A Theory of "Yes Men"

By CANICE PRENDERGAST*

This paper illustrates an incentive for workers to conform to the opinion of their supervisors when firms use subjective performance evaluation. This desire to conform arises endogenously from the firm's need to induce the worker to exert effort. I show that firms may optimally eschew the use of incentive contracts to retain workers' incentives for honesty. I illustrate that the incentive to conform implies inefficiencies, even when workers are risk-neutral, and is likely to lead to more centralized decision-making than in the absence of the desire to conform. (*JEL* D20, J30)

Hamlet: *Do you see yonder cloud that's almost in the shape of a camel?*
Polonius: *By the mass, and 'tis like a camel, indeed.*
Hamlet: *Methinks it is like a weasel.*
Polonius: *It is backed like a weasel.*
Hamlet: *Or like a whale?*
Polonius: *Very like a whale.*
—William Shakespeare, *Hamlet*, Act 3, Scene 2

The profitability of most organizations depends on how cheaply and efficiently they can collect information. For example, managers routinely make decisions on whether to undertake projects, purchase new machinery, hire new workers, or evaluate a worker's performance or promotion prospects. The quality of the manager's information clearly determines the efficiency of any of these decisions. This paper analyzes incentive issues associated with the collection of information and argues that subordinates have an incentive to conform to the opinions of their superiors.

Assume that a manager assigns a subordinate, whom I call "the worker," to collect information on the value of some parameter. For example, the worker's role may be to determine the profitability of a new project or the performance of the worker's subordinates. Assume further that the collection of information involves some effort on the part of the worker, which he would like to avoid. More effort by the worker is assumed to improve the accuracy of the worker's observation. How can the worker be provided with incentives to collect information?

One possibility is to obtain more objective information on the true parameter before rewarding the worker. For example, a manger could wait to see if the project is profitable before deciding how well to reward the worker. This has three typical problems. First, for some types of activities, there may be little more information revealed. For example, if the worker's role is monitoring the performance of a subordinate, obtaining an objective measure of the worker's input to the subordinate's performance may be difficult. Second, there may be a delay between the worker's observation on the parameter and more information becoming available. Hence, in the absence of commitment power (and good memory), contracts that explicitly depend on future information may be difficult to enforce. Finally, whatever information arrives may be tainted with noise which was unpredictable at the time the worker's report was made. For example, the profitability of a project may depend on competitors' activities, technological change, the state of the economy, weather, and so forth, in a way that cannot

*University of Chicago, Graduate School of Business, 1101 E. 58th St., Chicago, IL 60637. Many thanks to Eddie Lazear, Raghu Rajan, Lars Stole, and an anonymous referee for helpful comments. All errors are my own.

be predicted when the initial inference is made. Hence the value of additional information may be limited.

Since the objective of the firm is to collect information, there is no obvious metric by which to measure whether a worker's report is correct. In this sense, the model is very different from standard principal–agency models where, for example, higher output is better. Here a high valuation of the parameter in question may indicate shirking. In the absence of reliable objective measures of performance, firms generally resort to subjective evaluation procedures, where managers compare their own findings to those of their subordinates. When the point of reference for adequate performance is the manager's opinion, an endogenous desire arises for the worker to conform to the opinion of the manager. This arises because the only way to induce the worker to exert effort is by comparing the findings of the manager with those of the worker. Hence subjective performance evaluation gives rise to the existence of "yes men," who attempt to second-guess the opinions of their monitor and mimic them.

In Section I, I describe a model in which a manager and a worker obtain an observation on the value of a normally distributed parameter, where the firm's profits are decreasing in the variance of the estimate of the parameter. Hence the quality of the information matters. Other than a (common) initial prior on the parameter and the observations of the manager and worker, there is no *useful* information available. I begin by showing that the firm can use a relative performance evaluation scheme to induce the worker to exert effort, where the worker's information is compared to the manager's opinion of the true value of the parameter. (If the worker's wage is uncorrelated with the manager's opinion, he has no incentive to exert effort.) Hence subjective performance evaluation requires that the worker benefits from guessing the opinion of the manager.

This does not imply that decision-making is inefficient unless the worker has information on what the manager observes. I assume that, in addition to his own observation, the worker receives a private noisy signal of the manager's observation. (For example, the manager may emphasize "quantity" over "quality," so that the worker may slant his opinion on a project's profitability to ignore quality aspects; or a manager may be known to be predisposed to give poor evaluations to certain employees being supervised by the worker, which the worker can reinforce by his recommendation.) When the worker observes a noisy signal of the manager's opinion, I show that the worker biases his report toward what he believes the manager wants to hear, *but the manager cannot infer the worker's information*. This reduces the value of the worker's information. One of the principal results of the paper is to illustrate that the uncertainty about the true observation of the worker (caused by this desire to conform) implies that information is not aggregated as efficiently as when there is no incentive to conform. Thus the existence of yes men leads to inefficiencies; yet the existence of yes men is a necessary implication of providing incentives to exert effort. Hence this paper highlights a trade-off between inducing workers to exert effort and encouraging them to be honest in reporting their findings.

So far, I have assumed that it is optimal to induce the worker to exert effort. It has often been argued that offering incentive contracts can be harmful to organizations, as it can imply that workers get involved in only those activities for which they are directly rewarded (see Bengt Holmstrom and Paul R. Milgrom, 1991). I show that offering incentive contracts may be dominated by offering a wage that is independent of the worker's actions for a somewhat different reason. With an incentive contract, workers do indeed exert effort, but it also gives them an incentive to be dishonest. By offering wages that are independent of perceived actions, workers shirk as in a standard principal–agency model, but at least it gives them an incentive to report their information honestly. I show that if the desire to conform is large enough, the firm does not offer incentive contracts, in order to retain incentives for honesty. This paper therefore

provides another reason why employees may be allowed considerable discretion over their behavior instead of using explicit performance measures.

Following the basic results, I consider how the desire to conform affects how decisions are made in organizations. For example, does the existence of yes men imply that workers have more or less input about whether a project is undertaken? I show that because the worker distorts his opinion, the value of the worker's information falls so that, in organizations with subjective performance evaluation, decisions are likely to be more centralized (in the sense that more weight is placed on the opinion of the manager in any decision) than with objective performance measures.

The basic model describes a process in which a worker reports his private information to a manager; the worker's incentives are to distort his report in the direction of what he believes that the manager wants to hear. In Section II, I extend this analysis to describe how the desire for conformity may more generally pervade communication within organizations. Assume that a group of individuals with private information combines to determine the true value of a parameter. As in Section I, the provision of incentives requires relative performance evaluation. However, if relative performance evaluation is used, I show that communication between individuals gives rise to a desire for conformity in exactly the same way as described above, such that workers report not only on the basis of their own observations, but also on what they believe the opinions of others to be. Hence interaction stifles "creativity," as all workers converge inefficiently to a similar conclusion. One implication of this is that organizations may insert "Chinese Walls" between departments, so that departments are excluded from knowledge of the behavior of other departments, to eliminate this conformity problem.

