Federation of Business Disciplines Journal
Volume 4, 2016 1 – 21

SUBJECTIVE BELIEFS AND MANAGEMENT CONTROL: A USER'S GUIDE AND A MODEL EXTENSION

Anthony Meder
Binghamton University
Binghamton, NY 13902
USA

Steven Schwartz
Binghamton University
Binghamton, NY 13902
USA

Eric Spires
The Ohio State University
Columbus, OH 43210
USA

Richard Young
The Ohio State University
Columbus, OH 43210
USA

ABSTRACT

Recently published research studies contracting between a superior and subordinate who have different beliefs regarding the subordinate's ability. The typical assumption is that the subordinate is overconfident, in that his belief regarding his productivity is higher than warranted. Maintaining this assumption, this paper employs easy-to-follow numerical illustrations that demonstrate overconfidence can lead to large increases in the riskiness of subordinate compensation and lower expected compensation costs for the superior. Also included are a brief analysis of the effects of underconfidence and a discussion of the plausible real world effects of manager overconfidence. The paper concludes with an extension of previous modeling to explore conditions under which the effects of overconfidence are predicted to be most pronounced.

Key Words: overconfidence, contracting, agency theory

INTRODUCTION

This paper uses numerical and graphical examples to illustrate and extend recent research modeling the effect of subjective beliefs in management control. The goal is to make important findings in management control accessible to a wider audience. In the setting, which is based on recent models by Gervais et al. (2011) and de la Rosa (2011), a superior and subordinate hold different beliefs over the subordinate's productivity. Most results of the modeling hold regardless of whether superior, subordinate or neither holds the correct beliefs; however, for expositional ease, we assume the superior holds the correct beliefs. Accordingly, a subordinate who holds identical beliefs to the superior is referred to as "realistic," while a subordinate who believes he is *more* productive than does the superior is referred to as "overconfident." We also take a brief look at underconfidence, where a subordinate believes he is *less* productive than does the realistic superior.

Modeling subjective beliefs in employment settings is especially important to the study of corporate control. Gervais et al. (2011) cite evidence that more confident individuals are more likely to reach a position of importance within the firm. Further, there is considerable research finding that high levels of confidence are associated with better performance. For example, Hales et al. (2015) find that optimism is associated with increased productivity— it seems reasonable to expect the same from overconfidence.[1] Hirshliefer et al. (2012) link managerial overconfidence with innovation. Mertins and Hoffeld (2015) find that more confident individuals are more cooperative.

While overconfidence about one's ability may be associated with increased productivity, there are other implications. This paper identifies and illustrates five notable observations that emerge from the modeling. First, overconfidence generally reduces the cost of compensation for the organization. Second, the actual riskiness of compensation to the subordinate is first decreasing then increasing in the subordinate's confidence. Third, a subordinate who is *underconfident* receives both greater expected compensation and riskier compensation from the organization. Fourth, some projects that would not be profitable to an organization if the subordinate had realistic expectations might be profitable if its subordinate is overconfident. Fifth, we demonstrate in our extension of these models that the effects of overconfidence are heightened in situations where success is actually either very likely or very unlikely. A discussion of the implications of the models, as well as of issues not captured by the models, is included in the paper.

Drilling deeper into the second observation, which is at the heart of the analysis by de la Rosa (2011), there are two ways the organization might best exploit an overconfident subordinate. First, if the subordinate is *slightly* overconfident the superior can *reduce* the power, and hence riskiness, of incentives. As an analogy, assume a realistic, but risk averse, baseball pitcher in an equilibrium contract is paid a $10,000,000 base salary, plus a $1,000,000 bonus if he wins 20 games. Now assume the pitcher is slightly overconfident, believing more than the team does that his hard work increases the likelihood of winning 20 games. The team can reduce the bonus to, say, $850,000. The pitcher's subjective expected value of the bonus is the product of the bonus amount and his subjective likelihood of receiving the bonus. A higher estimate of the likelihood

[1] Overconfidence is distinguished in the literature from "optimism" where the latter is an over-estimation of the likelihood of a favorable outcome, irrespective of the subordinate's choice, or even of whether the subordinate makes choices at all. Although the results on overconfidence and optimism are similar, this paper focuses on overconfidence because the intersection of subjective beliefs and decision making offers interesting implications.

of receiving the bonus (relative to a realistic pitcher) implies the team can lower the bonus amount and still keep the expected value of the bonus the same, and hence still have sufficiently powerful incentives to induce hard work. This is what we mean by reducing the power of incentives.

Now suppose the pitcher is *significantly* overconfident. It is optimal to *increase* the power of incentives and the riskiness of the contract. The optimal contract from the team's perspective pays the pitcher the league minimum base salary of, say, \$1,000,000 plus a \$20,000,000 bonus for winning 20 games. The pitcher is so confident in the value of his hard work he is willing to "wager" more on himself than he otherwise would and as we shall so later the highly risky contract would be optimal even if it weren't necessary to motivate the pitcher's high effort. The team benefits despite placing more risk on the risk averse pitcher because the bonus compensation is "cheap" in their eyes. A \$20,000,000 bonus is a lot of money, but only if it is paid, which the team believes is much less likely than the pitcher does.

The rest of this paper is follows. Section 2 presents the model found in previous literature. Section 3 offers an extension to the model. Section 4 concludes the paper.

MODEL

General Employment Model

In this simple model of a firm, a superior hires a subordinate to make decisions.[2] The superior has the following attributes: (1) she maximizes her expected utility, (2) she is risk neutral, (3) she wishes the subordinate to choose an act, which we refer to as H, regardless of the additional compensation cost, and (4) she can commit to do things in the future that she may at a later time wish to avoid, such as paying the subordinate a bonus for good performance. The subordinate is employed by the superior and has the following attributes: (1) he maximizes his expected utility, (2) he is risk averse, (3) all else being equal he prefers an action different than that desired by the superior, which we refer to as L, and (4) he has a next-best offer of $\overline{U}$ in expected utility terms. [3] Together these assumptions ensure an interesting setting where the superior's and subordinate's interests are not naturally aligned and where, due to the subordinate's risk aversion, it is inefficient for the subordinate to simply buy out the superior and work for himself.

As mentioned, the subordinate is faced with the choice of whether to work for the superior, and if so, whether to choose H or L. The subordinate's utility is increasing in the compensation received from the superior but reflects an aversion to risk and effort. These attributes are modeled by assuming his utility function is $\sqrt{w}$ – c(a), where w is his compensation and c(a) is his disutility for the action. Let c(H) = c > 0 and let c(L) be normalized to 0.[4]

A jointly observable performance measure is available whose realization is affected by both the subordinate's input H or L and a random component. For simplicity, the performance measure can take on one of two values, interpreted as either good or bad, denoted g and b, respectively. Given this simplification, consistent with the interpretation of H being higher effort

[2] For expositional clarity we refer to the superior as "she" and the subordinate as "he".

[3] H could be interpreted as "working hard" but that is only one of many interpretations. The characteristic that matters here is the superior and subordinate disagree about which action is best. We are also implicitly assuming that by choosing to motivate H the firm has the resources to provide the necessary compensation.

[4] Any concave utility function for money will capture risk aversion and risk aversion is all that is necessary to produce the qualitative model results. The square root function is chosen because it is simple to work with. Also, it is implicitly assumed in the examples that the only wealth the subordinate has is his earnings from the superior. Including wealth effects would be a distraction and would not affect the main results illustrated.

than L, the conditional probability of good performance is greater under H than under L. The superior can commit to condition pay on the performance measure. Let w_g and w_b be the superior's pay to the subordinate for good and bad outcomes, respectively.

The true probability of a good outcome given H, is p_H, where $0 < p_H < 1$, and the true probability of a bad outcome given H is $1-p_H$. The probabilities of the good and bad outcomes given L are p_L and $1-p_L$, respectively. For simplicity, in this and all subsequent examples, p_L is set to zero and is common knowledge. These assumptions are sufficient to prevent the superior from inferring the subordinate's action after observing a bad outcome, as the bad outcome could be the result of the subordinate choosing L or it could be the result of the subordinate choosing H but being unlucky. If the observed outcome is good, the superior can infer that the subordinate chose H.

