

Discussion Paper No. 1143R

**INCENTIVE CONTRACTS IN TWO-SIDED MORAL
HAZARDS WITH MULTIPLE AGENTS***

by

Nabil I. AL-NAJJAR**

November 1993

Revised: June 1996

(*) Forthcoming in *Journal of Economic Theory*

(**) Department of Managerial Economics and Decision Sciences, J. L. Kellogg
Graduate School of Management, Northwestern University, 2001 Sheridan
Road, Evanston, IL, 60208.

[Phone: 847-491-5426; Fax: 847-467-1220; e-mail: al-najjar@nwu.edu]

This is a revised version of an earlier paper titled "Incentive Contracts and Statistical Inference in Large Team Moral Hazards". I thank Michael Riordan, Dilip Mookherjee, Doug Allen, John Ledyard, two referees of this Journal and seminar participants at Laval, Northwestern, Princeton, Penn and MIT for their comments. This research was partially funded by a grant from the Social Sciences and Humanities Research Council of Canada. All remaining errors and shortcomings are my own.

Abstract:

The paper studies a contracting problem in which a Principal enters in two-sided moral hazards with N independent agents. There are no technological or informational linkages between the N agency problems. Despite this independence, optimal incentive schemes essentially eliminate the Principal's incentive problem when team size N is large enough. Reputation-like effects appear in our static setting through an improved aggregation of information about the actions of the Principal. One implication of this result is that it is generally suboptimal to require each agent's compensation to depend only on his own outcome. Another implication is the existence of purely *informational economics of scale* to increasing team size. Thus, the concentration of otherwise unrelated transactions in a single 'firm' creates wealth through a more efficient use of information about the Principal's actions.

The paper shows that extremely simple *statistical contracts* are approximately optimal in large teams. The outcome of such contracts is observationally indistinguishable from standard Principal-Agent contracts. This provides a theoretical justification for using standard Principal-Agent contracts in environments that involve two-sided hazard in a fundamental way.

Journal of Economic Literature Classification Numbers: D21, D23.

1. INTRODUCTION

This paper examines a class of contracting problems in which a team consisting of a Principal and N agents produces a vector of outcomes y_1, \dots, y_N . The Principal and agent n take unobserved actions a_n and b_n that determine the distribution $\pi(a_n, b_n)$ of the outcome of the n th agency problem, y_n . The Principal is free to vary his actions across agents.

Two-sided moral hazards with multiple agents may be viewed as a special type of partnership games in which the Principle plays the role of a central partner.¹ Examples and potential applications of our model abound. One is that of a franchisor operating a network of N franchisees. The outcome y_n of the n th franchise outlet is a vector of performance measures. This outcome is random with a distribution jointly determined by the franchisor's and franchisee's inputs a_n and b_n . Another example is a firm selling a product to N consumers. The firm's choice of quality and the consumer's effort in caring for the product jointly determine a probability distribution over the possible levels of product performance (*e.g.*, frequency of breakdowns and cost of repairs). A third example is that of a firm dealing with N employees. Each employee chooses a level of effort, while the firm chooses such things as the employees' working conditions and the safety and productivity of the equipment they use. For example, the employer may be a retailer who can influence the productivity of his salespersons through investment in marketing research and the development of new products.

Optimal incentive contracts in two-sided moral hazards with multiple agents can be quite complex. One reason is that it is not a priori obvious whether the optimal incentive scheme s_n used with agent n should condition this agent's compensation on the entire team outcome or on his outcome alone. Since the optimal contract cannot be guaranteed to take the simpler form $s_n(y_n)$, one has to examine, in principle, all incentive schemes of the form $s_n(y_1, \dots, y_N)$ where agent n 's compensation might depend on the outcomes of all other agents, possibly in a very complex way.²

In the model of this paper, the N agency problems are identical and independent: (1) The Principal's costs are additive across agents; (2) The outcomes y_n are independent given the vector of Principal's and agents' actions; and (3) The Principal can freely and costlessly vary his actions a_n across agents. Assumption (1) eliminates technological economies of scale, while (2) says that there is no common uncertainty that can be used to improve the agents' incentives through relative performance evaluation.

The absence of any physical or informational linkage between the N agency problems might give the impression that optimal incentive schemes in a large team are merely N -replicas of the optimal schemes in the single-agent problem. If this were true, then there can be no gains to conditioning the agents' payoffs on the other agents' outcomes and no advantages or disadvantages to changing team size N .

The results of this paper show that this conclusion is incorrect: Contracts in which agent n 's compensation depends only on his own outcome y_n are not optimal in general. Conditioning on the entire vector of team outcomes is valuable because it improves the aggregation of information about the actions of the Principal, hence reducing (and, in the limit, eliminating) the cost of providing him with incentives. The surprising aspect of this result is the fact that it can be obtained using contracts of very simple structure and that it holds even though the Principal has complete freedom to vary his actions across agents.

To describe the results in more details, let V_1 be the Principal's expected payoff in the optimal contract in the one-agent case ($N = 1$). Imagine now that the Principal is somehow able to credibly and costlessly commit to taking any action of his choice (action a^+ , say) with all agents. Given that action a^+ has been taken, the Principal's side of the moral hazard problem is eliminated, so what was originally a two-sided moral hazard is effectively transformed to a standard Principal-Agent problem with outcome function $\pi(a^+, \cdot)$. Let $s^+(y_n)$ be an optimal Principal-Agent contract when the Principal takes the action a^+ , and let $V^+ > V_1$ denote the corresponding expected payoff of the Principal. The paper's main result is that the average expected payoff of the Principal in an optimal incentive scheme is approximately V^+ provided N is sufficiently large. Thus, the Principal's incentive problem is essentially eliminated through an improved aggregation of information about his actions. This frees the incentive contract to deal exclusively with the agents' moral hazard, much like in the standard Principal-Agent model.

One implication of this result is that it is generally incorrect to assume that the naive contract consisting of replicating the optimal contract of the single-agent problem will be optimal when the Principal is dealing with multiple agents.³ The result also implies that the conclusions of models with two-sided moral hazard and a single agent are not robust to increasing the number of agents. This is significant because models with two-sided moral hazard are often intended as models of environments with a large number of agents, as in the franchising and product quality examples mentioned earlier.

The contrast is particularly apparent when agents are risk-neutral because our results would then imply that with multiple agents first-best allocation can be arbitrarily closely approximated. In particular, the inefficiencies that drive the qualitative conclusions of two-sided moral hazard models in the single agent case are significant only for low values of N .⁴

Another implication is the existence of purely *informational economics of scale* to increasing team size.⁵ Viewing N as a measure of firm size, our results suggest that concentrating otherwise unrelated transactions in a single large ‘firm’ creates wealth through a more efficient structure for generating and using information about the firm’s actions. This is true despite the absence of any obvious technological advantages to organizing production in larger firms. The informational economics identified in this paper also constitute a source of competitive pressures to increase firm size (*e.g.*, through mergers, expansions or predation), suggesting that asymptotic arguments where N increases have economic content in our context.

