

Risk, Incentives, and Moral Hazard

This note considers the economics of incentive contracts, and the fundamental tension between pecuniary motivation and risk sharing. Prior to class, please prepare your answers to the problem on pages 7-8.*

Consider the following three economic problems:

- An insurance company sells a policy that insures a factory against fire and wonders whether the factory owner, now that she has insurance, will take care not to leave oily rags lying around.
- A venture capitalist, considering whether to finance an entrepreneur, wonders whether the entrepreneur will work diligently to bring a new idea to market, or with will spend her new-found capital imprudently.
- A company that manufactures large capital equipment employs salespersons who deal with the firm's clients. The sales effort takes place outside the view of top management of the manufacturer, which wonders how much effort its salespersons put into selling the product.

These are examples of a problem called *moral hazard*. These are situations in which (1) the actions of one party affect the economic welfare of others, (2) the interests of the parties diverge to some extent, and (3) the actions of each party are not completely controllable by the other.

The solution to problems of moral hazard is, broadly speaking, incentives. This is the logic of pay-for-performance compensation in employment relationships. However, as the examples above indicate, the range of contexts in which moral hazards can be found is wide. This note explores how direct financial incentives can overcome moral hazard problems and improve business performance. The key economic consideration in these situations is how to address the fundamental trade-off between pecuniary motivation (i.e., monetary gain) and desirable risk sharing.

* Much of this material is based on the insightful notes and examples of David M. Kreps, to whom I am deeply indebted.

Theory and Concepts

In theory, there are two simple solutions to problems of incentives. *The first is to determine what choice of action is desirable and fix this contractually.* In the case of a factory owner, the insurance contract specifies that no insurance is paid if the owner does not store oily rags safely. A salesperson is provided no compensation if he does not exert himself. The entrepreneur loses control of her business to the venture capitalist if she does not take prudent decisions.

Does this work? Not really. The problem with each of these is that it may be impossible, *ex post*, to verify that the contractual terms were met. It might be possible to monitor compliance with a contractual provision on the storage of flammable items, with (say) with spot checks. But how can a salesperson's level of effort be monitored? How can the level of prudence of an entrepreneur be measured? And, even if these things can be monitored or measured, can this be done in a manner that a court could enforce? For example, a commercial banker who loans money to an entrepreneur may know that decisions were imprudent, but can a judge or jury in a civil trial be convinced? The problem, generally speaking, is that often the desired "inputs" cannot be adequately measured or monitored (or cannot be made part of a legally-enforceable contract).

In theory again, we still have a second simple solution. *Construct an arrangement that puts the onus entirely on the party choosing the action.* That is, do not offer any fire insurance to the factory owner, so she bears fully the consequences of not storing oily rags carefully. Have the salesperson bear fully the impact of his effort decisions, by giving him a payment that equals (on the margin) the full effect of whether he makes the sale. Have the entrepreneur bear completely the consequences of her choice, by making a loan that guarantees a fixed amount of repayment to financier.

Will this work? It might. *But:* It is generally not an efficient, nor profitable, strategy for the parties involved. The problem with these extreme solutions is that they imply *no risk sharing* between the parties involved. If the marginal profit from a sale is \$10 million and a sale is uncertain even if the salesperson tries his best, will he be willing to bear that much risk in his level of compensation? Will the entrepreneur be willing to bear all the financial risk of her venture? The problem is perhaps clearest in the insurance context: If the solution to moral hazard in an insurance context is not to offer any insurance, then the insurance business is dramatically curtailed.

In sum, there is a fundamental trade-off concerning incentives. One can view these sorts of situations as involving a *hidden action* that one party can take, but the other cannot observe and verify. For example, a salesperson's effort away from the office is (typically) hidden from view of her employer. But the salesperson's success also depends on things well beyond her control, such as the customer's willingness to pay, the quality of the product she must sell, and so on. To share risk efficiently typically means shielding the

salesperson choosing the hidden action (effort) from at least some consequences of that choice. Doing so, however, lowers the incentive to put in long hours, research the customers' needs, and generally exert a high sales effort. The trick is to balance these competing effects of *risk sharing* and *motivation* through the sensible design of incentives.

To understand these ideas and how we can resolve them, it is useful to consider a few detailed examples.

Example 1: *Insurance*

As an insurance underwriter, you have been asked to write a policy that insures a factory against loss by fire for a period of 1 year. If the factory has a fire, it will be a total loss of \$8 million. The owner of the factory is an expected utility maximizer, with (gross) utility function $U(X + 1 \text{ million})^{1/2}$ where X is the value of the factory at year's end; that is, $X = \$8$ million if there is no fire and $X = \$0$ if there is a fire. Your insurance company is risk neutral.

The chance of a fire depends on whether the owner of the factory takes due care. If she does not, the chance of a fire over the 1-year period is 0.05. If she does take due care, the chance of a fire over the 1-year period is 0.01. To take due care will raise the owner's cost of doing business by \$150,000 per year. That is, the factory owner's overall utility depends on both X and on her decision whether to take due care or not, according to

$$U(X, \text{no due care}) = (X + 1 \text{ million})^{1/2}$$

$$U(X, \text{due care}) = (X + 1 \text{ million} - 150K)^{1/2}.$$

We want to understand what the market will offer in the way of insurance to the factory owner. To do this, there are three scenarios we must consider.

Case 1: *No insurance.* Consider first what the factory owner would do without any insurance. In this case, her expected utility if she does not take care is

$$EU \text{ if no care: } EU = (.95)(8 \text{ million} + 1 \text{ million})^{1/2} + (.05)(1 \text{ million})^{1/2} = 2900$$

If she does take due care,

$$EU \text{ if takes care: } EU = (.99)(9 \text{ million} - 150K)^{1/2} + (.01)(1 \text{ million} - 150K)^{1/2} = 2954.4$$

So she will take care to reduce the likelihood of a fire, even though it will cost her \$150K to do so.

Case 2: *Full Insurance.* Now let's consider the case where the insurance company offers

to insure the building fully; that is, it will write a policy that pays the owner \$8 million in the event of a fire. Suppose that the insurance company cannot verify whether the owner ensures her workers to take care to never leave oily rags around, to always put flammables back in a fire-proof locker, and so on—that is, the insurer cannot verify whether or not “due care” is taken by the factory owner.

