IS IT RIGHT TO REWARD?

21 pages

By Alfie Kohn (1993) from *Punished by Rewards: The Touble with Gold Stars, Incentive Plans, A's, Praise, and Other Bribes*

The interest of the behaviorist in man's doings is more than the interest of the spectator — he wants to control man's reactions as physical scientists want to control and manipulate other natural phenomena.

— John B. Watson,

Behaviorism

What a fascinating thing! Total control of a living organism!

— B. F. Skinner, 1983

When two people find themselves at odds over an issue like capital punishment, the disagreement may concern the intrinsic right ness or wrongness of the policy as opposed to its empirical effects. An opponent of the death penalty may argue, for example, that there is something offensive about the idea of killing people in the name of justice. Evidence regarding the effect of executions on the crime rate probably would not be seen as relevant to this objection.

The same distinction can be made with respect to a discussion about pop behaviorism. Separate from the question of whether rewards do what we want them to do is the question of whether there is something fitting or troubling about their use. Some believe it is inherently desirable to give rewards, that people ought to get something for what they do quite apart from the consequences this may bring. Others believe there is something objectionable about the whole idea of giving rewards. Lest these opposing values get buried under a mound of studies (and become confused with factual findings), this chapter will care fully examine each of them in turn.

*This position, I should note, was not taken by Skinner, nor is it offered as a rule by other behaviorists.

Saving Room for Just Deserts

It is an integral part of the American myth that anyone who sets his mind to it can succeed, that diligence eventually pays off. It seems to follow, then, that people who do not succeed can be held responsible for their failure. Failure, after all, is prima facie evidence of not having tried hard enough. This doctrine has special appeal for those who are doing well, first because it allows them to think their blessings are deserved, and second because it spares them from having to feel too guilty about (or take any responsibility for) those who have much less.

The belief that rewards will be distributed fairly, even if it takes until the next lifetime to settle accounts, is one component of what is sometimes referred to as the "just world" view. Social psychologists have found that those who hold this position are indeed likely to assume that apparently innocent victims must have done something to deserve their fate; to face the fact that suffering is visited upon innocent people is, of course, to

recognize that the world is not particularly just at all. It does not take much imagination to see where this sort of thinking can lead: one group of children, after watching a film about the Nazis, were reported to have said, "But the Jews must have been guilty or they wouldn't have been punished like that."

The belief I have been describing can be summarized as follows: deserving people will be rewarded. Underpinning this idea is an even more basic and widely held premise: deserving people *should* be rewarded. In theory, these two views can be separated, but in practice the latter often drives the former. Many people assume, at least with respect to important issues, that things eventually work out the way they ought to. It is terribly unsettling, after all, to acknowledge that our society, much less life itself, is not especially fair. The sheer wish that it were can produce in some people a belief that things are, or in time will become, what they should be.

Let us look a little more closely at the idea that rewards should be bestowed on those who merit them. For many people, the moralistic corollary to this assumption is that bad things should be bestowed on, or good things withheld from, those who are undeserving. Many of us have watched people become uneasy, if not positively furious, when they believe some offense — including one committed by a child — has not been punished severely enough. Later in this book I will argue that a child's misbehavior is best construed as a "teachable moment," a problem to be solved together rather than an infraction that calls for a punitive response. I will try to show that this approach is not only more respectful and humane but also much more effective over the long haul at helping children develop a sense of responsibility. But I have seen people brush aside such arguments, sometimes becoming visibly disturbed at the prospect that a miscreant not have to suffer any consequences for her action. *Consequences* may be a code word for punishment, and punishment may produce resentment rather than responsibility, but never mind. The important thing, on this view, is that Justice is served, and cosmic balance restored, by cracking down on a wrongdoer.*

[*A popular way to express this position is to say that the offender must be made to "pay" for what he did — a locution suggesting that we often conceive of justice in economic terms.]

The entirely reasonable ideal of personal responsibility has been transformed in our culture into a terror of permissiveness that extends beyond child-rearing to a general fear of social laxity. We see it in outraged reactions to prisons that are judged too comfortable, or even to organizations that compensate employees on any basis other than achievement. When pay is not conditioned on performance we are sometimes said to be rewarding incompetence (or laziness) and giving some people a free lunch — a prospect that sends shudders through executive dining rooms.

When stripped of this harsh rigidity, of course, the basic idea that people should get what they deserve, which social scientists refer to as the equity principle, seems unremarkable and, indeed, so intuitively plausible as to serve for many people virtually as a definition of fairness Rarely do we even think to question the idea that what you put in should determine what you take out.