The superior's assessment of the probability of good outcome is always accurate while the subordinate's assessment of the probability of a good outcome given H is not necessarily accurate. To capture this difference in beliefs about the probability of a good outcome given H, we assume the subordinate believes the probability of g given H, is $p_H + p_O$. If the subordinate is realistic then $p_O = 0$ whereas if the subordinate is overconfident then $p_O > 0$. Finally it is assumed that the subordinate is rational in that the probability of bad and good performance adds to one, which requires that the subordinate's subjective assessment of the probability of b given H is $1 - (p_H + p_O)$. Finally, we assume that the superior is aware of the subordinate's beliefs; i.e., she knows p_O.

The optimal contract from the superior's perspective must satisfy three conditions. First, the contract must ensure that the subordinate's expected net utility from employment is at least as great as $\overline{U}$, or he will seek employment elsewhere. Second, the contract must ensure that the subordinate's expected utility if he chooses H is at least as great as if he chooses L. This aligns the incentives of the subordinate with the wishes of the superior.[5] Finally, while meeting the first two objectives, the contract minimizes the expected payments to the subordinate. Because the superior is risk neutral, and given it is best for her to motivate H, minimizing the expected payments to the subordinate is the equivalent of maximizing her expected utility. The superior's program that satisfies these conditions and can be used to solve for the optimal contract is as follows.

General contracting program (GP):

$$\min_{w_g, w_b \geq 0} \; p_H w_g + (1 - p_H) w_b$$

subject to:

$$(p_H + p_O)\sqrt{w_g} + (1 - p_H - p_O)\sqrt{w_b} - c \geq \overline{U} \qquad \text{(P)}$$

$$(p_H + p_O)\sqrt{w_g} + (1 - p_H - p_O)\sqrt{w_b} - c \geq \sqrt{w_b} \qquad \text{(IC)}$$

Realistic Subordinate

This section analyzes the case of a realistic subordinate in which the subordinate and superior have identical beliefs about the subordinate's chances of success if he chooses H. That is, we solve GP where $p_O = 0$.

First best benchmark with realistic subordinate

As a benchmark, it is useful to establish the optimal contract absent an incentive problem. This involves solving (GP) without (IC). This reduced program would be valid if the subordinate was willing to follow the superior's instructions even though his own self-interest would dictate

[5] The program implicitly assumes that if the subordinate is indifferent between two alternatives he will choose the alternative that the superior desires.

otherwise. Another way this could be valid is if the superior could costlessly observe the subordinate and levy a sufficiently harsh penalty for not choosing H such that the subordinate would always choose H and the penalty would never be levied.

Table 1. Effects of over(under) confidence on optimal contracts

parameters	Example 1 realistic expectations	Example 2 slight overconfidence	Example 3 significant overconfidence	Example 4 underconfidence
p_H	0.50	0.50	0.50	0.5
p_L	0	0	0	0
c	50	50	50	50
p_0	0	0.05	0.15	- 0.05
$\overline{U}$	300	300	300	300
w_g	160	152.81	175.25	169.01
w_b	90	90	50.22	90
E_R (w)	125	121.405	112.39	129.51
E_{SUB} (Utility|H)	300	300	300	300
E_R (Utility|H)	300	295.45	271.10	300
E_{SUB} (Utility|L)	300	300	224.77	305.56
E_R (Utility|L)	300	300	224.77	300

p_H = probability of g, given H
p_L = probability of g given L
p_O = over(under) confidence parameter
c = disutility of H
$\overline{U}$ = expected utility from next-best alternative
w_i = payment (in thousands) conditional on outcome i, i = g, b
E [w] = expected wage based on superior's realistic beliefs
E_R[ ] = expectations based on superior's realistic beliefs
E_{SUB}[ ] = expectations based on subordinate's beliefs

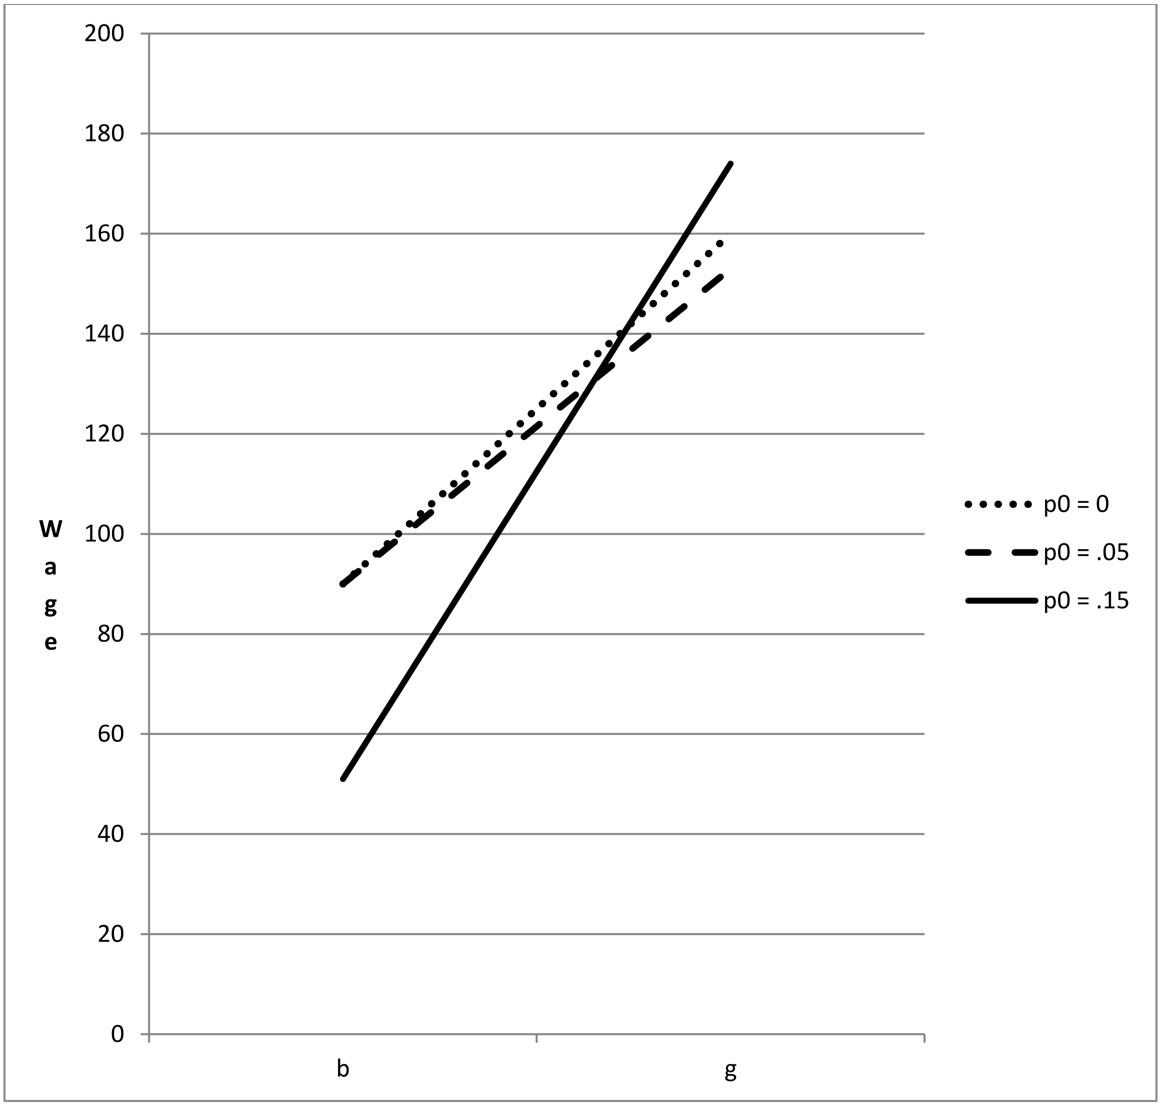

Figure 1. Performance-based wages in Examples 1, 2 and 3
Parameters:
$\overline{U} = 300$
$c = 50$
$p_H = 0.5$
$p_L = 0$
p_O = overconfidence parameter

A numerical example will be illustrative of the general solution to the first best problem. Table 1 provides the parameters for Example 1, wherein p_H is set equal to 0.5, c is set equal to 50, and $\overline{U}$, is set to 300. The solution is $w_b = w_g = 350^2 = 122{,}500$. The solution illustrates that absent incentive considerations, it would be optimal for the superior to bear all the risk, meaning the *realistic* subordinate is paid a flat wage. This contract is referred to as "first best" because (by assumption) it motivates the subordinate to choose H and it addition it optimally assigns risk.

