Distortion and Risk in Optimal Incentive Contracts.

by

George Baker

Forthcoming, *Journal of Human Resources*

Abstract

Performance measurement is an essential part of the design of any incentive system. The strength and value of incentives in organizations are strongly affected by the performance measures available. Yet, the characteristics of valuable performance measures have not been well explored in the agency literature. In this paper, I use a multi-task model to develop a two-parameter characterization of performance measures and show how these two parameters—distortion and risk—affect the value and use of performance measures in incentive contracts. I show that many complex issues in the design of real world incentive contracts can be fruitfully viewed as trade-offs between these two features of performance measures. I also use this framework to analyze the provision of incentives in several specific environments, including R&D labs and non-profit organizations.

I would like to thank Nancy Beaulieu, Bob Gibbons, Brian Hall, Ed Lazear, Sherwin Rosen, Jasjit Singh, Lars Stole, an anonymous referee, and participants at the NAS Conference on “Designing Incentives for the Production of Human Capital” for comments on this paper. This research was funded by the Harvard Business School Division of Research. The data used in this article can be obtained beginning October 2002 through October 2005 from George P. Baker, Harvard Business School, Baker Library 283, Soldiers Field Road, Boston, Massachusetts 02163.

I. Introduction

The provision of incentives to individuals and groups in organizations is one of the central problems in the economics of the firm. A long and varied literature considers the question of what optimal incentive contracts look like (see Gibbons 1998, for a review). Most of this literature examines the use of “risky” performance measures, and as a result focuses on what Gibbons calls “the much studied trade-off between incentives and insurance.” Yet, in most incentive contracts in the real world, risk is not a central issue: with the exception of stock-based plans for top executives, most compensation arrangements in fact impose very little risk on employees. In addition, as Predergast (2000) points out, the data do not confirm the existence of a trade-off between risk and incentives.

In many incentive contracts, the central issue is not risk, but what Steven Kerr calls "The Folly of Rewarding for A While Hoping for B" (Kerr 1975). Consider the incentive plan tried by Lincoln Electric to motivate typists in its secretarial pool: they paid a piece rate for each key stroked. (Fast and Berg 1975) The plan was abandoned when it was discovered that secretaries were spending their lunch hours tapping the same key over and over. Understanding this aspect of incentive contracting has been the objective of the so-called “multi-tasking” literature. This literature has been concerned not with risk, but with “distortion.” The papers in this literature get the same result as those in the insurance literature: the firm reduces the strength of the incentive contract. But the reason is not to avoid imposing risk on the agent, but to avoid rewarding the wrong behavior.

The multi-tasking literature has evolved along two almost independent paths, one in economics and one in accounting. In economics, Holmstrom-Milgrom (1991) and Baker (1992) each develop models that show why principals avoid strong incentives for agents who face many effort margins. Holmstrom-Milgrom shows that the principal distorts incentives when differences in measurement accuracy lead them to focus risk averse agents more on some tasks than others. Baker shows that even when agents are risk neutral, the constraint that the principal generally cannot pay for what she really cares about leads to effort distortion.

In accounting, Feltham-Xie (1994) build a model (very similar to the one developed in this paper) that demonstrates how the incompleteness of most managerial compensation contracts leads to distorted incentives. In their model, the fact that agents take many more actions than the firm can measure leads to inefficiency in agent effort levels across all tasks. They also show how additional performance measures can mitigate, but generally not eliminate, this problem. A recent paper in this literature, Datar, Kulp and Lambert (forthcoming) asks explicitly about how principals should weight multiple, distorted performance measures.

In this paper, I build on both strands in this multi-tasking literature to develop a simple two-parameter characterization of performance measures that captures many of the problems, and clarifies much of the intuition, in the use of incentive contracting in organizations. The contribution of the paper is twofold. One is an intuitive geometric interpretation of and trigonometric expression for distortion in performance measures. The second contribution is to show how this two-parameter characterization and its interpretations can capture and explain many of the issues that plague actual incentive plans.

A. Distortion Relative to What?

In order to characterize a distorted performance measure, it is necessary to define an undistorted performance measure. One particular performance measure plays a crucial role in this paper: the total value of the organization. This performance measure (labeled V) should be thought of as capturing the present value of all future net benefits to the residual claimant of the organization. For publicly traded companies, V represents the market value of equity; in this case, firm value is an observable, contractible, and frequently used performance measure in incentive contracts. All types of stock-based compensation plans, including bonuses based on stock price, as well as executive stock grants, option grants, Employee Stock Ownership Plans, phantom stock, and stock appreciation rights are examples of incentive plans based on firm value. I make two assumptions about the nature of firm governance and equity markets that affect the performance measure V: (1) the objective of a publicly traded firm is to maximize the value of the firm to shareholders; and (2) the stock price accurately reflects this value. Thus, in this paper, incentives plans based on equity value are undistorted: by definition they provide incentives that are perfectly aligned with the organization's objective.[1]

In many organizations, the total value of the organization cannot be used in an incentive plan. In privately held firms, not only is the financial value of the organization unknown, it is likely not be the owner's objective. Rather, the objective is to maximize the owner's utility, and her utility is clearly not a contractible performance measure.

This problem of non-contractible organization value is even more acute in non-profit organizations and government agencies. In organizations of this type, it may not even be possible for managers to agree on and specify the organization's objective. Such organizations are often characterized by particularly difficult incentive problems; they cannot use stock based incentive devices, and the absence of a well-articulated organizational objective hampers the design of an efficient performance measurement system. If the top level objectives are not known, then how is the organization to measure the performance of individual employees? I will argue below that the difficulty in defining "good" performance measures in non-profit organizations is one reason for the weak incentives that so often characterize organizations of this type, and for the dysfunctional consequences that often arise when these types of organizations try to use strong incentives.

The paper is organized as follows. In Section 2, I develop the model and derive the optimal slope first for a risky but undistorted performance measure, and then for a risky and distorted performance measure, when no other performance measure is available. The main intuition of the model is evident from this derivation: incentives are optimally weaker when performance measures are either riskier or more distorted. In Section 3, I derive the optimal incentive contract when both a distorted performance measure and a risky but undistorted measure are used. I derive many of the same results as in the single performance measure case, along with some new results on relative performance evaluation. Section 4 examines the trade-off presented by the choice between risky and distorted performance measures, showing how many issues in the design of incentive contracts are fruitfully viewed as a choice between distortion and risk in performance measurement. Section 5 concludes with a discussion of future work, as well as a discussion of incentives for innovation and the design of incentive programs in organizations without well-defined objectives, such as non-profit institutions.

II: Single Measure Incentive Plans

A. Incentive Contracts Based on Firm Value

Consider the following simple linear production function:

$V(\mathbf{a},\varepsilon) = \mathbf{f}\bullet\mathbf{a} + \varepsilon$,

where:

V is firm value,

a is an n-dimensional vector of actions that the employee can take = $\{a_1,a_2,...a_n\}$

f is an n-dimensional vector of marginal products of those actions on firm value = $\{f_1,f_2,...f_n\}$. The elements of this vector can take on positive or negative values, as can the action choices of the employee.

ε represents effects which are, from the employee's perspective, random. These might be acts of nature or the uncontrollable actions of others in the firm. It has variance σ^2_ε.

Firm value is thus a function of controllable actions (the employee's choice vector **a**) and uncontrollable events (ε). The firm's problem is to design a linear compensation scheme for the employee—a base salary (s) plus some share (b_v) of V—that induces the employee to take actions that maximize the net value, subject to the employee's participation in the scheme. The firm is risk neutral, while the employee is risk averse with a utility function that places positive utility on expected compensation, and negative utility on the variance of compensation and on effort in each of the tasks in his action space. If V is interpreted as the total value of the equity of the firm, then b_v can be thought of as the share of the firm owned by the employee.

