compensation – is agreed upon by the principal and the agent. The agent supplies productive effort $a \in A$ and allocates the available capital $\bar{q}$ between a risky investment q in a productive technology and a risk-free investment $q_f = \bar{q} - q$ in a financial type asset yielding the risk-free rate, which we normalize to zero throughout the paper. Subsequently both parties observe the outcome $x \in X$, and finally the agent receives remuneration $s(x)$ from the principal. Thus, we have the time line illustrated in Figure 1. Both the action set A, and the outcome set X are assumed convex.

[Figure 1 about here.]

The principal is risk neutral and thus maximizes the ex-ante market value which equals total expected outcome net of compensation to the agent. The agent's utility $U(s, a)$ depends both on effort a and on consumption, which we assume equals compensation $s = s(x)$. We assume that $U_a(s, a) < 0$, $U_{aa}(s, a) < 0$, $U_s(s, a) > 0$, $U_{ss}(s, a) < 0$ so that the agent is risk and effort averse. In addition we make the usual assumption that the agent's preferences are either multiplicatively or additively separable in a and s, that is

$$U(s, a) = -H(a) + J(a)V(s), \tag{1}$$

where either $H(\cdot)$ or $J(\cdot)$ is assumed constant, see, e.g., Grossman and Hart (1983).

The outcome x depends on productive effort and investments as well as on an unobservable state of nature. Let $g(x|a, q, \bar{q})$ denote the probability density function over x

conditioned on productive effort a, productive investment q, and the capital available for investments $\bar{q}$. We henceforth suppress $\bar{q}$ and denote the density $g(x|a,q)$. In most of the subsequent analysis we specify outcome as

$$x = \bar{q} - q + f(a,q) + h(q)^{1/2}\epsilon, \tag{2}$$

i.e., outcome is the sum of the payoff $q_f = \bar{q} - q$ from the risk-free investment and the risky payoff $f(a,q) + h(q)^{1/2}\epsilon$. We assume $\epsilon \sim N(0,\sigma^2)$, the production-induced outcome variance multiplier $h(q)$ is weakly increasing, and the mean payoff from the risky investment $f(\cdot)$ is strictly increasing and weakly concave.[10] One interpretation of the technology is that the agent's productive investment affects scale – and thus both mean and variance – while effort affects only the mean outcome. These assumptions imply that the density of outcome is

$$g(x|a,q) = \frac{1}{h(q)^{1/2}\sigma\sqrt{2\pi}} \exp\left\{-\frac{1}{2}\frac{(x - f(a,q) - (\bar{q}-q))^2}{\sigma^2 h(q)}\right\}. \tag{3}$$

In Section 3 we derive the optimal unrestricted compensation contract as a function of outcome and compare to the case where the investment level is contractible. We further derive optimal compensation contracts under various restrictions regarding the slope of the contract. In Section 3.6 we reformulate to compensation contracts depending on stock price and discuss the role of stock options in compensation.

3 Outcome-based compensation

The principal maximizes expected outcome net of compensation conditional on the effort and investment levels chosen by the agent to maximize his expected utility given the compensation contract. The program describing the Pareto optimal compensation

[10] That is, $f_a(\cdot)$, $f_q(\cdot) > 0$, $f_{aa}(\cdot)$, $f_{qq}(\cdot) \leq 0$, and $f_{aa}(\cdot)f_{qq}(\cdot) - f_{aq}(\cdot)^2 \geq 0$.

contract can be formulated as follows:

$$\max_{s(x),a,q} \int_{-\infty}^{\infty} [x - s(x)]g(x|a,q)\,dx \tag{4}$$

$$\text{s.t.} \qquad \int_{-\infty}^{\infty} U(s(x),a)g(x|a,q)\,dx \geq R, \qquad \text{(IR)}$$

$$(a,q) \in \operatorname*{argmax}_{\hat{a}\in A, \hat{q}\leq \bar{q}} \int_{-\infty}^{\infty} U(s(x),\hat{a})g(x|\hat{a},q)\,dx, \qquad \text{(IC)}$$

$$s(x) \geq \underline{s}, \qquad \text{(LL)}$$

where (IR) is the agent's individual rationality constraint given a reservation utility of R, (IC) is the agent's incentive compatibility constraint concerning effort and productive investment, and (LL) expresses the limited liability of the agent. We initially assume the standard first-order approach applies so that the (IC) constraint can be replaced by the associated first-order conditions

$$\int_{-\infty}^{\infty} \left[U_a(s(x),a) + U(s(x),a)\frac{g_a(x|a,q)}{g(x|a,q)}\right] g(x|a,q)\,dx = 0, \qquad (\text{IC}, a)' \tag{5}$$

$$\int_{-\infty}^{\infty} U(s(x),a)g_q(x|a,q)\,dx = 0, \qquad (\text{IC}, q)' \tag{6}$$

where subscripts on U and g indicate partial derivatives. We let λ, μ_a, and μ_q denote the Lagrange multipliers associated with the (IR), (IC,a)$'$, and (IC,q)$'$ constraints, respectively. The appropriateness of the first-order approach is discussed in Section 4.

3.1 The benchmark case: Investment level contractible

In standard agency models, investments are usually not modeled explicitly or are assumed to be contractible so that the principal effectively sets the productive investment level q. Consequently, (4) is solved without the incentive compatibility constraint concerning q. With separable preferences (1) and unspecified outcome distribution, pointwise maximization shows that the optimal compensation scheme $s(x)$ must satisfy

$$\frac{1}{U_s(s(x), a)} = \lambda + \mu_a \left[\frac{J'(a)}{J(a)} + \frac{g_a(x|a, q)}{g(x|a, q)} \right], \tag{7}$$

where this does not conflict with the limited liability bound. This compensation scheme is referred to as the *second-best* contract.[11] Compensation depends on the likelihood ratio g_a/g which reflects how strongly the effort choice affects the outcome distribution.

It follows that utility functions and distributional assumptions can be combined to yield convex compensation schemes. For example, if $1/U_s(s(x), a)$ is concave (convex) in $s(\cdot)$ and the right-hand side of (7) is affine in x, then $s(x)$ is convex (concave).

With normally distributed production (2) with density given by (3), we see that

$$\frac{g_a(x|a, q)}{g(x\,|a, q)} = \frac{x - \mathrm{E}\,[x|a, q]}{\sigma h(q)^{1/2}} \frac{f_a(a, q)}{\sigma h(q)^{1/2}}, \tag{8}$$

so that the first-order condition on compensation becomes

$$\frac{1}{U_s(s(x), a)} = \lambda + \mu_a \left[\frac{J'(a)}{J(a)} + \frac{x - \mathrm{E}[x|a, q]}{\sigma h(q)^{1/2}} \frac{f_a(a, q)}{\sigma h(q)^{1/2}} \right]. \tag{9}$$

Note the likelihood ratio concerning effort g_a/g is increasing in x and thus satisfies the monotone likelihood ratio property (MLRP). Assume $1/U_s(s(x), a)$ is convex in $s(x)$, which is satisfied by CARA utility (which we employ later) and CRRA utility exhibiting a relative risk aversion above 1. With this assumption we have maximized the odds against convex compensation. When $1/U_s$ is convex the compensation scheme $s(x)$ is an increasing and concave function as long as the lower bound on compensation is not binding.[12] This leads to two observations. First, higher outcome is considered good news

[11] The *first-best* contract is the optimal contract for the problem where the incentive compatibility constraint (IC) is not imposed (or is not binding).

[12] Without the limited liability constraint an optimal solution might not exist. According to Mirrlees (1999) it is possible to get arbitrarily close to first-best, the socalled Mirrlees problem. Being casual the problem is the right-hand side of (9) is affine in x. It follows μ_a cannot be positive, since if it is, values of x exist for which the right hand side of (9) and thus marginal utility becomes negative.

(follows from MLRP), which is natural as high outcomes indicate more strongly that the agent did behave. Second, compensation is – where this is interior – concave in outcome x and local convexity is present only due to the limited liability constraint.

For low level employees the limited liability constraint might be of first-order importance, yet we doubt this is the case for high level employees let alone for executives in anything but the smallest of companies. Empirical evidence supports this conjecture as most executive compensation packages consist of a relatively high base with option incentives on top (Dittmann and Maug 2007). Hence, the convexity induced by a limited liability constraint is not a compelling case for convexity in executive compensation. In the remainder of the paper we will allow compensation to be unbounded.[13] Thus, instead of relying on limited liability, we will show below that convexity in compensation can be optimal when induced moral hazard via an unobservable investment decision is present.

3.2 Optimal unrestricted compensation: Investment level not contractible

From now on we assume the investment level is not contractible so that the incentive compatibility constraint concerning productive investment must be taken into account when determining the compensation scheme. With normally distributed outcome (2),

$$\frac{g_q(x|a,q)}{g(x|a,q)} = \frac{x - \mathrm{E}[x|a,q]}{\sigma h(q)^{1/2}} \frac{f_q(a,q) - 1}{\sigma h(q)^{1/2}} + \left\{ \left[\frac{x - \mathrm{E}[x|a,q]}{\sigma h(q)^{1/2}} \right]^2 - 1 \right\} \frac{1}{2} \frac{h'(q)}{h(q)}. \tag{10}$$

That is, the likelihood ratio concerning investment g_q/g is a second-order polynomial and a strictly convex function of outcome x. In this case the first-order approach leads to the following characterization of the optimal compensation scheme, which we will also refer to as the second-best contract.