This paper has much in common with the principal–agency literature, such as Holmstrom (1979) and Holmstrom and Milgrom (1991); the literature on influence activities, such as Milgrom and D. John Roberts (1988); and the recent literature on herding and conformity, such as David Scharfstein and Jeremy Stein (1990), B. Douglas Bernheim (1991), and Abhijit V. Banerjee (1992). Most similar in spirit is the work of Bernheim, who illustrates that conformity can arise when agents wish to be well perceived by others. Here "popularity" is endogenously derived by the relative-performance nature of the incentive contract.

I. The Model

Assume that there are two individuals in a firm, a manager (m) and a worker (w). The function of both agents is to observe the true value of a parameter, η. The manager and the worker get imperfect estimates of η. In particular, the manager observes

$$\eta_m = \eta + \varepsilon_m \tag{1}$$

where ε_m is a normally distributed error with mean 0 and variance σ_m^2. The worker observes

$$\eta_w = \eta + \varepsilon_w \tag{2}$$

where ε_w is a normally distributed with mean 0 and variance σ_w^2. The errors are uncorrelated. Both agents have a common prior that η is normally distributed with mean η_0 and variance σ_0^2.

Throughout the paper I wish to emphasize the importance of knowing the true value of the random variable. As a result, I assume that profits of the manager are normalized to the negative of the variance of the estimate of η, where the estimate of the manager is made after observing his own report, the initial prior and a report by the worker on the parameter. The manager also acts as principal, designing the incentive contract and paying the worker.

There are two natural interpretations of η. First, η could refer to the profitability of a potential project which is only undertaken if its true value exceeds some critical level η^*. Then the profits of the manager depend on the variance of η because for any η^*, as $\mathrm{Var}(\hat{\eta})$ rises (where $\hat{\eta}$ is the estimate of η), the project is incorrectly undertaken (and

eschewed) more often.[1] An alternative interpretation of η is the effort level of a subordinate to the worker, whom the worker monitors. High values of η imply that the subordinate exerted effort. Here the firm's profits are likely to be tied to the variance of the estimate of η, as the variance of η determines the efficiency of the monitoring mechanism. High levels of $\text{Var}(\hat{\eta})$ imply that shirking by the subordinate is observed less often than with low $\text{Var}(\hat{\eta})$. Hence ceteris paribus the manager would like to reduce $\text{Var}(\hat{\eta})$ to increase the subordinate's incentives.

Two further ingredients are necessary for my results. First, I assume that the acquisition of information is costly to the worker and the manager. I assume that the variance of σ_i^2 is given by $h_i(e_i)$, where e_i is the effort level of agent i, $i = \text{m,w}$. I assume $h_i'(e_i) < 0$, $h_i''(e_i) > 0$. Hence, by exerting effort, the agents get better information on the variable of interest. I further assume that the cost of exerting effort is given by $C_i(e_i)$, where $C_i'(e_i) > 0$, $C_i''(e_i) > 0$, $C_i(0) = 0$, $C_i'(0) = 0$, and $C_i'(\infty) = \infty$, for $i = \text{m,w}$. Neither employee's effort can be observed.

The second ingredient required is that the worker gets an observation on what the manager has seen. I assume that, in addition to η_{w}, the worker also observes η_λ, with

$$\eta_\lambda = \eta_{\text{m}} + \lambda \tag{3}$$

where λ is a normally distributed error with mean 0 and variance σ_λ^2. All variances are assumed to be common knowledge.

The assumption that the worker has information on the manager's opinion does not seem unreasonable. For example, the worker may know that the manager is predisposed to liking certain types of projects; or that he has a dislike for certain employees. Alternatively, the manager may only be able to carry out a cursory examination of the evidence of η, so that the worker has information on what the manager's first impression is likely to be. Hence, I assume that workers have information on the manager's opinion, but that observation is unknown to the manager.

Two points are worth noting here. First, the information that the worker has on η_λ has no value for determining η in the sense that η_{m} is statistically sufficient for η_λ (η_λ can be generated by η_{m} plus an auxiliary randomization). This implies that, if information on η_{m} is available, η_λ should be ignored. As will be seen, however, this is not the case in equilibrium if the worker is offered incentives to exert effort. Second, no effort is required by the worker to obtain information on the manager's opinion concerning the project. Introducing this effect implies that the manager must design incentives to induce the worker not to exert effort on these unproductive activities.[2]

I assume that the worker and the manager are risk-neutral. The worker has reservation utility of r, and the manager is assumed to be residual claimant.[3] I assume risk-neutrality to illustrate that the efficiency problems which arise here are not due to costs associated with workers facing risk to induce them to exert effort. This allows me to concentrate on the desire for conformity.

As mentioned in the introduction, the manager could potentially induce the worker to exert effort by waiting for further objective measures on the true value of the parameter, η. However, for the reasons stressed above, this may not be possible. As a result, I assume that no further information becomes available on the true value of

[1] It is also true that the profitability of the firm depends on the level of η, but this does not affect any of the marginal decisions here, as I assume that the true value of η is independent of actions.

[2] My conjecture on this is that introducing this incentive will reduce the sensitivity of incentives to the worker's report in equilibrium as the manager seeks to reduce the incentive for the worker to get involved in these rent-seeking activities. The reason for this is similar to that proposed by Milgrom and Roberts (1988), to reduce unproductive rent-seeking.

[3] Note that I do not allow sale of the company to the worker. If this is allowed, then a similar problem potentially arises about providing the manager with incentives.

η. Instead, the worker must be compensated on three parameters, η_0, η_m, and $\hat{\eta}_w$, where $\hat{\eta}_w$ is the report made by the worker on the value of η.

