WORKING PAPERS

N° 1421

JANUARY 2024

“Should They Compete or Should They Cooperate? The View of Agency Theory”

Pierre Fleckinger, David Martimort and Nicolas Roux

Toulouse School of Economics

Should They Compete or Should They Cooperate? The View of Agency Theory*

Pierre Fleckinger,† David Martimort,‡ Nicolas Roux§

January 19, 2024

Abstract

What is the most efficient way of designing incentives in an organization? Over the past five decades, agency theory has provided various answers to this crucial question. This line of research suggests that, depending on the organizational context, the optimal approach to providing incentives may involve either relying on collective compensations or, conversely, employing relative performance evaluations. In the first scenario, cooperation among agents is the key aspect of the organization. In the second, competition prevails. This paper provides a comprehensive overview of this extensive literature, with the aim of understanding the conditions under which one or the other type of incentive schemes is more desirable for the principal of the organization. To this end, we use a flexible and versatile model capable of addressing a wide range of scenarios characterized by different technologies, information constraints, and behavioral norms.

JEL classification: D20, D86, J33, L23, M12, M50.

Keywords: Incentive contracts, moral hazard, teams, competition vs. cooperation, collusion, free riding, tournaments, peer effects, organizational design.

---

*We are extremely grateful to four referees and the Editors of this journal Steven Durlauf and David Romer for their useful suggestions and comments, which greatly improved the exposition and clarity of this article. We thank Javier Gonzales-Morin and Pol Morvan for excellent research assistance. The second author acknowledges financial support from the ERC (MARKLIN). This paper has also been funded by the Agence Nationale de la Recherche under grant ANR-17-EURE-0010 (Investissements d'Avenir program). The usual disclaimer applies.

†Mines Paris - PSL & Paris School of Economics. Email: pierre.fleckinger@minesparis.psl.eu.

‡Toulouse School of Economics. Email: david.martimort@tse-fr.eu.

§TWS Partners. Email: nicolas.roux@tws-partners.com.

Contents

1 Introduction

What is the most efficient way of designing incentives for a group of agents? Whether it applies to the internal organization of firms or to the more global architecture of markets and societies, the issue of identifying how to reward *'good performance'* is of overwhelming importance. Should incentives be based on collective performance, taking the whole group of agents within the organization as a team? Or, should instead incentives be driven by inducing more competition among those very same agents?

Fostering competition and inducing cooperation are indeed two alternative ways of providing incentives in complex organizations. Competition is frequently employed to enhance incentives. For example, it is a common practice to compensate retailers or salespersons based on the sales they have made, often accompanied by implicit incentives like *'best salesperson of the week'* rewards. Compensation schemes for CEOs may either be explicitly benchmarked against the performances reached by other firms in the same industry, or indirectly compared when CEOs are paid with stocks whose values may be correlated with those of competing firms. The various forms of yardstick competition that have been implemented by regulators in several countries and several regulated sectors throughout the world (for instance among hospitals or public electricity facilities) are also aimed at boosting incentives of their top management to improve services, cut costs and lower prices. In other circumstances, such as movers working as teams or researchers in large-scale research units, individual contributions are much harder to disentangle and the use of collective bonuses certainly becomes more attractive, especially to foster cooperation and tasks coordination. In each of the examples above, incentive schemes, sometimes in conjunction with other elements of organizational design such as job characteristics and task assignments, are tailored to reduce agency costs and efficiently respond to a number of specific features of both the organization and the economic environment in which it operates.

While early contributions emphasized performance comparison and competition as motivators, recent research in agency theory has shifted its focus toward fostering cooperation. To investigate whether incentives should be designed either to promote a more competitive environment or to induce more cooperation within the organization, the extant literature offers a variety of scattered models. This diversity is of course welcome as such, but unfortunately, also makes it particularly difficult for modelers new to the field to isolate the key factors that drive the choice of an incentive system. In other words, what is currently missing is an encompassing unified framework allowing researchers to clearly map the characteristics of the organization they are interested in and the various constraints that pertain to this organization (be they behavioral, technological, institutional or market-induced) with the kind of incentives that should be implemented.[1]

The primary aim of this paper is to fill this gap. To achieve this goal, this essay surveys and unifies contributions from agency theory that have examined the costs and benefits of relying on either competitive or collective incentives within a versatile model.

Our analysis rests on a generic formulation of a standard agency model in the context of complex organizations. The principal can be thought of as representing the interest of the government in a policy-making framework, the shareholders of a firm in a corporate context, or the manager of a nonprofit organization or a public administration service. Agents can be thought of respectively as regulated firms, productive units within a given firm, or employees in a hierarchical structure.

OVERVIEW AND ORGANIZATION. Our simple moral hazard model of a complex multi-agent organization is laid down in Section 2. Agents within an organization exert costly efforts that affect their performances. Effort is non-observable. Incentive schemes are designed to induce agents to exert efforts so as to promote the objectives of the organization. In a complex organization, incentive schemes may be linked across agents. In this respect, a key distinction must be made between compensation schemes based on Relative Performance Evaluation (RPE) and those based on Collective Performance Evaluation (CPE). The former rewards individual agents based on their performance relative to others, while the latter rewards all agents based on the collective performance of the organization, and all agents benefit from other agents' better performance. Sometimes, organizations may also choose to employ Independent Performance Evaluation (IPE), in which case agents operate under bilateral contracts with the principal of the organization. Importantly, RPE or CPE are not needed when the first-best can be achieved with IPE. They become strictly valuable in a second-best world where contracts cannot costlessly overcome moral hazard.

Our analysis uncovers several fundamental principles that inform the choice of compensation schemes. Table 1 summarizes these principles.[2] The rows list the main topics

[1]Social-psychologists have also been interested in this question for a long time. See Deutsch (1949*a*,*b*) for seminal contributions. This subject is obviously a key point for the management literature–see DeMatteo, Eby and Sundstrom (1998) and Cohen and Bailey (1997)–and for the growing literature on personnel economics initiated by Lazear (1995).

[2]Section 4, not referenced in this table, essentially features a game theoretic analysis laying the ground for the subsequent sections of the paper, along with additional material on sequential production processes

that we chose to cover, while the cells contain the determinants of the incentive schemes listed in the columns.

Table 1: Main factors affecting the incentive schemes

	RELATIVE	INDEPENDENT	COLLECTIVE
PRODUCTION Section 3.1	Substitutability	Independence	Complementarity Help and sabotage
INFORMATION Section 3.2	Positive correlation	Independence	Negative correlation
BEHAVIORAL Section 3.5	Status-seeking preferences Self-overconfidence		Inequity aversion Overconfidence in others
COLLUSION Section 5			Side-transfers feasible Mutual monitoring Repeated interactions
PRINCIPAL'S OPPORTUNISM Section 6	Principal's moral hazard Subjective evaluation		
MARKETS Section 7		Career concerns Agents' holdup	

Section 3 discusses the fundamental new phenomena that arise when providing incentives to agents in complex organizations. Although the basic drivers of incentive compensations are similar to those found in much simpler and more familiar bilateral principal-agent relationships, a prime notion that emerges from our analysis is that incentive schemes should exploit the various linkages that may exist between agents. These linkages may be either technological or informational (see the corresponding definitions below) and it is their precise nature that determines how the incentives of a given agent impact those of their peers.

As an example of technological linkages (Section 3.1), consider two retailers who distribute two different products of the same company. If those products are complements, the demand for good 2 might increase if the retailer of good 1 puts more effort into promoting its own product; a scenario with positive production externalities. In contrast, if the products are substitutes, retailers exert negative externalities on each other. Irrespective of the scenario behind these linkages, efficiency requires agents to internalize those externalities when choosing their own effort. In the case of positive externalities, this goal is reached with CPE, while RPE is preferred when externalities are negative.

and discriminatory treatment of otherwise symmetric agents.

Informational linkages occur when the performance of one agent provides insight into the effort exerted by their peer (Section 3.2). Traders whose performance depends on the same aggregate stock market index constitute a standard example. Another example includes salespeople serving similar customers in different geographic areas. These informational externalities ought to be exploited by indexing one agent's compensation on the performances of their peers. Whether the resulting compensations are competitive or collective depends on the correlation between those performances. In cases of positive correlation, such as in the trader example, competitive incentives are preferred, for instance implemented through tournaments (Section 3.3). Conversely, with negative correlation, collective incentives tend to dominate.

In contrast, when agents work as a team, individual outputs can no longer be observed. In this scenario, incentives can only be based on an aggregate measure and only CPE schemes are available to the principal. Free riding is now pervasive and under-provision of effort may follow. Moral hazard in teams is studied in Section 3.4.

Another, somewhat different, form of incentive linkage arises when agents are no longer purely selfish. We cover the topic of social preferences in Section 3.5. The fact that agents are concerned about their payments relative to those of others, because of altruism or envy, creates externalities similar to production externalities. To illustrate, status-seeking agents are incentivized by *ex post* differences in wages, and RPE magnifies their incentives, while if agents exhibit a strong taste for equity, CPE is preferable.

Incentive schemes induce a game among agents in the organization at the time they choose their efforts. Section 4.1 analyzes these games and unveils the specific strategic issues that arise with RPE and CPE. RPE creates a prisoner's dilemma between agents, while CPE suffers from free riding and multiple equilibria. In Section 4.2, we modify the information structure within the team, either by allowing mutual monitoring (Section 4.2.1) or by considering sequential production processes with observable stages (Section 4.2.2). This analysis paves the way to a discussion of the joint choice of incentive schemes and information structure by the principal; a problem that is addressed in Sections 4.2.3 and 4.2.4.

In Section 5, we depart from the assumption of agents adopting Nash behavior and instead focus on the implications of more collusive behavior. Interdependent schemes introduce strategic externalities among agents, and adopting a more cooperative approach enables them to internalize these externalities. To illustrate, RPE schemes are more likely to be gamed by agents–much in the same way as rival firms are tempted by market collusion when their competition would otherwise be fierce. The impact of collusion or cooperation on incentives depends on the organization of work. Agents may have the ability to monitor each other's efforts, enabling coordination. Coalitional behavior can be established through enforceable side-contracts or reputation-based mechanisms. If formal contracts between agents are available, full cooperation becomes possible, particularly regarding joint effort choices. When reputation mechanisms are at work, the extent of cooperation depends on how agents perceive the future of their relationships. Mutual monitoring not only mitigates moral hazard but also expands the range of contracting possibilities. The principal can leverage this dynamics to implement more cost-effective incentives through CPE.

Even when agents cannot monitor each other's efforts, they can still increase their joint payoffs with side-payments. For example, under an RPE scheme, a high-performing agent may share his bonus with a low-performing agent if reciprocal behavior is ensured. In this case, however, the principal cannot be better off, as the agent's side-contracting imposes additional collusion-proof constraints on the set of incentive schemes.

Section 6 gives a more active role to the principal of the organization. Her own incentive issues impose additional restrictions on compensations, which in turn may alter the incentives for the agents. We make this point in a variety of contexts. First, the principal may contribute to the success of the agents' production processes through her own effort. Balancing incentives between the principal and the agents using CPE creates tension, which is relieved by using RPE. Secondly, the principal may manipulate workers' performance evaluations to avoid paying bonuses. This can be prevented through the use of tournaments, which are a specific form of RPE. In this case, the principal pays a fixed prize, regardless of the subjective evaluation of individual performance. Thirdly, despite being observable by all players, the agents' performance may not be verifiable. Performance bonuses can now be enforced through relational contracts. The fear of terminating the relationship makes it credible for the principal to pay bonuses for high performances, as long as the bonuses are not too large. In a multi-agent setting, this enforcement constraint applies to the sum of individual bonuses, as long as the agents can threaten to leave collectively following any deviation by the principal. This constraint is, again, easier to satisfy with RPE, as such schemes permit to smooth the overall payment made by the principal.

In many circumstances, the provision of incentives within the firm is constrained by opportunities on adjacent markets. Whether collective or competitive incentives remain effective depends on the pressure from these markets on the organization. This topic is addressed by means of two examples in Section 7. In a nutshell, market environments introduce various forces that tilt incentive schemes towards independent evaluations. Consider for example a scenario in which managers' incentives are driven by career concerns on the labor market. By exerting effort and reaching high performance, a manager may induce the labor market to believe he has a high productivity and secure higher wages in the future. When their talents have a common component, managers free ride in effort provision if each potential firm on this market can condition their future wages on the whole array of observed past performances of managers. Instead, IPE do not suffer from that sort of free riding.

A second example of market interactions constraining the set of feasible incentive schemes arises when property rights on individual outputs are weak. The threat of agents directly selling their output on the product market (or leaving, taking their clients with them) becomes a significant concern for firms. To counteract this hold-up problem, agents must be paid at least as much as what they would receive on the market. This limits the firm's ability to rely on the information contained in a peer's performance when determining an agent's own wage, making IPE more appealing.

2 A Model of Incentives in Multi-Agent Organizations

2.1 Basics

Our workhorse model depicts an organization composed of two agents (thereafter "he/they") working on behalf of a risk-neutral principal (thereafter "she"). The moral hazard literature leans in two directions for making incentives costly and studying meaningful trade-offs in such contexts. A first strand of the literature supposes that agents are risk-averse, an assumption shared by earlier seminal contributions. The focus here is on the trade-off between incentives and insurance as stressed in Ross (1973), Mirrlees (1999), Holmström (1979), Shavell (1979) and Grossman and Hart (1983). A second, more recent body of works assumes instead that agents are risk-neutral but protected by limited liability. Here the works of Sappington (1983), Innes (1990), Holmstrom and Tirole (1997) and Poblete and Spulber (2012) stand out as important references. In these circumstances, the limited scope for punishing agents' bad performances implies that inducing effort requires giving them some limited liability rent. The trade-off is now between incentives and rent extraction (see Laffont and Martimort, 2002, Chapter 4). This second paradigm is somewhat less prone to technicalities, and we choose limited liability as the sole source of agency frictions unless otherwise stated. All players are thus assumed to be risk neutral, but we also discuss the case of risk aversion when relevant. Focusing on limited liability as the source of agency costs is almost costless in terms of economic insights. We can still recover most of the results of the extant literature that have been developed when assuming risk aversion instead. The reason for this similarity between the two paradigms is simple: both limited liability and risk aversion tend to create frictions by limiting the variability in payments, although for different reasons. Yet, the tractability of limited liability environments also allows us to present original results that would have required a much more tedious analysis had we instead assumed risk aversion.

For simplicity, agents (indexed by $i \in \{1,2\}$) are assumed to be symmetric, but most qualitative results are unchanged otherwise. As in Itoh (1991), Ramakrishnan and Thakor (1991), Che and Yoo (2001) and Fleckinger (2012) among others, each agent i is in charge of a single project, whose (expected) benefit to the principal increases with the effort e_i exerted by this agent. For simplicity, effort is binary. Each agent privately chooses whether or not to exert effort, $e_i \in \{0,1\}$, at a cost ce_i where $c > 0$. Efforts are not observed by the other players, unless specified otherwise. Later, moral hazard will stem from the fact that agents' efforts are not directly observable and are thus not verifiable by a Court of Law.

For most of our results below, we do not need to specify the principal's benefit function. Only the minimization of the cost of providing incentives will matter to evaluate the best compensation schemes. Nevertheless, we will explicitly introduce the benefit part whenever needed.

The principal relies on a pair of signals (sometimes referred to as outcomes or performances depending on the precise context) $\boldsymbol{R} = (R_1, R_2)$ to regulate the agents' behavior. For simplicity, each signal may only take two possible values, high (H) or low (L). For most of the results discussed, it is assumed that the signals are verifiable and serve as the sole variables available for contracting. We denote by $Prob(\boldsymbol{R}|e_1, e_2)$ the probability of observing signals $\boldsymbol{R}$ when agent's efforts are e_1 and e_2.

The most natural interpretation of the model is that signal R_i directly reflects agent i's performances, which itself is a garbled measure of effort. Other interpretations for this model are nevertheless available. For instance, the actual realizations of the value of the project might be too distant in the future to be contracted upon (e.g. a firm's stream of profits over an infinite time horizon) and only some intermediate signals can be subject to a contract. Or, the signal directly reflects effort while performances remain non-verifiable.

2.2 Incentive Schemes

If efforts were observable and subject to a contract, the principal could use forcing contracts to implement her most preferred effort pair $(1,1)$. Forcing contracts requires each agent i to exert effort, and be compensated for the cost incurred, c, so that they agree to work for the principal. Hence, each agent i faces a simple bilateral compensation scheme. In this first-best world, there is no need to write contracts on realized outcomes, nor to introduce any sort of linkages between the incentives provided to different agents. The direct control of effort suffices. This scenario stands as a crucial benchmark, demonstrating that intricate compensation schemes linking agents only become attractive when agency considerations matter.

When efforts are not observable, the principal has to rely on incentive contracts which can only be based on observable signals. Since the signals that pertain to different agents are generally related to each other, the compensation scheme should be made contingent on the whole signal pair $\boldsymbol{R}$.

Although Section 4.2.3 below will show that treating symmetric agents in a discriminatory manner might be attractive in some structured environments, we will suppose, for the purpose of our general presentation and without loss of generality for most results below, that symmetric agents operate under symmetric incentive schemes. A symmetric incentive scheme is thus an array of rewards

$$\boldsymbol{w} = \{w(H,H), w(H,L), w(L,H), w(L,L)\}$$

that represents the wage received by an agent as a function of his own performance–the first variable–and the other agent's performance–the second variable. The notation $w(\boldsymbol{R})$ will represent the element of $\boldsymbol{w}$ when the pair of signals $\boldsymbol{R}$ is realized. Given an outcome-contingent scheme $\boldsymbol{w}$ and a pair of efforts (e_i, e_{-i}) (with the usual convention), agent i's expected payoff can be defined as

$$U_i(\boldsymbol{w}|e_i, e_{-i}) = \mathbb{E}_{\boldsymbol{R}}\left[w(\boldsymbol{R})|e_i, e_{-i}\right] - ce_i. \tag{1}$$

Our central question is to determine under which circumstances the principal prefers to use relative incentives, or on the contrary to use team bonuses. To give a precise formal meaning for these notions, we shall adopt the following typology throughout the paper:[3]

[3] The vector-inequalities below should be understood as component-wise comparisons with at least one strict inequality.

Definition 1. *(Standard incentive schemes)*

Relative Performance Evaluation *(RPE) arises when:*

$$(w(H,L), w(L,L)) > (w(H,H), w(L,H)).$$

Collective Performance Evaluation *(CPE) arises when:*

$$(w(H,H), w(L,H)) > (w(H,L), w(L,L)).$$

Independent Performance Evaluation *(IPE) arises when:*

$$(w(H,H), w(L,H)) = (w(H,L), w(L,L)).$$

With RPE, holding fixed his own performance, an agent is better off when his peer fails, while the opposite is true with CPE. While RPE organizes a competitive playing field between agents, CPE instead provides collective incentives and may foster cooperation. With IPE, the organization is run by means of simple bilateral contracts with each agent being paid as a function of his own performance only. Note that these three types of schemes do not exhaust all possible configurations. As we discuss below, adopting more mixed solutions might sometimes be attractive.

2.3 The Principal's Problem

Since the principal's objective function is separable in terms of the benefits and costs of implementing any given effort pair, we follow Grossman and Hart (1983) in first looking for the agency cost of implementing any effort pair. This relegates to a second step the question of determining whether it is worth paying those costs or simply opting for zero effort. Throughout, we will assume that the principal always wants both agents to exert effort. The principal's benefit of inducing the effort pair $(1,1)$ minus the agency cost of doing so is thus assumed greater than the principal's payoff with any other effort configuration.[4] The first step of the optimization problem, and our focus thereafter, thus consists in minimizing the cost of implementing the effort pair $(1,1)$.

[4]It is worth emphasizing again that under this assumption, our results then do not depend on the shape of the principal's benefit function. Later on, we explicitly define the principal's payoff whenever needed. The results presented until Section 6 are based solely on implementation costs.

In abstract game theoretic terms, contracts define a normal form game played by the agents. When making their effort choices, we first assume that agents adopt a non-cooperative behavior and play a Nash equilibrium for that game. In general, this Nash equilibrium may not be unique. Implementation is said to be *partial* when the effort pair $(1,1)$ is implemented in at least one Nash equilibrium. In a first pass, we shall neglect this multiplicity problem and postpone its analysis to Section 4.1.2. Other implementation concepts can also be entertained. For instance, modelers may be interested in *full* implementation (the effort pair $(1,1)$ is implemented in all Nash equilibria), subgame-perfect implementation (the effort pair $(1,1)$ is implemented as a subgame-perfect Nash equilibrium of an extensive-form game) or collusive behavior (the effort pair $(1,1)$ is implemented when agents collude). We devote attention to these alternative concepts in Section 5.

For the time being, observe that (partial) Nash implementation requires that the following incentive constraint be satisfied for each agent:

(2) $$U_i(\boldsymbol{w}|1,1) \geq U_i(\boldsymbol{w}|0,1) \quad \text{for } i = 1,2.$$

This incentive constraint requires that exerting effort is a best-response for agent i when the other exerts effort. Other constraints must be added to complete the characterization of the whole set of incentive-feasible allocations. First, agents are subject to limited liability. Assuming that agents have no assets prior to contracting, the limited liability condition can be written in vector terms as:

(3) $$\boldsymbol{w} \geq 0.$$

Next, the principal should also ensure that agents prefer to participate, rather than obtaining some reservation utility, such as seeking employment elsewhere in the economy. For simplicity, this outside opportunity will be fixed at zero. As is standard in moral hazard contexts (Laffont and Martimort, 2002, Chapter 4), the zero-limited liability constraint (3) implies this zero-participation constraint, which is thus ignored in the sequel.[5]

[5]This combination of assumptions provides a classic description of an employment contract. We could alternatively entertain the possibility that agents possess pledgeable wealth, denoted as ω, prior to contracting. The limited liability constraint (3) becomes

$$\boldsymbol{w} \geq -\omega.$$

The principal is therefore able to punish agents for bad performances, which helps to reduce agency costs

Observe that the incentive constraint of each agent only features his own wage, and that the cost of implementation borne by the principal also depends linearly on those wages. Therefore, the principal's agency cost minimization program is additive in the incentive schemes given to each agent. Hence, this program can be expressed on a *per capita* basis so as to find the cheapest way of providing incentives for an individual agent:

$$\min_{\boldsymbol{w}} \mathbb{E}_{\boldsymbol{R}} \left[w(\boldsymbol{R}) | 1, 1 \right] \text{ subject to (2) and (3).} \tag{4}$$

This problem is thus very close to that found in a simple bilateral relationship between the principal and a single agent. It is nevertheless richer, partly because the signal structure is itself richer, and partly because agent i's incentive constraint depends on his peer's effort.

The limited liability constraint (3) is key to study the potential optimality of CPE and RPE. Indeed, in absence of such a constraint, a well calibrated IPE scheme implements the first-best: the collective incentive problem can freely be solved, without resorting to any complex linking of agent's rewards.