An innovation of this paper is the method used to derive the results which is of interest both for its simplicity and for generating substantial detailed information about how optimal allocations may be implemented. Instead of going through the daunting task of calculating and characterizing the fully-optimal contracts in the N -agent problem, I focus instead on extremely simple and tractable class of contracts that bypass the Principal’s optimization problem. These *statistical contracts*, motivated by Green [7], divide output on the basis of the value of a simple test statistic. While naive statistical contracts are typically suboptimal, the Principal cannot do worse in the optimal contract than in any statistical contract. The main result then follows by showing that there is a statistical contract that yields an average payoff close to V^+ when N is sufficiently large.

One nearly optimal statistical contract has an equilibrium in which there is high probability that all agents will receive the payments predicted by the standard Principal-Agent model with one-sided moral hazard. This equilibrium is sustained by penalties triggered by a drop in a test statistic, which is interpreted as evidence that the Principal deviated. Such a drop has an arbitrarily small probability when equilibrium strategies are played. Consequently, an outside observer is likely to find strong evidence in favor of the conclusion that the agency relationship is organized as a standard Principal-Agent relationship, even though two-sided moral hazard is present in a fundamental way.

Stated differently, the standard Principal-Agent model may perform well even when it in fact fails to capture a fundamental aspect of the incentive problem. This point is significant because Principal-Agent models are often applied to relationships in which the Principal's moral hazard is an important factor (as in franchising, product quality, and employment relationships). Yet a modeler who applies the standard Principal-Agent model to such relationships is effectively focusing on the agent's incentives and implicitly assuming that the Principal's incentive problem can be ignored. The analysis of the paper makes it possible to justify and appropriately qualify this assumption.

An important qualification has to do with the value of the Principal-Agent model in comparative statics. Imagine an outside observer who uses the Principal-Agent model to predict how the contract will change in response to exogenous changes in the parameters defining the agency relationship. The prediction reached will generally be incorrect. The problem is that by overlooking the two-sided nature of the moral hazard problem, this outside observer will focus exclusively on changes in the terms of the contracts and will not take into account the Principal's freedom to change his actions in response to changes in the environment.

The results of this paper have a dynamic, reputation-like flavor. This is surprising in view of the fact that our framework considers a static relationship where outcomes are revealed once and for all at the moment of contract execution, rather than gradually during the play of a repeated game. Of direct relevance to our analysis is Radner's [13] model of cooperation in undiscounted repeated games. Radner considers an infinite horizon model of a repeated partnership in which players maximize average payoffs. A cooperative outcome can be sustained because, roughly, a player can benefit only by deviating an infinite number of periods. But such deviations would generate massive evidence that permits accurate detection and punishment.⁶ In our model, a deviation from a^+ with a large fraction of agents will similarly generate massive information when N is large. This information can be used to deter a Principal maximizing average expected payoff from making 'large deviations'. The incentive scheme we find may be consistent with the Principal occasionally deviating by playing something other than a^+ with some agents. However, the scheme ensures that the Principal will have the incentives to play a^+ with high probability with all but a vanishing fraction of agents.⁷

2. THE ONE-AGENT MODEL

2.1. Actions, Outcomes and Payoffs

The Principal and the agent have finite action sets A and B with generic elements denoted a and b respectively.⁸ Let Δ_A and Δ_B denote the simplexes defined on the finite sets A and B respectively. A Principal's mixed strategy is a probability distribution $\alpha \in \Delta_A$, while an agent's mixed strategy is $\beta \in \Delta_B$. Let $\alpha(a)$ and $\beta(b)$ denote the probabilities with which actions a and b are chosen under the mixed strategies α and β respectively.

There is a finite set of possible outcomes: $Y = \{y^1 = 0, y^2, \dots, y^M\} \subset [0, \infty)$, which is interpreted as a set of possible output or profit levels. The outcome y is a random variable whose distribution is determined by the actions of the Principal and the agent. Using Δ_Y to denote the simplex in \mathbb{R}^M representing all probability distributions on Y , the distribution of y is given by the *outcome function* $\pi : A \times B \rightarrow \Delta_Y$, with $\pi_m(a, b)$ denoting the probability of $y = y^m$ and $\pi(a, b)$ denoting the vector of probabilities. To make the problem interesting, I impose the standard assumption:

A1: $\pi(a, b)$ has full support in the sense that $\pi_m(a, b) > 0$ for every a, b and y^m .

The Principal and the agent randomize independently from each other and from y . The distribution of y under the mixed strategies α and β is given by

$$\pi_m(\alpha, \beta) = \sum_{a,b} \pi_m(a, b) \alpha(a) \beta(b).$$

The Principal is risk-neutral and has a cost function $c : A \rightarrow \mathbb{R}$. The agent has a von Neumann-Morgenstern utility function u and a cost of effort function $g : B \rightarrow \mathbb{R}$. The expected costs under the mixed actions α and β are $c(\alpha) = \sum_a c(a) \alpha(a)$ and $g(\beta) = \sum_b g(b) \beta(b)$.

2.2. Contracts

A *contract* is a function $s : Y \rightarrow \mathbb{R}$, with $0 \leq s(y) \leq y$, representing a binding agreement in which the Principal receives $s(y)$ and the agent receives $y - s(y)$. Note that the definition imposes budget

balancing as well as a rather strong form of limited liability. The agent is assumed to have a reservation value of 0. A contract s defines the payoffs of a non-cooperative game $\Gamma(s)$ where the action set of the Principal is A and that of the agent is B plus the option of rejecting the contract.

The Principal and the agent are assumed to take their actions simultaneously and independently. Their expected payoffs under the mixed strategies α and β are:

$$V(s, \alpha, \beta) = E_{(\alpha, \beta)} s(y) - c(\alpha),$$

$$U(s, \alpha, \beta) = E_{(\alpha, \beta)} u(y - s(y)) - g(\beta).$$

2.3. Two-Sided Moral Hazard with a Single Agent

A mixed profile (α, β) is a (Nash) equilibrium in $\Gamma(s)$ if:

$$V(s, \alpha, \beta) \geq 0 \quad (PC)$$

$$V(s, \alpha, \beta) \geq V(s, \alpha, \beta') \quad \forall \beta', \quad (IC_A)$$

$$V(s, \alpha, \beta) \geq V(s, \alpha', \beta) \quad \forall \alpha', \quad (IC_P)$$

IC_P and IC_A are the familiar incentive constraints for the Principal and the agent, while PC is the agent's participation constraint.⁹ The problem of the Principal can now be formulated as:

$$\max_{s, (\alpha, \beta)} V(s, \alpha, \beta) \quad \text{such that } (\alpha, \beta) \text{ is an equilibrium for } \Gamma(s)$$

Let V_1 denote the Principal's maximum expected payoff in the one-agent problem.