What happens if the insurance company offers full insurance? First, consider how much the factory owner would be willing to pay for full insurance. If the factory owner has insurance, X in her utility function is adjusted down by the up-front price she pays for insurance coverage, and X is adjusted up in the event of a fire by the compensation she receives (called “indemnification” in insurance lingo). So if she pays a price P for full insurance and does not take care, she will be fully compensated and have \$9 million – P for sure:

$$EU \text{ if no care:} \quad EU = 1 \cdot (9 \text{ million} - P)^{1/2}$$

If she does take care, she’ll have \$9 million – P , less the \$150K cost of taking care:

$$EU \text{ if takes care:} \quad EU = 1 \cdot (9 \text{ million} - P - 150K)^{1/2}$$

Because she faces no risk with full insurance, she gets this utility with probability 1 (thus the ‘1’ in front of each EU calculation).

Note something important here: *No matter what the price of full insurance coverage is, she is better off not taking due care to avoid a fire.* If she is fully insured, her utility if she takes due care is always less than her utility if she does not take due care. That much should be intuitive: When someone else is bearing all the risk, you have little incentive to take a costly action to reduce that risk. This is the fundamental tension between risk sharing and motivation when effort is not observable to everyone.

Now back to how much she will pay. With no insurance (case 1 above), we found that her best action is to take due care and that gives her $EU = 2954.4$. So she would pay for insurance coverage up to the point where it leaves her with less utility than her next best alternative—which is taking care with no insurance coverage. But with full insurance she has no incentive to take care, and utility $(9 \text{ million} - P)^{1/2}$ for sure. Solving

$$(9 \text{ million} - P)^{1/2} = 2954.4$$

for P gives $P = \$271,724$. This is the maximum she would pay for full insurance coverage, given her risk preferences and the cost of taking care to reduce the risk of a fire.

Will the market provide full insurance coverage to the factory owner? The insurance company will make its decision on whether or not to offer full insurance coverage knowing what we just discovered: Firms that are completely insured do not have a strong incentive to take

costly actions to reduce the risk of a fire. That moral hazard problem hurts the market's efficiency. Here's why.

The factory owner is willing to pay \$271,724 for full insurance. If she takes due care, the expected monetary loss from the risk of fire is $(.01)(\$8 \text{ million}) = \$80,000$. So she would be willing to pay a *risk premium* of \$191,724 ($= \$271,724 - \$80,000$) to avoid the risk of a fire. For a risk neutral insurance company, that's a pretty healthy premium over the expected claims loss—if the factory owner would still take care once insured.

But the insurer knows better. If it offers full insurance, the factory owner has a strong monetary incentive not to take due care. And if she does not take care, the probability of a fire becomes .05 and the expected monetary loss to the insurer is $(.05)(\$8 \text{ million}) = \$400,000$. So the insurance company will not provide full insurance for less than \$400,000. And, we already discovered, the factory owner is not willing to pay more than \$271,724 for full insurance coverage. In short, the market will not provide this insurance service.

This is a key lesson. When there is moral hazard, markets will not necessarily yield efficient outcomes. Since the factory owner is risk averse and the insurer is risk neutral, the (only) efficient outcome is for the two parties to transfer all risk to the insurer in exchange for an up front payment by the factory owner. But because the insurer does not control (nor can it verify for contract enforcement purposes) whether or not the factory owner took care, *and* because both parties know that the factory owner has a strong incentive to *not* take care if fully insured, the insurance company will not accept all the risk. *Moral hazard prevents markets from efficiently sharing risk, and value is lost.*

Case 3: Partial Insurance. How can this problem be fixed? One can balance the competing goals of risk sharing (which creates value for everyone) and motivation through an incentive contract with *partial* risk sharing. Let's suppose that the insurance company, being well versed in this problem, decides to offer a "co-insurance" policy to the factory owner. This policy will cost the factory owner \$75,000 up front. It will then compensate the owner for 50% of the \$8 million loss in the event of a fire. The owner bears the remaining \$4 million loss if the factory has a fire.

Will the factory owner accept this policy and still take due care? In the factory owner's utility function $U(X + 1 \text{ million})^{1/2}$, X now includes the value of the building, less the \$75,000 price of purchasing insurance, plus any payment received from the insurance company in the event of fire, less the \$150,000 cost of taking care if the factory owner takes due care.

Do the expected utility calculations as usual: If she does not take due care,

$$\begin{aligned} EU \text{ if no care: } \quad EU &= (.95)(8 \text{ million} + 1 \text{ million} - 75\text{K})^{1/2} \\ &+ (.05)(4 \text{ million} + 1 \text{ million} - 75\text{K})^{1/2} = 2949.1 \end{aligned}$$

If she does take due care, however,

$$\begin{aligned} \text{EU if takes care: } \text{EU} &= (.99)(8 \text{ million} + 1 \text{ million} - 75\text{K} - 150\text{K})^{1/2} \\ &+ (.01)(4 \text{ million} + 1 \text{ million} - 75\text{K} - 150\text{K})^{1/2} = 2954.5 \end{aligned}$$

(NB: Make sure you understand where all the numbers in the parentheses come from).

This incentive contract clears two hurdles: First, the factory owner would prefer to incur the \$150K cost of taking care in order to reduce the risk of a fire from 5% to 1%, given that she retains \$4 million of the loss if the factory has a fire. This comes from comparing the two EU values above. Second, her EU if she takes care and buys this partial insurance contract is *higher* than her EU if she takes care and does not buy any insurance at all (Case 1, which had EU = 2954.4). It is a near thing, but she'll prefer to buy partial insurance over no insurance on these terms, given her risk preferences.

Will the insurance company finally make a profit with this partial-insurance contract? Let's check. Knowing that the co-insurance contract provides a sufficient incentive to take care, the insurer will evaluate the risk of a fire as a 1% chance of a loss of \$4,000,000 to the insurer. Their gain is the payment of \$75,000 for certain. So the insurance company has an expected profit under this contract of $\text{EMV} = \$75,000 - (.01)(4 \text{ million}) = \$35,000$. Since it is risk neutral, it will happily take on its share of the fire risk to receive this positive expected payoff. The market works again.