To show this, consider the following IPE scheme $w^*(\boldsymbol{R})$ defined as

$$w^*(H, H) = w^*(H, L) = c + \frac{1 - Prob(H|1, 1)}{Prob(H|1, 1) - Prob(H|0, 1)} c > 0$$

and

$$w^*(L, H) = w^*(L, L) = c - \frac{Prob(H|1, 1)}{Prob(H|1, 1) - Prob(H|0, 1)} c < 0,$$

where the probability that R_1 realizes is denoted $Prob(R_1|e_1, e_2) \equiv Prob(R_1, H|e_1, e_2) + Prob(R_1, L|e_1, e_2)$ for any R_1 in $\{H, L\}$ and for any effort pair (e_1, e_2). With this IPE scheme, each agent gets zero expected utility when exerting effort, while (2) holds as an equality. The principal can thus reach the same payoff as if effort was observed and forcing contracts were available. Hence, RPE and CPE cannot constitute strictly better options: linking agents' compensations can only become valuable when the principal cannot achieve the first-best in the bilateral relationship. Unfortunately, and in contrast with forcing contracts, this IPE scheme now violates the limited liability constraint. An agent is punished with a negative payment when a bad outcome occurs.[6]

but does not alter any fundamental insights, as long as ω is not too large. When ω is sufficiently large, contracts that make agents residual claimants for their performance can achieve first-best efforts, offering agents precisely their reservation utility. However, this scenario is of less interest.

[6] It is easy to see that neither any RPE nor any CPE scheme could provide incentives to exert effort, extract all surplus from the agents and keep positive wages under all circumstances.

Pursuing the comparison with a simple bilateral relationship, we might follow Laffont and Martimort (2002) and define agent i's *liability rent* r as the excess payoff he gets on top of the disutility of effort he would be paid were effort observable. Formally, we have:

$$r \equiv \mathbb{E}_{\boldsymbol{R}}\left[w(\boldsymbol{R})|1,1\right] - c.$$

Minimizing the expected payments needed to implement high effort as a Nash equilibrium amounts then to minimizing the agents' liability rent.

The structure of the optimal incentive scheme for this problem shall depend in fine details on how observed signals are related to non-observable efforts. To understand this dependence, the following definition is useful.

Definition 2. *The* ***incentive efficiency*** *at signals* $\boldsymbol{R}$ *is defined as*

$$I(\boldsymbol{R}) \equiv \frac{Prob(\boldsymbol{R}|1,1) - Prob(\boldsymbol{R}|0,1)}{Prob(\boldsymbol{R}|1,1)}.$$

The incentive efficiency is the ratio between the coefficient of the wage $w(\boldsymbol{R})$ in the incentive constraint (2) and the probability of paying this wage in equilibrium. It is therefore a measure of how an agent's marginal incentives to exert effort are related to the principal's marginal cost of inducing such effort. The incentive efficiency can be identified with the likelihood ratio that a positive effort has been chosen by agent i when signal $\boldsymbol{R}$ has been observed and conjecturing that agent $-i$ has also exerted a positive effort. Note that the incentive efficiency is bounded above by one. At this upper bound, the wage is fully effective in inducing effort, because signal $\boldsymbol{R}$ indicates that agent i has exerted effort for sure, since we must then have $Prob(\boldsymbol{R}|0,1) = 0$.

As a preliminary step towards solving Problem (4) for the principal, the following lemma, the proof of which is featured in the Appendix, turns out to be very useful.

Lemma 1. *Under risk neutrality and limited liability, an optimal incentive scheme that solves Problem (4) entails a positive wage only for a signal pair* $\boldsymbol{R}$ *with the highest incentive efficiency* $I(\boldsymbol{R})$*.*

This lemma formalizes the intuitive idea that all wages should be concentrated on those signals that are the most efficient in inducing effort. To illustrate, consider the usual case where the monotone likelihood property holds, i.e., a higher signal on agent i's performance is more informative on the fact that this agent has exerted effort. This condition

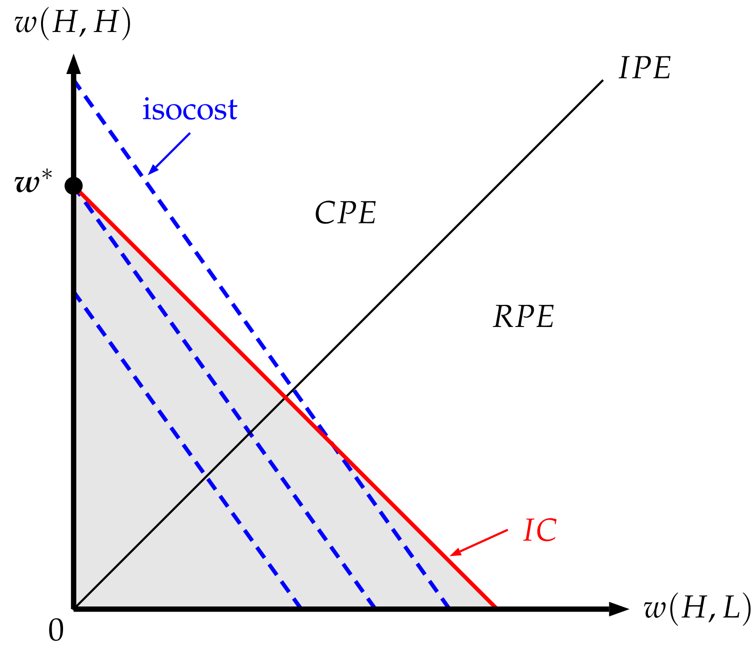

Figure 1: The Principal's Problem in which CPE is Optimal

writes as

(5)
$$\frac{Prob(H, R_{-i}|1,1) - Prob(H, R_{-i}|0,1)}{Prob(H, R_{-i}|1,1)} \geq \frac{Prob(L, R_{-i}|1,1) - Prob(L, R_{-i}|0,1)}{Prob(L, R_{-i}|1,1)} \quad \forall R_{-i}.$$

Under those circumstances, and as an immediate corollary of Lemma 1, agent i should never be rewarded following a bad signal $R_i = L$:

$$w^*(L, H) = w^*(L, L) = 0.$$

In the sequel, with a view to simplify the presentation, we shall assume that Condition (5) holds, without always being explicit about it.[7]

Turning now to the characterization of $w(H, H)$ and $w(H, L)$, the solution to Problem (4) boils down to finding which signals pair has the highest incentive efficiency. Solving this problem tells us whether CPE (when $w(H, H) > 0 = w(H, L)$) or RPE (when $w(H, L) > 0 = w(H, H)$) should be preferred.

Figure 1 provides a convenient representation of the principal's problem. On this

[7]See Fleckinger (2012, Section 4) for a complete discussion of the implications of relaxing this Condition.

figure, we have drawn in the $(w(H,H), w(H,L))$ space the incentive constraint (IC) and three iso-agency cost curves (isocost). Clearly, those agency costs increase in the North-East direction, i.e., when either wage is increased. The set of incentive compatible schemes lies above the incentive constraint (IC), while the shaded area contains schemes under which an agent shirks. It is then straightforward to observe that the optimal scheme is determined at the intersection of the incentive constraint and the lowest possible isocost curve. Whether the optimal scheme is CPE or RPE (that is, whether the optimal incentive scheme lies above or below the 45° line) depends on the comparison of the slopes of the IC and the isocost straight lines. In turn, these slopes both depend on the stochastic mapping $Prob(\boldsymbol{R}|e_i, e_{-i})$. Lemma 1 shows that this comparison actually amounts to checking whether the incentive efficiency is higher at (H,H) or at (H,L), respectively. In the first case (the slope of the IC constraint being steeper than the slope of the isocost), CPE is preferred while RPE is preferred otherwise. When the incentive efficiencies at (H,H) and (H,L) are equal, a whole range of incentive schemes are optimal. In this knife-edge case, there is no loss of generality in opting for IPE compensations.

3 Linking Incentives Across Agents

We now examine, in turn, the implications of technological and informational externalities for the use of CPE and RPE. A key issue is to determine when those externalities force the organization to give up IPE and to link the two agents' incentives, providing more cooperation through CPE or more competition by means of RPE.

3.1 Technological Externalities

Consider two retailers selling products for a manufacturer on the same market place. If retailers sell the same product, an increase in one agent's effort makes it more difficult for his peer to attract customers and so the latter's performance decreases. On the other hand, if retailers sell complementary products, one retailer's effort to promote his own product might also make his peer benefit from a larger clientele. In such a context, the optimal incentive scheme should be designed so as to make each agent internalize the impact of his own action on his peer.

In order to first highlight the pure role of technological externalities, we shall assume that projects are *informationally independent*, i.e., signals on each agent's performances are conditionally independent. Formally, this means that the probability of a given pair of signals $\boldsymbol{R} = (R_i, R_{-i})$ has the following product structure:

$$Prob(\boldsymbol{R}|e_i, e_{-i}) = Prob(R_i|e_i, e_{-i})Prob(R_{-i}|e_i, e_{-i}). \tag{6}$$

Now observe that the incentive efficiency in state (H, H) is greater than in state (H, L) whenever

$$\frac{Prob(H, H|0, 1)}{Prob(H, H|1, 1)} \leq \frac{Prob(H, L|0, 1)}{Prob(H, L|1, 1)}.$$

Using the fact that projects are *informationally independent* as in Condition (6), this inequality becomes

$$\frac{Prob(R_{-i} = L|1, 1)}{Prob(R_{-i} = H|1, 1)} \leq \frac{Prob(R_{-i} = L|0, 1)}{Prob(R_{-i} = H|0, 1)}. \tag{7}$$

This expression makes it explicit that the probabilities involved are those concerning signal R_{-i} on agent $-i$'s performance. The inequality then simply means that a high performance signal from agent $-i$ is now more informative about agent i's shirking. Further simplifying using $Prob(R_{-i} = H|e_i, e_{-i}) + Prob(R_{-i} = L|e_i, e_{-i}) = 1$, and flipping the identity of the two agents, this condition can also be expressed as

$$Prob(H|1, 1) \geq Prob(H|1, 0). \tag{8}$$

In other words, the condition requires that agent $-i$'s effort increases the probability of a good signal for agent i. Under those circumstances, agents' efforts are complementary. When the inequality in (8) is reversed or is an equality, efforts are either substitutes or projects are independent.

From Lemma 1, and still assuming that Condition (5) holds, we know that a positive payment is given only in state (H, H). Instead, whenever the inequality in (7) is reversed, rewards are only used in state (H, L). The next proposition summarizes these findings.

Proposition 1. *Suppose that projects are informationally independent, i.e., Condition (6) holds. The optimal scheme exhibits CPE whenever efforts are complements, RPE when they are substitutes, and IPE when they are independent.*

This proposition echoes broad and intuitive principles often found in organizational economics. In a celebrated paper, Alchian and Demsetz (1972) explicitly defined team production as relying on technological complementarities when they wrote *"Team production [...] is production in which 1) several types of resources are used and 2) the product is* *not a sum of separable outputs of each cooperating resource"* (Alchian and Demsetz, 1972, p. 779); a feature referred to as *"labor input complementarity"* by Marschak and Radner (1972). Incentives within the firm's boundaries are thus expected to always feature some degree of CPE. Tools such as stock-options, team and division bonuses, and even more informal practices like celebrations of collective goal achievements all share this feature to some extent. In contrast, incentives across organizations often feature elements of RPE, especially when those organizations compete on the market place by, for instance, selling substitute products. Market competition and competition for rewards then go hand in hand. Taking a broader perspective, the optimal match between technology and incentives and their consequences for job design was a concern of the early developments of agency theory in the fields of labor and managerial economics (e.g. Drago and Turnbull, 1987, 1988; Choi, 1993; Itoh, 1994; Lin, 1997).

3.2 Informational Externalities

The previous section showed how technological externalities called for some linkage in agents' compensation schemes. We shall now consider a stochastic structure that allows us to focus on the sole role of informational externalities. To this end, we now assume that the probability distribution over outcomes can be written as follows

$$Prob(R_i = R_{-i}|e_i, e_{-i}) = Prob(R_i|e_i)Prob(R_{-i}|e_{-i}) + \gamma, \quad (9)$$

$$Prob(R_i \neq R_{-i}|e_i, e_{-i}) = Prob(R_i|e_i)Prob(R_{-i}|e_{-i}) - \gamma, \quad (10)$$

where the parameter γ is actually a correlation parameter (whose absolute value is supposed to be small enough to ensure that the probabilities of all signals remain in the unit interval under all circumstances). A higher value of γ means that agents are more likely to jointly generate good (or bad) signals on their performances while $\gamma = 0$ corresponds to the case of independent projects. Observe also that

$$Prob(R_i|e_i, e_{-i}) = Prob(R_i|e_i), \quad (11)$$

i.e., the effort of one agent has no influence whatsoever on his peer's performance. Technological externalities are absent. To shorten expressions, we shall use another piece of notation in the sequel and thus denote

$$p(e_i) \equiv Prob(R_i = H|e_i) \quad \forall e_i \in \{0,1\}. \tag{12}$$

One has for instance $Prob(H, L|e_i, e_{-i}) = p(e_i)(1 - p(e_{-i})) - \gamma$ where, of course, $p(1) > p(0)$.

Correlation in performances is an important driver of the shape of the incentive schemes found in practice. The use of benchmarking and other sorts of yardstick mechanisms in many regulatory contexts, especially in the electricity and water industries in the UK, the Netherlands and Germany offered clear illustrations of this issue over recent years. For instance, OFWAT, the water regulator in the UK, uses a variety of econometric models to benchmark price caps for each utility under scrutiny.

In his seminal paper, Holmström (1979) demonstrated the so called *"Informativeness Principle"*, also referred to as the *"Sufficient Statistics Result"*, which gives the precise conditions for an additional signal to be valuable to calibrate incentives. This principle, which was demonstrated in the context of a simple bilateral relationship, has two parts. First, it states that including an additional signal in the compensation scheme of a given agent is optimal whenever that signal is informative on the agent's effort. In our multi-agent setting, this additional signal consists of the performance of the agent's peer. Secondly, the use of any other signal, which would *not* be informative on the agent's effort, would only add noise to the compensation. This would be costly and thus suboptimal when the agent is risk-averse.

The first insight remains available in the context of a multi-agent organization, with the extra subtlety that the additional signal used to better incentivize a given agent is generated by the effort of his peer. In our context, this result states that signal R_{-i} should be used in the contract of agent i if and only if R_i is not a sufficient statistic for $\boldsymbol{R}$.

The second take-away brought by the *"Informativeness Principle"* cannot be *stricto sensu* replicated in our model which assumes risk neutrality. Yet, it remains true that paying agent i depending on his peer's performance is useless when $\gamma = 0$. The signal R_{-i} does not influence the incentive efficiency of signal R_i in this case, since then

$$I(\boldsymbol{R}) = \frac{Prob(R_i|1)Prob(R_{-i}|1) - Prob(R_i|0)Prob(R_{-i}|1)}{Prob(R_i|1)Prob(R_{-i}|1)} = \frac{Prob(R_i|1) - Prob(R_i|0)}{Prob(R_i|1)}.$$

Relying on IPE schemes is thus (weakly) optimal in this context.

Proposition 2. *Suppose that the information structure satisfies (9) and (10). The optimal scheme exhibits RPE when $\gamma > 0$, CPE when $\gamma < 0$, and IPE when $\gamma = 0$.*

The proof again follows on from Lemma 1. With our previous notations, observe that state (H, L) has a higher incentive efficiency than (H, H), i.e., RPE is optimal, whenever

$$\frac{(1-p(1))(p(1)-p(0))}{p(1)(1-p(1))-\gamma} > \frac{p(1)(p(1)-p(0))}{p(1)^2+\gamma} \Leftrightarrow \gamma > 0.$$

The intuition here is that, when performances are positively correlated, state (H, H) is less informative on the fact that agent i has made a positive effort than state (H, L). RPE schemes then filter any common random component that could jointly affect individual performances. Because the common component due to factors beyond the agents' control has been corrected for, incentives can be better calibrated and offered at a lower cost when an agent is paid following mixed performances (H, L).

For future reference, let us specify the optimal positive bonus in the various scenarios. When $\gamma < 0$, the agents' incentive constraint in a CPE is as follows:

$$(p(1)^2+\gamma)w(H,H) - c \geq (p(0)p(1)+\gamma)w(H,H).$$

The optimal positive bonus, which is obtained when this incentive constraint binds, is

$$w^*(H,H) = \frac{c}{p(1)(p(1)-p(0))}, \tag{13}$$

and the limited liability rent left *per capita* is therefore

$$r^*_{CPE} = (p(1)^2+\gamma)w^*(H,H) - c = \frac{p(0)p(1)+\gamma}{p(1)(p(1)-p(0))}c. \tag{14}$$

As performances become more negatively correlated (γ negative and decreasing), the *per capita* rent under CPE diminishes.

Alternatively, when $\gamma > 0$, the agents' incentive constraint for an RPE scheme becomes

$$(p(1)(1-p(1))-\gamma)w(H,L) - c \geq (p(0)(1-p(1))-\gamma)w(H,L)$$

and he optimal positive bonus is now

$$w^*(H,L) = \frac{c}{(1-p(1))(p(1)-p(0))}, \tag{15}$$

while the corresponding limited liability rent is

$$r^*_{RPE} = (p(1)(1-p(1)) - \gamma)w^*(H,L) - c = \frac{p(0)(1-p(1)) - \gamma}{(1-p(1))(p(1)-p(0))}c. \tag{16}$$

As performances become more positively correlated (γ positive and increasing), the *per capita* rent under RPE diminishes.

Let us consider how our results would have been modified had we instead supposed risk aversion rather than limited liability as a source of agency frictions, while keeping the same informational structure. Preferences are separable and of the form $u(w) - ce$ where u is increasing and (strictly) concave. It is routine to check that the optimal contract would be such that[8]

$$w^*(H,H) > w^*(H,L) > w^*(L,H) > w^*(L,L) \text{ when } \gamma < 0,$$

while

$$w^*(H,L) > w^*(H,H) > w^*(L,L) > w^*(L,H) \text{ when } \gamma > 0.$$

In other words, RPE schemes are again optimal when signals are positively correlated while CPE dominates otherwise. In both scenarios, good signals on performance are rewarded while bad signals are punished. Of course, independent schemes based on individual performances are optimal only when $\gamma = 0$.

The use of the *"sufficient statistics result"* in (and its generalization to) multi-agent models goes back to Holmström (1982), Gjesdal (1982) and Mookherjee (1984). Other more recent contributions include not only a large literature on corporate finance studying the incentives of CEOs and money managers (Admati and Pfleiderer, 1997; Aggarwal and Samwick, 1996, 1999; Abowd and Kaplan, 1999, among many others), but also some specific applications to the theory of the firm (Salas-Fumas, 1992; Luporini, 2006). Taking a broader perspective, the fact that correlation pleads for competition through some sort of benchmarking is also a feature often found in models of adverse selection instead of moral hazard; a result first shown by Demski and Sappington (1984). The use of *"Yardstick Competition"* mechanisms in a regulatory framework (Shleifer, 1985) can be thought of as a version of Sappington and Demski (1983) in the special case of perfect correlation. Riordan and Sappington (1988), Crémer and McLean (1985, 1988), McAfee and Reny (1992)

[8]In a nutshell, with risk-averse agents, optimal wages follow the same ordering as incentive efficiencies, as observed in the case of risk-neutral agents (see Fleckinger, 2012).

and Kosenok and Severinov (2008) have also developed models showing how correlated information can be useful in various mechanism design environments. The provocative and controversial view of those models is that fine details on the correlation structure can sometimes be used by the mechanism designer to fully extract all of the agents' information rents. The moral hazard setup we consider here does not yield such an extreme result. The preferred incentive modes are robust to small changes in performance correlation. Yet, it is still true that using correlation reduces the agent's rent.

3.3 Tournaments as RPE

A tournament, or contest, is an extreme form of RPE which relies only on ordinal information: an agent receives a "prize" when he outperforms his peers. In other words, in a tournament, an agent's performance is benchmarked against those of his peers. In practice, tournaments are often found to reward innovation. They can also be viewed as a reasonable proxy for the process of promotion within both public and private organizations, or as stylized representations of political processes that select political leaders through elections.

Of course, relying only on ordinal information could mean throwing away valuable information. On the other hand, as argued by Lazear and Rosen (1981), using only ordinal measures of performance may be justified by the greater cost of measuring performances in absolute terms as compared to ranking. In fact, the advantage of tournaments is precisely to overcome (part of) the costs imposed by noisy performance measures that arise when contracting individually with agents. Although Proposition 2 has already shown us how optimal RPE schemes dominate independent contracts when performances are positively correlated, the above limitations of tournaments leave open the possibility of a less stark trade-off.

The purpose of this section, in the spirit of Lazear and Rosen (1981), is to understand why the principal could benefit from using a tournament instead of a pair of independent schemes based on the sole performance of each agent. At this stage, this comparison might seem somewhat odd since neither form of contracts is likely to be optimal. We postpone to Section 6.2.1 further arguments that might demonstrate the optimality of tournaments in specific contexts.

To compare tournaments and IPE, we now slightly enrich the information structure of (9) and (10) to account for the difficulty in measuring performance in absolute terms.

In the spirit of Baker, Gibbons and Murphy (1994), we hereafter assume that the principal can only rely on a biased measure of performances.[9] Formally, the probabilities of measuring similar performances, namely (H,H) and (L,L), are now respectively given by

$$(17) \qquad Prob(H,H|e_i,e_{-i}) = p(e_i)p(e_{-i}) + \gamma + \beta((1-p(e_i))(1-p(e_{-i})) + \gamma),$$

$$(18) \qquad Prob(L,L|e_i,e_{-i}) = (1-\beta)((1-p(e_i))(1-p(e_{-i})) + \gamma),$$

where $\beta \in [0,1]$ and we assume that all probabilities lie between 0 and 1. In words, the principal might observe (H,H) when the true underlying performance is actually (L,L). The parameter β measures this observational bias.

Instead, underlying performances, when asymmetric, are perfectly observed and the corresponding probabilities remain

(19)

$$Prob(H,L|e_i,e_{-i}) = p(e_i)(1-p(e_{-i})) - \gamma \text{ and } Prob(L,H|e_i,e_{-i}) = (1-p(e_i))p(e_{-i}) - \gamma.$$

In a tournament, the agents' performances are compared to determine their respective rewards. The agent with the highest assessed performance receives a given prize W, while his losing peer gets nothing. Because of risk neutrality and limited liability, the losing agent's reward should optimally be set at zero and we streamline the exposition accordingly. Importantly, the reward is independent of the absolute performance difference. Finally, we assume that ties are broken according to a fair coin. A tournament is thus akin to imposing restrictions on the (expected) wage structure, namely

$$w(H,L) = W,\ w(L,L) = w(H,H) = \frac{1}{2}W, \text{ and } w(L,H) = 0.$$

By symmetry, both agents have an equal *ex ante* probability of winning when they choose to exert effort. In turn, let q be the probability that agent i wins the tournament when shirking while his peer exerts effort. This event occurs either upon receiving a signal $R_i = H > L = R_{-i}$, or with probability $1/2$ when $R_i = R_{-i}$. The overall probability

[9] To save on notations, we shall use the same variables to denote the biased measure and the "true" underlying performance, even though these are technically different signals.

of winning can thus be expressed as:

$$q \equiv Prob(H, L|0, 1) + \frac{1}{2}\left(Prob(H, H|0, 1) + Prob(L, L|0, 1)\right).$$

Simplifying, we obtain

$$q = \frac{1}{2} - \frac{p(1) - p(0)}{2}. \tag{20}$$

Importantly, the correlation parameter γ disappears from Equation (20). This captures in a nutshell an important idea: tournaments wash out the common shocks that affect performances. Moreover, q does not depend on the bias parameter β either. Indeed, either the signal perfectly ranks the agents with distinguishable performances, or a fair coin breaks the tie so that whether performances are actually high or low does not matter.

