

Incentives

WHY DID THE socialist economic systems fail so miserably? The best-laid Five-Year Plans of Stalin and his successors “gang agley” because the workers and the managers lacked adequate incentives. Most importantly, the system offered no reward for doing a good job rather than a merely adequate one. People had no reason to show initiative or innovation, and every reason to cut corners wherever they could—fulfilling quantity quotas and slacking on quality, for example. In the old Soviet economy, this was captured by the quip: They only pretend to pay us, so we only pretend to work.

A market economy has a better natural incentive mechanism, namely the profit motive. A company that succeeds in cutting costs or introducing a new product makes a greater profit; one that lags behind stands to lose money. But even this does not work perfectly. Each employee or manager in a company is not fully exposed to the chill wind of competition in the market, and the top management of the firm has to devise its own internal carrots and sticks to obtain the desired standards of performance from those below. When two firms join forces for a particular project, they have the added problem of designing a contract that will share the incentives between them in the right way.

We develop the components required for a smart incentive scheme through a series of examples.

INCENTIVES FOR EFFORT

Of all the steps involved in writing a book, surely the most tedious for an author is the correction of the printer’s proofs. For readers unfamiliar with the process, let us briefly explain what that entails. The printer sets the type from the

final manuscript. These days this is done electronically and is therefore relatively error-free, but strange errors—missing words and lines, chunks of text shifted to wrong places, bad line and page breaks—may still creep in. Moreover, this is the author's last chance to correct any slips of writing or even of thinking. Therefore the author has to read the printer's copy in parallel with his own manuscript, and locate and mark all errors for the typesetter to correct.

The author is reading the same text for the umpteenth time, so it is not surprising that his eyes glaze over and he misses several errors. Therefore it is better to hire someone, typically a student in our case, to do the proofreading. A good student can not only catch the typographical errors but also spot and alert the author to more substantive errors of writing and thinking.

But hiring a student to read proofs brings its own problems. The authors have a natural incentive to make the book as error-free as possible; the student is less motivated. Therefore it becomes necessary to give the student the correct incentives, usually by paying him in relation to how well he does the job.

The author wants the student to catch all the typesetting errors there may be. But the only way the author can tell whether the student has done a perfect job is to do a perfect check himself, which defeats the whole purpose of hiring the student. The student's effort is unobservable—he takes away the materials and comes back a week or so later to give the author a list of the errors he has found. What is worse, even the outcome cannot be observed immediately. The author will find out about any error that the student failed to catch only when some other reader (such as yourself) finds and reports the errors, which may not happen for several months or even years.

So the student has the temptation to shirk—just hold on to the materials for a few days and then say that there were no errors. Thus it won't do to offer the student a fixed flat sum for the job. But if offered a piece rate (so much per error he finds), he may worry that the typesetter has done a perfect job, in which case he will have to spend a week or more on the work and get no money at the end of it. He will be reluctant to take the job on these terms.

We have a problem of information asymmetry, but it is different from the ones we considered in chapter 8. The author is the informationally disadvantaged party: he cannot observe the student's effort. This is not something innate to the student; it is the student's deliberate choice. Therefore the problem is not one of adverse selection.* Rather, it is similar to the problem that an insured homeowner may be less careful about locking all doors and windows. Of course insurance companies regard such behavior as almost immoral, and they have coined the term *moral hazard* to describe it. Economists and game theorists take a more relaxed attitude. They think it perfectly natural that people would

respond in their own best interests to the incentives they face. If they can get away with shirking on the job, they will do so. What else should one expect from rational players? The onus is on the other player to design the incentive scheme better.

Although moral hazard and adverse selection are different issues, there are some similarities in the methods for coping with them. Just as screening mechanisms have to consider the restrictions of incentive compatibility and participation, so do incentive payment schemes that cope with moral hazard.

A fixed sum does not handle the incentive aspect well, whereas a pure piece-rate payment does not handle the participation aspect well. Therefore the compensation scheme has to be a compromise between the two extremes—a flat sum plus a bonus per error the student discovers. This should give him enough assurance of the total compensation to make the job sufficiently attractive and also enough incentive to attempt a thorough reading.