But the value of the equity principle is not nearly as self-evident as it may seem. Once we stop to examine it, questions immediately arise as to what constitutes deservingness. Do we reward on the basis of how much effort is expended (work hard, get more goodies)? What if the result of hard work is failure? Does it make more sense, then, to reward on the basis of success (do well, get more goodies)? But "do well" by whose standards? And who is responsible for the success? Excellence is often the product of cooperation, and even individual achievement typically is built on the work of other people's earlier efforts. So who "deserves" the reward when lots of people had a hand in the performance?

These questions lead us gradually to the recognition that equity is only one of several ways to distribute resources. It is also possible for each person to receive an equal share of the goods — or for need to determine who gets what. Different circumstances seem to call for different criteria. Few school principals hand out more supplies to the teachers who stayed up longer the night before to finish a lesson plan; rather, they look at the size and requirements of each class. Few parents decide how much dinner to serve to each of their children on the basis of who did more for the household that day. Few policymakers and moral theorists, struggling with the knotty question of how to distribute scarce health care resources, automatically assume that the most productive contributors to society (whatever that means) should get the most care.

In short, the equity model, as social psychologist Melvin Lerner put it, "applies to only a limited range of the social encounters that are affected by the desire for justice." Specifically, it is the favored mode of "impersonal, economic relations." To assume that fairness always requires that people should get what they "earn" — that the law of the marketplace is the same thing as justice — is a very dubious proposition indeed. What's more, as Morton Deutsch warns, "the danger of conceiving of personal relations in terms appropriate to marketplace exchanges is that it hastens the depersonalization of personal relations by fostering the intrusion of economic values into such relations."

Just as important as the realization that principles other than equity could legitimately be invoked in many situations is the fact that principles other than equity are invoked in many situations. If we want to predict how people will choose to distribute resources, the most important thing we need to know is what kind of relationship exists among those involved. The equity principle, not surprisingly, is more likely to be the first choice of strangers. (This is why it is a little suspicious that assumptions about the universality of that principle are largely based on contrived experiments in which the subjects have never met each other before.) Other factors also help to determine which principle is used. For example, cultural background matters: where people are accustomed to thinking in communal rather than individualistic terms, they are more likely to distribute rewards equally rather than on the basis of who performed better. Women are more likely than men to share this preference for equality as the basis for distribution. Finally, there are differences on the basis of individual personality. It is interesting to reflect on what kind of person might be expected to insist that what someone gets must be based on what he produced.

Edward E. Sampson, a psychologist who frequently writes about American culture, observed that we have been led to "take equity as the natural state and deviations from it as unnatural." However, the assumption that people should be rewarded on the basis of

what they have done is "not as much a psychological law about human nature as it is a psychological outcome of a culture's socialization practices." This doesn't mean that it is impossible to defend the view that people who have done something should be rewarded; rather, it suggests that this view must be defended, as opposed to taken for granted as obviously true.

To this point, I have been referring to rewards as resources to be distributed, which may be an appropriate way to think about, say, what to do with a company's profit at the end of the year. But this does not accurately describe many other kinds of rewards, such as grades or gold stars or praise. Many goodies have been invented for the express purpose of rewarding certain kinds of behavior. If the equity model applies here, it cannot be assumed on the basis of rules for deciding how much to pay employees.

Not long ago, I heard a teacher in Missouri justify the practice of handing out stickers to her young students on the grounds that the children had "earned" them. This claim struck me as an attempt to deflect attention away from — perhaps to escape responsibility for — the decision she had made to frame learning as something one does in exchange for a prize rather than as something intrinsically valuable. How many stickers does a flawless spelling assignment merit? One? Ten? Why not a dollar? Or a hundred dollars? After the fact, one could claim that any reward was "earned" by the performance (or performer), but since these are not needed goods that must be handed out according to one principle or another, we must eventually recognize not only that the size of the reward is arbitrarily determined by the teacher but that the decision to give any reward reflects a theory of learning more than a theory of justice.