Second best solution with realistic subordinate

The second best solution assumes the subordinate must be motivated to perform H. To obtain the second best solution, the (IC) constraint is reinstated. Of special interest is that motivating a risk averse subordinate to take an unobservable action is not free. The optimal contact where the subordinate's choice is not observable requires risk to be placed on the risk averse subordinate. To see this, substitute the solution above, $w_b = w_g = 122{,}500$ and note that it violates (IC). Therefore it would not motivate H. The second best contract sets $w_g = 160{,}000$ and $w_b = 90{,}000$. One way to look at the optimal second best contract is it supplies a wage of 90,000 as base pay, with an additional 70,000 bonus for good performance. This contract imposes significant risk on the risk averse subordinate.

We see the difference in the expected wage between the first best and second best cases is equal to 2,500; this difference is due to the fact that the subordinate must be compensated extra for receiving risky pay. This *risk premium* is the cost of placing risk on a subordinate who is risk averse.[6]

It is useful to generalize the numerical solutions above, which we derive in the Appendix. The optimal first best contract is to set a flat wage at $(\bar{U} + 2c)^2$. The optimal second best contract is $w_b = \bar{U}^2$ and $w_g = (\bar{U} + 2c)^2$ and the expected wage = $(\bar{U} + c)^2 + c^2 > (\bar{U} + c)^2$, indicating that providing incentives to a realistic subordinate when his actions are unobservable is generally costly from a risk sharing perspective.

Overconfident Subordinate

First best benchmark with overconfident subordinate

Now consider an overconfident subordinate, who believes his productivity when choosing H is greater than it actually is. Recall from above that the subordinate's belief about the probability of g given H is $p_H + p_O$, with p_H equal to the true probability and overconfidence captured by assuming $p_O > 0$. Unlike the first best setting with a realistic subordinate, it is no longer efficient for the subordinate to receive a flat wage – the subordinate is paid more for an outcome of g than for b. This occurs even though there is no need to place incentives and the superior is risk neutral and the subordinate is risk averse! De la Rosa (2011) refers to this as the "wagering effect" — this concept is fundamental to the understanding of overconfidence.

As in the previous section, we locate the first best solution by solving (GP) while suspending (IC). A numerical example is illustrative. Suppose that $p_O = 0.05$. In the first best solution $w_g = 145{,}298$ and $w_b = 97{,}269$. To see why it is optimal to assign risk to a risk-averse subordinate, first consider the expected utility using the subordinate's subjective beliefs. His expected utility is $0.55\sqrt{145{,}298} + 0.45\sqrt{97{,}269} - 50 = 300$. This is equivalent to the expected utility from the first best payment of 122,500 for a realistic subordinate, $\sqrt{122{,}500} - 50 = 300$. However, from the perspective of the realistic superior the expected compensation, 0.5(145,298) + 0.5(97,269) = 121,284, is less than 122,500. So the risky contract is equal to the flat wage contract in the eyes of the subordinate but better than the flat wage contract in the eyes of the superior. The key here is that the subordinate overestimates the probability of receiving w_g, so the superior increases the size of that payment. In order to reduce the expected compensation to the

[6] Notice that if the superior could costlessly monitor the subordinate (or if the subordinate was trustworthy), she could pay the subordinate a fixed wage that is greater than 122,500 but less than 125,000 and make *both* the superior and the subordinate (the "organization") better off than if the superior could not costlessly monitor.

subordinate, she decreases the payment for a bad outcome, w_b. Hence, an overconfident subordinate is optimally assigned risk even though there is no incentive problem to solve. Figure 2a displays the payments to the subordinate under the first best contract as a function of the parameter of overconfidence parameter p_O.

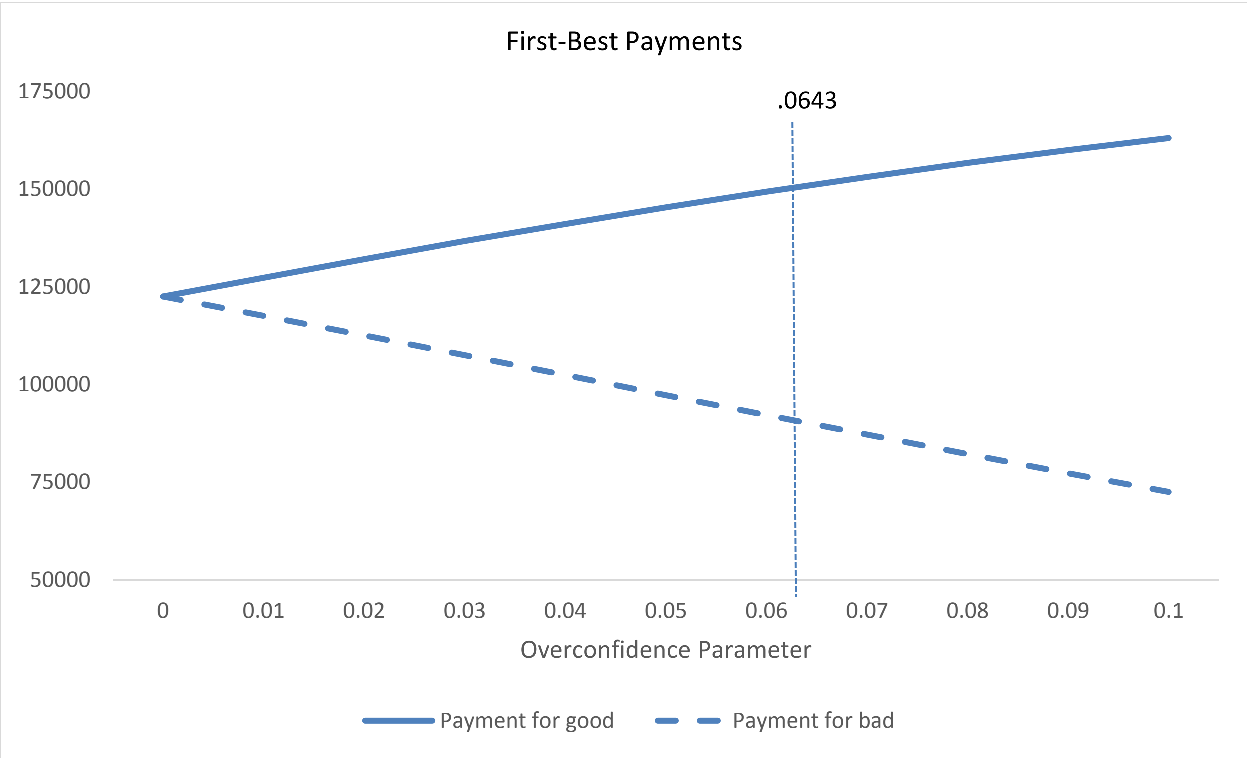

Figure 2a. First best payments to the subordinate as a function of overconfidence