Proposition 1 *Assuming the first-order approach is valid and the incentive constraints on the moral hazard and the capital allocation constraints are binding, the compensation*

[13] Had we carried along the assumption of limited liability, we would have added an 'where this does not conflict with the limited liability constraint' in each proposition where we characterize optimal contracts.

scheme $s(x)$ must satisfy

$$\begin{aligned}\frac{1}{U_s(s(x),a)} &= \lambda+\mu_a\left[\frac{J'(a)}{J(a)}+\frac{x-\mathrm{E}[x|a,q]}{\sigma h(q)^{1/2}}\frac{f_a(a,q)}{\sigma h(q)^{1/2}}\right] \qquad (11)\\ &+\mu_q\frac{x-\mathrm{E}[x|a,q]}{\sigma h(q)^{1/2}}\frac{f_q(a,q)-1}{\sigma h(q)^{1/2}}+\mu_q\left\{\left[\frac{x-\mathrm{E}[x|a,q]}{\sigma h(q)^{1/2}}\right]^2-1\right\}\frac{1}{2}\frac{h'(q)}{h(q)}.\end{aligned}$$

As $1/U_s(s(x),a)$ is convex, the compensation scheme is an increasing concave function of the weighted likelihood ratios, and if – as we assume – $h'(q)>0$, the right-hand side of (11) is quadratic in x. For positive μ_q, the right-hand side of (11) is a convex second-order polynomial and thus symmetric around a global minimum $\hat{x}$. The compensation scheme $s(x)$ satisfying (11) is a function of the right-hand side and therefore also symmetric around $\hat{x}$, where compensation obtains its minimum. Hence, low as well as high outcomes are rewarded, and the compensation scheme is locally convex (also, the Mirrlees problem no longer exists). Furthermore, we show in the Appendix that $s(x)$ is increasing at the mean outcome, $\mathrm{E}[x|a,q]=f(a,q)+(\bar{q}-q)$. We summarize these properties as follows.

Proposition 2 *Assume that the first-order approach is valid and that outcome is given by* (2) *with h strictly increasing. The optimal unrestricted compensation scheme $s(x)$ has the following properties:*

1. *compensation has a minimum at some outcome $\hat{x}$, is decreasing for all $x<\hat{x}$ and increasing for all $x>\hat{x}$,*
2. *compensation is symmetric around $\hat{x}$, i.e., $s(\hat{x}-\delta)=s(\hat{x}+\delta)$ for all $\delta>0$,*
3. *compensation is increasing at the mean outcome, i.e., $s'(\mathrm{E}[x|a,q])>0$.*

As demonstrated by Holmström (1979), it is the informativeness of the outcomes that determines pay. Locally – around the mean – lower outcomes signal the agent slacked off and/or underinvested and thus, locally, lower outcomes are considered bad news regarding effort and investment decisions. In contrast, low as well as high outcomes

– far from the mean – signal that the agent undertook the appropriate risky investment. An interpretation is that the agent is granted insurance in order to induce sufficient risk taking. This is in line with Ross (2004) who suggests that risk-taking behavior might optimally be induced through granting of put as well as call options. In our setting the sign of $s''(x)$ is determined by a second-order polynomial in x, thus as $1/U_s$ is sufficiently convex the optimal compensation contract takes the form of a "short butterfly:" $s(x)$ is symmetric, convex in a neighborhood around the minimum, and concave in the tails.

Let us be more explicit about the induced moral hazard problem in our setting. If both effort and investment are contractible, the principal can write a forcing contract and thus the optimal compensation contract is independent of outcome, i.e., $s(x)$ is constant. As the agent has no direct preferences over investments, the same will hold assuming effort, but not investment, is contractible. Regardless of the fact the agent has no direct preferences over investments, there is no guarantee that dealing with the effort incentive problem will lead to efficient investment decisions. The optimal compensation scheme ignoring the investment, but not the effort, decision is characterized by (9). However, this scheme might lead the agent to undertake undesirable investment decisions. In fact, if the compensation schemes characterized in (9) and (11) are not identical almost everywhere, we have an induced moral hazard problem concerning the investment decision.

In general it is difficult to determine whether the agent – if offered a contract satisfying (9) – has an incentive to reallocate investments, that is whether we have an induced moral hazard problem. However, if investment affects neither marginal productivity of effort nor variance, then induced moral hazard is absent as summarized by the following proposition.[14]

Proposition 3 *Assume that outcome is given by* (2)*. If* $f_{aq}(a, q) = 0$ *and* $h'(q) = 0, \forall a, q$*, then there is no induced moral hazard.*

[14]Demski and Sappington (1987) also derive sufficient conditions to ensure the absence of induced moral hazard.

The lack of a variance effect (so that $h'(q) = 0$) is generally not sufficient to avoid induced moral hazard. Formally, this follows from (11) as only the last term related to unobservable investment vanishes, while the first term (containing μ_q) is still present. The intuition is as follows: One component of agency costs is the second-best action loss encountered due to the agent choosing second-best effort different from first-best effort. Hence, if investment increases marginal productivity of effort, then overinvestment might aid in controlling effort and thus reduce the second-best action loss. Therefore, it is not optimal for the principal to insist on "optimal" investment, i.e. $f_q(a;q) - 1 = 0$, and induced moral hazard prevails.[15] Thus, unobservable productive investments affecting variance generally imply induced moral hazard.

3.3 Non-decreasing compensation

When the investment level is unobservable, low as well as high outcomes are rewarded in order to induce risky investments. However, when the agent can affect outcome by other means than effort, it might be problematic to implement compensation schemes which are decreasing. First, if the agent can destroy or underreport outcome, rewarding low outcomes is not a viable approach (Dye 1988). Second, rewarding high as well as low outcomes may well invite overinvestment. When second-best compensation takes the form of a "short butterfly," the agent obviously has an incentive to overinvest if that leads to extreme variance; overinvestment can be interpreted as a form of asset substitution. If, in addition, overinvestment leads to (very) low expected outcome, i.e., if overinvestment – in expectation – constitutes destruction of outcome, it is only an added benefit.

Technically, the problem is that the agent's expected utility with such a "short butterfly" contract can be non-monotonic in the productive investment q, so that the first-order approach will lead to only a local optimum while expected utility will become even higher

[15] The combination of weakly increasing compensation – as implied by (9) with investment level contractible – and investments ranked according to first-order stochastic dominance leads to value maximizing investments so that $f_q(a;q) - 1 = 0$, cf. Demski and Sappington (1987). However, with monotonous compensation it is impossible to induce overinvestment and, hence, induced moral hazard persists.

for investment levels above some $\hat{q}$ and then keep growing as q is increased. If $\hat{q}$ is below the total capital made available by the principal, $\bar{q}$, or the agent can obtain sufficient additional funds from the capital market without the principal observing it, the agent will deviate from the actions that the "short butterfly" contract was supposed to induce. In the following, we will generally assume that this problem does not arise with non-decreasing contracts, but we return to the discussion in Section 4.

We will refer to the optimal non-decreasing compensation scheme for the case with unobservable investment as the *third-best* contract. Note that this has to be computed by numerical optimization, but a natural conjecture is that it is flat for outcome levels up to some specific level x' near the point $\hat{x}$ where the optimal unrestricted compensation is minimal and then fairly close to the unrestricted optimal consumption for higher levels and, in particular, increasing in outcome. Intuitively, such a compensation scheme will bear some resemblance to an option contract. We elaborate on this below.

3.4 Numerical illustrations

In our numerical illustrations we assume CARA (constant absolute risk aversion) utility

$$U(s,a) = -e^{-r[s-C(a)]}, \tag{12}$$

corresponding to $H(a) = 0$, $J(a) = e^{rC(a)}$, and $V(s) = e^{-rs}$. Here r is the degree of absolute risk aversion, and $C(a)$ is personal (monetary) cost of effort, assumed increasing and weakly convex, i.e., $C'(a) > 0$ and $C''(a) \geq 0$.

The CARA-normal setting is standard in the principal-agent literature. An alternative used in some papers on stock-based compensation is the CRRA-lognormal setting where the agent has CRRA (constant relative risk aversion) utility and outcome or stock price is lognormally distributed. While tradition, perceived realism, or computational convenience may lead one to prefer one version over the other, the fact remains that the two versions are very close. A model with lognormally distributed stock prices is easily transformed to its normal version by contracting on the log of price instead of contracting on the price directly. We choose the model with is normally distributed outcome as the effects of effort and investment on mean and variance are more distinct in that setting than in the lognormal equivalent. The lognormal formulation is presented in Appendix B.

CARA utility is computationally convenient as it ensures that optimal compensation is unaffected by the agent's market alternative as stated in the following lemma.[16]

Lemma 1 *Assuming the agent has CARA utility* (12) *and the solution is interior, then the optimal contract is unique up to a positive constant.*

As there are no wealth effects, we can rewrite the compensation function:[17]

$$s(x) = \frac{1}{r}\ln\lambda + \frac{1}{r}\ln F(x,a,q) + C(a) + \frac{1}{r}\ln r,$$

where

$$\begin{aligned} F(x,a,q) = 1 &+ \frac{\mu_a}{\lambda}\left[rC'(a) + \frac{(x - \mathrm{E}[x|a,q])f_a(a,q)}{\sigma^2 h(q)}\right] \\ &+ \frac{\mu_q}{\lambda}\left[\frac{x - \mathrm{E}[x|a,q]}{\sigma h(q)^{1/2}}\frac{f_q(a,q) - 1}{\sigma h(q)^{1/2}} + \left\{\left[\frac{x - \mathrm{E}[x|a,q]}{\sigma h(q)^{1/2}}\right]^2 - 1\right\}\frac{1}{2}\frac{h'(q)}{h(q)}\right], \end{aligned}$$

which is convenient for our numerical implementation as we can concentrate on $\hat{\mu}_a = \mu_a/\lambda$ and $\hat{\mu}_q = \mu_q/\lambda$ and subsequently add a constant to satisfy individual rationality.

In our numerical examples, we consider a production function of the form

$$f(a,q) = ka^\theta q^{1-\theta},$$

where k is a general productivity scaling parameter, and $\theta \in (0,1)$ governs the relative productivity of effort and investment. We assume that the production-induced variance

[16]The result holds regardless of whether the first-order approach is valid or not.

[17]$\frac{1}{r}\ln(\cdot)$ is the inverse of $1/U_s(\cdot)$

of outcome is of the form $h(q) = q^\zeta$ for some $\zeta > 0$ and that the agent's cost of effort is $C(a) = a^2$. Our base case parameters are listed in Table 1.

[Table 1 about here.]