In essence, the partial insurance contract reduces the level of risk sharing in this market, but by doing so provides a sufficient incentive for the factory owner to take care to minimize these risks. These sorts of more complex incentive arrangements are the economic solution to providing people with the right motivation to undertake costly efforts, when those efforts can't be observed directly.

About this problem. Not to belabor things, but let's be sure about the economic lesson in this example. Suppose an insurance company insures a business owner against fire. The chance of a fire depends on how careful the business owner is, although even the most careful individual can face a risk of fire. Efficient risk sharing, assuming the insurance company is risk neutral and the business owner is risk averse, means insuring the business owner entirely against any loss. That is, the business owner buys full insurance coverage, the insurance company bears entirely the risk of a fire, and the business owner is completely compensated for any loss. But, if taking care is costly for the business owner and she is entirely insured against loss, she has little incentive to take care (and, in extreme cases, might have an incentive for arson). This is where the term *moral hazard* originated. The insurance company, to motivate the business owner to take care, must "reward" her if there is no fire; that is, she must be better off if there is no fire than if there is. Typically, an insurer does not leave this to the business owner's innate moral unrighteousness, but offers only partial insurance coverage (insurance with a substantial deductible). The

prospect of losing the deductible if there is a fire motivates the business owner to take care. But it necessarily means less than efficient risk sharing.

Problem to Prepare for Class

Please be prepared to offer your best answers to this problem in class. As usual, you do not need to turn in your work on class-session problems, and you may work with others if you wish. I recommend you use a spreadsheet, although parts A & B can be done by hand.

Structuring Incentive Financing

An entrepreneur has identified a risky venture that requires an up-front capital investment of \$1 million to get going. The venture will have one of three possible outcomes. If it is a flop, it will make \$0, in which case the entire investment of \$1 million will be lost. If it is modestly successful, it will make \$5 million. And if it is wildly successful, it will make \$100 million. These venture payout figures do not include the up-front \$1 million investment.

The level of success depends on the entrepreneur's effort. If she exerts a low level of effort, the chance of the \$100 million outcome is 2%, the chance of the \$5 million outcome is 5%, and there is a 93% chance it will be a flop. However, if she exerts a high level of effort, the chance of the \$100 million outcome is 10%, the chance of the \$5 million outcome is 12%, and there is a 78% chance the venture will be a flop. The entrepreneur is risk averse, with a utility function over wealth of

$$U(W) = (W)^{1/2}$$

The entrepreneur suffers an additional disutility of 500 (utils) if she exerts a high level of effort. The disutility from exerting low effort is zero.

The entrepreneur cannot undertake the \$1 million initial investment herself, as she has initial wealth of only \$200,000. However, a risk-neutral venture capitalist is prepared to support this venture.

A. A Share Financing Arrangement

In practice, venture capitalists typically structure their deals to take a *share* of the total earnings of the venture, however successful it may turn out to be. Suppose the venture capitalist will pay the entrepreneur a base amount of \$900,000 up front, in return for which the venture capitalist will retain 67% of the gross proceeds generated by the venture. The entrepreneur must contribute the remaining \$100,000 to get the venture going.

To be clear: This means the VC's net earnings are either: a total loss of \$900,000 if the venture flops, or \$2.45 million (= $-\$900,000$ up-front + 67% of 5 million) if it is modestly successful, or \$66.1 million (= $-\$900,000$ + 67% of 100 million) if it is wildly successful. The entrepreneur faces final wealth levels of \$100,000 if the venture flops, \$1.75 million if it is a modest success, or \$33.1 million if it is a wild success.

Suppose the venture capitalist cannot contractually specify the effort level of the entrepreneur. Assume (for the moment) the entrepreneur's only alternative to taking the venture capitalist's offer is to walk away, leaving her with wealth of \$200,000 for sure.

- (a) Will the entrepreneur accept this deal and exert high effort, or not?
- (b) Assuming she will exert her preferred level of effort, find the entrepreneur's certainty equivalent wealth for this deal.
- (c) What is the venture capitalist's expected net profit from the deal?

B. *An Option Financing Arrangement*

Suppose, instead of the share financing arrangement above, that the venture capitalist arranges the contract so as to take a fixed payoff out of the venture only if the venture's gross earnings reach a very high threshold. Specifically, suppose that the venture capitalist will pay the entrepreneur a base amount of \$925,000 up-front, in return for which the venture capitalist will retain \$75 million out of the \$100 million the venture generates if it is wildly successful. The entrepreneur must contribute the remaining \$75,000 to get the venture going out of her initial wealth.

To be clear: This means the entrepreneur loses her \$75,000 initial contribution if the venture flops, leaving her with a final wealth of \$125,000 in that event. She keeps all \$5 million of the venture's gross proceeds if it is a modest success, for a final wealth of \$5.125 million (= \$5 million in gross proceeds + \$200,000 initial wealth - \$75,000 initial investment). And she keeps \$25 million of the venture's gross proceeds if it is a wild success, for a final wealth of \$25.125 million. The VC takes a net loss of \$925,000 if the venture is a flop or a modest success, but has a net gain of \$74.075 million ($\75 million - $\$925,000$) if it is a wild success.

- (a) Does the entrepreneur like this deal structure better? What is her CE wealth?
- (b) What is the venture capitalist's expected net profit now?

C. *Advising the Venture Capitalist*

If you were advising the VC, can you come up with better deal terms (higher expected profit for you)? You can use either the share or the option deal structure. If you use a share structure, choose a new base payment to the entrepreneur *and* a new share of the

gross proceeds (in percent) to the VC. If you use an option financing structure, choose a new base payment to the entrepreneur *and* a new amount to retain (in \$) if wild success occurs.

Can you find new deal terms that give the VC expected profit of \$6.6 million or more? (Note that you must to give the entrepreneur at least \$800K up front for the venture to get going, and her CE wealth from the deal must beat her \$200K in initial wealth or she'll turn you down).

Example 2. *Performance Pay*

A second leading example of the ideas in this note is the design of compensation contracts. Suppose that you employ a salesperson who is going to try to make a particular sale for you. If the person makes the sale, you will earn a profit of \$60,000. If the person does not make the sale, you earn \$0. These figures do not include the compensation you pay the salesperson.