<u>Superior's program:</u>

$$\min_{w_g w_b} 0.5w_g + (1 - 0.5)w_b$$

subject to:

$$(0.5 + p_O)\sqrt{w_g} + (1 - 0.5 - p_O)\sqrt{w_b} - 50 \geq 300$$

p_O = overconfidence parameter

wg = payment for good outcome

w_b = payment for bad outcome

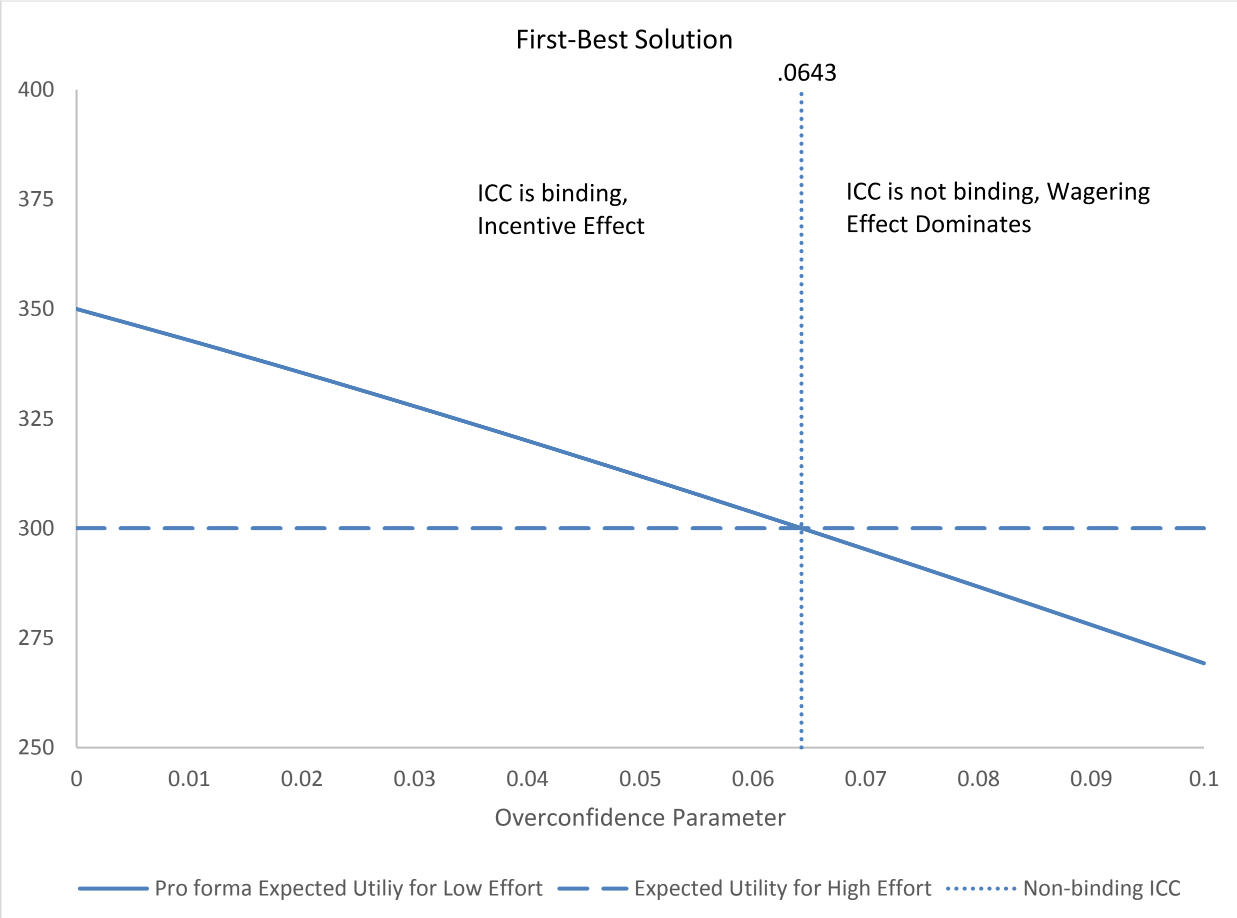

Figure 2b. First best solution, subordinate's expected utility

<u>Superior's program:</u>

$\min_{w_g w_b} 0.5w_g + (1 - 0.5)w_b$
subject to:
$(0.5 + p_O)\sqrt{w_g} + (1 - 0.5 - p_O)\sqrt{w_b} - 50 \geq 300$
p_O = overconfidence parameter
Subordinate's expected utility from choosing H: $(0.5 + p_O)\sqrt{w_g^*} + (1 - 0.5 - p_O)\sqrt{w_b^*} - 50$
Subordinate's pro forma expected utility from choosing L: $\sqrt{w_b^*}$

w_g^*, w_b^* are the wages from the solution to the superior's program.

The assumption underlying a first best contract is the superior can costlessly induce the choice of H through its observation or other means, so in equilibrium there is no need to provide risky compensation as an incentive. However, as an artifact of optimal risk sharing, an overconfident subordinate is paid more for an outcome of g than for b, and hence has partial incentives to choose H even if the program is not designed to provide any. Setting aside the cost of H, the subordinate's expected utility from choosing H is $0.55\sqrt{145{,}298} + 0.45\sqrt{97{,}269} =$

350 and his expected utility from choosing L is $\sqrt{97{,}269}$ = 311.9. The difference, 350 – 311.9 = 38.1, is not enough to offset the incremental private cost of H, 50, but it is closer than provided by the first best contract to a realistic subordinate. In Figure 2b are displayed the expected utilities, net of the subordinate's private cost of H, for the subordinate under the first best contract as a function of p_O. For p_O < .0643, as above, the first best contract does not provide sufficient incentives to induce the subordinate to choose H – hence, (IC) is tight. However, for $p_O \geq 0.0643$ the first best contract is sufficient to induce the subordinate to choose H, even though it was designed only with risk sharing in mind. In summary, to induce the proper action a realistic subordinate must receive incentives to choose H, whereas a sufficiently overconfident subordinate does so naturally.

The general first best solution with a subordinate who is not realistic is derived in the Appendix. The solution is $w_b = \left[(\bar{U} + c)\frac{(0.5-p_O)}{(0.5+2p_O^2)}\right]^2$ and $w_g = \left[(\bar{U} + c)\frac{(0.5+p_O)}{(0.5+2p_O^2)}\right]^2$. Note that if $p_O = 0$, w_b and w_g are equal to the first best payments provided in the previous section. For $p_O > 0$, the subordinate is paid more for a good outcome than a bad outcome, even though this puts risk on the risk averse subordinate.

In the appendix we further demonstrate that (IC) is not tight under the following condition.

$$\frac{c}{\bar{U} + c} \leq \frac{p_O + 2p_O^2}{0.5 + 2p_O^2}$$

Denote the value of p_O that makes this an equality by $\overline{p_O}$. Using the same parameters as in Example 1, where p_H is set equal to 0.5, c is set equal to 50, and $\bar{U}$ is set to 300, $\overline{p_O}$ = .0643. Thus, at p_O > 0.0643 the first best contract, the one that efficiently shares risk, also is able to induce the subordinate to choose H. For this reason de La Rosa (2011) would say that when p_O > 0.0643 the *wagering effect* dominates the *incentive effect*. Conversely, when p_O < 0.0643, the incentive effect dominates. We shall hereafter refer to a subordinate such that the wagering effect dominates as significantly overconfident and to a subordinate such that the incentive effect dominates as slightly overconfident.

It may be useful to interpret this example using the analogy of the pitcher discussed in the Introduction. Assume the pitcher does not like to work hard in practice. For a slightly overconfident pitcher, the team must include a constraint that makes high practice effort at least as good as low effort in the eyes of the pitcher, or else the pitcher will exert low effort. The constraint is costly because it imposes risk on the pitcher he would otherwise rather not face. The team takes advantage of the pitcher's overconfidence by lowering the bonus for good performance, while keeping the base salary constant. In this way the "spread" in payments to the pitcher is reduced as well as the expected payments, yet necessary incentives are kept in place for high effort; hence, the label "incentive effect."

In the case of the significantly overconfident pitcher, the pitcher's desire to wager on his ability to win 20 games is driving the solution. The cheapest way for the team to ensure the pitcher (using his subjective beliefs) obtains $\bar{U}$ in expectation is to let him wager. However, because the wagering on winning 20 games provides such strong incentives for high effort, the team need do nothing else in order to induce the pitcher to exert high effort; hence the label "wagering effect." In fact, de la Rosa (2011) in an extension to his model shows that effort is increasing in the subordinate's overconfidence even if the superior cannot observe effort. This is attributable to the fact that it is cheaper to motivate any (non-zero) level of effort for an overconfident subordinate than for a realistic one.

