Moral Hazard in Risk-Averse Teams

Eric Rasmusen

The RAND Journal of Economics, Volume 18, Issue 3 (Autumn, 1987), 428-435.

Stable URL:
http://links.jstor.org/sici?sici=0741-6261%28198723%2918%3A3%3C428%3AMHIRT%3E2.0.CO%3B2-G

Moral hazard in risk-averse teams

Eric Rasmusen*

Holmström (1982) has shown that a non-budget-balancing contract induces a team of risk-neutral agents to choose the first-best effort levels. This is not generally true when agents are risk averse. Furthermore, a "massacre" contract, which punishes all but one agent when the outcome is low, can attain the first best over a wider range of parameters than any other budget-balancing contract.

"No, they have no railroad accidents to speak of in France. But why? Because when one occurs, *somebody* has to hang for it! Not hang, maybe, but be punished at least with such vigor of emphasis as to make negligence a thing to be shuddered at by railroad officials for many a day thereafter. 'No blame attached to the officers'—that lying and disaster-breeding verdict so common to our soft-hearted juries, is seldom rendered in France."

—Mark Twain, *The Innocents Abroad*, chap. 12

1. Introduction

■ Holmström's 1982 article, "Moral Hazard in Teams," begins with a theorem that shows that no contract that is "budget-balancing," allocating all of the team's output to the members of the team, can induce the members to choose the efficient effort levels. Instead, the efficient output is obtained by a contract giving each member of the team a payoff of zero if output is lower than if all members had chosen the efficient effort levels. Such a contract is not entirely satisfactory from a modelling point of view because it requires a commitment to discard an output that is even slightly below the efficient level. Once a low output is observed, all of the team members would like to repudiate their contract and share the output among themselves. To commit to losing it, they must introduce an outsider to whom they agree to surrender all of the output whenever it is insufficient.

Holmström's theorem, however, depends on the agents' utility functions' being linear in money. If agents are risk averse, they can use their risk aversion to write the efficient budget-balancing contract that I shall describe below. Holmström is not in error, but readers of his article, accustomed to models in which risk aversion prevents, rather than permits, the first best to be obtained, might be misled. Fortunately, the validity of the results in Holmström's later sections is unaffected, since by relaxing the requirement that all contracts be budget-balancing, he does not exclude such contracts from his later propositions.

I shall also show that of the many efficient budget-balancing contracts, the "massacre" contract, in which one randomly selected agent benefits and all the others are punished when an out-of-equilibrium outcome is observed, is feasible more often than any other

contract. In particular, the massacre contract is feasible for a larger set of parameters than the "scapegoat" contract, in which one agent is punished and the others benefit.

Section 2 describes the model and Holmström's original result. Section 3 constructs a budget-balancing contract of the scapegoat type and shows when it can induce the first-best effort level. This is followed by general comments on budget-balancing. Section 4 compares different kinds of budget-balancing contracts and shows that the massacre budget-balancing contract is most often feasible. Concluding remarks appear in Section 5.

2. The Holmström model

■ I shall use Holmström's notation. Each of n agents indexed by i takes an unobservable action or effort level $a_i \in [0, \infty)$. Write

$$a_{-i} = (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n) \quad \text{and} \quad a = (a_i, a_{-i}).$$

The value of output is a function $x(a)$, which is strictly increasing, concave, and continuously differentiable. Output depends only on the effort levels; there is no random error.

The compensation rule specifies $s_i(x)$ as agent i's compensation if the output is x. Effort is not observable, so that compensation cannot be a direct function of a_i. Moreover, agent i has limited liability, and his compensation cannot be less than the liability bound of $s_i = -\omega_i$, where we assume that $\omega_i \geq 0$. Agent i's utility function is separable in money and effort and can be written in the form

$$u_i(s_i, a_i) = m_i(s_i) - v_i(a_i). \tag{1}$$

The disutility of effort v_i is continuously differentiable, strictly convex, and increasing. In Holmström's article utility is linear in money, so that $m_i = s_i$. Here I assume that m_i is given by the function

$$m_i(s_i) = -e^{-\theta_i s_i}, \tag{2}$$

which has constant absolute risk aversion equal to $\theta_i > 0$ for agent i. The particular form of function (2) is not important except for its convenient parameterization and its differentiability.