This salesperson must decide on a level of effort to devote to selling for you. He can *work flat out*, he can *work hard*, he can *try but not hard*, or he can *loaf*. His level of effort affects the probability that he makes the sale. If he works flat out, he will make the sale with probability 0.5. If he works hard, the probability that he will make the sale is 0.4. If he tries but not hard, he makes the sale with probability 0.25. And if he loafs, he will make the sale with probability 0.05.

This salesperson has a utility function that depends on his wages and the amount of effort he puts into making the sale. If he is receives compensation W , his utility is

$$U = (W)^{1/2} - \text{disutility of effort,}$$

where his disutility of effort is 40 if he works flat out, 20 if he works hard, 10 if he tries but not hard, and 0 if he loafs.

The salesperson's best alternative to working for you is a job in which, for the length of time this sales call will take, he will make \$10,000 with no disutility of effort. So, to get this salesperson to work for you, you must provide him a contract where his expected utility is at least $(10,000)^{1/2} = 100$. You are risk neutral, and you want to maximize your net expected profit (net the compensation you pay to the salesperson, that is).

Before going further, I should comment on the assumption that you are risk neutral, while the salesperson is risk averse. This assumption is more than a matter of analytic convenience. If we think of "you" in this story as a sizeable, publicly-held company, with shares widely dispersed among many shareholders, then it is reasonable to suppose that you are risk neutral. On the other hand, the salesperson is likely to bear entirely the risk in his

own compensation, for which he is invariably somewhat risk averse.

An initial incentive contract. Of course, you cannot observe the salesperson's true level of effort directly, and it seems hard to imagine how an employment contract could be written to specify "high employee effort" so precisely. What do you do? To motivate the salesperson, we have to give him some incentive to try, which means making the amount he is paid if a sale occurs greater than what he makes if he comes back without a sale. But we do not want to go too far in this direction, since that means loading too much risk on him. It is a matter of trading off incentives against risk sharing, and finding the right balance.

Suppose you write a contract for the salesperson that specifies that he gets a base wage of \$9500, regardless of whether a sale is made or not, and in addition a bonus of \$15,000 if he makes a sale. On these terms, would he take the job? If so, what level of effort would he put in? What would be your net expected profit?

He has the following five choices:

1. *If he chooses not to work for you at all*, he nets utility 100 in his next best alternative.
2. *If he works for you and loaf*s, he receives income \$9500 with probability 0.95 and \$24,500 with probability 0.05. He has zero disutility of effort, so his overall expected utility is $(.95)(9500)^{1/2} + (.05)(24,500)^{1/2} - 0 = 100.42$.
3. *If he tries but not too hard*, he receives income \$9500 with probability 0.75 and \$24,500 with probability 0.25. He has disutility of effort equal to 10, so his overall expected utility is $(.75)(9500)^{1/2} + (.25)(24,500)^{1/2} - 10 = 102.23$.
4. *If he tries hard*, he receives income \$9500 with probability 0.6 and \$24,500 with probability 0.4. He has disutility of effort equal to 20, so his overall expected utility is $(.6)(9500)^{1/2} + (.4)(24,500)^{1/2} - 20 = 101.09$.
5. *If he works flat out*, he receives income \$9500 with probability 0.5 and \$24,500 with probability 0.5. He has disutility of effort equal to 40, so his overall expected utility is $(.5)(9500)^{1/2} + (.5)(24,500)^{1/2} - 40 = 87.00$.

Among his five options, taking the job and trying but not too hard maximizes his expected utility (net of the disutility of effort). So we conclude that this is what he would do, if offered this incentive contract. Therefore, your expected profit, net of the cost of his wages, is

$$(\$60,000)(0.25) - \$9500 - (\$15,000)(0.25) = \$1750.$$

This is your expected gross profit, less his base wage, less the bonus times the probability that he earns the bonus.

	A	B	C	D	E
1	Base wage	\$9,500		Gross profit	\$ 60,000
2	Bonus	\$15,000		if sale	
3					
4	Gross utility if sale	156.52	= $\text{sqrt}(\text{Base wage} + \text{bonus})$		
5	Gross utility if no sale	97.47	= $\text{sqrt}(\text{Base wage})$		
6					Expected
7	Effort Choice	Disutility	Prob Sale	EU	Net Profit
8	loafing	0	0.05	100.42	\$ (7,250)
9	tries but not hard	10	0.25	102.23	\$ 1,750
10	works hard	20	0.40	101.09	\$ 8,500
11	works flat out	40	0.50	87.00	\$ 13,000
12	<i>outside option</i>			100	

FIGURE 1

A better incentive contract. That is not bad, but can you do better? To help answer this question a bit faster, consider the (very) simple spreadsheet in Figure 1. Here is what it does: The base wage and bonus are in cells B1 and B2. The spreadsheet computes the salesperson's gross (of the disutility of effort) utility if a sale is made and if not in cells B4 and B5. Then the table computes the salesperson's EU and your expected net profit for each of the four possible effort levels, using the same formulas on the previous page.

With this compensation offer, you (the employer) have an expected net profit of \$1750. Note that, with this base wage and bonus level, if the salesperson works flat out then your expected net profit rises to \$13,000. But the nature of the moral hazard problem here is that you do not choose the effort level, the salesperson does. And he chooses his effort level (or whether to work for you at all) based on the incentives you put before him. With this contract, his choice is to try but not hard, and the only way you can induce a higher level of effort from him, intuition tells us, is to give him a bigger bonus if he makes the sale.

Will it be worthwhile for you to induce more effort from him? With the spreadsheet in hand, we can play with the base wage and bonus to find out. Suppose, for instance, you increase his bonus to \$20,000. See Figure 2. Now his choice is to work hard. And having induced this level of effort, your expected net profit rises to \$6500.

