Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Relational Contracts, Multiple Agents, and Correlated Outputs

Ola Kvaløy, Trond E. Olsen

informs
http://pubsonline.informs.org/journal/mnsc/

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–11
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Relational Contracts, Multiple Agents, and Correlated Outputs

Ola Kvaløy,[a,b] **Trond E. Olsen**[b]

[a] University of Stavanger Business School, 4036 Stavanger, Norway; [b] Department of Business and Management Science, Norwegian School of Economics, 5045 Bergen, Norway

Contact: ola.kvaloy@uis.no, http://orcid.org/0000-0003-2004-5807 (OK); trond.olsen@nhh.no (TEO)

Received: July 6, 2017
Revised: January 10, 2018; May 14, 2018
Accepted: June 25, 2018
Published Online in Articles in Advance: May 14, 2019

https://doi.org/10.1287/mnsc.2018.3162

Abstract. We analyze relational contracts between a principal and a set of risk-neutral agents whose outputs are correlated. When only the agents' aggregate output can be observed, a team incentive scheme is shown to be optimal, where each agent is paid a bonus for aggregate output above a threshold. We show that the efficiency of the team incentive scheme depends on the way in which the team members' outputs are correlated. The reason is that correlation affects the variance of total output and thus, the precision of the team's performance measure. Negatively correlated contributions reduce the variance of total output, and this improves incentives for each team member in the setting that we consider. This also has implications for optimal team size. If the team members' outputs are negatively correlated, more agents in the team can improve efficiency. We then consider the case where individual outputs are observable. A tournament scheme with a threshold is then optimal, where the threshold depends on an agent's relative performance. We show that correlation affects both the efficiency and design of the optimal tournament scheme.

History: Accepted by Shivaram Rajgopal, accounting.

Supplemental Material: The online appendix is available at https://doi.org/10.1287/mnsc.2018.3162.

Keywords: relational contracts • performance evaluation • incentives • teams

1. Introduction

Within organizations, employees' outputs are often correlated: many times positively, such as when they are exposed to the same business cycles, and other times negatively, such as when they compete for the same resources or meet different sets of demands from customers or superiors. For instance, professionals within the same partnerships who specialize in different industries can have rather asymmetric income shocks.

The way in which outputs correlate is potentially important for incentive design. Correlation affects the gains from risk sharing if agents are exposed to joint performance evaluation, and it filters out common noise if agents are exposed to relative performance evaluation (RPE) (Holmström 1982, Mookherjee 1984). Recent contributions within the accounting literature show how these insights affect more practical questions regarding bonus arrangements and organizational design. For instance, Huddart and Liang (2005) show that (positively) correlated performances reduce the optimal size of partnerships, and Rajan and Reichelstein (2006) show the importance of correlated signals for the design and efficiency of discretionary bonus pools.

A common feature in this literature is that correlated performances affect the efficiency of incentive systems via risk considerations. The literature is based on models where at least some of the performance measures are verifiable, and hence, first best contracts are achievable if agents are comfortable with bearing risk. In this paper, we show that correlated performances are highly important for incentive design, even in the absence of risk considerations. In contrast to previous literature, we study how correlated performances affect optimal incentives in situations where no verifiable performance measures are available.

In practice, incentive contracts are often based on performance measures that are difficult to verify by a third party (e.g., Gibbs et al. 2004). The quality or value of the agents' performance may be observable to the principal but cannot easily be assessed by a court of law. The parties must then rely on self-enforcing relational contracts. Through repeated interactions, the parties can make it costly for each other to breach the contract by letting breach ruin future trade. However, relational contracts cannot fully solve the principal's incentive problem, because the agents' monetary incentives (bonuses) are limited by the value of the future relationship. If bonuses are too large (or too small), the principal (or agents) may deviate by not paying as promised, thereby undermining the relational contract. The principal must thus provide as efficient incentives as possible under the constraint that the feasible bonuses are limited.

In this paper, we analyze optimal relational contracts between a principal and a set of agents whose outputs are either positively or negatively correlated. We focus exclusively on the effects of stochastic dependencies and therefore, exclude any "technological" dependencies (e.g., complementarities) between the agents. This is not to deny that the latter can be important, but their effects are reasonably well understood (e.g., Levin 2002). We consider two cases: (a) where only aggregate output can be observed and (b) where individual outputs can be observed.

We first show that the optimal contract under (a) is a team incentive scheme, where each agent is paid a maximal bonus for aggregate output above a threshold and a minimal (no) bonus otherwise.[1] This parallels the characterization of Levin (2003) for the single-agent case. We then show, for a parametric (normal) distribution, that the efficiency of the team incentive scheme depends on the way in which the team members' outputs are correlated. The reason is that correlation affects the variance of total output and thus, the precision of the team's performance measure. Variance is important, not because it affects risk (because all agents are risk neutral by assumption) but because it affects, for any given bonus level, the incentives for each team member to provide effort. The lower the variance and thus, the more precise the performance measure, the stronger the marginal effect of each agent's effort is on the probability to obtain the bonus and thus, the stronger the marginal incentives (MIs) are for effort. A team composed of agents with negatively correlated outputs has this effect. It reduces the variance of total output and thus, improves incentives.[2]

This also has implications for optimal team size. We show that the team's efficiency decreases with its size (number of agents n) when outputs are nonnegatively correlated but that efficiency may increase considerably with size if outputs are negatively correlated.

In case (b), where individual output is observable, Levin (2002) showed that, for independent outputs, the optimal relational contract entails a stark RPE scheme: a form of tournament where at most, one agent is paid a bonus. We point out that the efficiency of this tournament scheme increases with the number of agents and hence, becomes progressively better compared with a team when the number of independent agents increases. Then, we extend the analysis to correlated variables and show for the parametric (normal) distribution that the optimal contract is an RPE scheme with a threshold, where the threshold depends on an agent's relative performance and where the conditions for an agent to obtain the (single) bonus are then stricter for negatively compared with positively correlated outputs. The reason for this is that the losing agent's output is informative about the winner's expected output when these outputs are correlated. Under negative correlation, a bad performance by one agent raises the conditional expected performance of the other agent. Hence, to maximize incentives, the bonus threshold should increase when outputs are negatively correlated.

The efficiency of the tournament contract is shown to improve with stronger correlation, both positive and negative. The latter aspect is noteworthy, because the efficiency of a standard tournament (i.e., a tournament where the winner gets a fixed bonus irrespective of performance) decreases with stronger negative correlation. In contrast, when the winning agent also needs to pass a hurdle that depends on the other agent's output, then correlation (both positive and negative) reduces the importance of luck and increases the importance of effort to achieve the bonus.

The main contribution of this paper is to analyze correlated performance measures in multiagent relational contracts. To the best of our knowledge, this has not yet been analyzed in the literature. Our secondary contribution is to vary the number of agents in the model and thus, also analyze the effects of team size. Although the literature on team incentives generally recognizes team size as an important determinant for team performance, questions concerning optimal team size have received limited attention.[3] Most notable are the contributions within the accounting literature: in particular, Huddart and Liang (2003, 2005) and Liang et al. (2008) show that team size can affect monitoring activities within teams as well as how teams respond to exogenous shocks. An interesting implication from Huddart and Liang (2005) is that partnerships are less likely to increase in size if outputs are positively correlated. It follows from the general idea of Holmström (1982) that larger teams can achieve more efficient risk sharing given that team members' contributions are not perfectly correlated. We point to a different mechanism: Adding more agents to the team when contributions are positively correlated leads to less precise performance measures.