We seek sharing rules $s_i(x) \geq -\omega_i$ such that we have budget balancing,

$$\sum_{i=1}^{n} s_i(x) = x, \quad \text{for all } x, \tag{3}$$

and the noncooperative game with payoffs

$$m_i(s_i(x(a))) - v_i(a_i), \quad i = 1, \ldots, n, \tag{4}$$

has a Nash equilibrium a^* that satisfies the following condition for Pareto optimality.[1]

Condition 1. There do not exist efforts $\tilde{a}$ and sharing rules $\tilde{s}$ such that: (a) for all i, $Eu_i(\tilde{s}_i, \tilde{a}_i) \geq Eu_i(s_i, a_i^*)$; and (b) for some agent j, $Eu_j(\tilde{s}_j, \tilde{a}_j) > Eu_j(s_j, a_j^*)$.

It may seem unnecessary to write Condition 1 in terms of expected utility, since there is no production uncertainty in the model, but doing so allows for the possibility of randomized sharing rules.

Having assumed linearity of the utility functions, Holmström can simplify the payoffs to

$$s_i(x(a)) - v_i(a_i), \quad i = 1, \ldots, n, \tag{4a}$$

which validates the following proposition.

---

[1] Holmström uses the linearity of $m_i(s_i)$ in his model to state the Pareto-optimality condition more simply than Condition 1.

Proposition 1 (Holmström). There do not exist sharing rules $\{s_i(x)\}$ that satisfy (3) and yield a^* as a Nash equilibrium in the noncooperative game with payoffs (4a).

I shall not repeat Holmström's proof here, but the result is intuitive.[2] If one agent shirks, he receives the full benefit of his diminished effort. The cost, on the other hand, which is lower output, is shared by all the agents.

We shall see that when the payoffs are not (4a), but (4), Proposition 1 is no longer true. When agents are risk averse, an efficient budget-balancing contract exists if the agents are either sufficiently wealthy or sufficiently risk averse.

3. An efficient budget-balancing contract

■ Under the following contract, a^* is a Nash equilibrium for some values of the parameters ω and θ. If output is $x(a^*)$, let each agent i receive a share b_i such that the budget is balanced and Condition 1 is satisfied. If output is greater than $x(a^*)$, split the surplus evenly among the agents after giving each agent i the amount b_i. If output is less than $x(a^*)$, choose one agent j and let him receive $-\omega_j$. Let each of the remaining $(n-1)$ agents i receive $b_i + (b_j + \omega_j - x(a^*) + x)/(n-1)$. Depending on the unlucky agent j's wealth and whether he is paid more than his marginal product in equilibrium, $(b_j + \omega_j - x(a^*) + x)$ is greater or less than zero, and the lucky agents are paid more or less than they would have been had no one shirked.

The sum of the rewards when output is below the Pareto-optimal level and agent j is punished is

$$\sum_{k=1}^{n} s_k = -\omega_j + (\sum_{i \neq j} b_i) + (n-1)\left(\frac{b_j + \omega_j - x(a^*) + x}{(n-1)}\right) = x, \tag{5}$$

so that the contract is budget-balancing.