	A	B	C	D	E
1	Base wage	\$9,500		Gross profit	\$ 60,000
2	Bonus	\$20,000		if sale	
3					
4	Gross utility if sale	171.76	= $\text{sqrt}(\text{Base wage} + \text{bonus})$		
5	Gross utility if no sale	97.47	= $\text{sqrt}(\text{Base wage})$		
6					Expected
7	Effort Choice	Disutility	Prob Sale	EU	Net Profit
8	loafing	0	0.05	101.18	\$ (7,500)
9	tries but not hard	10	0.25	106.04	\$ 500
10	works hard	20	0.40	107.18	\$ 6,500
11	works flat out	40	0.50	94.61	\$ 10,500
12	<i>outside option</i>			100	

FIGURE 2

Note that the salesperson's expected utility in the arrangement in Figure 2 is 107.18. You need to give him an expected net utility of only 100 to get him to work, so you are paying him more than his next best outside alternative. That means there is still some room for maneuver. We continue to get the desired incentive effect (motivating hard work) if the bonus is decreased a bit. While we do this, we must add a bit to the base wage, to keep his net expected utility at or above the reservation utility level of 100. Figure 3 shows that, by hunting around, with a base wage of \$8700 and a bonus of \$17,000, the salesperson is still willing to work for the firm (100.089 vs. 100 for his next best alternative) and still prefers working hard to trying but not hard (100.089 vs. 100.033). This gives an expected net profit of \$8500.

	A	B	C	D	E
1	Base wage	\$8,700		Gross profit	\$ 60,000
2	Bonus	\$17,000		if sale	
3					
4	Gross utility if sale	160.31	= $\sqrt{\text{Base wage} + \text{bonus}}$		
5	Gross utility if no sale	93.27	= $\sqrt{\text{Base wage}}$		
6					Expected
7	Effort Choice	Disutility	Prob Sale	EU	Net Profit
8	loafing	0	0.05	96.63	\$ (6,550)
9	tries but not hard	10	0.25	100.03	\$ 2,050
10	works hard	20	0.40	100.09	\$ 8,500
11	works flat out	40	0.50	86.79	\$ 12,800
12	outside option			100	

FIGURE 3

There is one final thing to try. Expected net profit is better still if we get the salesperson to work flat out. Indeed, we know that this is the effort level to shoot for if we could directly monitor and verify that effort level. Can we increase the bonus enough that this is his choice, leaving us with a higher expected net profit? Alas, getting him to work flat out requires an enormous bonus. The best one can do is a base wage of \$1530 and a bonus of \$56,500, which yields a paltry expected profit of \$220. The optimal incentive scheme will not involve motivating him to work flat out.

Hunting with a spreadsheet like this for successively better schemes is a fast and simple thing, in contrast to real-life implementation. But there are several important points to all this. Mechanically, we are applying two general procedures: Find what is the best one can do in terms of expected net profit if we give the salesperson sufficient incentive to take each of the four actions. And for each level of effort, fine-tune the incentive scheme so that (1) the salesperson is brought as close to his reservation utility of 100 as possible and (2) he is shielded from risk to the maximum extent possible, consistent with keeping him at the effort level being targeted. Shielding him from risk is good for both of you, since he is risk averse and will therefore require a lower expected compensation if he faces less income risk.

Time for a Reality Check

Wait, you say. How practical is all this really, since you rarely know with precision the utility function of salespersons who work for you?

The point of this analysis is twofold. First, what we are doing here is mimicking how one goes about finding profit-maximizing incentive contracts for employees in general: It is a trial and error process. It takes time, and you learn from your initial trials how much people produce (via the profit you earn, how often you pay out bonus payments, and so on). That process gives you the same general feedback that a spreadsheet model does. Carrying out this trial-and-error process in reality can be *far more profitable* than paying a straight wage and leaving the employee with no economic incentive to exert effort, or just picking an initial wage and bonus structure and never using the feedback it provides to fine-tune the structure. This much any successful sales manager can attest.

The second, equally important insight generated by hunting for the best incentive contract here is the basic trade-off between risk sharing versus motivation. To motivate higher levels of effort, you have to pay a higher bonus. But as you pay a higher bonus, you load more and more risk on the salesperson (because the uncertainty over his total income increases). The optimal solution involves a compromise of these two opposing forces: You want a high enough bonus to motivate the level of effort you target, and you want to be sure that this target is the best effort level to aim for.

Two Recent Observations on Incentive Pay

In all sorts of situations where individuals must be given incentives to take desired actions, incentives are based on imprecise assessments of the individual's actions and we see the same economic trade-off. Two articles attached at the end of this note summarize two firms' recent experiences with changing performance pay within their organizations. One article ("Paying by the Piece") clearly establishes that when output can be observed, bonuses tied to output (called *piece-rate pay*) enhance productivity. The second article ("Seasoning Compensation Stew") shows how extensive bonus plans are within some corporations. Question: Why, as you move down the hierarchy of the firm, are bonuses tied to individual performance rather than the firm's overall performance?

Application: Stock Options for Top Managers

Top managers have jobs that are multi-task in character, so providing them with good direct incentives is especially tricky. On grounds that the value of shareholder equity is a good proxy for what is good for the organization, compensation for top managers is often tied to the price of the firm's equity through stock and stock options, which allow top managers to buy shares in the company at a fixed price.

Stock options given to executives have been controversial, in terms of how they should be reported in a firm's financial accounting statements and for the level of income they provide top executives. But their basic *raison d'être* is not so controversial. They are meant to tie the compensation of top management to the fortunes of the enterprise, thus aligning the interests of top managers with those of equity holders. Stock options are used, the story goes, because the market price of equity most accurately captures how the enterprise is doing. Stock market prices capture the value of the firm better than any other formula, because market prices are set by savvy investors interested in predicting the future value of the firm.

This is a nice story, but as events in recent years have shown, the assumption on which this story is based—that the market price of equity reflects the true value of the firm—may be flawed. To be clear, the market price of a firm's stock reflects the information that equity markets possess. However, to the extent that top management can manipulate financial statements, top management may be able to manipulate (at least for a while) the price of their firm's stock. Given options keyed to the market price of the firm's equity, top management is sometimes motivated to inflate that market price.

Of course, top management might respond to their incentives by improving the economic fortunes of the firm, as intended. But they can also serve that incentive by engaging in accounting practices that paint a rosier picture of the firm's prospects than is the truth, hoping to cash in the options and get their money before their misleading accounting practices have been discovered. The bottom line: *Providing incentive compensation can create a problem of mal-aligned incentives that is worse than providing no incentive compensation at all—if the firm's true performance is not monitored accurately by its shareholders.*

The *NYT Economic Scene* article attached at the end of this note explores this tension—and the dangers that might arise from policy reforms in Washington that ignore it.