When such individuals are pressed on their insistence that it is simply right to reward people for what they do, it sometimes turns out that their real concern is with the results they fear would follow the abolition of rewards. One business consultant, for example, writes that he was horrified to learn about a company that allocated bonuses equally to all employees; "mediocrity would receive the same rewards as excellence," he exclaims. But as we read on, we find that what at first appears to reflect a moral stance (you should pay for what you get) is ultimately based on expectations about the consequences (you'll get what you pay for) — a very different sort of objection. His suspicion is that workers will come to ask, "Why work harder" if there is no tangible benefit to doing so? This, of course, is a question that can be addressed by looking at evidence about what actually motivates people and what happens when rewards are (and are not) used.

Treating People Like Pets

As behaviorists cheerfully admit, theories about rewards and various practical programs of behavior modification are mostly based on work with rats and pigeons. The underlying assumption, according to one critic, seems to be that "the semistarved rat in the box, with virtually nothing to do but press on a lever for food, captures the essence of virtually all human behavior.'"

But it is not only researchers who make this assumption. We join them in taking "one giant leap toward mankind'" when we import the principles and techniques used to train the family pet to the realm of raising children. The way we sometimes talk about (or to) our daughters and Sons reflects a view of parent-child relationships quite congenial to a committed behaviorist. Discussions about how to "handle" our kids are a case in point;

on reflection, this seems a rather peculiar verb to use in the context of a relationship with another human being. Likewise, when we call out a hearty "Good girl!" in response to a child's performance, the most appropriate reply would seem to be "Woof!" With respect to the workplace or public policy, we talk casually about the use of "carrots and sticks," and there is food for thought here, too. Before these words came to be used as generic representations of bribes and threats, what actually stood between the carrot and the stick was, of course, a jackass.

Presumably most of us do not intend to compare ourselves — or more precisely, the people to whom we are administering these inducements — to poodles or donkeys. Surely we know that human beings can reflect on rewards and develop complicated expectations and opinions about them (and about the activities for which they are being dispensed) in a way that animals cannot. Yet it is not an accident that the theory behind "Do this and you'll get that" derives from work with other species, or that behavior management is frequently described in words better suited to animals.

My claim is that pop behaviorism is by its very nature dehumanizing. But I do not mean by that word merely that we are treated or understood as being on a par with other species, this is just a symptom. In the case of Skinnerian theory, the human self has been yanked up by its roots and the person reduced to a repertoire of behaviors. It is hard to imagine what could be more dehumanizing than the removal of what defines us as human. In fact, even to suggest that we learn or work only in order to obtain rewards — an assumption held by behaviorists less extreme than Skinner — is not only inaccurate but demeaning as well.

Some observers think that to manipulate workers with incentives is to treat them like children. In a way this is true, but there is some thing problematic about treating people of any age this way. For other critics, the more apt comparison is to how we train animals. But again, this characterization does not go far enough because the assumption that an organism's behavior is wholly dependent on, and controlled by, reinforcements has been shown to be inaccurate even for rodents. Perhaps, then, as sociologist William Foote Whyte proposed, what reward systems finally suggest is an implicit comparison to nonliving things:

> Management also seems to assume that machines and workers are alike in that they are both normally passive agents who must be stimulated by management in order to go into action. In the case of the machines, management turns on the electricity. In the case of workers, money takes the place of electricity.

The behaviorist's conception of humans as passive beings whose behavior must be elicited by external motivation in the form of incentives is, by any measure, outdated. Although the work done by some modern psychologists continues to rely implicitly on this assumption, more and more researchers have come to recognize that we are beings who possess natural curiosity about ourselves and our environment, who search for and overcome challenges, who try to master skills and attain competence, and who seek to reach new levels of complexity in what we learn and do. This is more true of some people than others, of course, and in the presence of a threatening or deadening environment, any of us may retreat to a strategy of damage control and minimal effort.

But in general we act on the environment as much as we are acted on by it, and we do not do so simply in order to receive a reward.

Within the discipline of psychology, the passive-organism view has faded along with the influence of behavior theory itself. But in everyday life, in the workplace and the classroom and the home, this view continues to make its presence felt through the practices of pop behaviorism. To put this the other way around, our everyday practices rest on an implicit theory of human nature that fails to do us justice. When we repeatedly promise rewards to children for acting responsibly, or to students for making an effort to learn something new, or to employees for doing quality work, we are assuming that they could not or would not choose to act this way on their own. If the capacity for responsible action, the natural love of learning, and the desire to do good work are already part of who we are, then the tacit assumption to the contrary can fairly be described as dehumanizing. *

The underlying theory of human nature, however, is not the only reason that handing out rewards (or, for that matter, punishments) is dehumanizing. That description also seems to apply because the practice is, at its core, neither more nor less than a way of trying to control people. Now there are circumstances, especially where children are involved, in which it is difficult to imagine eliminating all vestiges of control. (I will say more about this later.) But anyone who is troubled by a model of human relationship founded principally on the idea of one person controlling another must ponder whether rewards are as innocuous as they are sometimes made out to be.

