Information Sharing and Moral Hazard in Teams
(Work in Progress)

Luis Garicano
LSE and CEPR

Adam Meirowitz
U. of Utah – Eccles School

Luis Rayo
U. of Utah – Eccles School, and CEPR

November 2, 2017

Abstract

We study team production with privately-informed agents and relational contracting between them. The team faces a joint moral-hazard and information-sharing problem. When private information is about the output technology – not the agents' effort costs – the first-best is achieved by means of a canonical "surplus-sharing" arrangement: a subset of players receive output shares and pay performance bonuses to their peers in exact proportion to these shares. This arrangement, of which a single owner/manager is a special case, is both necessary and sufficient for agents to internalize all externalities. When instead private information is about the agents' effort costs, the first-best can only be achieved with a budget breaker.

To derive applications, we impose minimal additional structure. When players' efforts vary independently from each other across different states (e.g. it is possible for a player's effort not to be needed while the other team members exert a lot of effort), then having a single owner/manager is efficient. When instead the optimal efforts of at least some subset of the team are in sync with the overall effort of the team (e.g. several players in team are essential), then a partnership, with profits divided among this subset, is efficient.

1 Introduction

To be successful, organizations must achieve two goals: motivating agents and aggregating their private information. A key obstacle is that dispersed information is crucial to determining what must be done, who should do it, and also perhaps how it should be done. While the literature has advanced in solving the motivation and information aggregation problems separately, little progress has been made in determining how organizations can deal with the two problems together.

For instance, consider a team of specialized doctors working on the diagnosis and treatment of a patient. The doctors may be well motivated to cure the patient, but this may be useless if they are not performing the right tests or using the right treatments. The choice of the next test or the most reasonable treatment might require pooling information or judgements across the specialists. The hematologist might be able to reasonably rule out certain basic neurological conditions as sufficiently unlikely to justify bringing in a neurologist, but since the hematologist does not absorb any of the costs of neurological testing, she may lack the incentives to share this information. Alternatively, a surgeon may derive prestige from certain invasive procedures and therefore may refrain from mentioning alternative options. As these examples suggest, there can be important and complicated dependencies between how incentive schemes create incentives for actions and how they create incentives for information sharing.

In a very different industry, professional cycling, similar organizational challenges emerge. In races like the Tour de France, 9-man teams may participate in different contests: the race for the yellow-jersey each day, the race for the overall title after nearly a month of racing, stage wins, the climbers jerseys, or the sprinters jersey. Each day the team must decide how to focus the efforts of each rider for that day. Moreover, the fitness and health of riders varies tremendously over three weeks of racing with each rider possessing significant private information about their performance potential for that stage. Famed sprinter Mark Cavendish (2013) writes about the perils of having the team devote substantial effort in support of a rider that is just not firing on all cylinders that day. In particular, he laments the decision of management to focus riders on supporting a leader with aspirations to win the overall title at the expense of supporting Cavendish's efforts to win sprints when, on that day, the overall contender was under the weather, and Cavendish was strong.

Economic theory has taught us that when only the moral-hazard problem is present, two simple solutions are available. First, an external "budget breaker" may use group rewards and punishments (Holmstrom, 1982). Secondly, the logic of the folk theorem tells us that repeated play achieves the first best by means of appropriate self-enforcing bonuses. This discussion leaves open two questions: can the moral-hazard and information-aggregation problems be solved simultaneously, and if so, what forms will the resulting organizations take?

In this paper, to answer both questions, we consider a model of repeated team production in which players possess private information. Efficiency requires aggregating this information and creating incentives for optimal state-contingent levels of effort in each period. We show that the nature of private information is key in two ways. First, it determines whether or not the first best can be obtained without a budget breaker. Second, when efficiency is possible, it determines the structure of optimal contracts (organizations). The need to solve both the moral-hazard and information-sharing problems places tight demands. As a result, we obtain tight characterizations of the efficient organizational structures. Moreover, the model has a rich set of primitives, namely the structure of information dispersion, and therefore provides insights into how optimal organizational design depends on the nature of the underlying informational environment.

In the case in which private information is about how efforts determine output the results are positive: the first best can be achieved. Efficiency requires the use of a canonical "surplus sharing" arrangement. In such an arrangement, a subset of players receive output shares and pay performance bonuses to their peers in exact proportion to these shares. This arrangement, of which a single owner/manager is a special case, is necessary for agents to internalize all externalities. In particular, the performance bonuses force agents to internalize the externalities of their reports on the effort of their peers. For example, if the hematologist's report causes additional effort by the neurologist, then the hematologist must internalizes the cost of this additional effort; if the contender for the yellow jersey stands to gain prestige and advertisement revenue from a good showing, he must also feel the costs he imposes on team mates that support his efforts.

Conversely, we show that provided there is sufficient relational capital (the value of the relationship is high) any surplus-sharing arrangement can be made to support first-best efforts. Although a large range of profit shares can support first-best efforts, different allocations of shares differ in how much relational capital is needed.

The efficient allocation of shares – i.e. the allocation that requires the least relational capital – depends on the fundamentals of the production technology. In settings in which players' efforts vary independently from each other across different states (e.g. it is possible for a player's effort not to be needed while the other team members exert a lot of effort), then having a single owner/manager is efficient. When instead the optimal efforts of at least some subset of the team are in sync with the overall effort of the team (e.g. several players in the team are essential), then a partnership, with profits divided among this subset, is efficient.

The intuition begins with the following observation. The self-enforcing bonus that each agent must pay involves compensating his peers for a fraction of their effort costs – in direct proportion to the agent's profit shares – and, simultaneously, receiving compensation for a fraction of his own cost – in proportion to his peers' profit shares. As a result, each agent's reneging temptation is highest in states in which he exerts low effort while at the same time his peers work hard.

Thus, when the efforts of a subset of players are in sync with the overall effort of the team, the total reneging temptation is minimized by sharing ownership – and therefore pooling bonus payments – between them: by doing so, since effort costs tend to cancel out, no single agent is ever called upon to make a large payment. When instead the efforts of all players vary independently from each other, bonus payments can no longer be pooled. In this case, it is best to assign all ownership to the player who exerts the highest level of effort across states: by doing so, the team at least avoids having to pay a large bonus to this player.