An important difference between the works of Huddart and Liang (2003, 2005) and our paper is that we have a principal who can withhold payments to agents. Our relational contracting approach may thus be more relevant for teams within corporations than for partnerships. Previous literature on relational contracts between a principal and multiple agents considers situations in which there exist observable signals about individual performances (Levin 2002; Kvaløy and Olsen 2006, 2008; Rayo 2007; Baldenius et al. 2016; Deb et al. 2016; Glover and Xue 2018). However, individual contributions to the firm's output are often unobservable, which was underscored by Alchian and Demsetz (1972). Surprisingly then, relational contracts between a principal and a team of agents, where only aggregate output is observable, have (to our best knowledge) not yet been studied.[4]

The rest of the paper is organized as follows. Section 2 presents the model and analyzes team incentives given that only total output can be observed. Section 3 deals with the case where individual outputs can be observed, whereas Section 4 concludes.

2. Model

We analyze an ongoing economic relationship between a principal and n (symmetric) agents. All parties are risk neutral. Each period, each agent i exerts effort e_i, incurring a private cost $c(e_i)$. Costs are strictly increasing and convex in effort (i.e., $c'(e_i) > 0$, $c''(e_i) > 0$, and $c(0) = c'(0) = 0$). Each agent's effort generates a stochastic output x_i, with marginal density $f(x_i, e_i)$. Expected outputs are given by $\bar{x}(e_i) = E(x_i | e_i) = \int x_i f(x_i, e_i) dx_i$, and total surplus per agent is $W(e_i) = \bar{x}(e_i) - c(e_i)$. First best is then achieved when $\bar{x}'(e_i^{FB}) - c'(e_i^{FB}) = 0$. Outputs are stochastically independent (given efforts) across time.

The parties cannot contract on effort provision. We assume that effort e_i is hidden and only observed by agent i. With respect to output, we consider two cases: either individual outputs x_i are observable, or only total output $y = \Sigma x_i$ is observable. In both cases, we assume that outputs are nonverifiable by a third party. Hence, the parties cannot write a legally enforceable contract on output provision but have to rely on self-enforcing relational contracts.

2.1. Team: Only Total Output Observed

We first consider the case where individual output is unobservable, and hence, the parties can only contract on total output provision. We focus here on team effects generated by stochastic dependencies among agents' contributions, and thus, we assume a simple linear "production structure" but allow individual outputs to be stochastically dependent.

Each period, the principal and the agents then face the following contracting situation. First, the principal offers a contract saying that agent i receives a noncontingent fixed salary α_i plus a bonus $b_i^T(y)$, $i = 1 \ldots n$ conditional on total output $y = \Sigma x_i$ from the n agents.[5] Second, the agents simultaneously choose efforts, and value realization y is revealed. Third, the parties observe y, and the fixed salary α_i is paid. Then, the parties choose whether to honor the contingent bonus contract $b_i^T(y)$.

Conditional on efforts, agent i's expected wage in the contract is then $w_i = E(b_i^T(y) | e_1 \ldots e_n) + \alpha_i$, whereas the principal expects $E(y | e_1 \ldots e_n) - \Sigma w_i = \Sigma_i E(x_i | e_i) - \Sigma w_i$. If the contract is expected to be honored, agent i chooses effort e_i to maximize his payoff; that is,

$$e_i = \arg\max_{e_i'} (E(b_i^T(y) | e_i', e_{-i}) - c(e_i')). \qquad \text{(IC)}$$

The parties have outside (reservation) values normalized to zero. In the repeated game that we consider, like Levin (2002), a multilateral punishment structure is used, where any deviation by the principal triggers punishment from all agents. The principal honors the contract only if all agents honored the contract in the previous period. The agents honor the contract only if the principal honored the contract with all agents in the previous period. Thus, if the principal reneges on the relational contract, all agents take their outside option forever after and vice versa: if one (or all) of the agents reneges, the principal takes her outside option forever after.[6] A natural explanation for this is that the agents interpret a unilateral contract breach (i.e., the principal deviates from the contract with only one or some of the agents) as evidence that the principal is not trustworthy (see discussions in Bewley 1999 and Levin 2002).

Now (given that (IC) holds), the principal will honor the contract with all agents $i = 1, 2, \ldots, n$ if

$$-\Sigma_i \alpha_i - \Sigma_i b_i^T(y) + \frac{\delta}{1-\delta}(E(y | e_1 \ldots e_n) - \Sigma w_i) \geq -\Sigma_i \alpha_i, \qquad \text{(EP)}$$

where δ is a common discount factor. This condition can be seen as an enforcement constraint for the principal. The left-hand side (LHS) of the inequality shows the principal's expected present value from honoring the contract, which involves paying out the promised bonuses and then receiving the expected value from relational contracting in all future periods. The right-hand side (RHS) shows the expected present value from reneging, which implies breaking up the relational contract and receiving the reservation value (zero) in all future periods.

Agent i will accept the bonus offered if

$$\alpha_i + b_i^T(y) + \frac{\delta}{1-\delta}(w_i - c(e_i)) \geq \alpha_i, \qquad \text{(EA)}$$

where this can similarly be seen as an enforcement constraint for agent i. The LHS shows the agent's expected present value from honoring the contract, whereas the RHS shows the expected present value from reneging.

Following established procedures (e.g., Levin 2002), we have the following.

Lemma 1. *For given efforts $e = (e_1 \ldots e_n)$, there is a wage scheme that satisfies (IC), (EP), and (EA) and hence, implements e if and only if there are bonuses $b_i^T(y)$ and fixed salaries α_i with $b_i^T(y) \geq 0$, $i = 1, \ldots, n$, such that (IC) and condition (EC) below hold:*

$$\Sigma_i b_i^T(y) \leq \frac{\delta}{1-\delta} \Sigma_i W(e_i). \qquad \text{(EC)}$$

The lemma implies that the enforcement constraints (EP) and (EA) for the principal and the agents, respectively, can be replaced by the aggregate enforcement constraint (EC). To see sufficiency, set the fixed wages α_i such that each agent's payoff in the contract equals his reservation payoff (i.e., $\alpha_i + E(b_i^T(y)|e) - c(e_i) = 0$). Then, (EA) holds, because $b_i^T(y) \geq 0$. Moreover, the principal's payoff in the contract will be $\Sigma_i W(e_i)$ (i.e., the surplus generated by the contract). Then, (EC) implies that (EP) holds. Necessity follows by standard arguments.

Unless otherwise explicitly noted, we will follow the standard assumption in the literature, and we will assume that the first-order approach (FOA) is valid and hence, that each agent's optimal effort choice is given by the first-order condition (FOC)

$$\frac{\partial}{\partial e_i} E(b_i^T(y)|e_1 \ldots e_n) = c'(e_i). \quad (1)$$

Given that FOA is valid, the agents' optimal choices are characterized by the condition (1), which we will refer to as a "modified" (IC) constraint. We will further assume that the "monotone likelihood ratio property" (MLRP) holds for aggregate output y in the following sense: its density is assumed to be of the form $g(y; l(e_1 \ldots e_n))$ with $l_{e_i}(e_1 \ldots e_n) > 0$ and such that $\frac{g_l(y,l)}{g(y,l)}$ is increasing in y.