Second best Contract

Examples 2 and 3, summarized in Table 1, provide numerical illustrations of the optimal contract for an overconfident subordinate when the subordinate's choice is not observable. The parameters in Example 2 are the same as in Example 1, except the overconfidence parameter, p_O, is set to 0.05 instead of zero. From Figure 2b, this falls within the region where the incentive effect dominates. Therefore, the superior must put in place incentives to motivate the slightly overconfident subordinate to join the firm and choose H, based on the subordinate's subjective beliefs on the probability of good performance.

The solution to Example 2 still has the base pay at 90,000, as it was with the realistic subordinate in Example 1. However, the bonus has decreased from 70,000 to 62,810, for total pay of 152,810 for good performance. Exactly as in the case of the slightly overconfident baseball pitcher, the value of the bonus is lowered but sufficient incentives to choose H are maintained given the subordinate's subjective beliefs. The superior is clearly saving money, lowering one payment (for good performance) while leaving the other unchanged (for bad performance). The expected compensation has gone from 125,000 for a realistic subordinate to 0.5(152,810) + 0.5(90,000) = 121,405.

In Example 3 p_O is increased to 0.15. As shown in Figure 2b, p_O now falls within the region where the wagering effect dominates. Inspecting Table 1, the constraint on incentive compatibility is no longer tight, as we would expect when the wagering effect dominates. Using the subordinate's subjective beliefs, there is more than enough spread in the payments to motivate him to choose H. From a realistic subordinate who receives payments of 160,000 and 90,000 for good and bad outcomes, respectively to a significantly overconfident subordinate who receives payments of 174,249 and 50,522 for good and bad outcomes, respectively.

The Appendix derives a closed-form solution for the second best contract as a function of p_O. As stated in the previous section, the characteristics of the optimal contract depend on the magnitude of p_O. In particular, both (P) and (IC) are tight when p_O is positive but small (slight overconfidence), but only (P) is tight when p_O is sufficiently large (significant overconfidence).

As with the significantly overconfident baseball pitcher, the subordinate is wagering on his performance and the wagering creates sufficient incentives for high effort. If it is assumed that the superior is indeed realistic with her beliefs while the subordinate is truly overconfident, then as is shown in Table 1 the subordinate's actual expected utility is 271.1, rather than his next-best opportunity of 300. It is small wonder then that Humphery-Jenner et al. (2016) refer to region where the wagering effect dominates as the "exploitation hypothesis." Note from Figure 1 that the pay to a realistic subordinate does not dominate the pay to a significantly overconfident subordinate; that is, relative to less confident subordinates, the pay to a significantly overconfident subordinate for an outcome of b is lower and the pay for an outcome of g is higher. In addition, the slope is steeper, indicating the significantly overconfident subordinate is bearing additional risk.

An interesting corollary observation from the modeling is that it is possible that a project would only be undertaken with a sufficiently overconfident individual. Return to the parameters in Examples 1 and 2. Assume that if the subordinate chooses L the expected benefit to the superior (net of all costs *except* compensation) is zero, so it is not worthwhile to hire a subordinate for positive pay and have him select L. Also, assume instead that if the subordinate chooses H the expected benefit to the superior (net of all costs *except* compensation) is 123,000. If the subordinate is realistic, as in Example 1, the expected cost of his compensation is equal to 125,000, which is less than the expected benefit of 123,000 and therefore it is not rational for the superior to hire the subordinate to manage this project. However, if the subordinate is overconfident as in Example 2, his expected compensation of 121,405 is less than the expected benefit of 123,000 and therefore it is worthwhile for the superior to hire the subordinate to manage the project and choose H.

The model being illustrated, while quite useful, considers only a one-dimensional choice for the subordinate. In a richer setting a subordinate will have many choices that Boards of Directors cannot fully anticipate and contract for.[7] As an example of the dangers of a high-risk contract, consider several reasonable factors that are outside the model in the instructive case of MF Global. MF Global in 2010 was mostly a commodity and futures brokerage. They executed trades for clients and often held client funds to secure leveraged positions. One of the most important ways the firm made money was from clients' "segregated" funds. The funds were segregated in the sense that they in no way could be used to fund the operations of the firm. The firm earned income from the spread between what they paid clients on segregated funds and what they earned investing those funds. However, in the extremely low interest rate climate after the 2008 financial crisis, the spread narrowed quite a bit and the firm was losing money.

Enter Jon Corzine, the former head of investment bank Goldman Sachs and the former governor of New Jersey. In his account of MF Global's bankruptcy, Skyrm (2013) describes Corzine as a man of extreme hubris and overconfidence. More importantly, Skyrm (2013) relates that Corzine expected he could earn a billion dollars turning around MF Global; this on a base salary of $4 million per year. This is basically an all-or-nothing proposition given the relative size of $4 million and $1 billion. What Corzine proceeded to do is make a series of highly leveraged bets outside the normal course of MF Global's business practices.

In a nutshell, Corzine made short-term bets on the sovereign debt of European periphery nations, such as Italy, Spain and Ireland. The firm would not actually lose money unless there was a default, and this was deemed unlikely. However, the rates on these nations' debt were significantly higher than those at which the firm could borrow. In fact, none of the debt in which Corzine invested actually suffered a default. However, while the bets were on, the debt suffered declines in market value. Because of the leverage used in the trade, the declines triggered margin calls. Unlike his former firm Goldman Sachs, which had very deep pockets, MF Global had shallow pockets, and ultimately members of the firm embezzled segregated client funds in an unsuccessful attempt to avoid insolvency. In fact, Corzine had already been told that margin calls could bankrupt the firm, but he dismissed the head of risk management who gave him those warnings.

The story of MF Global is reminiscent of Long Term Capital Management (LTCM), another organization, this time a hedge fund, which was described as having extremely confident management (Lowenstein 2000). The hedge fund, formed with highly successful traders and academics in 1994, earned spectacular returns in its first three years. But in 1997 the returns started to decline, as others learned some of their strategies, and in response the hedge fund, in essence, doubled its bets when it returned to investors half of their investment. This left the hedge fund much more vulnerable to a liquidity squeeze. When LTCM was hit with massive margin calls in the wake of the Russian debt default of 1998, they were forced to accept a bailout from the other

[7] To be clear, in the models of Gervais et al. (2011) and de la Rosa (2011), the Board of Directors (superior) perfectly foresees the effects of the contract they offer the subordinate, so the vignettes we describe should not occur. Banerjee et al. (2015) provide some evidence than firms can select CEOs based on their overconfidence. It may be reasonable to assume that there will be unanticipated consequences to offering highly incentivized contracts.

investment banks, brokered by the NY Federal Reserve. The partners, having kept most of their own capital in the fund, lost heavily.

As indicated by the modeling, overconfident subordinates make large wagers on their abilities, and Jon Corzine and the partners of LTCM fit this description. Although we have not modeled the potential externalities that arise from such wagering, there is empirical evidence to suggest that significant overconfidence can be dangerous. In support of these notions Campbell et al. (2011) find that CEOs with relatively low and relatively high overconfidence are most likely to be fired, reflecting the fact that some overconfidence is helpful in aligning the risk preferences of CEOs and investors, but too much overconfidence can lead to dangerous bets by the CEO.

Finally, it should be noted that although our illustration follows more closely the standard principal-agent model of de la Rosa (2011), there is an extensive empirical literature on CEO overconfidence and investment decisions. In the basic model an overconfident CEO is assumed to have incentives naturally aligned with shareholders unlike the de la Rosa set-up. The CEO tends to overinvest due to his or her misperception of their ability to execute the investment. That is he makes mistakes. However, this is predicted to occur more often when the firm has free cash flows, because the CEO also over-estimates the value of the firm's equity and hence does not want to raise capital on secondary markets, which in his view tend to "under-value" the firm (Malmandier and Tate 2005). In empirical tests, as a proxy for overconfidence, researchers use the degree to which CEOs do not exercise vested options. The idea is that CEOs are under-diversified and should exercise sufficiently "in the money" options (and then sell the shares). Using these measures, research has provided empirical support (Malmandier and Tate 2015) for the theory of Malmandier and Tate (2005).