Throughout most of our analysis, we focus on fairly simple arrangements in which output shares (which are externally enforced) do not depend on the realized level of output, or on the agent's internal communications. If there are large gains from added complexity, one might expect more elaborate arrangements to emerge. To illustrate the potential size of these gains, we show that sufficiently complex court-enforced contracts, provided they can be enforced, achieve the first best without any need for self-enforcing bonuses. These schemes involve a complex selection of shares off the equilibrium path; but, importantly, on the equilibrium path, they look just like the simpler systems of fixed shares.

Finally, we consider settings in which private information is about the players' individual effort costs. We show that, in this case, relational contracts alone do not lead to first best. As a result, a Holmstrom-like result obtains: only through a budget breaker can the team achieve efficiency. In particular, aligning incentives requires that each team member becomes a *full* residual claimant of surplus. Since there is only one surplus to distribute among team members, a budget breaker is needed. In other words, when the budget is balanced, it is not possible to create the correct incentives for information sharing without simultaneously damaging effort incentives.

2 Related Literature

Although the problem of coordination and motivation has long been considered the crucial problem of organization (at least since Arrow, 1974 masterful treatment), the literature studying these issues has branched out into two separate fields. One area, the principal agent literature, has been concerned (following Holmstrom, 1979) with the problem of motivating agents to exert unobservable effort. When multiple tasks and agents are involved, this literature has emphasized the need for low-powered incentives in contexts in which coordination matters (Holmstrom and Milgrom, 1991, 1994, Holmstrom, 1999).

A separate literature has been concerned with coordination problems due to bounded rationality and limited information, absent incentive conflicts. For example, Cremer (1980) and Vayanos (2002) study the optimal grouping of subunits into units in the presence of interdependencies; Dessein and Santos (2006) study the trade-off between ex-ante coordination, through rules, and ex-post coordination, through communication; Cremer, Garicano, and Prat

(2007) study how organizational codes allow for improve coordination, and how they place limits on firm scope.

More recently, some authors have studied problems that involve both coordination and incentive problems. Segal (1999) studies agents who exert externalities on each other through their participation decisions. Bernstein and Winter (2012), in a complete information setting, show that a subset of the players must be subsided to ensure that they invest even if no one outside the set does. Given that these early adopters adopt, the rest of the players then proceed to invest, and can even be taxed.

Dessein et al. (2010) study a model where optimal contracts must ensure both effort incentives and coordination among agents. They show that optimal incentives leads to biased decision making, as agents must (optimally) be made to care about their own units profits. Edmans et al. (2012) study a model in which effort by one agent reduces the cost of effort by other agents. As a consequence it is optimal to "over incentivize" synergistic agents. Sakovics and Steiner (2011) study a coordination problem in global game setting where agents have asymmetric information. Their objective is to study towards whom should taxes and subsidies be directed. Subsidies to encourage adoption must be directed to the agent on whom others impose less externalities, who relies less on others adopting the new technology .

Like these more recent papers, in our paper agents must both be coordinated and motivated. Like some of those papers (e.g. Dessein et al., 2011) information about the right combination of efforts is dispersed, and in the hands of the agents. Unlike in any of them, relational contracts are allowed, as agents work together repeatedly. Moreover our model is broader than the ones in those papers. While we sharply expand the range of the contracts allowed, the problem is surprisingly tractable.[1]

This paper contributes to the growing literature on relational contracts

[1] A less closely related recent literature has studied the problem of pricing externalities on a network. In these papers, consumers are situated on a graph, and exert externalities on each other: for instance, one consumer's adoption decision increases the likelihood that another consumer will adopt. A monopolist must determine the prices charged to maximize profits, internalizing adoption externalities. For instance, in some papers, such as Candogan et al. (2012), Bloch and Querou (2013) and Fainmesser and Galeotti (2015), this pattern of interactions is specified as a deterministic graph. While these papers worry about interactions and incentives, there is no unobservable effort nor relational contracting.

– a few examples include Bull, 1987, MacLeod and Malcomson, 1989, 1998, MacLeod, 2003, Baker, Gibbons, and Murphy, 1994, Levin, 2003, Fuchs, 2007, Halac, 2012). This literature focuses on settings with an exogenous principal who is not directly involved in production. In our model, in contrast, all agents to contribute to output, share the team's profits, and pay each other's bonuses. As result, the model addresses questions of optimal organizational design.[2]

In this literature, Rayo (2007) also studies a team moral-hazard problem. He is interested in how the ability to monitor team members impacts the optimal allocation of profit shares. In contrast, we are interest resolving a joint information-sharing and incentive problem. This richer problem allows us to tightly characterize the team's optimal organizational structure.

Finally, Skrzypacz and Toikka (2015) study the problem of inducing an optimal level of trade, when parties are privately informed, in a dynamic environment. There are also interested in understanding when the first best can be achieved. The difference between the two approaches follows from the underlying difference in the two settings: in their case, a buyer-seller relationship in which players have exogenously assigned roles; in our case, a multi-sided effort-provision problem in which the role of each player depends on her endogenously chosen profit shares.

3 Model

Consider a team of N players who live forever. Players are risk neutral and discount future payoffs using interest rate r. In each period, output is given by $y = f(e, \theta)$, where $e = (e_1, ..., e_N)$ is a profile of effort choices and $\theta = (\theta_1, ..., \theta_N)$ is a profile of private types (only player i observes θ_i). We refer to θ as the state. Player i's cost of effort is $c(e_i, \theta_i)$. (Throughout, to avoid clutter, we omit time subscripts.) Output is observed by the courts; efforts are observed only inside the team.