2.4. The Extended Principal-Agent Game

The main issue addressed in this paper is whether increasing team size creates new commitment opportunities for the Principal. If the Principal were able to commit to a given action a , the two-sided moral hazard problem will be effectively transformed to a classical Principal-Agent problem with outcome function $\pi(a, \cdot)$. Given a , the Principal selects a Principal-Agent contract s_a that is optimal

relative to $\pi(a, \cdot)$. One can therefore think of the Principal's choice of a as selecting among a family of Principal-Agent problems parametrized by the outcome functions $\pi(a, \cdot)$.

To formalize this, consider the modification of the last problem with the Principal's incentive constraint removed:

$$\max_{s, (\alpha, \beta)} V(s, \alpha, \beta) \quad \text{subject to } PC, IC_A$$

Assume that (s^+, α^+, b^+) is a solution to this problem and define $V^+ = V(s^+, \alpha^+, \beta^+)$. Note that s^+ must be an optimal contract in the standard Principal-Agent with outcome function $\pi(\alpha^+, \cdot)$. To make the problem interesting, I make the assumption that commitment is valuable to the Principal:

A2: $V^+ > V_1$.

Define the distance between two contracts by $|s - s'| = \max_y |s(y) - s'(y)|$. That is, two contracts are close if they yield approximately the same payments at all states. The optimality of s^+ implies that the agent's best response to (s^+, α^+) is not single-valued (Grossman and Hart [8], Proposition 6). The next assumption guarantees the existence of contracts s_ϵ^+ arbitrarily close to the optimal contract s^+ and to which b^+ is the agent's unique best response:¹⁰

A3: *For every $\epsilon > 0$ there is a contract s_ϵ^+ such that $|s_\epsilon^+ - s^+| < \frac{\epsilon}{M}$ and such that b^+ is the unique best response to (s_ϵ^+, α^+) .*

Thus, s_ϵ^+ is approximately optimal since it guarantees the Principal an expected payoff arbitrarily close to V^+ . If A2 is not satisfied, slightly weaker versions of the results would still be valid. In particular, one can redefine the extended Principal-Agent program by replacing the weak inequalities in IC_A with a strict inequality, limiting the agent's actions to pure actions, and replacing the max with a sup. This would give a limiting expected payoff \hat{V}^+ which will in general be lower than V^+ . On the other hand, if $\hat{V}^+ > V_1$, then the main conclusions about the value of aggregation of information and the benefits from conditioning agents' compensations on the entire vector of outcomes would still be valid.

3. THE N -AGENT MODEL

3.1. Actions, Outcomes and Payoffs

There are N identical agents. Let A^N , B^N and Y^N , denote the sets of all N -vectors of A , B and Y , with generic elements denoted \vec{a} , \vec{b} and \vec{y} respectively. Two action vectors \vec{a} and \vec{b} determine a probability distribution on the set of team outcomes Y^N denoted:

$$P(\vec{y} | \vec{a}, \vec{b}).$$

Individual outcomes are assumed to be conditionally independent in the sense that the random variables (y_1, \dots, y_N) are independent for any given profiles \vec{a} and \vec{b} . Under this assumption, P has a simple decomposition into the probabilities $\pi_y(a_n, b_n)$.

Agent n 's mixed strategy is a probability distribution $\mathcal{J}_n \in \Delta_B$. Agents are assumed to randomize independently from each other and from the Principal, so the product mixed strategies determine the joint distribution ν on agents' action profiles B^N .

A Principal's mixed strategy is a probability distribution μ over the set A^N of all action vectors \vec{a} . Note that there is a fundamental asymmetry between the Principal and the agents because the Principal is free to correlate his randomizations across agents while the agents randomize independently. Since the Principal randomizes independently from the agents, the mixed profile (μ, ν) defines a probability distribution on team outcomes given by:

$$P(\vec{y} | \mu, \nu) = \sum_{\vec{a}, \vec{b}} P(\vec{y} | \vec{a}, \vec{b}) \mu(\vec{a}) \nu(\vec{b}).$$

The expected payoff of agent n is the same as in the one-agent case. Note that the outcome y_n depends on μ only through its marginal μ_n on a_n , so the distribution of y_n is given by $\pi(\mu_n, \mathcal{J}_n)$. Note, however, that since the Principal can correlate his randomizations across agents, μ is not necessarily the product of its marginals.

The Principal is assumed to be risk neutral and has cost function $C : \Delta_A \rightarrow \mathbb{R}$ that is *additive*: If $\alpha \in \Delta_A$ represents the frequency distribution of actions in \vec{a} , then $\frac{C(\vec{a})}{N} = c(\alpha) = \sum_a c(a)\alpha(a)$. This

says that the Principal's cost of effort function displays constant returns to scale. See Section 4.7 for further discussion of this point. I will often use the average cost (per-agent) c instead of the total cost function C . If the Principal uses the mixed strategy μ , then his expected average cost is $c(\mu) = E_\mu c(\alpha)$.

3.2. Contracts

A contract is a vector $\vec{s} = (s_1, \dots, s_N)$ where $s_n(y_n, \vec{y})$ is the share of the Principal in the n th agency relationship as a function of y_n and the team outcome \vec{y} . As in the single agent case, a contract determines the payoffs in a non-cooperative game $\Gamma_N(\vec{s})$ between the Principal and the N agents:

$$V(\vec{s}, \mu, \nu) = \frac{1}{N} E \sum_{n=1}^N s_n(y_n, \vec{y}) - c(\mu).$$

$$U(s_n, \mu, \nu) = E u(y_n - s_n(y_n, \vec{y})) - g(\beta_n).$$

The problem facing the Principal is

$$\max_{\vec{s}, (\mu, \nu)} V(\vec{s}, \mu, \nu) \quad \text{such that } (\mu, \nu) \text{ is an equilibrium for } \Gamma_N(\vec{s})$$

Let V_N denote the per-agent average value of this program (that is, the optimal value of the program divided by N), so V_N is measured on the same scale regardless of N .¹¹

3.3. Anonymous Strategies: Motivation

Equilibria in our contracting setting will involve randomization by the Principal in an essential way. The reason is that it may be impossible to give the Principal an incentive to play exactly α^+ with all agents. Thus, the best one can hope for is to design the incentive scheme in such a way that the Principal will deviate with at most a small fraction of agents, and this behavior may require randomization.