One of us (Dixit) recently hired a student to read the proofs of a 600-page book. He offered a fixed payment of \$600 (a dollar per page), plus the outcome-based incentive of \$1 per error found. (The student found 274.) The work took the student about 70 hours, so the average per hour, \$12.49, was quite a decent wage by undergraduate student standards. We don't claim that the scheme was fully optimal or the best deal Avinash could have gotten. And the outcome of the job was quite good but not perfect: about 30 errors the student missed have come to light since.* But the example illustrates the general idea of a mixture and how it works in practice.†

We see the same principle applied in many jobs and contracts. How do you pay a software designer or a copywriter? It is hard to monitor their hours. Is the time spent playing foosball, web surfing, or doodling part of the creative process or just slacking off? Even more important, it is even harder to measure how hard they have worked. The answer is to base part of their compensation on the success of the project and the success of the company, and that can be done using the company's stock or stock options. The principle is to combine a basic salary and an incentive bonus tied to the outcome. The same applies with even greater force to the compensation of higher-level management. Of course, like everything else, these schemes can be manipulated, but the general principle behind their use as incentive payment mechanisms remains valid.

Many extensions and applications of this principle have been studied by game theorists, economists, business analysts, psychologists, and others. We will give you a brief taste of some of this work and point you to references where you can pursue the subject as far as you wish.¹

HOW TO WRITE AN INCENTIVE CONTRACT

The main problem of moral hazard is the unobservability of a worker's action or effort. Therefore payments cannot be based on effort, even though more or better effort is what you as an employer want to achieve. Payments must be based on some observable metric, such as the outcome or your profit. If there were a perfect and certain one-to-one relationship between such observable outcomes and the unobservable underlying action, perfect control of effort would be possible. But in practice outcome depends on other chance factors, not just the effort.

For example, the profits of an insurance company depend on its sales force and claim agents, the pricing, and Mother Nature. In a season with many hurricanes, profits will be depressed no matter how hard people work. In fact, they will often have to work harder due to the increased number of claims.

The observable outcome is only an imperfect indicator of the unobservable effort. The two are related, and therefore incentive payments based on the outcome are still useful in influencing effort, but they don't work to perfection. Rewarding an employee for a good outcome in part rewards him for good luck, and penalizing him for a poor outcome in part penalizes him for bad luck. If the chance element is too large, the reward is only poorly related to effort, and therefore the effect of outcome-based incentives on effort is weak. Realizing this, you would not offer such incentives to any great extent. Conversely, if the chance element is small, then stronger and sharper incentives can be used. This contrast will appear repeatedly in what follows.

Nonlinear Incentive Schemes

A special feature of many incentive schemes, for example payment per error found by the proofreader, or a fixed percentage of the sale paid to a salesperson, or payment in stock that constitutes a given fraction of the firm's profit, is their linearity: the incremental payment is strictly proportional to the improvement in the outcome. Other commonly used payment schemes are distinctly nonlinear. The most obvious is a bonus that is paid if the outcome exceeds a specified threshold or an assigned quota. What are the relative merits of a quota and bonus scheme, as compared to a linear or proportional compensation scheme?

Take the context of a salesman, and consider a pure quota bonus scheme: the salesman is paid a low fixed sum if he fails to meet the quota during the year and a higher fixed sum if he does. First suppose the quota is set at such a level that

by exerting a hard level of effort he stands a good chance of meeting it, but the chance drops substantially if he slacks off even a little. Then the bonus provides a powerful incentive: the salesman stands to gain a lot of income or lose a lot, depending on whether he decides to work hard or shirk.

But suppose the quota is set at such a demanding level that the salesman has little chance of meeting it even with superhuman effort. Then he will see no point in exerting himself for the sake of that unlikely bonus. And circumstances may change during the year, turning what you thought was a well-set quota into something too demanding and therefore ineffective.

For example, suppose that the quota for the full year is not absurdly high, but the salesman has some bad luck in the first half of the year, making it unlikely that he would meet the year's quota in the six months that remain. That would cause him to give up and take things easy for the rest of the year, which is surely not what the employer wants. Conversely, if the salesman gets lucky and meets the quota in June, again he can relax for the rest of the year: there is no added reward to any further effort in that year. In fact, the salesman could conspire with some customers to postpone their orders to the next year to give him a good start toward next year's quota. Again that is hardly in your interests.

TRIP TO THE GYM NO. 9

The typical real estate agent commission is 6%, which is a linear incentive. How much of an incentive does your agent have to get a higher price? What would an extra \$20,000 in the price bring in? Hint: The answer is not \$1,200. How would you design a better incentive scheme? What might be some of the issues with your alternative scheme?