We assume that $e_i \in [0, \overline{e}]$ and $\theta_i \in [0, 1]$. We also make standard assumptions on the output technology and cost functions: f and c are continuously-

[2] Levin (2003), MacLeod (2003), Fuchs (2007) or Halac (2012) consider prncipal-agent problems with asymmetric information. The first three papers have private information about the outcome variable, while Halac (2012) has persistent private information about the players' outside payoffs.

differentiable in θ and f is increasing in e, θ. In addition, θ raises the marginal productivity of e, namely, for all $e' > e$ the difference $f(e', \theta) - f(e, \theta)$ is (weakly) increasing in θ. The state θ is drawn, at the beginning of each period, from a c.d.f. $H(\cdot)$ that has full support (and is otherwise unrestricted).

Denote the team's overall surplus by

$$S(e, \theta) = f(e, \theta) - \sum_i c\left(e_i, \theta_i\right).$$

Denote the first-best effort choices, for any given θ, by

$$e^*(\theta) = \arg\max_e \; S(e, \theta).$$

(Which we assume exist and are unique.) Note that the increasing-differences assumption implies that $e^*(\theta)$ is (weakly) increasing in θ.

Throughout, we focus on first-best arrangements that induce players to select $e^*(\theta)$ for all states.

Participation in the team is voluntary. At the beginning of each period, players can exit the relationship an obtain a per-period outside option v_i. We assume that $\underline{S} = \inf_\theta S(e^*(\theta), \theta) \geq \sum_i v_i$, which grantees that, provided the team operates efficiently, it provides a surplus large enough to simultaneously cover all of the players' outside options.

The team's goal is to select the first-best effort profile $e^*\left(\theta\right)$ period by period while also guaranteeing that each player wishes to participate in the team.

We impose a minimal level of structure on the team's interactions. During each period, after players learn their private types, we have:

1. A reporting stage in which players simultaneously send public messages $m_i \in M_i$.

2. An effort-provision stage in which players simultaneously select efforts e_i.

3. A money transfer stage in which players execute formal contracts – described below – and make self-enforcing money transfers to each other.

Per the Revelation Principle, we restrict attention to direct-revelation mechanisms in which players are asked to report their types ($M_i = [0,1]$) and, we focus on equilibria in which they do so truthfully ($m_i = \theta_i$ for all θ_i).

In studying the reporting incentives it is convenient to define payoffs that anticipate equilibrium play at the effort and transfer stages. Let $V_i(m_i;\theta)$ denote player i's payoff in period t when she reports type m_i, the true state is θ, and her peers report truthfully and all players follow all prescribed actions subsequently. It is convenient to let $V_i(\theta) = V_i(\theta_i;\theta)$ denote player i's equilibrium payoff.

Contracts

The team has two types of contracts at its disposal. First, players enter into a court-enforced contract that specifies, for each player, a share of output α_i and a wage w_i. Second, players promise self-enforcing money transfers t_i to each other, as a function of the reports θ and the chosen efforts e (both of which are observed by all members of the team).

For the time being, we focus on court-enforced contracts with a simple (empirically-relevant) form: the shares α_i and wages w_i are constants (i.e. they are independent of the realization of output and the players reports). Shares are constrained to be non-negative (which prevents budget-breaking schemes) and must add up to one; wages must add up to zero, as team members pay these wages to each other. (In Section ??, we consider more complex arrangements.)

Without loss of generality, the self-enforcing money transfers t_i are only paid when players follow their prescribed efforts. Namely,

$$t_i(e,\theta) = \begin{cases} b_i(\theta) & \text{if } e = e^*(\theta), \\ 0 & \text{otherwise.} \end{cases}$$

for some function $b_i(\theta)$. We refer to $b_i(\theta)$ as the bonus paid by player i (with a negative bonus meaning that player i gets paid).

Player i's payoff is

$$\underbrace{\alpha_i y + w_i}_{\text{court enforced}} \; - \; \underbrace{b_i(\theta)}_{\text{self enforced}} \; - \; c(e_i,\theta).$$

A *relational contract*, or more compactly, a *contract*, is a court-enforced contract together with a set of self-enforcing bonuses. Following Levin (2003) we focus without loss of generality on *stationary* agreements in which this contract remains constant across time.

Solution concept

It is standard in the study of contracting with private information to appeal to Bayesian Nash equilibrium or Bayesian mechanism design. This approach, however, has two important drawbacks. First, as observed by Wilson (1985), Bayesian Nash mechanism design typically imposes major requirements on the information possessed by the designer(s) of the institution. Second, these equilibria may not be robust to perturbations of what agents believe about each other (Bergemann and Morris 2005). To avoid these pitfalls, we focus on the more demanding problem of designing mechanisms in which truth telling is a mutual best response *regardless* of what players believe about their peers. (Formally, we focus on the problem of (partial) ex-post implementation of the first best). This concept requires that, given that all her peer's report truthfully, each player i finds reporting her type truthfully to be a best response to *every* profile of types her peers may have.[3]

Preliminaries

The team faces the following constraints:

1. Each player must be willing to report her type truthfully, regardless of her type and regardless of the types of her peers:

$$V_i(\theta_i; \theta) \geq V_i(m_i; \theta) \text{ for all } \theta, m_i. \qquad (IC)$$

2. Given truthful reporting, players must be willing to exert first-best efforts. In other words, selection of $e_i^*(\theta)$ for each i a mutual best response

[3] This condition implies that truth-telling is optimal for any possible belief that i may hold about θ_{-i} and is thus stronger than partial robust implementation. Accordingly, the concept is robust to information leakage or heterogeneous beliefs about the underlying type profile. We do not focus on the related question of full ex post implementation (Bergemann and Morris 2009).

at every θ, taking the effect of effort on transfers and future play into consideration.[4]

3. Each player must be willing to honor the self-enforcing transfer, regardless of the state:

$$\underbrace{\sup_{\theta} \; b_i(\theta)}_{i\text{'s maximum reneging temptation}} \leq \underbrace{\frac{1}{r}\left[\mathbb{E}V_i(\theta) - v_i\right]}_{i\text{'s relational capital}}. \qquad (DE)$$

(In addition, the budget must be balanced: $\sum_i b_i(\theta) = 0$ for all θ.)

We say that a relational contract achieves the first-best if it satisfies all constraints for a finite level of relational capital.