Additional information about the provision of incentives and the structure of contracts in large teams can be obtained by considering Principal's strategies that are, in some sense, anonymous. Roughly, these are Principal's strategy vectors which treat the agents symmetrically while still allowing the Principal to

take different actions with different agents. As we discuss formally in Section 3.4, these will correspond to Principal's randomizations μ that are *exchangeable* across agents. The strategy of our proofs will be to first examine equilibria in the game in which the Principal is restricted to anonymous strategies, then show that these continue to be equilibria in the original game without this restriction.

To motivate anonymous strategies, consider the example of a firm selling a product to $N = 100$ identical consumers. For each unit of the product, the firm can either choose low quality a^l or high quality a^h . A strategy for the firm is a vector of actions (a_1, \dots, a_{100}) with $a_n \in \{a^l, a^h\}$. An anonymous strategy, on the other hand, is a vector of the fractions of the population with which particular actions were taken. In our example, with only two actions, we may think of $\Delta_A = [0, 1]$ as representing the fractions of low quality products. The set of population fractions consistent with $N = 100$ is $\Delta_A^* = \{0, \frac{1}{100}, \frac{2}{100}, \dots, \frac{99}{100}, 1\} \subset \Delta_A$. The interpretation of choosing $\alpha = \frac{3}{100}$, for example, is that the Principal produces 3% low quality units (*i.e.*, chooses action a^l with 3 agents) and 97% high quality units.

An anonymous strategy is one in which the Principal first chooses a population fraction (possibly at random) of low quality products from Δ_A^* , then distributes these products symmetrically across agents. Thus, each agent has the same chance of getting a low quality product as any other agent. In the last example, this means that each consumer knows that 3 out of 100 consumers will end with a^l -products and that his chance of getting one is the same as everyone else. The reason for calling such strategy anonymous is that the Principal only determines the fraction rather than the *names* of the agents with which a particular action is chosen. Such strategy rules out, for example, profiles in which a particular agent believes that he will be singled out for a special punishment or reward by the Principal.

As we hinted earlier, the firm can further randomize over the population fractions α . For example, the firm might choose $\alpha = 3\%$ and $\alpha' = 5\%$ with probability $\frac{1}{2}$ each. This is different from choosing the average frequency $\alpha'' = 4\%$ with probability 1. The idea is randomness in the choice of α represents *aggregate uncertainty* from the perspective of the agents. In our example, each consumer may be uncertain both about whether he will receive a low quality product for a given α , and about the value of α itself.

The restriction to anonymous strategies may be interpreted in a number of ways. In the product quality example, one might think that the firm produces the 100 units, with α percent low quality units, then each consumer selects one unit randomly. Alternatively, one might think of the firm as having a list of the 100 consumers, selecting 3 names at random and producing 3 low quality units for them. Finally, one might think of the firm as knowing from the outset the names of the three consumers to play a^l against, but any particular consumer is not sure about the rule used to make this choice and assumes that his chance of facing a^l is the same as everyone else.

3.4. Anonymous Strategies and Exchangeable Randomizations:

Formal Definitions

Let a^1, \dots, a^L be a list of the actions of the Principal. An *admissible action distribution* for the Principal in the game with N players has the form $\alpha = \frac{1}{N} (k^1, \dots, k^L)$ where each k^l is a non-negative integer and $\sum_{l=1}^L k^l = N$. Let $\Delta_A^* \subset \Delta_A$ denote the set of admissible action distributions and note that this set is finite.

Each action vector \vec{a} has a unique frequency distribution of actions $\alpha = (\alpha^1, \dots, \alpha^L) \in \Delta_A^*$, where α^l denotes the percentage of times a^l is played. Two action vectors \vec{a} and \vec{a}' are equivalent (written $\vec{a} \sim \vec{a}'$) if they have the same frequency distribution of actions. The idea is that we can obtain \vec{a}' from \vec{a} through a permutation of actions labels across agents, keeping fixed the frequency with which any particular action is taken.

An *anonymous strategy* is a probability distribution μ on A^N that is *exchangeable* (or symmetrically dependent). That is, μ has the property that $\mu(\vec{a}) = \mu(\vec{a}')$ whenever $\vec{a} \sim \vec{a}'$. It is easy to see that the marginal distribution μ_n on the action taken with agent n is the same for all agents and is equal to α .¹² An anonymous strategy μ defines a probability distribution on Δ_A^* by letting $\mu(\alpha)$ be the probability of the set of action vectors \vec{a} with frequency distribution α . The restriction to exchangeable randomizations defines a new game $\Gamma_N^*(\vec{s})$ in which the strategy set of the Principal is Δ_A^* instead of Δ_A . Intuitively, in the game with exchangeable randomizations the Principal first randomizes over aggregate frequency distributions over actions (*i.e.*, randomizes over Δ_A^*) then, given the chosen frequency $\alpha \in \Delta_A^*$, he symmetrically assigns actions to agents so that the marginals are equal.

4. RESULTS

This section states and interprets the main results of the paper. Section 5 provides the proofs.

4.1. The Main Result

PROPOSITION 1: *Under assumptions A1-A3, the Principal's expected average payoff in any optimal incentive scheme is approximately V^+ when N is sufficiently large. That is, for any $\epsilon > 0$, there is \tilde{N} such that $V_N > V^+ - \epsilon$ whenever $N \geq \tilde{N}$.*

Note that if the Principal is restricted to incentive schemes of the form $s_n(y_n)$ (i.e., where agent n 's payment depends only on his own outcome), then the best the Principal can hope to achieve is an average expected payoff of V_1 . Proposition 1 therefore implies that in an optimal scheme the compensation of at least some agents must depend on the outcomes of other agents.¹³

An optimal contract might link the compensations of various agents in a potentially very complex way. To gain a better understanding of incentives in large teams, I will focus instead on extremely simple *statistical contracts* and note that the Principal will do at least as well in the optimal contract as in any statistical contract. This approach will generate substantial information about optimal incentive schemes without actually having to compute them.

4.2. Statistical Contracts

A statistical contract \bar{s}^{sc} in the N -agent problem is defined by a one-agent contract $s(y)$ and a test statistic $T : Y^N \rightarrow \{0, 1\}$ such that

$$s_n^{\text{sc}}(y_n, T(\vec{y})) = \begin{cases} s(y_n) & \text{if } T(\vec{y}) = 1 \\ 0 & \text{if } T(\vec{y}) = 0 \end{cases}$$

In this contract, agent n 's payment depends on his output only when the value of the test statistic T is high. A low value of T , on the other hand, will be interpreted as evidence that the Principal has deviated, so he is punished by giving all the output to the agent.