Application: Corporate Takeovers

The traditional corporate governance structure is straightforward. Shareholders own the firm's assets and assume the risks of doing business. Any residual earnings from the actions of managers are split among the shareholders. Shareholders hire managers (we'll call them *agents*) to perform the duties of running the business. Agents choose an action from a number of alternative possibilities (e.g., allocating resources within the firm). Actions affect both the welfare of the agents and the shareholders. However, the interests of the two parties are usually not identical.

In this context, we can see the implications of moral hazard problems in two ways. First, think about the decentralized firm. Because it is difficult to centralize knowledge about productivity or cost-saving opportunities, the divisional managers have information about

production possibilities not available to corporate staff. The productive units may in fact have incentives not to reveal their full potentiality, because it is then easier to operate with less taxing requirements.

Second, the owners of the firm (we'll call them *principals*) try to reduce the moral hazard problem by monitoring their managers. With monitoring, principals try to get additional information about an agent's actions in excess of what is revealed by standard financial measures of firm performance. Monitoring, however, is often costly and difficult even for a firm's owners. How then do managers get monitored?

One method is outsiders who evaluate when a firm is being inefficiently managed, its resources under-utilized, and its assets poorly deployed. Basically, outside investors and managerial teams monitor the market for *managerial control*, and look to replace inefficient managers. Such a market is called *the takeover market*. In viewing the monitoring process in this way, the active players are competing managerial teams (or their institutional investor supporters). Smaller shareholders play a passive role, providing the funds for takeovers.

Let's see how this model changes the corporate governance structure. Shareholders are willing to let other managers monitor the management of their firm. They are willing to accept the management team that brings the highest price for their shares in the firm. Managers are disciplined not directly by shareholders, but by other teams of managers who, if present management is inefficient, will offer shareholders a better deal.

There are several mechanisms through which this market for managerial control takes place. Some of the most common are:

Mergers and tender offers. Here a team of outsiders offers to buy a firm's common stock for a premium above market price. Mergers are generally negotiated by the boards of the competing teams. Tenders are taken right to the shareholders.

Proxy fights. Here a competing team attempts to gain control of a firm's board by convincing shareholders the new team can do a better job, and getting shareholder votes for the new team's candidates. This is more like a political contest, with incumbents usually holding several advantages.

Research indicates that target firms in successful takeovers realize abnormally higher share-price changes of about 20% in mergers, and 30% in tender offers. Bidding firms realize abnormally higher price changes of only 4% in tender offers, and 0% in mergers. The total gain to bidders *and* targets appears to be positive. This means that in spite of reports in the popular press, research indicates that on average, takeovers create value for shareholders. Several reasons have been advanced for why value is created:

1. Elimination of inefficient managers in the target firm.
2. Efficiency improvements because the acquiring firm has resources that the target firm lacks. The assets of the target firm are worth more to the bidding firm, so shareholders of the target firm realize a profit. (This is what is meant when executives use the term *synergies* to describe a merger or takeover.)
3. Tax savings due to increased leverage or other tax shields.
4. Increased market power (mainly in horizontal mergers).

Studies seem to indicate the market power hypothesis is not common. There are plausible reasons for this: mergers that are expected to increase market power are often opposed by the U.S. or E.U. anti-trust authorities, which results in the merger being scuttled (or a long and costly legal battle with the government, as occurred in the Oracle-PeopleSoft merger in 2005). The third explanation (taxes) arises in some contexts, but they are often limited. This leaves the first two explanations for the main reason why value is created through corporate takeovers. There is a market for corporate control, and this market mechanism may help reduce moral hazard among corporate executives who are not maximizing profit for shareholders.

Managerial Implications

1. Incentive contracts are useful when the outcome you care about depends on the efforts of others, but you can't verify that effort directly. The task in creating value through an incentive contract is to trade off the tension between *value creation through risk sharing* and *pecuniary motivation to take effort*.
2. Moral hazard implies less-than-efficient risk sharing. The more it can be reduced through direct monitoring, the more value can often be created.
3. How do you devise good incentive contracts? Typically it is a trial and error process, informed by the logic of reducing the risk faced by risk-averse individuals *up to the point* where they would not undertake desirable levels of effort.
4. Performance contracts tend to have little effect, and thus little value, if the outcome you care about is barely affected by the agent's effort.
5. Implementation requires two seemingly obvious, but often overlooked points: (1) You can measure the outcomes you care about (e.g., sales per salesperson), and (2) these measures are not subject to manipulation or misrepresentation by the agent. The latter is one of the lessons from Enron, WorldCom, and other accounting scandals. If you can't accurately verify the outcomes you base your incentives on, then your incentives can create more harm than good.

Paying by the Piece Increases Productivity

Economic theory holds that if a firm switches from paying hourly wages to paying piece rates, the average productivity of workers will rise, the firm will attract a more able work force, and variations in output among workers also will rise. However, there have been few attempts to test this theory, mostly because of a lack of suitable data. But NBER Research Associate **Edward Lazear** found the data at the nation's largest installer of auto-glass, Safelite Glass Corporation, based in Columbus, Ohio. During

1994 and 1995, after the introduction of new management, that company gradually changed the compensation method of its work force, moving from hourly wages to piece rate pay. "The theory is borne out almost to the letter," Lazear writes in **Performance Pay and Productivity** (*NBER Working Paper No. 5672*). "Claims by sociologists and others that monetizing incentives may actually reduce output are unambiguously refuted by the data." The effects were dramatic: First, the switch boosted the average level of

output for the company as a whole by about 35 percent. Second, for a given individual, the use of piece rates increased productivity by about 20 percent. The firm shared the gains in productivity with its workforce; the piece rate was set so that a worker captured in extra pay about half of his increase in productivity. Thus, the average incumbent worker's wage rose by 9.6 percent as a result of the switch.

Third, the new pay system reduced turnover among the most productive workers, but increased

turnover overall. The system also enabled the firm to attract more able workers as employees.

Fourth, the switch in systems resulted in a greater range in workers' abilities and output. On average,

that guaranteed range; now, fewer than 30 percent are in the guaranteed range. Fifth, the piece rate pay reduced the amount of paid sick leave, probably by making sick leave relatively more costly to take.