The optimal contract now maximizes total surplus $(\Sigma_i W(e_i) = \Sigma_i(E(x_i|e_i) - c(e_i)))$ subject to (EC) and the modified (IC) constraint (1). Then, we have the following.

Proposition 1. *The optimal symmetric scheme pays a maximal bonus to each agent for output above a threshold ($y > y_0$) and no bonus otherwise. The threshold is given by $\frac{g_l(y_0,l(e))}{g(y_0,l(e))} = 0$. For $l(e_1 \ldots e_n) = \Sigma_i e_i$, no asymmetric scheme can be optimal.*

The maximal symmetric bonus is, by (EC), $b_i^T(y) = b^T(y) = \frac{\delta}{1-\delta} W(e_i)$ when efforts e_i are equal for all i. This result parallels that of Levin (2003) for the single-agent case. The threshold property comes from the fact that incentives should be maximal (minimal) where the likelihood ratio is positive (negative). Because this ratio is monotone increasing, there is a threshold y_0 where it shifts from being negative to positive, and hence, incentives should optimally shift from being minimal to maximal at that point.[7]

2.2. Team Size, Correlation, and Efficiency

We will here study how team size and correlations among individual contributions to team output affect efficiency. To see how efficiency is affected, note from Proposition 1 that the (IC) constraint (1) can now be written

$$c'(e_i) = b^T \frac{\partial}{\partial e_i} \Pr(y > y_0 | e_1 \ldots e_n),$$

where the term on the RHS is agent i's marginal revenue from effort. The latter is determined by the bonus and the extent to which higher effort for the agent affects the probability of obtaining the bonus. The strength of this effect is, for given efforts by the other agents, determined by the distribution of the aggregate team output y, and it is given by

$$\frac{\partial}{\partial e_i} \Pr(y > y_0 | e_1 \ldots e_n) = \int_{y > y_0} g_i(y; e_1 \ldots e_n) dy, \quad (2)$$

where g_i denotes the partial derivative of the density with respect to e_i. The optimal solution $e_i = e_i^*$ (the maximal effort per agent that can be implemented) is thus given by

$$\frac{c'(e_i^*)}{\int_{y > y_0} g_i(y; e^*) dy} = b^T = \frac{\delta}{1-\delta} W(e_i^*). \quad (3)$$

The first equality shows the required bonus (per agent) to implement effort e_i^* (from the (IC) constraint). The second equality shows the feasible (maximal) bonus.

When team size n increases, a single agent's marginal influence on his expected bonus payment will be affected. In the case of independent outputs, this marginal influence is reduced. Hence, for a fixed bonus (to each agent), every agent will provide less effort. This outcome is similar to the classical free-rider problem, but the mechanism is different. It is not that a given bonus has to be divided between more agents, it is rather that each agent's influence on the team's probability of reaching the bonus threshold becomes lower. Moreover, this lower effort will in turn reduce the surplus and hence, lower the maximal feasible bonus. This will further reduce effort; thus, it is clear that equilibrium effort will be reduced in such a case.

However, with correlated outputs, more agents in the team will not necessarily reduce effort. To see this, we first analyze the effects of correlated outputs for fixed team size, and then, we consider the effects of varying the number of agents.

Regarding stochastic dependencies (correlations) among individual outputs, we see that their effects on efficiency will be determined by their effects on the marginal probability of obtaining the bonus (2). To make the analysis tractable, we will assume that outputs are (multi-)normally distributed and correlated. Given this assumption and (by symmetry) with each x_i being $N(e_i, s^2)$, then total output $y = \Sigma x_i$ is also normal with expectation $Ey = \Sigma e_i$ and variance

$$s_n^2 = var(y) = \Sigma_i var(x_i) + \Sigma_{i \neq j} cov(x_i, x_j)$$
$$= ns^2 + s^2 \Sigma_{i \neq j} corr(x_i, x_j).$$

It follows from the form of the normal density that the likelihood ratio is linear and given by $\frac{g_i(y,e_1 \ldots e_n)}{g(y,e_1 \ldots e_n)} = (y - \Sigma e_i)/s_n^2$. As shown above, the optimal bonus is maximal (minimal) for outcomes where the likelihood ratio is positive (negative), and hence, it has a threshold $y_0 = \Sigma e_i^*$ in equilibrium. Applying the normal distribution, it then follows (as shown in the online appendix) that the marginal return to effort for each agent in equilibrium is given by

$$b^T \int_{y>y_0} g_i(y; e^*) dy = b^T/(Ms_n), \quad M = \sqrt{2\pi}. \tag{4}$$

The marginal return to effort is thus inversely proportional to the standard deviation of total output in this setting. This implies that a team composition that reduces this standard deviation and thus, increases the precision of the available performance measure (total output) will improve incentives and be beneficial here.[8]

An intuition for this is the following. In equilibrium, the team members obtain a bonus when team output y exceeds a hurdle set at the expected output, an event that occurs with probability 1/2 in this setting. For a fixed bonus scheme and thus, a fixed hurdle, additional effort by an individual team member will move the mean of the y distribution and by that increase, the probability of obtaining the bonus. The effect on this probability is stronger the more narrow the distribution (i.e., the lower the variance of team output). Thus, the lower this variance is, the stronger the marginal effect of more individual effort is on the probability to obtain the bonus and thus, the stronger the MIs are for effort.

The (IC) condition (1) for each agent's (symmetric) equilibrium effort is now $c'(e_i) = b^T/(Ms_n)$. It then follows from (3) that the maximal effort per agent that can be sustained is given by

$$c'(e_i^*)s_n M = b^T = \frac{\delta}{1-\delta} W(e_i^*). \tag{5}$$

When all agents' outputs are fully symmetric in the sense that all correlations as well as all variances are equal across agents (i.e., $var(x_i) = s^2$ and $corr(x_i, x_j) = \rho$ for all i, j), then the variance in total output will be

$$s_n^2 = ns^2 + s^2 \Sigma_{i \neq j} corr(x_i, x_j) = ns^2(1 + \rho(n-1)).$$

For fixed team size ($n \geq 2$), the variance of total output will increase with increasing correlation (ρ). This will then be detrimental for individual incentives, because the marginal return to effort is (for a fixed bonus) inversely proportional to the standard deviation of output. Individual efforts must then be reduced in equilibrium, and increased correlation thus reduces efficiency for the team.

Practitioners and empirical researchers may be interested in how the threshold in the bonus scheme varies with correlation. Recall that this threshold is (for the case of normally distributed output) given by $y_0 = \Sigma_i e_i^*$ (i.e., the optimal threshold is specified as the equilibrium expected value of total output). The scheme thus awards each agent a bonus if the team's output realization is higher than expected. Because increased correlation reduces equilibrium efforts, it consequently reduces the threshold for the bonus scheme as well. Increased correlation will thus lead to a lower and hence, less demanding threshold in the bonus scheme.