A. *The First Best*

In order to determine the first best, assume that the worker truthfully reports his observation. Then the manager computes the posterior distribution of η, which is normally distributed with mean

$$(4)\quad \hat{\eta} = \frac{\sigma_0^2\sigma_m^2\eta_w + \sigma_w^2\sigma_m^2\eta_0 + \sigma_0^2\sigma_w^2\eta_m}{\sigma_0^2\sigma_m^2 + \sigma_m^2\sigma_w^2 + \sigma_0^2\sigma_w^2}$$

with variance

$$(5)\quad \mathrm{Var}(\hat{\eta}) = \frac{\sigma_0^2\sigma_m^2\sigma_w^2}{\sigma_0^2\sigma_m^2 + \sigma_m^2\sigma_w^2 + \sigma_w^2\sigma_0^2}.$$

Then,

$$(6)\quad \frac{\partial \mathrm{Var}(\hat{\eta})}{\partial e_i} = h_i'(e_i)\frac{\left[\sigma_0^2\sigma_j^2\right]^2}{\left[\sigma_0^2\sigma_m^2 + \sigma_m^2\sigma_w^2 + \sigma_w^2\sigma_0^2\right]^2} \qquad i \neq j.$$

The first-best level of effort is characterized by

$$(7)\quad C_w'(e_w^*) = -h_w'(e_w^*)\frac{\left[\sigma_0^2\sigma_m^2\right]^2}{\left[\sigma_0^2\sigma_m^2 + \sigma_m^2\sigma_w^2 + \sigma_w^2\sigma_0^2\right]^2}$$

and

$$(8)\quad C_m'(e_m^*) = -h_m'(e_m^*)\frac{\left[\sigma_0^2\sigma_w^2\right]^2}{\left[\sigma_0^2\sigma_m^2 + \sigma_m^2\sigma_w^2 + \sigma_w^2\sigma_0^2\right]^2}.$$

There is a unique solution to this pair of equations by the convexity of the cost function and the concavity of $\mathrm{Var}(\hat{\eta})$ with respect to e_i.

This determines the optimal *control mechanism* in the organization; in the first best, the manager exerts effort e_m^* and the worker exerts effort e_w^*, so that from (4), the weight placed on the opinion of the manager is given by $[\sigma_0^2\sigma_w^2]/[\sigma_0^2\sigma_m^2 + \sigma_w^2\sigma_m^2 + \sigma_m^2\sigma_0^2]$, and the weight on the worker's opinion is $[\sigma_0^2\sigma_m^2]/[\sigma_0^2\sigma_m^2 + \sigma_w^2\sigma_m^2 + \sigma_m^2\sigma_0^2]$. These weights are endogenously induced by the effort choices of the manager and the worker.

The first best may not be possible, as the worker has an incentive not to exert effort. For example, if the worker is offered a wage that is independent of his report, he exerts zero effort so that the first best cannot be obtained $[h_w(0) > h_w(e_w^*)]$. Hence the worker's wage must be conditioned on his report for the first best. However, there is only noisy evidence on η, so that the manager can only use his own observation, the initial prior, and the worker's report in order to determine whether or not the worker has exerted effort.

Contracts consist of a comparison of the worker's report with the manager's opinion of η. Because the worker is risk-neutral, one contract among the set of optimal contracts where $e_w > 0$ is to offer the worker a wage w_1 if the worker's report lies within an interval k of the manager's observation and w_0 otherwise, where $w_1 > w_0$. Therefore the worker is paid w_1 only if $|\hat{\eta}_w - \eta_m| < k$. This provides the worker with an incentive to exert effort, as exerting effort gives a better estimate on η and, hence, η_m.[4]

When the worker is risk-neutral and faces no liquidity constraints, it is well known that the manager induces the level of effort by the worker which minimizes $\mathrm{Var}(\hat{\eta}) + C_w(e_w) + C_m(e_m)$ in equilibrium. To see this,

[4]It is not necessarily the case that the worker can only be repaid when he is close to the manager's opinion. In fact, exactly the same analysis holds for any transformation of the manager's opinion. For example, assume that the worker is only paid well if his report is $\phi(\eta_m)$, for some monotonic function ϕ. Then the manager can invert the worker's report by simply inverting ϕ. The incentive to compute the manager's opinion is just as useful as when the objective is to second-guess the manager exactly.

note that the worker chooses e_w to

$$(9)\quad \max_{e_w}[\text{prob}(|\hat{\eta}_w-\eta_m|<k)]w_1 + [1-\text{prob}(|\hat{\eta}_w-\eta_m|<k)]w_0 - C_w(e_w)$$

subject to $h_w(e_w)=\sigma_w^2$. This implies a first-order condition

$$(10)\quad h_w'(e_w)\frac{\partial[\text{prob}(|\hat{\eta}_w-\eta_m|<k)]}{\partial\sigma_w^2}[w_1-w_0] = C_w'(e_w)$$

where $\partial(\text{prob}(\cdot))/\partial\sigma_w^2<0$, which determines the wage spread w_1-w_0. Then the worker's individual rationality constraint is satisfied by setting $[\text{prob}(|\hat{\eta}_w-\eta_m|<k)]w_1 + [1-\text{prob}(|\hat{\eta}_w-\eta_m|<k)]w_0 = r + C_w(e_w)$ for whatever e_w the manager chooses to induce. Thus, the provision of incentives costs no more than the cost of the effort exerted. (This is, of course, due to the assumption of risk-neutrality.) Hence I can restrict attention to contracts that minimize $\text{Var}(\hat{\eta})+C_w(e_w)+C_m(e_m)$, where $\hat{\eta}$ is the estimate based on the η_m, η_0, and $\hat{\eta}_w$. Thus the inefficiencies described here do not arise because the ability to provide incentives is more expensive than in the first best, as occurs in standard principal–agency models with risk-averse workers.

Consider the worker's incentives. The worker has observed η_w and must decide what value of $\hat{\eta}_w$ he should report to the manager. The worker has three pieces of information to use; η_0, η_λ, and η_w. Consider how the worker computes his opinion of the manager's observation. After observing η_0 and η_w, the posterior on η_m is given by a normal distribution with mean

$$(11)\qquad \eta_1=\frac{\sigma_0^2\eta_w+\sigma_w^2\eta_0}{\sigma_0^2+\sigma_w^2}$$

and variance

$$(12)\qquad \frac{\sigma_0^2\sigma_w^2+\sigma_w^2\sigma_m^2+\sigma_0^2\sigma_m^2}{\sigma_0^2+\sigma_w^2}.$$

(Note that this is the posterior distribution on η_m, not on η.) The worker also observes η_λ which is an unbiased normally distributed estimator of η_m with variance σ_λ^2. Hence, given η_w, η_0, and η_λ, the worker's estimate of η_m is given by a normal distribution with mean

$$(13)\quad \hat{\eta}_m=\frac{\frac{\sigma_0^2\eta_w+\sigma_w^2\eta_0}{\sigma_0^2+\sigma_w^2}\sigma_\lambda^2+\eta_\lambda\frac{\sigma_0^2\sigma_w^2+\sigma_0^2\sigma_m^2+\sigma_w^2\sigma_m^2}{\sigma_0^2+\sigma_w^2}}{\frac{\sigma_0^2\sigma_w^2+\sigma_0^2\sigma_m^2+\sigma_w^2\sigma_m^2}{\sigma_0^2+\sigma_w^2}+\sigma_\lambda^2}.$$