The first-best solution (assuming effort is contractible) is $a = 0.844$ and $q = 4.271$, leaving the principal an expected profit of 0.712. The second-best solution (inducing both effort and productive investment with an unrestricted compensation scheme) entails an effort and investment choice of respectively 0.413 and 2.420, leaving the principal an expected profit of 0.379237. The optimal compensation scheme inducing the specified effort and investment decision is depicted as the dashed curve in Figure 2 and has the short butterfly form. The third-best solution (inducing effort and productive investment with a non-decreasing compensation scheme) entails effort and investment choice of respectively 0.413 and 2.420, leaving the principal an expected profit of 0.379182. Virtually nothing – 0.01% – is lost by restricting the contract to be nondecreasing. The optimal non-decreasing compensation scheme is also depicted in Figure 2 which confirms that the two schemes are very close as long as outcome is within two standard deviations from the mean. Depending on where the endpoint of the flat part of the compensation function is located, the strictly increasing part can be both locally convex and locally concave. We emphasize that the option-resembling compensation is not due to the limited liability constraint, but is caused either by the fact the agent can destroy outcome or by the fact the agent can overinvest – or both.

[Figure 2 about here.]

3.5 Piecewise linear compensation

Figure 2 suggests that, for outcomes near the mean, the optimal non-decreasing compensation scheme is well approximated by a piecewise linear contract, i.e., a "fixed pay plus options" contract. This suggests that little is lost by optimizing only over piecewise linear contracts,

$$s(x) = \kappa_0 + \kappa_1 \max\{x - \kappa; 0\}, \tag{13}$$

for some constants $\kappa_1 > 0$, κ_0, and κ, so that compensation consists of a fixed payment and a bonus payment equal to a fraction κ_1 of the outcome exceeding κ. In our base case with the parameters listed in Table 1, it turns out that the following contract is optimal in the class of piecewise linear contracts:

$$s(x) = -1.764 + 0.550 \max\{x - (-2.755); 0\}. \tag{14}$$

This contract is depicted alongside the optimal non-decreasing contract in Figure 3. This juxtaposition reveals that the compensation schemes are very similar for outcomes up to two standard deviations around the mean outcome. The similarity is also evident when effort levels, investments and expected payments are considered.

[Figure 3 about here.]

The piecewise linear contract induces effort and investment levels of $a = 0.413$ and $q = 2.417$, which are close to the effort and investment levels of $a = 0.413$ and $q = 2.420$ induced by the optimal non-decreasing contract. Likewise, the piecewise linear contract will provide the principal with an expected net payoff of 0.378 (net of expected compensation cost of 0.308), whereas the optimal non-decreasing contract yields an expected net payoff of 0.379 (net of expected compensation cost of 0.290). At least in this particular example not much – 0.3% – is lost by restricting the contractual form.

One may erroneously infer that restricting the contract to the class of linear contracts, $s(x) = \ell_0 + \ell_1 x$, will not lead to any significant loss of welfare either. This is not so. Under the optimal linear contract the principal's surplus is reduced by more than 12% to 0.331 (the optimal linear contract is $-0.025 + 0.498x$, induces effort $a = 0.366$ and investment

$q = 1.540$, and the expected compensation cost is 0.279). If the endogenous variance function is $q^{0.9}$ as opposed to $q^{1/4}$, then the principal's payoff is reduced by 36% (from 0.321 to 0.205) when a linear contract is employed instead of a piecewise linear contract. The effect of restricting contracts to the linear class is surprisingly large considering that the normal distribution is characterized by thin tails and the fact the optimal piecewise linear contract is well approximated by a linear contract near the mean outcome.[18] The disadvantage of linear contracts is that the agent is unwilling to undertake the appropriate investments. Our numerical illustrations indicate that for more severe induced moral hazard problem, the loss caused by restricting the contract to be linear becomes higher.

3.6 Stock-based compensation

The preceding section discussed the design of compensation schemes as a function of outcome, which could be sales, cost reductions, cash flow, some variation of earnings, or some other informative variable. Exchange-listed companies frequently implement schemes where compensation depends on the stock price of the company, typically by allocating stocks and/or stock options to the manager making the scheme piecewise linear. Moreover, the literature counts numerous papers discussing stock-based compensation. To compare both with that literature and the wide-spread practice, we will now study the consequences of our setting and assumptions for stock-based remuneration.

First, we address the interesting question of how to summarize the incentive effects of a compensation contract involving stocks and/or options. It is standard in the literature to measure the incentives provided by a compensation contract involving stocks and/or options by the sensitivity of the Black-Scholes value or the agent's certainty equivalent with respect to the current stock price, see, e.g., Hall and Murphy (2000) and Dittmann and Maug (2007). Others, e.g., Core and Guay (2010), add a measure of risk-taking

[18]If one were to study the agent's incentives to alter variance under the piecewise linear scheme and under the linear scheme, one may erroneously infer the agent would have identical incentives to alter investment under the linear scheme, for example if the agent's expected utilities are close to being equally sensitive to variance under the two compensation schemes. This can be highly misleading as the incentives for changing investment can be substantially different in the two situations.

incentives namely the sensitivity of either the value of stock and option holdings or the agent's certainty equivalent with respect to volatility.

In our setting both tasks favorably affect expected future stock price and hence also current stock price. One may thus believe that altering a contract to increase the sensitivity of the certainty equivalent to current stock price will lead to an increase in effort as well as investment and as a consequence in intrinsic value.[19] This intuition may hold, but we also show that increasing the sensitivity of the certainty equivalent with respect to the current stock price can lead to the agent decreasing both effort and investment (we can also provide examples of the reverse as well as examples of the agent increasing effort and decreasing investment). The problem, of course, is that this measure of incentives ignores the agent's ability to affect variance.

The intuition is best developed by assuming that stock price equals outcome gross of compensation such that $P_1 = x$ and $P_0 = \mathrm{E}[P_1]$; the latter equality is due to our assumptions of investor risk-neutrality and a zero discount rate. This implies $P_1 = P_0(a,q) + h(q)^{1/2}\varepsilon$ with probability density function

$$\hat{g}(P_1|P_0(a,q),q) = \frac{1}{h(q)^{1/2}\sigma\sqrt{2\pi}} \exp\left\{-\frac{(P_1 - P_0(a,q))^2}{2h(q)\sigma^2}\right\}.$$

For any density function fully characterized by P_0 and q, the first-order condition wrt. effort is[20]

$$\begin{aligned}\int_{-\infty}^{\infty} U(\hat{s}(P_1),a)\hat{g}_{P_0}(P_1|P_0)\frac{\partial P_0}{\partial a}dP_1 + \int_{-\infty}^{\infty} U(\hat{s}(P_1),a)rC'(a)\hat{g}(P_1|P_0)dP_1 \\ = \frac{\partial}{\partial P_0}\,\mathrm{E}_{P_1}[U(\hat{s}(P_1),a)|P_0]\frac{\partial P_0}{\partial a} + \mathrm{E}_{P_1}[U(\hat{s}(P_1),a)|P_0]rC'(a) = 0,\end{aligned}$$

and it thus makes sense to measure incentives for *effort* as the derivative of the agent's certainty equivalent wrt. current stock price, $\partial\,\mathrm{E}_{P_1}[U(\hat{s}(P_1),a)|P_0]/\partial P_0$. Note that if the agent is risk neutral, if P_1 is lognormally distributed, and if compensation consists of

[19] As in Feltham and Xie (1994) we are trying to control two tasks with a single performance measure. Hence it should come as no surprise there are combinations of effort and investment levels, which cannot be induced with a piecewise linear contract or with any other contract for that matter.

[20] As we are taking expectations over P_1, the first term cannot be written as $\mathrm{E}_{P_1}[\partial U(\cdot)/\partial P_0]\partial P_0/\partial a$.

a fixed component and a number of options with exercise price $\hat{\kappa}$, i.e., $\hat{s}(P_1) = \kappa_0 + \Delta \max\{P_1 - \hat{\kappa}; 0\}$, then the derivative of the certainty equivalent wrt. current stock price is just $\Delta N(d_1)$ where $N(d_1)$ is the option delta from the Black-Scholes formula.

The agent's first-order condition wrt. investment,

$$\frac{\partial}{\partial P_0} \mathrm{E}_{P_1}[U(\hat{s}(P_1), a)|P_0] \frac{\partial P_0}{\partial q} + \frac{\partial}{\partial h(q)\sigma^2} \mathrm{E}_{P_1}[U(\hat{s}(P_1), a)|P_0] \frac{\partial h(q)\sigma^2}{\partial q} = 0,$$

depends on the derivative of the expected utility (or certainty equivalent) wrt. stock price P_0, but also on the derivative of the expected utility (or certainty equivalent) wrt. variance, the "vega" of the compensation when the agent is risk neutral (and stock price is lognormal). Thus, when the agent can control investments which affect both mean and variance, the derivative of the certainty equivalent wrt. variance is an important (partial) measure of investment incentives. This term clearly depends on the agent's risk aversion and the extent to which the agent can affect variance.

If incentives could be measured as the sensitivity of the certainty equivalent or the value of the contract to changes in current price, then the agent would exert higher effort and invest more under the contract $\tilde{s}(P_1)$ than under the contract $\hat{s}(P_1)$ if

$$\frac{\partial \, \mathrm{E}_{P_1}[U(\hat{s}(P_1), a)|P_0]}{\partial P_0} < \frac{\partial \, \mathrm{E}_{P_1}[U(\tilde{s}(P_1), a)|P_0]}{\partial P_0}.$$

Using a numerical example, we will demonstrate the futility of measuring incentives as the sensitivity of the certainty equivalent/value to changes in current price. Consider our original example but let endogenous variance be given by $h(q) = q^{0.9}$. The optimal piecewise linear contract, denoted s_{pwl} in Figure 4, induces the manager to undertake $(a, q) = (0.318, 1.836)$. The resulting expected outcome is 0.533, the standard deviation is 4.259, and the principal's expected payoff is 0.321.