A tournament with prize W implements an equilibrium with high efforts provided that the following incentive constraint holds:

$$\frac{W}{2} - c \geq qW.$$

The optimal prize, which minimizes agency costs, is thus

$$W^T \equiv \frac{2c}{1 - 2q} = \frac{2c}{p(1) - p(0)}.$$

From a practical point of view, this tournament has the desirable property that it does not need to be tailored to fine (and hard to assess) details of the information structure. Another important feature, pointed out among others by Malcomson (1984), is that tournaments entail a fixed incentive budget–the prize W–which the principal might find attractive. Under the optimal tournament, the *per capita* liability rent is

$$r^T = \frac{W^T}{2} - c = \frac{(1 - p(1) + p(0))c}{p(1) - p(0)}.$$

We now turn to the analysis of individual contracts. Such schemes pay agent i only on the basis of the observed (but biased) signal on his own performance, R_i. Clearly, these compensations are not optimal, as signals are correlated (both through the correlation parameter γ and the measure of bias β). Nevertheless, considering those schemes corresponds fairly closely to the approach of Lazear and Rosen (1981), who study contracts that are in addition restricted to be linear.

With this observation in mind, we are now in position to derive the optimal individual contract. This scheme pays agent i a bonus $w(H)$ only when $R_i = H$ and zero otherwise (again due to the limited liability requirement). In this scenario, the agents' incentive constraint writes as

$$(Prob(H,H|1,1) + Prob(H,L|1,1))w(H) - c \geq (Prob(H,H|0,1) + Prob(H,L|0,1))w(H).$$

Agency costs are minimized when this constraint is binding. Taking into account the expressions of probabilities given in (17) and (19), we obtain the optimal bonus

$$w^*(H) = \frac{c}{(p(1) - p(0))(1 - \beta + \beta p(1))}.$$

An agent is paid a higher wage when the signal on his high performance is biased ($\beta > 0$), since he is sometimes paid even when his true underlying performance is low. The *per capita* liability rent that an agent receives with an individual compensation is thus

$$\begin{aligned} r^*_{IPE} &= (Prob(H,H|1,1) + Prob(H,L|1,1))w^*(H) - c \\ &= \frac{p(0) + \beta((1 - p(0))(1 - p(1)) + \gamma)}{(p(1) - p(0))(1 - \beta + \beta p(1))} c. \end{aligned}$$

We can now assess the relative advantage of individual contracts and tournaments by comparing liability rents in the two scenarios.

Proposition 3. *With no production externalities, the optimal tournament strictly dominates individual contracts if and only if the signal on performances is sufficiently biased, i.e.,*

$$\frac{\beta}{1-\beta} > \frac{1 - p(1)}{(1 - p(1))^2 + \gamma}. \tag{21}$$

If mistakes in evaluating good performances are sufficiently likely, individual contracts become too costly because they require a large bonus that is often wasted. A tournament is then preferred since it instead rewards an agent *more precisely* when his true performance can be distinguished as being above that of his peer.

In our limited liability context, tournaments are more attractive when the correlation parameter γ increases as the right-hand side of (21) is then lower. Following Lazear and Rosen (1981), the early literature on tournaments (Green and Stokey, 1983; Nalebuff and

Stiglitz, 1983; Mookherjee, 1984; Milgrom, 1988) has stressed the role played by the correlation of performances when agents are risk-averse. The value of tournaments is then in filtering noise, thereby reducing the risk premium that the principal must pay to risk-averse agents when designing incentives. At the same time, however, tournaments, especially involving many agents, make rewards rare, and highly discontinuous, thereby creating another form of risk. This explains why they do not systematically dominate piece-rate contracts. The idea that tournaments can be used to filter common risk has been extended to preferences exhibiting ambiguity aversion by Kellner (2015).

An important application of tournaments regards the internal labor of the firm and the incentives induced by the promotion system (Lazear and Rosen, 1981; Belzil and Bognanno, 2008; Lazear, 2018; DeVaro, 2006). For reasons beyond the scope of the present analysis, a principal may want to organize the firm into a pyramidal structure (a hierarchy).[10] Such an organizational form induces agents to seek promotion opportunities, and the implied incentive structure is that of a multi-stage tournament, as emphasized in Rosen (1986). Malcomson (1984) and Prendergast (1999) underline other desirable properties of tournaments, in particular related to additional incentive constraints on the principal's side; a point we discuss in Section 6 below. It has also been argued that competition for promotion allows the principal to sort the workers who are best qualified to occupy a job into a higher layer and thus have efficient dynamic allocation properties (Lazear and Rosen, 1981; Gibbons and Waldman, 1999). Broadly speaking, tournament theory has impregnated the organizational and managerial literatures (see Connelly et al., 2014, for a review).

On the empirical side, Lavy (2009) presents evidence about the positive effect of tournaments on teachers' efforts and performances in Israel as compared to providing no incentives at all.[11] Several authors have studied empirically the incentive effects of tournaments and the composition of pools of contestants (e.g. Ehrenberg and Bognanno, 1990; Brown, 2011).

Tournaments stand as specific examples of contests. The literature on contests has

[10]For instance, Sah and Stiglitz (1986) picture organizations as groups of decision-makers prone to mistakes, Bolton and Dewatripont (1994) study costly communication flows, and Garicano (2000) analyzes the knowledge structure of firms. They all show that a hierarchy is an efficient way to process and aggregate information.

[11]In this context, notice that it is not clear why a tournament is used since objective performance measures, namely students' test-scores, are used to rank teachers so that there is no additional value in the comparison. One reason could be that a principal does not know well what should reasonably be expected for the test-score. Using a tournament washes out this common uncertainty.

grown over the recent years. We refer the reader to Corchon (2007) and Konrad (2009) for a comprehensive overview which is beyond the scope of this paper.[12] This literature takes as given that contests are the sole vehicles for incentives and its focus is on the properties of the agents' equilibrium efforts.

3.4 *CPE* and the Free Riding Problem

Alchian and Demsetz (1972) observed that, when individual contributions to the firm's output cannot be disentangled and only an aggregate measure of performance can be used, the provision of incentives may be undermined by a significant free riding problem. This free riding problem in moral hazard environments echoes the early mechanism design literature on public good provision (Groves, 1973; Groves and Ledyard, 1977; Green and Laffont, 1979) which also discusses similar issues and provides solutions in structured environments. In this mechanism design literature, agents may conceal their willingness to pay for the public good, resulting in underprovision. In the moral hazard scenario we study, underprovision is manifested as limited effort. To align individual incentives with the firm's objectives, a solution would be to make each agent a residual claimant for the firm's output. However, this requires distributing the firm's proceeds n times, which would violate the budget balance.

Solutions to this *"moral hazard in teams"* problem, a term introduced by Holmström (1982), may nevertheless exist. Alchian and Demsetz (1972) found there a *raison d'être* for the hierarchical structure of the firm. Indeed, a principal could act as a monitor for the agents' individual contributions. Unfortunately, such a solution might require setting up a costly monitoring technology. It also fails to explain why corporations are owned by outside shareholders with limited information on actual operations, who may face, on the contrary, difficulties in metering individual contributions. Accordingly, Holmström (1982) suggested that this principal could simply act as a budget-breaker even when individual performances are not observed. In other words, the separation between ownership and control could just be a response to the free riding problem.

In this vein, Holmström (1982) also argued that the principal might use additional information on individual performances to enhance incentives The *"Informativeness Principle"* also operates in a multi-agent context. As discussed above, this principle still stands even though such extra information is related to the other agents' individual output con-

[12] See also Siegel (2009) and Olszewski and Siegel (2016) for recent contributions.

tributions and is thus endogenously and strategically produced, a point made by Gjesdal (1982) and Mookherjee (1984).

In this respect, an interesting question is whether the principal suffers any loss from not observing individual contributions. Surprisingly, the answer is sometimes negative. To see why, suppose that the principal could only observe the aggregate output of a team so that the outcomes (H, L) and (L, H) cannot be distinguished. This constraint on the observability of individual performances implies that, for any feasible incentive scheme, $w(H, L) = w(L, H)$. In fact, there are circumstances under which this requirement does not constrain the optimal contract. As seen in Propositions 1 and 2, there is indeed no loss in not using individual performances provided the optimal scheme stipulates zero payment when either outcomes (H, L) or (L, H) is realized.

Proposition 4. *If the production technology exhibits complementarity and/or individual performances are negatively correlated, so that a CPE scheme is optimal, the fact that individual performances are non-observable entails no additional agency costs for the principal.*

This proposition, even though it appears quite extreme and specific to our setting, suggests that free riding might be less of a concern when there are production complementarities. In the extreme case of a Leontieff production analyzed in Vislie (1994), each agent is necessary to the production process, and hence perceives the full benefit of his effort. Incentives are not undermined compared to the social optimum. In a model with multidimensional inputs and multidimensional outputs, Battaglini (2006) also shows that free riding disappears provided the output dimension is sufficiently large compared to the input dimension. In a series of important papers, Legros and Matsushima (1991), Legros and Matthews (1993), d'Aspremont and Gerard-Varet (1998), Rahman and Obara (2010) and Rahman (2012) presented conditions on the information structure of generic partnerships that allow efficiency, individual incentives and budget balance to be reconciled.

The free riding problem may also be undermined in repeated relationships. For instance, MacLeod (1988) demonstrated that implicit contracts can provide correct incentives, even with budget-balanced sharing rules. This line of research again culminated in the characterization of information structures that allow efficiency to be reached in various repeated game environments. The literature in this field is very substantial, and although it is clearly outside the scope of this paper, let us mention the seminal contributions by Fudenberg, Levine and Maskin (1994) for the case of public monitoring, while Kandori and Obara (2006) and Deb, Li and Mukherjee (2016) investigated the case of private monitoring. By assumption, these papers focus on scenarios without a third-party principal to provide incentives. In this regard, Rayo (2007) endogenized the role of a principal in partnerships in a model of relational contracts. Let us also mention the original contribution by Gershkov, Li and Schweinzer (2009), which bridges the gap between partnerships and contests. In their partnership model, these authors show that an information structure providing only ordinal information on performances can be used to design an efficient contest, thereby overcoming the moral hazard in teams problem.

In more complex environments involving both moral hazard and adverse selection, Picard and Rey (1987) and McAfee and McMillan (1991) observed that a principal can perform equally well by monitoring the overall team's performance as by tracking individual contributions.

Halac, Kremer and Winter (2021) argue that the irrelevance of monitoring individual contribution in Proposition 4 is due to the focus on partial implementation. Exerting effort is an equilibrium, but other equilibria might exist with a CPE scheme, as we discuss in Section 4.1.2 below. Starting from this important observation, these authors insist on full implementation (i.e., exerting effort is the unique equilibrium) to study the optimal design of monitoring teams. In this respect, Mookherjee and Reichelstein (1992) already observed that the value of monitoring by the principal might be that it avoids bad equilibrium outcomes. The design of a monitoring structure then complements pure wage incentives.

3.5 Agent's Preferences

This section addresses the robustness of our earlier findings to various changes in the specifications of agents' preferences.

3.5.1 Heterogenous Agents

The fact that agents are heterogenous in terms of their ability might have a perverse impact on the provision of incentives. In the context of golf tournaments, Brown (2011) reported evidence that the participation of high-skilled players has a negative effect on the performance of less talented ones. In contrast, Hansen (1997) studied service representatives of financial institutions and found that the performance of highly technological agents declines (while low technological agents' performance increases) when individual compensation is replaced with team rewards. A simple extension of our model helps explain such findings.

For simplicity, consider a setting with no production externalities. With CPE, an agent's incentives are boosted when his peer succeeds. Instead, RPE generates the opposite effect. Therefore, a low-skilled agent is more incentivized when facing a high-skilled agent under CPE, while the opposite holds with RPE. Thus, whenever ability is observable, low-skilled workers should receive higher bonuses under RPE while more talented agents should receive higher bonuses under CPE. Since ability cannot always be contracted upon, the motivation of either high- or low-skilled workers (depending on the context) might be undermined with interdependent incentives. Production externalities may modify these results. As an illustration, observe that agent i's incentive constraint under CPE could become less stringent if having a less able teammate means that $Prob(R_i = H|1,0)$ falls by more than $Prob(R_i = H|1,1)$. This could be the case if success on the weak teammate's project critically depends on agent i's effort.

A recent literature has started exploring further the question of heterogeneity in teams through the lens of incentive theory (Glover and Kim, 2021; Au and Chen, 2021; Upton, 2023), in the setting of relational contracting that we develop in Section 5.4.

3.5.2 Other-Regarding Preferences

The issue of designing group incentives has also been approached from a behavioral perspective (see e.g. Baker, Jensen and Murphy, 1988; Fehr and Fischbacher, 2002). One approach is to introduce considerations for others' well-being into each agent's utility function, allowing for interdependent individual utility functions. An immediate consequence is that optimal incentive schemes should also be interdependent, even in the absence of any technological or informational externalities. As Milgrom and Roberts (1992) put it, *"a given level of pay may be viewed as good or bad, acceptable or unacceptable, depending on the compensation of others in the reference group, and as such may result in different behavior. [...] This is a constraint on the use of any sort of incentive pay."* This point was already recognized by Frank (1984). Interestingly, this concern was also raised on more normative grounds by Meyer and Mookherjee (1987) long before the practical implications of inequity-aversion were high on the research agenda.

Itoh (2004) extensively discusses behavioral modifications to standard utility maximization that are important for group incentives. Considerations of fairness, inequity-aversion, altruism, and similar factors are highly relevant in workplace interpersonal relationships. Various theoretical papers address these issues, and the terminology varies somewhat (see, among others, Grund and Sliwka, 2005; Rey-Biel, 2008; Chillemi, 2008;

Dur and Sol, 2010; Englmaier and Wambach, 2010; Bartling, 2011). Dhillon and Herzog-Stein (2009) modeled status-seeking as rank-dependent utility, while Dur and Tichem (2015) focused on altruism in hierarchies. Notably, Bartling and von Siemens (2011) did not find significant experimental evidence of the impact of wage inequality, but Babcock et al. (2015) discovered important social effects in teams. In the context of informational feedback, Mohnen, Pokorny and Sliwka (2008) connected peer pressure to inequity aversion.

Beyond inequity-averse preferences, status and envy also call for incentive distortions. Status-seeking preferences arise when agents enjoy being "ahead" in terms of wages (see for instance Auriol and Renault, 2008, and Charness, Masclet and Villeval, 2014, for experimental evidence), while envy corresponds to a scenario in which receiving a lower wage than a peer brings disutility.

Following Itoh (2004), the (ex post) utility of an inequity averse agent i can be expressed in such contexts as:

$$U_i = w_i - \alpha \max\{w_{-i} - w_i, 0\} - \lambda\alpha \max\{w_i - w_{-i}, 0\}. \quad (22)$$

The non-negative parameter α reflects the intensity of other-regarding preferences while $|\lambda| \leq 1$ captures either inequity-aversion ($\lambda \geq 0$) when the agent is ahead ($w_i \geq w_{-i}$), though at a lower rate than when he is behind, or concerns for status ($\lambda \leq 0$). Finally, we assume that projects are completely independent to focus on the pure effect of other-regarding preferences. The next proposition follows from Itoh (2004).

Proposition 5. *When agents have other-regarding preferences, the optimal scheme entails the following properties.*

- *If $\lambda > p(1)(1 - p(1))$, CPE is preferred as in the absence of other regarding preferences.*
- *If $\lambda < \min\{p(1)(1 - p(1)), \frac{1}{\alpha} - \frac{p(1)(1-p(0))}{p(0)(1-p(1))}\}$, RPE is preferred and the corresponding optimal reward $w^*(H, L)$ decreases in α and increases in λ.*

The optimal incentive scheme responds to social preferences in an intuitive way. When agents are sufficiently inequity-averse, CPE is optimal. Instead, when status-seeking concerns are strong enough (λ sufficiently small, or α relatively small), RPE is preferred.

3.5.3 Overconfidence

Another important behavioral aspect that recently came high on the research agenda comes from the fact that agents may be overconfident about their own performances. Following a large literature in psychology, the effects of overconfidence biases have been explored both from theoretical and experimental viewpoints. That motivation and self-confidence interact is uncontroversial, but modeling attempts are still relatively sparse.[13] When overconfident agents tend to provide more effort, it may be desirable for the principal to boost their self-image. One aspect that is of particular interest in the design of multi-agent incentive schemes is how the principal can take advantage of agents' mistaken beliefs and/or manipulate these beliefs. With mistaken beliefs, Santos-Pinto (2008) shows that the principal may benefit from using interdependent schemes even in absence of any technological or informational externalities.[14] The intuition for the argument is clear. If an agent is overconfident about his own ability compared to that of a peer, an RPE scheme will require a lower payment than IPE. Similarly, if this agent overestimates the ability of his peer, he is ready to accept a lower collective bonus in any CPE scheme. Note that this line of reasoning needs not even rely on an agent's *own* overconfidence *per se*; it is mostly a matter of biased beliefs concerning relative ability.

4 Incentive Schemes as Game Forms

4.1 Simple Environments

Incentive schemes define normal-form games played by agents within the contours of the organization. Understanding the strategic properties of these games certainly informs us on what sort of behavior is expected in the workplace. To explore these issues and for the sake of simplicity, we will first consider a case featuring only informational externalities. Indeed, we saw in Section 3.2 that the corresponding information structures lead to particularly sharp predictions on the optimality of either RPE or CPE. The game-theoretic and efficiency properties of these schemes are very different from the agents' viewpoint. This section explores these properties.

[13] For two interesting contributions based on different premises, see Bénabou and Tirole (2003) and de la Rosa (2011).

[14] Other theoretical contributions include Gervais and Goldstein (2007) regarding teams and Santos-Pinto (2010) regarding tournaments. Field and laboratory experiments seem to confirm the importance of overconfidence in tournaments (Camerer and Lovallo, 1999; Park and Santos-Pinto, 2010) and teams (Vialle, Santos-Pinto and Rullière, 2011).

4.1.1 Prisoner's Dilemma with RPE

First, recall that the probabilities of the different signal pairs available are given by (9) and (10). When the correlation parameter γ is positive, we already know that the optimal incentive scheme entails RPE; an agent receives a bonus when he succeeds while his peer fails. The optimal bonus is given by (15), while the *per capita* limited liability rent is given by (16). Note that this rent is strictly lower than the rent under IPE (namely $r^*_{IPE} = \frac{p(0)c}{p(1)-p(0)}$) where agents are treated independently, as soon as $\gamma > 0$. RPE induces the game form described in Table 2 below.

Table 2: Strategic form under optimal RPE (positive correlation)

	$e_2 = 1$	$e_2 = 0$
$e_1 = 1$	r^*_{RPE}, r^*_{RPE}	$r^*_{RPE} + \frac{p(1)}{1-p(1)}c, r^*_{RPE}$
$e_1 = 0$	$r^*_{RPE}, r^*_{RPE} + \frac{p(1)}{1-p(1)}c$	$r^*_{RPE} + \frac{p(0)}{1-p(1)}c, r^*_{RPE} + \frac{p(0)}{1-p(1)}c$

The strategic structure of this game is that of a *Prisoner's Dilemma* with exerting effort ($e_i = 1$) being a (weakly) dominant strategy. Observe that the payoff vector reached when agents shirk (effort pair $(0,0)$) Pareto dominates the $(1,1)$ equilibrium outcome from the agents' viewpoint. Yet, $(0,0)$ is not a Nash equilibrium of that normal form game. As Section 4.1.2 shall demonstrate, the existence of such a Pareto-dominating outcome makes RPE schemes vulnerable to collusion between agents. Finally, observe that the game is not a strict *Prisoner's Dilemma* since $(1,0)$ and $(0,1)$ are also Nash equilibria, and they both Pareto-dominate $(1,1)$. As noted by Ishiguro (2002), it is however easy to knock out these asymmetric equilibria by increasing $w(H,L)$ by an arbitrarily small amount.

4.1.2 Multiple Equilibria with CPE

Consider now the scenario $\gamma < 0$. In this case, we demonstrated above that CPE is optimal. The positive bonus, which is only given when both agents succeed, is defined in (13). The *per capita* limited liability rent is then given by (14). The normal form game so induced by CPE has the structure given in Table 3.

This normal form game, which is essentially a *Stag Hunt* game, has multiple equilibria. Beyond the high-effort equilibrium $(1,1)$, the shirking outcome $(0,0)$ is also an equilibrium (and a third equilibrium in mixed strategies exists as well). This multiplicity is a serious concern from an implementation viewpoint. The principal might indeed be unsure

Table 3: Strategic form under optimal CPE (negative correlation)

	$e_2 = 1$	$e_2 = 0$
$e_1 = 1$	r^*_{CPE}, r^*_{CPE}	$r^*_{CPE} - c, r^*_{CPE}$
$e_1 = 0$	$r^*_{CPE}, r^*_{CPE} - c$	$r^*_{CPE} - \frac{p(0)}{p(1)}c, r^*_{CPE} - \frac{p(0)}{p(1)}c$

as to whether agents will coordinate on her most preferred equilibrium $(1, 1)$. Of course, it is tempting to focus on this outcome because it is also Pareto-dominating for the agents. Yet, the *bad* equilibrium $(0, 0)$, although dominated, is robust to small trembles. To illustrate this point, consider indeed the optimal CPE scheme as defined in Section 3.2. For such a scheme, the incentive constraint of agent i is binding if his peer also exerts effort. This indifference implies that the equilibrium $(1, 1)$ is not robust to small trembles. If for some exogenous reasons, agent $-i$ does not work with a small probability, then agent i's incentive constraint no longer holds and shirking becomes the unique best-reply.

Taking a broader perspective, the multiplicity of equilibria raises a least three issues. First, the principal may want to design, or at least benefit from, an information structure that would help to suppress this multiplicity problem. In Section 4.2.1 below, we discuss how the mutual observability of efforts might help to coordinate agents on the most-preferred effort pair, but also to ensure that the first-best outcome is the one reached. Secondly, with multiple equilibria, agents may want to switch to their most preferred equilibrium which might differ from the principal's most preferred option. A collusive agreement within the organization, be it explicit or implicit, may allow agents to enforce their most preferred outcome at the expense of the organization's objectives. Section 5 addresses this collusion issue. Again, whether efforts are mutually observable (Section 5.3) or not (Section 5.2) has different consequences on the impact of collusion for the organization. Finally, the multiplicity problem also raises concerns on how agents learn to play the principal's most preferred outcome; especially when the latter is not a strict equilibrium as in the case of CPE schemes we discussed above. This theme has received less attention in the literature. An exception is MacLeod (1987) who argued that collective incentives require coordination between agents and that such coordination is likely to necessitate learning about each other's behavior. Unfortunately, learning is not so easy, especially when the $(1, 1)$ equilibrium is not strict as CPE schemes illustrate. As a result, MacLeod (1987) argued that collective incentive schemes might lose some of their strength when learning is made more difficult because workers are mobile across organizations.