Therefore, given the worker's information, the worker reports $\hat{\eta}_w$ rather than η_w (i.e., a convex combination of the three signals). For notational convenience, let

$$(14)\quad \mu_0=\frac{\sigma_0^2\sigma_\lambda^2}{\sigma_0^2\sigma_w^2+\sigma_0^2\sigma_m^2+\sigma_w^2\sigma_m^2+\sigma_\lambda^2[\sigma_w^2+\sigma_0^2]}$$

$$\mu_1=\frac{\sigma_w^2\sigma_\lambda^2}{\sigma_0^2\sigma_w^2+\sigma_0^2\sigma_m^2+\sigma_w^2\sigma_m^2+\sigma_\lambda^2[\sigma_w^2+\sigma_0^2]}$$

$$\mu_2=\frac{\sigma_0^2\sigma_w^2+\sigma_0^2\sigma_m^2+\sigma_w^2\sigma_m^2}{\sigma_0^2\sigma_w^2+\sigma_0^2\sigma_m^2+\sigma_w^2\sigma_m^2+\sigma_\lambda^2[\sigma_w^2+\sigma_0^2]}.$$

This implies that, conditional on any effort level, the worker reports something other than that which he sees (i.e., he reports $\hat{\eta}_m=\mu_0\eta_m+\mu_1\eta_0+\mu_2\eta_\lambda$). This in itself is not necessarily a problem. Consider the case in which $\sigma_\lambda^2=\infty$, so that the worker gets no indication of the manager's opinion. Then as $\mu_2=0$, the worker reports

$$(15)\qquad \hat{\eta}_w=\mu_0\eta_w+\mu_1\eta_0$$

where here $\mu_0+\mu_1=1$. Once again, the worker does not report honestly; instead, he distorts towards the initial prior. However, as η_0 is known, this equation has only one unknown and can be solved for η_w. Hence, even when there is a desire to mimic the manager, it is not necessarily the case that there are efficiency losses. (For example, if the worker says that a project is "not bad," he really means that it is "good." Once

unique inversions of the worker's report are possible, there are no efficiency losses.)

A problem arises when $\sigma_\lambda^2 < \infty$. Then the worker's report depends on two variables which are unknown to the manager, η_w and η_λ, but only η_w is relevant for determining the true value of η (as η_m is a sufficient statistic for η_λ). To see this, note that the manager knows η_0 and μ_i, for all i, in equilibrium. Hence the uncertainty facing the manager over the true value of η_w is given by the random variable

$$z = \mu_0 \eta_m + \mu_2 \eta_\lambda. \tag{16}$$

The conditional distribution of η_w given z is normal with mean η_w and variance $[\mu_2/\mu_0]^2\sigma_\lambda^2$. Since $[\mu_2/\mu_0]^2\sigma_\lambda^2 > 0$, this implies that the existence of yes men introduces greater uncertainty than when the worker reports honestly. The key parameter for this analysis, then, is $\mu_2/\mu_0 = [\sigma_0^2\sigma_w^2 + \sigma_w^2\sigma_m^2 + \sigma_0^2\sigma_m]/\sigma_\lambda^2\sigma_0^2$.

When the worker reports honestly, the variance of $\hat{\eta}$ is given by

$$V^* = \frac{\sigma_w^2\sigma_m^2\sigma_0^2}{\sigma_m^2\sigma_0^2 + \sigma_m^2\sigma_w^2 + \sigma_0^2\sigma_w^2} \tag{17}$$

with effort given by (7) and (8). However, as the manager can only invert the worker's report to determine $z = \eta_w + [\mu_2/\mu_0]\eta_\lambda$, the estimate of the variance of $\hat{\eta}$ is given by

$$\frac{\left[\sigma_w^2 + \left(\frac{\mu_2}{\mu_0}\right)^2 \sigma_\lambda^2\right]\sigma_m^2\sigma_0^2}{\left[\sigma_w^2 + \left(\frac{\mu_2}{\mu_0}\right)^2 \sigma_\lambda^2\right]\left[\sigma_m^2 + \sigma_0^2\right] + \sigma_m^2\sigma_0^2} \tag{18}$$

$$> \frac{\sigma_w^2\sigma_m^2\sigma_0^2}{\sigma_w^2\sigma_0^2 + \sigma_m^2\sigma_w^2 + \sigma_0^2\sigma_m^2}.$$

Thus, because the worker has a preference to mimic the manager, information becomes noisier, so that for any given effort level, the firm's profits fall.[5] For example, if η is the profitability of a project, the project is incorrectly undertaken more frequently than is optimal and gets eschewed more often than when the worker reports honestly. Therefore, the existence of yes men where the worker conforms (i.e., where $\sigma_\lambda^2 < \infty$ and the manager monitors the worker) hinders the revelation of information.[6]

PROPOSITION 1: *If the worker exerts positive effort, he does not report his information honestly. This implies lower profits for the manager.*

There is one key point that must be stressed here. I have not assumed that workers wish to mimic their supervisors per se. Instead, this arises endogenously as the only way to induce the worker to exert effort so that the existence of yes men is a natural implication of subjective performance evaluation.

B. *The Efficiency of Incentive Contracts*

It has been informally argued that incentive contracts can be harmful to efficiency, as they give workers an incentive to carry out actions that are not in the firm's interest. However, so far, I have assumed the optimality of an incentive contract, where the worker exerts positive effort. In this section, I show that this is not necessarily the case; offering incentive contracts can reduce the firm's profits in some cases. I have assumed that $C_w'(0) = 0$, so that the cost of exerting a small amount of effort is small. Notwithstanding this, I now show that if σ_λ^2 is small enough, the firm's profits are higher with no incentive contract. The reason is simple. With incentive contracts, the

[5]Note that it may still be efficient to hire the supervisor if $\sigma_\lambda^2 > 0$, as the supervisor's information has some value.

[6]So far, I have stressed that the worker gets an observation on the manager's observation but knows the initial prior. Identical results arise when the worker gets no idea of the manager's observation but has a prior which is unknown to the manager.

worker lies about what he sees, but he exerts positive effort. With no incentive contract, the worker exerts zero effort but tells the truth, as the worker has no incentive but to tell the truth if his wage is not dependent on it. Thus, there is a discrete "honesty" cost to introducing incentive contracts. If σ_λ^2 is small enough (or, in other words, the ability to conform is strong enough), the firm prefers to induce the worker to tell the truth at a cost of zero effort.