Clearly, punishments are harsher and more overt; there is no getting around the intent to control in "Do this or else here's what will happen to you." But rewards simply "control through seduction rather than force. In the final analysis, they are not one bit less controlling since, like punishments, they are "typically used to induce or pressure people to do things they would not freely do" — or rather, things that the controller believes they would not freely do. This is why one of the most important (and unsettling) things we can recognize is that the real choice for us is not between rewards and punishments but between either version of behavioral manipulation, on the one hand, and an approach that does not rely on control, on the other.

[*To the extent that we sometimes do seem to be driven by rewards, this may be attributed, at least in part, to the way pop behaviorism creates a dependence on itself (seepage 17).]

In the workplace, there is no getting around the fact that "the basic purpose of merit pay is manipulative." One observer more bluntly characterizes incentives as "demeaning" since the message they really convey is, "Please big daddy boss and you will receive the rewards that the boss deems appropriate." The use of treats at home, meanwhile, offers exactly the same message, except here "big daddy boss" may literally be Big Daddy.

Sometimes the controlling nature of rewards is too obvious to miss. Consider the so-called token economy, which is used primarily with captive, dependent populations such as patients in psychiatric hospitals or children in school. The idea — again, explicitly derived from work with laboratory animals — is that when the people in charge notice the patients or children engaging in the "correct" sort of behaviors, chips or other markers are handed out that can be exchanged later for privileges or treats. Even at the height of

their popularity,* these programs offended a number of people for a number of reasons. But specific objections aside, it is difficult to imagine a more flagrant example of control than one person' giving another a token redeemable for candy or privileges to reward him for being "cooperative."

[*The first institutional token economies in the United States were developed in the 1960s; after being in vogue for perhaps a decade, their use declined to the point that few hospitals now have them in place. How widely these programs are still used in schools is difficult to determine, but probably the single most popular program of classroom management (Assertive Discipline) is clearly a behavior modification plan even if it is not strictly a token economy.]

We don't need critics of this approach to teach us this, however; the reliance on crude control is a point made more convincingly (albeit unintentionally) by the vocal proponents of token economies. In an article for school psychologists, a pioneer of such plans writes that "children need to be reminded frequently that they are working for reinforcers" and that "a teacher must always keep in mind that the teacher is the manager of the classroom." If a child is sneaky enough to save up tokens rather than feeling driven to keep earning new ones, we are warned that "the child and not the teacher is in control" of her behavior (a prospect evidently regarded as appalling on its face). Any complaints from children who think the administration of the rewards is unfair "can be easily handled by ignoring or redirecting"; if such concerns are "simply not reinforce[d] . . . they will extinguish."

Just as threats are simply a more blatant version of control than bribes are, so token economies merely exaggerate the manipulation that describes other, less systematic, applications of rewards. The point to be emphasized is that all rewards, by virtue of *being* rewards, are not attempts to influence or persuade or solve problems together, but simply to control. In fact, if a task is undertaken in response to the contingency set up by the rewarder, "the person's initial *action* in choosing the task is constrained."

This feature of rewards is much easier to understand when we are being controlled than when we are doing the controlling. This is why it is so important to imagine ourselves in the other position, to take the perspective of the person whose behavior we are manipulating. It is easy for a teacher to object to a program of merit pay — to see how patronizing it is to be bribed with extra money for doing what some administrator decides is a good job. It takes more effort for the teacher to see how the very same is true of grades or offers of extra recess when she becomes the controller. Exactly the same is true of the worker, chafing under the burden of a manipulative compensation plan, who comes home and manipulates his child with a Skinnerian system that differs only in the type of reward.

By definition, it would seem, if one person controls another, the two individuals have unequal status. The use of rewards (or punishments) is facilitated by this lack of symmetry but also acts to perpetuate it. Naturally, the impact of this fact is different in the relationship between two adults than it is in that between an adult and a child, but the fact itself is worth pondering. If you doubt that rewarding some one emphasizes the rewarder's position of greater power, imagine that you have given your next-door neighbor a ride downtown, or some help moving a piece of furniture, and that he then

offers you five dollars for your trouble. If you feel insulted by the gesture, consider why this should be, what the payment implies. Again, this feeling of resentment in response to the status differential between giver and receiver should be kept firmly in mind when the roles have been reversed and we are the ones doing the rewarding. *

[*Another example was provided by accounts of a magazine editor who, after disagreements with others on the staff, was given to handing out bottles of wine or gift certificates. This, according to one former associate editor, "made us feel like a McDonald's Employee of the Month." Perhaps we should ask why anyone, even employee at McDonald's, should be made to feel that way.]