The detrimental effect of increased correlation is caused by the second best nature of the relational bonus contract. If team output was verifiable, the first best could be implemented when all parties are risk neutral, and the principal can function as a budget breaker. Stochastic dependencies would then play no role. If team output is verifiable and the agents are risk averse, the first best cannot generally be achieved, and increased correlation may again be detrimental through its effect on the agents' exposure to risk. This would be the case (e.g., in a setting like that of Holmström and Milgrom 1990), where increased correlation will increase the variance of the performance measure (y); in turn, this would increase the risk costs of providing incentives and lead to reduced incentives and efforts in equilibrium. The detrimental effect in our setting does not operate via risk costs (because there are none) but exclusively via lower-powered incentives for effort. More risk does not lead to higher costs of providing incentives but rather, leads to a lower incentive effect from a given bonus. The principal cannot compensate for this negative effect by providing higher monetary incentives, because such bonus payments are bounded by the self-enforcement constraint.

Consider now a variation in team size. If $\rho \geq 0$, the variance will increase with n, and this will be detrimental for efficiency.[9] Optimal size n should, therefore, be smaller with larger ρ. Moreover, the standard deviation of total output (s_n) increases rapidly with n when $\rho \geq 0$ (at least of order $\sqrt{n}$); hence, the effort per agent that can be sustained will then decrease rapidly with n . Large teams are, therefore, very inefficient if all agents' outputs are nonnegatively correlated.

For negative correlations, the situation is quite different. If $\rho < 0$, one can in principle reduce the variance to (almost) zero by including sufficiently many agents. The model then indicates that adding more agents to the team is beneficial, at least as long as $1 + \rho(n-1) > 0$ and the conditions for FOA to be valid are fulfilled. As shown by Hwang (2016), this is the case as long as the variance of the performance measure (here s_n^2) is not too small. Adding agents is not beneficial because of any technological complementarities—because there are none by assumption—but it is beneficial, because adding agents provides a more precise performance measure; this, in turn, improves individual incentives.

Note that assuming symmetric pairwise negative correlations among n stochastic variables only makes sense if the sum has nonnegative variance and hence, $1 + \rho(n-1) \geq 0$.[10] Given $\rho < 0$, there can thus only be a maximum number n of such variables (agents). Also, given $n > 2$, we must have $\rho > -\frac{1}{n-1}$.

Note also that, for given negative $\rho > -\frac{1}{2}$, the variance is first increasing and then decreasing in n (it is maximal for $n = \frac{1}{2}(1 - \frac{1}{\rho})$). Hence, the optimal team size in this setting is either very small ($n = 2$) or "very large" (includes all of the relevant agents).

Proposition 2. *For normally distributed outputs and symmetric agents and for a fixed team size (n), efficiency decreases with increasing correlation (ρ) among outputs. For fixed correlation, efficiency decreases with team size if outputs are nonnegatively correlated. For negatively correlated outputs, efficiency first decreases (for $n > 2$) and then increases with increasing team size.*

The intuition can be summarized as follows. The agent's marginal influence on expected bonus payments (i.e., his or her marginal benefit of effort) depends essentially on the variance of the performance measure s_n^2. Adding one more agent then has two effects. First, it increases the variance of total output (even if there is no correlation in output). This means that team output becomes more "spread out," and hence, each agent's marginal benefit of effort is reduced. Second, the correlation in output (induced by the extra agent) again affects the variance. Positive correlation increases the variance and hence, reduces the marginal benefit of effort. Negative correlation, in contrast, decreases the variance and therefore, increases the marginal benefit of effort. When the correlation is positive, both effects go the same way and reduce the marginal benefit of effort. When the correlation is negative, these two effects go the opposite way, and the effect from negative correlation dominates for large-enough n, implying that the effect is U shaped in the number of agents.

The assumption of equal pairwise correlations among all involved agents is somewhat special, but it illustrates in a simple way the forces at play when the team size varies. In reality, correlations among agents may vary; there might, for example, be positive correlations among some agents and negative correlations among others. Such features may in fact be straightforwardly incorporated in our team model (i.e., we may allow correlation coefficients to vary across agents). This follows from Proposition 1, which justifies that a symmetric bonus scheme is optimal, and the ensuing analysis leading to the equilibrium condition (5). The only required modification is that the output variance should then be given by the general expression $s_n^2 = ns^2 + s^2\Sigma_{i\neq j}\rho_{ij}$, where the correlation coefficients ρ_{ij} may vary across agent pairs. A procedure to pick agents to obtain the most precise performance measure would then be, for each n, to pick those n agents that yield the smallest variance for the team's output.

Remark (On Multitasking).[11] We finally note that the model in this section can alternatively be interpreted as a model of a single agent with n tasks, where task i yields an unobservable contribution x_i to aggregate output $y = \Sigma_i x_i$. Only y is observable and can be a basis for effort incentives. These incentives and the resulting equilibrium will then be identical to those derived for the team setting when the agent's cost structure is additive (of the form $\Sigma_i c(e_i)$) and thus, has no interaction effects among efforts. In this setting, the model predicts that the agent's efficiency is decreasing in the level of correlation (ρ) among tasks and that adding more tasks will be beneficial only when this correlation is negative.

3. Observable Individual Outputs

Consider now the case where individual outputs are observable but still nonverifiable. The principal can then offer a bonus contract $b_i^I(x_1 \dots x_n)$ to each agent $i = 1 \dots n$ conditional on all individual outputs. Now, if the contract is expected to be honored, agent i's expected wage is then, for given efforts, $w_i = E(b_i^I(x_1 \dots x_n) | e_1 \dots e_n) + \alpha_i^I$, whereas the principal expects $\Sigma\bar{x}(e_i) - \Sigma w_i$. The agent then chooses effort

$$e_i = \arg\max_{e_i'} \left(E(b_i^I(x_1 \dots x_n) | e_i', e_{-i}) - c(e_i')\right). \tag{6}$$

In the repeated relationship, we still assume that the principal honors the contract only if all agents honored the contract in the previous period and that the agents honor the contract only if the principal honored the contract with all agents in the previous period.

Now (given that the (IC) condition (6) holds), the principal will honor the contract with all agents $i = 1, 2, \dots, n$ if

$$-\Sigma_i b_i^I(x_1 \dots x_n) + \frac{\delta}{1-\delta}(\Sigma_i E(x_i | e_i) - \Sigma_i w_i) \geq 0. \tag{7}$$

Agent i will accept the bonus offered if

$$b_i^I(x_1 \dots x_n) + \frac{\delta}{1-\delta}(w_i - c(e_i)) \geq 0. \tag{8}$$

It is now straightforward to show (as in the previous case, where only $y = \Sigma_i x_i$ is observed) that we have the following lemma.