For the environments we consider below, the effort constraint turns out to be slack. For ease of exposition, we begin by ignoring this constraint. Then, in the Appendix, we show that doing so is in fact valid.

As we shall see, a theory of incentives – and organizational form – emerges from the interaction of (IC) and (DE).

4 Productivity is private information

Here we study the case in which the state θ affects the output function $f(e, \theta)$ (but not the players' effort costs). This setting allows players to be privately informed, for example, about the demand for the team's product, about the marginal value of each other's efforts, or about the degree to which their efforts complement each other. In this section, we normalize $c(e_i) = e_i$.

We begin by showing that there is a unique family of relational contracts – which we call "surplus-sharing" agreements – that implement the first best (each one requiring a potentially different level of relational capital).

A surplus-sharing arrangement is a contract in which each player's payoff $V_i(\theta)$ has the following canonical form:

$$V_i(\theta) = \alpha_i S(\theta) + W_i(\theta_{-i}) \text{ for all } i,$$

[4] Joint deviations at the reporting and effort exertion stages must also be deterred This joint constraint, however, is always slack.

where $\alpha_i \in [0,1]$ is the player's court-enforced output share and $W_i(\theta_{-i})$ is independent of player i's type. In other words, other than the wage (comprised of the court enforced w_i and possibly transfers), player i receives a share of (net) surplus equal to her share of (gross) output.[5]

Interpretation: The team has a set of owners (according to shares α). These owners promise each player in the team (owners included) a self-enforcing bonus equal to her effort cost. When output arrives, owners first pay all promised bonuses (in exact proportion to their output shares) and then split the output that remains. (In addition, players may transfer fixed payments $W_i - w_i$ amongst themselves.)

Theorem 1 *If a relational contract achieves the first best, then it is a surplus-sharing agreement. Conversely, for any profile of shares α there exists a surplus-sharing agreement with shares α that achieves the first best (for a finite level of relational capital).*

Proof. We first show that every first-best contract must be a surplus-sharing one. Notice that, for any given contract,

$$V_i(\theta) \;=\; \underbrace{\alpha_i f(e^*(m),\theta) + w_i}_{\text{court-enforced payment}} \;-\; b_i(m) \;-\; e_i^*(m) \text{ for } m = \theta.$$

We now use the Envelope Theorem twice:[6]

1. Truth telling requires that

   $$V_i(\theta) = \int_0^{\theta_i} \alpha_i \frac{\partial}{\partial\theta_i} f(e^*(z,\theta_{-i}), z, \theta_{-i})\, dz + V_i(0,\theta_{-i})$$

   (where z is a dummy variable that takes the place of θ_i). The integrand is the *direct* derivative of the player's payoff with respect to his true type θ_i – which, crucially, affects her payoff only through its direct effect on output.

[5] Since the shares α_i add up to one, the wages $W_i(\theta_{-i})$ add up to zero.

[6] The Envelope Theorem we use is Theorem 2 of Milgrom and Segal (2002). We must verify two conditions to use this theorem (absolute continuity and a bounding condition): By assumption $f(e,\theta)$ is continuously differentiable in θ and is therefore absolutely continuos. Second, the function, $f_{\theta_i}(e,\theta)$ of e is bounded above by $f_{\theta_i}(\bar{e},\theta)$ for each θ.

2. The fact the surplus is maximized for all θ requires that

$$S(\theta) = \int_0^{\theta_i} \frac{\partial}{\partial \theta_i} f\left(e^*\left(z, \theta_{-i}\right), z, \theta_{-i}\right) dz + S\left(0, \theta_{-i}\right)$$

The integrand is the *direct* derivative of surplus with respect to θ_i – which, crucially, affects surplus only through its direct effect on output.

By combining these two expressions we obtain $V_i(\theta) = V_i(0, \theta_{-i}) + \alpha_i S(\theta) - \alpha_i S(0, \theta_{-i})$. Namely, each agent captures a constant plus share α_i of her type's contribution to team surplus. Therefore, by setting $W_i(\theta_{-i}) = V_i(0, \theta_{-i}) - \alpha_i S(0, \theta_{-i})$, we obtain the desired result.

Conversely, fix α. We may construct a surplus-sharing arrangement that achieves the first best as follows. For all i, set $W_i(\theta) = w_i$ such that $\sum_i w_i = 0$ and $\alpha_i \mathbb{E}S(\theta) + w_i - v_i > 0$ (which is possible because $\mathbb{E}S(\theta) > \sum_i v_i$).[7] On the one hand, (IC) follows from the fact that each player's payoff is proportional to the team's surplus. On the other hand, (DE) follows from the fact that player i's bonus is

$$b_i(\theta) = \alpha_i \sum_j e_j^*(\theta) - e_i^*(\theta).$$

This bonus is bounded above by $\alpha_i \sum_{j \neq i} \overline{e}_j$. As a result, a finite level of relational capital suffices to overcome the player's maximum reneging temptation.

■

This Theorem is best understood by describing the players' bonuses. Consider first a team with two players. Under any given surplus-sharing agreement, player must pay a bonus

$$b_1(\theta) = \underbrace{\alpha_1 \cdot e_2^*(\theta)}_{\text{2's externality on 1}} - \underbrace{\alpha_2 \cdot e_1^*(\theta)}_{\text{1's externality on 2}}$$

(and player 2's bonus is symmetric). In other words, player 1 must pay for player 2's effort cost in proportion to 1's profit share and, simultaneously, player 2 must pay for player 1's effort cost in proportion to 2's profit share. (And, to minimize the reneging temptation, only the resulting net payment changes hands). This arrangement guarantees that players internalize all ex-

[7] In this case, $b_i(\theta) = \alpha_i \sum_j e_j^*(\theta) - e_i^*(\theta)$ and therefore $\sup_\theta b_i(\theta) \geq 0$.

ternalities they cause on each other.

For example, when player 1 is the sole owner ($\alpha_1 = 1$), she must pay 100% of player 2's effort cost, while receiving no payment in return. This arrangement is needed so that player 1 does not overstate (through her report) the level of effort that player 2 must exert and, at the same time, player 2 does not understate the level of effort that she herself must exert. (Player 1, being the full residual claimant of output, internalizes all benefit of her own effort, and therefore requires no bonus to work hard). In addition player 1 must pay 2 a constant to offset the outside option v_2.