"...the switch [to piece rate compensation] boosted the average level of output for the company as a whole by about 35 percent."

glass installers are paid \$20 per unit installed, but the new system included a guaranteed wage of about \$11 an hour (if a worker's weekly pay on the basis of units installed fell below that guarantee). Initially, many workers ended up in

The data from Safelite came from a sophisticated computerized information system which keeps track of how many units of each kind each installer in the company installed in a given week. Since the performance pay plan was phased in over an

18-month period, many workers were employed under both hourly and piece-rate regimes. Thus, data on individual output are available for installers under both systems. With an average 2,040 workers employed during the average month, the result was 38,764 monthly measures of individual output. Installers averaged 2.96 glass units per day over the entire period. Under the performance pay formula, the average worker earned \$2,250 a month when vacation, holiday, sick pay, and other factors are included with the performance pay.
DRF

Seasoning Compensation Stew


Varying the Recipe Helps TV Operation Solve Morale Problem

Star System

To reward its top achievers, Discovery Communications has introduced a bonus program that is a leading example of the move toward pay for performance. The value of a potential bonus is much higher for upper-level managers; lower-level employees are rewarded more directly for their own performance.

BONUS DEPENDENT ON PERFORMANCE OF:

- OVERALL COMPANY
- CORPORATE DIVISION
- INDIVIDUAL EMPLOYEE



By JONATHAN D. GLATER

BETHESDA, Md. — A few years ago, Mark Kozaki and other executives at Discovery Communications Inc. had a morale problem on their hands.

Some of the company's top achievers were earning the same salaries as slackers in the same position. In other cases, employees doing the same work as colleagues in other parts of the company, which is based in Bethesda, were being paid less.

The only way for managers to reward the workplace stars was to give

them a bonus or promote them to another position; Discovery's compensation structure did not allow for big raises for people who remained in the same job.

"People felt that it was unfair," said Mr. Kozaki, who is senior vice president for administration and operations in the United States for Discovery, which owns the Discovery Channel, Animal Planet and 31 other cable networks.

The company found that its compensation structure was not just unfair, it was often counterproductive. Some people accepted promotions to jobs they were not equipped for, while others tried to switch to more generous de-

partments in the company.

And so, in early 1998, Discovery began an overhaul of its compensation policy, switching to a pay-for-performance system that allowed for both big raises and bonuses.

Even though the traditional justification for maintaining a relatively fixed pay structure is to promote a sense of equity among employees, by abandoning it, Mr. Kozaki said, dissension in the ranks soon dissipated.

"It has eliminated many of the concerns about whether there is across-the-board fairness," he said. "There is

not as much discussion or wondering or suspicion" as before.

Experts in employee compensation say more companies are shifting away from fixed pay structures that have long been the norm to more flexible performance-driven arrangements like Discovery's.

Though the old practice of keeping pay within narrow boundaries for workers in the same job classification still prevails, an increasing number of managers are realizing that they have to come to grips with the challenge of rewarding their best workers without raising labor costs significantly.

The solution for many companies is to hold automatic pay raises to a minimum, thus punishing mediocre employees, while lavishing merit raises and bonuses on the superstars.

The bonuses are increasingly based on measurements of both the individual's attainments and the company's overall profits and other results.

More than half of the 2,400 companies surveyed in 2000 by William M. Mercer Inc., a human resources consulting firm in New York, reported that they had performance-based incentive programs in place for both management and nonmanagement employees. And 49 percent said that they had increased the number of employees eligible for such programs since 1997.

"It's due to two issues," said Steven E. Gross, who runs the compensation consulting practice at Mercer. "Attraction and retention of good people, and companies trying not to raise their fixed costs."

And by linking bonuses to corporate results as well as individual attainments, he said, companies can share the wealth in good times, but can cut bonuses before jobs in bad times. "It self-corrects some of your costs," he said.

Mr. Gross said some studies suggested that every \$1 invested in an employee incentive program could lead to a \$2 increase in revenue, but he cautioned that so many variables had to be considered in making that calculation that conclusions were extremely tentative.

Managers, of course, have always been rewarded for doing well, but the expansion of financial incentives into the middle and lower ranks of the organization is a relatively recent phenomenon that has accelerated

over the last decade, compensation experts say.

Late in 2000, as just one illustration, the New York law firm of **Stroock & Stroock & Lavan** announced that it would begin paying bonuses to junior lawyers based on their individual performance, a bounty once reserved for partners only.

At Discovery, which grew so quickly that the compensation structure long went unexamined, Mr. Kozaki said that attrition rates had fallen and complaints from employees were less frequent since the new policy went into effect. "Employees think they are being treated fairly," he said.

Under the old system, an information-technology specialist who was doing good work could not have received a raise when competitors' wages rose for comparable employees unless she was also promoted, said Anthony R. Amato, Discovery's vice president for compensation, benefits and human resources.

Under the new system, the specialist could receive a raise — rewarding good performance and also matching competitors — without being burdened with new management responsibilities that she was not ready for, he said.

"Historically, we weren't separating performance" of the individual employee, that employee's job and the wages paid by competitors from evaluation of that employee's contribu-

tion for more responsibility, Mr. Amato said. That meant that employees might be promoted too quickly if they did well or if it was necessary to give them a raise to keep them from defecting.

Now, Discovery provides two types of incentives to workers. First, all employees are eligible for a merit increase in salary. A computer program considers an employee's current base salary and an evaluation by that employee's manager, then recommends a minimum and maximum raise.

Second, every employee is eligible for a year-end bonus, which is a percentage of the employee's base pay. Again, the bonus is based in part on a manager's evaluation, but it incorporates the company's performance, too.

Sherner Sumter is an executive assistant who has worked at Discovery for a year and a half. Under the new pay structure, she is eligible for both a raise based on her manager's evaluation of her job performance and for a bonus of up to 10 percent of her salary. Her boss's evaluation of how well she does her job and how eagerly she shows initiative and takes on new responsibilities accounts for 80 percent of the potential bonus, while 10 percent is based on how her division does, and a final 10

More companies are shifting away from fixed pay structures.

percent is based on how the company as a whole does.

"It helps you to be motivated," Ms. Sumter said. The lure of higher pay has prompted her to take on "a lot of different projects" that she might otherwise not seek out, she said, like developing a system that allowed her department to respond more quickly and efficiently to requests from different networks. It has also inspired her to speak up more on ways her department could streamline some of its operations.

"I might go in and say, 'What can we do to improve the way we manage the data?'" she said.