Lemma 2. *For given efforts $e = (e_1 \dots e_n)$, there is a wage scheme that satisfies (6), (7), and (8) and, hence, implements e if and only if there are bonuses $b_i^I(x_1 \dots x_n)$ and fixed salaries α_i^I with $b_i^I(x_1 \dots x_n) \geq 0$, $i = 1 \dots n$, such that (6) and condition (9) below hold:*

$$\Sigma_i b_i^I(x_1 \dots x_n) \leq \frac{\delta}{1-\delta}\Sigma_i W(e_i). \tag{9}$$

Here, $W()$ denotes (as before) surplus per agent: $W(e_i) = E(x_i | e_i) - c(e_i)$. Assuming that the FOA is valid, we can replace the (IC) constraint (6) with the FOC

$$\frac{\partial}{\partial e_i} E(b_i^I(x_1 \dots x_n) | e_1 \dots e_n) = c'(e_i). \tag{10}$$

The optimal contract then maximizes total surplus ($\Sigma_i W(e_i)$) subject to (9) and (10). All results in the following assume that the FOA is valid.

3.1. Independent Outputs

Consider first independent outputs. These were analyzed by Levin (2002), who showed that the optimal contract entails RPE with a bonus paid to at most one agent, namely the agent whose outcome yields the highest likelihood ratio. Moreover, the bonus is paid to this agent only if the likelihood ratio is positive. Given symmetric agents and strictly increasing likelihood ratios, this means that the agent with the largest output wins the bonus, provided that his output exceeds some threshold x_0 (where the likelihood ratio is positive for $x_i > x_0$).

The intuition for this result is that, because the bonus pool is bounded because of the enforcement constraint (9) and because the agents are not averse to risk, the optimal scheme entails maximizing individual incentives by letting the agents compete for a single bonus. We will now use this result to analyze how the efficiency of this scheme varies with the number of agents (for independent outputs). The next section considers correlated outputs.

With n agents, agent i's probability of winning the bonus b^I, given own output $x_i = x > x_0$ and given symmetric efforts e_j from all others, is now $\Pr(\max_j x_j < x) = F(x; e_j)^{n-1}$. Hence, the expected bonus payment to agent i is $b^I \int_{x_0}^{\infty} F(x_i; e_j)^{n-1} f(x_i; e_i) dx_i$, and for symmetric efforts, the (IC) condition (10) takes the form

$$b^I \int_{x_0}^{\infty} F(x_i; e_i)^{n-1} f_{e_i}(x_i; e_i) dx_i = c'(e_i). \qquad (11)$$

In passing, it is worth noting that the integral here extends only over values of x_i, where $f_{e_i}(x_i; e_i) > 0$. In a standard tournament, where agent i would obtain a bonus when he had the largest output, the integral would extend over all values of x_i. The payment scheme here, which we may call a modified tournament, thus provides stronger incentives (for a given bonus b^I) than a standard tournament scheme.

The optimal RPE bonus is maximal (i.e., $b^I = \frac{\delta}{1-\delta} \cdot \Sigma_i W(e_i)$), where $W(e_i)$ is total surplus (for agent i). Hence, from (11), we have, in symmetric equilibrium,

$$\frac{c'(e_i)}{\int_{x_0}^{\infty} F(x; e_i)^{n-1} f_{e_i}(x; e_i) dx} = b^I = \frac{\delta}{1-\delta} n W(e_i). \qquad (12)$$

Consider now variations in the number of agents. Higher n increases the competition to obtain the bonus (the probability of winning is reduced), and therefore, the bonus must be increased to maintain effort; this is captured by the first equality in (12). The second equality shows how much the bonus can be increased, namely by the increased total surplus. The question is then whether the latter is sufficient to compensate for the reduced probability of winning.

The answer is affirmative, and the reason is essentially that, although the surplus on the RHS increases proportionally with n, the marginal probability (in the denominator) on the LHS decreases less rapidly. This allows a higher effort per agent to be implemented, and therefore, we have Proposition 3.

Proposition 3. *For observable and independent individual outputs, effort per agent in the RPE scheme (the modified tournament) increases with the number of agents.*

When individual output measures are available and these outputs are independent, we thus see that efficiency in the (modified) tournament is improved by including more agents. This is in sharp contrast to efficiency in a team for independent outputs: as we saw above, the team efficiency rapidly decreases under such conditions. The reason for the difference is as follows. Under both team and tournament incentives, the MI effect from a given bonus is reduced when adding more agents. However, under a tournament scheme, only one agent (at most) is awarded the bonus, and hence, the firm can increase the bonus without violating the self-enforcement constraint if more agents are included. In contrast, team bonuses are awarded to all of the agents, and hence, the firm cannot compensate for the lower incentive effect by increasing the bonus.

3.2. Stochastically Dependent Outputs

Consider now stochastically dependent outputs. As before, we limit attention to symmetric agents and thus, symmetric efforts in equilibrium. The basic insight from Levin (2002) that at most one agent should be rewarded a bonus extends to this environment; thus, a type of modified tournament is still optimal. However, the tournament is not necessarily one based on raw outputs x but rather, one based on what we may call indexes: one for each agent. The relevant index for agent i is the likelihood ratio

$$l_i(x; e_1 \dots e_n) = \frac{f_{e_i}(x | e_1 \dots e_n)}{f(x | e_1 \dots e_n)} \qquad (13)$$

evaluated at (symmetric) equilibrium efforts. Denote this as $l_i^*(x)$; thus, $l_i^*(x) = l_i(x; e_1 \dots e_n)$ with $e_i = e^*$, the common equilibrium effort, for all i.[12] As we show in the online appendix, the optimal symmetric scheme pays a maximal bonus to the agent with the highest such likelihood ratio, provided that this ratio is positive, and no bonus to the other agents.

Lemma 3. *There are indexes $l_i^*(x_1, \dots x_n)$, $i = 1 \dots n$, one for each agent and given by the respective likelihood ratios, such that the optimal symmetric scheme pays a single and maximal bonus to the agent with the highest index value, provided that this value is positive.*

For a given vector of output realizations x, the agents are thus compared in terms of the indexes $l_i^*(x)$, and the agent with the highest index value is awarded the bonus, provided that this value exceeds a threshold, which here is zero. The index for agent i will generally depend on the whole vector x of individual output realizations. The special feature of stochastically independent outputs is that each agent's index depends only on his own realization in that case.

It is again instructive to consider the *multinormal distribution*, in particular because the indexes then take a simple form. A very convenient feature of this distribution is that likelihood ratios—and therefore, the relevant indexes—are linear functions of the variables, and this considerably simplifies comparisons of these entities.

Therefore, assume now that $x = (x_1 \ldots x_n)$ multinormal with $Ex_i = e_i$, $var(x_i) = s^2$ and (identical) correlations[13] $corr(x_i, x_j) = \rho$. From the form of the multinormal distribution (see the online appendix), the likelihood ratio for agent i is now

$$l_i(x; e_1 \ldots e_n) = k_1 \cdot (x_i - e_i) + k_2 \Sigma_{i \neq j}(x_j - e_j), \qquad (14)$$

where $k_i = k_i(n, \rho, s^2)$, $i = 1, 2$ are coefficients with $k_1 > 0$, $k_1 > k_2$.