When the team has N players, the non constant part of bonuses are a straightforward generalization of the bonus above:

$$b_1(\theta) = \underbrace{\alpha_1 \sum_{i \neq 1} e_1^*(\theta)}_{\text{peers' externality on 1}} - \underbrace{(1-\alpha_1)e_1^*(\theta)}_{\text{1's externality on her peers}}$$

(and her peer's bonuses are symmetric.) This arrangement guarantees that players internalize the multi-sided externalities they cause on each other.

4.1 Efficient Ownership Structure

Now that we have established that the first-best can be achieved (provided the team has sufficient relational capital) we turn to the problem of finding the *efficient* ownership structure, defined as the profile of output shares α, and associate bonus payments, that minimize the relational capital needed to implement the first best.

Recall that the players' combined reneging temptation is

$$\sum_i \sup_\theta b_i(\theta) = \sum_i \sup_\theta \left[\alpha_i \sum_j e_j^*(\theta) - e_i^*(\theta)\right]$$

(where, crucially, the sup is taken separately player by player). The efficient profile α solves the problem of minimizing this combined reneging temptation:

$$\min_{\alpha} \quad \sum_i \sup_{\theta} \left[\alpha_i \sum_j e_j^* (\theta) - e_i^* (\theta) \right] \tag{1}$$
$$s.t.$$
$$\alpha \geq 0 \text{ and } \sum_i \alpha_i = 1.$$

Throughout, to avoid knife-edge cases in which more than one arrangement is efficient, we assume that the players' effort ranges $[0, \overline{e}_i]$ are generic in the sense that $\overline{e}_i$ differs across players.

We are ready to characterize the efficient ownership structure for specific environments. Below, we consider two classes of environments of interest. These are described, respectively, by Conditions 1 and 2.

Condition 1 (Full effort range.) An environment exhibits **full effort range** if, for any given player i, there exists a state θ in which player i exerts zero effort and all of her peers exert maximum effort – namely $e_i^* (\theta) = 0$ and $e_{-i}^* (\theta) = \overline{e}_{-i}$.

This condition implies that no single player is needed in all states and importantly requires that in some state in which everyone other than i is working maximally, i is not needed.

Proposition 1 *Under Condition 1 (full effort range) the efficient arrangement is a single owner/manager who receives all profit shares and pays all bonuses of her peers.*

Proof. See Appendix. ■

Intuitively, for any surplus-sharing arrangement the worst case scenario – in terms of bonus payments – for player i is a state in which she exerts no effort and her peers exert maximum effort. In this case, player i's reneging temptation (i.e. the size of the bonus she must pay) is $\alpha_i \sum_{j \neq i} \overline{e}_j$. As a result, the efficient profile of shares α, which minimizes the players' combined reneging temptation, solves

$$\min_{\alpha} \quad \sum_i \alpha_i \sum_{j \neq i} \overline{e}_j$$
$$s.t.$$
$$\alpha \geq 0 \text{ and } \sum_i \alpha_i = 1.$$

Since the reneging temptation of each player increases linearly with her share α_i, it is best to concentrate all shares in the hands of the player with the largest effort range $[0, \overline{e}_i]$, as this is also the player for whom the value of $\sum_{j \neq i} \overline{e}_j$ is smallest.

In other words, by concentrating all shares in the hand of one player, the teams avoids having to pay a bonus to the player for whom this bonus would have to be largest.

In what follows, we assume (without loss) that it is player 1 who has the highest effort range $\overline{e}_1 > \overline{e}_i$ or all $i > 1$.

Note. The condition under which Proposition 1 is derived is that no single player is essential. However, all that is needed for this results is the weaker condition that no single player *other than* player 1 is essential.

An alternative condition on the environments is:

Condition 2 (Joint efforts.) An environment exhibits **joint efforts** if there is a subset of players M containing $m \geq 2$ agents such that

$$\sum_{j \in M} \sup_{\theta} \left[\frac{1}{m} \sum_i e_i^* (\theta) - e_j^* (\theta) \right] < \sum_{i \neq 1} \overline{e}_i.$$

This condition says that players j in the set M exert an effort that is not too small when the total team effort $\sum_i e_i^* (\theta)$ is high. The LHS is the relational capital needed for a surplus-sharing arrangement that splits ownership between the agents in M and the RHS is the relational capital needed for an arrangement where 1 is the owner.

Notice that this condition is quite week. For example, when there are two players, it requires only that the distance between their two efforts $|e_1^* (\theta) - e_2^* (\theta)|$

is no greater than the highest possible effort $\overline{e}_2$ of player 2. In other words, player 2's is at least somewhat important for the team, regardless of the state. When there are more than two players, the condition is weaker still: it suffices that there is a subset of players whose combined effort is not too small whenever the team as a whole works hard.

Proposition 2 *Under Condition 2 (joint efforts) the efficient arrangement is a partnership in which ownership is dispersed across a subset of players.*

Proof. See Appendix. ■

The intuition behind this result is as follows. Suppose we have two players. When shares are split 50-50, player 1 pays for 50% of player 2's effort and vice versa As a result, player 1 's maximum reneging temptation is

$$\sup_{\theta} \frac{1}{2} \left[e_2^* \left(\theta \right) - e_1^* \left(\theta \right) \right]$$

(and player's 2's maximum reneging temptation is symmetric). The smaller the distance between the two efforts, the smaller this reneging temptation: when both players work hard, none of them must pay a large bonus to the other. When instead shares are concentrated in the hands of player 1, her maximum reneging temptation is

$$\sup_{\theta} e_2^* \left(\theta \right),$$

which exposes the fact that player 1 must pay the full effort cost of player 2 without ever receiving money in return, for her own effort. It follows that when players exerts similar efforts (regardless of their overall effort level), shared ownership is best.