The advantage of these contracts is their simplicity: Any such contract is completely determined by two components, the test statistic T and the one-agent contract s . The next result establishes the existence of an incentive scheme supported by a statistical contract of a particularly simple form. First, define $\gamma(\vec{y}) = (\gamma_1(\vec{y}), \dots, \gamma_M(\vec{y})) \in \Delta_Y$ to be the vector of empirical frequencies corresponding to \vec{y} (that is, $\gamma_m(\vec{y})$ is the percentage of times outcome y^m occurs in the vector \vec{y}).

PROPOSITION 2: *For every $\epsilon > 0$ there is an integer \tilde{N} and a statistical contract s_n^{sc} defined by (s_ϵ^+, T) such that:*

- i) $|s_\epsilon^+ - s^+| < \epsilon$;
- ii) T depends on \vec{y} only through $\gamma(\vec{y})$
- iii) $\Gamma_N(s^{sc})$ has a Nash equilibrium (μ, ν) with the properties:
 - 1- μ is exchangeable
 - 2- $P\{\vec{y} : T(\vec{y}) = 1 \mid \mu, \nu\} > 1 - \epsilon$ whenever $N \geq \tilde{N}$;
 - 3- $|W^+ - V_N^{sc}| < \epsilon$ whenever $N \geq \tilde{N}$;

Note that T may be too coarse to be a sufficient statistic for the team outcome \vec{y} so the statistical contract is likely to ignore some potentially valuable information about the team performance and may therefore fail to be optimal. Proposition 2 may be interpreted as saying that the value of the information overlooked by the coarse test statistic becomes small so T becomes, in a sense, approximately sufficient, and the statistical contract becomes approximately optimal. The optimal contract, on the other hand, takes into account and optimizes over every bit of information even when these make a negligible contribution in improving incentives.

4.3. The Form of Incentive Contracts in Large Teams

To the extent that executing complex contracts involves some unmodeled costs, the simplicity of the statistical contract s_n^{sc} found in Proposition 2 suggests that this nearly optimal contract might be a reasonable model for the incentive schemes actually implemented in large teams.

Suppose that the statistical contract s_H^{SC} is indeed the one being used. Think of the conclusions an outside observer might reach by observing the outcomes and compensations for each agent, but without knowing whether or not the underlying contracting problem involves two-sided moral hazard. When $T = 1$, an event with high probability under (μ, ν) , all of the N agents will be compensated according to a contract arbitrarily close to a standard Principal-Agent contract. Since the contract is symmetric, each output level has positive probability, so the N realized output-payment pairs will provide information about compensations at all possible outcomes (with high probability when N is large).

Thus, an outside observer will find overwhelming evidence supporting the conclusion that the contracting relationship is governed by a standard Principal-Agent model with one-sided moral hazard. This observer will be able to fairly accurately predict compensation in the N -agent two-sided moral hazard problem using a standard Principal-Agent contract s^+ , ignoring all the subtleties and problems caused by the Principal's moral hazard. This provides a theoretical justification for using Principal-Agent incentive schemes even in environments where the Principal is likely to have the ability to take actions that impact on the agents' productivity and compensation.

Our framework also provides an important qualification for the intuition that the Principal's incentive problems can be ignored by pointing out the role of the special safeguards built into the overall contractual arrangement between the Principal and the agents. In particular, the success of a standard Principal-Agent model in environments with two-sided moral hazards is highly dependent on the presence of a large number of agents.

Finally, the model explains that an outside observer using a standard Principal-agent model to explain compensation in a two-sided moral hazard environment as in the last paragraph will systematically miscalculate the implications of exogenous parameter changes on contract and performance structures. The reason why the standard Principal-agent model will not do very well in comparative static experiments is that it ignores the fact that some adjustments to exogenous changes can be made through changes in the Principal's actions rather than in the observed contractual terms. The role played by these adjustments is taken into account in the extended Principal-agent model of Section 2.4 in which the Principal can respond to exogenous changes using two instruments s and a rather than just s .

4.4. Informational Economies of Scale

An obvious implication of Propositions 1 and 2 is the existence of economies of scale to increasing team size, measured by $V^+ - V$ per agent. The source of these economies in our model is that a large team makes it possible to provide the Principal with incentives at a vanishingly small cost.

There are, of course, other reasons that make increasing team size beneficial to the Principal. Two potential sources of economies of scale are relative performance evaluation¹⁴ and the possibility of using the large team size to implement group penalties through the transfer of output across agents. To isolate and highlight the role played by the aggregation of information about the actions of the Principal, the construction of Proposition 2 ensures that other sources of economies of scale do not arise in the model. Relative performance evaluation is eliminated by the assumption that outcomes are conditionally independent, so the outcome of one agent conveys no useful information about the action taken by another agent. Transfers, on the other hand, are eliminated by construction:

PROPOSITION 3: *The nearly optimal statistical contract \bar{s}^{SC} satisfies the no-transfer condition:*

$$0 \leq s_n^{SC}(y_n) \leq y_n, \quad \text{for every agent } n \text{ and outcome } y_n.$$

The point is that while the optimal contract might involve such transfers (*e.g.*, to spread risk, or to execute group penalties), the fact that one can design a nearly optimal contract without such transfers implies that the net benefit to using them in an optimal contract must become negligible in a large team.

4.5. Special Case: Risk-Neutral Agents

To illustrate the results, consider the case where agents are risk-neutral. Under this assumption, the solution of the one agent case is very simple. In particular, when the Principal cannot commit to a given action, then the relationship is a simple partnership game with two players. In such game, the optimal contract supporting V_1 will generally involve a sharing of output because making the Principal's

payment contingent on output is the only way to give him incentives to choose a high level of effort. The resulting outcome is inefficient in general because increasing the share of the Principal under budget balancing constraint necessarily reduces the agent's incentive to choose a high level of effort.

On the other hand, the solution of the one-agent problem when the Principal can commit to a given action is quite different. From standard Principal-Agent theory, the value of this problem, V^+ , is supported by an optimal contract s^+ in which the Principal receives a fixed payment and the agent absorbs all the risk. By contrast with V_1 , the outcome corresponding to V^+ is in general efficient. Essentially, the classic partnership problem that causes V_1 to be inefficient disappears because the Principal's ability to commit to a given action leaves the contract complete freedom to deal exclusively with the agent's incentives.

Under agents' risk-neutrality, our result that V_N is approximately equal to V^+ means that the average outcome in a large enough team is approximately first-best efficient. This may be viewed as providing an approximate solution to the partnership problem (Legros and Matsushima [10]) for the special class of partnership games with a central partner.

4.6. Failure of Output-Sharing-Under-Risk-Neutrality

An attractive feature of two-sided moral hazard models with one agent is their ability to explain output sharing between the Principal and the agent even when both are risk-neutral. By contrast, non-trivial output sharing in the standard Principal-agent model can be obtained only by introducing further complications, such as agents' risk-aversion, that often make it difficult to compute the optimal contract explicitly.