Assume the principal offers the agent the alternative contract s_{high}, illustrated in Panel (a) of Figure 4, with a slope twice the slope of the optimal contract, s_{pwl}. The new contract is designed such that (i) both contracts become increasing at the same level of the stock price and (ii) *if* the agent undertakes $(a, q) = (0.318, 1.836)$, then the agent's expected utility is the same under the two contracts (and equal to the reservation utility of -1). If the derivative of the certainty equivalent or value wrt. current stock price measures incentives, the agent has stronger incentives under the new contract and is expected to increase effort as well as productive investment. The new contract results in lower effort $a = 0.243$, and significantly lower productive investment $q = 0.344$. Thus, the immediate intuition concerning the effects of slope sometimes fails. The key is that the agent can dramatically affect the variance, which turns out to be desirable for the agent, as can also be seen from the plot of the two probability density functions in Panel (a) of Figure 4. With this parametrization, the benefit the agent derives from reducing standard deviation to 2.005 outweighs the consequences of a reduction in mean outcome to 0.287.[21] Further, measuring incentives for risk taking by sensitivity of value to volatility, one would also conclude that the agent would increase investment as a volatility increase of 10% would lead to a higher increase of the value under s_{high} than under s_{pwl}. If instead one were to measure incentives for risk taking by the change in certainty equivalent, then one would conclude the agent would decrease investment under s_{high} as the certainty equivalent falls when volatility is increased by 10%.

[Figure 4 about here.]

Next, consider the alternative contract s_{out}, illustrated in panel (b) of Figure 4, which has the same slope as s_{pwl} but is closer to being at the money – the strike output κ is -3.666 as opposed to -5.666. As above, s_{out} is adjusted such that *if* the agent undertakes $(a, q) = (0.318, 1.836)$, then the agent's expected utility is the same under the two contracts and equal to the reservation utility ($= -1$). Given effort and investment of $(a, q) = (0.318, 1.836)$, one could compare the change in certainty equivalent or value as a function of a 10% increase in stock price or in volatility under the two contracts. The change in

[21] If endogenous variance is insensitive to changes in investment, then the agent will increase both effort and investment.

certainty equivalent is higher under s_{pwl} than under s_{out} for a 10% increase in stock price, simply because s_{pwl} is issued further in the money. If volatility is increased by 10%, the certainty equivalent increases by 0.07% under s_{pwl} whereas the certainty equivalent increases 0.27% under s_{out}. As 0.07% and 0.27% are both quite close to zero, one might believe risk-taking incentives are similar across the two contracts. If incentives for risk taking are similar across the two contracts while incentives for effort are stronger under s_{pwl}, one would expect the agent to decrease effort and, as effort affects the marginal productivity of investments, to decrease investment. The opposite turns out optimal for the agent. The agent both exerts more effort and invests more under s_{out} than under s_{pwl}. Under s_{out} the agent chooses $(a, q) = (0.417, 5.211)$ and thus almost triples investment – stock price volatility increases by 160%. In our view, the fact that contracts which give seemingly identical incentives for risk taking lead to very different behavior is part of the same phenomenon that linear and piecewise linear contracts give rise to very different outcomes, even if the kink in the piecewise linear contract is located several standard deviations from the mean. Thus, seemingly negligible differences in contracts may lead to very different outcomes.

Assessing risk-taking incentives by comparing how the values or certainty equivalents of CEOs' compensation packages react to a percentage increase in stock price or stock price volatility is problematic. As demonstrated above, these measures neither measures the direction nor the force of incentives well. Hence, comparisons of risk-taking incentives across industries with very different production technologies can be grossly misleading. For example, providing a banker with the 'same' risk-taking incentives as a manufacturer may lead to entirely different results. Similarly, comparing large and small banks could be problematic unless they face the same risk-return tradeoffs.

Although common in the literature, the assumption price equals outcome gross of compensation can be problematic. In practice and in theory compensation plays a large role in the valuation of firms – think for example of the hidden dirty surplus adjustments related to equity-based compensation which are standard in equity valuation. Hence, compensation impacts future and thus current stock price both directly and indirectly. We assume the distribution of the outcome is affected by the agent's actions, which in turn are affected by the compensation scheme. That is, for fixed effort and investment the distribution of outcome x is independent of the compensation scheme, however, the distribution of net payoff to the principal $x - s(x)$ is not.

If the agent's remuneration is based on stock price but settled with cash, matters are fairly simple: let y be the number of shares outstanding, P_0 the stock price at $t = 0$ after the contract is set up, and P_1 the stock price immediately after the outcome is realized and the agent is paid. Given owners are risk neutral, it follows that $P_0 = \mathrm{E}[x - s(x)]/y$ and $P_1 = [x - s(x)]/y$. As long as P_1 is a monotonic function of x, any outcome-based contract $s(x)$ can be replaced by an equivalent stock-based contract $\hat{s}(P_1)$ leading to the same decisions and values to both the agent and the principal. A linear outcome-based contract clearly has an equivalent linear stock-based contract. Similarly for piecewise linear contracts (see Lemma 2 below). When P_1 is not a linear or piecewise linear function of x, the contract $\hat{s}(\cdot)$ may depend on P_1 differently from how the contract $s(\cdot)$ depends on x, but if compensation is small compared to outcome the difference between the *shapes* of the two compensation schemes is limited.

When the agent's remuneration takes place in the form of physical delivery, matters are slightly more complicated. We focus on stock-based contracts that are piecewise linear with a single kink, i.e., the stock-based equivalent to the piecewise linear outcome-based contract studied in Section 3.5. In principle there is no difference whether the agent is issued options or warrants. However, if the agent's share of the outcome $\kappa_1$is larger than 0.5, then option contracting poses some conceptual problems (the agent must be awarded options on more than 100% of the outstanding shares). Assume the agent is issued Δ options or warrants with strike price $\hat{\kappa}$. The following lemma is easily verified.

Lemma 2 *Let $s(x)$ be defined as in* (13) *and let*

$$\hat{s}(P_1) = \kappa_0 + \Delta \max\{P_1 - \hat{\kappa}; 0\}.$$

If $\Delta = \frac{\kappa_1}{1-\kappa_1} y$ and $\hat{\kappa} = [\kappa - \kappa_0]/y$, then $\hat{s}(P_1(x))$ and $s(x)$ are equivalent, i.e., for any outcome x the principal and the agent receive the same under $\hat{s}(\cdot)$ and $s(\cdot)$. Further, the *distribution of* P_1, $\breve{g}(P_1\,|a,q)$, *is given by*

$$\breve{g}(P_1\,|a,q) = \begin{cases} g(\kappa_0 + P_1|a,q) & ; P_1 \le \hat{\kappa}, \\ g(\kappa_0 + (1+\Delta)P_1 - \Delta\hat{\kappa}|a,q) & ; P_1 > \hat{\kappa}. \end{cases}$$

Any piecewise linear outcome-based contract can thus be reinterpreted as a contract with a fixed pay plus Δ options with strike price $\hat{\kappa}$ (we will not explicitly distinguish between options and warrants, one can freely substitute warrant for option).[22] The lemma also has implications for our analysis of incentive measures: in our first example above where we made the assumption stock price equaled outcome gross of compensation we altered incentives by holding fixed the point κ where compensation starts to increase by lowering fixed compensation κ_0 and by increasing variable compensation κ_1. When price equals outcome net of compensation, this corresponds to an increase in the strike price $\hat{\kappa}$, a lowering of fixed compensation κ_0 and an increase in Δ, i.e., more options issued further out of the money.[23] Moreover, the distribution of future stock price P_1 also changes. Hence, if one believes that 1) price equals outcome net of compensation and 2) the basic production technology yielding x is independent of compensation, then the evaluation of alternative incentive schemes formulated in price is considerably more complicated than standard option analysis suggests.

The optimal piecewise linear contract with benchmark parameters corresponds to

$$\hat{s}(P_1) = -1.764 + 1.220y \max\{P_1 - (-0.991/y); 0\}.$$

The moneyness of the option depends on the relation between P_0 and $\hat{\kappa}$ and, as $P_0 = 0.378/y$ and $\hat{\kappa} = -0.991/y$, the option is issued in the money ($P_0 > \hat{\kappa}$). As options are being issued in, at, and out of the money, an interesting question is what characteristics of the problem affect the moneyness of the optimal option contract. Regardless of

[22] We assume $0 < \kappa_1 < 1$, which includes the optimal piecewise linear contract. Then $\Delta > 0$, while $\hat{\kappa}$ can be positive or negative, as can the stock price, due to the assumed normally distributed outcome.

[23] If price equals outcome net of compensation and we change incentives by holding fixed the strike price $\hat{\kappa}$ by lowering fixed compensation κ_0 and by increasing the number of options Δ, then formulated in outcome this manoeuver corresponds to increasing the point κ where compensation starts to increase to lowering fixed compensation κ_0 and to increasing variable compensation κ_1.

the specifics of the problem, the agent must be given incentives to exert effort and to invest. Incentives to provide effort are best served by an in-the-money contract, whereas investment incentives benefit from an out-of-the-money contract – relative to an out-of-the-money contract the incentives to take on risk are reduced under an in-the-money contract. Intuitively, it thus seems that the more risk averse the agent and the more sensitive (less concave) the variance function $h(\cdot)\sigma^2$, the more likely it is that the optimal contract is issued out of the money.

3.7 Comparative Statics

Holding fixed a piecewise linear contract, one would expect an increase in the agent's risk aversion r to have a minor direct effect on the agent's incentives to increase the mean, but a relatively large and negative direct effect on the agent's willingness to increase variance. The expected net effect of an increase in risk aversion is thus a reduction in the agent's willingness to undertake risky investments, a worsening of the induced moral hazard problem. Also, for fixed investment, an increase in risk aversion makes it more costly to induce effort and one would expect a decrease in equilibrium levels of both effort and investment.[24] Our numerical examples confirm this. As illustrated in Table 2, increased risk aversion leads to the issuance of fewer options (a flatter contract) and an increase in the distance between current stock price and the strike price. The distance is increased both when measured in absolute terms and when measured in relative terms (the number of standard deviations of the induced stock-price distribution).[25]

[Table 2 about here.]

If we express the contract in terms of outcome, an increased risk aversion leads to a flatter contract and a decrease in the distance between the mean outcome and the outcome level κ at which compensation starts to increase. The distance is reduced both in

[24]Effort and investment are substitutes in the production function $f(\cdot)$ and as reduced investment lowers marginal productivity of effort, lower investment naturally leads to lower effort.