When we have more than two players the argument is very similar. Now, for a partnership to be optimal, it suffices to find a subset of players such that the effort of each one of these players is never two far away from the average effort level of the team – namely, all players in this subset are essential to team output.

Note. The above argument also tells us which players are the ideal owners of the team: they are the subset of players who work hardest in those states in which the team as a whole works hard.

4.2 Complex mechanisms

The simple contracts described above, in which the output shares α are constants, are satisfying for several reasons. First, they limit the difficulties faced by a court who wishes to enforce these payments. Second, they are immune to players manipulating the exact time in which revenues are generated. Finally, they provide incentives regardless of the team's baseline level of output. That being said, if there are large gains from added complexity, one might expect more elaborate arrangements to emerge.

Here we show that sufficiently complex court-enforced contracts, provided they can be enforced, achieve the first best without any need for self-enforcing bonuses. Moreover, these contracts are identical, on the path, to the surplus-sharing agreements discussed above.

Consider the following family of court-enforced contracts, which we call complex surplus-sharing. Select an arbitrary profile α of constant shares. Now, define the following function of the players reports θ' and realized output y:

$$\sigma_i(\theta', y) = \alpha_i + \frac{e_i^*(\theta') - \alpha_i \sum_j e_j^*(\theta')}{y}.$$

Next, after players report θ' and output y is realized, have the court distribute y according to shares $\sigma_i(\theta', y)$ and force have players to pay fixed wages w_i amongst themselves (with $\sum_i w_i = 0$). As a result, when $\theta' = \theta$, player i receives a court-enforced payment equal to

$$\sigma_i(\theta, y) \cdot y + w_i = \alpha_i S(\theta) + w_i.$$

Theorem 2 *For any profile α, there exists a complex surplus-sharing agreement that achieves the first best without any self-enforcing bonuses.*

Proof. Fix α. For all i, divide output according to shares $\sigma_i(\theta, y)$ and set w_i such that $\sum_i w_i = 0$ and $\alpha_i \mathbb{E} S(\theta) + w_i - v_i > 0$ (which is possible because

$\mathbb{E}S(\theta) > \sum_i v_i$).[8] (IC) follows from the fact that each player's payoff is proportional to the team's surplus and (DE) is met automatically because no bonus payments are called for. ■

Notice that, on the path, these complex contracts are identical to the simple contracts considered above. The only difference is that, instead of internalizing externalities by adjusting the players' payments with self-enforcing bonuses, these contracts adjust their payments directly. In other words, output shares are fine-tuned so that the effect of a misreport on the output share mirrors the effect of a misreport on the self-enforcing bonuses in the simple surplus-sharing agreements. For example, when i's report causes j to work a lot, i's share of output is reduced. As a result, the incentive to cause others to overwork is negated.

5 Effort costs are private information

Here we assume that the state θ to affects the players' effort costs $c(e_i, \theta_i)$ (but not the output function).[9] Types may now capture, for example, private opportunity costs.

We show that when the types of all players interact, in the sense that the cross partial $S_{\theta_1 \ldots \theta_N}$ is nonzero, no relational contract (with a balanced budget) achieves the first best.

Theorem 3 *Suppose effort costs are private information and* $S_{\theta_1 \ldots \theta_N} \neq 0$. *Then, there is no mechanism that respects the teams' budget and implements the first best.*

Proof. Notice that, for any given contract,

$$V_i(\theta) = \underbrace{\alpha_i f(e^*(m)) + w_i}_{\text{court-enforced payment}} - b_i(m) - c(e_i^*(m), \theta_i) \text{ for } m = \theta.$$

[8] Players with positive α_i typically receive a negative wage and vice versa. We assume that players have sufficient liquidity to make these payments.

[9] Theorem 2 generalizes to the case in which θ affects f as well.

We now use the Envelope Theorem twice:[10]

1. Truth telling requires that

$$V_i(\theta) = -\int_0^{\theta_i} \frac{\partial}{\partial \theta_i} c\left(e^*\left(z, \theta_{-i}\right), z\right) dz \; + \; V_i(0, \theta_{-i})$$

   (where z is a dummy variable that takes the place of θ_i). The integrand is the *direct* derivative of the player's payoff with respect to his true type θ_i – which, crucially, affects her payoff only through its direct effect on cost.

2. The fact the surplus is maximized for all θ requires that

$$S(\theta) = -\int_0^{\theta_i} \frac{\partial}{\partial \theta_i} c\left(e^*\left(z, \theta_{-i}\right), z\right) dz \; + \; S(0, \theta_{-i})$$

   The integrand is the *direct* derivative of surplus with respect to θ_i – which, crucially, affects surplus only through its direct effect on cost.

By combining these two expressions we obtain $V_i(\theta) = V_i(0, \theta_{-i}) + S(\theta) - S(0, \theta_{-i})$. Namely, a constant of integration plus the contribution of θ_i to surplus. It follows that

$$\sum_i V_i(\theta) = NS(\theta) + \sum_i \left[V_i(0, \theta_{-i}) - S(0, \theta_{-i})\right].$$

As a result, budget balance (namely, $\sum_i V_i(\theta) = S(\theta)$) requires that

$$\sum_i \left[V_i(0, \theta_{-i}) - S(0, \theta_{-i})\right] = -(N-1)\, S(\theta) \text{ for all } \theta.$$

Since $\frac{\partial^N}{\partial \theta_1 \ldots \partial \theta_N} \sum_i \left[V_i(0, \theta_{-i}) - S(0, \theta_{-i})\right] = 0$ (by construction) and $\frac{\partial^N}{\partial \theta_1 \ldots \partial \theta_N} S(\theta) \neq 0$ (by assumption) this equality cannot hold. ■

[10] The Envelope Theorem we use is Theorem 2 of Milgrom and Segal (2002). We must verify two conditions to use this theorem (absolute continuity and a bounding condition): By assumption $c(e_i, \theta_i)$ is continuously differentiable in θ and is therefore absolutely continuos. Second, the function, $c_{\theta_i}(e_i, \theta_i)$ of e_i is bounded above by $c_{\theta_i}(\bar{e}, \theta)$ for each θ.