The technology that Discovery has developed to link individual bonuses

to divisionwide and corporatewide performance sets the company apart, said Russell H. Miller, a partner at S.C.A. Consulting, a compensation adviser that helped Discovery come up with the structure — though not the computer system — for its compensation.

"They're on the leading edge of integrating those systems fully," he said, adding that S.C.A. did not have the technical skill to develop Discovery's system. "They are more sophisticated than I typically come across."

Discovery used Lotus Notes to develop a database that could take into consideration managers' evaluations of employees, overall corporate performance and employees' responsibilities to recommend a minimum and maximum bonus for individual employees.

Before the adoption of the new program, employees received an annual cost-of-living adjustment to their wages, usually a little above the rate of inflation for that year, said the company's president and chief operating officer, Judith A. McHale.

Although employees could receive bonuses, the only way to receive a raise was to be promoted.

Such promotions are not necessarily a good thing, Ms. McHale said, because doing one thing well does not necessarily mean that an employee is ready to be promoted to manage several people doing the same thing.

"We needed to be able to, within a job category, acknowledge that there are different levels of performance," she said.

The hard part of overhauling compensation is deciding how loud a message to send, Mr. Miller said. Overrewarding star performers could backfire by causing jealousy among co-workers and eroding their productivity, he said. And it is probably still too early to know whether Discovery's new incentive structure, which is still in its infancy, will have an adverse effect.

"The key is to make sure that you have a system or process that gives you the highest likelihood of success of accurate differentiation," Mr. Miller said. "The less arbitrary it is and the more objective it is, the better."

Some resentment is probably inevitable, though, in any system that ranks some people as more valuable than others, he said. Whether such a system is fair, he added, "is a philosophical decision."

Business 'reforms' should not ignore the positive roles of incentives and competition.

IN a recent interview with Wired magazine, the technology pundit George Gilder reflected on how he had gone from the world's best stock picker for several years to the worst in 2001. Most telecommunications companies he promoted have lost almost all their value.

Newsletter readers who bought and held the stocks he recommended lost big, as did Mr. Gilder himself. He's now broke.

When prices got high, why didn't he suggest selling?

"In retrospect, it's obvious that I should've subtly said, 'Hey, things have gotten out of hand at JDS Uniphase, and it's not worth what you'd have to pay for it,'" he told Wired. But he didn't, for fear that his influence would make the stock price crash.

His newsletter's picks moved prices up. Presumably, his advice would also move them down.

"If I had said, 'Hey, this is a top, you should all sell,' it would've been a cataclysmic event," he said, adding, "Half of my subscribers would have been eternally grateful, but the other half — the new ones — would've been enraged because they had just come in."

There, in a nutshell, is the fundamental problem of a market bubble: Everyone who owns stock wants prices to stay high.

When reality finally breaks the bubble, stockholders cry that they weren't warned. But when prices are high, it seems almost irresponsible to deliver bad news. If influential analysts or executives tell people to sell, they can cost current shareholders millions.

Hence the pressures to keep earnings high, leading some companies to adopt devious or fraudulent accounting practices. Stockholders weren't just the victims of such shady practices. They were also, in many cases, the beneficiaries.

Making executives look out for shareholder value is, after all, why companies skewed their compensation systems toward stock options and other equity compensation.

The goal was to give decision makers a strong incentive to keep the stock price high. They would then act in the owners' interests rather than pursuing less-profitable pet projects, avoiding painful spending cuts, or lavishing money on their personal perks.

In some cases, that incentive worked all too well. Stock prices grew, executives became rich, and so did shareholders. But for some companies, the earnings supporting those prices weren't real.

Fraud didn't cause the bubble, but the bubble encouraged fraud.

Cooking the books — or, legally but problem-

atically, always making judgment calls in the direction of higher earnings — helped support otherwise insupportable stock values. If it's the job of management to represent the interest of current shareholders, arguably that's what aggressive accounting did.

Like Mr. Gilder's newer subscribers, current shareholders are hurt by bad news. In the short term at least, they're better off with cheerful lies.

The fundamental problem is that the interest of current shareholders is not the same as the interest of a well-functioning market.

The policies Congress is rushing through to demonstrate Washington's concern with corporate shenanigans won't do much to address this problem. Some will actually make things worse.

Take the idea, pushed by Senator John McCain, the Arizona Republican, and others; that executives shouldn't be able to sell their stock until they've left the company.

Rather than encourage sound management that builds good companies, that would actually punish long-term commitments. Executives who devoted their careers to a single business would be financial fools, never able to diversify their personal portfolios.

Managers who really did build value would need to move on after a few years to reap the rewards of their tenure. Corporate America would increasingly be run by job-hopping generalists, not long-term chief executives with deep knowledge of their particular business. And in exchange for assuming more financial risk, top managers would demand even higher salaries.

At start-ups, which have more promise than cash for salaries, the effect would be devastating. Someone like Margaret C. Whitman, who brought management experience at an early stage to

eBay, would probably decline to head a start-up altogether or would leave and cash out her stake as soon as possible. Before it went public, eBay lured Ms. Whitman with options to buy 14.4 million shares for 3.5 cents each.

Successful founders would have to retire early to cash out their equity. There'd be no more long careers like the ones that built Intel, Federal Express, Microsoft or Southwest Airlines.

Such a strong financial incentive for executives to change jobs would mean growing companies suffer constant turnover in their management. Neither current nor future shareholders, or a healthy marketplace, would be served by that.

But what about the problem of fly-by-night chief executives and the shareholders who benefit from inflated results? If the shareowners themselves don't have an interest in truthful financial reporting, who guards the corporate guardians?

One possibility is the stock exchange. While artificially high earnings may help a company, accurate results serve the market as a whole. An exchange that independently monitored its listed companies' results would benefit from investor confidence at the expense of laxer competitors.

If auditors were hired by the exchange, rather than the company, and given a forensic mandate, they would enjoy both more freedom and more prestige. Indeed, turning audits into independent investigations might help solve one of accounting's biggest problems: Auditing is too boring to attract and keep the best accountants.

Even now, the mob waving pitchforks and torches finds the details of accounting, compensation and corporate governance too tedious to take seriously. But "reforms" that ignore the role of incentives and competition will turn out to be monsters themselves.