Returning to the agency model of de la Rosa (2011) wherein subordinate and agent preferences are not naturally aligned, the prediction from his continuous effort model is that effort from overconfident agents will exceed that from realistic agents. This prediction, however, rests on the assumption that principals are aware of agents' overconfidence and optimally contract to "exploit" these erroneous beliefs. Because an overconfident agent is less costly to motivate for *any* level of effort, in equilibrium effort is higher. A simpler hypothesis to test is, all else being equal, for any fixed level of bonus contracts, overconfident agents will contribute less effort because they believe less effort is needed to reach the bonus level than a realistic agent would. Effort is difficult to measure using archival data, which is why experiments may be helpful. Recent experiments on overconfidence and optimism include those by Hales et al. (2015), Mertins and Hoffeld (2015) and Thoma (2016). Such experimental procedures can be extended to look specifically at issues of effort.

Underconfident Subordinate

As pointed out by de la Rosa (2011), the model is perfectly suited to examine underconfidence as well as overconfidence. To do this, allow p_O to be negative. Before proceeding, recall that at p_O close to 0 the "incentive effect" dominated, wherein small increases in confidence led to less risky and lower expected compensation. With underconfidence the incentive effect always dominates the wagering effect. Therefore, a decrease in p_O will result in increased risk and increased expected compensation. Example 4, shown in Table 1, uses the same parameters as Example 3 but with $p_O = -0.05$.

Payments for outcomes of good and bad are 169,012 and 90,000, respectively, whereas for a realistic subordinate they are 160,000 and 90,000, respectively. Expected compensation is 129,506 as opposed to 125,000. The intuition behind the results is as follows. If a subordinate's subjective beliefs on his chances for success conditional on H are less than they truly are, that subordinate will assess the expected value of a bonus for success less than it truly is. Therefore, that subordinate will require a higher bonus in order to prefer H, resulting in higher expected compensation to the superior. Given that underconfidence is unlikely to confer advantages outside the model (such as higher effort and more cooperation that are associated with overconfidence), this result suggests that individuals are unlikely to find employment in areas where they are less confident. Further, individuals who are generally underconfident may tend to find employment in areas where wages are fixed, either due to low information asymmetry (perhaps because direct monitoring is inexpensive) or less of a need to provide motivation (because the cost of motivating H exceeds its benefits).

MODEL EXTENSION

In the preceding section, numerical examples illustrated recent findings on overconfidence and expected compensation. In this section, an extension of the model is offered through additional numerical illustrations. Up to this point, the numerical examples have shown that expected compensation to a subordinate is decreasing in his overconfidence. This is true independent of whether there is a control problem, that is, in both the first best and second best contracts. However, because expected compensation to the subordinate is higher when a control problem exists, it may seem intuitive that the worse the control problem, the greater the "benefit" of overconfidence to the superior. This is not the case, however. Below are two illustrations that help explain the relationship between overconfidence and the severity of the control problem. They show that subordinate overconfidence is *most* beneficial to the superior (assuming she has the correct assessment of outcome probabilities) when the control problem is *most* or *least* severe, and is *least* beneficial when the control problem is *moderately* severe.

The management control problem at the heart of the model is brought on by the information asymmetry between the superior and subordinate. The subordinate is aware of his action and the superior is not. If the superior could observe the subordinate's action, she would optimally write a contract to "force" the subordinate to take the desired action. Instead, the superior must use a publicly observable signal to compensate the subordinate and motivate H. The lack of observability causes an increase in expected compensation to the subordinate to compensate him for taking on excessive risk. Logically, then, the more informative the signal, the less severe the control problem. A signal that is extremely informative about the subordinate's choice leads to a much less severe control problem than a signal that is nearly uninformative regarding the subordinate's choice.

Given the use of a square root utility function, the informativeness of signal can be measured by the variance of the likelihood ratios: the higher the variance, the more informative the signal (Kim and Suh 1991). The likelihood ratio of each outcome is the probability of the outcome given the undesired action divided by the probability of the outcome given the desired action. So for the outcome of good the likelihood ratio is $\frac{p_L}{p_H}$, and for the outcome of bad the likelihood ratio is $\frac{1-p_L}{1-p_H}$. The mean likelihood ratio, given the subordinate is choosing H, is $p_H \frac{p_L}{p_H} + (1-p_H)\frac{1-p_L}{1-p_H} = 1$. Therefore the variance of the likelihood ratios is $p_H\left(\frac{p_L}{p_H}-1\right)^2 + (1-p_H)\left(\frac{1-p_L}{1-p_H}-1\right)^2$. Notice that as p_H goes to p_L, the variance of the likelihood ratio goes to zero, implying a severe control problem. Further, in our examples $p_L = 0$, so that the variance of the likelihood ratios simplifies to $\frac{p_H}{(1-p_H)}$, which is increasing as p_H approaches 1 from below. Figure 3 plots the variance of the likelihood ratios as a function of p_H for our examples wherein p_L = 0.

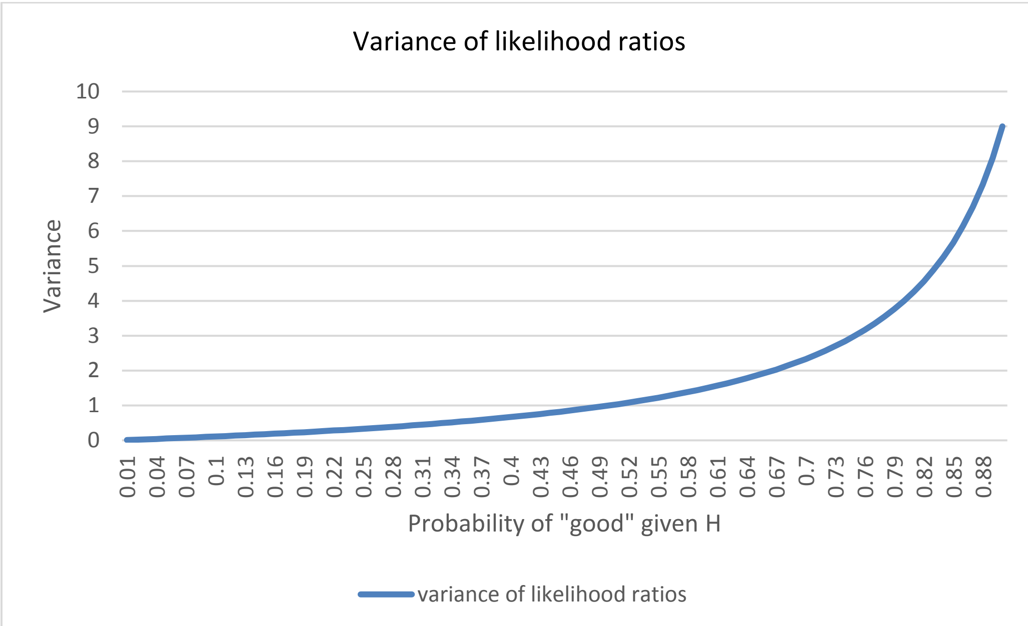

Figure 3. Variance of likelihood ratios as a function of p_H when p_L = 0

<u>Likelihood ratios:</u>

good outcome (g) = $\frac{p_L}{p_H}$

bad outcome (b) = $\frac{(1-p_L)}{(1-p_H)}$

<u>Variance of likelihood ratios:</u>

$$p_H\left(\frac{p_L}{p_H}-1\right)^2+(1-p_H)\left(\frac{1-p_L}{1-p_H}-1\right)^2$$

with $p_L = 0$: $\frac{p_H}{(1-p_H)}$

Figure 4 displays expected compensation for the subordinate for values of p_H ranging from 0.1 to 0.9, with p_O = 0.10 for the overconfident subordinate and all other parameters as set in the previous examples. Expected compensation is decreasing in p_H for the realistic subordinate. What is surprising is that the compensation to an *overconfident subordinate* is not monotonically decreasing in p_H. Expected compensation for the overconfident subordinate has an inverted U shape; expected compensation is lowest for severe and mild control problems, and highest for moderate control problems. To help with the intuition for this feature of the extended model, Table 2 includes the second best solution to the superior's program for both realistic and overconfident subordinates at p_H = 0.1, 0.5 and 0.9.