Our results point out that strong qualifications must be imposed when a model with two-sided moral hazard is intended for applications in multiple agent settings such as franchising and product quality. Proposition 2 shows that the value V_1 and the contract supporting it will be suboptimal when N is large, and that the optimal scheme will approximately yield an average value of V^+ that involves no sharing at all.

To take an example, consider a finite-action version of the two-sided moral hazard model of product

warranties proposed by Cooper and Ross [5]. If the intended application is a manufacturer selling a large number of units of the product, then our results point out that the optimal contract for the one agent case is not robust to increases in team size, and that a rather simple incentive scheme can achieve a nearly efficient outcome. This conclusion is similar to the one derived in a dynamic model of product warranties (Al-Najjar [1]). In that model, the driving force is the firm's reputation rather than explicit contracts in a team of large size. It is interesting to note that explicit contracting in a static problem with many agents has reputation-like implications.¹⁵

4.7. Non-additive Cost and Uniform Actions across Agents

In many situations with two-sided moral hazard and multiple agents, some of the Principal's actions are agent-specific while others are by nature uniform across agents. Thus, in the manufacturer-consumer product quality example discussed earlier, the quality of material and after-sale services are consumer-specific while a more careful design of the product that improves its safety represents an action that is uniform across agents. Similar examples in the contexts of employment and franchising relationships can be easily found.

In our basic model the assumption of additive cost structure was maintained for analytical and notational simplicity. The model can be easily modified to take into account broader definitions of cost structures, including the case where some of the Principal's actions are agent-specific while others are uniform across agents.

To see this, consider the N -agent model, and recall our identification of an action a with the action distribution in which the Principal takes action a with every agent (see Section 3.1). In this case, $c(a)$ represents the average cost of the Principal when he takes action a with all agents.¹⁶ An α which is not a vertex of Δ_A represents a situation in which the Principal uses different actions with different agents with frequencies given by α . The average cost function c is additive if $c(\alpha)$ is just the average of the costs of the underlying pure actions weighted by the components of α .

More generally, c may be *concave*, reflecting technological advantages to taking uniform actions across agents. As an illustration, in the quality example of Section 3.3, $c(\alpha)$ represents the average cost

when the percentage of low quality products is α . Suppose that $c(a^h) = 2$, $c(a^l) = 1$ and $N = 100$. If c is linear and the firm chooses a mixture of qualities $\alpha = 0.03$ so 3 cars are of low quality and 97 are of high quality, then the average cost of a car is $c(\alpha) = 1(0.03) + 2(0.97) = 1.97$. On the other hand, a strictly concave c implies $c(\alpha) > 1.97$, so while the firm can vary the quality of each car, producing a mixture of qualities is relatively more costly compared with averaging over uniform quality outputs.

Actions like advertising, product design and research and development which, by their nature, have to be chosen uniformly across agents can be viewed as ones for which c sharply increases as we move away from the vertices of Δ_A , so it is prohibitively costly to choose different actions with different agents. In this case, our results go through and will, in fact, be much easier to prove. The reason is that the only deviations we need to be concerned with are those in which the Principal switches from taking a^+ with all agents to taking some other action a also with all agents. Such deviations generate massive evidence that makes detection all but certain.

Appropriate versions of our results on economies of scale and the benefits to conditioning an agent's compensation on other agents' outcomes also hold for the intermediate cases in which c is (mildly) strictly concave. The analysis would be notationally more complex, however, because it is no longer the case that the optimal mixture α^+ in the one-agent problem remains optimal in the N -agent problem.

5. PROOFS

The main contribution of the following lemma is ensuring that \bar{N} can be chosen so that it works uniformly for all $\alpha \in \Delta_A$.

Lemma 5.1: *For any $\epsilon > 0$ there is N such that for every $\alpha \in \Delta_A$ and $N > N$:*

$$P\{\bar{y} : |\gamma(\bar{y}) - \pi(\alpha, b^+)| < \epsilon \mid \alpha, \bar{b}^+\} > 1 - \epsilon.$$

Before proving the lemma, we need additional notation. Suppose that the set of N agents is divided into L subpopulations with N_l agents in the l th population. Given a vector of outcomes \bar{y} , let $\#y_m^l$ denote the number of times outcome y_m appears in the l th subpopulation. Let $\gamma_m^l(\bar{y})$ denote the empirical frequency of the m th outcome in the l th subpopulation, and let $\gamma^l(\bar{y})$ denote the vector of such frequencies. Recall that $\gamma_m(\bar{y})$ and $\gamma(\bar{y})$ denote the corresponding frequencies for the entire population. Then $\gamma_m(\bar{y}) = \frac{1}{N} \sum_l \#y_m^l = \sum_l \frac{N_l}{N} \frac{\#y_m^l}{N_l}$, so we have $\gamma(\bar{y}) = \sum_l \alpha_l \gamma^l(\bar{y})$. It should be noted that $\gamma^l(\bar{y})$ depends only on the portion of \bar{y}^l in which a_l was played. This redundancy in the definition of $\gamma^l(\bar{y})$ simplifies the notation.

Proof: Fix \bar{a} that is consistent with α . Let N_l be the number of times a_l is played in \bar{a} and note that $\alpha = \frac{1}{N}(N_1, \dots, N_L)$. To simplify notation, write $\pi^l = \pi(a_l, b^+)$ and $\alpha_l = \alpha(a_l)$. Note also that $\pi(\alpha, b^+) = \sum_l \alpha_l \pi^l$ by definition, so that

$$\begin{aligned} |\pi(\alpha, b^+) - \gamma(\bar{y})| &= \left| \sum_l \alpha_l \pi^l - \sum_l \alpha_l \gamma^l(\bar{y}) \right| \\ &= \left| \sum_l \alpha_l (\pi^l - \gamma^l(\bar{y})) \right| \\ &\leq \sum_l \alpha_l |\pi^l - \gamma^l(\bar{y})|. \end{aligned}$$

Using Chebyshev's inequality, the formula for the variance of binomials, and recalling that there are M possible outcomes, we have

$$P\left\{|\gamma^l(\bar{y}) - \pi^l| > \frac{\epsilon}{L} \mid a_l, b^+\right\} < \frac{M}{4N_l(\frac{\epsilon}{L})^2}.$$

Choosing $N \geq \tilde{N} > \frac{2L}{\epsilon} \frac{ML^3}{2\epsilon^3}$ and noting that $N_l = \alpha_l N$, we conclude that either:

- 1) $\alpha_l < \frac{\epsilon}{2L}$; or
- 2) $N_l > \frac{ML^3}{2\epsilon^3}$, in which case $P\{|\gamma^l(\bar{y}) - \pi^l| > \frac{\epsilon}{L} \mid a_l, b^+\} < \frac{\epsilon}{2L}$ for every l .