This illustrates in an extreme form the drawbacks of many nonlinear payment schemes. They have to be designed just right, otherwise their incentives may be ineffective. And they are prone to manipulation. Linear schemes may not give extra incentives at just the right points, but they are much more robust to changing circumstances and misuse.

In practice, combinations of linear and nonlinear schemes are often used. For example, salespeople usually get a percentage commission in addition to a bonus for achieving a given quota. And there may be larger bonuses for achieving further thresholds, such as 150 percent or 200 percent of the base quota. Such

mixtures can achieve some of the useful purposes of the quota without risking the larger drawbacks.

Carrots versus Sticks

An incentive payment scheme has two key aspects: the average payment to the worker, which must be enough to fulfill the participation constraint, and the spread of payments in good versus bad outcomes, which is what provides the incentive to exert more or better effort. The bigger the spread, the more powerful the incentive.

Even holding the spread constant, the incentive scheme can be designed as either a carrot or a stick. Imagine that the spread is 50 (and the average payment is 100). Under a carrot reward scheme, the employee would get, for example, 99 most of the time and 149 in the event of exceptional performance. The required output for the exceptional performance reward is set so high that the chance of hitting this target is only 2%, assuming the worker puts forth the desired amount of effort. Conversely, under a stick reward scheme, the employee would get 101 almost all the time, but a punishment of 51 in the event of exceptionally poor performance. Here, the bar for poor performance is set so low that the chance of hitting it is only 2%, assuming the desired level of effort is provided. Although the schemes feel quite different, both the spread and the average payments are the same.

The average is determined by the participation constraint, and therefore in turn by how much the worker could have earned in opportunities other than this employment. The employer wants to keep the payment to the worker low so as to increase his own profit. He may deliberately seek agents with poor alternatives, but workers who come for such low average payments may have low skills; an adverse selection problem may raise its ugly head.

In some cases the employer may have strategies that deliberately reduce the worker's alternatives. That is exactly what Stalin attempted. The Soviet state did not pay its workers very well even when they performed well, but if they didn't, they went to Siberia. They could not escape from the country; they had no outside opportunities.

This could have been a good incentive scheme, in the sense that it provided powerful incentives and was cheap to operate. But it failed because the punishment scheme wasn't closely tied to effort. People found that they might be denounced and punished almost as easily if they worked hard or shirked, so they did not have the incentive to work hard after all. Luckily, private employers and

even the state in modern democracies lack these arbitrary powers to reduce the workers' alternative opportunities.

But think of the compensation schemes of CEOs in this light. It seems that CEOs gain huge sums as incentive rewards if their companies do well on their watch, but they get only slightly less huge sums if the companies just do okay, and a “golden parachute” if the company actually fails on their watch. The average of these large sums, calculated using the probabilities of the possible outcomes, must be greatly above what is truly needed to induce these people to take the jobs. In the theory's jargon, their participation constraints appear to be grossly overfulfilled.

The reason for this is competition for the CEO candidates. Compared to the alternative of driving a cab or retiring early to play golf, the pay is well more than what is needed to keep the person working. But if another company is willing to guarantee \$10 million no matter what happens, then the participation constraint with your company is not compared with driving a taxi or playing golf but with taking a \$10 million CEO package elsewhere. In Europe, where CEO pay is generally much lower, companies are still able to hire and motivate CEOs. The lower pay still beats playing golf, and because many candidates are unwilling to move their families to the United States, the relevant participation constraint is with the other European companies.

THE MANY DIMENSIONS OF INCENTIVE PAYMENT SCHEMES

We have so far emphasized a single task, such as reading the proofs of a book, or selling a product. In reality, each context where incentive schemes are used has multiple dimensions—many tasks, many workers—even many workers engaged simultaneously in the same or similar tasks—and many years before the outcome is fully known. Incentive schemes must take into account the interactions among all these dimensions. This makes for some difficult analysis, but some simple ideas emerge. Let us look at some of them.

Career Concerns

When the job is expected to last for several years, the worker may be motivated in the early years not by immediate cash (or stock) payments, but by the prospect of future salary increases and promotions—that is, by incentives that extend over the full career. Such considerations are stronger for people who have longer futures with the company ahead of them. They are less useful for

workers near retirement, and not very useful for youngsters who are just starting out in the labor market and expect to move and change jobs a few times before settling on a career. Promotion incentives are most useful for younger employees at lower and middle levels. To use an example from our experience, assistant professors are motivated to do research by the prospects of tenure and promotion much more than they are by immediate salary increases within the rank.