Table 2. Effect of p_H on optimal contracts

	Realistic subordinate ($p_O = 0.0$)			Overconfident subordinate ($p_O = 0.10$)		
p_H	w_g	w_b	E(wage)	w_g	w_b	E(wage)
0.1	640,000	90,000	145,000	396,900	78,400	110,250
0.5	160,000	90,000	125,000	163,092	72,485	117,789
0.9	126,420	90,000	122,778	122,500	0	110,250

$\bar{U}$ = 300
c = 50
p_L = 0
w_g , w_b are the wages from the solution to the superior's program.

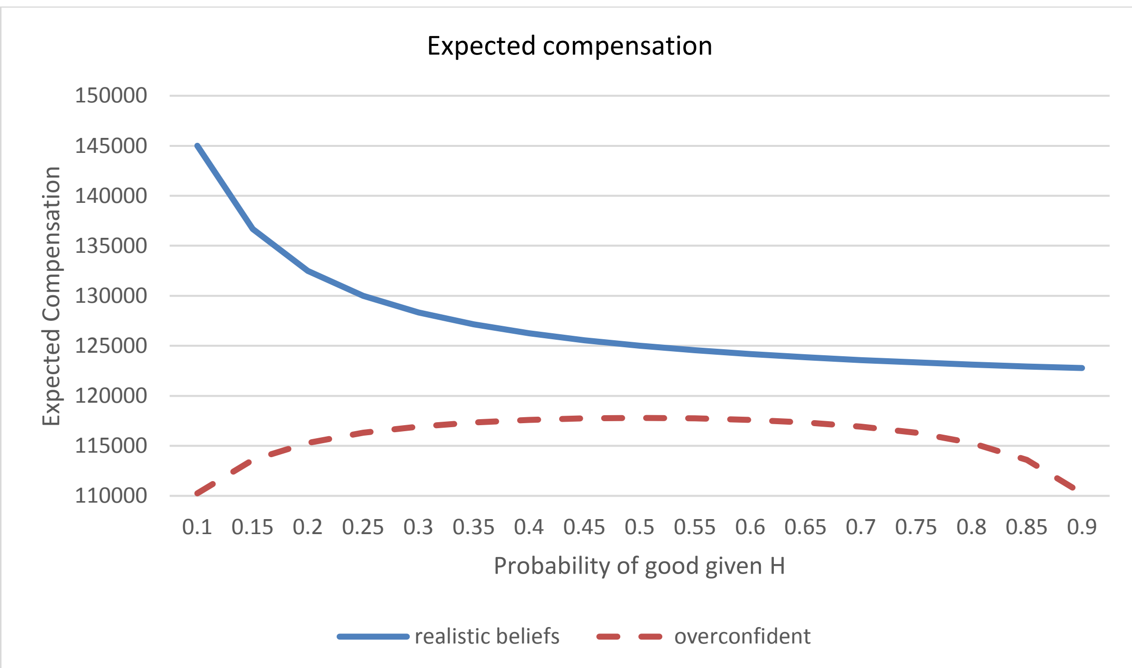

Figure 4. Expected compensation for realistic and overconfident ($p_O = 0.10$) subordinates as information asymmetry changes

Parameters:
$\bar{U}$ = 300
c = 50

$p_L = 0$
Expected compensation = $p_H w_g^* + (1 - p_H) w_b^*$
w_g^*, w_b^* are the wages from the solution to the superior's program.
For the realistic (overconfident) subordinate, p_O = 0.0 (0.10).

From Table 2, for a low likelihood of success (p_H = 0.1), the overconfidence parameter, p_O, is large relative to the true probability of success. The subordinate believes he will be successful (achieve an outcome of good) 20% of the time when in fact he will be successful 10% of the time. The superior exploits this by significantly reducing the payment for an outcome of good, from 640,000 to 396,900. Combined with the reduced payment for an outcome of bad, from 90,000 to 78,400, the superior saves 0.1(640,000 – 396,900) + 0.9(90,000 – 78,400) = 34,750 in expected compensation. The subordinate believes the probability of good and bad outcomes is 0.2 and 0.8 respectively, so that the expected utility of H is $0.2\sqrt{396{,}900} + 0.8\sqrt{78{,}400} - 50$ = 300, which is strictly greater than the expected utility of L, $\sqrt{78{,}400}$ = 280.

In contrast, for a high likelihood of success, such as p_H = 0.9, the subordinate severely underestimates the likelihood of failure. The subordinate believes there is a 0% chance of failure when there is in fact a 10% chance of failure. The superior exploits the subordinate's "erroneous" beliefs by reducing the payment for bad performance, in the example reducing it to 0. The subordinate's subjective beliefs are that there is no chance of receiving the payment of zero for a bad outcome. The superior saves 0.1(90,000 - 0) = 9,000 in expected compensation for bad performance, where 90,000 is the payment to a realistic subordinate for an outcome of bad. (In fact, if not for the implicit non-negativity constraints with the use of square root utility, the payment for bad performance can be set arbitrarily low.) In addition, the superior saves 0.9(126,420 – 122,500) = 3,528 in expected compensation for an outcome of good, for a total savings of 9,000 + 3,528 = 12,528.

Compared to extreme values for p_H, for intermediate levels, such as p_H = 0.5, the subordinate's subjective beliefs on the likelihoods of success and failure are relatively closer to the true likelihoods and there is less opportunity for the superior to exploit these beliefs. In the example, the superior saves 125,000 – 117,789 = 7,211 in expected compensation when p_H = 0.5, which is less than the amount saved for either extreme belief. (Recall that when p_H = 0.1, the amount saved is 34,750 and when p_H = 0.9, the amount saved is 12,528.)

To put the model extension in perspective, consider some examples from history. Gold rushes might be characterized by high subjective beliefs in success relative to the true probability. This is a case where the underlying probability of success, p_H, is low, but rewards are high. Small absolute increases in beliefs on the likelihood of success may result large differences in the estimate of the expected value of the venture. For example, if a prospector has a 1% chance of finding 10,000,000 in gold, the expected value is 100,000. But if the prospector (erroneously) believes there is a 5% chance of success the expected value is 500,000. Small errors in beliefs could have a major impact on behavior.

In contrast to gold rushes, the large leveraged bets made at MF Global mentioned earlier are a good example of where the underlying probability of success is high but the rewards are moderate (hence the need for leverage). Leverage increases risk, yet a trader who severely underestimates the likelihood of failure may not be aware of the risk. Taleb (2005) speaks at length about traders making large leveraged bets that very low likelihood events would not occur. He also discusses the confidence of these traders and how they over-estimated their ability to predict the future. If a trader can make frequent high-probability-of-success bets that yield modest profits if successful, but ruinous, career-ending failure if unsuccessful, would that trader make the bet? What if the trader assessed the likelihood of failure at 0.1% (1 in 1,000)? Then perhaps yes. What if the true probability of failure is 2% (1 in 50)? Then perhaps no. Again, small errors in assessing likelihoods can have large effects on behavior. These effects occur at either end of the probability-of-success distribution.

CONCLUSION

This paper outlines and provides intuition for the recent findings regarding how overconfidence and underconfidence of subordinates affect management control systems. Specifically, overconfidence eases control problems, enabling the superior to pay lower expected compensation. A very overconfident subordinate will, in a sense, be exploited in that he will receive lower expected compensation and riskier compensation than what he perceives. Under confidence, in contrast, always exacerbates the control problem, because it worsens the optimal sharing of risk. The existing literature is extended by demonstrating an interaction between the severity of the control problem (relative to the realistic subordinate) and the degree of subordinate overconfidence. Specifically, the reduction in expected compensation due to overconfidence is greatest when the control problem is extremely severe or minor. Conversely, this reduction in expected compensation is least when the control problem is of intermediate severity.