To a single agent i, who expects all of the other agents to choose the efficient effort level, the contract appears as contract (6):

$$s_i(x) = \begin{cases} b_i + (x - x(a^*))/n & & \text{if} & x \geq x(a^*), \\ b_i + z_i & \text{with probability } (n-1)/n & \text{if} & x < x(a^*), \\ -\omega_i & \text{with probability } 1/n & \text{if} & x < x(a^*), \end{cases} \tag{6}$$

where z_i is a random variable taking the value $(b_j + \omega_j - x(a^*) + x)/(n-1)$ with probability $1/(n-1)$ for $j = 1, \ldots, n, j \neq i$. The agent chooses either the Pareto-optimal effort level and the reward b_i, or some lower effort level and a gamble in which with probability $1/n$ he receives $-\omega_i$ and with probability $(n-1)/n$ he receives not only his own b_i but also an amount depending on the unlucky agent's wealth and equilibrium share.

Choose the b_i's in (6) so that $\sum_i b_i = x(a^*)$ and Condition 1 is satisfied when b_i replaces s_i.

Agent i does not want to exert an effort greater than a_i^*, given that the other agents are exerting the efficient level. Under the contract just described, increasing his effort beyond a_i^* raises every other agent's utility, and if it raises i's also, then a^* is not the Pareto-optimal effort. We shall implicitly carry this result through the article, and in demonstrating that an efficient Nash equilibrium is attained, we shall be concerned only about low effort levels.

With each agent i is associated a "cheating effort" $\hat{a}_i \in [0, a^*]$ that represents the deviation from equilibrium most tempting to him. Agent i's cheating effort maximizes his utility, given that the other agents choose a_{-i}^* and the contract is replaced by the "deviation lottery" characterized in (6). If agent i selects the lowest effort possible when he decides to cheat, then $\hat{a}_i = 0$, but he might choose a higher cheating effort because with probability

---

[2] Proposition 1 in this article is Holmström's Theorem 1 (1982, p. 339).

$(n - 1)/n$ his monetary compensation rises with effort and output. The cheating effort for contract (6) solves the problem,

$$\max_{a_i} \left[\frac{(n-1)}{n} Em_i(b_i + z_i) + \frac{1}{n} m_i(-\omega_i) - v_i(a_i) \right]. \tag{7}$$

The objective function in (7) is strictly concave in effort, because we have assumed that $m_i'' \leq 0$, $v_i'' > 0$, and $x'' < 0$. Given the concavity of problem (7), classical optimization tells us that the solution $\hat{a}_i$ exists and is unique.

To induce the agents to select the efficient effort levels, the utility difference Y_i between the efforts a_i^* and $\hat{a}_i$, under the Nash assumption that the other agents choose a_{-i}^*, must be positive for each agent i:

$$Y_i \equiv [m_i(b_i) - v_i(a_i^*)] - \left[\frac{(n-1)}{n} Em_i(b_i + z_i) + \frac{1}{n} m_i(-\omega_i) - v_i(\hat{a}_i) \right] > 0. \tag{8}$$

Proposition 2. If agents are risk averse, then provided that (a) punishments can be great enough (ω_i is large enough for every i) or (b) risk aversion is great enough (θ_i is large enough for every i), an efficient budget-balancing contract exists.

Proof. We need to show that inequality (8) is true under either of the two conditions. We start by holding the levels of risk aversion fixed and showing that if the ω's are large enough, then a^* is a Nash equilibrium under contract (6).

The first derivative of the utility difference Y_i with respect to the liability bound ω_i is (by invoking the envelope theorem to disregard the change in $\hat{a}_i$)

$$\frac{dY_i}{d\omega_i} = \frac{1}{n} m_i' > 0, \tag{9}$$

and the second derivative is

$$\frac{d^2Y}{d\omega_i^2} = -\frac{1}{n} m_i'' > 0. \tag{10}$$

Both expressions (9) and (10) are positive, because $m_i' > 0$ and $m_i'' < 0$ for a concave utility function. Hence, Y_i is increasing in ω_i and increasing at an increasing rate, so that it does not converge to an asymptotic value, and if ω_i is chosen large enough, inequality (8) is true. This applies to any agent i. For every θ_i there is some level of ω_i that allows a_i^* to be the equilibrium effort level, which proves the first part of Proposition 2.