The indexes $l_i^*(x)$ for the tournament in Lemma 3 are thus linear functions of the output realizations. Moreover, from symmetry (including symmetric efforts in equilibrium; $e_i = e^*$ for all i), we see that $l_i^*(x) - l_j^*(x) = (k_1 - k_2)(x_i - x_j)$, which implies that the agent with the highest output will here also have the highest index value. This agent will thus win the tournament and obtain the bonus if the index value is positive. As shown in the online appendix, the index value is positive if and only if

$$x_i > e^* + \frac{\rho}{(n-2)\rho + 1} \Sigma_{j \neq i}(x_j - e^*) = E(x_i | x_{-i}). \qquad (15)$$

This condition says that agent i's performance must exceed his expected performance, conditional on the performance of all other agents. Thus, we have the following proposition.

Proposition 4. *For normally distributed outputs, the optimal symmetric scheme pays a maximal bonus to the agent (say i) with the highest output, provided that this output satisfies* $x_i > E(x_i | x_{-i})$.

For $n = 2$ agents, we now have that agent 1 gets the bonus if and only if he has the highest output ($x_1 > x_2$) and $x_1 - e^* > \rho(x_2 - e^*)$. This is illustrated in Figure 1 for $\rho = \frac{1}{2}$ (Figure 1(a)) and $\rho = -\frac{1}{2}$ (Figure 1(b)). Agent 1 is to get the bonus for outcomes in the shaded region in Figure 1 (to the right of the broken lines in Figure 1).

In both cases, the agent with the highest output gets the bonus if both of them have outputs that are above average ($x_1, x_2 > Ex_i = e^*$). If agent 2 has below-average output ($x_2 < Ex_i = e^*$), the requirement for agent 1 to get the bonus is less strict when there is positive correlation than when there is negative correlation. In the latter case, agent 1 must have an output well above average to obtain the bonus, and more so, the worse is the output for agent 2. Under negative (positive) correlation, a bad performance by agent 2 raises (lowers) the expected *conditional* performance of agent 1 and thus, raises (lowers) the requirement—the hurdle (threshold)—for agent1 to get the bonus.[14]

Having characterized the optimal scheme, we will now consider its incentive properties. To make the analysis tractable, we restrict attention to $n = 2$ agents. Consider then agent 1's incentives in this scheme, with "reference point" (equilibrium) $e_1^* = e_2^*$. His probability of obtaining the bonus is

$$\Pr(x_1 > \max[x_2, e_1^* + \rho(x_2 - e_2^*)] | e_1, e_2^*) \equiv \Pr(B) = \int_{x \in B} f(x | e_1, e_2^*). \qquad (16)$$

Figure 1. Bonus Regions for Correlation Coefficients 0.5 (Panel A) and −0.5 (Panel B)

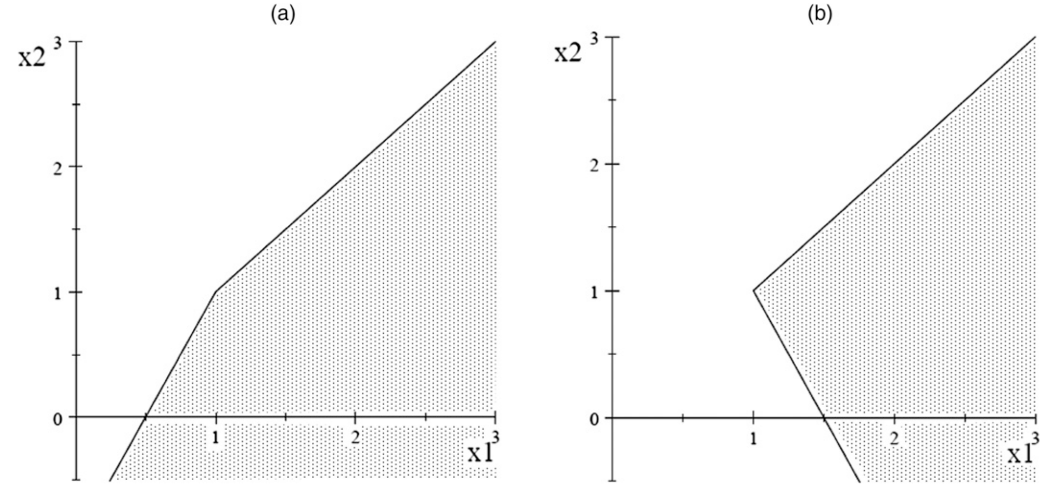

Therefore, the marginal gain from effort is $\int_B f_{e_1}(x|e_1, e_2^*)$, and in symmetric equilibrium $e_1^* = e_2^* = e^*$, we will then have

$$b^I \int_B f_{e_i}(x|e^*, e^*) - c'(e^*) = 0.$$

An interesting question is then as follows. For given effort e^* to be implemented, how do MIs vary with correlation ρ? For example, do these MIs become stronger when ρ increases, implying that a lower bonus is required to implement the same effort? It is well known that a standard tournament scheme performs well for positive correlation but poorly for negative correlation, and this might indicate that similar features should be present here. However, the optimal scheme here is not a standard tournament; it is modified by a relative performance element associated with the hurdle that must be passed to win the bonus. Because the hurdle is related to relative performance, the optimal scheme is thus an RPE scheme, and we also know that such schemes generally work well for both positive and negative correlations in other settings.[15] It turns out that the latter property also holds here.

Proposition 5. *For normally distributed variables and $n = 2$, the agent's FOC for (symmetric) equilibrium effort is*

$$b^I \frac{1}{\sqrt{2\pi}s} \frac{1}{2} \left(\frac{1}{\sqrt{1-\rho^2}} + \frac{1}{\sqrt{1-\rho}} \frac{1}{\sqrt{2}} \right) = c'(e^*). \quad (17)$$

The MI (i.e., the expression on the LHS) is increasing in the correlation coefficient ρ for $\rho > \rho_0 \approx -0.236$ and decreasing in ρ for $\rho < \rho_0$. Hence, implementing a given effort requires a lower (higher) bonus when the correlation ρ increases for $\rho > \rho_0$ (for $\rho < \rho_0$).[16]

This is illustrated in Figure 2, which depicts the MI as a function of ρ for the RPE scheme and a standard tournament (Figure 2, dashed line).

As a function of ρ, the MI for effort is thus U shaped in the optimal scheme, which again is a modified tournament. In comparison, in a standard tournament, the MI is monotone increasing in ρ (as shown by Figure 2, dotted line; this MI is given by $\frac{d}{de_1}\Pr(x_1 > x_2) = \frac{1}{\sqrt{2\pi}s_d}$, where $s_d = \sqrt{2(1-\rho)}s$ is the standard deviation of $x_1 - x_2$, and the formula follows from the normal distribution). In comparison the modified tournament yields higher MI for effort for every ρ (which allows a higher effort to be implemented with the same bonus), and the MI is high for both strongly positive correlated and strongly negative correlated outputs.