These findings may suggest there is a potential benefit to the organization from seeking out employees who are overconfident. In addition, the organization might benefit from manipulating the beliefs of its employees about their probability of success. However, in the short run this benefit may be at the (unknown) expense of subordinates.

REFERENCES

Banerjee, S., Dai, L., Humphery-Jenner, M., & Nanda, V. (2015). Top dogs: Overconfident executives and new CEO selection. Working paper. University of Wyoming.

Campbell, T., Gallmeyer, M., Johnson, S., Rutherford, J. & Stanley, B. (2011). CEO optimism and forced turnover. *Journal of Financial Economics*. 101, 695-712.

de la Rosa, L. (2011). Overconfidence and moral hazard. *Games and Economic Behavior*. 73(2), 429-451.

Gervais, S., Heaton, J., & Odean, T. (2011). Overconfidence, compensation contracts and capital budgeting. *Journal of Finance*. 66(50), 1735-1777.

Hales, J., Wang, L., & Williamson, M.W. (2015). Selection benefits of stock-based compensation for the rank-and-file. *The Accounting Review*. 90(4), 1497-1516.

Hirshleifer, D. A., Low, A. & Teoh, S. H. (2012). Are overconfident CEOs better innovators? *Journal of Finance.* 67, 1457–1498.

Humphery-Jenner, M., Lisic, L., Nanda, V., & Silveri, S. (2016). Executive overconfidence and compensation structure. *Journal of Financial Economics*. Forthcoming.

Kim, S., & Suh, Y. (1991). Ranking of accounting information systems for management control. *Journal of Accounting Research.* 29(2), 386-396.

Lowenstein, R. (2000). *When Genius Failed: The Rise and Fall of Long-Term Capital Management.* Random House: New York.

Malmendier, U., & Tate, G. (2005). CEO overconfidence and corporate investment. *Journal of Finance.* 60, 2660-2700.

Malmendier, U., & Tate, G. (2015). Behavioral CEOs: The role of managerial overconfidence. *Journal of Economic Perspectives*. 29(4), 37-60.

Mertins, V. & Hoffeld, W. (2015). Do overconfident workers cooperate less? The relationship between overconfidence and cooperation in team production. *Managerial and Decision Economics*. 36(4), 265-274.

Skyrm, S. (2013). *The Money Noose: Jon Corzine and the Collapse of MF Global*. Brick Tower Press: New York.

Taleb, N. (2005). *Fooled by Randomness: The Hidden Role of Chance in Life and Markets*. Random House: New York.

Thoma, C. (2016). Under- versus overconfidence: an experiment on how others perceive a biased self-assessment. *Experimental Economics*. 19(1), 218-239.

APPENDIX
Derivation of optimal solutions

The general program to find an efficient incentive compatible contract is as follows.

$$\min_{w_g w_b} \; p_H w_g + (1 - p_H)\, w_b$$

subject to:

$$(p_H + p_O)\sqrt{w_g} + (1 - p_H - p_O)\sqrt{w_b} - c \geq \bar{U} \qquad \text{(P)}$$

$$(p_H + p_O)\sqrt{w_g} + (1 - p_H - p_O)\sqrt{w_b} - c \geq \sqrt{w_b} \qquad \text{(IC)}$$

As in most of the examples, p_H is set equal to 0.5.

1. First best solution

The first best solution implies (IC) will not be tight. Let $U_g = \sqrt{w_g}$ and $U_b = \sqrt{w_b}$ (and so $w_g = U_g^2$ and $w_b = U_b^2$). Then using the fact that (P) is tight one can solve for U_g:

$$U_g = \frac{\bar{U} + c}{0.5 + p_O} - \frac{0.5 - p_O}{0.5 + p_O}\, U_b$$

Substituting for U_g in the objective function and setting the derivative with respect to U_b equal to 0 leads to the following expressions for U_g and U_b.

$$U_g = (\bar{U} + c)\frac{0.5 + p_O}{0.5 + 2p_O^2}$$

$$U_b = (\bar{U} + c)\frac{0.5 - p_O}{0.5 + 2p_O^2}$$

1.a. Realistic subordinate

If the subordinate is realistic, $p_O = 0$, which leads to the solution that the subordinate receives a constant pay.

$U_g = U_b = (\bar{U} + c)$

$$w_g = w_b = (\bar{U} + c)^2$$

Also, the expected wage paid by the superior to the subordinate is $(\bar{U} + c)^2$.

Notice this solution would violate (IC). To see this, observe that (IC) requires that $p_H(\sqrt{w_g} - \sqrt{w_b}) \geq c$ and because $c > 0$, (IC) requires that $w_g > w_b$.

1.b. Overconfident subordinate

If the subordinate is overconfident, $p_O > 0$, and inspecting the solution, $w_g > w_b$. The superior and subordinate both end up bearing a portion of the risk, due to their different beliefs about the outcome of the performance measure given H. Further, w_g is *greater* than w_b, which could provide enough incentive for the subordinate to choose H if c were sufficiently small.

2. Second best solution

In this section it can no longer be assumed that (IC) is loose.

2.a. Realistic subordinate

As explained above, if the agent is realistic he would be paid a constant wage if incentive compatibility were not a concern. Therefore, now *both* (P) and (IC) are tight. Solving both (P) and (IC) as an equality yields $w_b = \bar{U}^2$ and $w_g = (\bar{U} + 2c)^2$, with the expected wage equal to $0.5\bar{U}^2 + 0.5(\bar{U} + 2c)^2 = (\bar{U} + c)^2 + c^2$, which exceeds the first best expected wage because the subordinate must be paid a premium in order to accept this risky contract.

2.b. Overconfident subordinate

There are two cases, one where the first best solution satisfies (IC) and one where it does not. (IC) requires that $(U_g - U_b) \geq \frac{c}{0.5+p_O}$. Substituting the first best solutions for U_b and U_g implies (IC) is satisfied if the following is true:

$$\frac{c}{\overline{U}+c} \leq \frac{p_O + 2p_O^2}{0.5 + 2p_O^2}$$

This inequality will be violated if p_O is sufficiently small, which is as shown surely true if $p_O = 0$. In this case both (P) and (IC) are tight, and the solution is found by satisfying the two inequalities. The solution is as follows:

$$U_b = \overline{U}$$

$$U_g = \overline{U} + \frac{c}{0.5 + p_O}$$

$$w_b = \overline{U}^2$$

$$w_g = \left[\overline{U} + \frac{c}{0.5 + p_O}\right]^2$$

If, however, p_O is large, only (P) is tight and the solution is as provided above under the first best scenario.

ABOUT THE AUTHORS:

Anthony Meder received his PhD. in Accounting from The Ohio State University in 2011 and his Bachelor's of Business Administration in 2002. He is currently an Assistant Professor in the School of Management, Binghamton University. His research interests include accounting regulation for the financial industry, corporate governance and debt contracting.

Steven Schwartz received his PhD. in Accounting from The Ohio State University in 1997 and his Bachelor's of Arts in Economics in 1987. He is currently an Associate Professor in the School of Management, Binghamton University. His research interests include other-regarding behaviors, agency contracting and financial history.

Eric Spires received his PhD. in Accounting from the University of Illinois in 1987 and his Bachelor's of Business Administration from Eastern Kentucky University in 1976. His is currently an Associate Professor in the Fisher College of Business, The Ohio State University. His research interests include, Decision making in both accounting and auditing contexts.

Richard Young received his PhD in Accounting from The Ohio State University in 1984 and his Bachelor's of Arts in Economics from Rutgers University in 1978. He is currently a professor in the Fisher College of Business, The Ohio State University. His research interests include the uses of accounting information for the purposes of management control.