In the example of the student reading the proofs of the professor's book, a longer-term interaction may exist because the professor is supervising the student's research, or the student will need letters of recommendation from the professor in the future for jobs where similar skills are relevant. These career concerns make immediate cash incentive less important. The student will do a good job for the sake of the implicit future rewards—closer attention to his research and better letters of recommendation. These aspects don't even have to be spelled out explicitly by the professor; everyone in the environment understands the larger game being played.

Repeated Relationships

Another aspect of ongoing employment is that the same worker performs similar actions repeatedly. Each time, there is an element of chance that makes the outcome an inaccurate indicator of the effort, so incentives cannot be perfect. But if luck at different times is independent, then by the law of large numbers the average output is a more accurate measure of average effort. This makes more powerful incentives possible. The idea is that the employer may believe the employee's hard luck story once; the claim of persistent bad luck is less plausible.

Efficiency Wages

You are considering hiring someone for a job in your firm. This job requires careful effort, and good work would be worth \$60,000 a year to you. The worker would rather take things easy; he suffers some mental or even physical cost in exerting effort. The worker's subjective valuation of this cost is \$8,000 per year.

You need to pay enough to entice the worker to your company and pay in a way that induces careful effort. There are dead-end jobs requiring no special effort that pay \$40,000. You need to beat that.

In terms of motivating effort, you cannot observe directly whether the worker is exerting careful effort or not. If he does not, there is some chance that

things will go wrong in a way you will notice. Suppose the probability of this is 25 percent. What incentive payment will induce the worker to put in the requisite effort?

You can offer a contract of the following kind: “I will pay you some amount above your other opportunities so long as no shirking on your part comes to light. But if that ever happens, I will fire you, and spread the word of your misbehavior among all the other employers, with the result that you will never earn anything more than the basic \$40,000 again.”

How high does the salary have to be so that the risk of losing it will deter the worker from cheating? Clearly, you will have to pay more than \$48,000. Otherwise, the worker will only take your job if he plans to shirk. The question is how much more. We will call the extra amount X , for a total salary of $\$48,000 + X$. That means the worker is X better off at your job compared to the alternative.

Suppose the worker does cheat one year. In that year he will not have to incur the subjective cost of effort, so he will have gained the equivalent of \$8,000. But he will run a 25 percent risk of being found out and losing $\$X$ this year and every year after that. Is the one-time gain of \$8,000 worth the prospective loss of $0.25X$ every year thereafter?

That depends on how money at different times is compared—that is, on the interest rate. Suppose the interest rate is 10%. Then getting an extra X annually is like owning a bond with a face value of $\$10X$ (which at 10% pays X annually). The immediate gain of the equivalent of \$8,000 should be compared with the 25 percent chance of losing $\$10X$. If $\$8,000 < 0.25 \times 10X$, the worker will calculate that he should not shirk. This means $\$X > \$8,000/2.5 = \$3,200$.

If you offer the worker an annual salary of $\$48,000 + \$3,200 = \$51,200$ so long as no shirking comes to light, he will not in fact shirk. It isn’t worth risking the extra \$3,200 forever in order to slack off and get a quick \$8,000 this year. And since good effort is worth \$60,000 per year to you, it is in your interest to offer this higher wage.

The purpose of the wage is to get the worker to put in the requisite effort and work more efficiently, and so it is called an *efficiency wage*. The excess above the basic wage elsewhere, which is \$11,200 in our example, is called the *efficiency premium*.

The principle behind the efficiency wage enters your daily life at many points. If you go to the same car mechanic regularly, you would do well to pay a little more than the lowest rate for the job. The prospect of the steady extra profit will deter the mechanic from cheating you.* You are paying him a premium, not for efficiency in this case but for honesty.

Multiple Tasks

Employees usually perform multiple tasks. To take an example close to home, professors teach and carry out research. In such cases, the incentives for the different tasks can interact. The overall effect depends on whether the tasks are substitutes (in the sense that when the worker devotes more effort to one task, the net productivity of effort on the other task suffers) or complements (in the sense that more effort to one task raises the net productivity of the effort devoted to the other task). Think of a farmhand working in the cornfields and in the dairy. The more he works on corn, the more tired he will get, and each hour he spends in the dairy will be less productive. Conversely, think of a farm worker who looks after beehives and apple orchards. The more effort he puts into bee-keeping, the more productive his effort in growing apples.