Let L_1 and L_2 denote the subsets of indices l for which conditions 1 and 2 above are satisfied. Then, for any values $L_1, L_2 \leq L$, we have

$$\sum_{l \in L_1} \alpha_l |\pi^l - \gamma^l(\bar{y})| < L \frac{\epsilon}{2L} = \frac{\epsilon}{2},$$

and

$$P\left\{\bar{y} : \sum_{l \in L_2} \alpha_l |\pi(\alpha, b^+) - \gamma(\bar{y})| > \epsilon\right\} \leq \sum_{l \in L_2} P\{\bar{y} : |\gamma^l(\bar{y}) - \pi^l| > \frac{\epsilon}{L}\} < L \frac{\epsilon}{2L} = \frac{\epsilon}{2}.$$

We may therefore conclude that

$$P\left\{\bar{y} : \sum_l \alpha_l |\pi(\alpha, b^+) - \gamma(\bar{y})| > \epsilon\right\} < \epsilon.$$

The result now follows by noting that

$$P\{\bar{y} : |\gamma(\bar{y}) - \pi(\alpha, b^+)| < \epsilon \mid \alpha, \bar{b}^+\} = \sum_{\bar{a}} P\{\bar{y} : |\gamma(\bar{y}) - \pi(\alpha, b^+)| < \epsilon \mid \bar{a}, \bar{b}^+\} \alpha(\bar{a}).$$

Q.E.D.

Propositions 1 and 3 are immediate consequences of the proof of Proposition 2.

Proof of Proposition 2: By A2, we can find a contract s_ϵ^+ such that $|s_\epsilon^+ - s^+| < \epsilon$, and such that b^+ is the agent's unique best response to (α^+, s_ϵ^+) . Let $U' \subset \Delta_A$ be the subset consisting of all α such that b^+ continues to be the agent's unique best response to (α, s_ϵ^+) . Note that for small N , we may have $U' \cap \Delta_A^* = \{\alpha^+\}$. The set U' is open in Δ_A and contains α^+ (because b^+ is a unique best response to (α^+, s_ϵ^+) , B is finite, and $U(s_\epsilon^+, \alpha, b)$ is continuous in α for all b). Thus b^+ is also the agent's unique best response to (μ, s_ϵ^+) for any exchangeable mixed strategy μ with support contained in U' . Let V_ϵ^+

denote the Principal's payoff under the contract s_t^+ and assuming that actions α^+ and b^+ were taken. Note that $V^+ - V_\epsilon^+ < \epsilon$.

Define $\gamma^+ = \pi(\alpha^+, b^+)$ and for any $\delta > 0$, let $B_\delta = \{\gamma \in \Delta_Y : |\gamma - \gamma^+| < \delta\}$. Since U' is open, we may choose δ small enough that $U_\delta = \{\alpha : \pi(\alpha, b^+) \in B_\delta\} \subset U'$. Finally, define T by $T(\vec{y}) = 0$ if $\gamma(\vec{y}) \notin B_{\delta/2}$, and $T(\vec{y}) = 1$ otherwise.

Abusing notation, I will use b^+ to denote the profile in which every agent plays b^+ . I now show that there is μ such that (μ, b^+) is an equilibrium of $\Gamma_N^*(\vec{s}^{s^c})$. To ensure that b^+ is best response, it is enough to find a μ with $\text{supp}(\mu) \subset U_\delta \subset U'$.

By Lemma 5.1, we can find N large enough that for any $\alpha \notin U_\delta$ and $N \geq \bar{N}$

$$P\{T(\vec{y}) = 0 \mid \alpha, \vec{b}^+\} > 1 - \epsilon$$

This means that for any such α , $V(\vec{s}^{s^c}, \alpha, b^+) < \epsilon y^M$. On the other hand, we can also use the lemma to ensure that $V(\vec{s}^{s^c}, \alpha^+, b^+) < (1 - \epsilon)V_\epsilon^+$. Choosing ϵ small enough implies that no best response to b^+ in $\Gamma_N^*(\vec{s}^{s^c})$ assigns positive probability to an $\alpha \notin U_\delta$. Finally, recall that the set of admissible action distributions Δ^* (relative to N) is finite, so b^+ must have at least one mixed best response in $\Gamma_N^*(\vec{s}^{s^c})$. We can now use Proposition 4 below to conclude that this is also a best response to b^+ in $\Gamma_N(\vec{s}^{s^c})$.

Finally, the various claims in the conclusion of the proposition follow directly from the construction above.

Q.E.D.

PROPOSITION 4: *The equilibrium (μ, ν) for the game $\Gamma_N^*(\vec{s}^{s^c})$ constructed in the proof of Proposition 3 above is also an equilibrium for $\Gamma_N(\vec{s}^{s^c})$.*

Proof: We must show that, in the game $\Gamma_N(\vec{s}^{s^c})$, we have $V(\vec{s}^{s^c}, \vec{a}, \vec{b}^+) \leq V(\vec{s}^{s^c}, \mu, \vec{b}^+)$ for every \vec{a} . Suppose that \vec{a} and \vec{a}' are two actions whose frequency distribution is some $\alpha \in \Delta_A$. Clearly, the symmetry of the agency problems implies $V(\vec{s}^{s^c}, \vec{a}, \vec{b}^+) = V(\vec{s}^{s^c}, \vec{a}', \vec{b}^+)$. Thus, viewing α in $\Gamma_N(\vec{s}^{s^c})$ as a uniform measure on the set of action vectors with frequency distribution α , we find

that $V(\vec{s}^{SC}, \vec{a}, \vec{b}^+) = V(\vec{s}^{SC}, \alpha, \vec{b}^+)$. But since μ was a best response to \vec{b}^+ in $\Gamma_N^*(\vec{s}^{SC})$, we must have $V(\vec{s}^{SC}, \alpha, \vec{b}^+) \leq V(\vec{s}^{SC}, \mu, \vec{b}^+)$.

Q.E.D.