4.2 More Complex Environments

An important feature of some organizations is that agents often have superior information about their peers' efforts compared to the principal (a theme extensively studied in the management literature, see for instance Edwards and Ewen, 1996). This section studies how the principal can use this additional information to reduce the cost of incentives.

4.2.1 Mutual Monitoring and First-Best Implementation

When agents mutually observe each other's effort, the principal may want to induce whistle-blowing so as to learn from one agent about his peers. Beyond their productive role, agents then additionally act as reciprocal monitors. Importantly, in this context of mutually observable efforts, the moral hazard problem can be completely overcome by designing an appropriate communication mechanism.

To make our analysis as simple as possible, we suppose that agents work on two independent projects, with neither technological nor informational externalities. This scenario provides a simple benchmark: an individual compensation scheme is optimal if agents do not observe each other's effort.

Suppose instead that agents observe each other's effort, while their efforts remain non-observable to the principal. Perhaps surprisingly, the principal can now construct a normal-form game that alleviates this non-observability. The logic is well known from the implementation literature (Maskin, 1999; see Moore, 1993, for an exhaustive survey). Private information, if it is observable by all parties, is (generally) not an obstacle to efficient contracting. In most of the implementation literature, information is given at the outset. In contrast, in the moral hazard context that is of interest to us, information, here the agents' efforts, is endogenous. A revelation mechanism should thus induce agents to report their mutually observable efforts without, at the same time, perturbating their incentives to exert effort. Ma (1988) has shown that reconciling those two goals is possible, and that the principal can implement first-best efforts at zero cost. A mechanism that reaches that goal unfolds as follows. After choosing their efforts, but before outcomes are realized, agents are asked to report their teammates' action. The mechanism awards a whistle-blower a bonus, above his own cost of exerting effort, for reporting a low effort by his peer. But if it turns out that his peer succeeds, the whistle-blower incurs a larger fine. When properly designed, the mechanism has each agent report truthfully what he has observed on his peer's effort.

To relate this construction to our general thrust on whether incentives should be collective or competitive within the organization, it is useful to understand the structure of the optimal compensation scheme that reaches first-best implementation. On the production side, an agent is compensated for his disutility of effort whenever his peer truthfully reports that he has exerted an effort and this compensation does not depend on his actual performance. To also induce truthful reporting on his peer's effort, this compensation is actually given in expectation. If an agent reports that his peer has exerted an effort and this announcement is confirmed by the favorable performance of the peer, the agent receives a bonus. He is however punished when this announcement is contradicted by a low performance. This *General Principle of Expertise*[15] captures how an agent's announcement, viewed as a monitor for his peer, is benchmarked against the latter's performance. Definitively, such a scheme bears the mark of an RPE.

The next proposition, whose proof is relegated to the Appendix, summarizes our findings.

Proposition 6. *Suppose that agents mutually observe each other's effort and that projects are independent. Then, the principal can implement the first-best levels of effort and fully extract the agent's liability rent at a Nash equilibrium.*

A first issue raised by this kind of mechanisms is that implementation is only partial. Often, a non-truthful equilibrium may exist that is Pareto-dominating for the agents. In this equilibrium, agents report a positive effort when having exerted no effort. Of course, this outcome is not the most preferred one for the principal. Ma (1988) showed however that the principal can design a multi-stage mechanism to eliminate this unwanted equilibrium.[16] In this respect, the simple revelation mechanism proposed in the proof of Proposition 6 also has this desirable property of knocking out the bad equilibrium in which agents shirk and misreport.[17]

[15]Gromb and Martimort (2007) coined this expression. See also, among others, Demski and Sappington (1987), Lambert (1986), and Malcomson (2009) for similar insights.

[16]This point was also noticed by Laffont and Rey (2003), Brusco (1998, 2002), Ma, Moore and Turnbull (1988) and Ishiguro and Itoh (2001) in related contexts. Complex revelation mechanisms have sometimes been questioned for being too involved to be used in practice or for relying excessively on the agents' rationality that supports Nash or subgame-perfect equilibrium behavior. Dewatripont (1993) provided a crisp and insightful discussion of these issues.

[17]Another equilibrium of the simple revelation mechanism proposed in the proof of Proposition 6 however exists, where both agents shirk and report truthfully. This equilibrium can be eliminated as well, if one allows virtual implementation (Abreu and Sen, 1991), in the spirit of Ishiguro (2002). Wages can be slightly adjusted (by arbitrarily small amounts) to uniquely implement the first-best.

The second, perhaps more problematic, issue is that such revelation mechanisms are prone to more explicit collusion.[18] Agents would certainly benefit from choosing zero effort but jointly reporting the opposite to the principal. As pointed out by Brusco (1997), such coordination could be reached by exchanging side-payments. As argued by Towry (2003), the reluctance of firms to base incentives on peer evaluations (Coates, 1998) might be justified by these concerns about collusion. We will come back to this issue of collusion in Section 5.

4.2.2 Sequential Production Processes

The principal may sometimes benefit from the nature of the production process itself to better structure incentives. For example, in R&D activities, which often involve a sequential production process with strong complementarities across different stages, each unit in the supply chain can assess whether its predecessors have exerted effort and act accordingly.

In a team context *à la* Holmström (1982), but with a sequential production process in which agents observe efforts exerted by their predecessors, Miller (1997) and Strausz (1999) demonstrated the existence of balanced sharing rules that uniquely implement first-best efforts. The intuition is straightforward. Each agent exerts effort only if he observes that his predecessors have also done so. The threat that followers on line may stop working when he shirks provides an agent with enough incentives to exert effort on his own.

The same logic also applies in hierarchical contexts where a principal runs the organization for her own benefit. When agents have only imperfect signals on their predecessors' effort in a multi-stage process with strong complementarities, a direct part of an agent's incentives at any given stage of the production process results from the principal's rewards for final completion; but another indirect part of those incentives also comes from the fact that final rewards are only obtained if all followers also exert effort. As demonstrated by Winter (2006), this feature of the production process puts agents moving early on line on lower powered monetary incentives.

To exemplify, consider the following version of our workhorse model where the only possible observable performances are (H, H) and (L, L). In other words, agents work on a single project that may either succeed or fail. To capture the strong complementarity

[18] A point already made in a different context by Tirole (1986) and the subsequent literature on collusion in organizations. See Tirole (1992) and Laffont and Rochet (1997) for overviews.

between the agents' efforts, let us further assume that agents have a symmetric impact on the probability of joint success, namely

$$Prob(H,H|e_1,e_2) = Prob(H,H|e_2,e_1),$$

and, more importantly, that the following property of increasing differences holds

$$Prob(H,H|1,1) - Prob(H,H|0,1) > Prob(H,H|1,0) - Prob(H,H|0,0). \tag{23}$$

This inequality simply means that agent 2's effort is all the more valuable if agent 1 has already exerted effort on his own.

SCENARIO 1: SIMULTANEOUS AND NON-MUTUALLY OBSERVABLE EFFORTS. Here, the optimal mechanism is clearly CPE. The fact that agents do not observe each other's effort precludes the use of revelation mechanisms as in Section 5.3. Denoting by $w(H,H)$ the common bonus in case of success, exerting effort is a Nash equilibrium when the following incentive constraint holds

$$Prob(H,H|1,1)w(H,H) - c \geq Prob(H,H|0,1)w(H,H).$$

The optimal bonus is thus equal to

$$w^*(H,H) = \frac{c}{Prob(H,H|1,1) - Prob(H,H|0,1)} \tag{24}$$

and the corresponding *per capita* limited liability rent is

$$r^*_{CPE} \equiv Prob(H,H|1,1)w^*(H,H) - c = \frac{Prob(H,H|0,1)}{Prob(H,H|1,1) - Prob(H,H|0,1)}c. \tag{25}$$

SCENARIO 2: SEQUENTIAL AND MUTUALLY OBSERVABLE EFFORTS. Consider a scenario where efforts are made sequentially and suppose that agent 2 can now observe agent 1's earlier effort and react optimally. In sharp contrast with Section 4.2.1, we assume that the principal cannot use any revelation mechanism to induce agents to truthfully report what they commonly know on e_1 in the second stage. Yet, while agent 2 is now induced to exert effort when he receives the bonus $w^*(H,H)$ given in (24) and agent 1 has exerted effort, but he has no incentive to do so if agent 1 has instead shirked at an earlier stage.

Indeed, (23) also implies

$$Prob(H, H|0, 1)w^*(H, H) - c < Prob(H, H|0, 0)w^*(H, H).$$

Agent 1, if he shirks, thus induces his follower to shirk as well. Agent 1's incentive constraint thus writes as

$$Prob(H, H|1, 1)w^1(H, H) - c \geq Prob(H, H|0, 0)w^1(H, H).$$

The optimal bonus for agent 1 is thus equal to

$$w^{1*}(H, H) = \frac{c}{Prob(H, H|1, 1) - Prob(H, H|0, 0)} < w^*(H, H).$$

In other words, agents are on higher powered incentives as they come later in the production process. While the limited liability rent received by Agent 2 remains unchanged with sequential timing, Agent 1's rent is now reduced to

$$r^{1*}_{CPE} \equiv Prob(H, H|1, 1)w^{1*}(H, H) - c = \frac{Prob(H, H|0, 0)}{Prob(H, H|1, 1) - Prob(H, H|0, 0)}c < r^*_{CPE}.$$

We can summarize our findings as follows.

Proposition 7. *Suppose that agents act sequentially on a project with strong complementarity and that efforts by earlier agents are observable by followers. Then, the principal optimally offers* CPE *with higher powered incentives to those agents who act later in line.*

4.2.3 Incentives and Discrimination

Consider again SCENARIO 1 as described in Section 4.2.2. Remember that, in this scenario, agents do not observe each other's effort. Unfortunately, the CPE scheme with bonus $w^*(H, H)$ also suffers from equilibrium multiplicity. Indeed, it is straightforward to check that, under the assumption of strong complementarity between the agents' efforts made in (23), $(0, 0)$ is also a Nash equilibrium since

$$Prob(H, H|1, 0)w^*(H, H) - c < Prob(H, H|0, 0)w^*(H, H).$$

Insisting on unique implementation of the high-effort Nash equilibrium imposes a cost on the principal. In response, Winter (2004) proposed an elegant solution of the multiplicity problem that relies on a discriminatory treatment of otherwise symmetric agents. Suppose indeed that agent 1 is given a high-powered bonus that induces his effort as a dominant strategy, i.e., irrespectively of agent 2's own effort. The minimal resulting bonus should satisfy

$$w^{1*}(H,H) = \frac{c}{Prob(H,H|1,0) - Prob(H,H|0,0)}.$$

In other words, if agent 1 is pessimistic and believes that agent 2 shirks, then he should be indifferent between exerting effort or not. If agent 2 instead works (and he does so at equilibrium), then work should also be the preferred option for agent 1, since it becomes more valuable when the technology features strong complementarities as in (23). Agent 1's work creates a positive externality on Agent 2. In turn, the corresponding bonus for the latter would remain as

$$w^{2*}(H,H) = \frac{c}{Prob(H,H|1,1) - Prob(H,H|0,1)}.$$

Such asymmetric treatment of agents ensures unique implementation. This technique can be easily generalized to more than two agents. In this case, the principal publicly commits to a ranking of agents and designs specific compensations for each of them; a *"divide-and-conquer"* strategy. An agent now finds it optimal to exert effort because he expects all higher-ranked peers to do so and does not care about the efforts of lower-ranked peers.

Proposition 8. *Suppose that symmetric agents work on a project with strong complementarity and that they do not observe each other's effort. Then, the principal can uniquely implement the effort by all agents by offering a discriminatory CPE scheme that publicly ranks agents and offers higher ranked agents higher powered incentives.*

4.2.4 Transparency

In Section 4.2.2, the monitoring structure was taken as given. The degree of transparency that should prevail within an organization is certainly an important design question both from a theoretical viewpoint and for practitioners.[19] Proper job design, for instance by having agents see each other's effort on the job, may thus affect not only incentives to exert effort but also the sort of social relationships that will develop within the organization. Several contributions have sought to endogenize an organization's degree of

[19] See for instance the empirical studies of Ichino and Maggi (2000) and Falk and Ichino (2006).

transparency. On this front, Mohnen, Pokorny and Sliwka (2008), Winter (2010) and Bag and Pepito (2012) have shown that, under complementarity between the agents' efforts, the cost of providing incentives decreases with the degree of transparency. Cato and Ishihara (2017) also offered a careful analysis of the link between transparency and the optimal incentive scheme in sequential production settings. Gershkov and Winter (2015) have shown that, again with complementarity and sequential provision in efforts, CPE schemes can use peer monitoring to substitute for direct monitoring by the principal himself. Again, more transparency among peers allows subsequent agents to punish shirking by their predecessors. This effect acts as a substitute for direct monitoring. A very specific piece of information that might be relevant for some production process is where agents rank in a line. A crucial assumption is indeed that the principal can publicly commit to ranking agents. Halac, Lipnowski and Rappoport (2021) showed that the principal may want to keep the ranking of agents secret to reduce agency costs. Finally, these authors also argued that a principal may use her monitoring ability to uniquely implement effort provision as a Nash equilibrium among agents. At a more abstract level, Jehiel (2014) showed that full transparency in an organization is in general not desirable when there is moral hazard.

5 Collusive and Cooperative Behaviors Call for CPE

A manufacturer using franchising contracts with retailers located in different cities may expect direct interaction between stores to be limited. In this case, our previous working assumption that agents adopt a Nash behavior seems reasonable. An upstream manufacturer can design incentive schemes for those retailers, who maximize their own payoffs independently from what others are doing. However, interdependent schemes tie agents' interests together and agents may thus benefit from coordinating actions. The extent to which they can do so and whether it is at the expense or to the advantage of the principal depends on the internal organization of the firm. Assuming that agents adopt a non-cooperative behavior becomes unreasonable when agents entertain close and repeated interactions. Designs of incentive schemes for sellers who work in the same store every day must take into account the possibility of collusion and/or ongoing cooperation among them. This section addresses these issues and roughly shows that whether collective behavior is something that the principal wants to combat (the case of collusion addressed in Sections 5.2 and 5.3 below) or from which it might benefit (the case of cooperation presented in Sections 5.4 and 5.5), incentive schemes should be tilted towards

CPE.

As an introduction to the theoretical investigation that follows, it is worth noting that empirical studies have shown cases where firms chose collective incentives despite independent agent technologies (Herries, Rees and Zax, 2003). Additionally, there are instances where the introduction of collective incentives improved overall performance (Chan, Li and Pierce, 2014; Hamilton, Nickerson and Owan, 2003). The underlying mechanisms are broadly referred to as *peer effects*, which essentially corresponds to the idea that collective schemes align agents' interests, and facilitate their cooperative behavior in various forms.

5.1 Modeling Collusion: An Overview

In this section, we investigate how agents may form a coalition to coordinate and promote their own objectives rather than those of the organization. In so doing, agents aim at reaching an outcome that they find more attractive than what they would get when acting non-cooperatively. To model such collusive behavior, we might assume that agents can design side-contracts stipulating how to trade side-transfers conditional on final outcomes. Such side-contracts would be agreed upon just after the contours of the organization and especially compensation schemes are known, but before efforts are exerted. On top, efforts may or may not be specified by side-contracts depending on whether agents observe each other's efforts or not. Writing effort targets into a side-contract may not be feasible either because they are not mutually observable or because, even if they are, there is no mechanism that make them internally verifiable by the coalition of agents.[20] Finally, side-contracts are supposed to be enforceable.

Each of these assumptions deserves some comment. First, modeling collusion as an enforceable side-contract is of course an extreme assumption and we refer to Tirole (1992) and Laffont and Martimort (1997) for some discussion on its limits. Sometimes, side-payments are not even feasible, but the simple repetition of their relationship may suffice to induce cooperative behavior between agents. Enforceable side-contracts can then be viewed as reduced forms for the collusive ongoing relationships that agents may entertain. Side-transfers are thus very similar to continuation payoffs in the self-enforcing scenario with the extra requirement that side-transfers are generally budget-balanced,

[20] Baron and Besanko (1999) discuss this notion of internal verifiability in the context of an adverse selection model.

while continuation payoffs are not bound to such constraints but must instead be self-enforcing.[21] Secondly, when efforts within the coalition cannot be observed or verified for inclusion in a side-contract, it implies that agents do not have a comparative advantage over the principal in terms of monitoring. This issue is further explored in Section 5.2. Additionally, the requirement for side-transfers to be budget-balanced means that colluding agents cannot create payment opportunities beyond what the principal originally had available. In essence, these assumptions place colluding partners and the principal on an equal footing as far as contracting is concerned.

5.2 Weak Collusion: Non-Mutually Observable Efforts

Suppose that colluding agents cannot mutually observe their efforts. By colluding, agents can actually change the game induced by the incentive schemes, and possibly its equilibria. These new strategic possibilities add extra constraints to the principal's problem since incentive schemes must be designed so that agents cannot improve their expected payoffs by colluding. In other words, incentive schemes should be robust to the threat of collusion.

Proposition 9. *When agents do not observe each other's effort, the principal can never benefit from the possibility of side-contracting between the agents.*

The intuition for this result is straightforward and somehow relies on the *Collusion-Proofness Principle* as stated in Tirole (1986) and Laffont and Martimort (1997) in related contexts (with adverse selection rather than moral hazard but the same logic applies). Even though a formal proof is out of the scope of this paper, we should here recall that this *Principle* claims that there is no loss of generality in restricting the analysis to compensation schemes that cannot be improved upon by the collusion process. To illustrate the logic, let us assume that agents have the option to create side-contracts and exchange side-transfers $t_i(R_i, R_{-i})$ based on observable signals on top of any incentive scheme $\tilde{w}_i(R_i, R_{-i})$ already offered by the principal. Following an approach initially introduced by Laffont and Martimort (1997) in the context of adverse selection, suppose that agents collaborate to maximize the total of their payoffs while ensuring that side-transfers remain budget-balanced. In this scenario, an optimal side-contract is one that implements an effort pair (e_i, e_{-i}) that maximizes the sum of the agents' payoffs, subject to the incentives constraints, since effort are not mutually observable. Additionally, it must satisfy

[21] On the explicit modeling of self-enforcing collusion, see also Tirole (1992) and Martimort (1999).

participation constraints and, notably, outperform non-cooperative behavior. Suppose the principal now offers in the first place the equivalent payments in each state of nature, i.e. she offers transfers $w_i(R_i, R_{-i}) = \tilde{w}_i(R_i, R_{-i}) + t_i(R_i, R_{-i})$. This scheme also implements the coalition-maximizing effort pair. The optimal side-contract that would be agreed upon by the agents if they were offered such a scheme initially is obviously the null side-contract, which proves the *Collusion-Proofness Principle*. Yet, the coordinated choice of side-payments and possibly the joint choice of efforts that it induces means that, at best, this allocation might have to satisfy further coalition-incentive compatibility constraints beyond the familiar individual incentive constraints. The discussion after Proposition 10 below nicely illustrates this fact.

Proposition 9 can be easily demonstrated if we momentarily depart from assuming limited liability as a source of agency frictions. Consider instead the case of risk aversion. We already alluded to the structure of the optimal compensation scheme in Section 2.2 above. Of particular importance is the fact that, in this scenario, the optimal wages satisfy $w^*(H, L) > w^*(L, H)$. This implies that agents might trade side-payments when performances are mixed so as to reach mutual insurance.[22] The requirement of mutual co-insurance thus imposes that a collusion-proof incentive scheme necessarily entails $w(H, L) = w(L, H)$. As a result, the cost of providing individual incentives to exert effort is necessarily increased. This point was made in a variety of contexts by Holmström and Milgrom (1990), Itoh (1993), Varian (1990), Ramakrishnan and Thakor (1991) and Macho-Stadler and Perez-Castrillo (1993).

Returning now to our bare-bone model with limited liability to further investigate the cost of collusion, consider the case of informational externalities and the optimal schemes found in Section 4.1. Because of the strategic differences in the game forms induced for agents by RPE and CPE, collusion has different effects on these incentive schemes.

Proposition 10. *Suppose $\gamma < 0$. The optimal CPE incentive scheme is immune to the possibility of side-contracting.*

Suppose $\gamma > 0$. The optimal RPE scheme is not immune to side-contracting. The possibility of collusion increases agency costs under RPE schemes.

Consider the case $\gamma < 0$. Note that the equilibrium effort pair $(1, 1)$ induced by CPE also maximizes the sum of the agents' payoffs. Therefore, agents have no incentives to sign a side-contract since reaching another effort pair would be Pareto-dominated.

[22] Such side-contract may also be signed after efforts are chosen but before performances are realized.

Instead, *RPE* schemes are strongly affected by the possibility of collusion. To illustrate, consider now the case $\gamma > 0$. From Section 3.2, the optimal *RPE* entails a positive bonus

$$w^*(H,L) = \frac{c}{(1-p(1))(p(1)-p(0))}$$

which is offered only in case of success combined with the failure of the peer. With this scheme, the incentive compatibility constraint is binding. Consider a side-contract that specifies a positive side-payment $\hat{t} \equiv t_1(H,L)$ from agent 1 when successful to agent 2 when unsuccessful, and vice-versa. This side-contract thus establishes a system of reciprocal favors. The side-payment $\hat{t}$ effectively eliminates any motivation to exert effort and prevents the $(1,1)$ equilibrium outcome. Indeed, with the above value of $w^*(H,L)$, inducing a unilateral deviation by one agent from a putative equilibrium with high efforts requires

$$\begin{aligned}(p(1)(1-p(1))-\gamma)(w^*(H,L)-\hat{t}+\hat{t})-c \\ < (p(0)(1-p(1))-\gamma)(w^*(H,L)-\hat{t})+((1-p(0))p(1)-\gamma)\hat{t},\end{aligned}$$

which boils down to $\hat{t} > 0$.

Instead, implementing $(0,0)$ as a collusive equilibrium requires that an agent never wants to deviate by unilaterally exerting effort, i.e.,

$$\begin{aligned}(p(0)(1-p(0))-\gamma)(w^*(H,L)-\hat{t}+\hat{t}) \\ \geq (p(1)(1-p(0))-\gamma)(w^*(H,L)-\hat{t})+((1-p(1))p(0)-\gamma)\hat{t}-c\end{aligned}$$

or $\hat{t} \geq \frac{c}{1-p_1}$. In particular, a side-payment that transfers all wage from the winning agent to the losing peer, namely $\hat{t} = w^*(H,L)$, implements the collusive equilibrium $(0,0)$ while keeping in check liability constraints.

We can now express the condition that ensures agents choose to exert effort rather than deviate towards the collusive shirking outcome. The corresponding coalition-incentive compatibility constraint is[23]

$$(p(1)(1-p(1))-\gamma)w(H,L)-c \geq (p(0)(1-p(0))-\gamma)w(H,L).$$

[23]The attentive reader will have certainly noticed that we have implicitly restricted the analysis to collusion-proof *RPE* schemes in writing this coalition-incentive compatibility constraint in its simplest form. Tedious computations show that this restriction is warranted.