Employee's Expected Utility = $E(s+b_vV) - h\,\mathrm{var}(s+b_vV) - \sum_{i=1}^{n} a_i^2/2$

The solution to this problem is an optimal slope:[2]

(1) $$b_v^* = \frac{F^2}{F^2 + 2h\sigma_\varepsilon^2}$$

where

F is the length of the vector of marginal products: $F = \|\mathbf{f}\| = \sqrt{\sum_{i=1}^{n} f_i^2}$,

h is the employee's coefficient of absolute risk aversion,

σ^2_ε is the variance of the noise in the production function.

Note that b_v^* is decreasing in the employee's risk aversion (h) and in the variance of outcomes (σ^2_ε). It is increasing in the "length" of the vector **f**. This length is a measure of the contribution of the employee's actions to firm value relative to the contribution of the noise in the production function. The greater the effect of the employee's actions on V, the higher is the optimal b_v.

These are very standard results in the literature. The optimal weight on the performance measure is a function of the signal (the importance of effort in determining V) and the noise (the variance of ε) (Banker and Datar 1989, Holmstrom and Milgrom 1991). Indeed, equation (1) can be rewritten as:

(2) $$b_v^* = \frac{S_V}{S_V + 2h}$$

where

$S_v = \frac{F^2}{\sigma^2_\varepsilon}$, which is the signal-to-noise ratio.

This expression makes it clear that, from the standpoint of optimal incentives, "controllability" and riskiness are opposite sides of the same coin. When an employee can take actions that have a large effect on outcomes (relative to the amount of uncontrollable noise), then he should have an incentive scheme with a high slope. When, on the other hand, there is a lot of riskiness (that is, controllability is low), then the slope should be small.

B. Incentive Contracts Based on a Distorted Performance Measure

As discussed above, in many situations, the total value of the organization is not contractible, and thus cannot serve as the basis for an incentive contract. Even in publicly traded companies, where V is contractible, Equations 1 and 2 make it clear that the optimal slope of an incentive contract for most employees will be quite small. For these reasons, most incentive contracts are based on measures other than the total value of the organization. I now assume the existence of a performance measure, P, that is also a linear function of the employee's actions.

$P(\mathbf{a},\varphi) = \mathbf{g}\bullet\mathbf{a} + \varphi.$

g is an n-dimensional vector of marginal products of action on the performance measure,

φ is the effect of uncontrollable events on the performance measure. It has variance σ^2_φ and may be correlated with ε. The correlation between φ and ε is ρ.

P can be any piece of information that the employee and the firm can observe and contract on. The critical difference between P and V is that the marginal product of each type of action on P may differ from the marginal product of this action on V. Thus, for instance, let V represent the stock price and P represent this year's accounting earnings for the firm. Consider an action that an executive might take, such as deciding to invest in an R&D project. Such an action might have a positive effect on the stock price, but a negative impact on this year's accounting earnings (since the spending must be booked in the current year, and the revenues may not materialize for years.) A manager that invested in such projects would increase firm value, but reduce (short run) accounting profits. Or consider certain accounting manipulations, such as changing the timing of revenue recognition or expense accrual policies. Such manipulations have no effect on the cash flows to the firm, and if they are transparent to the market, no effect on the stock price. But they affect accounting profits. In both of these cases, paying executives on the basis of accounting profits may induce them to engage in wasteful or even counterproductive activities. These are example of distortion in performance measurement.

Examples of distorted performance measures are not, of course, restricted to for-profit companies. Consider an incentive program for elementary school teachers that uses student test scores as the performance measure. Suppose that the true objective of the school is to educate students to be good citizens and successful members of the society. There are many things that a teacher can do to achieve this objective, some of which might also improve test scores. However, there are also things that a teacher can do (such as "teaching to the test,") which will improve the performance measure while having little effect on the true objective of the school system, and things that the teacher can do to achieve the true objective (such as developing an innovative curriculum) that may not improve test scores. The fact that some activities affect the schools' true objective differently from how they affect the performance measure is what makes the performance measure distorted.

The effect of this sort of distortion is evident in the solution to the problem of the optimal use of performance measures in incentive contracts. Assume that the firm pays the employee based only on P, in the following way:

$$\text{Pay} = s + b_p P.$$

The optimal slope of an incentive contract based on P is:

$$(3) \quad b_p^* = \frac{FG\cos\theta}{G^2 + 2h\sigma_\phi^2}$$

where

F is the length of the vector of marginal products on firm value,

G is the length of the vector of marginal products on the performance measure,

h is the measure of the employee's risk aversion,

σ^2_φ is the variance of the noise in the performance measure,

θ is the angle between the vectors **f** and **g**.

Equation (3) can be rewritten as:

$$(4) \quad b_v^* = \left(\frac{F}{G}\right)\frac{S_p\cos\theta}{S_p + 2h}$$

where

$S_p = \frac{G^2}{\sigma_\varepsilon^2}$, which is the signal-to-noise ratio of the performance measure.

Equation (4) is very similar to equation (2). (F/G) is a normalization term (discussed below) that adjusts b_v^* if the units of P are different from the units of V. The other difference is cosθ in the numerator.

Cosθ is a measure of how distorted the performance measure is with respect to firm value. Higher cosθ means lower distortion. A performance measure will be relatively undistorted if a set of actions taken by the employee that increases P also increases V.[3] If cosθ = 1, then P provides the same effort incentives that V provides. Higher distortion in the performance measure leads to a lower weight placed on that performance measure in an incentive contract.

Figure 1 provides a simple illustration of how cosθ captures distortion in a performance measure. Consider a model with two tasks, a_1 and a_2, where $V = f_1a_1 + \varepsilon$ and $P = g_1a_1 + g_2a_2 + \varphi$. Thus, the marginal product of task 1 on firm value is f_1, while the marginal product of task 2 on firm value is zero. The employee can take an action (a_2) that increases the performance measure without increasing firm value.

The figure shows vectors **f** and **g**, and θ, the angle between them. If g_2 is not zero, the performance measure is distorted, because some elements of **g** are different from some elements of **f**. Cosθ is a summary measure that captures the extent to which a performance measure generates good signals to the employee about what actions to take. The discussions above (about R&D spending, accounting manipulation, and teacher incentives) all describe situations in which the marginal products on the performance measure are different from the marginal products on organizational value.

A comment is in order here about how to interpret a change in cosθ, and how such a change affects the scaling of P. In the comparative static analyses presented below, cosθ will be changed without changing the length of the vector **g**. Being able to conduct this type of comparative static is an important innovation of this paper, since it allows me to isolate the effects of changes in the distortion of the performance measure alone. Reference to Figure 1 shows why the intuition of this analysis is not obvious. Note that a small increase in g_1 will, *ceteris paribus*, decrease θ and thus increase cosθ, but will also result in an increase in the length of the vector P. Thus, the signal-to-noise ratio of P increases, and the scaling of P to V changes. The overall effect of such a change on b_p^* is ambiguous. However, Equation (3) shows that a change to **g** that decreases θ (increases cosθ) but preserves the length of the vector has an unambiguous effect on the strength of incentives.

A simple set of comparative statics will help clarify this intuition. Consider a change in P that changes only its units (for example, a farmer pays a piece rate on pounds of tomatoes picked, rather than on kilograms). Such a change has no effect on the actual productivity of labor, nor on the distortion or the signal-to-noise ratio of the performance measure, but it will force the optimal piece rate downward, since the new performance measure is 2.2 times greater than the old. The piece rate must decline because a given level of effort has a greater effect on P, but no greater effect on V. Equation (4) shows how b_p^* changes.