The proof of Theorem 3 shows that a necessary condition to achieve the first-best is that each player's payoff has the form

$$V_i(\theta) = S(\theta) + W_i(\theta_{-i}).$$

Therefore, unlike a surplus-sharing agreement in which player i receives only a fraction of surplus, each play must receive 100% of surplus. But since there are N players, and only one actual surplus to divide among them, the team lacks the necessary budget (i.e. is short $N-1$ multiples of surplus).

The only remedy, therefore, is for the team to bring in a budget breaker.

6 Discussion

Our theory allows us to rationalize and understand some key features of a wide range of contracts. Specifically, sharing contracts have been prevalent for a long time in history. Consider, for instance, fishermen (Arrunada and Gonzalez, 1997):

> "Firstly, a small quantity of fish is distributed as compensation in-kind among crewmembers. Gross revenue obtained from selling the rest of the catch is then reduced by the commission paid to the market (between 3.5 and 4.5%). Next to be paid are common expenditures, which are collectively supported, such as food provisions, Social Security charges, fishing tackle (nets, rods, lines, etc.), bait, salt and ice. The shipowner finances these common expenditures, recovering them only after the catch is sold. If the value of the catch is lower than the expenditures, the difference is accumulated and deducted from future catches.
>
> At this time, a fixed compensation is also paid to the crew. Fishermen thus receive a small wage coming directly from the catch value, as do common expenditures before the main partition of net revenue between capital and labor is made. This is not strictly a fixed compensation since it is defined most commonly as a percentage of gross revenue or as a fixed amount per month whichever is lower, and fishermen do not earn it if the catch is not valuable

> enough to meet all the common expenditures. In practice, however, it can be considered as a fixed wage given that catches small enough not to recover such expenditures are very rare and even in this case payment of expenses might be postponed.
>
> Once the common expenditures and this fixed compensation are deducted, the resulting net revenue is divided into two parts, one for the ship-owner (the "ship's part") and the other for the crew ("the people's part").... "The 49% allotment received by the crew is then allocated among crewmen according to a previously agreed formula. Frequently, it is divided by the number of fishermen plus 1.5 (so that the value of one share is 0.49 [CV-SE] / [n+1.5], where n is the total number of people in the boat, CV is the catch value and SE the total shared expenditures). Every crewman earns one of these shares except for the fishing skipper who receives two and the coast skipper who takes one and a half. Traditionally, the sharing is made openly, so that every crewman can monitor the prices and the shares earned by others. From his 51% participation in net revenue the shipowner pays two additional shares, one to the fishing skipper and another one to the machinist, as well as one half or one quarter of a share"

Our analysis shows that this apparently arbitrary agreement is an optimal one. In fact, as Theorem 1 shows, the first-best can only be achieved with a surplus-sharing rule of this form.

Our theory also helps us understand the structure of the sharing contracts. Consider, instead of fishing, the law, where similar formulas are used. Roughly, those working in law firms are rewarded in three ways. Many receive a wage, some receive bonuses, a few enjoy a partnership stake. Partners compensate associates (and themselves) by first, paying them their hourly wage. What is left over (which is the surplus) is then divided up among the partners. The components of this mechanisms are very much in alignment with the ones our theory leads us to expect.

This type of relational and formal contracts are also relevant between firms. Consider shopping malls. The sales of each store depend on the decisions of all other stores. For instance, the cleanliness and quality of the display of a store makes each store more attractive. As Gould et al. (2005) show, externalities between stores are internalized by subsidizing the rents of the stores that generate mall traffic for other stores (which requires charging rent premia to pay for them). Anchor stores, for instance, occupy 58% of the space in an average mall, but only pay 10% of total rent. What complicates the problem, as our model highlights, is that the impact of each store, the ability that each store has to increase its own soles and those of others, is only known to the store itself.

When should ownership be concentrated or dispersed?

Our analysis also allows us to derive empirically-testable predictions by exploiting variations in the team production technology. We obtains two main results in this respect:

a) Single ownership (Proposition 1): When players' efforts vary independently from each other across different states (e.g. it is possible for a player's effort not to be needed while the other team members need to exert a lot of effort), then having a single owner/manager is efficient. Intuitively, that would be the player for whom bonuses must be largest. By concentrating shares on her, the reneging temptation is minimized at the lowest cost. Moreover, to prevent this player from misreporting her information, she must pay bonuses of her peers in proportion to the effort they are asked to exert. As a result, this player emerges, endogenously, as a principal. As solution that mirrors Alchian and Demsetz's (1972) classic arrangement.

b) Dispersed ownership (Proposition 2): When instead the optimal efforts of at least some subset of the team are in sync with the overall effort of the team (e.g. several players in team are essential). Shares should be awarded to those players who work hardest in those states in which the team as a whole works hard.

To understand these two results, consider a simple example from a well known business, the music industry. Consider a single star musician with a group of backup singers and musicians. Our analysis tells us that ownership to be concentrated on the main musician, whereas others would receive a fixed wage and a bonus. Instead, when the performers are a band, so that each one of their efforts is necessary, and so these efforts co-move, our theory tells us that they will be co-owners. By canceling out bonuses in each direction, this formula minimizes the reneging temptation of all players.

References

Alchian, A. A., & Demsetz, H. (1972). Production, information costs, and economic organization. The American economic review, 62(5), 777-795.

Arrow, K. E. (1974). The limits of organization.

Arruñada, B., & González Díaz, M. (1997). How competition controls team production: The case of fishing firms.

Baker, G., R. Gibbons, and K. Murphy (1994). "Subjective Performance Measures in Optimal Incentive Contracts." *Quarterly Journal of Economics* 109(4): 1125-1156.

Bernstein, S., & Winter, E. (2012). Contracting with heterogeneous externalities. American Economic Journal: Microeconomics, 4(2), 50-76.

Bergemann, D., & Morris, S. (2009). Robust implementation in direct mechanisms. The Review of Economic Studies, 76(4), 1175-1204.

Bloch, F., & Quérou, N. (2013). Pricing in social networks. Games and economic behavior, 80, 243-261.