Even if the agents' liabilities are limited, if they are sufficiently risk averse, an efficient budget-balancing contract can exist. Rewriting expression (8) by using the full form of the function m_i from equation (2) and substituting for z_i from contract (6), we obtain

$$Y_i = -e^{-\theta_i b_i} - v_i(a_i^*) + \frac{1}{n}(\sum_{j \neq i} e^{-\theta_i(b_i + 1/(n-1)[b_j + \omega_j - x(a^*) + x(\hat{a}_j, a_{-j}^*)])}) + \frac{1}{n}(e^{\theta_i \omega_i}) + v_i(\hat{a}_i). \tag{11}$$

As θ_i increases, the first and third terms on the right-hand side of (11) approach zero, the second is unaffected, and the fifth term may change as $\hat{a}_i$ changes, but it is bounded by $v_i(0)$ and $v_i(a_i^*)$. The fourth term of (11) increases exponentially with θ_i, which implies that Y_i can be made arbitrarily large. In particular, if θ_i is large enough, then Y_i is greater than zero, and (6) is an efficient budget-balancing contract. *Q.E.D.*

Contract (6), like Holmström's non-budget-balancing contract, does not necessarily yield a^* as the unique Nash equilibrium. An agent will not choose a_i^* if he expects another agent to choose an inefficiently low effort. Agent i's response might be either also to choose a low effort or to choose a high effort to compensate for the shirker and to avoid the random punishment, and there might exist other Nash equilibria in which some efforts are either insufficient or excessive. This does not mean that a^* is not a strong Nash equilibrium: a player's solitary deviation to any other effort level would lower his utility. Note also that

contract (6) is just one of the many possible contracts that rely on risk aversion to obtain the efficient equilibrium a^*, and efficient budget-balancing contracts can differ in their equilibrium outcomes as well as in the out-of-equilibrium punishments.

□ **General comments.** Under the contract in the preceding subsection, if the outcome shows that the wrong effort level was chosen, the agents are subjected to risk. *Ex post,* shirking makes only one of the agents worse off, and all the others better off, unless the shirking decreased output by more than the $b_i + \omega_i$ lost by the unlucky agent. We may find random punishments morally distressing, but they are similar to the punishments in tournament contracts under uncertainty, a context in which most people feel that punishments conditioned on random events are fair.

If the efforts of the agents are observable, even if only with error, the observed efforts can be used as the criteria for punishment. When the observation error has a high variance, the wrong agent would often be punished if output were low, but that does not detract from the contract's efficiency. Low observed output is merely an excuse to punish, although the shirker would have at least a slightly greater probability than the other agents of being punished. Depending on the variance of the observation error, the punishment is more or less random. This article describes a model in which the variance of the observation error is infinite, and hence the punishment seems completely capricious.

If efforts are observable with error, but we continue to assume that there is no production uncertainty, then no punishment occurs in equilibrium. Although the observed individual efforts might be used to allocate punishment, the only punishment trigger needed to attain efficiency is the team's output. In equilibrium output is $x(a^*)$, and if at the same time the observed efforts are low, it is clear they must have been observed with error.

If, however, individual efforts are observable with error, the first best may sometimes be achieved with budget balancing, even if the agents are risk neutral. The exact form of the contract depends on the specification of the observation error, but one possibility is to give the entire output to the agent whose observed effort is closest to his efficient effort. Such a plan resembles more a tournament contract than a team contract, and it relies heavily on risk neutrality since the equilibrium compensation (not just the out-of-equilibrium compensation) varies widely.

I mentioned earlier that one reason for dissatisfaction with non-budget-balancing contracts is the perfectness problem of having to discard output if it is insufficient. One might wonder whether the random punishment contract addresses the perfectness problem any better. Both contracts rely on lowering the utility of all the agents after certain outcomes are observed. The random punishment contract differs in that some agents may be better off *ex post,* after the punishment actually occurs, and they would vote against any recontracting. Between observing the output and choosing the agent to be punished, however, all agents could raise their utility by abandoning the punishment scheme and reverting to a nonstochastic sharing rule. The contract must somehow prevent this.