If rewards not only reflect differences in power but also contribute to them, it should not be surprising that their use may benefit the more powerful party — that is, the rewarder. This point would seem almost too obvious to bother mentioning except for the fact that, in practice, rewards are typically justified as being in the interests of the individuals receiving them. We claim to reinforce people to teach them things that they need to be taught. But one writer, after ticking off the specific objectives of behavior modification programs, asks, "In whose interests is it for a prisoner, a student, or a patient to be less complaining, more attentive, submissive, and willing to work?" Who really benefits when a child quiets down and sits still?

To be sure, some behavior managers — the parent who tries to reinforce a child's display of good values, or the teacher who attempts to interest students in doing research by offering extra credit for a report — may genuinely be concerned to improve the lives of those they reward. In such cases we can proceed directly to ask whether these rewards have the intended effects. But it is possible that others who use rewards or punishments are being less than honest with themselves when they insist they are only trying to help whomever they are controlling. It may be their own convenience (or continued capacity to control) that is really at stake.

Cui bono? — Who benefits? — is always a useful question to ask about a deeply entrenched and widely accepted practice. In this case, it is not merely the individual rewarder who comes out ahead; it is the institution, the social practice, the status quo that is preserved by the control of people's behavior. A pair of psychologists who reviewed token economies and similar plans in classroom settings observed that those peddling such systems "have used their procedures to serve the goals and values of the existing school system." More generally, these psychologists encourage us to ask, "To what extent is behavior modification. . . helping the existing institutional system achieve its present goals, e.g., goals of control for the sake of control, order and (misleading) tranquility, thus preventing rather than producing needed change?"

But we do not have to rely on critics to make this point. The father of behaviorism, John Watson, made it himself, candidly acknowledging that he and his colleagues are constantly manipulating stimuli, dangling this, that and the other combination in front of the human being in order to determine the reactions they will bring forth — hoping. that the reaction will be "in line with progress," "desirable," "good." (And society really means by "desirable," "good," "in line with progress," reactions that will not disturb its recognized and established traditional order of things.)

If rewards bolster the traditional order of things, then the psychologist Mihaly Csikszentmihalyi is right to warn (or promise) that "to deemphasize conventional rewards threatens the existing power structure." The thrust of a book like this one is rightly viewed as political insofar as it raises questions about systems that support the status quo, but it is no more political than pop behaviorism, which turns out to be a profoundly conservative doctrine posing as a value-free technique.

In one sense, this conclusion is paradoxical: most people who call themselves conservative emphasize the agency and responsibility of the individual, whereas Skinner spent his life denying the idea of choice and urging us to control reinforcers in the environment since they, in turn, control us. * Indeed, some social reformers have been enamored of the behaviorist vision (derived by Watson from John Locke) of human beings as blank slates that can be written on as one pleases.

[* According to Skinner, the reason we have no cause to fear abuse by behaviorists and their surrogates — the reason Walden Two will not come to resemble the world of Nineteen Eighty-four — is that the Chief Reinforcer "doesn't control others; he designs a world that controls others." This line of reasoning has not reassured everyone.]

Paradoxical or not, though, it is difficult to deny the conservatism of behavior control. In a very practical sense, applied behaviorists are anxious not to offend their clients, and they therefore offer a system of control that helps sustain the institutions and programs that currently exist. But I am making a more fundamental point. While it may seem that reward-and-punishment strategies are inherently neutral, that any sort of behavior could, in principle, be encouraged or discouraged, this is not completely true. If it were, the fact that these strategies are invariably used to promote order and obedience would have to be explained as a remarkable coincidence.

More realistically, we must acknowledge that because pop behaviorism is fundamentally a means of controlling people, it is by its nature inimical to democracy, critical questioning, and the free exchange of ideas among equal participants. Rewarding people for making changes in the existing order (which might include the very order that allows some individuals to be controllers and others controlled) is not merely unlikely but a contradiction in terms. 'The master's tools will never dismantle the master's house," as one writer put it. This point may seem far removed from the act of promising a child a weekend trip to the zoo if he gets to bed on time all week, but that is all the more reason to think seriously about the connection, which is to say, about the implications of any system of control.