Let $\underline{\sigma}^2$ be the variance of the worker's observation of η where $e_w = 0$ $[h_w(0) = \underline{\sigma}^2]$. Then if the manager offers the worker a contract that is independent of his report, the worker reports η_w which is normally distributed with mean η and variance $\underline{\sigma}^2$. Hence the manager's utility is

$$(19)\quad \underline{\pi} = -\frac{\underline{\sigma}^2\sigma_m^2\sigma_0^2}{\sigma_w^2\sigma_0^2 + \underline{\sigma}^2\sigma_m^2 + \underline{\sigma}^2\sigma_w^2} - C_m(\underline{e}_m)$$

where

$$(20)\quad C_m'(\underline{e}_m) = -h_m'(\underline{e}_m)\frac{\left[\sigma_0^2\underline{\sigma}^2\right]^2}{\left[\sigma_0^2\sigma_m^2 + \sigma_m^2\underline{\sigma}^2 + \underline{\sigma}^2\sigma_0^2\right]^2}.$$

On the other hand, if the firm offers an incentive contract, the optimal level of effort $(\tilde{e}_m, \tilde{e}_w)$ is characterized by (7) and (8) with $\tilde{\sigma}_w^2 = \sigma_w^2 + (\mu_2/\mu_0)^2\sigma_\lambda^2$, so that profits are given by

$$(21)\quad \hat{\pi} = -\frac{\left[\sigma_w^2 + \left(\frac{\mu_2}{\mu_0}\right)^2\sigma_\lambda^2\right]\sigma_m^2\sigma_0^2}{\left[\sigma_w^2 + \left(\frac{\mu_2}{\mu_0}\right)^2\sigma_\lambda^2\right](\sigma_m^2 + \sigma_0^2) + \sigma_m^2\sigma_0^2} - C_m(\tilde{e}_m) - C_w(\tilde{e}_w).$$

Note that $\hat{\pi}$ is continuously increasing in σ_λ^2, as the desire to conform falls. This has a lower bound of $-\sigma_m^2\sigma_0^2/(\sigma_m^2 + \sigma_0^2) - C_m(\underline{e}_m)$ as $\sigma_\lambda^2 \to 0$ and an upper bound of $-V^* - C_m(e_m^*) - C_w(e_w^*)$ when $\sigma_\lambda^2 = \infty$; but

$$(22)\quad -V^* - C_m(e_m^*) - C_w(e_w^*) > \underline{\pi} > -\frac{\sigma_m^2\sigma_0^2}{\sigma_m^2 + \sigma_0^2}C_m(\underline{e}_m).$$

Then, as profits vary monotonically and continuously in σ_λ^2, there exists a unique $\sigma_\lambda^{2*} > 0$ such that the profits from offering an incentive contract exactly equal those from not doing so. This implies Proposition 2.

PROPOSITION 2: *If $\sigma_\lambda^2 < \sigma_\lambda^{2*}$, the manager's profits are higher from offering no incentive contract. If $\sigma_\lambda^2 > \sigma_\lambda^{2*}$, an incentive contract is offered.*

The parameter σ_λ^{2*} determines the critical level of noise above which the firm wants to encourage effort over honesty. For parameter values lower than this critical level, the returns to inducing effort exertion are lower than those from encouraging honesty. Hence, in organizations where workers know what their managers want to hear, incentive contracts should not be used in order to encourage honesty.

Note that this section provides another reason why firms may not offer incentive pay to workers. In many ways, this result is similar to that of Holmstrom and Milgrom (1991), who argue that low-powered incentives should be used when there is a negative externality from "effort" on one task to "effort" on another. Here the two activities are effort and honesty, where providing incentives to exert effort has a harmful effect on honesty.

C. *Effort Levels*

Proposition 1 illustrates that the desire to coordinate with the manager implies that information is not aggregated as efficiently as when the worker reports his information honestly. Proposition 2 illustrates that it may be optimal to offer the supervisor no incentives. This desire to conform is now shown more generally to change the optimal level of effort by the manager and the worker, which changes their relative input to decisions based on η. I first show that, for any fixed effort level by the worker, the reduced value of the worker's information caused by conformity increases the optimal level of managerial effort. Hence, more weight is likely to be placed on the manager's opinion than when information is reported honestly,

so that for example the manager has more input over whether a project is undertaken or not.

The optimal effort level given the desire to conform is given by a pair of equations similar to (7) and (8), taking into account the added variance in the worker's report. First consider the manager's choice of effort level, fixing the worker's effort level and beliefs. Note that

$$(23)\quad \frac{\partial \operatorname{Var}(\hat{\eta})}{\partial \sigma_{\mathrm{m}}^2} = \frac{\left[\tilde{\sigma}_{\mathrm{w}}^2 \sigma_0^2\right]^2}{\left[\tilde{\sigma}_{\mathrm{w}}^2 \sigma_{\mathrm{m}}^2 + \sigma_{\mathrm{m}}^2 \sigma_0^2 + \sigma_0^2 \tilde{\sigma}_{\mathrm{w}}^2\right]^2}$$

where $\tilde{\sigma}_{\mathrm{w}}^2 = \sigma_{\mathrm{w}}^2 + [\mu_2/\mu_0]^2 \sigma_\lambda^2$; but also note that

$$(24)\quad \frac{\partial^2 \operatorname{Var}(\hat{\eta})}{\partial \sigma_{\mathrm{m}}^2 \partial \tilde{\sigma}_{\mathrm{w}}^2} = \frac{2\left[\sigma_{\mathrm{m}}^2 \tilde{\sigma}_{\mathrm{w}}^2 \sigma_0^6\right]}{\left[\tilde{\sigma}_{\mathrm{w}}^2 \sigma_{\mathrm{m}}^2 + \sigma_{\mathrm{m}}^2 \sigma_0^2 + \sigma_0^2 \tilde{\sigma}_{\mathrm{w}}^2\right]^3} > 0.$$

Hence the marginal return to exerting effort (reducing the variance of η) increases as $\tilde{\sigma}_{\mathrm{w}}^2$ increases. Proposition 3 follows.

PROPOSITION 3: *For any fixed effort level by the worker, the worker's incentive to conform increases the manager's effort, implying that more weight is placed on the manager's opinion than in the absence of yes men.*

The implication of this is that when subjective evaluation is used, for any fixed effort level by the worker, more control is placed in the hands of managers (as $[\tilde{\sigma}_{\mathrm{w}}^2/\sigma_{\mathrm{m}}^2]$ rises) as their opinions carry more weight in the determination of the true value of η than in the absence of yes men. Therefore, the existence of yes men (caused by the desire to conform when incentives are offered to the supervisor) implies less delegation within organizations, in the sense that the workers have less input to decision-making. It should be stressed here that this result does not arise because the provision of incentives is more costly with the noisier signal, as would be the case in a standard moral-hazard problem such as in Holmstrom (1979). The firm can always induce the worker to exert effort for the expected cost of the worker's effort. Instead, it is the *reduced value* of the worker's information that causes the result, since the worker's information is contaminated with noise.