ENDNOTES

- 1- On partnership games see, for example, Legros and Matsushima [10].
- 2- The complexity of the spaces of outcomes and incentives schemes grows exponentially with N . Thus, with M possible outcomes in each agency problem, there are M^N vectors of team outcomes and the space of possible incentives schemes for a single agent is a subset of a linear space of dimension M^N . For example, in a medium size franchise network of 50 outlets and 10 possible outcomes per outlet, there will be 10^{50} team outcomes that must be taken into account in designing the optimal contract for an individual agent.
- 3- This mistake is made in a recent paper by Bhattacharyya and Lafontaine [4, Section 4]. They calculate the optimal two-sided moral hazard contract when there is a single agent and assert that the optimal N -agent contract consists of a simple N -replica of the optimal single-agent contract. Proposition 2 below points out that the naive N -replica contract is strictly dominated by an extremely *simple* incentive scheme based on a binary test statistic. The proposition allows for arbitrarily large number of actions and outcomes and makes minimal distributional assumptions on the way actions are mapped to outcomes. The suboptimality of the naive contracts is particular apparent in the special case where the Principal must take a uniform action across agents (see Section 4.7 below). Of course, one can simply *force* the N -replica contract to be optimal by assuming that it is impossible to write or enforce a contract in which an agent's compensation may depend on an easy to compute binary test statistic. On the other hand, this is an assumption which would have to be incorporated explicitly in the model and its use be defended on a case-by-case basis.
- 4- All we need here is that the Principal's payoff is bounded away from his average payoff in the single agent case. This is much weaker than the claim that the Principle can approximately get V^+ (and generally requires a smaller N). The proofs of the results allow, in principle, to compute tight bounds on N , and in simple settings (*e.g.*, two-action/two-outcome cases) the N needed appears to be relatively small.
- 5- Since there is no common uncertainty in the agents' performance, Holmstrom's [9] analysis implies that conditioning agent n 's compensation on the entire vector of outputs cannot be justified on the grounds that this additional information can be used to improve the agents' incentives through relative performance evaluation.

- 6- To find a counterpart for the agents in Radner's model, one should modify his model so that one of the two players is myopic (short-run player).
- 7- For example, our result that the Principal gets an average payoff close to his commitment payoff is in a similar spirit as Fudenberg and Levine's [6] results for reputation games where a long-run player deals with a sequence of short-run players (see also Section 4.6 below). A recent paper (Al-Najjar and Ye [3]) explores a dynamic model which can be specialized to one with a Principal interacting with a team of agents repeatedly. Repeated play creates two new effects (absent in the static model of the present paper): explicit reputation effects on the part of the Principal, and the potential for agents to collude against the Principal.
- 8- We will think of the actions as potentially representing vectors of levels of various attributes of the efforts or the investments made by the Principal and the agent, as in the case of multi-task agency problems. This should be contrasted with the more common (and more restrictive) assumption that effort is one-dimensional variable. Subject to some caveats, in some cases one may interpret actions as contingent plans in a multi-stage setting.
- 9- To simplify the exposition, I will (somehow loosely) ignore the Principal's participation constraint because A can be redefined in such a way as to enable him to opt out of the relationship.
- 10- Since standard Principal-Agent setting is not the main focus in this paper, it is more convenient to state this assumption in this form rather than in terms of the primitives of the agency problem (*i.e.*, π , g , c and u).
- 11- Note the assumption implicit in the definition of the optimal value V_X that if $\Gamma(s)$ has multiple equilibria, the Principal can choose his most preferred equilibrium. This is a weak implementation criterion which is nevertheless standard in the literature. Stronger implementation criteria (*e.g.*, requiring the result to hold for any choice of a Nash equilibrium of $\Gamma(s)$) are possible under stronger restrictions on the primitives of the game. See Mookherjee [12] and Ma [11] for a discussion of this problem, and Fudenberg and Levine [6] for the sort of conditions needed to ensure that deviations can be statistically identified. Mookherjee [12] and Ma [11] also pointed out, in the context of other multiple agent settings, that weak Nash implementation might require agents to play weakly dominated strategies. It is worth mentioning that the Nash equilibrium constructed in the proof of

Proposition 2 below has the property that agents play strict best responses, so our framework does not suffer from this problem.

- 12- To see this, fix $\alpha = (\alpha^1, \dots, \alpha^M) \in \Delta_A^*$ and let k^l be the number of times with which action a^l must be chosen in order to ensure that the action distribution is equal to α . There are $\frac{N!}{k^1! \dots k^L!}$ action vectors which are consistent with α , and $\frac{(N-1)!}{k^1! \dots (k^l-1)! \dots k^L!}$ action vectors in which a^l is chosen against agent n . A simple calculation shows that the probability that action a^l is played against agent n is equal to $k^l/N = \alpha^l$. This shows that the marginal distribution for agent n is α for each n .
- 13- Note that the result is in terms of the *average* payoff per transaction rather than in terms of total payoff. There is a number of questions for which the average payoff may be the more appropriate measure of the Principal's performance. For example, restrictions on average payoffs is enough to deal with the issues of informational economies of scale, optimality of conditioning compensation on the agent's own outcome only (see footnote 3), and approximate efficiency. More crucially, restrictions on average payoffs is all that is relevant in evaluating the advantages of a centralized structure of organizing transactions compared to a more fragmented structure with N independent Principals.
- 14- See footnote 5 above.
- 15- See also Al-Najjar [2] for a related model which combines reputation with increases in firm size to show the existence of informational economies of scale to reputation building.
- 16- Note that in our definition of $c : \Delta_A \rightarrow \mathbb{R}$ there is an implicit assumption of constant returns in the sense that $c(a)$ remains constant as the number of agents N increases.

REFERENCES

1. Al-Najjar, N. I., Reputation, Product Quality and Warranties, *Journal of Economics & Management Strategy* **3** (1995), 605-37.
2. Al-Najjar, N. I., "On the Role of Reputation in Aggregating Information," mimeo, 1993.
3. Al-Najjar, N. I. and T. Ye, "Collusion and Reputation in Games with a Large Player," MEDS Department, Kellogg GSM, Northwestern University, 1996.
4. Bhattacharyya, S. and F. Lafontaine, Double-Sided Moral Hazard and the Nature of Share Contracts, *The RAND Journal of Economics* **26** (1995), 761-81.
5. Cooper, R. and T. W. Ross, Product Warranties and Double Moral Hazard, *The RAND Journal of Economics* **16** (1985), 103-13.
6. Fudenberg, D. and D. Levine, Maintaining Reputation when Strategies are Imperfectly Observed, *Review of Economic Studies* **59** (1992), 561-79.
7. Green, E., Classical Statistics as a Theory of Incentives, *Games and Economic Behavior* **2** (1990), 13-28.
8. Grossman, S. and O. Hart, An Analysis of the Principal-Agent Problem, *Econometrica* **51** (1983), 7-45.
9. Holmstrom, B., Moral Hazard in Teams, *The Bell Journal of Economics* **13** (1982), 324-40.
10. Legros, P. and H. Matsushima, Efficiency in Partnerships, *Journal of Economic Theory* **55** (1991), 296-322.
11. Ma, C.-T., Unique Implementation of Incentive Contracts with Many Agents, *Review of Economic Studies* **55** (1988), 555-71.

12. Mookherjee, D., Optimal Incentive Schemes with Many Agents, *Review of Economic Studies* **21** (1984), 433-46.
13. Radner, R. Repeated Partnership Games with Imperfect Monitoring and no Discounting, *Review of Economic Studies* **LIII** (1986), 43-57.