Now consider a change in P that increases G, the length of the **g** vector, without changing either the real productivity of effort (F), the distortion (cosθ) or the noise (σ^2_φ). The effect on b_p^* is ambiguous. The signal-to-noise ratio of P has increased, but there is also a scaling effect, similar to the one that occurred with the change from kilograms to pounds: again, a given level of effort has an increased effect on P with no change in its effect on V.[4] In all of the results presented below, these scaling effects will be held constant by doing comparative statics on cosθ and σ^2_φ and leaving G, the length of the **g** vector, unchanged.

Three results are easily derived from equation (3).

(R1) $\dfrac{d\,b_p^*}{d\cos\theta} > 0$: the weight placed on a performance measure in an optimal linear incentive contract decreases with distortion;

(R2) $\dfrac{d\,b_p^*}{d\,\sigma_\phi^2} < 0$: the weight placed on a performance measure in an optimal linear incentive contract decreases with the noise in the performance measure;

(R3) $\dfrac{d\,b_p^*}{d\rho} = 0$: the weight placed on a performance measure in an optimal linear incentive contract is independent of the correlation between V and P.

(Proofs are by inspection of Equation (3).)

As noted above, the more distorted is the performance measure, the lower is the optimal b_p^*. The intuition behind this result is simple: the costs of distortion (in terms of wasted effort and/or actions that actually reduce firm value) leads the firm to optimally reduce the weight placed on P. Note also that it is possible for b_p^* to be zero; if the vector **f** is orthogonal to the vector **g**, then cosθ is zero, and b_p^* is zero.[5] Thus the firm places no weight on such a performance measure.

Equation (3) shows the main contribution of this paper by highlighting, in simple terms, the role of two "problem parameters" that determine the use of any performance measure: distortion and risk. The second of these problem parameters, risk, has been the focus of agency theory research for decades. Yet, as discussed in the introduction, risk does not seem to be the main factor limiting the use of incentives in many situations. Organizations are full of low noise performance measures that nonetheless are not used as part of any incentive contract. And the reason is obvious: the use of such a performance measure would induce useless or counterproductive behavior.

Result 3 above is important. It says that the weight placed on a performance measure (when used alone) is unrelated to the correlation between V and P.[6] What matters is not whether *the measure* is correlated with firm value, but whether the *margins on the measure* are similar as captured by cosθ. Much literature in this area has failed to recognize this subtle but crucial distinction, and much confusion about the use of performance measures has ensued as a result. Consider first the large literature in accounting that has attempted to measure the correlation between various accounting measures and stock price. One of the objectives of this research has been to determine whether incentive contracts should be based on accounting measures. The model presented here shows that measuring the correlation between accounting numbers and stock price is measuring the wrong thing. The correct measure (which is unfortunately much harder to assess) is whether accounting profits move with managers' actions in the same way that stock prices do. A firm may optimally place a very small weight on accounting numbers, even if they are highly correlated with the stock price, if it is concerned that such an incentive contract would lead to gaming, manipulation, or other dysfunctional side effects.[7]

Heckman, Heinrich, and Smith (1997) (HHS) provide another example of the importance of the distinction between the *correlation* of performance measures and their *distortive effects*. They examine the effectiveness of performance standards in Job Training Partnership Act (JTPA) centers. These centers are mandated to provide job training for the disadvantaged, and are rewarded (through increases in funding) for their performance in successfully placing their clients. The problem, as described by HHS, and Courty and Marschke (1997) is that the performance standards give JTPA bureaucrats incentives to game the system by “cream skimming” (accepting into the program only those clients most likely to place well) and strategically terminating clients (both early and late) from the program. The root of this problem, of course, is that while the true objective of the JTPA is to *increase the employability* of the disadvantaged, what is being rewarded is the eventual *employment* of those who complete the program. One of the HHS hypotheses is that the degree of gaming in a center should depend on how highly correlated employment is with increases in employability. However, (as HHS point out) this is not quite correct. The amount of bureaucratic gaming is not predicted by the correlation of measured outcomes with desired outcomes, but rather by the correlation of the effects of employee actions on measured outcomes with the effects on true objectives. Unfortunately, this latter correlation is much harder to estimate.

The discussion above has focused on how the optimal weight on a performance measure varies with the performance measure's risk and distortion. Figure 2 illustrates the main results: it shows lines of constant b_p^* as a function of σ_φ^2 and (1-cosθ). The strength of incentives placed on a distorted and noisy performance measure decreases with both risk (σ^2_φ) and distortion (1-cosθ). If the margins vector on P is very different from that on V, so that cosθ is negative, b_p will be negative.

The point high on the σ^2_φ axis represents V, firm value. Firm value (assumed so far to be unavailable as a performance measure) has high risk, but zero distortion (by definition). If it were available as a performance measure, it would receive low weight because of its noise. If the firm has available a contractible performance measure that combines low noise with low distortion, it can use strong incentives; if it does not, incentives will be optimally weak.

C. Issues in the Use of Distorted Measures

Several issues and puzzles in the literature, and in the design of incentive contracts are easily understood using this two-parameter characterization of performance measures.

1. Controllability and Incentive Strength

Traditional agency theoretic treatments of the incentive problem tend to predict that "controllable" (that is, high signal-to-noise ratio) performance measures will be used more extensively as they allow stronger incentives without requiring a high risk premium for the employee. But the theory in this paper suggests that this simple prediction will not hold, since distortion in the performance measure may also drive down incentive strength. Some high signal-to-noise performance measures (pieces produced in a machine shop) may get great more in an incentive scheme, while others (keys stroked by a secretary) may receive low weight. The distinction between these two examples is simply whether employees can take actions that increase the performance measure without simultaneously increasing organizational value. If not, then the performance measure has low distortion, and can be profitably given high weight. If so, then using the performance measure will lead to significant dysfunctional behavior. If, as I argue below, risk and distortion tend to be negatively correlated in the performance measures observed in the world, then the there will be no negative correlation between risk and incentive strength, and indeed the correlation could be positive. [8]

2. Paying for Inputs versus Paying for Outputs

Another issue which this framework helps to clarify is the distinction between paying for "inputs" versus paying for "outputs." Lazear 1995 argues that:

> Workers are sometimes paid on the basis of some time unit, such as an hour, day, week, or year, and sometimes on the basis of some measured output, such as number of pieces of fruit they pick. . . . I define salary as compensation based on input and contrast it with a piece rate, which is compensation based on output. Hourly wages are salaries, by this definition. . . .”

The theory developed above suggests a different interpretation of the distinction between inputs and outputs. I treat all data on which the organization and the employee can contract as potential performance measures. *Almost any piece of data can be used as a performance measure:* profits reported, widgets produced, hours worked, keys stroked, etc. The relevant considerations for any performance measure are how risky it is, and how much it distorts incentives. Those performance measures that are relatively undistorted but risky (profits reported, widgets produced) tend to be called "outputs;" those that are low risk but narrow or distortionary (hours worked, keys stroked) tend to be called inputs. The weight placed on a performance measure (and thus whether the pay scheme will look like salary or like a piece rate) will be determined by the tradeoff between risk and distortion. Some “inputs” (for example, billable hours in a law firm) might receive high weight, while some “outputs” (student test scores) might receive low weight.