I now consider how the conformity problem affects the optimal contract to the worker, conditional on incentives being offered. To illustrate the worker's incentives, first fix the effort level of the manager. Then, the optimal level of effort for the worker is characterized by

$$(25)\quad C_{\mathrm{w}}'(e_{\mathrm{w}}) = -h_{\mathrm{w}}'(e_{\mathrm{w}}) \frac{\partial \operatorname{Var}(\hat{\eta})}{\partial \sigma_{\mathrm{w}}^2}.$$

However,

$$(26)\quad \frac{\partial \operatorname{Var}(\hat{\eta})}{\partial \sigma_{\mathrm{w}}^2} = \frac{\left[1 + 2\frac{\mu_2}{\mu_0}\left(\frac{\sigma_0^2 + \sigma_{\mathrm{m}}^2}{\sigma_0^2}\right)\right]\sigma_0^4 \sigma_{\mathrm{m}}^4}{\left\{\left[\sigma_{\mathrm{w}}^2 + \left(\frac{\mu_2}{\mu_0}\right)^2 \sigma_\lambda^2\right](\sigma_0^2 + \sigma_{\mathrm{m}}^2) + \sigma_{\mathrm{m}}^2 \sigma_0^2\right\}^2}.$$

In order to consider the impact of the desire to conform on the level of effort exerted by the worker in equilibrium, I differentiate (26) with respect to σ_λ^2 as higher values of that parameter reduce the desire to conform ($[\mu_2/\mu_0]^2 \sigma_\lambda^2$ is decreasing in σ_λ^2). Differentiating (26) with respect to σ_λ^2 gives an expression which cannot be signed without knowledge of σ_λ^2. More specifically, it is straightforward to show that the marginal return to effort by the worker is decreasing in σ_λ^2 if and only if $\sigma_\lambda^2 \leq 2$. For $\sigma_\lambda^2 > 2$, increasing σ_λ^2 (reducing the degree of conformity) increases the marginal product of effort for the worker. Hence the marginal return to increasing the worker's effort is lower in the absence of yes men only if the degree of certainty about the manager's opinion is sufficiently high, so that in general I cannot predict the effect of yes men on the effort level of the worker.

The reason for this is as follows. At first blush, it would appear that the worker should exert less effort as the value of his

information is lower. This is indeed the case if the extent to which the worker distorts ($[\mu_2/\mu_0]^2\sigma_\lambda^2$) is independent of σ_m^2; but the distortion in the worker's report falls as the worker exerts more effort as he trusts his own observation more ($\partial[\mu_2/\mu_0]/\partial e_w < 0$). Hence there is a return to increasing effort which is not evident in the absence of the desire to conform. These two conflicting effects imply that the effect of conformity on the optimal effort level of the worker cannot be signed.

For example, if $\partial^2 \mathrm{Var}(\hat{\eta})/\partial\tilde{\sigma}_w^2\partial\sigma_\lambda^2 > 0$, which occurs if the *marginal* product of the worker's effort is lower than in the absence of yes men (where σ_λ^2 is ∞), then the equilibrium level of effort by the worker falls, and the equilibrium level of effort by the manager rises [from (24)]. However, if $\partial^2 \mathrm{Var}(\hat{\eta})/\partial\tilde{\sigma}_w^2\partial\sigma_\lambda^2 < 0$, the marginal return to exerting effort rises with yes men. Then the effect of yes men on equilibrium effort and control over decisions becomes less clear. To see this, note that

$$(27)\qquad \frac{\partial e_w}{\partial \sigma_\lambda^2} = -h'_w C_w^{-1\prime\prime}\frac{\partial^2 \mathrm{Var}(\hat{\eta})}{\partial\tilde{\sigma}_w^2\partial\sigma_\lambda^2}$$

and

$$(28)\qquad \frac{\partial e_m}{\partial \sigma_\lambda^2} = -h'_m C_m^{-1\prime\prime}\frac{\partial^2 \mathrm{Var}(\hat{\eta})}{\partial\sigma_m^2\partial\tilde{\sigma}_w^2}\frac{\partial\tilde{\sigma}_w^2}{\partial\sigma_\lambda^2}.$$

Then, if $\partial^2 \mathrm{Var}(\hat{\eta})/\partial\tilde{\sigma}_w^2\partial\sigma_\lambda^2$ is positive, or negative and small, the variance of the worker's estimate, $\tilde{\sigma}_w^2$, is higher with yes men than with truth-telling in equilibrium. Therefore, the desire to conform leads to more centralized decision-making. However, it is possible that $\partial^2 \mathrm{Var}(\hat{\eta})/\partial\tilde{\sigma}_w^2\partial\sigma_\lambda^2$ is sufficiently large and negative such that $\partial\tilde{\sigma}_w^2/\partial\sigma_\lambda^2$ is negative. If this is the case, in equilibrium the worker exerts so much more effort that the manager exerts less [from (24)] as the employees' efforts are substitutes, implying less centralization of decision-making.

This outcome seems unlikely. The economic intuition for this perverse result is that, although the value of the worker's information is worse with yes men in the sense that the variance is higher for any effort level by the worker, the *marginal* return to getting information is so high that it swamps any reduction in the value of the worker's information. This implies higher precision of the worker's estimate in equilibrium than with truth-telling. It seems most likely that the outcome of reducing the value of the worker's information for any effort level is that the manager exerts more effort to compensate for this, resulting in more centralized decision-making.

D. *Comparative Statics*

The model has been chosen to be simple to provide comparative-static results in order to provide predictions on when the incidence and costs of conformity are likely to be high. I now consider how the inefficiencies associated with conformity are affected by various parameters of the model, assuming incentives are offered. This is measured by the noise associated with the worker's report, $[\mu_2/\mu_0]^2\sigma_\lambda^2$. As $[\mu_2/\mu_0]^2\sigma_\lambda^2$ rises, the costs associated with yes men rise.