Figure 2. MIs as Functions of ρ

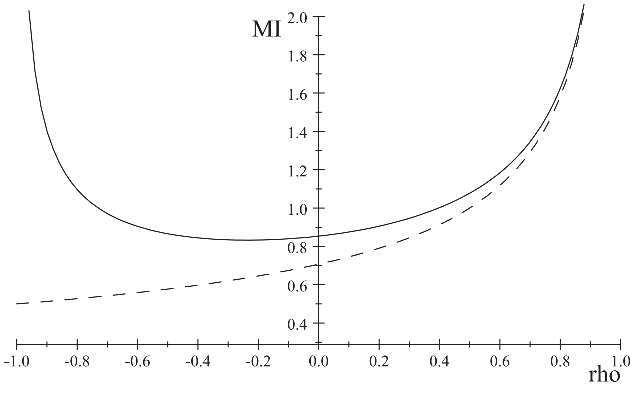

The latter property is caused by the specific criteria to obtain the bonus in the modified tournament as illustrated in the figures. In a standard tournament, agent 1 wins and gets a bonus if $x_1 > x_2$, whereas in the modified tournament, he gets a bonus only if $x_1 > x_2$ and $x_1 - e^* > \rho(x_2 - e^*)$. Therefore, the probability of obtaining the bonus is (all else equal) higher in a standard tournament, but the marginal effect of own effort on the probability (the MI) is higher in the modified tournament.

The last proposition shows that, as a function of correlation ρ, the MIs in the modified tournament are minimal not for $\rho = 0$ but for $\rho_0 < 0$. This can be explained by taking into account the two aspects of this incentive scheme: the pure tournament aspect (largest output wins) and the hurdle aspect. The former yields MIs that are monotone increasing in ρ, such as just illustrated, but the second yields (in isolation) MIs[17] that are increasing in ρ^2 and therefore, symmetric around $\rho = 0$. The combination of the two effects thus yields incentives that are increasing at $\rho = 0$ and minimal for some negative ρ.

4. Concluding Remarks

We investigate how stochastic dependencies between employees affect optimal incentive schemes in situations where performance measures are nonverifiable. We show that the way that the employees' outputs correlate is an important determinant for the efficiency of team-based incentives and the efficiency and design of tournament rewards.

With respect to teams, we derive testable theoretical predictions on team size and team composition. We do so by analyzing optimal self-enforcing (relational) contracts between a principal and a set of agents, where only aggregate output can be observed, and we show that the principal can use team size and team composition as instruments to improve incentives. In particular, the principal can strengthen the agents' incentives by composing teams that utilize stochastic dependencies between the agents' outputs.

Our model predicts that teams are more efficient when the team members' outputs are negatively correlated. This relates to questions concerning optimal team composition. A central question is whether teams should be homogenous or heterogeneous with respect to tasks (functional expertise, education, and organizational tenure) as well as biodemographic characteristics (age, gender, and ethnicity). One can conjecture that negative correlations are more associated with heterogeneous teams than homogenous teams and also more associated with task-related diversity than with biodemographic diversity. Interestingly, a comprehensive metastudy by Horwitz and Horwitz (2007) finds no relationship between biodemographic diversity and performance but a strong positive relationship between team performance and task-related diversity. An explanation is that task-related diversity can both reduce risk and create positive complementarity effects. We point to an alternative explanation, namely that diversity may create negative correlations that reduce variance and thereby, increase MIs for effort. The team members "must step forward when others fail." Diversity and heterogeneity among team members can thus yield considerable efficiency improvements.[18]

Team incentives are generally not optimal when individual outputs are observable. For a parametric (normal) distribution, we have shown that the optimal relational contract is then an RPE scheme: a form of a tournament where the conditions for an agent to obtain the (single) bonus are stricter for negatively compared with positively correlated outputs. The efficiency of the RPE contract is shown to increase with the number of agents and improve with stronger correlation, both positive and negative.

As a final remark, it should be noted that our main results are shown in the parametric setting of a normal distribution and that some properties derived in this setting may well not be generally valid. Among other things, a convenient feature of the normal distribution is that aggregate team output satisfies the MLRP, and this leads to a simple structure of optimal bonus schemes. However, even if some features of our model are not generally valid, it is a general fact that stochastic dependencies do affect performance measures and by that, incentives and efficiency in relational contracts. Our model forcefully illustrates this point and provides interesting and testable implications for settings where normal distributions can be taken as a reasonable assumption.

Acknowledgments
The authors thank the editor and referees for valuable comments and suggestions. They also thank Eirik Kristiansen, Steve Tadelis, Joel Watson, and seminar participants at the Norwegian School of Economics; University of California, Berkeley; University of California, San Diego; and University of Konstanz for comments and suggestions on earlier versions of this paper.

Endnotes

[1] Although a team's aggregate output may be easier to verify than individual outputs, there is still a range of situations in which a team's output is nonverifiable. Teams are also, like individuals, exposed to discretionary bonuses and subjective performance evaluation, which by definition, cannot be externally enforced.

[2] In our model, negatively correlated signals reduce the moral hazard problem even with risk-neutral agents. In this respect, our finding relates to insights from Diamond (1984), who showed that correlated signals may reduce output variance and thus, reduce entrepreneurs' moral hazard opportunities toward investors.

[3] Economists studying teams, beginning with Alchian and Demsetz (1972), have mainly focused on the free-rider problem, in particular the conditions under which the first best outcome will be achieved (Holmström 1982, Rasmusen 1987, Legros and Matthews 1993). If individual signals are observable, the literature has also shown how the principal can foster cooperation (Holmström and Milgrom 1990; Itoh 1991, 1992, 1993; Macho-Stadler and Perez-Castrillo 1993) or exploit peer effects (Kandel and Lazear 1992, Arya et al. 1997, Che and Yoo 2001).

[4] Although we focus on the multiagent case, our paper is indebted to the seminal literature on *bilateral* relational contracts. The concept of relational contracts was first defined and explored by legal scholars (Macaulay 1963; Macneil 1974, 1978), whereas the formal literature started with Klein and Leffler (1981), Shapiro and Stiglitz (1984), and Bull (1987). MacLeod and Malcomson (1989) provide a general treatment of the symmetric information case, whereas Levin (2003) generalizes the case of asymmetric information.

[5] We thus assume stationary contracts, which have been shown to be optimal in settings like this (Levin 2002, 2003).

[6] See Miller and Watson (2013) for alternative strategies and "disagreement play" in repeated games.

[7] The assumption that total team output y has a distribution that satisfies MLRP is not entirely innocuous. It holds true for the case of normally distributed individual contributions (x_i) assumed below, but it may not hold true for other cases. If MLRP does not hold for output y, then the optimal bonus scheme may be nonmonotonous; thus, it pay a bonus for, for example, low and high realizations of y but no bonus for intermediate realizations. Although theoretically feasible, such schemes are rarely observed in reality.

[8] There is, however, a caveat, because the FOA is valid only if the standard deviation (s_n) is not too small (e.g., Kvaløy and Olsen 2014, Hwang 2016). For numerical indications, these papers show that, for isoelastic effort costs ($c(e) = ke^m,\ m \geq 2$), the FOA is valid for parameters such that $e_i^*/s_n < k_0\sqrt{m-1}$, with $k_0 \cong 2.2$. For sufficiently small s_n, the FOA is not valid, and the analysis must be modified. It turns out that the optimal bonus scheme is still a hurdle scheme and that lower variance also improves incentives and increases efforts in that case (Chi and Olsen 2018).

[9] For verifiable team output, the higher variance induced by increased size n will have similar effects as those just discussed for increased correlation and fixed n.

[10] Indeed, $1 + \rho(n-1) > 0$ is the condition for the covariance matrix to be positively definite and hence, the multinormal model to be well specified.

[11] We thank a referee for suggesting this interpretation of the model.