3. Degrading Performance Measures

Another puzzle elucidated by this model is the tendency for seemingly informative performance measures to "degrade" or become dysfunctional when they are used for incentive purposes. John Darley, a psychologist who has studied dysfunction in control systems, calls this problem "How Good Numbers Go Bad” (Darley 1991). He says:

> The essence of my claim about control systems can be quickly stated. In any organizational system, it is likely that some “numbers” exist that provide reasonable evidence about the success of that organization. Profits for a business, cases closed for a police department, or test scores for children graduating from

> high school are all reasonable candidates for such standing; they tell us a good deal about the success of the organization that produces those outcomes.
>
> Naturally, there is a temptation to use those numbers to control that organization. . . .The problem is this. The criteria we took measured the success of the organization. Or at least they did, before we . . . announced that henceforth those were the numbers by means of which we would measure the success of . . . those in the organization. Corporate executives would be paid bonuses for reaching certain profit figures, police departments that failed to close a reasonable number of cases would be subject to organizational shakeups, schools that did well on students' SAT scores would do well when teacher raises were granted.
>
> When we did this, we unleashed a complex set of social processes. . . . [We] caused all the players in the system to turn their attention to maximizing those numbers. Sometimes—often—they did so in ways that we did not anticipate, and that caused actions that "maximized the numbers" without maximizing what we really care about. In fact, the act of imposing the control system often destroyed what we really cared about.

The model in this paper helps to explain Darley's observation in economic rather than social psychological terms. Consider a school system that administers standardized tests to its students, but does not use the scores to motivate its teachers. It is likely that, in both cross sectional and time series data, the teachers whose students are better prepared for life will also, on average, receive higher scores on the standardized tests. But since teachers have no incentive to "teach to the test," they will not do so. Now consider a (seemingly reasonable) decision to use the test scores as a performance measure in an incentive system. Teachers now have incentives to engage in dysfunctional behavior that increases the performance measure without increasing the school's real objective. This behavior will clearly not be what the school system was intending, and could even result in the performance measure no longer being correlated with what the school system cares about.

Consider another example: using customer satisfaction as a performance measure in an employee bonus plan. While high levels of customer satisfaction may be a by-product of a well-run organization, and carefully designed studies may show that customer satisfaction is correlated with firm performance, it is nonetheless dangerous to develop an incentive program that pays for happy customers. There are simply too many ways to increase customer satisfaction without increasing profits. Should employees begin to be rewarded explicitly for customer satisfaction, they may learn to exploit these non-profit-maximizing ways to curry favor with customers, and drive the performance measure up while driving performance down. Customer satisfaction and firm performance will thus be highly correlated only so long as incentives are not placed on customer satisfaction.[9]

III: Incentive Plans Using Multiple Measures

While firm value may not be a contractible performance measure in private companies, non-profits, or government agencies, in publicly traded firms it is commonly used. In this section, I solve for the optimal weight on a distorted performance measure when the simultaneous use of incentives based on firm value is possible.

Recall that:

$V(\mathbf{a},\varepsilon) = \mathbf{f}\cdot\mathbf{a} + \varepsilon,$

and

$P(\mathbf{a},\varphi) = \mathbf{g}\cdot\mathbf{a} + \varphi.$

When using both V and P in an incentive contract, the covariance between ε and φ becomes relevant. It seems quite unrealistic to assume that ε and φ are uncorrelated: much of the noise in any one performance measure is also reflected in others, and is part of the noise in the stock price. In what follows, I assume that *all* of the noise in V is also in P: that is $\varepsilon=\varphi+\mu$, where φ and μ are uncorrelated.[10] This implies that:

$\sigma^2_\varphi \leq \sigma^2_\varepsilon,$

$\text{cov}(V,P) = \sigma^2_\varphi .$

Let pay equal:

$\text{Pay} = s + b_v V + b_p P.$

This specification of the reward scheme can be thought of as modeling the firm's decision about how steep a slope to place on firm value (for example, how large a share of the firm the employee should hold) while simultaneously determining the slope of his non-stock-based incentive contract. In this model the choice is a function of the marginal products (**f** and **g**), the risk in the stock price σ^2_ε, as well as the characteristics of the distorted performance measure: its risk (σ^2_φ) and distortion ($-\cos\theta$).

$$(3) \quad b_p^* = \frac{2hF(G\cos\theta\sigma_\varepsilon^2 - F\sigma_\phi^2)}{(G^2 + 2h\sigma_\phi^2)(F^2 + 2h\sigma_\varepsilon^2) - (FG\cos\theta + 2h\sigma_\phi^2)^2}$$

(4) $$b_v^* = \frac{F^2G^2(1-\cos^2\theta) + 2h\sigma_\phi^2 F(F - G\cos\theta))}{(G^2 + 2h\sigma_\phi^2)(F^2 + 2h\sigma_\varepsilon^2) - (FG\cos\theta + 2h\sigma_\phi^2)^2}$$

Several factors about these expressions are noteworthy. Note first that if **g** = **f** (that is, F=G and $\cos\theta = 1$) then $b_v^* = 0$. This result is sensitive to the assumption about the structure of the noise in the two performance measures: when **g** = **f**, V is a pure garbling of P, and so receives no weight. However, the intuition is valid even when V is not a pure garbling: when firms have excellent performance measures available to motivate employees, they can use less stock ownership as part of an incentive plan.

Note also that the sign of b_p^* depends on the relative riskiness of the stock price and the performance measure, and the distortion in P. If the ratio of the noise terms is greater than $\cos\theta$ times the ratio of the magnitudes of the marginal products, then b_p^* will be negative: it will be used to "back out" noise as part of a relative performance evaluation scheme. I return to this below.

Most of the insights about the combined use of these two performance measures can be determined by doing comparative statics on the ratio of these two slopes. I will refer to this ratio as the relative weight of b_p (or of b_v).

Define $R = \frac{b_p^*}{b_v^*}$.

For the case when $\cos\theta$ is positive[11]:

(R4) $\frac{dR}{d\cos\theta} > 0$: the relative weight on the distorted performance measure falls as its distortion increases. Equivalently, the relative weight on firm value increases as the distortion of the performance measure increases.

(R5) $\frac{dR}{d\sigma_\phi^2} > < 0$: the relative weight on the distorted performance measure is not monotonic in its noisiness. This is in part because the performance measure can be used as an element of a relative performance measurement scheme.

(R6) $\frac{dR}{d\sigma_\varepsilon^2} > 0$: the relative weight on the distorted performance measure increases as the noisiness of the stock price increases.

(Proofs are by inspection of equations (3) and (4).)

Results R4 and R6 are similar to results in the single performance measure case. R4 restates the intuitive argument given above: when a non-distortive performance measure exists, it will receive more weight and less weight will be placed on the (noisier) stock price. R6 shows that when the stock price has low noise, less weight will be placed on distorted performance measures. Result (5) is less obvious, and is discussed below.

A. Relative Performance Evaluation

When G or $\cos\theta$ are small, b_p^* is negative. In these situations, the model suggests that firms will use P to “back out” noise, in effect holding the employee accountable for her relative performance, and reducing the amount of risk borne by the employee while maintaining incentives. This model thus helps clarify the situations in which relative performance evaluation (RPE) will be used.

In order to explore the intuition on RPE, assume that $||\mathbf{g}||=0$. In this case, $P=\varphi$, and it is easily shown that $b_p^*=-b_v^*$. This case represents a standard one in the literature on the use of relative performance evaluation. φ might represent a "common shock," which can be filtered out of the performance measure in order to reduce the risk to the employee. (Tournaments and contests are an often-cited example of this sort of RPE.) However, as Gibbons and Murphy (1992) point out, "Paying workers based on relative performance . . . distorts the workers incentives whenever the worker can take actions that affect the average output of the reference group." The present model provides a new way to see this effect. So long as the employee cannot affect P, then the firm will exploit RPE as predicted in the standard model. However, if $||\mathbf{g}||$ is not zero—that is, if the employee can affect the reference group—b_p^* moves closer to zero. Thus, this model provides a precise way to determine whether to use relative performance measures to back out noise.