When tasks are substitutes, giving strong incentives to each kind of effort hurts the outcome of the other. Therefore both incentives have to be kept weaker than you would think if you considered each task in isolation. But when tasks are mutual complements, incentives to induce better effort on each task help the outcome of the other. Then the owner can make both incentives strong, to take advantage of these synergies without having to worry about any dysfunctional interactions.

This has ramifications for organization design. Suppose you want to get many tasks performed. You should try, to the extent possible, to assign them to the employees in such a way that each performs a set of tasks that complement one another. In the same way, a big enterprise should be structured into divisions, each of which is responsible for a subset of mutually complementary tasks, and substitute tasks should be assigned to different divisions. That way you can take advantage of strong incentives for each employee, and within each division.

The consequences of failing to follow this rule can be seen by everyone who has ever flown to or through London's Heathrow Airport. The function of an airport is to receive departing passengers at the curbside and deliver them to their airplanes, and receive arriving passengers from their airplanes and deliver them to their ground transport. All the activities that occur in each of these processes—check-in, security, shopping, and so on—are complements. Conversely, multiple airports serving a city are mutual substitutes (albeit not perfect ones—they differ in their locations relative to the city, the ground transport facilities to which they are connected, and so on). The principle of grouping together complementary activities and separating substitute activities says that all the functions within one airport should be under the control of one

management, and then different airports should compete with each other for the airlines' and the passengers' business.

The UK government has done exactly the opposite. All three airports serving London—Heathrow, Gatwick, and Stanstead—are owned and managed by one company, the British Airports Authority (BAA). But within each airport, different functions are controlled by different bodies—the BAA owns and leases the shopping areas, the police are in charge of security but the BAA provides the physical setup of the security checks, a regulatory agency sets landing fees, and so on. No wonder incentives are dysfunctional. The BAA profits from the leases on shops and so provides too little space for security checks, the regulator sets landing fees too low with the aim of benefiting consumers, but this causes too many airlines to choose Heathrow, which is nearer central London, and so on. Both authors have suffered from this, along with millions of other “users” of these airports.

Turning to an application even closer to the authors' experience, are teaching and research substitutes or complements? If they are substitutes, the two should be performed in separate institutions, as is done in France, where universities do mostly teaching and the research is done in specialized institutes. If they are complements, the optimal arrangement is to combine research and teaching within one institution, as is the case in major U.S. universities. The comparative success of these two organizational forms is evidence in favor of the complements case.

Are the multiple tasks performed within the new Department of Homeland Security substitutes or complements? In other words, is that department a good way to organize these activities? We do not know the answer, but it is surely a question deserving careful attention from the nation's top policymakers.

Competition between Workers

In many firms and other organizations, there are many people simultaneously performing similar or even identical tasks. Different shifts work the same assembly line, and investment fund managers deal with the same overall market conditions. The outcome of each task is a mixture of the person's effort, skill, and an element of chance. Because the tasks are similar and performed under similar conditions at the same time, the chance part is highly correlated across the workers: if one has good luck, it is likely that all the others have good luck, too. Then comparisons of the outcomes produced by different workers yield a good indication of the relative efforts (and skill) exerted by

them. In other words, the employer says to a worker who pleads bad luck to explain his poor outcome: “Then how come the others performed so much better?” In this situation, incentive schemes can be designed based on relative performance. Investment fund managers are ranked based on how well they do relative to their peers. In other cases, incentives are provided by a competition, with prizes for the top performers.

Consider the professor who employs a student to correct the proofs of his book. He can hire two students (who do not know each other) and split the work between them, but with some overlap of pages to be corrected by both of them. A student who finds only a few errors in the overlap segment will be proved to have been slacking if the other finds many more. So the payment can be based on the “relative performance” on the overlap, to sharpen the incentives. Of course the professor should not tell each student who the other is (else they can collude), nor should he tell either what pages overlap between them (else they will be careful with that part and careless with the rest).

Indeed, the inefficiency caused by the overlap can be more than offset by the improved incentives. This is one of the advantages of dual sourcing. Each supplier helps create the baseline by which to judge the other.