One way to prevent recontracting is to allocate the punishment immediately and automatically after low output is observed and before recontracting can occur.[3] Once a victim has been chosen, the other agents block any recontracting. In some situations there is a time lag between the date of choosing effort and the date of observing output. Committing to punishment is then simple: the lottery for allocating possible future punishments is held between the two dates. Any agent who objects to the lottery must have shirked, since in equilibrium the lottery is harmless as its punishment is never imposed.

Holmström has suggested that if the team adds a manager, whose sole function is to serve as the residual claimant, the perfectness problem for the non-budget-balancing contract

---

[3] One way to allocate shares, for example, would be to adopt a plan for allocating sealed-bid orders suggested, but not adopted, during the "Electric Conspiracy." This plan was to use stock market results published in the *Wall Street Journal* to allocate bids randomly, an interesting application of the theory of efficient markets (Sultan, 1974, p. 47).

can be solved. Output below the efficient level triggers a large payment to the manager, who therefore has the incentive to oppose any recontracting. This idea is particularly interesting because it depicts the manager not as a principal, but as an agent of the agents.[4] He is truly a "public servant." But this introduces the possibility of new agency problems: the manager has an interest in seeing that the output is below the efficient level so that the punishment is triggered. Randomization schemes are free from this danger, despite the drawback of their greater complexity.

4. Scapegoats versus massacres

■ Under the randomizing contract described above, one agent is chosen to be the "scapegoat" when output is low, and the others benefit at his expense. Another extreme in the class of randomizing contracts is the "massacre" contract in which all of the agents are punished when output is low except for one, randomly chosen, who receives the entire output. We shall see that the massacre is a better contract in the sense that it is feasible for a strictly larger set of liability bounds and risk-aversion parameters. Since it should be clear now that heterogeneity of the agents is unimportant except for the amount of notation, we shall simplify the model of Sections 2 and 3 by assuming that the agents have identical utility functions and liability bounds.

In the model with identical agents the scapegoat contract appears to the individual agent as

$$
s_i(x) = \begin{cases} x/n & \text{if} \quad x \ge x(a^*), \\ (x+\omega)/(n-1) & \text{with probability } (n-1)/n & \text{if} \quad x < x(a^*), \\ -\omega & \text{with probability } 1/n & \text{if} \quad x < x(a^*). \end{cases} \qquad (12)
$$

The effort-level vector a^* is a Nash equilibrium under the scapegoat contract if and only if the utility of choosing a_i^* is greater than that of choosing a lower effort level; that is, if

$$
\left[m\left(\frac{x(a^*)}{n}\right) - v(a_i^*) \right] - \left[\frac{(n-1)}{n} m\left(\frac{x_s+\omega}{n-1}\right) + \frac{1}{n} m(-\omega) - v(\hat{a}_s) \right] > 0, \qquad (13)
$$

where $\hat{a}_s$ is the effort level chosen under the scapegoat contract by the single cheating agent and $x_s \equiv x(\hat{a}_s, a^*_{-i})$.

The massacre contract appears to the individual agent as

$$
s_i(x) = \begin{cases} x/n & \text{if} \quad x \ge x(a^*), \\ -\omega & \text{with probability } (n-1)/n & \text{if} \quad x < x(a^*), \\ x+(n-1)\omega & \text{with probability } 1/n & \text{if} \quad x < x(a^*). \end{cases} \qquad (14)
$$

The vector a^* is a Nash equilibrium under the massacre contract if and only if the utility of choosing a_i^* is greater than that of choosing a lower effort level; that is, if

$$
\left[m\left(\frac{x(a^*)}{n}\right) - v(a_i^*) \right] - \left[\frac{(n-1)}{n} m(-\omega) + \frac{1}{n} m(x_m + [n-1]\omega) - v(\hat{a}_m) \right] > 0, \qquad (15)
$$

where $\hat{a}_m$ is the effort level chosen under the scapegoat contract by the single cheating agent and $x_m \equiv x(\hat{a}_m, a^*_{-i})$. The existence and uniqueness of $\hat{a}_m$ follow from the concavity of the deviation lottery (14) by the same argument made for the existence of $\hat{a}_s$ in Section 3.