Consider the following simple model, motivated by Dye (1992) who argued that using industry-adjusted performance could distort managerial incentives. Let V be firm value, and P be industry performance. The question is whether the incentive contract should adjust firm performance for industry performance by setting $b_p^*<0$. Assume that the CEO can take actions that affect her industry benchmark. This could be by choosing the industry (through acquisition, for instance) against which she will be compared. Being in a good industry may help firm performance, but it also raises the benchmark against which the CEO will be evaluated.

In this case, firm performance and industry performance are:

$V = f_1a_1 + f_2a_2 + \varepsilon + \mu$,

$P = g_2a_2 + \mu$

where

f_1 is the effect of actions (and a_1 are actions) that affect the value of the firm directly ($f_1>0$),

f_2 is the effect of actions (and a_2 are actions) that affect the value of the firm through activities that the CEO undertakes to choose her industry ($f_2\geq 0$),

g_2 is the effect on the industry benchmark of managerial actions taken to choose her industry ($g_2\geq 0$),

ε represent firm-specific shocks,

μ represent industry shocks.

In this case, R, the ratio of b_p to b_v is:

$$R = \frac{Fg_2 2h\sigma_\varepsilon^2 \cos\theta - F^2 2h\sigma_\phi^2}{F^2 g_2^2(1-\cos^2\theta) + F2h\sigma_\phi^2(F - g_2\cos\theta)},$$

where

$$F = \sqrt{f_1^2 + f_2^2}\,.$$

Consider first the case in which $g_2=0$. In this case, the CEO can do nothing to affect the benchmark and R=-1; that is $b_p^* = -b_v^*$ and the firm fully adjusts for industry performance in the CEO's incentive contract. However, as g_2 becomes greater than 0, R becomes greater than –1: the amount of RPE used is less. The firm avoids full RPE if the manager can take actions that affect the reference group. If g_2 becomes too large—that is, if the manger has too much scope to choose her industry benchmark—the firm avoids RPE altogether.

IV: The Trade-Off Between Distortion and Risk

It is a contention of this paper that the trade-off between distortion and risk modeled here is at the core of the problem of incentive design in many organizations. Viewed in this way, the objective of incentive system design is to discover or create low distortion, low risk performance measures. In what follows, I discuss several examples of incentive plan design problems, showing how the choice of performance measures can be usefully viewed as a trade-off between distortion and risk.[12]

A. Timing of Measurement

Decisions about performance measurement often revolve around issues of timing: should employees be evaluated on short run or long run results? Two examples help to illustrate how this choice involves trading off distortion and risk.

1. Incentives for bank loan officers and project managers

The typical incentive plan for loan officers in a bank involves "origination fees," in which the loan officer is paid for lending money. A feature of this type of incentive is that it gives the loan officer no incentive to search for and write "good" loans—that is high interest rate loans that are likely to be repaid. Instead, loan officers have incentives to make any loan, and banks typically have credit committees (made up of higher-level bank officers) whose job it is to determine the credit-worthiness of the potential debtor, and approve or deny the loan.[13]

The question in this scheme is why loan officers are not paid on the eventual profitability of the loan, rather than on its origination.[14] Bonuses based on loan profitability would have the advantage of giving loan officers incentives to search out good credit risks, and sell loans with higher expected value. In other words, such a performance measure would provide less distorted incentives to the loan officers. However, such a scheme would also give the loan officers greater risk, since many things can happen to debtors that are essentially unknowable when a loan is written. In this case, it appears, the trade-off between risk and distortion is made in favor of lower risk and higher distortion.

The opposite choice is often made in the design of bonus plans for project managers in large construction projects. Construction managers often leave one project and move to a second before the first project is completed. Frequently, the project manager will be paid a bonus based the final profitability of the project when it is completed. Thus the project manager might have two or three "contingent" unpaid bonuses to his credit, whose payment awaits the completion of a project that he worked on months or even years earlier.

The problem with such a bonus plan is clear: the project manager's bonus for a particular project depends on many factors over which the project manager has no control, not the least of which is the performance of his successor(s). Yet the benefits are also clear: such a scheme gives the manager incentives to be concerned about the *long run* profitability of his decisions. Trying to evaluate the project manager on the profitability of the project when he leaves would encourage him to (perhaps literally) bury problems that would not become clear until long after he had left the project. In this case, the benefits of low incentive distortion outweigh the costs of high risk for the project manager.

In both of these examples, the choice of performance measures involves trading off risk and distortion. In both cases, the choice is between a higher risk, lower distortion performance measure (loan performance, final profitability) versus a lower risk, higher distortion measure (loan origination, short term profitability). Which measure is chosen depends on the relative costs of distortion and risk.

B. Level of Aggregation

Compensation design problems frequently involve choosing the level of aggregation at which to measure performance. I examine two examples, one involving the choice of the group size over which to measure performance, and the other involving responsibility accounting.

1. Group Size

A key decision in determining an employee's incentive package is the weight to place on individual versus group performance, and if group performance, how large a group. Consider the design of a performance measurement system for a worker who is a member of a work group. Each worker engages in tasks that affect his own measured performance, as well as engaging in cooperative activities that improve the performance of the entire group. Attempts to reward the worker for individual performance may thus reduce teamwork and destroy cooperation. On the other hand, rewarding individuals on the basis of group performance makes their rewards depend on the performance of the entire group, including all of the uncontrollable events (and actions of others in the group) that affect group output.

Once again, the choice of whether to use group or individual performance in the incentive contract depends on the trade-off between risk and distortion. Group rewards subject group members to risk by making their rewards depend on uncontrollable events; individual rewards distort incentives to cooperate. How this trade-off gets resolved depends mainly on the value of cooperation (and thus the distortion induced by an individual reward scheme) and the riskiness of group (relative to individual) output.

2. Responsibility Accounting

The tools used to examine rewards based on individual versus group performance can be taken to a higher level, and used to explore a whole range of possible systems for evaluating managers in large organizations. This question of what to hold managers accountable for is sometimes referred to as responsibility accounting.

Consider the design of an incentive contract for a divisional plant manager in a multi-divisional firm. The plant manager works for a division manager, who works for the CEO of the company. The division has multiple plants, and there is a divisional sales organization that reports to the division manager.

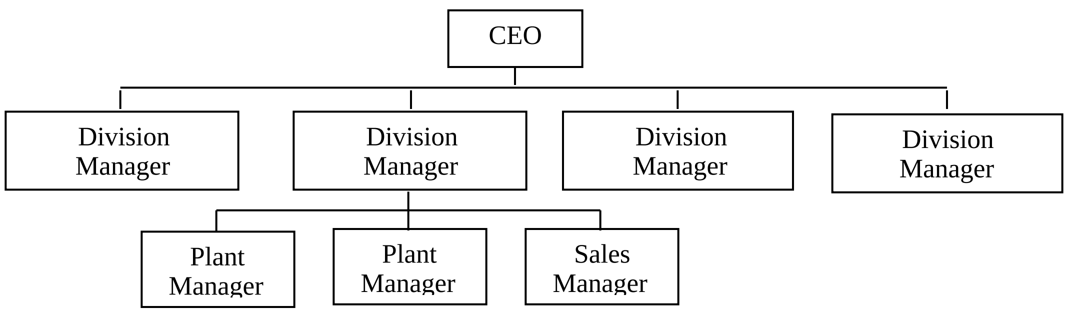

Costs and revenues are measured at the plant, division, and company levels. Consider the following five possible performance measures and associated incentive systems.