[25]As stock price is measured net of compensation, a decrease in the incentive rate Δ may lead to increased stock price variance even when the variance of output decreases. This phenomenon tends to occur when the change in incentive rate (and/or strike price) is relatively large relative to the change in variance of outcome – as in our examples. Further, when Δ and changes in Δ are relatively large, then effects may well move in opposite directions when the contract is formulated in outcome relative to when it is formulated in stock price. Therefore, we also report effects for outcome-based contracts.

absolute terms and in terms of the number of standard deviations of the outcome. With unrestricted contracts, compensation is minimal for outcome $\hat{x}$, i.e., when the standardized outcome $z(x) = [x - \mathrm{E}[x|a,q]]/\sigma h(q)^{1/2} \sim N(0,1)$ equals $\hat{z}$, where

$$\hat{z} = -\left[\frac{\mu_a}{\mu_q} f_a(a,q) + f_q(a,q) - 1\right] \frac{h(q)}{h'(q)} \frac{1}{\sigma h(q)^{1/2}}. \tag{15}$$

For fixed effort and investment the only effect of risk aversion is through μ_a/μ_q. In the numerical examples $|\hat{z}|$ is decreasing in r for fixed effort and investment, hence μ_q must be increasing relative to μ_a. As the non-decreasing contracts starts to increase at an outcome level of $x' \approx \hat{x}$, this effect carries over to non-decreasing contracts, and our examples suggest the effect carries over to piecewise linear contracts as well. Thus, our examples indicate the induced moral hazard problem is getting relatively more serious as the agent is getting more risk averse.

We already know that when investments do not affect variance (e.g., when $\zeta = 0$), then the strike price approaches $-\infty$, whereas the strike price is closer to current stock price when investments affect variance (e.g., for $\zeta = 0.25$). Thus, when $h'(q)$ or σ^2 is increased, one would expect the principal to counter the agent's increased aversion against undertaking investments by offering more insurance. That is, an increase in $h'(q)\sigma^2$ is expected to lead to options being issued with an exercise price closer to or even exceeding current stock price – options closer to being out of the money. An interesting question is whether it is a universal phenomenon that an increase in the marginal effect on variance from investment $h'(q)\sigma^2$ leads to options being issued further out of the money. That is, whether the contracts employed in firms where investments significantly impacts volatility tend to be less in the money/more out of the money than options employed in firms where investments have a limited effect on volatility. In our numerical examples, this depends on the level from which ζ or σ^2 is increased, and it makes a difference whether ζ or σ^2 is increased. When we increase the exogenous part of the variance σ^2, the options move further into the money. This holds in absolute but not in relative terms where the contract first moves in the direction of becoming out of the money and subsequently, for σ^2 sufficiently high, moves in the direction of becoming back in the money, see Table 2.[26] In our examples, an increase in ζ leads to options being issued less in the money when ζ is raised from 0.0625 to 0.25, whereas an increase in ζ from 0.25 to 0.9 results in options being issued further in the money. The observations pertaining to changes in ζ hold regardless of whether the contract is formulated in P or in x and regardless of whether the distance is measured absolutely or in terms of standard deviations. Increasing ζ leads to decreasing investment and effort levels, which allows for a lower slope on the increasing part of the compensation scheme. A lower slope in turn lessens the agent's aversion against undertaking investments and works to move the options further into the money. In some of our examples, this effect dominates and the resulting contract is issued further in the money.

Next, we consider the relation between effort and investment. The first-best ratio of investment to squared effort is given by $q/a^2 = 2\frac{1-\theta}{\theta}$. In our numerical examples $\theta = 1/4$ so $q/a^2 = 6$. When subjected to a linear contract with a low slope, the agent naturally tends employ too much capital relative to effort as the agent receives only part of the benefit but carries the whole cost of effort. If the slope approaches unity, the agent receives the whole benefit and carries the whole cost of effort. Hence, in that case one would think the ratio of investment to squared effort would approach $2\frac{1-\theta}{\theta}$, but in fact the agent typically underinvests such that the ratio is below $2\frac{1-\theta}{\theta}$. The reason for this is that investment increases variance, which is a cost to the agent. Now, as r, ζ, or σ^2 are increased, so is the ratio of investment to squared effort q/a^2. This is natural as the contract is becoming less steep, however, the ratio is going up considerably also compared to what the ratio would have been had a linear contract with a slope equal to Δ (or κ_1) been employed (or under the optimal linear contract). Under the optimal piecewise linear contract, the agent invests such that $f_q(a, q) - 1 < 0$, which also tends to increase the risk-premium. The upside of the induced overinvestment is that it facilitates the agent's incentives to undertake effort as investment increases the marginal productivity of effort. The tradeoffs involved in setting the optimal contract are non-trivial, which is why one's immediate intuition doesn't always hold, for example when ζ is increased from 0.25 to

[26] If formulated in outcome, the increase in exogenous variance in our numerical examples leads to a contract more in-the-money when measured in absolute terms. If, on the other hand, we measure in relative terms, then contracts are moving out of the money as σ^2 is increasing.

0.5 or 0.9.

When r, σ^2, or ζ is going up the equilibrium levels of effort and investment are going down. As the level of effort and investment determines the contract (given r, σ^2, and ζ) it follows from the fact that q/a^2 is going up relative to a linear contract with the same slope that more insurance is offered – or that relatively more investment is induced in equilibrium. It is also apparent that the gain of using a piecewise linear contract compared to a linear contract is increasing in r, σ^2, and ζ, that is, as the induced moral hazard problem is getting more serious, the greater the advantage of a piecewise linear contract – we report the percentage loss form going to a linear contract in the last column of Table 2. It does not follow, however, that the distance – measured absolutely or in terms of standard deviations – between the strike price and the current stock price is decreasing as the relative advantage of a piecewise linear contract is increasing. This is counterintuitive and may lead to the erroneous conclusion that the benefit of options relative to restricted stock is slim when options are issued far in the money.

When incentives for risk taking is a concern, optimal contracts resemble a combination of a fixed salary and call options with a relatively high strike price. Hence, even if the comparative statics are not monotone, we still expect research/investment heavy businesses like pharmaceuticals or exploration companies to be more prolific users of options than for example utilities. As it is primarily the endogenous effect of investments on variance and to a lesser extent the resulting (total) variance which influences the optimal contract, one should *not* expect firms with high stock price volatility to make more use of options (which are more out of the money) than firms with low stock price volatility, unless of course high marginal effect and high variance go hand in hand.

Our model incorporates a single period, no systematic risk, and a zero interest rate, so current stock price is simply the expected value of future stock price. As moneyness describes the relation between current stock price and future exercise price, our assumptions introduces a bias against out-of-the-money contracts. Had we allowed for systematic risk, a non-zero interest rate and a substantial vesting period, our analysis would have lead to option contracts which at the grant date are at or out of the money even if they are expected to be in the money upon vesting. Regardless, our results still speak to the effect of various factors on the moneyness of options granted.

4 The trouble with options

Above we demonstrated that the agent could have an incentive to reduce investments as the benefit related to a reduction in future stock price volatility more than outweighed the associated reduction in the current value of the agent's stock options. A potentially more serious problem is that the reverse tradeoff may instead be beneficial. Even when granted in-the-money options, the agent may overinvest as the increase in stock price volatility will outweigh the reduction in expected future stock price and lead to a higher certainty equivalent and, of course, a higher market value of the agent's stock options. Until now the assumption has been that the agent could not access the capital market without the principal observing it, and that the available capital was low enough the overinvestment problem did not occur. We now relax these assumptions and allow for the agent to invest without bounds.

As pointed out in Section 3.3, the optimal unrestricted contract – the "short butterfly" contract – *might* lead to the agent overinvesting with potentially severe consequences for the principal. As also pointed out, the fact that the unrestricted contract rewards extreme outcomes, regardless of whether these are high or low, implies that the agent *will* overinvest if the variance function is unbounded, e.g., $h(q) = q^{\zeta}$, or if mean outcome is unbounded below in investment such that large investments leads to large decreases in mean outcome.

Non-decreasing contracts eliminate the incentive to decrease the mean, but not necessarily the incentive to increase variance. In our base case example, the overinvestment problem does not arise under the piecewise linear contract derived using the first-order approach as the agent's expected utility is globally concave.[27] The overinvestment problem resurfaces, however, under piecewise linear contracts for other parameter values where concavity of the agent's decision problem is only local and the agent's expected utility starts increasing in the investment level above some point, i.e., the agent has an incen-

[27]Regarding our previous examples we have studied a substantial convex and compact set of effort and investment levels and have found the agent's problem concave within the set.

tive to dramatically increase investments to exploit the upside potential of the contract. Under such circumstances the agent will increase investment without bounds unless the principal can somehow control the agent's access to capital markets. This is for example the case under the piecewise linear contract derived using the first-order approach when the endogenous variance function is $h(q) = q^{5/4}$ instead of $q^{1/4}$ as in the base case.[28] It is somewhat surprising the agent has such incentives as the options awarded are deep in the money and the normal distribution are characterized by thin tails (with other distributions having thicker tails the overinvestment problem could potentially be much more severe).

The overinvestment problem occurs when the endogenous variance function is sufficiently convex so that variance increases rapidly relative to the fall in the mean. One cure against this malady is to employ linear contracts, which eliminates overinvestment as the agent shares equally in gains and losses. As previously noted, this solution is very costly, and a cheaper alternative is to cap compensation. A capped contract can be created by writing a contract in which the manager is awarded a number of options with a low exercise price and in which the agent issues an equal number of call options with a high exercise price to the firm; the contract should also specify that the firm must exercise concurrently with the agent if at all. Such a contract is similar to the optimal non-decreasing contract which is concave in its upper tail. Capping any contract derived using the first-order approach will eliminate the overinvestment problem provided the cap is sufficiently low. Of course, capping a contract leads to different effort and investment decisions, but for caps sufficient to curb overinvestment incentives the effect of this is often limited.[29]

As an example of the consequences of capping contracts, consider our base case, but let the endogenous variance function be $h(q) = q^{5/4}$. The 'optimal' piecewise linear contract seems to induce $a = 0.283$ and $q = 1.514$ yielding an expected payoff of 0.302 to

[28] For the production and variance functions we have implemented numerically, the first-order approach is valid for the non-decreasing and the piecewise linear compensation schemes as long as the endogenous variance function is sufficiently concave, that is, when ζ is sufficiently small.