Before concluding this discussion, I want to consider three objections that might be offered to these remarks concerning control. The first is the relatively modest point that it may be misleading to speak of rewards as inherently controlling, since some rewards are more controlling than others. To some extent this is true. In fact, I would expand on this comment by proposing a rough guideline for determining degrees of control: we need to look at the intention of the rewarder, the perception of the rewardee, and various characteristics of the reward itself.

Suppose, therefore, that we wish to present someone with a reward but also to reduce the extent to which the transaction is controlling. The first step is to examine our own motivation: are we ultimately trying to teach a skill, promote a value, boost self-esteem,

or are we mostly interested in making someone do what we want? Next, we might try to put ourselves in the reward recipient's shoes and imagine whether she might *feel* manipulated, irrespective of our intentions. (An expression of positive feedback might be construed as useful information by one person and as a clever attempt to control what she does tomorrow by another.) Finally, we ought to look at various objective features of the reward experience — how much emphasis the incentive has been given, how large or attractive it is, how closely it is tied to the quality of performance, and so on — with an eye to minimizing the extent to which the recipient will see the reward as driving his actions.

While I think it makes sense to attend to these features, we have to be careful about assuming that doing so can eliminate the problem entirely. Inflating the manipulative features of a reward may make a bad thing worse, but there is no getting around the fact that any time we say "Do this and you'll get that," we are attempting to control the behavior of the person we are addressing. Probably all a rewarder can do is minimize — or more disturbing, distract people from noticing — what is really going on.

This may suggest a second, more sweeping objection, one that has been made not only by Skinner and Skinnerians but also by social theorists with whom they have little in common: control is an unavoidable feature of human relationships; all that actually varies is the subtlety of the system of reinforcement. A brief smile and nod are just as controlling as a dollar bill — more so, perhaps, since social rewards may have a more enduring effect than tangible rewards. Just because we cannot readily identify the operative operant doesn't mean it isn't there.

The introduction to a book entitled *Man Controlled* pretty well captures this perspective. Those who raise concerns about what the title suggests, we are told, simply have a "fear of new knowledge" that has been cultivated by "alarmists." Realists recognize that "the technology of behavior control is not good or bad, but neutral" — therefore "not even an issue" — for the simple reason that there is no freedom (in fact, this word appears only within quotation marks) to be lost. Whether we like it or not, "all behavior is controlled. . . . The world is, in a sense, one large 'Skinner box."

Here, I think, we have to be very careful to tease apart what are actually quite different claims. That subtler reinforcers too can be controlling is quite true; in fact, I argued earlier that rewards are just as controlling as punishments, and delicate rewards as controlling as heavy-handed ones. But to conclude from this that all human interaction is therefore best described as an exercise in control seems to me a grave mistake. People who believe this have done one of two things: either they have taken on faith that selfhood and choice are illusions and we do only what we have been reinforced for doing, or they have stretched the word *control* until it encompasses many other kinds of interaction, such as trying to convince someone of the value of one's point of view. At this point, the word has become so broad and imprecise as to be of little use. If I decide that whenever two people talk about something, each is really trying to control the other, then this is true only because of my rather contrived definition of *control*, which doesn't help us understand very much.

A far more defensible position, it seems to me, is that some forms of human interaction are controlling and some are not. The line might not be easy to draw in practice, but the

distinction is still meaningful and important. Consider an analogy: the line between truth and falsehood isn't always easy to draw either (as in the case where a possibly relevant true statement is omitted). Likewise, lots of people tell small lies. But we are not entitled to conclude from this that all human communication is fraudulent and that it makes no sense to oppose categories of talk that are inherently misleading. My point, of course, is that we can say the same about control.

The last objection we might anticipate is that even if it is possible to avoid controlling other people, control is sometimes an appropriate, even desirable, mode of interaction, whether we use rewards or some other technique. It could even be argued that parents who fail to control their children are not living up to their responsibilities

I will have a number of things to say on this subject in chapters 9 and 12, but a few words at this point may not be out of place. To begin with, when people talk about the need to control children, they very often mean that children cannot be left entirely to their own devices. It is hard to imagine how anyone could disagree with this. But to say that children need structure or guidance is very different from saying they have to be controlled. In everyday conversation, we tend to confuse these different kinds of interventions; once we clarify them, it is not clear that many situations actually call for measures that most of us would see as controlling or manipulative.