For this book, Barry distributed copies to students enrolled in Yale’s undergraduate game theory class. The reward was \$2/typo, but you had to be the first to find it. This obviously led to massive duplication of effort, but in this case, the students were reading the book as part of the course. While many students did well, the big money winner was Barry’s assistant, Catherine Pichotta. The reason she did best wasn’t just that she found the greatest number of typos. The reason was that, unlike the Yale students, she thought ahead and started at the back of the book.

Motivated Workers

We have assumed that workers don’t care about doing the job well for its own sake or about the employer’s success other than through how it directly affects their pay and careers. Many organizations attract workers who care about the job itself and about the success of the organization. This is especially true in nonprofits, health care, education, and some public sector agencies. It is also true of tasks that require innovation or creativity. More generally, people are intrinsically motivated when performing activities that improve their self-image and give them a sense of autonomy.

Going back to the example of the student reading proofs, a student willing to

do an academic-related job on campus for relatively low payment instead of more lucrative outside jobs, for example as a software consultant for local businesses, may be more genuinely interested in the subject of the book. Such a student has intrinsic motivation to do a good job of proofreading. And such a student is also more likely to want to become an academic, and therefore will be more aware of, and more strongly motivated by, the “career concerns” mentioned above.

Intrinsically rewarding tasks and do-good organizations need fewer or weaker material incentives. In fact, psychologists have found that the “extrinsic” monetary incentives can diminish the “intrinsic” incentives of workers in such settings. They come to feel that they are doing it just for the money, rather than for any warm glow that comes from helping others or from the achievement itself. And the existence of material penalties such as lower pay or dismissal for failure may undermine the enjoyment of doing a challenging or worthwhile task.

Uri Gneezy and Aldo Rustichini conducted an experiment where subjects were given fifty questions from an IQ test.² One group was asked to do the best they could. Another was given 3¢ per correct answer. A third group was rewarded with 30¢ per correct answer and a fourth was paid 90¢ per correct answer. As you might have predicted, the two groups being paid 30¢ and 90¢ both outperformed the ones with no bonus—on average, they got 34 questions right compared to 28. The surprise was that the group with only 3¢ payment did the worst of all, getting only 23 right on average. Once money enters the picture, it becomes the main motivation, and 3¢ just wasn’t enough. It may also have conveyed that the task wasn’t that important. Thus Gneezy and Rustichini conclude that you should offer significant financial rewards or none at all. Paying just a little might lead to the worst of all outcomes.

Hierarchical Organizations

Most organizations of any size have multiple tiers—companies have a hierarchy of shareholders, board of directors, top management, middle management, supervisors, and line workers. Each is the boss of those lower down in the hierarchy, and in charge of providing them with the appropriate incentives. In these situations, the boss of each level must be aware of the danger of strategic behavior among those lower down the hierarchy. For example, suppose the incentive scheme for a worker depends on the quality of the work as certified by the immediate supervisor. Then the supervisor passes shoddy work in order to meet a target to get his own bonus. The supervisor can’t punish the

worker without also hurting himself in the process. When the upper-level boss designs incentive schemes to mitigate such practices, the effect is generally to weaken incentives at those levels, because that reduces the potential benefit from deception and fraud.

Multiple Owners

In some organizations the control structure is not a pyramid. In places the pyramid gets inverted: one worker is responsible to several bosses. This happens even in private companies but is much more common in the public sector. Most public sector agencies have to answer to the executive, the legislature, the courts, the media, various lobbies, and so on.

The interests of these multiple owners are often imperfectly aligned or even totally opposed. Then each owner can try to undermine the incentive schemes of the others by placing offsetting features in his own scheme. For example, a regulatory agency may be in the executive branch while the Congress controls its annual budget; the Congress can threaten budget cuts when the agency responds more to the wishes of the executive. When the different bosses offset each other's incentives in this way, the effect is weakness of incentives in the aggregate.

Imagine that one parent gives a reward for good grades and the other a reward for success on the athletic field. Instead of working synergistically, each reward is likely to offset the other. The reason is that as the kid spends more time studying, this will take some time away from athletics and thus reduce the chance of getting the sports award. The expected gain from an extra hour hitting the books won't be, say, \$1, but \$1 minus the likely reduction in the sports prize. The two rewards might not totally offset each other, as the kid could spend more time studying and practicing with less time for sleeping and eating.

In fact, mathematical models show that the overall strength of incentives in such situations is inversely proportional to the number of different bosses. That may explain why it is so difficult to get anything done in international bodies like the United Nations and the World Trade Organization—all sovereign nations are separate bosses.