[29] Note that (1) if the piecewise linear contract derived using the first-order approach leaves the agent with incentives to overinvest, the contract is suboptimal in the first place, and (2) the effect of curbing the agent's incentives to overinvest through a linear contract is potentially substantial.

the principal. This contract is infeasible as the endogenous variance function is convex enough that the agent will overinvest if he can gain access to capital in excess of 2.0. If we cap this contract such that the contract converts from increasing to being flat at the mean plus 2.5 standard deviations and adjusts the slope of the increasing part just slightly, then the agent's problem is concave and the agent no longer overinvests. This contract induces $a = 0.296$ and $q = 1.454$ and yields an expected payoff of 0.299 to the principal, that is a loss of roughly 1.3% relative to the feasible non-decreasing contract. The alternative of offering a linear contract results in an expected payoff of 0.207 to the principal, a loss of 30.8% relative to the capped contract.

[Table 3 about here.]

Further, if there is some uncertainty regarding the exact specification of the production and variance functions, a capped contract can be an effective means against unexpected overinvestment. If, for example, the agent is presented with the piecewise linear contract with one kink which is optimal for the variance function $h(q)\sigma^2 = 10.5\,q^{0.9}$, while the true variance function (known by the agent) is $10.5\,q^{1.1}$, then the agent has an incentive to overinvest. As can be seen from Table 3, the agent will triple investments relative to what was expected (5.377 compared to 1.836). If instead the agent is presented with the same contract capped for outcomes more than 2.5 standard deviations above the mean, the principal's expected payoff is close to the principal's conjectured payoff regardless of the variance function. With a capped contract, the agent will still invest more when the variance function is $10.5\,q^{1.1}$ than when the variance function is $10.5\,q^{0.9}$, however as indicated in Table 3 the differences are much smaller. In Table 3 we list the agent's effort and investment decisions and the resulting price given ζ is either 0.9 or 1.1 for three different caps (in these examples we only cap the contract we do not optimize the contract given it is capped). In general, the agent will choose actions that are both unforeseen and inoptimal from the principal's point of view when the principal errs in specifying the parameters of the production function and this, of course, implies an associated loss. However, if the misspecification leads to an overinvestment problem, we believe the associated loss is considerably smaller when a capped contract is employed.

As explained in the introduction, regulation of executive compensation has recently been discussed and to some extent implemented both in the US and Europe. The results presented in the preceding section show that stock options may very well be part of a (near-)optimal executive compensation plan, which suggests that a general ban on option compensation should be avoided. The observations made in the present section indicate that a cap on compensation is useful in avoiding excessive risk taking that may emerge under some circumstances. Of course, the owners of a company will generally have an incentive to curb excessive risk taking, and thus introduce a cap on the manager's compensation by themselves if they find it relevant. The overinvestment problem tend to occur when the manager can increase risk substantially without hurting the expected future outcome and thus stock price. If raising additional capital is necessary to increase risk, creditors will under normal circumstances increase credit costs due to the increased default risk. The increased credit cost will reduce the expected future stock price and perhaps in turn eliminate the problem of excessive risk taking. Thus, our analysis suggests that compensation regulation should encompass businesses that can (1) increase leverage without increasing credit costs, because lenders know the business will not be allowed to default on its debt, and (2) covertly increase risk without risking being caught by regulators – perhaps by creating new types of investments/securities. That is, large financial institutions.[30]

5 Conclusion

Stock options are ubiquitous in executive compensation but are claimed to provide perverse incentives for risk taking eventually adding, if not directly leading, to the recent crisis. This paper demonstrates that if the investment decisions of the manager are unobservable and subject to induced moral hazard, it may be optimal to offer an incentive contract rewarding both low and high outcomes. But such short butterfly contracts are infeasible if the agent can increase variance without bounds or destroy outcome. We also show that if the agent has such opportunities, an option contract might emulate the

[30] Of course, the optimal regulation from a society's point of view must be studied by formalizing some social welfare function depending on the compensation of managers.

optimal non-decreasing contract quite closely but – at least under some circumstances – capping the option payoff can be necessary to avoid excessive risk taking.

We also demonstrate that when the manager controls investments – or more generally variance – the standard measure of the incentives provided by a compensation scheme can be grossly misleading. The optimal design of the compensation contract and its inherent incentives are highly depending on the extent to which the manager can affect the variance of the earnings or stock price. In particular, the options included in the near-optimal compensation package can be issued either in-, at-, or out-of-the-money and should potentially be capped.

Our analysis has implications for the current initiatives regarding the regulation of compensation. We demonstrate that options can be optimal and that confining variable compensation to restricted stock can be detrimental to shareholders (and to society) as the manager's induced preferences over risky projects will lead to too little risk taking. On the other hand, we show that options may incentivize the agent to undertake excessively risky investments, but that this behavior can be curbed by appropriately capping compensation while maintaining incentives for appropriate risk taking. Financial institutions that are considered too-big-to-fail do not bear all the costs associated with excessive risk taking and may therefore not impose such a cap on manager compensation, unless they are forced to do so by regulators.

A Proofs

Proof of Proposition 2. From (11) it follows that

$$s'(x) = \frac{U_s(\cdot)^2}{-U_{ss}(\cdot)} \left(\mu_a \frac{f_a(a,q)}{\sigma^2 h(q)} + \mu_q \left\{ \frac{f_q(a,q)-1}{\sigma^2 h(q)} + \frac{x - \mathrm{E}[x|a,q]}{\sigma^2 h(q)} \frac{h'(q)}{h(q)} \right\} \right).$$

This implies that the slope of the compensation scheme, when evaluated at the mean, $\mathrm{E}[x|a,q] = f(a,q) + (\bar{q} - q)$, is

$$s'(\mathrm{E}[x|a,q]) = \frac{U_s(\cdot)^2}{-U_{ss}(\cdot)} \left(\mu_a \frac{f_a(a,q)}{\sigma^2 h(q)} + \mu_q \frac{f_q(a,q)-1}{\sigma^2 h(q)} \right).$$

Assume the compensation scheme is decreasing at the mean or alternatively that the mean outcome is the stationary point $\hat{x}$. Assuming θ belongs to the range of $s(x)$, then $\Pr(s(x) \leq \theta) = \Pr(x \in [\hat{x} - \eta, \hat{x} + \eta])$ some $\eta > 0$.

$$\Pr(x \in [\hat{x} - \eta, \hat{x} + \eta]|a,q) = \int_{\hat{x}-\eta}^{\hat{x}+\eta} g(x|a,q)\, dx.$$

Now

$$\begin{aligned} \frac{\partial}{\partial a} \int_{\hat{x}-\eta}^{\hat{x}+\eta} g(x|a,q)\, dx &= \int_{\hat{x}-\eta}^{\hat{x}+\eta} \frac{x - f(a,q) - (\bar{q} - q)}{\sigma^2 h(q)} f_a(a,q) g(x|a,q)\, dx \\ &= \frac{f_a(a,q)}{\sigma^2 h(q)} \{ \mathrm{E}[x|a,q, x \in [\hat{x} - \eta, \hat{x} + \eta]] - \mathrm{E}[x|a,q] \} \\ &\quad \times \Pr(x \in [\hat{x} - \eta, \hat{x} + \eta]|a,q) \\ &\geq 0, \end{aligned}$$

with the inequality being strict if $\mathrm{E}[x|a,q] < \hat{x}$. Thus, by first-order stochastic dominance, the agent weakly prefers to reduce effort. As there is a strict utility gain from doing so, the agent would reduce effort to zero under a scheme, which is decreasing at the mean. (The marginal monetary cost of effort $C'(a)$ is positive.) □

Proof of Proposition 3. Solving the agency problem ignoring investment incentives yields a compensation scheme, which from (9) is (weakly) increasing in x. As investment does not affect variance, the agent will choose investment to maximize expected outcome (first-order stochastic dominance). As investment only affects the first moment of the distribution, this is preferred by the principal as well. □

Proof of Lemma 1. Let $s(x)$ be an optimal solution for a reservation utility level of $R = -e^{-rR^{CE}}$. Assume the agent's reservation certainty equivalent is instead $R^{CE} + \Delta$. The agent's certainty equivalent of the compensation scheme $s(x) + \Delta$ is $R^{CE} + \Delta$ and thus the individual rationality constraint is satisfied. As the agent's preferences display no wealth effects, the incentive compatibility constraints are both satisfied. Hence, $s(x) + \Delta$ is feasible. Assume the compensation scheme $\hat{s}(x)$ is feasible and entails a higher expected net payment to the principal relative to $s(x) + \Delta$. As $\hat{s}(x) - \Delta$ is feasible for reservation certainty equivalent R^{CE}, $s(x)$ cannot be optimal, a contradiction. □

B A CRRA-lognormal version of the model

In the literature on executive compensation an alternative version of the principal-agent model combines preferences exhibiting constant relative risk aversion with a lognormally distributed outcome or stock price, see, e.g., Dittmann and Maug (2007) and Armstrong, Larcker, and Su (2008). On the surface the normal and the lognormal versions appear quite different but in reality, the models are close. As in the before-mentioned papers we now assume the stock price P_1 is lognormal and, to carry over the ideas of our original model, we assume more precisely that

$$\ln P_1 = \bar{q} - q + f(a, q) + h(q)^{1/2}\epsilon, \qquad \epsilon \sim N(0, \sigma^2),$$

as assumed for outcome x in the original model, cf. (2). Recall, however, that P_1 reflects the end-of-period outcome *after* compensating the agent, so modeling the stock price exogenously blurs the impact of compensation on the stock price and, consequently, the value left for the principal.[31] The distributional assumption implies a mean stock price of $\mathrm{E}[P_1] = \exp\{\bar{q} - q + f(a, q) + \frac{1}{2}h(q)\sigma^2\}$ and a variance of $\mathrm{Var}[P_1] =$

[31] Outcome and stock price cannot both be lognormal. Modeling outcome as lognormal would clarify the impact of compensation on the stock price, but then the stock price would not be lognormal and, in particular, the Black-Scholes option pricing formula would not apply.