Performance Measure	**Incentive System**
Firm value	Employee stock ownership
Firm-wide accounting profits	Firm-wide profit sharing
Divisional profits	Divisional profit sharing
Plant-level profits	Plant profit center
Plant-level costs	Plant cost center

Each of these performance measures (or a combination of them) represents a possible basis for the plant manager's incentive contract. Moving down this list represents a reduction in the amount of risk that would be borne by the manager (an increase in controllability), but simultaneously an increase in the distortion associated with the performance measure. Consider a switch from divisional profits to plant-level profits. Clearly, holding him accountable for divisional profits subjects him to risks (from events affecting the other plant) that he would not bear if he were only accountable for his own plant’s profits. However, if he can take actions that might affect (positively or negatively) the profitability of other plants, holding him accountable only for his own plant’s performance eliminates any incentive to consider the effects of such actions. Figure 3 illustrates how noise and distortion differ across these five performance measures.

A similar argument holds for the switch from plant-level profits to plant-level costs. To the extent the manager has any influence over the quality of the product produced by the plant, holding him accountable for plant costs rather than plant profits will tend to distort his incentives with respect to a cost-quality trade-off. On the other hand, holding him accountable for plant profits will make him bear the risk of sales variation over which he may have little control.

The choice of which performance measure to use (or the weights to place on them, if they are used in combination) depends on how the amount of distortion and the amount of risk change as one moves from one performance measure to another. In the case of a diversified conglomerate trying to decide whether to use firm-wide profits or divisional profits in an incentive plan for division managers, the amount of inter-divisional cooperation or synergy will be critical. If there is no scope for such cooperation, then the trade-off between distortion and risk will be simple: little will be lost by holding the division manager accountable only for divisional profits, since the amount of distortion in incentives will be small. On the other hand, when the scope for cooperative activities across units is large, the trade-off will tend to favor a riskier, but less distorted performance measure.

V. Extensions and Conclusions

This paper provides a simple and intuitive structure for understanding the choices organizations face in the design of incentive contracts. I argue that a performance measure's usefulness in an incentive contract will depend on its distortion and risk: the more distorted and the riskier the measure, the less valuable it will be to the organization and the less it will be used in an incentive contract. Furthermore, organizations rarely have available low risk, low distortion measures, and so are generally making trade-offs between measures that are high risk and low distortion, or low risk and high distortion. As discussed above, many problems in incentive system design can be fruitfully analyzed using this framework.

Of course, much work remains to be done. Questions raised by this analysis include:

- What is the underlying structure (of information, incentives, etc.) that requires organizations to choose between distortion and risk in performance measures? Is there some sort of “performance measurement possibility frontier? If so, what determines its efficiency? How do organizations choose where to locate on this frontier?

- How do distorted, risky measures combine into “portfolios?” What are the characteristics of linear combinations of performance measures? How should firms combine them?

Answers to these questions are important for developing a fuller understanding of the forces that drive the use of different performance measures in incentive contracts.

Other extensions of this framework permit analysis of some specific incentive problems that organizations face. Consider first the problem of designing incentive contracts to encourage innovation in technology-based firms. Large firms often struggle to deliver incentives to scientists and engineers involved in research and development. On what objective basis can such contracts be based? At the root of the difficulty in designing such an incentive contract is the problem that, in general, the desired outcomes cannot be known in advance, and the value of any given breakthrough is extremely hard to predict. The value of the breakthrough to the firm may not be known for many years, or perhaps may never be distinguishable from other causes of firm success or failure. In this context, how can research scientists be rewarded?

Firms (and economies) have several solutions to this performance measurement problem, none of them ideal. One is to pay research scientists flat wages, with modest rewards (in the form of career advancement and prestige) for good work as determined by subjective evaluations and peer recognition. While this solution is quite common it relies, to a large extent, on the intrinsic motivation of scientists to do interesting and personally rewarding research, and often results in weak incentives to invent profitable products. A second solution is to reward researchers with significant stock-based compensation, so that they will share in their value creation to the extent that it increases the value of the firm. The efficacy of this second solution, of course, is highly dependent on the size and diversity of the firm. The larger and more diverse the firm, the lower will be the signal-to-noise ratio of the stock price with respect to the scientist's actions. In very large firms, this type of reward is likely to have little effect on the scientist’s behavior, since the optimal weight on such a noisy performance measure is quite small.

One other solution is to have the R&D done outside the firm, in small companies whose only activity is R&D. In these small companies, the total value of the firm (the current stock price, or the future price in an IPO or buyout) will be quite sensitive to the actions of the research staff, making it a more powerful incentive instrument than an equity stake in a large firm. Such firms are common in technology-intensive industries, and the high incentive strengths made possible by their small size is often cited as an important reason for their success in generating innovation. Large firms, whose only choice is to rely on distorted performance measures or very noisy stock prices, cannot replicate these small firm incentives.

More generally, it is evident from this analysis that larger firms will face more difficult incentive problems than will small firms because, for any given employee, the optimal amount of stock ownership is lower for larger firms, and the reliance on distorted performance measures is likely to be greater. But what of organizations with no stock price at all, and no prospect of ever selling out? This must lead to even more challenging incentive design problems.

The problems involved in designing incentive plans for organizations without well-defined residual claimants—non-profit companies, government agencies, state-run service providers—are very difficult. However, I argue that these problems do not have their origins in several well-worn explanations about the difficulties with non-profit management. First, these problems do not arise from a lack of *available* performance measures, but from a lack of *undistorted* performance measures. Almost every performance measure available to for-profit firms (with the exception of a stock price) is available to these organizations. Non-profits can measure their assets, revenues, costs, "profits," and just about any other financial or non-financial measure just as easily as a for-profit firm can. And these performance measures are no riskier for non-profits than they are for for-profit organizations. The problem is that such performance measures are likely to be even more distorted for non-profits than they are for for-profit firms.

The problem with performance measurement in non-profits is also not due to any *inherent* difficulty in the measurement of the value of what non-profits produce or consume. Many for-profit firms produce highly intangible goods and services (think about sports cars, fashion items, or legal services) whose intrinsic value is very difficult to measure. The difference between the for-profits (for example, law firms) and the non-profits (for example, school systems) is that the non-profits rarely sells all their outputs on markets (they often give away or subsidize their services), and often acquire key inputs from non-market sources (that is, they hire volunteers, or receive tax funding or charitable donations). What this means is that the decentralized value assessments that are generally performed by market participants—customers deciding what a product is worth, suppliers deciding what an input is worth—are disabled by the fact that non-profits engage in only limited market transactions. This means that their revenues may not accurately measure the value of what they produce, and their costs may not accurately measure the value of what they consume. Thus, the available performance measures are distorted with respect to the true value that the organization creates, in a way that they are not for for-profit enterprises.

This analysis suggests some important insights about the design of incentive programs in non-profit organizations. First, when these organizations do engage in market-based activities, they should be evaluated on the basis of traditional (for profit) performance measures. An example of this arose in the early 1980s within Harvard University. The University decided to consolidate all of its non-academic property and place it in a separate organizational unit—Harvard Real Estate (HRE)—that would manage this property more like a for-profit landlord, charging market rental rates, properly maintaining the properties to ensure that they could charge these rates, and generally managing the properties to maximize the return to the University from these assets. All of the commercial and residential property of the University was thus collected together and managed in this way. Issues of housing subsidies (for graduate students or recruited faculty) were handled by the unit in the university that wanted to grant the subsidy (the financial aid office for students, the recruiting department for faculty). As a result of this reorganization, HRE was given a clear and reasonably undistorted performance measure, and was thus able to provide incentives to its employees to manage Harvard's non-academic property in a value maximizing way.