In the extreme situation where the owners' interests are totally opposed, aggregate incentives may be totally without power. That is like the admonition in the Bible: "No man can serve two masters...God and mammon."³ The idea is that interests of God and mammon are totally opposed; when the two are joint bosses, the incentives provided by one exactly cancel those provided by the

other.

HOW TO REWARD WORK EFFORT

We have illustrated the key elements of a well-designed incentive scheme. Now we want to flesh out some of these principles through more developed examples.

Imagine you are the owner of a high-tech company trying to develop and market a new computer chess game, Wizard 1.0. If you succeed, you will make a profit of \$200,000 from the sales. If you fail, you make nothing. Success or failure hinges on what your expert player-programmer does. She can either put her heart and soul into the work or just give it a routine shot. With high-quality effort, the chances of success are 80 percent, but for routine effort, the figure drops to 60 percent. Chess programmers can be hired for \$50,000, but they like to daydream and will give only their routine effort for this sum. For high-quality effort, you have to pay \$70,000. What should you do?

As shown in the following table, a routine effort will get you \$200,000 with a 60 percent chance, which comes out to \$120,000 on average. Subtracting the \$50,000 salary leaves an average profit of \$70,000. The corresponding calculation if you hire a high-effort expert is 80 percent of \$200,000 minus \$70,000—that is, \$90,000. Clearly you do better to hire a high-effort expert at the higher salary.

	Chance of success	Average revenue	Salary payments	Average profit
Low-quality effort	60%	\$120,000	\$50,000	\$70,000
High-quality effort	80%	\$160,000	\$70,000	\$90,000

But there is a problem. You can't tell by looking at the expert's work whether she is exerting routine effort or quality effort. The creative process is mysterious. The drawings on your programmer's pad may be the key to a great graphics display that will ensure the success of Wizard 1.0 or just doodles of pawns and bishops to accompany her daydreaming. Knowing that you can't tell the difference between routine effort and quality effort, what is to prevent the expert from accepting the salary of \$70,000 appropriate for high effort but exerting routine effort just the same? Even if the project fails, that can always be blamed on chance. After all, even with genuine quality effort, the project can fail 20 percent of the time.

When you can't observe the quality of effort, we know that you have to base your reward scheme on something you can observe. In the present instance the only thing that can be observed is the ultimate outcome, namely success or failure of the programming effort. This does have a link to effort, albeit an imperfect one—higher quality of effort means a greater chance of success. This link can be exploited to generate an incentive scheme.

What you do is offer the expert a remuneration that depends on the outcome: a larger sum upon success and a smaller sum in the event of failure. The difference, or the bonus for success, should be just enough to make it in the employee's own interest to provide high-quality effort. In this case, the bonus must be big enough so that the expert expects a high effort will raise her earnings by \$20,000, from \$50,000 to \$70,000. Hence the bonus for success has to be at least \$100,000: a 20 percent increase (from 60 to 80 percent) in the chance of getting a \$100,000 bonus provides the necessary \$20,000 expected payment for motivating high-quality effort.

We now know the bonus, but we don't know the base rate, the amount paid in the event of a failure. That needs a little calculation. Since even low effort has a 60 percent chance of success, the \$100,000 bonus provides an expected \$60,000 payment for low effort. This is \$10,000 more than the market requires.

Thus the base pay is—\$10,000. You should pay the employee \$90,000 for success, and she should pay you a fine of \$10,000 in the event of failure. Thus, with this incentive scheme, the programmer's incremental reward for success is \$100,000, the minimum necessary for inducing quality effort. The average payment to her is \$70,000 (an 80 percent chance of \$90,000 and a 20 percent chance of—\$10,000).

This pay scheme leaves you, the owner, an average profit of \$90,000 (an 80 percent chance of \$200,000, minus the average salary of \$70,000). Another way of saying this is that on average your revenue is \$160,000 and your average cost is what the worker expects to earn, namely \$70,000. This is exactly what you could have gotten if you could observe quality of effort by direct supervision. The incentive scheme has done a perfect job; the unobservability of effort hasn't made any difference.

In essence, this incentive scheme sells 50 percent of the firm to the programmer in exchange for \$10,000 and her effort.* Her net payments are then either \$90,000 or—\$10,000, and with so much riding on the outcome of the project it becomes in her interest to supply high-quality effort in order to increase the chance of success (and her profit share of \$100,000). The only difference between this contract and the fine/bonus scheme is in the name. While the name may matter, we see there is more than one way to achieve the same

effect.