[12] In this section, it is convenient to let e^* be a scalar and denote the symmetric equilibrium effort level.

[13] To guarantee full symmetry among agents, we consider here only the case where all pairwise correlations are identical.

[14] To illustrate these points, if $\rho = 0.5$ and agent 2 has output 10% below expected ($x_2/e^* = 0.9$), agent 1 can only win if his output is no more than 5% below expected. However, if $\rho = -0.5$, agent 1 must perform at least 5% better than expected to be eligible for the bonus (if in addition, he wins).

[15] Fleckinger (2012) provides a general treatment of stochastic dependencies and RPE for verifiable outputs and shows that greater correlation in outcomes does not necessarily call for RPE schemes.

[16] This rests on the FOA being valid. It can be verified analytically that this will be the case as long as the variance s^2 is not too small. For the

case of quadratic effort costs, it can be verified numerically that FOA holds if $\frac{s}{e_1^*} \geq K(\rho)$, where $K(\rho)$ is a U-shaped function with $K(-0.75) \simeq 0.737$, $K(0) \simeq 0.514$, and $K(0.75) \simeq 0.742$.

[17] By this, we mean the marginal effect of effort on the probability of $x_1 - e^* > \rho(x_2 - e^*)$, which can be seen to be increasing in ρ^2.

[18] Hamilton et al. (2003) provide one of the very few empirical studies on teams within the economics literature. They find that more heterogeneous teams (with respect to ability) are more productive (average ability held constant).

References

Alchian AA, Demsetz H (1972) Production, information costs, and economic organization. *Amer. Econom. Rev.* 62(5):777–795.

Arya A, Fellingham J, Glover J (1997) Teams, repeated tasks, and implicit incentives. *J. Accounting Econom.* 23(1):7–30.

Baldenius T, Glover J, Xue H (2016) Relational contracts with and between agents. *J. Accounting Econom.* 61(2-3):369–390.

Bewley T (1999) *Why Wages Don't Fall During a Recession* (Harvard University Press, Cambridge, MA).

Bull C (1987) The existence of self-enforcing implicit contracts. *Quart. J. Econom.* 102(1):147–159.

Che Y-K, Yoo S-W (2001) Optimal incentives for teams. *Amer. Econom. Rev.* 91(3):525–540.

Chi C-K, Olsen TE (2018) Relational incentive contracts and performance measurement. Discussion Paper No. 2018/6, Department of Business and Management Science, Norwegian School of Economics, Bergen, Norway.

Deb J, Li J, Mukherjee A (2016) Relational contracts with subjective peer evaluations. *RAND J. Econom.* 47(1):3–28.

Diamond DW (1984) Financial intermediation and delegated monitoring. *Rev. Econom. Stud.* 51(3):393–414.

Fleckinger P (2012) Correlation and relative performance evaluation. *J. Econom. Theory* 147(1):93–117.

Gibbs M, Merchant KA, Van der Stede WA, Vargus ME (2004) Determinants and effects of subjectivity in incentives. *Accounting Rev.* 79(2):409–436.

Glover J, Xue H (2018) Team incentives and bonus floors in relational contracts. Research Paper No. 18-43, Columbia Business School, New York.

Hamilton BH, Nickerson JA, Owan H (2003) Team incentives and worker heterogeneity: An empirical analysis of the impact of teams on productivity and participation. *J. Political Econom.* 111(3):465–497.

Holmström B (1982) Moral hazard in teams. *Bell J. Econom.* 13(2):324–340.

Holmström B, Milgrom P (1990) Regulating trade among agents. *J. Inst. Theoret. Econom.* 146(1):85–105.

Horwitz SK, Horwitz IB (2007) The effects of team diversity on team outcomes: A meta-analytic review of team demography. *J. Management* 33(6):987–1015.

Huddart S, Liang PJ (2003) Accounting in partnerships. *Amer. Econom. Rev.* 93(2):410–414.

Huddart S, Liang PJ (2005) Profit sharing and monitoring in partnerships. *J. Accounting Econom.* 40(1–3):153–187.

Hwang S (2016) Relational contracts and the first-order approach. *J. Math. Econom.* 63(C):126–130.

Itoh H (1991) Incentives to help in multi-agent situations. *Econometrica* 59(3):611–636.

Itoh H (1992) Cooperation in hierarchical organizations: An incentive perspective. *J. Law Econom. Organ.* 8(2):321–345.

Itoh H (1993) Coalition incentives and risk sharing. *J. Econom. Theory* 60(2):410–427.

Kandel E, Lazear EP (1992) Peer pressure and partnerships. *J. Political Econom.* 100(4):801–817.

Klein B, Leffler K (1981) The role of market forces in assuring contractual performance. *J. Political Econom.* 89(4):615–641.

Kvaløy O, Olsen TE (2006) Team incentives in relational employment contracts. *J. Labor Econom.* 24(1):139–169.

Kvaløy O, Olsen TE (2008) Cooperation in knowledge-intensive firms. *J. Human Capital* 2(4):410–440.

Kvaløy O, Olsen TE (2014) Teams and tournaments in relational contracts. CESifo Working Paper Series No. 4783, CESifo Group, Munich, Germany.

Legros P, Matthews SM (1993) Efficient and nearly-efficient partnerships. *Rev. Econom. Stud.* 60(3):599–611.

Levin J (2002) Multilateral contracting and the employment relationship. *Quart. J. Econom.* 117(3):1075–1103.

Levin J (2003) Relational incentive contracts. *Amer. Econom. Rev.* 93(3):835–857.

Liang PJ, Rajan MV, Ray K (2008) Optimal team size and monitoring in organizations. *Accounting Rev.* 83(3):789–822.

Macaulay S (1963) Non-contractual relations in business: A preliminary study. *Amer. Sociol. Rev.* 28(February):55–67.

Macho-Stadler I, Perez-Catrillo JD (1993) Moral hazard with several agents: The gains from cooperation. *Internat. J. Indust. Organ.* 11(1):73–100.

MacLeod WB, Malcomson J (1989) Implicit contracts, incentive compatibility, and involuntary unemployment. *Econometrica* 57(2):447–480.

Macneil IR (1974) The many features of contracts. *Southern Calif. Law Rev.* 47(May):691–816.

Macneil IR (1978) Contracts: Adjustment of long-term economic relations under classical, neoclassical, and relational contract law. *Northwestern Univ. Law Rev.* 72(6):854–904.

Miller D, Watson J (2013) A theory of disagreement in repeated games with bargaining. *Econometrica* 81(6):2303–2350.

Mookherjee D (1984) Optimal incentive schemes with many agents. *Rev. Econom. Stud.* 51(3):433–446.

Rajan M, Reichelstein S (2006) Subjective performance indicators and discretionary bonus pools. *J. Accounting Res.* 44(3):585–618.

Rasmusen E (1987) Moral hazard in risk-averse teams. *RAND J. Econom.* 18(3):428–435.

Rayo L (2007) Relational incentives and moral hazard in teams. *Rev. Econom. Stud.* 74(3):937–963.

Shapiro C, Stiglitz JE (1984) Equilibrium unemployment as a worker discipline device. *Amer. Econom. Rev.* 74(3):433–444.