A second insight about incentive provision in non-profits arises from the fact that, because these organizations so often lack undistorted performance measures, they are generally characterized by (optimally) weak incentives. For this reason, employee's intrinsic motivations and personal reasons for working at these organizations loom large in their choices about what actions they will take. Indeed, because these firms cannot use strong explicit incentives, they will tend to select employees who have strong preferences for the type of work in which the organization engages. This has implications for how incentive designers should use the limited incentive instruments that they have. An example of this arises in the performance standards system implemented at the JTPA (discussed earlier). The JTPA's objective is to improve the training and employability of the disadvantaged. In a 1996 paper, Heckman, Smith and Taber cite evidence that less disadvantaged clients may benefit more from the offered job training than do the most disadvantaged. But the employees who work at the centers have preferences to help the most disadvantaged. Thus, they will have a tendency to want to serve those who may get less out of the program. Heckman et al. argue that the performance standards, which encourage cream skimming, may thus actually increase the value added of the program by giving JTPA center employees an incentive to attract clients who will benefit more. Thus a distortionary performance measurement system is combined with a system that allows employees to (inefficiently) indulge their preferences to generate more valuable outcomes.

I conclude with one final observation. What can organizations do when forced to choose among performance measures that are very distorted and noisy? They avoid exclusive reliance on objective performance measurement, and turn to subjective assessments (Baker, Gibbons, Murphy 1994; Predergast and Topel 1996). Subjectivity allows supervisors to examine, *ex post*, the actions that an employee took and decide whether it represented dysfunctional behavior or gaming. It also allows supervisors to subjectively "back out" noise in an objective performance measure, reducing the riskiness of a measure without adding distortion. As I and others have argued, subjectivity comes with its own set of problems. Contracts based on subjective assessments are non-enforceable, and thus rely on the trustworthiness and reputation of the evaluator. Subjective assessments require a large amount of evaluator time and effort. In addition, they create the opportunity for supervisors to engage in favoritism. However, the prevalence of subjectivity in the performance measurement systems of virtually all organizations suggests that exclusive reliance on distorted and risky objective measures is not an efficient alternative.

References

Baker, George. 1992. "Incentive Contracts and Performance Measurement," Journal of Political Economy 100(3):598-614.

______, Robert Gibbons, and Kevin J. Murphy. 1989. "Subjective Performance Measures in Optimal Incentive Contracts." Quarterly Journal of Economics 109(4):1125-1156.

Banker, Rajiv, and Srikant Datar. 1989. "Sensitivity, Precision, and Linear Aggregation of Signals for Performance Evaluation." Journal of Accounting Research 27(1):21-39.

Courty, Pascal, and Gerald Marschke. 1997. "Measuring Governmental Performance: Lessons from a Federal Job Training Program." American Economic Review 87(2):383-388

Darley, John. 1991. "Setting Standards Seeks Control, Risks Distortion." Institute of Government Studies Public Affairs Report. University of California at Berkeley 32(4).

Datar, Srikant, Susan Kulp, and Richard Lambert. "Balancing Performance Measures." Journal of Accounting Research, forthcoming.

Dye, Ronald. 1992. "Relative Performance Evaluation and Project Selection." Journal of Accounting Research 30:27-52.

Fama, Eugene, and Michael Jensen. 1983 "Separation of Ownership and Control." Journal of Law & Economics 26.

Fast, Norman, and Norman Berg. 1975. "The Lincoln Electric Company." Harvard Business School Case 376-028.

Feltham, Gerald, and Jim Xie. 1994."Performance Measure Congruity and Diversity in Multi-Task Principal/Agent Relations." The Accounting Review 69(3):429-453.

Gibbons, Robert. 1998. "Incentives and Careers in Organizations." Journal of Economic Perspectives 12(4): 115-132.

_______, and Kevin J. Murphy. 1990. "Relative Performance Evaluation for Chief Executive Officers." Industrial and Labor Relations Review 43(3):305-515.

Heckman, James, Carolyn Heinrich, and Jeffrey Smith. 1997. “Assessing the Performance of Performance Standards in Public Bureaucracies.” American Economic Review 87(2):389-395.

Heckman, James, Jeffrey Smith, and Christopher Taber. 1996. "What Do Bureaucrats Do? The Effects of Performance Standards and Bureaucratic Preferences on Acceptance Into the JTPA Program," in Studies in Bureaucratic Behavior, ed. G. Liebcap JAI Press, 1996.

Holmstrom, Bengt, and Paul Milgrom. 1991.“Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design.” Journal of Law, Economics, and Organization 7:24-52

Ittner, Chris, David Larcker, and Marshall Meyer. 1997. “Performance, Compensation and the Balanced Scorecard.” Working Paper, The Wharton School.

Kerr, Steven. 1975“On the Folly of Rewarding for A while Hoping for B.” Academy of Management Journal 18(4):769-783.

Lazear, Edward. 1995. Personnel Economics. MIT Press.

Prendergast, Canice. 2000. "What Tradeoff of Risk and Incentives?" American Economic Review 90(2):421-425.

___________, “The Tenuous Trade-Off Between Risk and Incentives,” Journal of Political Economy, forthcoming.

___________ ,and Robert Topel. 1996. “Favoritism in Organizations.” Journal of Political Economy 104:958-978.

van Praag, Mirjaam, and Kees Cools. 2001. Performance Measure Selection: Aligning the Principal’s Objective and the Agent’s Effort.” Working Paper. University of Amsterdam.