I think it is true that very young children may sometimes require controlling responses; the bottom line is that a three-year-old cannot be permitted to toddle out into the street at will. But before we resort to control, we should be absolutely certain that less intrusive, more respectful, interventions cannot work. We should also think about how an act of control is exercised: Do we justify it with a reasonable explanation ("Cars come by here very fast sometimes and I love you so much that I have to make sure you don't get hurt")? Do we pause to ask whether what we are getting the child to do (or stop doing) is really necessary? Are we thinking about how best to help the child become a responsible person (as opposed to just how to get her to obey)?

Parents and teachers who defend the use of control without reservation do not, as a rule, pause to ask these sorts of questions. Often this is because they see the world in dichotomous terms: either you are controlling or you are permissive; either you crack down hard or you let kids get away with anything. To devise flexible and reasonable rules for children, preferably by working with them to solve problems rather than imposing these rules on them, is very different from control on the one hand and a laissez-faire approach on the other.

Some who support more coercive strategies assume that children will run wild if they are not controlled. However, the children for whom this is true typically turn out to be those accustomed to being controlled — those who are not trusted, given explanations, encouraged to think for themselves, helped to develop and internalize good values, and so on. Control breeds the need for more control, which then is used to justify the use of control.

The thoughtful pursuit of reasonable ends requires far less use of control for children, let alone adults, than many of us assume. If someone persists in controlling others, something else may be at work — a set of values and a view of relationship that no argument or evidence will suffice to challenge. In the end, we may just have to take our

stand with one or another vision of human life. Hugh Lacey and Barry Schwartz put it well:

> There has always been a moral impulse behind Skinner's driven and unrelenting commitment to behaviorism, a belief that the implementation of systematic behavioral controls will contribute quickly to solving the big social problems of the modern world. And much of his philosophical writing has been devoted to arguing that persons are the kind of beings defined by relations of control. We too have a moral motive. It is that relations of dialogue in all aspects of life are better for everyone than relations of control.

The thrust of this chapter, then, has been that giving people rewards is not an obviously fair or appropriate practice across all situations; to the contrary, it is an inherently objectionable way of reaching our goals by virtue of its status as a means of controlling others. Some readers will respond to this by saying that regardless of whether rewards are good, bad, or neutral from a moral point of view, the most important reason we use them is that they work.

Is it Effective to Reward?

[Rewards] have effects that interfere with performance in ways that we are only beginning to understand. – Janet Spence, 1971

Leon is on his way out the door to take a walk when Pam calls out to him. If you help me clean the kitchen this afternoon, she says, I'll take you to your favorite restaurant tonight. Leon closes the door and finds a sponge.

On Nora's list of favorite activities, working on math homework ranks just below having a root canal. So Phil announces that if she finishes the problem set on page 228 before eight o'clock, he will give her five dollars. Nora pulls out her book.

What has happened here? Both Leon and Nora complied with someone else's wishes, engaging in an activity they were otherwise not planning to do (at least not at the moment) in order to obtain something they valued. In each case, one person used a reward to change another's behavior. The plan worked, and that, most of us say, is all we need to know.

But let's probe further. Rewards are often successful at increasing the probability that we will do something. At the same time, though, as I will try to show in this chapter and the two that follow, they also change the *way* we do it. They offer one particular *reason* for doing it, sometimes displacing other possible motivations. And they change the *attitude* we take toward the activity. In each case, by any reasonable measure, the change is for the worse. Most behaviorists are not fond of punishment; as one Skinnerian has written, the trouble "may be not that it doesn't work but that it works only too well." I think exactly the same thing can be said of rewards: we pay a substantial price for their success.

However, even that statement concedes too much because the success of rewards is, in truth, widely misunderstood and vastly overrated. Here is where we begin our investigation, then: with a close look at the belief that rewarding people produces changes

in behavior, and (in the following section) at the belief that rewards improve performance on a variety of tasks.

Do Rewards Change Behavior?

To examine the claim that rewards are effective at altering behavior, we pose three questions: First, for whom are they effective? Second, for how long are they effective? And third, at what, exactly, are they effective? (I have already hinted at a fourth question At what cost are they effective? — but this we will set aside for the time being.)