PROPOSITION 4: *(i) As σ_w^2 rises, the costs of conformity rise; (ii) as σ_λ^2 rises, the costs of conformity fall; (iii) as σ_0^2 rises, the costs of conformity fall; (iv) as σ_m^2 rises, the costs of conformity rise.*

PROOF:

I simply show how $[\mu_2/\mu_0]^2\sigma_\lambda^2$ varies with the appropriate parameters. For (i),

$$(29)\qquad \frac{\partial(\mu_2/\mu_0)}{\partial\sigma_w} = \frac{2\sigma_0^2\sigma_w^2}{\sigma_0^2\sigma_\lambda^2} > 0.$$

For (ii),

$$(30)\qquad \frac{\partial(\mu_2/\mu_0)^2\sigma_\lambda^2}{\partial\sigma_\lambda} = -2\left[\frac{\mu_2}{\mu_0}\right]^2\sigma_\lambda < 0.$$

For (iii),

$$(31)\quad \frac{\partial(\mu_2/\mu_0)}{\partial\sigma_0} = \frac{-2\sigma_0\sigma_\lambda^2\sigma_w^2\sigma_m^2}{\left[\sigma_0^2\sigma_\lambda^2\right]^2} < 0.$$

For (iv),

$$(32)\quad \frac{\partial(\mu_2/\mu_0)}{\partial\sigma_m} = 2\sigma_m\frac{\sigma_0^2+\sigma_w^2}{\sigma_0^2\sigma_\lambda^2} > 0.$$

Part (i) of Proposition (4) arises because as σ_w^2 gets smaller, the value of the worker's observation increases so that he places more weight on what he sees himself rather than on what he believes the manager to have seen. Hence conformity becomes less costly with more talented workers. On the other hand, as σ_λ^2 gets smaller, workers get a better indication of what the manager wants. This implies that they use their own information less and rely on the observation of the manager more. As the variance on the common prior falls, both the precision of the worker's estimate and the λ estimate get better, but the speed of variance reduction is higher for the λ estimate, so that yes men are more likely (in the sense of $[\mu_2/\mu_0]^2\sigma_\lambda^2$ getting larger) as σ_0^2 gets smaller. Finally, as σ_m^2 is reduced, the observation on η_w and η_0 is a better predictor of η_m, and hence the worker relies less on η_λ.

What does this imply for organizational practices? First, high-ability workers are likely to have lower variances in their estimate of η. Consequently, Proposition 4 suggests that "toadyism" is likely to be concentrated among the less able, who trust their own opinion less than do able workers. The natural corollary of this is that yes men are a more costly problem with incompetent (high-variance) managers, simply for the reason that managers with high variance are liable to come out with bizarre suggestions, implying that mimicking these opinions imposes significant costs. Third, the model predicts that as σ_λ^2 rises, the costs of conformity fall. In other words, the more interaction there is between managers and workers, the greater the costs of conformity, since greater interaction is likely to lead to workers getting a better opinion on their superiors' beliefs (reducing σ_λ^2). This suggests that a policy a company may follow to reduce these costs is to encourage aloofness in its managers, thereby retaining a distinction between line and staff, in order to reduce the ability of workers to infer a manager's wishes.

Similar comparative statics apply for the optimality of the use of incentive contracts over flat incentive schemes. More specifically, (i) less able workers are less likely to be offered contracts than talented workers, and (ii) where there is much interaction between the manager and the worker, an incentive contract is unlikely. Therefore, the results also suggest that conformity is likely to be a problem when firms use high-powered incentives (with no incentive contract, there is no incentive to lie).

II. Communication in Organizations

So far, I have couched the model in terms of a worker imitating his superior. Yet there is nothing specific to managerial opinion which gives rise to this incentive. In this section of the paper, I consider how workers who have access to the information of others tend to conform to one another's opinions if their compensation depends on the quality of their information.

There are two obvious interpretations of communication here. First, team production has attracted much recent attention in the incentives literature, illustrating a variety of issues that do not arise with individual production (see Holmstrom, 1982; Edward Lazear, 1989; Hideshi Itoh, 1991). Team members routinely communicate with one another, so that information is likely to be shared among individuals. Second, perhaps the most common example of interaction between employees within organizations occurs when workers seek the opinions of others. For example, different departmental heads may consult on a project. This can obviously act as an efficient means of transmitting information and ironing out problems. In this section, I wish to consider a

disadvantage of interaction; namely, if an individual's compensation is dependent on his own contribution, there is an incentive for individuals to conform to the opinions held by others.

Assume that there is a group of N (risk-neutral) workers who are assigned to determine the true value of η. Member i of the group of workers exerts effort e_i at a cost $C_i(e_i)$ where $C_i'(e_i) > 0$, $C_i''(e_i) > 0$, $C_i'(0) = 0$, and $C_i'(\infty) = \infty$. By exerting effort, worker i observes

$$\eta_i = \eta + \varepsilon_i \tag{33}$$

where ε_i is distributed normally with zero mean and variance $\sigma_i^2 = h_i(e_i)$, with $h_i' < 0$ and $h_i'' > 0$. All error terms are uncorrelated. The manager makes no observation.

As above, the manager induces the workers to exert effort by a relative performance evaluation scheme; each worker is paid in relation to the reports of other workers and his own report. If the members of the group have no information other than their own, they report their estimate of η, which can be inverted exactly as in (15) to determine η_i. Then the manager simply computes the least-squares estimate of η given all η_i, which is given by

$$\hat{\eta} = \frac{\sum_{i=0}^{N} \prod_{\substack{j=0 \\ j \neq i}}^{N} \sigma_j^2 \eta_i}{\sum_{k=0}^{N} \prod_{\substack{j=0 \\ j \neq k}}^{N} \sigma_j^2} \tag{34}$$

where the σ_i^2 depend on the equilibrium levels of effort exerted by the workers. As above, if reports are made honestly, the manager will induce the first-best level of effort exertion where

$$C_i'(e_i^*) = -h_i'(e_i^*) \frac{\partial \operatorname{Var}(\hat{\eta} \mid \sigma_1^2, \sigma_2^2, \ldots, \sigma_N^2)}{\partial \sigma_i^2}. \tag{35}$$

I assume, however, that individuals who work together interact in such a way that they develop ideas about their colleagues' impressions. (For example, after seminars, academics typically discuss papers presented.) I assume that worker i does not observe j's opinion exactly but does so with error. In particular, worker i observes η_j^i on what worker j believes, where

$$\eta_j^i = \eta_j + \lambda_j^i \tag{36}$$

and where λ_j^i is normally distributed with mean zero and variance $\sigma_{\lambda,ij}^2$. I assume that $\sigma_{\lambda,ij}^2 > 0$ for all i, j and that all error terms are uncorrelated.