$\mathrm{E}[P_1]^2\,(\exp\{h(q)\sigma^2\} - 1)$. Further assume that the agent's preferences can be represented by a power utility function:

$$\hat{U}(\hat{s}(P_1), a) = V(\hat{s}(P_1)) - \hat{C}(a) = \frac{(\hat{w}_0 + \hat{s}(P_1))^{1-\gamma}}{1-\gamma} - \hat{C}(a),$$

where $\hat{w}_0$ is initial wealth and γ is the relative risk aversion. Provided the first-order approach is valid, an optimal compensation scheme is – in its interior – characterized by[32]

$$\begin{aligned}
[\hat{w}_0 + \hat{s}(P_1)]^{\gamma} = \hat{\lambda} &+ \hat{\mu}_a \frac{\ln(P_1/\,\mathrm{E}[P_1]) + \frac{1}{2}h(q)\sigma^2}{\sigma h(q)^{1/2}} \frac{f_a(a,q)}{\sigma h(q)^{1/2}} \\
&+ \hat{\mu}_q \frac{\ln(P_1/\,\mathrm{E}[P_1]) + \frac{1}{2}h(q)\sigma^2}{\sigma h(q)^{1/2}} \frac{f_q(a,q) - 1}{\sigma h(q)^{1/2}} \\
&+ \hat{\mu}_q \left\{ \left[\frac{\ln(P_1/\,\mathrm{E}[P_1]) + \frac{1}{2}h(q)\sigma^2}{\sigma h(q)^{1/2}} \right]^2 - 1 \right\} \frac{1}{2} \frac{h'(q)}{h(q)}.
\end{aligned} \tag{16}$$

In the benchmark case ignoring the investment decision, $\mu_q = 0$ so that the right-hand side is affine in $\ln(P_1)$ and thus – in its interior – optimal compensation is increasing and, when $\gamma > 1$, concave in P_1. If the agent controls investments and $\mu_q > 0$, the right-hand side is a second-order polynomial in $\ln(P_1)$, hence, for positive μ_q compensation is decreasing in P_1 for P_1 close to 0 and increasing in P_1 for P_1 sufficiently large.[33] Thus, even though the compensation scheme is no longer symmetric around a minimum, optimal contracts in the lognormal formulation possesses the same basic properties as optimal contracts in the normal formulation. In particular, it still holds that when $\mu_q > 0$ low as well as high outcomes are rewarded. Furthermore, when the agent has unrestricted access to the capital market or the agent can destroy value (reduce price) it is still problematic to handsomely reward low outcomes, and as in the normally distributed version options are more likely to be optimal, when investments are subject to induced moral hazard.

[32]Depending on the parameters, the Mirrlees problem may well apply to this version of the benchmark case.

[33]Irrespective of the value of γ.

The lognormal benchmark case, where the agent cannot control investments, is the model of Dittmann and Maug (2007) and Armstrong, Larcker, and Su (2008).[34] Feltham and Wu (2001) and Hemmer, Kim, and Verrechia (1999) demonstrate that optimal contracts are more likely to display convexity, when more than one moment of the outcome distribution is affected by effort. However, even though the agent's effort affects both the mean and the variance of the price distribution in the lognormal benchmark case, the optimal contract is concave in its interior. Hence, a search for the best contract within the class of contracts consisting of fixed salary and nonnegative holdings of restricted stock and call options, would not be expected to result in a contract including both stocks and options (much less options with more than one exercise price). In fact, Dittmann and Maug (2007) find that the optimal contract includes positive holdings of stock combined with negative holdings of stock options, i.e., a concave contract.

[34]Though Lambert, Larcker, and Verrecchia (1991) and Hall and Murphy (2000, 2002) do not explicitly model the agency problem, the model employed in these papers is consistent with the lognormal formulation.

References

Admati, A. and P. Pfleiderer (1997). Does It All Add Up? Benchmarks and the Compensation of Active Portfolio Managers. *Journal of Business* 70(3), 323–350.

Armstrong, C., D. F. Larcker, and C.-L. Su (2008). Stock Options and Chief Executive Officer Compensation. Available at SSRN: `http://ssrn.com/abstract=987693`.

Bebchuck, L. A. and H. Spamann (2009). Regulating Bankers' Pay. *The Georgetown Law Journal* 98, 247–287.

Berk, J. and R. Green (2004). Mutual Fund Flows and Performance in Rational Markets. *Journal of Political Economy* 12(6), 1269–1295.

Blinder, A. S. (2009). Crazy Compensation and the Crisis. Newspaper article available at: `http://online.wsj.com/article/SB124346974150760597.html`.

Carpenter, J. (2000). Does Option Compensation Increase Managerial Risk Appetite? *Journal of Finance* 55(5), 2311–2331.

Cassidy, J. (2002). The Greed Cycle: How the Financial System Encouraged Corporations to go Crazy. Magazine article available at: `http://www.csun.edu/~dgw61315/Greedcycle.doc`.

Coles, J. L., D. D. Naveen, and L. Naveen (2006). Managerial Incentives and Risk Taking. *Journal of Financial Economics* 79, 431–468.

Core, J. and W. Guay (2010). Is There a Case for Regulating Executive Pay in the Financial Services Industry? In *After the Crash: The Future of Finance*, pp. 115–140. Brookings Institution Press.

Cuoco, D. and R. Kaniel (2011). Equilibrium Prices in the Presence of Delegated Portfolio Management. *Journal of Financial Economics* 101(2), 264–296.

Demski, J. and D. Sappington (1987). Delegated Expertise. *Journal of Accounting Research* 25, 68–89.

DeYoung, R., E. Peng, and M. Yan (2010). Executive Compensation and the Business Policy Choices at U.S. Commercial Banks. Working paper RWP 10-02, Federal

Reserve Bank of Kansas City.

Dittmann, I. and E. Maug (2007). Lower Salaries and No Options? On the Optimal Structure of Executive Pay. *Journal of Finance* 62(1), 303–343.

Dittmann, I., E. Maug, and O. Spalt (2010). Sticks or Carrots? Optimal CEO Compensation when Managers are Loss-Averse. *Journal of Finance* 65(6), 2015–2050.

Dye, R. (1988). Earnings Management in an Overlapping Generations Model. *Journal of Accounting Research* 26, 193–235.

Fahlenbrach, R. and R. M. Stulz (2011). Bank CEO Incentives and the Credit Crisis. *Journal of Financial Economics* 99(1), 11–26.

Feltham, G. and M. Wu (2001). Incentive Efficiency of Stock versus Options. *Review of Accounting Studies* 6(1), 7–28.

Feltham, G. and J. Xie (1994). Performance Measure Congruity and Diversity in Multi-Task Principal/Agent Relations. *Accounting Review* 69(3), 429–453.

Grossman, S. J. and O. D. Hart (1983). An Analysis of the Principal-Agent Problem. *Econometrica* 51, 7–45.

Guay, W. R. (1999). The Sensitivity of CEO Wealth to Equity Risk: An Analysis of the Magnitude and Determinants. *Journal of Financial Economics* 53, 43–71.

Hall, B. J. and K. J. Murphy (2000). Optimal Exercise Prices for Executive Stock Options. *American Economic Review* 90(2), 209–214.

Hall, B. J. and K. J. Murphy (2002). Stock Options for Undiversified Executives. *Journal of Accounting and Economics* 33(1), 3–42.

Hemmer, T., O. Kim, and R. Verrechia (1999). Introducing Convexity into Optimal Compensation Contracts. *Journal of Accounting and Economics* 30(3), 307–327.

Holmström, B. (1979). Moral Hazard and Observability. *Bell Journal of Economics* 10, 74–91.

Lambert, R. (1986). Executive Effort and the Selection of Risky Projects. *RAND Journal of Economics* 17, 77–88.

Lambert, R. A. and D. F. Larcker (2004). Stock Options, Restricted Stock, and Incentives. Available at SSRN: `http://ssrn.com/abstract=527822`.

Lambert, R. A., D. F. Larcker, and R. E. Verrecchia (1991). Portfolio Considerations in Valuing Executive Compensation. *Journal of Accounting Research* 29(1), 129–149.

Madrick, J. (2003). A Theory on Corporate Greed. Newspaper article available at: `http://latrobefinancialmanagement.com/Research/Governance/ATheoryonCorporateGreed.pdf`.

Mehran, H. and J. V. Rosenberg (2007, October). The Effect of Employee Stock Options on Bank Investment Choice, Borrowing, and Capital. FRB of New York Staff Report no. 305. Available at SSRN: `http://ssrn.com/abstract=1022972`.

Mirrlees, J. A. (1999). The Theory of Moral Hazard and Unobservable Behaviour: Part I. *Review of Economic Studies* 66(1), 3–21.

Murphy, K. J. (2010). Executive Pay Restrictions for TARP Recipients: An Assessment. Working paper FBE 23-10, USC Marshall School of Business. Available at SSRN: `http://ssrn.com/abstract=1698973`.

Murphy, K. J. (2011). The Politics of Pay: A Legislative History of Executive Compensation. Available at SSRN: `http://ssrn.com/abstract=1916358`. Forthcoming in *The Research Handbook on Executive Pay* (edited by J. Hill and R. Thomas).

Ou-Yang, H. (2003). Optimal Contracts in a Continuous-Time Delegated Portfolio Management Problem. *Review of Financial Studies* 16(1), 173–208.

Rajgopal, S. and T. Shevlin (2002). Empirical Evidence on the Relation Between Stock Option Compensation and Risk Taking. *Journal of Accounting and Economics* 33, 145–171.

Ross, S. A. (2004). Compensation, Incentives, and the Duality of Risk Aversion and Riskiness. *Journal of Finance* 59(1), 207–225.

Sanders, W. and D. Hambrick (2007). Swinging for the Fences: The Effects of CEO Stock Options on Company Risk Taking and Performance. *Academy of Management Journal* 50(5), 1055–1078.