We shall consider the class of contracts that treat identical agents identically as do the scapegoat and massacre contracts. Otherwise there is some agent whose change in expected utility from deviating is highest, and he is the only relevant agent for determining whether

---

[4] Anyone interested in the problem of one agent serving several principals should see Bernheim and Whinston (1986).

the first best can be supported. That agent's compensation in equilibrium should be increased to make his loss from deviation greater, or his punishment after deviation should be increased while lightening that of the other agents. Carried to the extreme, this results in a symmetric contract, which gives each agent the same compensation in equilibrium and the same expected disutility from deviation.

We shall now prove that the massacre contract attains efficiency for a larger set of parameters than the scapegoat contract or any other budget-balancing contract.

Proposition 3. If any budget-balancing contract can achieve a Pareto optimum, the massacre contract can.[5]

Proof. Under a Pareto-optimal contract, no agent has the incentive to choose an excessive effort level, so that output is no greater than the efficient level x^*. Each agent faces a choice between the desired effort level and a compensation of x^*/n, or the cheating effort level and a deviation lottery with an expected value of some smaller x/n and a distribution that depends on the form of the contract. The expected values of different lotteries are different because the cheating efforts are different, but that will not be relevant to this proof.

The advantage of the massacre contract is that its deviation lottery is the riskiest of any contract. Let us suppose for the moment that the cheating efforts, and thus the expected values, are the same for the massacre contract's deviation lottery and for some arbitrary contract k's deviation lottery. We shall see that if the expected values are the same, the massacre lottery can be obtained from contract k's deviation lottery by the addition of a series of mean-preserving spreads of the kind formalized by Rothschild and Stiglitz (1970). The massacre contract creates the two-point deviation lottery that places the highest possible probability, $(n - 1)/n$, on the lowest possible payoff, $-\omega$, and the probability of $1/n$ on the highest possible payoff, $(x + (n - 1)\omega)$. Any other lottery puts smaller probability on $-\omega$ and positive probability on payoffs less than $(x + (n - 1)\omega)$. If the means of the lotteries are equal, the massacre lottery is riskier than lottery k, because it takes probability mass away from the payoffs between $-\omega$ and $(x + (n - 1)\omega)$ and puts it on those extreme points.

More precisely, if contract k has a discrete probability function for its deviation lottery, then we can obtain the massacre lottery from k's deviation lottery by a series of mean-preserving spreads of the form

$$f(m) = \begin{cases} \alpha \geq 0 & \text{for} & m = -\omega, \\ -\alpha \leq 0 & \text{for} & m = -\omega + d, \\ -\beta \leq 0 & \text{for} & m = x + (n-1)\omega - t, \\ \beta \geq 0 & \text{for} & m = x + (n-1)\omega, \\ 0 & \text{otherwise,} \end{cases} \tag{16}$$

where the values of d, t, α, and β are chosen so that

$$-\omega \leq -\omega + d \leq x + (n-1)\omega - t \leq x + (n-1)\omega \tag{17}$$

and

$$\alpha d = \beta t.$$

If contract k's deviation lottery has a continuous density, the notation of (16) is inappropriate, but it can easily be adapted to find the desired mean-preserving spreads.