The wage that implements the effort pair $(1, 1)$ at minimal agency cost is thus[24]

$$w^c(H,L) = \frac{c}{(1-p(1)-p(0))(p(1)-p(0))}.$$

It is straightforward to check that $w^c(H, L) > w^*(H, L)$ since coalitional incentive compatibility is now harder to satisfy than individual incentive compatibility. Intuitively, by not exerting effort, an agent exerts a positive externality on his peer and their collusion internalizes this externality.

The basic thrust here is that competitive environments, as defined by RPE schemes, are also the most prone to collusive behavior between agents to avoid such competitive stances. This logic also mirrors similar findings from the adverse selection literature as, for instance, in Laffont and Martimort (2000).

To conclude this section, let us mention that a number of other organizational issues studied in the multi-agent moral hazard framework are outside of the scope of this survey because they are not directly related to the question of whether collective or relative compensations should emerge. For instance, Macho-Stadler and Perez-Castrillo (1998) and Baliga and Sjöström (1998) study the incentive issues that arise when the principal can structure her organization vertically, i.e., a contract with agent 1 who himself subcontracts with agent 2. This hierarchy of contracts is one way for the principal to institutionalize existing collusion; a point also made elsewhere by Faure-Grimaud, Laffont and Martimort (2003).

5.3 Strong Collusion: Mutually Observable Efforts

We now come back to the model in Section 4.2.1 that entails neither technological nor informational externalities. Again, the competitive environment that is put in place by the principal when agents monitor each other's effort is fragile when collusion is a concern. Collusion is now facilitated by the fact that agents observe each other's effort and can design a side-contract that stipulates not only some (balanced) side-transfers but also which efforts they should collectively abide to. There is ample experimental evidence showing that agents develop strong reciprocity in the workplace (Fehr and Gächter, 2000; Fehr, Gächter and Kirchsteiger, 1997) and that such a behavioral norm certainly helps them to enforce side-agreements. Any collusive side-contract can now condition payments not

[24]Note that the assumption $p(1)(1-p(1)) > p(0)(1-p(0))$ is needed to implement $(1, 1)$ as a collusion-proof allocation. This assumption in turn requires $p(1) + p(0) < 1$.

only on outcomes but also on efforts. Such a strong collusion allows agents to behave as a syndicate, following the terminology coined elsewhere by Wilson (1968). We might ask how the optimal incentive scheme and the principal's expected payoff are modified by these new side-contracting possibilities. These questions were first investigated by Varian (1990) and Ramakrishnan and Thakor (1991) assuming that agents are risk-averse. CPE becomes optimal. Of course, having agents behave as a syndicate remains more costly than relying on the individual incentives involved had agents stuck to their non-cooperative behavior, since we know from Section 4.2.1 that reciprocal monitoring allows first-best implementation.

We now make a similar point in a model with risk-neutral and symmetric agents protected by limited liability. Recall that projects are assumed to be independent. Since agents can now agree on a side-contract that stipulates which efforts they should exert, they jointly choose those efforts to maximize the sum of their expected gains. Moreover, this side-contract can also redistribute wealth across agents for each possible realized outcome $\boldsymbol{R} = (R_i, R_{-i})$. Considering that a strong coalition forms, the principal now provides collective bonuses conditional on the aggregate output only. Let us denote by $W(\boldsymbol{R})$ the non-negative collective bonus when $\boldsymbol{R}$ is realized. That bonuses are based only on the aggregate output imposes $W(H,L) = W(L,H)$. Given the symmetry of the problem, there is no loss of generality in assuming that agents *ex ante* agree on sharing the sum of their expected gains equally. Taking into account the obvious fact that the principal will set $W(L,L) = 0$, the expected compensation of the strong coalition associated with the effort pair (e_1, e_2) is:

(26)

$$\mathbb{E}_{\boldsymbol{R}}\left[W(\boldsymbol{R})|e_1,e_2\right] = p(e_1)p(e_2)W(H,H) + (p(e_1)(1-p(e_2)) + p(e_2)(1-p(e_1)))W(H,L).$$

The optimal incentive scheme must induce the coalition to choose the effort pair $(1,1)$; that is, it must at the same time prevent an individual deviations towards $(0,1)$ and a global deviation towards $(0,0)$. The first condition boils down to a familiar individual incentive compatibility

$$\mathbb{E}_{\boldsymbol{R}}\left[W(\boldsymbol{R})|1,1\right] - \mathbb{E}_{\boldsymbol{R}}\left[W(\boldsymbol{R})|0,1\right] \geq c, \tag{27}$$

while the second condition is truly a *coalition-incentive compatibility constraint*, namely

$$\mathbb{E}_{\boldsymbol{R}}\left[W(\boldsymbol{R})|1,1\right] - \mathbb{E}_{\boldsymbol{R}}\left[W(\boldsymbol{R})|0,0\right] \geq 2c. \tag{28}$$

The principal's problem consists in minimizing the expected compensation of the coalition (26) for the effort pair $(1, 1)$, subject to the incentive compatibility constraints (27) and (28), together with the usual limited liability constraints

$$W(\boldsymbol{R}) \geq 0 \quad \forall \boldsymbol{R}. \tag{29}$$

Comparing the incentive efficiencies, it is easily shown that using the collective wage $W(H, H)$ for joint performance always dominates using $W(H, L) = W(L, H)$ for mixed realizations. Intuitively, the most informative signals on the fact that the agents have collectively exerted effort is precisely when both projects succeed. Therefore, the optimal incentive scheme is a CPE scheme. The coalition-incentive compatibility (28) is binding and the optimal incentive scheme for the coalition is thus

$$W^*(H, H) = \frac{2c}{(p(1) + p(0))(p(1) - p(0))} \quad \text{and} \quad W^*(H, L) = W^*(L, H) = W^*(L, L) = 0. \tag{30}$$

The *per capita* limited liability rent under this collective scheme is thus

$$r^*_{CPE} \equiv p(1)^2 \frac{W^*(H, H)}{2} - c = \frac{p(0)}{p(1) + p(0)} \frac{p(0)c}{p(1) - p(0)}. \tag{31}$$

This rent is lower than what an agent receives with IPE, namely

$$r^*_{IPE} = \frac{p(0)c}{p(1) - p(0)}.$$

Therefore, the principal is better off when agents act as a strong coalition.

Proposition 11. *Suppose that agents work on two independent projects, observe each others' effort and behave as a strong coalition. The optimal incentive scheme is a CPE scheme.*

The principal is better off with such a collective scheme than when agents behave non-cooperatively under IPE.

To understand the optimality of CPE, it is useful to draw a parallel with the optimal incentive scheme that an agent receives when working simultaneously on two tasks. Laux (2001) shows that if the principal finds both tasks valuable, the agent should receive a positive bonus only when both tasks are successful.[25] In other words, limited liability

[25] On similar agency issues of task assignments, see Schmitz (2005) and Iossa and Martimort (2015) among others.

rents exhibit economies of scope. The fact that the principal is better off when facing a coalition than when facing two isolated agents can be understood as follows. When tasks are independent and agents make decisions separately the optimal incentive scheme can be collective, competitive or independent. A collective (resp. competitive) scheme makes agents efforts complementary (resp. substitutes). These positive (resp. negative) externalities associated with efforts matter if agents can make decisions cooperatively because they will be internalized. This makes the incentives provided by a CPE (resp. RPE) scheme more (resp. less) effective when agents behave as a coalition.

Various results related to Proposition 11 have been expressed in contexts where the source of agency costs is the trade-off between insurance and incentives (Holmström and Milgrom, 1990; Varian, 1990; Itoh, 1992; Macho-Stadler and Perez-Castrillo, 1993). The benefits of coalitional behavior then come from a decreased risk premium in comparison with IPE.

Of course, the best outcome from the principal's viewpoint remains when agents (be they risk-neutral and protected by limited liability or risk-averse) maintain a non-cooperative behavior *and* the principal can use revelation mechanisms to have agents report each other's effort as in Section 4.2.1. However, a strong coalition allows agents to collude on their messages on each other's effort as well as on their efforts, and this makes those revelation mechanisms irrelevant.

5.4 Repeated Interactions and Induced Cooperation Between Agents

The analysis of Sections 5.2 and 5.3 relies on the strong assumption that agents are able to sign side-contracts in order to enforce their coalitional behavior. As discussed earlier on, this assumption is essentially a shortcut. An alternative approach is to view the agents' collective behavior as being achieved through repeated interactions in long-term ongoing relationships. Following the lines of the arguments in Section 5.3, inducing agents to rely on those self-enforcing side-contracts may be viewed as an attractive outcome for the principal, at least in comparison to the opportunities offered by IPE.

Che and Yoo (2001) study a model of long-term interaction in which cooperation can emerge as a subgame-perfect equilibrium of a repeated game.[26] Suppose that the princi-

[26]See Levin (2002) for a related contribution that we discuss in the Section 6.2.2. We follow them by keeping the assumption of the previous section that agents perfectly monitor each other, but assume now that they cannot sign side-contract. The idea that teamwork can be facilitated in repeated relationships can also be found in MacLeod (1988). More recently, Kvaløy and Olsen (2019) have extended the analysis of

pal commits to a stationary incentive scheme once and for all at the beginning of the relationship. Notice that history-dependent contracts are thereby ruled out, maybe because they are viewed as too complex to implement. Agents then play an infinitely repeated version of the game form so induced. Che and Yoo (2001) demonstrate that CPE implements efforts at lower costs than RPE by exploiting the long-term relationship between agents. Even under circumstances in which RPE would be optimal in a static framework, if agents are sufficiently patient, CPE can become optimal for the principal in a dynamic setting.

This result finds its roots in the strategic properties of CPE and RPE presented in Section 4.1. Under CPE, $(1,1)$ is the unique Pareto-optimal effort pair for the agents, and agents coordinate in the long run on the desirable equilibrium for the principal. Under RPE, the agents are trapped into a *Prisoner's Dilemma* when playing $(1,1)$ while they would prefer to coordinate on $(0,0)$, and they can do so in a collusive equilibrium of the infinitely repeated version of the game. This possibility is costly to prevent for the principal. As a result, RPE in a dynamic setting tend to increase agency costs.

We now develop more formally this argument in the framework of Section 4.1.1 which features informational externalities with positive correlation. We already know that, under those circumstances, RPE is optimal in a one-shot version of the game. In contrast with Section 4.1.1, we shall assume that efforts are mutually observable, but that the principal does not use a revelation mechanism like the one discussed in Section 4.2.1 to extract such information.

Consider any arbitrary compensation scheme $\boldsymbol{w}$ and the stage game it defines. When played only once, $(1,1)$ is a Nash equilibrium when the familiar static incentive compatibility constraint holds:

$$\mathbb{E}_{\boldsymbol{R}}\left[w(\boldsymbol{R})|1,1\right] - c \geq \mathbb{E}_{\boldsymbol{R}}\left[w(\boldsymbol{R})|0,1\right]. \tag{32}$$

Consider now the infinitely repeated version of this stage game and denote by δ the discount factor common to all players. The principal seeks to implement the infinite repetition of $(1,1)$ as a subgame-perfect Nash equilibrium. The first question is whether the mere repetition of the relationship reduces (per period) agency cost compared to the static context.

Che and Yoo (2001) by enriching the information structure and have shown that some forms of RPE can emerge as optimal relational contracts.

By shirking, an agent can always secure in any period a payoff equal to $\min_{e_{-i}} \mathbb{E}_{\boldsymbol{R}}\left[w(\boldsymbol{R})|0,e_{-i}\right]$. Hence, the following *dynamic incentive constraint* must hold in any putative subgame-perfect Nash equilibrium implementing $(1,1)$:[27]

$$\text{(33)} \qquad \mathbb{E}_{\boldsymbol{R}}\left[w(\boldsymbol{R})|1,1\right] - c \geq (1-\delta)\mathbb{E}_{\boldsymbol{R}}\left[w(\boldsymbol{R})|0,1\right] + \delta \min_{e_{-i}} \mathbb{E}_{\boldsymbol{R}}\left[w(\boldsymbol{R})|0,e_{-i}\right].$$

The left-hand side represents the agent's expected utility stream over an infinite time horizon within the equilibrium where both agents are working. The first term on the right-hand side of (33) is agent i's payoff from deviating for the current period. The second term is the payoff that agent i can guarantee himself by always shirking in the continuation. The incentive constraint (33) is only a necessary condition for $(1,1)$ to be a subgame-perfect equilibrium. Indeed, we have not yet confirmed that agent $-i$ would credibly choose an effort that minimizes agent i's payoff following the latter's deviation to shirking. Nevertheless, comparing the right-hand sides of (32) and (33) already provides some insight into whether the principal can reduce agency costs in a repeated setting.

Consider first the case of an *RPE* scheme with $w(H,L) > 0 = w(H,H) = w(L,H) = w(L,L)$. The right-hand side of (33) then takes a very simple form. When agent $-i$ exerts an effort, it reduces the likelihood of agent i being rewarded for outperforming him compared to when he shirks. Consequently, agent $-i$ minimizes agent i's payoff in the continuation when exerting effort:

$$\min_{e_{-i}} \mathbb{E}_{\boldsymbol{R}}\left[w(\boldsymbol{R})|0,e_{-i}\right] = \mathbb{E}_{\boldsymbol{R}}\left[w(\boldsymbol{R})|0,1\right].$$

Furthermore, since the stage game with *RPE* has a (non-strict) *Prisoner's Dilemma* structure, exerting an effort for agent $-i$ together with shirking for agent i constitute a subgame-perfect equilibrium in the continuation. Hence, the dynamic incentive compatibility constraint (33) boils down to its static counterpart. The corresponding wage that implements effort at minimal agency cost is thus

$$w^{\infty}(H,L) = w^{*}(H,L)$$

and each agent obtains the same per-period liability rent as in a static framework:

$$r^{\infty}_{RPE} = r^{*}_{RPE}.$$

[27]Payoffs in the dynamic relationship are average discounted payoffs so that they are comparable to payoffs achieved in the one-shot version of the game.

In other words, with *RPE*, the principal cannot save on agency costs compared to the static scenario.

In fact, the principal could even do worse. Indeed, under *RPE*, the pair of efforts $(1,1)$ is Pareto-dominated from the agents' viewpoint. In this repeated setting where the *Folk Theorem* applies, agents could coordinate on a collusive equilibrium $(0,0)$ if δ is sufficiently close to 1. To do so, it would be enough that, if an agent exerts effort while equilibrium play requires not to do so, his peer retaliates by also exerting effort in all future periods.

Consider now a *CPE* scheme with $w(H,H) > 0 = w(H,L) = w(L,H) = w(L,L)$. As seen in 4.1.2, the effort pair $(1,1)$ is a Pareto-dominant Nash equilibrium of the stage game, and it is also the principal's preferred option. However, $(0,0)$ is also a Nash equilibrium, and the worst one from the agents' viewpoint. In the repeated setting, agents can thus credibly revert to playing $(0,0)$ to sustain coordination on the Pareto-dominant situation. The corresponding dynamic incentive compatibility constraint takes the form

$$\mathbb{E}_{\boldsymbol{R}}\left[w(\boldsymbol{R})|1,1\right] - c \geq (1-\delta)\mathbb{E}_{\boldsymbol{R}}\left[w(\boldsymbol{R})|0,1\right] + \delta\mathbb{E}_{\boldsymbol{R}}\left[w(\boldsymbol{R})|0,0\right]. \tag{34}$$

Under *CPE*, when agent $-i$ shirks, it is less likely that agent i is rewarded, and shirking by agent $-i$ thus minimizes agent i's payoff in the continuation. Consequently,

$$\mathbb{E}_{\boldsymbol{R}}\left[w(\boldsymbol{R})|0,1\right] < \mathbb{E}_{\boldsymbol{R}}\left[w(\boldsymbol{R})|0,0\right].$$

Hence, (34) is easier to satisfy than its static counterpart (33). Therefore, a *CPE* scheme implements the principal's preferred equilibrium at a lower agency cost in a dynamic setting. Saturating the dynamic incentive constraint (34), the optimal wage must satisfy

$$(p(1)^2+\gamma)w^\infty(H,H) - c = (1-\delta)(p(1)p(0)+\gamma)w^\infty(H,H) + \delta(p(0)^2+\gamma)w^\infty(H,H),$$

which yields the following expression for the optimal bonus under *CPE*:

$$w^\infty(H,H) = \frac{c}{(p(1)+\delta p(0))(p(1)-p(0))}.$$

This bonus is always lower than the static *CPE* bonus given in (13). The *per period and* *capita* limited liability rent is now

$$r^{\infty}_{CPE} \equiv (p(1)^2 + \gamma) w^{\infty}(H,H) - c = \frac{p(0)((1-\delta)p(1) + \delta p(0)) + \gamma}{(p(1) + \delta p(0))(p(1) - p(0))} c \leq r^{*}_{CPE},$$

with a strict inequality when $\delta > 0$.

While relying on RPE cannot reduce agency costs compared to the static benchmark, CPE becomes even more attractive in the dynamic setting, especially as δ increases. Comparing the two schemes leads to the following proposition.

Proposition 12. *In a repeated relationship where agents mutually observe each other's efforts, CPE is preferred to RPE without collusion, even with positive correlation, if*

$$0 \leq \gamma \leq \frac{\delta p(0)}{1 + \delta p(0)} p(1)(1 - p(1)).$$

With CPE, the principal benefits from the fact that the agents' interests are aligned with his own (because $(1,1)$ is a Pareto-dominating profile for the three players) and observe each other's effort to implement a credible threat of reverting to a low-effort equilibrium if anyone deviates. Since payoffs are low in such continuation, effort is induced at a lower cost. Put differently, CPE allows the principal to somewhat delegate the disciplining role of contracts to the agents themselves. The associated cost reduction is higher that the benefits of RPE–paying a positive wage only when the most informative outcomes is realized–provided performances are not too positively correlated.

Economists have recently found empirical evidence of how collusive and/or coalitional behavior affect the efficiency of incentive schemes. Using personnel data from a fruit farm, Bandiera, Barankay and Rasul (2005) stressed that moving from competitive to independent incentives increased the productivity of the average worker by 50%. They showed that the weak effectiveness of the competitive incentive scheme was due to the fact that agents were able to monitor each other's efforts and to enforce collusive behavior detrimental to the organization. Symmetrically, Knez and Simester (2001) argued that the increase in employee performance at Continental Airlines following the introduction of collective incentives was imputable to a raise in mutual monitoring among employees within work groups. In both cases, the long-term relationship leveraged the benefits of CPE.

5.5 Further Approaches on Cooperation

5.5.1 Cooperation as an Organizational Choice

Some authors (see for instance Holmström and Milgrom, 1994; Wageman and Baker, 1997; Wageman, 2001; Mukherjee and Vasconcelos, 2011) have argued that incentives should be jointly designed with other features of an organization such as job design, degree of autonomy or level of transparency. In this respect, we now investigate how the principal should choose an incentive system that determines not only compensation schemes but also whether agents should cooperate or not. From our earlier findings, it follows that, if the principal relies on RPE, she should keep agents apart so as to avoid collusion. A *separated system* is preferable. On the contrary, with CPE, the principal should make agents work together in order to induce beneficial coalitional behavior. An *integrated system* is preferred. In other words, the degree of coordination between agents and the structure of incentives should be optimized altogether.

In this context, Ramakrishnan and Thakor (1991) argued that the optimal choice is somewhat biased towards CPE. To see why, let us come back to the case of informational externalities as developed in Section 3.2. Remember that the liability rent r^*_{RPE} in the separated system when run with RPE writes as in (16) which is a decreasing function of γ. Consider instead a CPE system in which agents cooperate when choosing effort. The corresponding *coalition incentive compatibility constraint* must now prevent a joint deviation where both agents shirk. The liability rent r^*_{CPE} now writes as in (31) which is increasing in γ. A CPE scheme is then optimal when $r^*_{CPE} \leq r^*_{RPE}$.

Proposition 13. *The principal prefers an integrated (resp. separated) system provided that the correlation parameter γ satisfies*

$$\gamma \leq \frac{p(0)}{1+p(0)} p(1)(1-p(1)) \text{ (resp. } \geq\text{)}.$$

An implication is that for γ positive but small, CPE is optimal if agents adopt a cooperative behavior even though RPE would have been optimal otherwise.

Maskin, Qian and Xu (2000) also addressed the relative value of separated and integrated systems to identify how well they take advantage of correlations when RPE is used.

5.5.2 Help and Sabotage

Beyond their own effort, agents might also undertake actions that affect the probability of success for their peers. These actions can be beneficial (*help*) or detrimental (*sabotage*). In contrast with our earlier analysis, the focus is now on how CPE and RPE affect incentives for agents to help or sabotage, rather than to exert productive effort on their own projects. In fact, CPE turns out to be more effective in both cases as noted by Itoh (1991), Milgrom (1988), Lazear (1989), Kandel and Lazear (1992), Macho-Stadler and Perez-Castrillo (1993), Drago and Garvey (1998) and Crama, Sting and Wu (2019). Quite intuitively, this stems from the fact that CPE provides incentives to undertake actions that are beneficial from a collective viewpoint. Agents will help teammates when it is optimal and they will not waste resources trying to undermine others' work. Competitive schemes provide the opposite incentives which can clearly be inefficient from the principal's point of view, as argued early on by Dye (1984). The rich literature on tournaments and contests has also examined the cost of sabotage (Chen, 2003; Münster, 2007). Experimental studies such as Harbring and Irlenbusch (2011) provide evidence of such sabotage. Social psychologists have long explored the relationship between what they term *"task interdependence"* (which corresponds to the idea that subjects can and should help each other) and the effectiveness of collective and competitive incentives in enhancing performance. Various experimental studies in the field have actually echoed the work of economists (Miller and Hamblin, 1963).

6 Incentivizing the Principal Calls for RPE

Thus far, we have assumed that incentives were only a concern for the agents. However, in a variety of settings, the principal's incentives matter as well for organizational performances. First, the principal herself can actively influence the outcome by taking actions that constitute moral hazard variables, subsequently impacting her agents' chances of success. Secondly, the principal's assessment of her agents' performance may not be based on objective, quantifiable measures incorporated into a contract, as previously assumed. In such cases, the principal might be tempted to manipulate these subjective evaluations to minimize incentive payments, subsequently undermining the credibility of incentive schemes and discouraging agents from exerting effort. This section explores the implications of these additional incentive-related challenges for the principal's optimal incentive schemes. We find that various forms of RPE can be beneficial in addressing these issues.

6.1 Double Moral Hazard

Suppose that the principal is able to affect agents' productivity through a non-verifiable effort. An example in order is given by Gould, Pashigian and Prendergast (2005) who studied contracts proposed by a mall developer (the principal) to the owners of indoor stores (agents) who sell products. These authors argued that contracts should incentivize both owners and developers since the latter's unobservable actions can impact store sales (e.g., cleanliness, renovation, parking design, etc.). Another example is that of a global franchise: the effort of the global franchisor, together with those of the local franchisees, determine the success in each local market (Nocke and Strausz, 2023). This two-sided moral hazard situation has received some attention in the extant literature with notable contributions by Carmichael (1983), Al-Najjar (1997), Gupta and Romano (1998), Itoh (1994) and Tsoulouhas (1999). These papers argue that moral hazard on the principal's side tilts the choice of compensation schemes towards RPE. To present this argument, we now adapt the framework of Gupta and Romano (1998) to our workhorse model.