Smith, G. and P. L. Swan (2008). The Incentive to "Bet the Farm": CEO Compensation and Major Investments. Available at SSRN: `http://ssrn.com/abstract=1009323`.

Stoughton, N. M. (1993). Moral Hazard and the Portfolio Management Problem. *Journal of Finance* 48(5), 2009–2028.

Stracca, L. (2006). Delegated Portfolio Management: A Survey of the Theoretical Literature. *Journal of Economic Surveys* 20(5), 823–848.

van Binsbergen, J. H., M. W. Brandt, and R. S. Koijen (2008). Optimal Decentralized Investment Management. *Journal of Finance* 63(4), 1849–1895.

Description	symbol	value
productivity parameter	θ	0.25
scale of production	k	2.0
agent's absolute risk aversion	r	0.1
variance of random shock	σ^2	10.5
power of production-induced variance	ζ	0.25

Table 1: **Base case parameters for the numerical examples.**

r	a	q_p	κ_0	Δ	$\hat{\kappa}$	P_0	σ_{P_0}	$\frac{P_0\text{–strike}}{\sigma_{P_0}}$	% loss
0.3	0.225	1.436	-0.819	0.440	-1.612	0.217	2.579	0.709	68.9
0.2	0.287	1.770	-1.105	0.641	-1.391	0.272	2.408	0.690	36.2
0.1	**0.413**	**2.417**	**-1.773**	**1.246**	**-0.963**	**0.378**	**2.003**	**0.669**	**12.3**
0.05	0.547	3.076	-2.649	2.487	-0.494	0.485	1.526	0.641	4.6
σ^2	a	q_p	κ_0	Δ	$\hat{\kappa}$	P_0	σ_{P_0}	$\frac{P_0\text{–strike}}{\sigma_{P_0}}$	% loss
13.5	0.372	2.254	-1.773	0.953	-1.430	0.342	2.465	0.719	18.9
12.0	0.390	2.325	-1.769	1.072	-1.205	0.359	2.244	0.697	15.6
10.5	**0.413**	**2.417**	**-1.773**	**1.246**	**-0.963**	**0.378**	**2.003**	**0.669**	**12.3**
9.0	0.438	2.511	-1.773	1.461	-0.742	0.400	1.754	0.651	9.3
7.5	0.469	2.626	-1.778	1.773	-0.523	0.427	1.486	0.639	6.7
5.0	0.541	2.894	-1.800	2.745	-0.160	0.484	0.991	0.651	2.7
2.5	0.655	3.344	-1.800	5.518	0.206	0.571	0.447	0.817	0.4
$h(q_p)$	a	q_p	κ_0	Δ	$\hat{\kappa}$ (strike)	P_0	σ_{P_0}	$\frac{P_0\text{–strike}}{\sigma_{P_0}}$	% loss
$q_p^{9/10}$	0.318	1.836	-2.170	0.602	-3.497	0.321	2.784	1.371	36.0
$q_p^{1/2}$	0.377	2.292	-1.993	0.836	-2.125	0.355	2.406	1.031	29.4
$\boldsymbol{q_p^{1/4}}$	**0.413**	**2.417**	**-1.773**	**1.246**	**-0.963**	**0.378**	**2.003**	**0.669**	**12.3**
$q_p^{1/8}$	0.430	2.287	-1.999	1.465	-0.970	0.385	1.714	0.791	3.3
$q_p^{1/16}$	0.458	2.277	-3.318	1.335	-2.344	0.392	1.482	1.847	0.1

Table 2: **Comparative statics:** The bold faced rows correspond to the base case. The column"% loss" contains the percentage loss from using a linear contract instead of the optimal piecewise linear contract. P_0 is the stock price net of compensation.

actual ζ	a	q_p	ψCap	E$[U]$	P_0
9/10	0.318	1.836	∞	-1	0.321
11/10	0.437	5.377	∞	-0.997	-0.144
9/10	0.295	1.494	2.5	-1	0.316
11/10	0.330	2.156	2.5	-1.0	0.313
9/10	0.289	1.423	2.4	-1	0.313
11/10	0.317	1.917	2.4	-1.0	0.317

Table 3: **Capped option contract.** For two cases of endogenous variance function (ζ), the "piecewise linear contract" from first-order conditions is capped ψCap standard deviations above the mean.

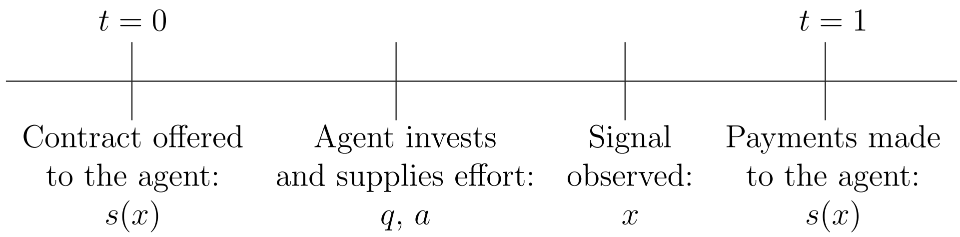

Figure 1: **Time line of the model.**

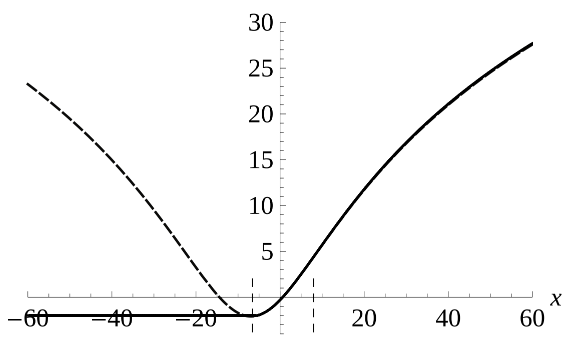

(a) Large output range

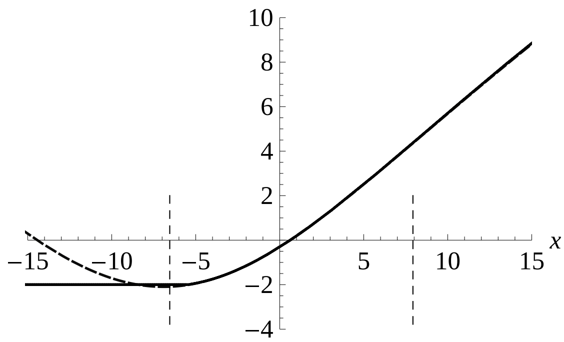

(b) Output range close to the mean

Figure 2: **Compensation contract in the second-best (dashed) and the third-best case (solid).** The vertical "|" lines mark two units of standard deviation.

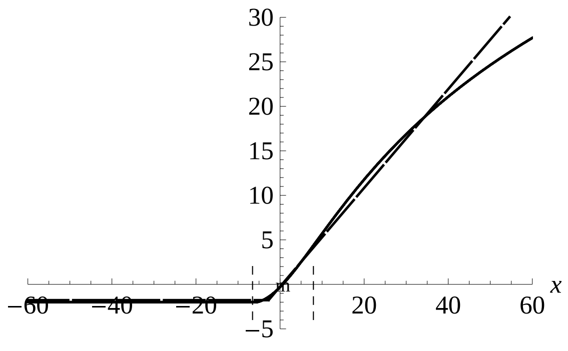

(a) Large output range

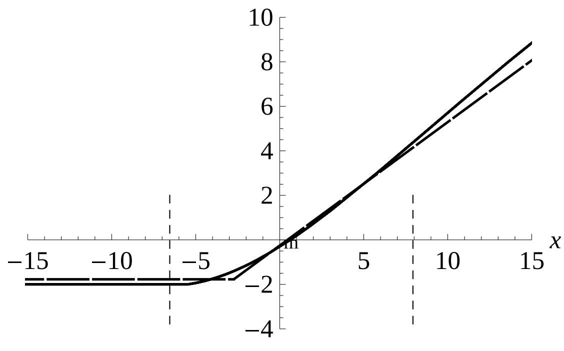

(b) Output close to the mean (m)

Figure 3: **Compensation contract in the second-best case (solid) and the best piecewise linear contract (long dashed).** The vertical "|" lines mark two units of standard deviation.

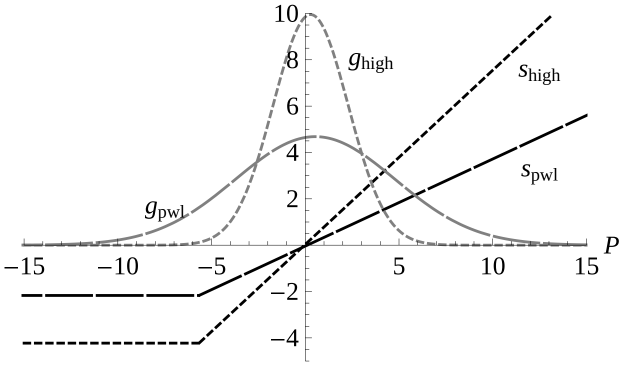

(a) The contract s_{high} has a twice as high slope and provides the agent with an expected utility equal to his reservation utility if he undertakes effort and investments as if s_{pwl} is offered.

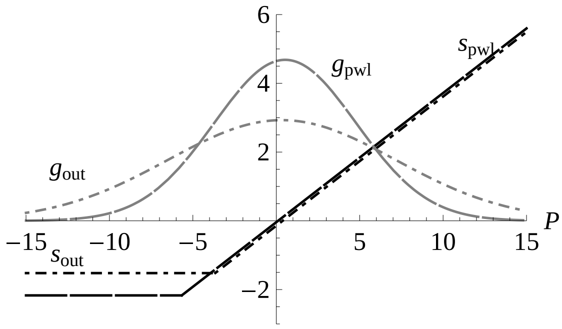

(b) The contract s_{out} has a strike price and provides the agent with an expected utility equal to his reservation utility if he undertakes effort and investments as if s_{pwl} is offered.

Figure 4: **Incentives and investment choice.** s_{pwl} is the optimal piecewise linear contract. The gray curves show the probability density function (scaled up by a factor of 50) in the respective cases