Rothschild and Stiglitz (1970) have shown that the expected utility of a risk-averse agent is lower with a lottery that is riskier in the sense of being obtained from other lotteries by a series of mean-preserving spreads. Since any deviation from a^* triggers the deviation lottery specified by the contract in force, an agent's expected utility is lower when a given

[5] I would like to thank Paul Milgrom for suggesting that I broaden Proposition 3 to its present form.

deviation occurs under a massacre contract than under any other contract. An agent who deviates to $\hat{a}_m$ will have lower utility under the massacre contract than under any other contract k.

Once we relax our temporary supposition that the cheating efforts are equal, contract k is even less able to deter deviation because the agent will reoptimize his cheating effort from $\hat{a}_m$ to a cheating effort preferred under contract k, and this raises his utility when he deviates under that contract. Although the massacre contract may not always be able to support a Pareto optimum by making the cost (the shift to a lottery) greater than the benefit (the reduced effort), it makes the difference as great as possible, and hence will support the Pareto optimum whenever any other feasible contract can. *Q.E.D.*

The intuition behind Proposition 3 is that the massacre contract puts greater risk on an agent if he chooses low effort, and less risk-averse agents can be deterred by this larger amount of risk. The massacre contract deters shirking in types of agents who would shirk under the scapegoat contract, and hence we might expect to see the massacre contract employed more often. The advantage of the massacre contract, however, also points up a major weakness of the model: the absence of production uncertainty.[6] If output were sometimes low, even if effort were high, the massacre contract would trigger greater accidental punishments than the scapegoat contract, and its superiority might well disappear. Moreover, I have said nothing about the output levels in the inefficient Nash equilibria that possibly exist even under efficient contracts, and the massacre contract may do badly in those inefficient equilibria.

5. Concluding remarks

■ We have seen that for the multiagent team, it is easier to find a first-best contract if agents are risk averse than if they are risk neutral. Although no budget-balancing first-best contract exists when agents are risk neutral, when they are risk averse such a contract does exist. The contract is similar to the non-budget-balancing contract that Holmström suggests for risk-neutral agents, because team output less than the efficient level triggers a punishment, but with risk-averse agents the punishment can take the form of a lottery rather than of the destruction of the output. If agents are sufficiently risk averse, or the lottery is sufficiently risky, each agent is unwilling to deviate from the efficient effort level, given that he believes his fellow agents are choosing the efficient level. The deviation lottery can take a number of forms, including the scapegoat lottery, in which one agent is punished and the others take his share, and the massacre lottery, in which one agent is rewarded by being granted the shares of all the others. Under some parameter values any of the deviation lotteries performs equally well, but the massacre lottery is able to attain the first best for less risk-averse agents or more tightly bounded punishments.

References

BERNHEIM, B. AND WHINSTON, M. "Common Agency." *Econometrica,* Vol. 54 (1986), pp. 923–942.
HOLMSTRÖM, B. "Moral Hazard in Teams." *Bell Journal of Economics,* Vol. 13 (1982), pp. 324–340.
NALEBUFF, B. AND STIGLITZ, J. "Prizes and Incentives: Towards a General Theory of Compensation and Competition." *Bell Journal of Economics,* Vol. 14 (1983), pp. 21–43.
ROTHSCHILD, M. AND STIGLITZ, J. "Increasing Risk I: A Definition." *Journal of Economic Theory,* Vol. 2 (1970), pp. 225–243.
SULTAN, R. *Pricing in the Electrical Oligopoly,* Vol. 1. Cambridge: Harvard University Press, 1974.

---

[6] The massacre contract is comparable to rewarding the winner of a tournament and the scapegoat contract is comparable to punishing the loser. Nalebuff and Stiglitz (1983) find that in tournaments with many players and production uncertainty, a punishment for the loser is superior to a prize for the winner. Their explanation is based on the combination of production uncertainty and a large number of players in their model, and does not rely on risk aversion. Tournaments are different from teams, because in tournaments the punishment is actually inflicted, and its direct disutility must be balanced against the effort it induces.