Suppose that the agents work on two independent projects. We know that if the only incentive problem pertains to the agents, IPE suffices in this context, and each agent receives a positive bonus $w(H,H) = w(H,L) = w(H)$ for good performances. The additional consideration here is that each agent's project returns depend on their own effort and a binary effort, denoted as $e_p \in 0,1$, which is exerted by the principal and affects both projects. Success is more likely when the principal also exerts effort. If the principal shirks and chooses $e_p = 0$, agent i's probability of success when choosing an effort e_i is $q(e_i)$, while it is $p(e_i) > q(e_i)$ if the principal undertakes effort and chooses $e_p = 1$. We denote by $c_p e_p$ the principal's disutility of effort.

Since this part of our analysis focuses on the principal's incentives, we need here to explicitly define her benefit from a project's success. Let us denote by B (resp. 0) the principal's *per project* benefit when it succeeds (resp. fails). With IPE, the principal's incentive constraint is thus

$$2p(1)(B - w(H)) - c_p \geq 2q(1)(B - w(H)).$$

The wage $w^*(H)$ that leaves an agent's indifferent between exerting effort or shirking, while anticipating the principal's exerting effort, is determined by the familiar condition

$$w^*(H) = \frac{c}{p(1) - p(0)}.$$

Hence, the principal exerts effort whenever

$$2(p(1) - q(1))B \geq c_p + \frac{2(p(1) - q(1))}{p(1) - p(0)}c. \tag{35}$$

This condition shows how providing incentives to both the principal and the agents are two conflicting goals for the organization. Indeed, all parties should be rewarded when both projects succeed, i.e., when (H, H) realizes. Incentivizing agents in that event requires giving them the wage $w^*(H)$ but, at the same time, doing so hardens the principal's incentive constraint as can be seen from the right-hand side of (35). This logic is even exacerbated if CPE is used instead of IPE.

This incentive conflict can nevertheless be mitigated with RPE. Because an agent is now rewarded only following mixed performances, the principal can partially decouple his own incentive problem from those of the agents. To confirm this point, consider an RPE scheme that pays an agent only when it outperforms his peer. The corresponding optimal wage $w^*(H, L)$ when anticipating effort from the principal is given by the familiar condition

$$w^*(H, L) = \frac{c}{(1 - p(1))(p(1) - p(0))}.$$

The principal's incentive constraint with RPE is thus

$$2p(1)\left(B - (1 - p(1))w^*(H, L)\right) - c_p \geq 2q(1)\left(B - (1 - q(1))w^*(H, L)\right),$$

or

$$2(p(1) - q(1))B \geq c_p + \frac{2(p(1) - q(1))(1 - p(1) - q(1))}{(1 - p(1))(p(1) - p(0))}c. \tag{36}$$

Comparing the right-hand sides of (35) and (36) is immediate: the principal's incentive constraint is easier to satisfy with RPE. The next proposition follows:

Proposition 14. *Suppose that agents work on independent projects and that the principal can exert a non-contractible effort that increases the projects' probability of success. Then, RPE provides more incentives to the principal than IPE.*

6.2 Subjective Evaluations

A major limit in the formal provision of incentives is the degree of verifiability of the agents' performances. In this respect, MacLeod (2003) pointed out that, when performances are non-verifiable by courts and parties may have subjective views of them, bonuses for high performances are hardly feasible. Indeed, principals would tend to always claim that agents have failed to avoid paying the promised bonus even following good performances. The agents' compensation schemes become less effective when such manipulations by the principal herself are a concern. An illustration, reported by Levin (2003), is given by law firms where discretionary bonuses are often manipulated.

6.2.1 Tournaments

In a multi-agent context, Section 3.3 has already shown how tournaments perform well when the relative measurement of performances is more accurate than its absolute counterpart. In addition, tournaments also overcome the subjective evaluation problem. Rewards are no longer based on a subjective measure of the agents' performances but on the ranking between those performances and such rankings are presumably harder to manipulate. In a tournament, or more generally in any promotion system, the aggregate bonus paid to the agents is independent of the principal's subjective evaluations of their individual performances. Under-evaluating those performances is no longer attractive for the principal and the agents' incentives are preserved as noted by Malcomson (1984, 1986) and Prendergast (1993).

In a bilateral context, MacLeod (2003) showed that burning money restores the principal's incentives to report her agent's performance truthfully while, at the same time, it preserves the latter's incentives to exert effort. In order for the principal to report truthfully, she must pay a fixed amount, irrespective of her report. In the optimal contract, the agent always receives this amount as a bonus, except following his worst performance, in which case money is burned. In a model with risk aversion, and two agents working on independent projects, Rajan and Reichelstein (2006) showed that, instead of burning money, the principal could simply reward a successful peer in the case where one agent fails. Doing so helps to satisfy the principal's incentive constraint while it also boosts this peer's effort. In other words, RPE schemes endogenously emerge with subjective evaluations.

Before closing this section, it is worth looking back at the comparison made by Lazear and Rosen (1981) between tournaments and individual compensation schemes through the lens offered by this section. The two contracting modes actually correspond to different degrees of manipulability concerning what is observable by the principal. Consider first IPE. With this contracting mode, the principal signs a bilateral contract with each agent that does not rely on any information that might be collected by the principal from observing performances by peers. This incompleteness assumption is akin to assuming that the peer's performances are non-verifiable by the court of law in charge of enforcing this bilateral contract. The reason may be that those bilateral contracts are signed sequentially and that the courts of law in charge of enforcing each of them have different auditing capabilities. Yet, the principal and each agent could write a bilateral contract contingent on the principal's report of what she learns from the other relationship.[28] The point is that the principal's ability to manipulate such announcements makes these message-contingent contracts irrelevant. An IPE scheme protects against the principal's manipulative behavior vis-à-vis each agent. Now consider tournaments. As we noted above, the principal's overall payment is now independent from the agents' precise array of performances, and only the ranking of those performances matters. This is another incompleteness assumption that limits what a multilateral contract can do, although a tournament remains a multilateral contract. The comparison between tournaments and individual compensation schemes then boils down to the comparison of two constraints on the kind of monitoring technology available and their consequences on manipulability.

6.2.2 Multilateral Relational Contracting

Consider now a scenario where the performance of agents cannot be verified, even though it is common knowledge. In this context, no formal contract can be enforced by a court. The only option is to rely on implicit promises for good performance. Unfortunately, the principal may still be tempted to renege on such promises. However, repeated interactions with her agents may somewhat discipline the principal. An agent can now use the threat of leaving the relationship as a means of enforcing the promise for a bonus. The logic behind such a *relational contract* is well known and has already been laid down elsewhere to explain ongoing contracting relationships on the market place (Macaulay, 1963; Klein and Leffler, 1981, among others). As pointed out by Baker, Gibbons and Murphy (2002), Levin (2003) and MacLeod (2003), relational contracts are often viewed as the

[28] See Dequiedt and Martimort (2015) for an analysis of such bilateral contracts in a mechanism design context.

glue of an organization in contexts where complete contracts cannot be written.[29] Effort is only rewarded when the principal's commitment to paying a bonus is credible. This means that the benefits of continuing the relationship should outweigh the immediate gains of not paying the bonus. This *self-enforcement constraint* limits the size of feasible bonuses. In turn, a lower bonus can have a negative impact on the agents' incentives.

Levin (2002) pointed out that the threat of ending the relationship is even stronger in a multi-agent context. Indeed, all agents may stop working following a deviation from the principal with only one of them, if it is publicly observable. In this case, the agents' bonuses must satisfy an aggregate self-enforcement constraint instead of bilateral self-enforcement constraints that would pertain to each of the bilateral relationship. This effect, akin to the famous multi-market contact phenomenon of the IO literature,[30] makes it easier to control the principal's opportunistic behavior. Levin (2002) then argues that the choice of incentive schemes is tilted towards RPE. Intuitively, with an aggregate self-enforcement constraint, the principal now faces an upper bound on the sum of bonuses she can credibly promise under all possible performance configurations. CPE exacerbates the principal's incentives to jointly deviate and not give any bonus when both agents succeed. With RPE, the incentives to deviate by not paying bonuses are similar to those achieved in bilateral relationships and are thus weakened. Levin (2002) showed that this effect calls for tournament-like incentive schemes, under which an agent receives a bonus if he outperforms his peer while, when both agents perform poorly, neither gets paid. In a setting closer to ours, Kvaløy and Olsen (2006, 2019) pursued this line of research, confirming that competitive incentives can be useful with multilateral relational contracts.

To illustrate those findings, consider an infinitely repeated version of our static model and suppose that projects are independent. The principal wants to induce both agents to work in each period. We focus on stationary incentive schemes.[31] As stated above, the incentive scheme must also satisfy an aggregate self-enforcement constraint so that the principal can credibly commit to pay bonuses following good performances. This constraint requires that the sum of individual bonuses at each pair $\boldsymbol{R} = (R_1, R_2)$ is less than the continuation value of the multilateral relationship to the principal. We assume as in the previous section that, for each project, B and 0 respectively denote the principal's

[29]See MacLeod (2007), Malcomson (2010) and Watson (2021) for recent surveys of this burgeoning literature.

[30]Bernheim and Whinston (1990).

[31]Our focus on stationarity precludes the possibility of disciplining agents by a threat of firing, i.e., of losing a stream of limited liability rents. Non-stationary review contracts as in Fuchs (2007) are useful in this respect.

return following either good or bad performance, and that B is large enough, so that the principal prefers to induce effort in the static case. The expected instantaneous net benefit per agent can thus be written as

$$p(1)B - \mathbb{E}_{\boldsymbol{R}}\left[w(\boldsymbol{R})|1,1\right] \tag{37}$$

where

$$\mathbb{E}_{\boldsymbol{R}}\left[w(\boldsymbol{R})|1,1\right] = p(1)^2 w(H,H) + p(1)(1-p(1))w(H,L)$$

denotes the expected wage and where expressions have been simplified by using the fact that agents are never paid following low performance, i.e. $w(L,L) = w(L,H) = 0$.

Let δ be the principal's discount factor. Following our earlier convention, we also normalize payoffs by the length of the period $1-\delta$. A multilateral relational contract should not only incentivize agents to exert effort in each period but also prevent the principal from cutting or reducing due bonuses. The first goal requires that the compensation scheme satisfies the familiar static moral hazard incentive constraint:

$$p(1)w(H,H) + (1-p(1))w(H,L) \geq \frac{c}{p(1)-p(0)}. \tag{38}$$

For the second objective, the principal should not gain by cutting bonuses, knowing that doing so induces both agents to terminate the relationship. We insist here on the fact that the relational contract is multilateral and both agents are ready to retaliate and leave the relationship if the principal does not fulfil his promises. We can thus write the corresponding aggregate self-enforcement constraint at (H,H) as

$$2(1-\delta)(B-w(H,H)) + 2\delta(p(1)B - p(1)^2 w(H,H) - p(1)(1-p(1))w(H,L)) \geq 2(1-\delta)B.$$

Simplifying yields

$$\delta p(1)B \geq (1-\delta+\delta p(1)^2)w(H,H) + \delta p(1)(1-p(1))w(H,L). \tag{39}$$

Proceeding similarly gives us the expression of the aggregate self-enforcement constraint at (H,L) as

$$(1-\delta)(B-w(H,L)) + 2\delta(p(1)B - p(1)^2 w(H,H) - p(1)(1-p(1))w(H,L)) \geq (1-\delta)B$$

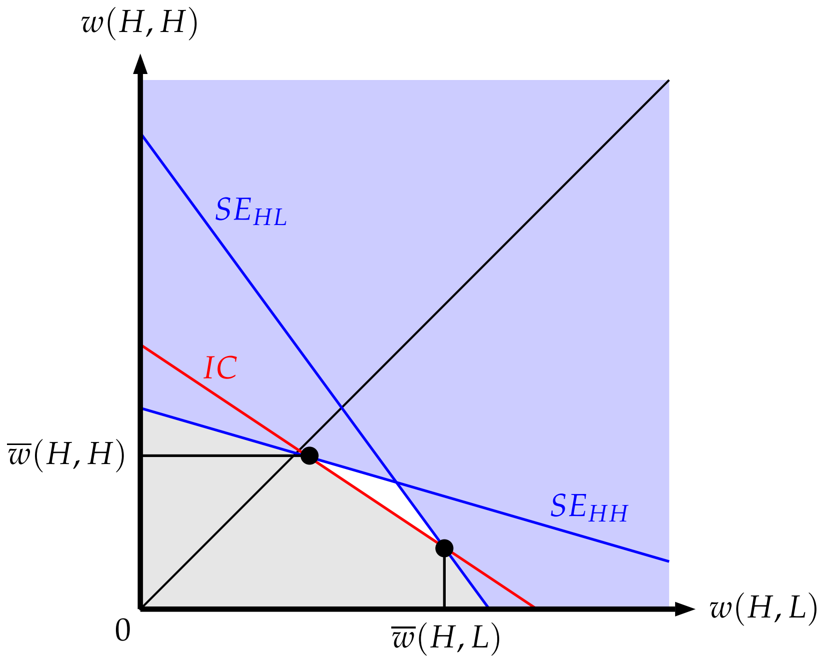

Figure 2: Self-enforceable allocations (unshaded area)

or, after simplifications,

$$(40) \qquad 2\delta p(1)B \geq 2\delta p(1)^2 w(H,H) + (2\delta p(1)(1-p(1)) + 1 - \delta)w(H,L).$$

Constraint (39) (resp. (40)) is denoted as SE_{HH} (resp. SE_{HL}) in Figure 2. We have also represented the agent's incentive compatibility constraint (38). Together, those constraints define the whole set of incentive-feasible bonuses. This set can readily be verified as non-empty provided that

$$\frac{\delta p(1)(2-p(1))B}{1-\delta+\delta p(1)(2-p(1))} \geq \frac{c}{p(1)-p(0)}.$$

Observe that this condition can be expressed as $(p(1)-p(0))B \geq c$ in the limit as δ approaches 1. This condition, which means that exerting effort is efficient, is met when B is sufficiently large. Conversely, for any given B, there exists a small enough δ such that the set of incentive-feasible bonuses becomes empty. In that case, the principal cannot incentivize agents, since her promise to pay wages is not credible if she heavily discounts the future.

An optimal contract should minimize the agent's expected payments. With independent projects, the principal's indifference curves as defined through (37) are parallel to the agent's incentive constraint (38). Figure 2 then shows that a whole range of possible compensation schemes are optimal as long as (38) is binding. Self-enforcement of the bonus requires that the wages $w(H,H)$ and $w(H,L)$ should not be too extreme, namely:

$$w(H,H) \leq \overline{w}(H,H) = \frac{\delta p(1)}{2(1-\delta)}\left(B - \frac{c}{p(1)-p(0)}\right),$$

and

$$w(H,L) \leq \overline{w}(H,L) = \frac{\delta p(1)}{1-\delta}\left(B - \frac{c}{p(1)-p(0)}\right).$$

The upper bound on the feasible collective bonus $w(H,H)$ is thus half the upper bound on the possible competitive bonus $w(H,L)$. This is a simple consequence of the fact that competitive bonuses are now spread out across states (H,L) and (L,H), whereas collective bonuses are both distributed when (H,H) is realized, which exacerbates the principal's incentives to cut bonuses. The next proposition follows.

Proposition 15. *Self-enforceability constraints are more severe with CPE than with RPE.*

Of course, the shape of the set of incentive-feasible schemes as shown in Figure 2 does not imply that RPE is always preferred at the optimum. It indicates that the set of feasible optimal schemes contains a whole range of combinations of RPE and CPE compensations, but they are tilted more towards RPE than in the static case.

7 Market Interactions Might Call for IPE

In the previous sections, our analysis focused primarily on the organization in isolation. However, when agents within the organization interact with labor or output markets, they may encounter various contracting opportunities beyond the organization's boundaries. This section explores how these external possibilities influence the level of competition or cooperation that can be established within the organization. We will examine a couple of scenarios to illustrate how optimal incentive schemes lean toward IPE when facing complex interactions in adjacent markets.

7.1 Career Concerns on the Labor Market

According to Fama (1980), career concerns on the labor market can bring enough discipline to alleviate moral hazard problems within the firm. Holmström (1999) presented a less optimistic view of this argument, arguing that reputational concerns might not suffice to induce efficient managerial effort. In line with these authors, we consider a scenario in which agents' incentives within an organization are solely derived from their career concerns. While formal incentive contracts are not feasible under this approach, it is instructive to analyze how the prospect of the future wages implicitly generates interlinked incentives. Current performances are not directly incentivized but are later used by firms competing with fixed wages to attract the agents. Though agents' own talents are unknown, even to themselves, they can be partially inferred from past performances. Therefore, by exerting effort in the first period, agents can signal high ability to the market and increase the competitive wage they receive later. In a multi-agent context, the inference process and its impact on incentives are complicated by possible informational externalities that arise when one agent's performance influences beliefs about his peers' talents.

To fix ideas and delve into this inference process, let us examine a scenario where the two agents work on independent projects. Using our previous notations, a good performance H yields a benefit $B > 0$ to the principal while a bad performance L yields 0. The added twist of a model with career concerns is that each agent now possesses an unknown talent. When agent i is talented, exerting effort guarantees the achievement of $R_i = H$, while shirking only results in this outcome with a probability of $p \in (0, 1)$. When agent i is unskilled, exerting effort generates $R_i = H$ with the same probability p, while the project always fails when shirking. Principals and agents share a common prior on the distribution of talents. Agents are equally likely to be talented or unskilled, and, importantly, talents are perfectly correlated. Finally, and to simplify modeling, we consider only two periods, and keep in line with our previous convention by assuming that their respective weights in intertemporal payoffs are $1 - \delta$ and δ.

In the first period, agents choose non-cooperatively to work or shirk, anticipating the consequences on their future wages.

In the second period, principals compete for agents with fixed wages $w(\boldsymbol{R})$, that depend only on the first-period performances $\boldsymbol{R}$. Whether IPE or CPE are considered below, incentives to exert effort thus vanish in the second period since wages do not depend on second-period performances. Competition drives these wages to the expected value of the second-period output, conditional on shirking at this date and conditional on what has been learned from first-period performances. The comparison between IPE and CPE thus boils down to comparing the conditions for exerting effort to be equilibrium behavior in the first period. Both agents exerting effort is an equilibrium when the following incentive constraint holds:

$$(1-\delta)(-c)+\delta\mathbb{E}_{\boldsymbol{R}}\left[w(\boldsymbol{R})|1,1\right]\geq\delta\mathbb{E}_{\boldsymbol{R}}\left[w(\boldsymbol{R})|0,1\right]. \tag{41}$$

The wages that competitive principals offer in the second period take into account how the agents' common talent is inferred from their past performances. To illustrate, when the market expects both agents to exert effort in the first period, the occurrence of a poor performance L by either of them early on certainly indicates that neither is talented. Hence, both agents will always produce L for sure in the second period which yields no benefit to principals. When the market uses both signals to infer the agents' talent, we thus have $w(H,L)=w(L,H)=w(L,L)=0$. Instead, the realization of (H,H) in the first period indicates that agents are likely to be talented. Both agents will produce H with some probability in the second period; which yields a positive benefit to principals. The wage $w(H,H)$ is positive and CPE emerges. Following the same logic, when the market instead evaluates each agent in isolation, i.e., IPE prevails, the second-period wage following a bad performance is $w(L)=0$, while $w(H)$ is positive.

Importantly, (H,H) is a stronger signal that agents are talented than a signal H taken in isolation. Applying Baye's rule, the posterior belief that agents' are talented following (H,H) is indeed $\frac{\frac{1}{2}}{\frac{1}{2}+\frac{p^2}{2}}$, while the posterior belief that an agent is talented following a single signal H is only $\frac{\frac{1}{2}}{\frac{1}{2}+\frac{p}{2}}$. Given that a talented agent will produce H in the second period with probability p when shirking, the wages under CPE and IPE are respectively $w(H,H)=\frac{p}{1+p^2}B$ and $w(H)=\frac{p}{1+p}B$. Remarkably, competition among principals drives the *expected* second-period market wage to $\frac{p}{2}B$ whether IPE or CPE prevails even though $w(H,H)>w(H)$. As a result, the left-hand side of (41) is the same under IPE and CPE.[32]

In turn, the gains from shirking, i.e., the right-hand side of (41), differ. Under IPE, shirking allows to obtain the bonus $w(H)$ with probability $\frac{1}{2}p$, i.e. when agents are talented and a high performance H is obtained with probability p. Under CPE, when an agent shirks, (H,H) happens if agents are talented and this deviating agent is lucky, i.e., with total probability $\frac{1}{2}p$ as well. The probabilities of obtaining a positive wage are identical under IPE and CPE, but the corresponding wages are different. Since

[32]The argument is not limited to the specific informational structure that generates here IPE or CPE. *Ex post* competition drives principals' profits to zero for any (common) belief on talent they may hold. Hence *expected* wages in equilibrium must be equal to the *ex ante* value brought by an agent shirking in the second period. The shape of the wage scheme influences only deviation payoffs.

$w(H,H) > w(H)$, shirking is more attractive under CPE than IPE.

Proposition 16. *When abilities are perfectly correlated and production shocks are independent, career concerns endogenously generate CPE and create weaker incentives than IPE.*

An immediate implication of this proposition is that employers would benefit from individually evaluating their (future) workers, although they might not be able to refrain themselves from using all market performances to better infer individual talents. The combination of career concerns with multiple agents and competing principal generates implicit incentives that contrast with those in the optimal incentive contract. Indeed, the fact that talents are positively correlated introduces a subsequent positive correlation in the agents' performance. A single principal would instead prefer to commit to an RPE scheme in a one-period scenario under those circumstances.[33]

7.2 Hold-Up Problems on the Product Market

When property rights are hard to enforce, agents may threaten to leave with the firm's proceeds. Incentive design within a firm can be compromised by this risk of hold-up. This is particularly relevant in the case of software developers, who create valuable intellectual property that they could easily sell to competitors. Similar concerns arise in consulting and law firms, where the threat of employees leaving with the clients they serve is pervasive. Compensation schemes must then also mitigate such opportunistic behavior, in addition to providing on-the-job incentives.

To illustrate how compensation schemes are here also tilted towards IPE, consider the generic version of the model sketched in Section 2. The argument made below is general and applies more broadly to any kind of technological or informational externalities. Suppose that each agent may now leave the firm and sell an output H for a price s on the market. The principal wants to prevent agents from leaving. The implicit assumption here is that the agent's output is more valuable inside the firm than on the market. For instance, the agent may have developed specific human capital or his innovation may be complementary to existing assets of the firm. Since upon obtaining a high output an agent can always quit and earn s on the market, the following *ex post* participation constraints

[33] Indeed, the incentive efficiencies are such that $I(H,H) < I(H,L) = 1$: the outcome (H,L) indicates that agent 1 has worked for sure, and should receive all the incentive weight, making RPE the optimal static contract.