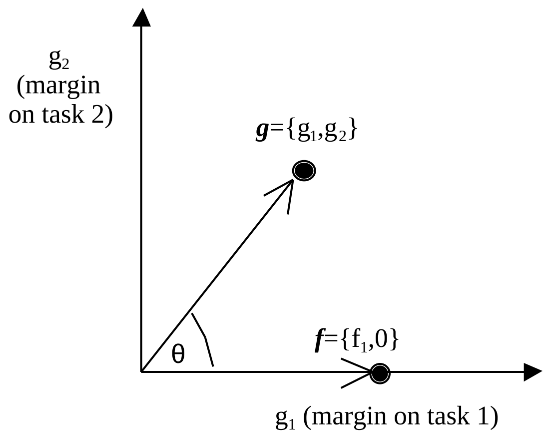

Figure 1

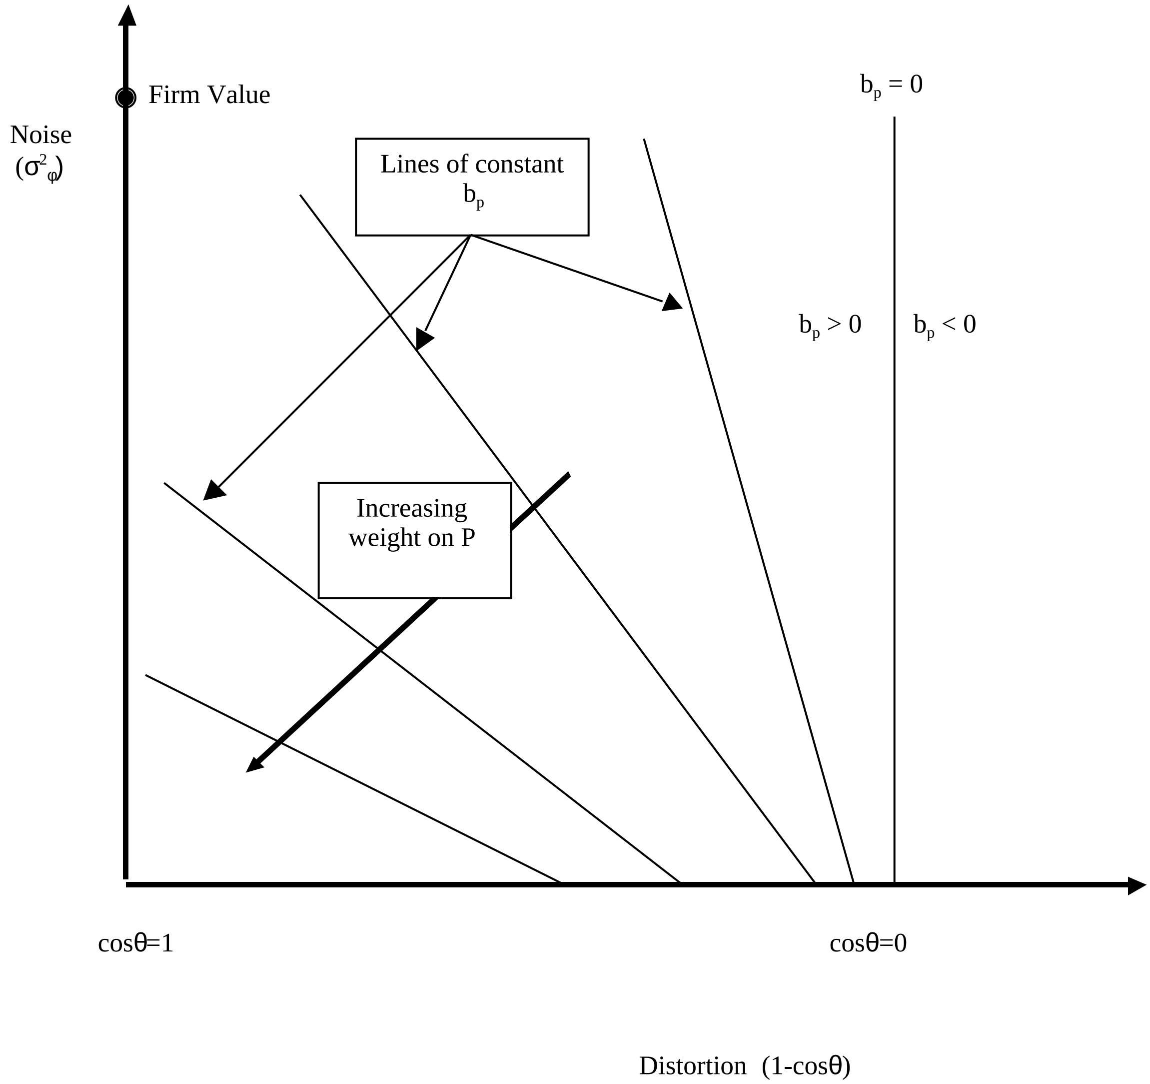
Firm Value
Noise
(σ^2_φ)
$b_p = 0$
Lines of constant b_p
$b_p > 0$
$b_p < 0$
Increasing weight on P
cosθ=1
cosθ=0
Distortion (1-cosθ)

Figure 2

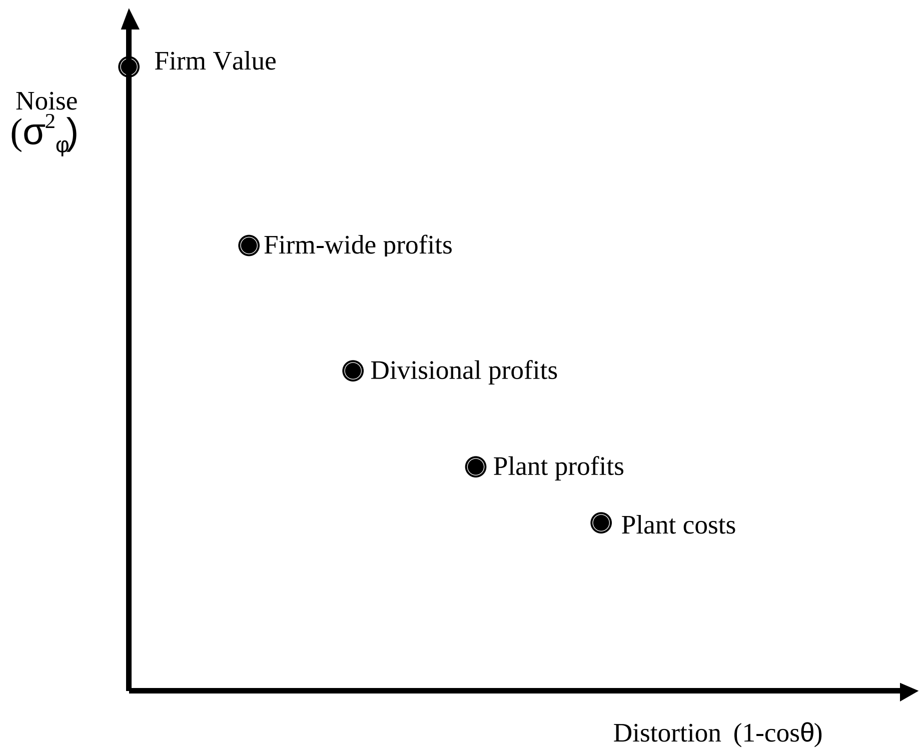

Figure 3
Trading off Risk and Distortion

---

1. To the extent that managers can manipulate the stock price of a publicly traded company through information disclosure, this assumption is unwarranted. Using the "long run" value of equity (that is, its value some time in the future when the effects of current actions will be known to the market) solves this problem to a large degree. In fact, firms use restricted stock grants and vesting schedules largely for this reason.

2. I assume that s is set so that the employee receives a compensation package whose expected value equals his reservation utility. The problem is thus to choose b_v to maximize surplus (V- h*var(pay) - distutility of effort) subject to the constraint that the agent chooses a’s to maximize his utility.

3. Feltham and Xie use the term "congruent" to capture the same concept, as do Datar, Kulp, and Lambert.

4. Note that, if P was noiseless, then the effect of the change in G on b_p^* would be unambiguous: only the scaling effect would operate, and b_p^* would fall.

5. Note also that if cosΘ is negative, then b_p^* is negative. In this case, actions that increase P reduce V, and the optimal contract punishes the employee for delivering P.

6. The strict independence of b* and the correlation between V and P in this model is driven by the fact that the marginal products are deterministic. If the margins depended on the state of the world (unknown at the time of contracting, but observed by the employee before acting) then b* would depend, to some extent on the correlation between V and P. However, the intuition that what matters is the correlation of the margins, rather than the correlations of the outcomes themselves, would still be valid. See Baker (1992).

7. The opposite is also possible: it may be optimal to place positive weight on a performance measure that is negatively correlated with value. Consider the following personal example. The neatness of my office is negatively correlated with my level of activity and productivity: when I am busy, my office is a mess, when I am not, my office is neat. But having a neat office clearly increases my productivity. Thus, a savvy incentive contract designer might give me incentives to keep a neat office, even though the performance measure is negatively correlated with productivity.

8. Prendergast (forthcoming) also provides a model that predicts a positive relationship between risk and incentive strength.

9. Ittner, Larcker and Meyer (1997) provide an example of this phenomenon at a large bank from which they got detailed data on performance and incentive programs. A system put in place to provide incentives for

branch managers to increase customer satisfaction quickly began to reward some of the most unprofitable branches.

10. Alternatively, I could write $\varphi=\varepsilon-\mu$, where φ and μ are uncorrelated, and $cov(\varepsilon,\mu)=\sigma^2_\mu$.

11. When $\cos\theta$ is negative, many of the comparative statics on R are hard to sign and hard to interpret. I do not examine this case here.

12. A recent paper by van Praag and Cools (2001) also examines these two aspects of performance measurement, and finds some empirical support for the proposition that these two factors are traded off in the design of incentive plans.

13. This separation of decision rights into “initiation” and “ratification” rights is labeled by Fama and Jensen (1983) a decision hierarchy, and is more colloquially called bureaucracy or “red tape.” It may foster specialization, in that the loan officer need not be expert in determining credit worthiness.

14. For long-term loans (say 15-20 years), this might be problematic because of turnover, retirement, etc. But for many shorter-term loans, these problems would not seem to be overwhelming.