As above, incentives can only be provided by relative performance evaluation. For example, the workers could be evaluated by a comparison of a worker's report with the manager's estimate of η after all reports are garnered. Hence, the worker has an incentive to report on the basis of $\hat{\eta}_i$, which is his opinion of the true value of η given his own observation and his observation on the opinions of others. The worker estimates $\hat{\eta}_i$ by equation (37), below.

$$\hat{\eta}_i = \left(\frac{\prod_{\substack{j=1 \\ j \neq i}}^{N} \left(\sigma_j^2 + \sigma_{\lambda,ij}^2\right)}{\sum_{k=1}^{N} \prod_{\substack{j=1 \\ j \neq i,k}}^{N} \left(\sigma_j^2 + \sigma_{\lambda,ij}^2\right)} \right) \frac{\sigma_0^2 \eta_i + \sigma_i^2 \eta_0}{\sigma_0^2 + \sigma_i^2} + \left(\frac{\sum_{j=1}^{N} \prod_{\substack{j=1 \\ j \neq i}}^{N} \left(\sigma_j^2 + \sigma_{\lambda,ij}^2\right) \eta_j^i}{\sum_{k=1}^{N} \prod_{\substack{j=1 \\ j \neq i,k}}^{N} \left(\sigma_j^2 + \sigma_{\lambda,ij}^2\right)} \right) \frac{\sigma_0^2 \sigma_i^2}{\sigma_0^2 + \sigma_i^2} \tag{37}$$

Note that $\hat{\eta}_i$ differs from η_i and $\hat{\eta}$ because no worker can aggregate information as well as the manager.

In exactly the same way as in previous cases, the manager can invert worker i's report to derive a random variable, z:

$$(38)\quad z_i = \eta_i + \sigma_i^2 \left(\frac{\sum_{\substack{k=1 \\ k\neq i}}^{N} \prod_{\substack{j=1 \\ j\neq i,k}}^{N} [\sigma_j^2 + \sigma_{\lambda,ij}^2]\eta_j^i}{\prod_{\substack{j=1 \\ j\neq i}} (\sigma_j^2 + \sigma_{\lambda,ij}^2)} \right).$$

This variable is normally distributed with mean η and variance

$$(39)\quad \sigma_i^2 + \sigma_i^4 \left(\frac{\sum_{\substack{k=1 \\ k\neq i}}^{N} \prod_{\substack{j=1 \\ j\neq k,i}}^{N} [\sigma_j^2 + \sigma_{\lambda,ij}^2]^2 \sigma_{\lambda,ij}^2}{\prod_{\substack{j=1 \\ j\neq i}} (\sigma_j^2 + \sigma_{\lambda,ij}^2)^2} \right) > \sigma_i^2.$$

Hence in the same way as in Section I, the worker reports a noisier signal than with truth-telling if he is offered incentives to exert effort. This implies that the firm's profits are lower than with truth-telling, as $\sigma_{\lambda,ij}^2 > 0$ for all i, j. Therefore, the firm's profits are higher in the absence of interaction between workers.

The purpose of this section is to illustrate how any interaction between individuals can stifle "creativity," as all workers distort by converging to the opinions believed to be held by others. This would suggest that organizations (or countries) that extensively use teamwork may find creativity stifled. Note that this issue is different from workers colluding with one another over reporting their findings. Here workers report their best guess given their information. Finally, one implication of this result is that companies may find it in their interests to construct "Chinese Walls" between departments, whereby one department does not know about the operations of another. In this way, conformity problems can be reduced.

It is straightforward to show that if the incentive to conform is large enough ($\sigma_{\lambda,ij}^2$ small enough), then the firm maximizes profits by not offering an incentive contract. Therefore, this paper suggests that team production may imply low-powered incentives, as interaction is obviously prevalent in teams. A similar result is found by Lazear (1989), who argues that low-powered incentives may aid cooperation between workers; here the idea is simply to give workers incentives for honesty.

It must be said that this is not designed to be a general theory of information-sharing between individuals. In general, there are advantages to individuals sharing information. However, a by-product of such information-sharing may be that individuals tend to conform to the opinions of others in a way that can be inefficient. Organizations must weigh these disadvantages against the traditional advantages of workers cooperating with one another.

III. Conclusion

The purpose of this analysis has been to illustrate a fundamental trade-off between inducing workers to tell the truth and inducing them to exert effort. If workers are rewarded on a subjective basis, they may distort their behavior toward what they feel their superiors want to hear. However, to induce the worker to act honestly requires that either incentives to exert effort are eliminated or more objective measures of performance are obtained.

Beyond simply illustrating the possible incentive for subordinates to conform to their superiors' opinions, this paper illustrates inefficiencies associated with yes men. First, information becomes noisier, implying for example less efficient project selection or monitoring of workers. This was shown to imply a reorganization of responsibilities within the organization, with increased emphasis on the manager's information becoming likely. It was also shown that, conditional on incentives being offered, yes men

are likely to be concentrated among the less able workers, among workers with less able managers, in organizations with much interaction between management and workers, and in organizations with high-powered incentives. Following the basic results, I extended the same ideas of conformity to a theory of communication in organizations. I showed that an organization that intensively uses teams may stifle creativity, as members have incentives to conform which do not arise in the absence of interactions.

This paper is incomplete in a number of dimensions. Perhaps the most important of these is that I allow managers no way to determine the veracity of the workers' reports. In reality, managers routinely ask subordinates why they came to a particular conclusion. Introducing this means of questioning will, I am sure, change the extent to which yes men can exit. However, once workers have the opportunity to distort their reports, the existence of yes men is likely to remain. Another limitation is that it does not address why managers may wish to have cronies who agree with them. In this paper, the manager prefers that the subordinate tells the truth rather than agree with him. In reality, this is open to question. Yet the basic insights are likely to remain, namely, that workers are likely to exercise such behavior when compensation is sensitive to subjective evaluation. Otherwise, why would the worker bother to agree with his superior? Hence the basic insights would appear to be robust to specifications in which managers receive utility from subordinates agreeing with them.

REFERENCES

Banerjee, Abhijit V., "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, August 1992, *107*, 797–819.

Bernheim, B. Douglas, "A Theory of Conformity," mimeo, Princeton University, July 1991.

Holmstrom, Bengt, "Moral Hazard and Observability," *Bell Journal of Economics*, Spring 1979, *10*, 74–91.

_____, "Moral Hazard in Teams," *Bell Journal of Economics*, Autumn 1982, *13*, 324–40.

_____ **and Milgrom, Paul R.,** "Multitask Principal–Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law, Economics, and Organization*, August 1991, *7*, 24–52.

Itoh, Hideshi, "Incentives to Help in Multi-agent Situations," *Econometrica*, May 1991, *59*, 611–36.

Lazear, Edward, "Pay Equality and Industrial Politics," *Journal of Political Economy*, June 1989, *97*, 561–80.

Milgrom, Paul R. and Roberts, D. John, "An Economic Approach to Influence Activities and Organizational Responses," *American Journal of Sociology*, July 1988, *94* (Supplement), S154–79.

Scharfstein, David and Stein, Jeremy, "Herd Behavior and Investment," *American Economic Review*, June 1990, *80*, 465–79.