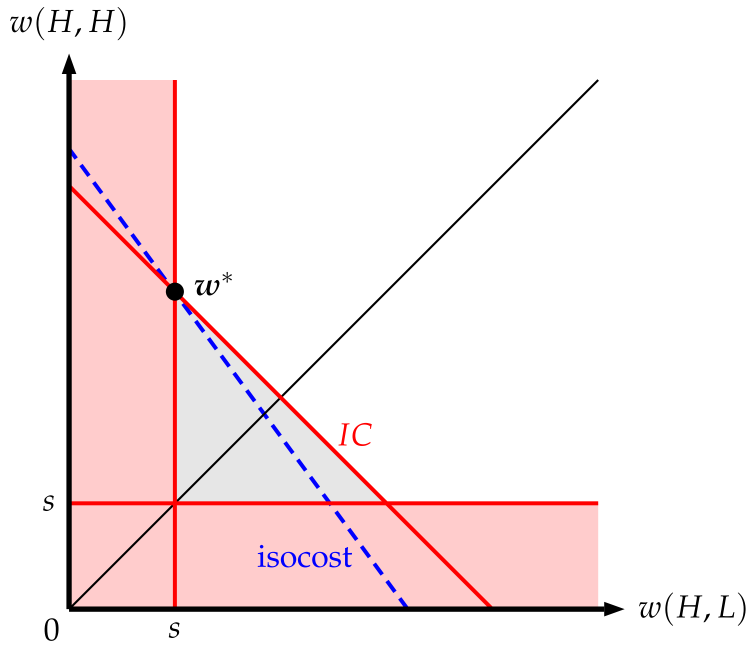

Figure 3: Agent's commitment problem.

must hold to prevent the agent's opportunism:

$$w(H,H),\ w(H,L) \geq s. \tag{42}$$

Of course, these constraints do not modify the incentive efficiency of any signal. Figure 3 gives a clue of the solution to the so constrained optimal scheme. Here, we take s to be small enough so that a fixed bonus $w(H,H) = w(H,L) = s$ would not suffice to ensure incentive compatibility.[34] Note also that we have depicted a scenario where the incentive efficiency at signal (H,H) is greater than at signal (H,L). Hence, pure CPE would be optimal in the absence of *ex post* constraints (42). Unfortunately, a successful agent would then walk away when his peer fails. Constraint (42) requires to smooth bonuses and choose the optimum wages $w^*(H,L)$ and $w^*(H,H)$ as

$$w^*(H,L) = s$$

[34]Formally: $(Prob(H,H|1,1) + Prob(H,L|1,1) - Prob(H,H|0,1) - Prob(H,L|0,1))s < c$.

and

$$(Prob(H,H|1,1) - Prob(H,H|0,1))w^*(H,H) + (Prob(H,L|1,1) - Prob(H,L|0,1))s = c.$$

Meeting the *ex post* participation constraints (42) naturally leads to higher agency costs. As s increases, the appeal of IPE grows, reaching a point where, if s is sufficiently large, IPE becomes the optimal choice, and the incentive constraint becomes slack:

$$w^*(H,H) = w^*(H,L) = s.$$

In this case, the principal can no longer benefit from linking incentive schemes, and only IPE can effectively prevent agents from engaging in hold-up behavior.

Proposition 17. *When agents can appropriate part of the ex post output they produce and walk away from the organization, IPE becomes more attractive.*

The attractiveness of IPE is also confirmed by Kvaløy and Olsen (2012) in a more complex model of repeated interactions à la Che and Yoo (2001) where any agent can leave the firm at any point in time with a share of his output (while Section 5.4 shows that CPE is optimal in the absence of such hold-up constraints).

8 Conclusion

Both scholars and practitioners agree that in complex organizations where agents interact, linking an agent's compensation to their peers' performance may be attractive. The optimal approach depends on the context. Inducing more cooperation or fostering competition between agents may be more effective for achieving the principal's goals.

Most of the mechanisms identified in the extant agency literature have studied heavily stylized environments that favor either purely competitive or purely collective schemes. In contrast, real-life examples abound in which mixed solutions that borrow from both modes are actually implemented. For example, most employees are motivated not only by stock participation and divisional bonuses, a form of collective incentive, but also by promotions, that are definitely more competitive. Sometimes, the organization of the firm itself exhibits this tension. To illustrate, an M-form multidivisional firm is organized through profit centers whose performances are easily comparable by the headquarters, while each profit center is itself organized through sub-teams whose incentives are necessarily more collective. Surprisingly, this issue of the optimal mix of collective and competitive incentives and their consequences for organization design have been by and large neglected in the theoretical literature.[35] We believe that understanding the use and properties of mixed incentive patterns, and uncovering systematic reasons why they could be optimal, remains an important research question.

Another important extension that would certainly deserve more work is related to the dynamics of organizational forms. Whether competition or cooperation prevail in an organization is rarely a permanent trait. Organizations evolve over time because of changing markets, technology or regulatory constraints in their surrounding environment, or because behavioral norms within the firm are changing. The cooperative incentives that apply to long-established organizations operating in stable environments can become vulnerable when these organizations face the threat of competition, whether in output or labor markets. Furthermore, it is not well-understood how the incentives generated by external competition influence internal dynamics within these organizations and how such changes are received by entrenched organizational cliques. This area remains largely unexplored.[36]

Although our discussion of the costs and benefits or collective and competitive incentives were formed with a particular normative criterion in mind, namely the principal's objective, other more positive perspectives are possible. Because the amount of social capital and the norms of reciprocity that might develop among agents also evolve over time, enjoying the benefits of moving towards a more efficient way of providing incentives takes time and organizational resources. Riding the road towards incentive efficiency, a long-term target for the organization, might require fine-tuned management. Entrenched groups within the organization may oppose changes in incentive modes. In other words, far from being based only on the minimization of agency costs, the choice of an incen-

[35] Carmichael (1983) recognized early on that absolute and relative evaluations are often simultaneously used in practice–e.g. *"salesmen who work on commission and also compete for Hawaiian vacations"* (Carmichael, 1983, p. 50), and that mixed patterns deserved attention. The choice of stick vs. carrots has been studied extensively, including in the tournaments literature (see e.g. Moldovanu and Sela, 2001; Moldovanu, Sela and Shi, 2012), but not the optimal mix of collective and relative incentives. Magill and Quinzii (2006) and Fleckinger (2012) have shown that general forms of informational dependence might generate an optimal mix of collective and competitive rewards.

[36] There is nevertheless a very small body of literature that studies how organizational forms respond to outside competitive pressures in the context of models of endogenous growth (François and Roberts, 2003; Martimort and Verdier, 2004). Unfortunately, these papers do not address whether more competitive outside environments have an impact on the degree of cooperation or competition that emerges within the firm.

tive mode may also respond to political considerations, sometimes within society as a whole. Studying this aspect might shed new light on the trade-off between collective and competitive incentives.

If anything, we hope that this survey will inspire more research on these fronts and stimulate readers to explore others.

A Proofs

A.1 Proof of Lemma 1

In the principal's program, let $\lambda > 0$ be the Lagrange multiplier associated with the incentive constraint, and $\mu_{\mathbf{R}} \geq 0$ that associated with the limited liability constraint $w(\mathbf{R}) \geq 0$. The first-order condition for each $w(\mathbf{R})$ is:

$$-Prob(\mathbf{R}|1,1) + \lambda(Prob(\mathbf{R}|1,1) - Prob(\mathbf{R}|0,1)) + \mu_{\mathbf{R}} = 0.$$

If $w(\mathbf{R}) > 0$ then $\mu_{\mathbf{R}} = 0$ and the last equation writes:

$$I(\mathbf{R}) = \frac{1}{\lambda}.$$

If $w(\mathbf{R'}) = 0$ for some outcome R', one obtains $I(\mathbf{R'}) = \frac{1}{\lambda}(1 - \frac{\mu_{\mathbf{R'}}}{Prob(\mathbf{R'}|1,1)}) < \frac{1}{\lambda}$. Hence, the conclusion.

A.2 Proof of Proposition 6

The mechanism presented implements $(e_i, e_{-i}) = (1,1)$, induces both agents to report the truth on their peer's effort, and extracts both agents' rent.

Let us denote by $\hat{e}_{-i}$ agent i's report on $-i$'s effort.

- When agent $-i$ has reported $\hat{e}_i = 1$, agent i receives in expectation (with respect to the realization of R_{-i}) a payment worth c regardless of the realization of R_i. Agent i's payment consists of the following lottery, where ε is positive but arbitrarily small:
  - When he has reported $\hat{e}_{-i} = 1$, agent i receives $c + (1 - p(1))\varepsilon$ if $R_{-i} = H$ and $c - p(1)\varepsilon$ if $R_{-i} = L$.
  - When he has reported $\hat{e}_{-i} = 0$, agent i receives $c - (1 - p(0))\varepsilon$ if $R_{-i} = H$ and $c + p(0)\varepsilon$ if $R_{-i} = L$.
- When agent $-i$ has reported $\hat{e}_i = 0$, agent i is paid nothing whatever his outcome R_i and report $\hat{e}_{-i}$.

Incentive compatibility constraints. Suppose that agent i has exerted effort $e_i = 1$ and agent $-i$ has truthfully reported such information. Telling the truth about agent $-i$'s effort is an equilibrium strategy for agent i that results in an expected payment of c, as claimed above, since the following incentive compatibility constraints hold:

$$\overbrace{p(1)(c+(1-p(1))\varepsilon)+(1-p(1))(c-p(1)\varepsilon)}^{c} \\ > \underbrace{p(1)(c-(1-p(0))\varepsilon)+(1-p(1))(c+p(0)\varepsilon)}_{c-(p(1)-p(0))\varepsilon} \tag{43}$$

and

$$\overbrace{p(0)(c-(1-p(0))\varepsilon)+(1-p(0))(c-p(0)\varepsilon)}^{c} \\ > \underbrace{p(0)(c+(1-p(1))\varepsilon)+(1-p(0))(c-p(1)\varepsilon)}_{c-(p(1)-p(0))\varepsilon}. \tag{44}$$

Suppose that agent i has shirked ($e_i = 0$) and agent $-i$ has truthfully reported such information. Agent i receives no payments and telling the truth on agent $-i$'s effort is thus again an equilibrium strategy that yields zero payoff.

From there, it follows that exerting effort and telling the truth is an equilibrium strategy for agent i and all his rent is extracted since

$$p(1)(c+(1-p(1))\varepsilon)+(1-p(1))(c-p(1)\varepsilon)-c=0.$$

Limited liability constraints. Observe that the above payments are either zero or close to c, and thus positive when ε is small enough.

Destroying the equilibrium with shirking and misreporting. When both agents have shirked, and agent $-i$ has instead reported $\hat{e}_i = 1$, the incentive compatibility constraint (44) implies that agent i strictly prefers to report $\hat{e}_{-i} = 0$. This incentive constraint therefore destroys the "bad" equilibrium with shirking and mutual covering up.

References

Abowd, John M., and David S. Kaplan. 1999. "Executive Compensation: Six Questions That Need Answering." *Journal of Economic Perspectives*, 13(4): pp. 145–168.

Abreu, Dilip, and Arunava Sen. 1991. "Virtual Implementation in Nash Equilibrium." *Econometrica*, 59(4): 997–1021.

Admati, Anat R., and Paul Pfleiderer. 1997. "Does It All Add Up? Benchmarks and the Compensation of Active Portfolio Managers." *Journal of Business*, 70(3): 323–50.

Aggarwal, Rajesh, and Andrew A. Samwick. 1996. "Executive Compensation, Strategic Competition, and Relative Performance Evaluation: Theory and Evidence." National Bureau of Economic Research, Inc NBER Working Papers 5648.

Aggarwal, Rajesh K., and Andrew A. Samwick. 1999. "The Other Side of the Trade-off: The Impact of Risk on Executive Compensation." *Journal of Political Economy*, 107(1): 65–105.

Alchian, Armen A., and Harold Demsetz. 1972. "Production , Information Costs, and Economic Organization." *American Economic Review*, 62(5): 777–95.

Al-Najjar, Nabil I. 1997. "Incentive Contracts in Two-Sided Moral Hazards with Multiple Agents." *Journal of Economic Theory*, 74(1): 174 – 195.

Au, Pak Hung, and Bin R. Chen. 2021. "Matching with peer monitoring." *Journal of Economic Theory*, 192: 105172.

Auriol, Emmanuelle, and Régis Renault. 2008. "Status and Incentives." *RAND Journal of Economics*, 39(1): 305–326.

Babcock, Philip, Kelly Bedard, Gary Charness, John Hartman, and Heather Royer. 2015. "Letting Down the Team? Evidence of Social Effects of Team Incentives." *Journal of the European Economic Association*, 13(5): 841–870.

Bag, Parimal Kanti, and Nona Pepito. 2012. "Peer Transparency in Teams: Does it Help or Hinder Incentives?" *International Economic Review*, 53(4): 1257–1286.

Baker, George P., Michael C. Jensen, and Kevin J. Murphy. 1988. "Compensation and Incentives: Practice vs. Theory." *Journal of Finance*, 43(3): 593–616.

Baker, George, Robert Gibbons, and Kevin J. Murphy. 1994. "Subjective Performance Measures in Optimal Incentive Contracts." *Quarterly Journal of Economics*, 109(4): 1125–1156.

Baker, George, Robert Gibbons, and Kevin J Murphy. 2002. "Relational Contracts and the Theory of the Firm." *Quarterly Journal of Economics*, 117(1): 39–84.

Baliga, Sandeep, and Tomas Sjöström. 1998. "Decentralization and Collusion." *Journal of Economic Theory*, 83(2): 196 – 232.

Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2005. "Social Preferences and the Response to Incentives: Evidence from Personnel Data." *Quarterly Journal of Economics*, 120(3): 917–962.

Baron, David P., and David Besanko. 1999. "Informational Alliances." *Review of Economic Studies*, 66(4): 743–768.

Bartling, Björn. 2011. "Relative Performance or Team Evaluation? Optimal Contracts for Other-Regarding Agents." *Journal of Economic Behavior & Organization*, 79(3): 183–193.

Bartling, Björn, and Ferdinand A. von Siemens. 2011. "Wage Inequality and Team Production: An Experimental Analysis." *Journal of Economic Psychology*, 32(1): 1–16.

Battaglini, Marco. 2006. "Joint Production in Teams." *Journal of Economic Theory*, 130(1): 138–167.

Belzil, Christian, and Michael Bognanno. 2008. "Promotions, Demotions, Halo Effects, and the Earnings Dynamics of American Executives." *Journal of Labor Economics*, 26(2): 287–310.

Bénabou, Roland, and Jean Tirole. 2003. "Intrinsic and Extrinsic Motivation." *Review of Economic Studies*, 70(3): pp. 489–520.

Bernheim, B Douglas, and Michael D Whinston. 1990. "Multimarket Contact and Collusive Behavior." *RAND Journal of Economics*, 1–26.

Bolton, Patrick, and Mathias Dewatripont. 1994. "The Firm as a Communication Network." *Quarterly Journal of Economics*, 109(4): 809–39.

Brown, Jennifer. 2011. "Quitters Never Win: The (Adverse) Incentive Effects of Competing with Superstars." *Journal of Political Economy*, 119(5): 982 – 1013.

Brusco, Sandro. 1997. "Implementing Action Profiles when Agents Collude." *Journal of Economic Theory*, 73(2): 395–424.

Brusco, Sandro. 1998. "Implementing Action Profiles with Sequential Mechanisms." *Review of Economic Design*, 3(3): 271–300.

Brusco, Sandro. 2002. "Unique Implementation of Action Profiles: Necessary and Sufficient Conditions." *International Economic Review*, 43(2): 509–532.

Camerer, Colin, and Dan Lovallo. 1999. "Overconfidence and Excess Entry: An Experimental Approach." *American Economic Review*, 89(1): pp. 306–318.

Carmichael, H Lorne. 1983. "The Agent-Agents Problem: Payment by Relative Output." *Journal of Labor Economics*, 1(1): 50–65.

Cato, Susumu, and Akifumi Ishihara. 2017. "Transparency and Performance Evaluation in Sequential Agency." *Journal of Law, Economics and Organization*, 33(3): 475–506.

Chan, Tat Y., Jia Li, and Lamar Pierce. 2014. "Compensation and Peer Effects in Competing Sales Teams." *Management Science*, 60(8): 1965–1984.

Charness, Gary, David Masclet, and Marie Claire Villeval. 2014. "The Dark Side of Competition for Status." *Management Science*, 60(1): 38–55.

Chen, Kong-Pin. 2003. "Sabotage in Promotion Tournaments." *Journal of Law, Economics and Organization*, 19(1): 119–140.

Che, Yeon-Koo, and Seung-Weon Yoo. 2001. "Optimal Incentives for Teams." *American Economic Review*, 91(3): 525–541.

Chillemi, Ottorino. 2008. "Competitive versus Collective Incentive Pay: Does Workers' Mutual Concern Matter?" *Economic Theory*, 35(1): 175–186.

Choi, Yoon K. 1993. "Managerial Incentive Contracts with a Production Externality." *Economics Letters*, 42(1): 37–42.

Coates, Dennis E. 1998. "Don't Tie 360 Feedback to Pay." *Training*, 25(9): 68–78.

Cohen, Susan G., and Diane E. Bailey. 1997. "What Makes Teams Work: Group Effectiveness Research from the Top Floor to the Executive Suite." *Journal of Management*, 23(3): 239–290.

Connelly, Brian L., Laszlo Tihanyi, T. Russell Crook, and K. Ashley Gangloff. 2014. "Tournament Theory: Thirty Years of Contests and Competitions." *Journal of Management*, 40(1): 16–47.

Corchon, Luis. 2007. "The Theory of Contests: A Survey." *Review of Economic Design*, 11(2): 69–100.

Crama, Pascale, Fabian J. Sting, and Yaozhong Wu. 2019. "Encouraging Help Across Projects." *Management Science*, 65(3): 1408–1429.

Crémer, Jacques, and Richard P. McLean. 1985. "Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist when Demands are Interdependent." *Econometrica*, 53(2): 345–361.

Crémer, Jacques, and Richard P. McLean. 1988. "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions." *Econometrica*, 56(6): 1247–1257.

d'Aspremont, Claude, and Louis-Andre Gerard-Varet. 1998. "Linear Inequality Methods to Enforce Partnerships under Uncertainty: An Overview." *Games and Economic Behavior*, 25(2): 311–336.

Deb, Joyee, Jin Li, and Arijit Mukherjee. 2016. "Relational Contracts with Subjective Peer Evaluations." *RAND Journal of Economics*, 47(1): 3–28.

de la Rosa, Leonidas Enrique. 2011. "Overconfidence and Moral Hazard." *Games and Economic Behavior*, 73(2): 429 – 451.

DeMatteo, Jacquelyn S., Lillian T. Eby, and Eric Sundstrom. 1998. "Team based Rewards: Current Empirical Evidence and Directions for Future Research." *Research in Organizational Behavior*, 20: 141–183.

Demski, Joel S., and David E. M. Sappington. 1987. "Delegated Expertise." *Journal of Accounting Research*, 25(1): 68–89.

Demski, Joel S., and David Sappington. 1984. "Optimal Incentive Contracts with Multiple Agents." *Journal of Economic Theory*, 33(1): 152–171.

Dequiedt, Vianney, and David Martimort. 2015. "Vertical Contracting with Informational Opportunism." *American Economic Review*, 105(7): 2141–82.

Deutsch, Morton. 1949*a*. "An Experimental Study of the Effects of Cooperation and Competition upon Group Process." *Human Relations*, 2(3): 199–231.

Deutsch, Morton. 1949*b*. "A Theory of Cooperation and Competition." *Human Relations*, 2(2): 129–152.

DeVaro, Jed. 2006. "Strategic Promotion Tournaments and Worker Performance." *Strategic Management Journal*, 27(8): 721–740.

Dewatripont, Mathias. 1993. "Implementation Theory: Discussion." *Advances in Economic Theory: Sixth World Congress*, , ed. Jean-Jacques Laffont Vol. 1 of *Econometric Society Monographs*, 324–327. Cambridge University Press.

Dhillon, Amrita, and Alexander Herzog-Stein. 2009. "Games of Status and Discriminatory Contracts." *Games and Economic Behavior*, 65(1): 105–123.

Drago, Robert, and Geoffrey K. Turnbull. 1987. "An Analysis of Non-insurance Work Incentives." *Economics Letters*, 23(3): 301–304.

Drago, Robert, and Geoffrey K. Turnbull. 1988. "The Incentive Effects of Tournaments with Positive Externalities among Workers." *Southern Economic Journal*, 55(1): 100–106.

Drago, Robert, and Gerald T Garvey. 1998. "Incentives for Helping on the Job: Theory and Evidence." *Journal of Labor Economics*, 16(1): 1–25.

Dur, Robert, and Jan Tichem. 2015. "Altruism and Relational Incentives in the Workplace." *Journal of Economics & Management Strategy*, 24(3): 485–500.

Dur, Robert, and Joeri Sol. 2010. "Social Interaction, Co-Worker Altruism, and Incentives." *Games and Economic Behavior*, 69(2): 293–301.

Dye, Ronald A. 1984. "The Trouble with Tournaments." *Economic Inquiry*, 22(1): 147–149.

Edwards, Mark R., and Ann J. Ewen. 1996. "How to Manage Performance and Pay with 360-Degree Feedback." *Compensation and Benefits Review*, 28(3): 41–46.

Ehrenberg, Ronald G., and Michael L. Bognanno. 1990. "Do Tournaments Have Incentive Effects?" *Journal of Political Economy*, 98(6): 1307–1324.

Englmaier, Florian, and Achim Wambach. 2010. "Optimal Incentive Contracts under Inequity Aversion." *Games and Economic Behavior*, 69(2): 312–328.

Falk, Armin, and Andrea Ichino. 2006. "Clean Evidence on Peer Effects." *Journal of Labor Economics*, 24(1): 39–57.

Fama, Eugene F. 1980. "Agency Problems and the Theory of the Firm." *Journal of political economy*, 88(2): 288–307.

Faure-Grimaud, Antoine, Jean-Jacques Laffont, and David Martimort. 2003. "Collusion, Delegation and Supervision with Soft Information." *Review of Economic Studies*, 70(2): 253–279.

Fehr, Ernst, and Simon Gächter. 2000. "Fairness and Retaliation: The Economics of Reciprocity." *Journal of Economic Perspectives*, 14(3): 159–181.

Fehr, Ernst, and Urs Fischbacher. 2002. "Why Social Preferences Matter – The Impact of Non-selfish Motives on Competition, Cooperation and Incentives." *Economic Journal*, 112(478): C1–C33.

Fehr, Ernst, Simon Gächter, and Georg Kirchsteiger. 1997. "Reciprocity as a Contract Enforcement Device: Experimental Evidence." *Econometrica*, 65(4): 833–860.

Fleckinger, Pierre. 2012. "Correlation and Relative Performance Evaluation." *Journal of Economic Theory*, 147(1): 93–117.

François, Patrick, and Joanne Roberts. 2003. "Contracting Productivity Growth." *Review of Economic Studies*, 70(1): 59–85.

Frank, Robert H. 1984. "Interdependent Preferences and the Competitive Wage Structure." *RAND Journal of Economics*, 15(4): 510–520.

Fuchs, William. 2007. "Contracting with Repeated Moral Hazard and Private Evaluations." *American Economic Review*, 97(4): pp. 1432–1448.