Bull, C. (1987). "The Existence of Self-Enforcing Implicit Contracts." *Quarterly Journal of Economics* 102(1): 147-160.

Candogan, O., Bimpikis, K., & Ozdaglar, A. (2012). Optimal pricing in networks with externalities. Operations Research, 60(4), 883-905.

Cavendish, M. (2013) At Speed: My Life in the Fast Lane. VeloPress.

Cremer, J., Garicano, L., & Prat, A. (2007). Language and the Theory of the Firm. The Quarterly Journal of Economics, 373-407.

Dessein, W., & Santos, T. (2006). Adaptive organizations. Journal of Political Economy, 114(5), 956-995.

Dessein, W., Garicano, L., & Gertner, R. (2010). Organizing for synergies. AEJ: Microeconomics, 2(4), 77-115.

Fainmesser, I. P., & Galeotti, A. (2015). Pricing network effects. The Review of Economic Studies, rdv032.

Fuchs, W. (2007). Contracting with repeated moral hazard and private evaluations. The American Economic Review, 97(4), 1432-1448.

Halac, M. (2012). "Relational contracts and the value of relationships." *American Economic Review* 102(2): 750-779.

Hölmstrom, B. (1979). Moral hazard and observability. The Bell journal of economics, 74-91.

Holmstrom, B. (1982). Moral hazard in teams. The Bell Journal of Economics, 324-340.

Holmstrom, B. (1999). The firm as a subeconomy. Journal of Law, Economics, and organization, 15(1), 74-102.

Holmstrom, B., & Milgrom, P. (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. Journal of Law, Economics, & Organization, 7, 24-52.

Holmstrom, B., & Milgrom, P. (1994). The firm as an incentive system. The American Economic Review, 972-991.

Bentley MacLeod, W. (2003). Optimal contracting with subjective evaluation. The American Economic Review, 93(1), 216-240.

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

MacLeod, W. B., and J. Malcomson (1989). "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment." *Econometrica* 57(2): 447-480

MacLeod, W. B., and J. Malcomson (1998). "Motivation and Markets." *American Economic Review* 88(3): 388-411.

Milgrom, P., & Segal, I. (2002). Envelope theorems for arbitrary choice sets. Econometrica, 70(2), 583-601.

Rayo, L. (2007). "Relational incentives and moral hazard in teams." *Review of Economic Studies* 74(3): 937-963.

Sakovics, J., & Steiner, J. (2012). Who matters in coordination problems?. The American Economic Review, 102(7), 3439-3461.

Segal, I. (1999). Contracting with externalities. Quarterly Journal of Economics, 337-388.

Skrzypacz, A., & Toikkaa, J. (2015). Mechanisms for repeated trade. American Economic Journal: Microeconomics, 7(4), 252-293.

7 Appendix: The effort provision constraint and joint deviations

We suppressed the effort provision constraint. When a player either has all shares, or has no shares, the effort constraint for this player is slack. In this case, the prospect of losing her bonus following an effort deviation is sufficient to induce her to select $e_i^* (\theta)$.

When instead a player is a partial owner, she my profit from the following composite deviation: at the reporting stage she reports her highest type (to maximize the effort chosen by her peers), in the effort selection stage she selects the effort level that maximizes her share of output $\alpha_i y$ minus her effort cost, and in the bonus payment stage she reneges on all bonuses, after which she is kicked out of the team. This deviation maximizes the player's "static" one-period payoff.

To rule out such deviations we assume that players have the ability to directly enforce the profile $e^* (m)$ for any given report m. One possibility is that players have the ability to post a "bond" that is lost following an effort deviation. This punishment may represent, for example, a loss of reputation in the labor market. Another possibility is that players commit to a group punishment such that any level of shirking leads to all players losing their output (e.g. Holmstrom, 1982). A final possibility is that players rely on peer pressure to ensure that, once a profile $e^* (m)$ has been agreed upon, all players are compelled to comply (e.g. Kandel and Lazear, 1992).

When the above options are not available, players must rely exclusively on their relational contract. In this case, a player who shirks is merely kicked out of the team, therefore losing all her continuation surplus. In the version of the model without private types, Rayo (2007) shows that the total relational capital needed to induce first-best efforts is minimized when shares are dispersed. As a result, when private types are introduced, the reporting constraints either leads to a single owner (Proposition 1) and the effort constraint is slack, or leads to dispersed shares, in which case both the reporting and effort constrains push the team toward dispersed shares.

8 Appendix: Proof of Propositions 1 and 2

Proof of Proposition 1. Under Condition 1,

$$\sup_{\theta} \underbrace{\alpha_i \sum_j e_j^*(\theta) - e_i^*(\theta)}_{b_i(\theta)} = \alpha_i \sum_{j \neq i} \overline{e}_j$$

Therefore, problem (1) simplifies to

$$\min_{\alpha} \quad \alpha_i \sum_{j \neq i} \overline{e}_j$$
$$s.t.$$
$$\alpha \geq 0 \text{ and } \sum_i \alpha_i = 1.$$

Since the objective is linear in α, to solution is to set $\alpha_k = 1$ for player $k = \arg\max_i \sum_{j \neq i} \overline{e}_j$ and $\alpha_i = 0$ for all other player. ■

Proof of Proposition 2. Select an arbitrary subset of players M containing $m \geq 2$ players. When profits shares are divided evenly among them (namely, $\alpha_i = \frac{1}{m}$ for all $i \in M$) the objective in problem (1) takes the value

$$\sum_i \left[\alpha_i \sum_j e_j^*(\theta) - e_i^*(\theta) \right] = \sum_{j \in M} \sup_{\theta} \left[\frac{1}{m} \sum_i e_i^*(\theta) - e_j^*(\theta) \right].$$

When instead all shares are concentrated in the hands of player 1 (the ideal candidate for single ownership), the objective in problem (1) takes the value

$$\sum_i \left[\alpha_i \sum_j e_j^*(\theta) - e_i^*(\theta) \right] = \sum_{j \neq 1} \overline{e}_j.$$

Under condition 1, the former value exceeds the latter. As a result, single ownership is dominated. ■