But these solutions may not be possible, either because assessing a fine on an employee may not be legal or because the worker does not have sufficient capital to pay the \$10,000 for her 50 percent stake. What do you do then? The answer is to go as close to the fine solution or equity-sharing as you can. Since the minimum effective bonus is \$100,000, the worker gets \$100,000 in the event of success and nothing upon failure. Now the employee's average receipt is \$80,000, and your profit falls to \$80,000 (since your average revenue remains \$160,000). With equity-sharing, the worker has only her labor and no capital to invest in the project. But she still has to be given a 50 percent share to motivate her to supply high-quality effort. So the best you can do is sell her 50 percent of the company for her labor alone. The inability to enforce fines or get workers to invest their own capital means that the outcome is less good from your point of view—in this case, by \$10,000. Now the unobservability of effort makes a difference.

Another difficulty with the bonus scheme or equity sharing is the problem of risk. The worker's incentives arise from her taking a \$100,000 gamble. But this rather large risk may lead the employee to value her compensation at less than its average of \$70,000. In this case, the worker has to be compensated both for supplying high-quality effort and for bearing risk. The bigger the risk, the bigger the compensation. This extra compensation is another cost of a firm's inability to monitor its workers' efforts. Often the best solution is a compromise; risk is reduced by giving the worker less than ideal incentives and consequently, this motivates less than an ideal amount of effort.

In other instances you may have other indicators of the quality of effort, and you can and should use them when designing your incentive scheme. Perhaps the most interesting and common situation is one in which there are several such projects. Even though success is only an inexact statistical indicator of the quality of effort, it can be made more precise if there are more observations. There are two ways in which this can be done. If the same expert works for you on many projects, then you can keep a record of her string of successes and failures. You can be more confident in attributing repeated failure to poor effort rather than to chance. The greater accuracy of your inference allows you to design a better incentive scheme. The second possibility is that you have several experts working on related projects, and there is some correlation in the success or failure of the projects. If one expert fails while others around her are succeeding, then you can be more confident that she is a shirker and not just unlucky. Therefore rewards based on relative performance—in other words, prizes—will generate suitable incentives.

CASE STUDY: TREAT THEM LIKE ROYALTY

The typical way that authors get paid for writing a book is via a royalty arrangement. For every book sold, the author gets a certain percentage, something like 15% of the list price on hardcover sales and 10% for paperback. The author might also get an advance against future royalties. This advance is usually paid in parts; one part upon signing of the contract, another upon delivery (and acceptance) of the manuscript, and the rest upon publication. How does this payment system create the right incentives, and where might it create a wedge between the interests of the publishing house and those of the author? Is there a better way to pay authors?

Case Discussion

The only good author is a dead author.

—Patrick O'Connor

An editor is one who separates the wheat from the chaff and prints the chaff.

—Adlai Stevenson

As these quotes suggest, there are many possible sources of tension in the relationship between authors and publishers. The contract helps resolve some of the problems and creates others. Holding back some of the advance gives the author an incentive to complete the book on time. The advance also transfers risk from the author to the publisher, who might be in a better position to spread the risk over a large number of projects. The size of the advance is also a credible signal that the publisher is truly excited about the prospects for the book. Any publisher can say that they love the book proposal, but actually offering a large advance will be much more costly if you don't believe that the book will sell a large number of copies.

TRIP TO THE GYM NO. 10

How big is the wedge between publishers and authors? Try to estimate how much more the publisher would like to charge compared to the

author.

One place where authors and publishers disagree is over the list price of the book. You might at first think that since authors are getting a percentage of the list price, they would want the price to be high. But what authors are really getting is a percentage of total revenue, say 15% in the case of the hardcover sales. Thus what authors really care about is total revenue. They would like the publishing house to pick a list price that maximizes total revenue.

The publisher, on the other hand, seeks to maximize its profits. Profits are revenue net of cost. What that means is that the publisher always wants to charge a higher price than would maximize revenue. If the publisher were to start at the revenue-maximizing price and go up a little bit, that would keep revenue almost constant but would reduce sales and thus cut costs. In our case, we anticipated this issue in advance and negotiated the list price as part of the contract. You are welcome. And thanks for reading the book.

There are two more cases on incentives in the next chapter: “Bay Bridge” and “But One Life to Lay Down for Your Country.”