Fudenberg, Drew, David Levine, and Eric Maskin. 1994. "The Folk Theorem with Imperfect Public Information." *Econometrica*, 62(5): 997–1039.

Garicano, Luis. 2000. "Hierarchies and the Organization of Knowledge in Production." *Journal of Political Economy*, 108(5): 874–904.

Gershkov, Alex, and Eyal Winter. 2015. "Formal versus Informal Monitoring in Teams." *American Economic Journal: Microeconomics*, 7(2): 27–44.

Gershkov, Alex, Jianpei Li, and Paul Schweinzer. 2009. "Efficient Tournaments Within Teams." *RAND Journal of Economics*, 40(1): 103–119.

Gervais, Simon, and Itay Goldstein. 2007. "The Positive Effects of Biased Self-Perceptions in Firms." *Review of Finance*, 11(3): 453–496.

Gibbons, Robert, and Michael Waldman. 1999. "A Theory of Wage and Promotion Dynamics inside Firms." *Quarterly Journal of Economics*, 114(4): 1321–1358.

Gjesdal, Frøystein. 1982. "Information and Incentives: The Agency Information Problem." *Review of Economic Studies*, 49(3): 373–390.

Glover, Jonathan, and Eunhee Kim. 2021. "Optimal Team Composition: Diversity to Foster Implicit Team Incentives." *Management Science*, 67(9): 5800–5820.

Gould, Eric D., B. Peter Pashigian, and Canice J. Prendergast. 2005. "Contracts, Externalities, and Incentives in Shopping Malls." *Review of Economics and Statistics*, 87(3): 411–422.

Green, Jerry R., and Jean-Jacques Laffont. 1979. *Incentives in Public Decision-making.* Amsterdam:North-Holland Publishing Company.

Green, Jerry R., and Nancy L. Stokey. 1983. "A Comparison of Tournaments and Contracts." *Journal of Political Economy*, 91(3): 349–364.

Gromb, Denis, and David Martimort. 2007. "Collusion and the Organization of Delegated Expertise." *Journal of Economic Theory*, 137(1): 271–299.

Grossman, Sanford J., and Oliver D. Hart. 1983. "An Analysis of the Principal-Agent Problem." *Econometrica*, 51(1): 7–45.

Groves, Theodore. 1973. "Incentives in Teams." *Econometrica*, 41(4): 617–31.

Groves, Theodore, and John Ledyard. 1977. "Optimal Allocation of Public Goods: A Solution to the "Free Rider" Problem." *Econometrica*, 45(4): 783–809.

Grund, Christian, and Dirk Sliwka. 2005. "Envy and Compassion in Tournaments." *Journal of Economics & Management Strategy*, 14(1): 187–207.

Gupta, Srabana, and Richard E. Romano. 1998. "Monitoring the Principal with Multiple Agents." *RAND Journal of Economics*, 29(2): 427–442.

Halac, Marina, Elliot Lipnowski, and Daniel Rappoport. 2021. "Rank Uncertainty in Organizations." *American Economic Review*, 111(3): 757–86.

Halac, Marina, Ilan Kremer, and Eyal Winter. 2021. "Monitoring Teams." *mimeo*.

Hamilton, Barton H., Jack A. Nickerson, and Hideo Owan. 2003. "Team Incentives and Worker Heterogeneity: An Empirical Analysis of the Impact of Teams on Productivity and Participation." *Journal of Political Economy*, 111(3): 465–497.

Hansen, Daniel G. 1997. "Work Performance and Group Incentives: A Case Study." *Industrial and Labor Relations Review*, 51(1): 37–49.

Harbring, Christine, and Bernd Irlenbusch. 2011. "Sabotage in Tournaments: Evidence from a Laboratory Experiment." *Management Science*, 57(4): 611–627.

Herries, Joshua, Daniel I. Rees, and Jeffrey S. Zax. 2003. "Interdependence in Worker Productivity." *Journal of Applied Econometrics*, 18(5): 585–604.

Holmström, Bengt. 1979. "Moral Hazard and Observability." *Bell Journal of Economics*, 10(1): 74–91.

Holmström, Bengt. 1982. "Moral Hazard in Teams." *Bell Journal of Economics*, 13(2): 324–340.

Holmström, Bengt. 1999. "Managerial Incentive Problems: A Dynamic Perspective." *Review of Economic studies*, 66(1): 169–182.

Holmstrom, Bengt, and Jean Tirole. 1997. "Financial Intermediation, Loanable Funds, and the Real Sector." *Quarterly Journal of Economics*, 112(3): 663–691.

Holmström, Bengt, and Paul Milgrom. 1990. "Regulating Trade among Agents." *Journal of Institutional and Theoretical Economics*, 146(1): 85–105.

Holmström, Bengt, and Paul Milgrom. 1994. "The Firm as an Incentive System." *American Economic Review*, 84(4): 972–91.

Ichino, Andrea, and Giovanni Maggi. 2000. "Work Environment and Individual Background: Explaining Regional Shirking Differentials in a Large Italian Firm." *Quarterly Journal of Economics*, 115(3): 1057–1090.

Innes, Robert D. 1990. "Limited liability and Incentive Contracting with Ex-ante Action Choices." *Journal of Economic Theory*, 52(1): 45 – 67.

Iossa, Elisabetta, and David Martimort. 2015. "The Simple Microeconomics of Public-Private Partnerships." *Journal of Public Economic Theory*, 17(1): 4–48.

Ishiguro, Shingo. 2002. "Optimal Incentives for Teams: Comment." *American Economic Review*, 92(5): 1711–1711.

Ishiguro, Shingo, and Hideshi Itoh. 2001. "Moral Hazard and Renegotiation with Multiple Agents." *Review of Economic Studies*, 68(1): 1–20.

Itoh, Hideshi. 1991. "Incentives to Help in Multi-Agent Situations." *Econometrica*, 59(3): 611–36.

Itoh, Hideshi. 1992. "Cooperation in Hierarchical Organizations: An Incentive Perspective." *Journal of Law, Economics and Organization*, 8(2): 321–45.

Itoh, Hideshi. 1993. "Coalitions, Incentives, and Risk Sharing." *Journal of Economic Theory*, 60(2): 410–427.

Itoh, Hideshi. 1994. "Job Design, Delegation and Cooperation: A Principal-Agent Analysis." *European Economic Review*, 38(3-4): 691–700.

Itoh, Hideshi. 2004. "Moral Hazard and Other-regarding Preferences." *Japanese Economic Review*, 55(1): 18–45.

Jehiel, Philippe. 2014. "On Transparency in Organizations." *Review of Economic Studies*, 82(2): 736–761.

Kandel, Eugene, and Edward P. Lazear. 1992. "Peer Pressure and Partnerships." *Journal of Political Economy*, 100(4): 801–17.

Kandori, Michihiro, and Ichiro Obara. 2006. "Efficiency in Repeated Games Revisited: The Role of Private Strategies." *Econometrica*, 74(2): 499–519.

Kellner, Christian. 2015. "Tournaments as a Response to Ambiguity Aversion in Incentive Contracts." *Journal of Economic Theory*, 159: 627–655.

Klein, Benjamin, and Keith B. Leffler. 1981. "The Role of Market Forces in Assuring Contractual Performance." *Journal of Political Economy*, 89(4): 615–641.

Knez, Marc, and Duncan Simester. 2001. "Firm-Wide Incentives and Mutual Monitoring at Continental Airlines." *Journal of Labor Economics*, 19(4): 743–72.

Konrad, Kai A. 2009. *Strategy and Dynamics in Contests. OUP Catalogue*, Oxford University Press.

Kosenok, Grigory, and Sergei Severinov. 2008. "Individually Rational, Budget-balanced Mechanisms and Allocation of Surplus." *Journal of Economic Theory*, 140(1): 126–161.

Kvaløy, Ola, and Trond E. Olsen. 2006. "Team Incentives in Relational Employment Contracts." *Journal of Labor Economics*, 24(1): 139–170.

Kvaløy, Ola, and Trond E. Olsen. 2012. "The Rise of Individual Performance Pay." *Journal of Economics & Management Strategy*, 21(2): 493–518.

Kvaløy, Ola, and Trond E. Olsen. 2019. "Relational Contracts, Multiple Agents, and Correlated Outputs." *Management Science*, 65(11): 5360–5370.

Laffont, Jean-Jacques, and David Martimort. 1997. "Collusion under Asymmetric Information." *Econometrica*, 65(4): 875–911.

Laffont, Jean-Jacques, and David Martimort. 2000. "Mechanism Design with Collusion and Correlation." *Econometrica*, 68(2): 309–342.

Laffont, Jean-Jacques, and David Martimort. 2002. *The Theory of Incentives: The Principal-Agent Model.* Princeton University Press.

Laffont, Jean-Jacques, and Jean-Charles Rochet. 1997. "Collusion in Organizations." *Scandinavian Journal of Economics*, 99(4): 485–495.

Laffont, Jean-Jacques, and Patrick Rey. 2003. "Moral Hazard, Collusion and Group Lending." Institut d'économie Industrielle (IDEI), Toulouse IDEI Working Papers 122.

Lambert, Richard A. 1986. "Executive Effort and Selection of Risky Projects." *RAND Journal of Economics*, 17(1): 77–88.

Laux, Christian. 2001. "Limited-Liability and Incentive Contracting with Multiple Projects." *RAND Journal of Economics*, 32(3): 514–26.

Lavy, Victor. 2009. "Performance Pay and Teachers' Effort, Productivity, and Grading Ethics." *American Economic Review*, 99(5): 1979–2011.

Lazear, Edward P. 1989. "Pay Equality and Industrial Politics." *Journal of Political Economy*, 97(3): 561–80.

Lazear, Edward P. 1995. *Personnel Economics.* Vol. 1 of *MIT Press Books*, MIT Press.

Lazear, Edward P. 2018. "Compensation and Incentives in the Workplace." *Journal of Economic Perspectives*, 32(3): 195–214.

Lazear, Edward P., and Sherwin Rosen. 1981. "Rank-Order Tournaments as Optimum Labor Contracts." *Journal of Political Economy*, 89(5): 841–64.

Legros, Patrick, and Hitoshi Matsushima. 1991. "Efficiency in Partnerships." *Journal of Economic Theory*, 55(2): 296–322.

Legros, Patrick, and Steven A. Matthews. 1993. "Efficient and Nearly-efficient Partnerships." *Review of Economic Studies*, 60(3): 599–611.

Levin, Jonathan. 2002. "Multilateral Contracting and the Employment Relationship." *Quarterly Journal of Economics*, 117(3): 1075–1103.

Levin, Jonathan. 2003. "Relational Incentive Contracts." *American Economic Review*, 93(3): 835–857.

Lin, Y. Joseph. 1997. "Division of Labor in Teams." *Journal of Economics & Management Strategy*, 6(1): 403–423.

Luporini, Annalisa. 2006. "Relative Performance Evaluation in a Multiplan Firm." *Economic Theory*, 28(1): 235–243.

Macaulay, Stewart. 1963. "Non-Contractual Relations in Business: A Preliminary Study." *American Sociological Review*, 28(1): 55–67.

Ma, Ching-To. 1988. "Unique Implementation of Incentive Contracts with Many Agents." *Review of Economic Studies*, 55(4): 555–72.

Ma, Ching-to, John Moore, and Stephen Turnbull. 1988. "Stopping Agents from "Cheating"." *Journal of Economic Theory*, 46(2): 355–372.

Macho-Stadler, Ines, and J. David Perez-Castrillo. 1993. "Moral Hazard with Several Agents : The Gains from Cooperation." *International Journal of Industrial Organization*, 11(1): 73–100.

Macho-Stadler, Ines, and J. David Perez-Castrillo. 1998. "Centralized and Decentralized Contracts in a Moral Hazard Environment." *Journal of Industrial Economics*, 46(4): 489–510.

MacLeod, W. Bentley. 1987. "Behavior and the Organization of the Firm." *Journal of Comparative Economics*, 11(2): 207–220.

MacLeod, W. Bentley. 1988. "Equity, Efficiency, and Incentives in Cooperative Teams." *Advances in the Economic Analysis of Participatory and Labor Managed Firms*, , ed. Derek C. Jones and Jan Svejnar Vol. 3, 5–23. JAI Press.

MacLeod, W. Bentley. 2003. "Optimal Contracting with Subjective Evaluation." *American Economic Review*, 93(1): 216–240.

MacLeod, W. Bentley. 2007. "Reputations, Relationships, and Contract Enforcement." *Journal of Economic Literature*, 45(3): 595–628.

Magill, Michael, and Martine Quinzii. 2006. "Common Shocks and Relative Compensation." *Annals of Finance*, 2(4): 407–420.

Malcomson, James M. 1984. "Work Incentives, Hierarchy, and Internal Labor Markets." *Journal of Political Economy*, 92(3): 486–507.

Malcomson, James M. 1986. "Rank-Order Contracts for a Principal with Many Agents." *Review of Economic Studies*, 53(5): 807–17.

Malcomson, James M. 2009. "Principal and Expert Agent." *B.E. Journal of Theoretical Economics*, 9(1): 0000102202193517041528.

Malcomson, James M. 2010. "Relational Incentive Contracts." University of Oxford, Department of Economics Economics Series Working Papers 508.

Marschak, Jacob, and Roy Radner. 1972. *Economic Theory of Teams.* Yale University Press.

Martimort, David. 1999. "The Life Cycle of Regulatory Agencies: Dynamic Capture and Transaction Costs." *Review of Economic Studies*, 66(4): 929–947.

Martimort, David, and Thierry Verdier. 2004. "The Agency Cost of Internal Collusion and Schumpeterian Growth." *Review of Economic Studies*, 71(4): 1119–1141.

Maskin, Eric. 1999. "Nash Equilibrium and Welfare Optimality." *Review of Economic Studies*, 66(1): 23–38.

Maskin, Eric, Yingyi Qian, and Chenggang Xu. 2000. "Incentives, Information, and Organizational Form." *Review of Economic Studies*, 67(2): 359–78.

McAfee, R. Preston, and John McMillan. 1991. "Optimal Contracts for Teams." *International Economic Review*, 32(3): 561–577.

McAfee, R. Preston, and Philip J. Reny. 1992. "Correlated Information and Mecanism Design." *Econometrica*, 60(2): 395–421.

Meyer, Margaret A., and Dilip Mookherjee. 1987. "Incentives, Compensation, and Social Welfare." *Review of Economic Studies*, 54(2): 209–26.

Milgrom, Paul R. 1988. "Employment Contracts, Influence Activities, and Efficient Organization Design." *Journal of Political Economy*, 96(1): 42–60.

Milgrom, Paul R., and John Roberts. 1992. *Economics, Organization and Management.* Prentice Hall.

Miller, L. Keith, and Robert L. Hamblin. 1963. "Interdependence, Differential Rewarding, and Productivity." *American Sociological Review*, 28(5): 768–778.

Miller, Nolan H. 1997. "Efficiency in Partnerships with Joint Monitoring." *Journal of Economic Theory*, 77(2): 285–299.

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

Mohnen, Alwine, Kathrin Pokorny, and Dirk Sliwka. 2008. "Transparency, Inequity Aversion, and the Dynamics of Peer Pressure in Teams: Theory and Evidence." *Journal of Labor Economics*, 26(4): 693–720.

Moldovanu, Benny, and Aner Sela. 2001. "The Optimal Allocation of Prizes in Contests." *American Economic Review*, 91(3): 542–558.

Moldovanu, Benny, Aner Sela, and Xianwen Shi. 2012. "Carrots and Sticks: Prizes and Punishments in Contests." *Economic Inquiry*, 50(2): 453–462.

Mookherjee, Dilip. 1984. "Optimal Incentive Schemes with Many Agents." *Review of Economic Studies*, 51(3): 433–46.

Mookherjee, Dilip, and Stefan Reichelstein. 1992. "Dominant Strategy Implementation of Bayesian Incentive Compatible Allocation Rules." *Journal of Economic Theory*, 56(2): 378–399.

Moore, John. 1993. "Implementation, Contracts, and Renegotiation in Environments with Complete Information." *Advances in Economic Theory: Sixth World Congress*, , ed. Jean-Jacques Laffont Vol. 1 of *Econometric Society Monographs*, 182–282. Cambridge University Press.

Mukherjee, Arijit, and Luis Vasconcelos. 2011. "Optimal Job Design in the Presence of Implicit Contracts." *RAND Journal of Economics*, 42(1): 44–69.

Münster, Johannes. 2007. "Selection Tournaments, Sabotage, and Participation." *Journal of Economics & Management Strategy*, 16(4): 943–970.

Nalebuff, Barry J., and Joseph E. Stiglitz. 1983. "Prices and Incentives: Towards a General Theory of Compensation and Competition." *Bell Journal of Economics*, 14(1): 21–43.

Nocke, Volker, and Roland Strausz. 2023. "Collective Brand Reputation." *Journal of Political Economy*, 131(1): 1–58.

Olszewski, Wojciech, and Ron Siegel. 2016. "Large Contests." *Econometrica*, 84(2): 835–854.

Park, Young, and Luís Santos-Pinto. 2010. "Overconfidence in Tournaments: Evidence from the Field." *Theory and Decision*, 69(1): 143–166.

Picard, Pierre, and Patrick Rey. 1987. "Incentives in Cooperative Research and Development." CEPREMAP CEPREMAP Working Papers (Couverture Orange) 8739.

Poblete, Joaquín, and Daniel Spulber. 2012. "The Form of Incentive Contracts: Agency with Moral Hazard, Risk Neutrality, and Limited Liability." *RAND Journal of Economics*, 43(2): 215–234.

Prendergast, Canice. 1993. "The Role of Promotion in Inducing Specific Human Capital Acquisition." *Quarterly Journal of Economics*, 108(2): 523–34.

Prendergast, Canice. 1999. "The Provision of Incentives in Firms." *Journal of Economic Literature*, 37(1): 7–63.

Rahman, David. 2012. "But Who Will Monitor the Monitor?" *American Economic Review*, 102(6): 2767–2797.

Rahman, David, and Ichiro Obara. 2010. "Mediated Partnerships." *Econometrica*, 78(1): 285–308.

Rajan, Madhav V, and Stefan Reichelstein. 2006. "Subjective Performance Indicators and Discretionary Bonus Pools." *Journal of Accounting Research*, 44(3): 585–618.

Ramakrishnan, Ram T.S., and Anjan V Thakor. 1991. "Cooperation versus Competition in Agency." *Journal of Law, Economics and Organization*, 7(2): 248–83.

Rayo, Luis. 2007. "Relational Incentives and Moral Hazard in Teams." *Review of Economic Studies*, 74(3): 937–963.

Rey-Biel, Pedro. 2008. "Inequity Aversion and Team Incentives." *Scandinavian Journal of Economics*, 110(2): 297–320.

Riordan, Michael H, and David E.M Sappington. 1988. "Optimal Contracts with Public Ex Post Information." *Journal of Economic Theory*, 45(1): 189–199.

Rosen, Sherwin. 1986. "Prizes and Incentives in Elimination Tournaments." *American Economic Review*, 76(4): 701–15.

Ross, Stephen A. 1973. "The Economic Theory of Agency: The Principal's Problem." *American Economic Review*, 63(2): 134–139.

Sah, Raaj Kumar, and Joseph E. Stiglitz. 1986. "The Architecture of Economic Systems: Hierarchies and Polyarchies." *American Economic Review*, 76(4): 716–727.

Salas-Fumas, Vicente. 1992. "Relative Performance Evaluation of Management." *International Journal of Industrial Organization*, 10(3): 473–489.

Santos-Pinto, Luís. 2008. "Positive Self-image and Incentives in Organisations." *Economic Journal*, 118(531): 1315–1332.

Santos-Pinto, Luís. 2010. "Positive Self-Image In Tournaments." *International Economic Review*, 51(2): 475–496.

Sappington, David. 1983. "Limited Liability Contracts between Principal and Agent." *Journal of Economic Theory*, 29(1): 1–21.

Sappington, David, and Joel S. Demski. 1983. "Multi-Agent Control in Perfectly Correlated Environments." *Economics Letters*, 13(4): 325–330.

Schmitz, Patrick W. 2005. "Allocating Control in Agency Problems with Limited Liability and Sequential Hidden Actions." *RAND Journal of Economics*, 36(2): 318–336.

Shavell, Steven. 1979. "Risk Sharing and Incentives in the Principal and Agent Relationship." *Bell Journal of Economics*, 10(1): 55–73.

Shleifer, Andrei. 1985. "A Theory of Yardstick Competition." *RAND Journal of Economics*, 16(3): 319–327.

Siegel, Ron. 2009. "All-Pay Contests." *Econometrica*, 77(1): 71–92.

Strausz, Roland. 1999. "Efficiency in Sequential Partnerships." *Journal of Economic Theory*, 85(1): 140–156.

Tirole, Jean. 1986. "Hierarchies and Bureaucracies: On the Role of Collusion in Organizations." *Journal of Law, Economics and Organization*, 2(2): 181–214.

Tirole, Jean. 1992. "Collusion and the Theory of Organizations." In *Advances in Economic Theory*. Vol. 2, , ed. Jean-Jacques Laffont, 151–206. Cambridge:Cambridge University Press.

Towry, Kristy L. 2003. "Control in a Teamwork Environment: The Impact of Social Ties on the Effectiveness of Mutual Monitoring." *Accounting Review*, 78(4): 1069–1095.

Tsoulouhas, Theofanis. 1999. "Do Tournaments Solve the Two-Sided Moral Hazard Problem?" *Journal of Economic Behavior & Organization*, 40(3): 275–294.

Upton, Harvey. 2023. "Implicit Incentives and Delegation in Teams." *Management Science*, forthcoming.

Varian, H.R. 1990. "Monitoring Agents with Other Agents." *Journal of Institutional and Theoretical Economics*, 146(1): 153–174.

Vialle, Isabelle, Luis Santos-Pinto, and Jean-Louis Rullière. 2011. "Self-Confidence and Teamwork : An Experimental Test." Groupe d'Analyse et de Théorie Economique (GATE), Centre national de la recherche scientifique (CNRS), Université Lyon 2, Ecole Normale Supérieure Working Papers 1126.

Vislie, Jon. 1994. "Efficiency and Equilibria in Complementary Teams." *Journal of Economic Behavior & Organization*, 23(1): 83–91.

Wageman, Ruth. 2001. "The Meaning of Interdependence." *Groups at Work: Theory and Research*, , ed. Marlene Turne, 197–217. Psychology Press.

Wageman, Ruth, and George Baker. 1997. "Incentives and Cooperation: The Joint Effects of Task and Reward Interdependence on Group Performance." *Journal of Organizational Behavior*, 18(2): 139–158.

Watson, Joel. 2021. "Theoretical Foundations of Relational Incentive Contracts." *Annual Review of Economics*, 13: 631–659.

Wilson, Robert B. 1968. "The Theory of Syndicates." *Econometrica*, 36(1): 119–132.

Winter, Eyal. 2004. "Incentives and Discrimination." *American Economic Review*, 94(3): 764–773.

Winter, Eyal. 2006. "Optimal Incentives for Sequential Production Processes." *RAND Journal of Economics*, 37(2): 376–390.

Winter, Eyal. 2010. "Transparency and Incentives among Peers." *RAND Journal of Economics*, 41(3): 504–523.