1. **For whom are rewards effective?** Perhaps we should instead ask, For what ... ? given that their most impressive successes take place in animal laboratories. But when we look at the probability that rewards will change human behavior, a pattern begins to emerge, as two management specialists have noticed:

> Many of the early (and highly successful) applications of the principles of behavior modification have involved animals (such as pigeons), children, or institutionalized adults such as prisoners or mental patients. Individuals in each of these groups are *necessarily* dependent on powerful others for many of the things they most want and need, and their behavior usually can be shaped with relative ease.

Notice that this is not a moral objection; it is a statement of fact about how behavior is easier to control when the organism you are controlling is already dependent on you. In part, this is true because a dependent organism can be kept in a state of need. Laboratory animals are typically underfed to ensure their responsiveness to the food used as a reinforcer. Likewise, "in order to make people behave in a particular way ... they must be ... needy enough so that rewards reinforce the desired behavior." People who have some degree of independence, such as Leon, will also respond to rewards on occasion, but it is more difficult to make this happen in a predictable, systematic way.

2. **For how long are rewards effective?** The short answer is that they work best in the short term. For behavior changes to last, it is usually necessary to keep the rewards coming. Assuming your child is reinforced by candy, you can induce him to clean up his room for as long as you keep providing sweets. In practice, however, this raises several problems. What if he becomes satiated with sugar so that the reward eventually stops being rewarding to him?* Alternatively, what if his demands to be paid off escalate (in frequency if not in quantity) beyond your desire or ability to meet them? Most important, do you really want him to help out around the house only as long as you have a supply of M&M's on hand?

[*The same point applies, of course, to other rewards, such as grades.]

In the real world, even if not in the laboratory, *rewards must be judged on whether they lead to lasting change --- change that persists when there are no longer any goodies to be gained.* This is the key question to pose to a manager who claims that performance in her division jumped after an incentive plan was introduced, or to a teacher who brags that his students read more books when they are given treats for doing so. We want to know what happens to productivity, or to the desire to read, once the goodies have run out.

In theory it is possible to keep handing out reward pellets forever. In practice, though, this is usually impractical, if not impossible, to sustain. What's more, most people with an interest in seeing some behavior change would say it is intrinsically better to have that change take root so that rewards are no longer necessary to maintain it. Even behaviorists generally accept this criterion. Virtually every behavior for which children are rewarded, from brushing their teeth to acting altruistically, is something we'd like them to keep doing when they are no longer rewarded. At work, people usually continue to get paid for what they do, but if the goal is to help people change their behavior — for example, by improving the quality of their work — a continued dependence on rewards can create a range of practical problems, including an increase in demands (for money rather than M&M's), as managers trying to implement incentive programs on a permanent basis have discovered.

If it does make sense to measure the effectiveness of rewards on the basis of whether they produce lasting change, the research suggests that they fail miserably. This news should not be shocking; most of us, after reflecting carefully; will concede that our own experience bears this out. However, what is not always recognized is, first, just how utterly unsuccessful rewards really are across various situations, and second, just how devastating an indictment is contained in this fact.

Token economy programs To start with, let us consider elaborate behavior modification plans such as token economies (where markers that can be redeemed for privileges or treats are dispensed when people act "appropriately"). Theoretically, these programs should have unusually high prospects for success since they are typically implemented in laboratory-like settings — closed environments with dependent subjects. * In the first systematic review of the research on token economies, conducted in 1972, two avid proponents of the idea stated that:

> the generalization of treatment effects to stimulus conditions in which token reinforcement is not given might be expected to be the *raison d'être* of token economies. An examination of the literature leads to a different conclusion. There are numerous reports of token programs showing behavior change only while contingent token reinforcement is being delivered. Generally, removal of token reinforcement results in decrements in desirable responses and a return to baseline or near- baseline levels of performance.

*Of course, it could be argued that the very dissimilarity between such artificial milieus and the rest of life makes it even more difficult to create lasting change.

Translation: when the goodies stop, people go right back to acting the way they did before the program began. In fact, not only does the behavior fail "to generalize to conditions in which [are not in effect" — such as the world outside the hospital — but reinforcement programs used each morning generally don't even have much effect on patients' behavior during the afternoon!

Ten years later, one of these authors, Alan Kazdin, checked back to see if anything had changed. Was the initial failure due only to inadequate implementation of a basically sound idea? After reviewing another decade's worth of research, the best he could offer was the rather tepid statement that the gains produced by token economies "are not inevitably lost." In some programs, "intervention effects are at least partially

maintained"; in others, they are not. On closer inspection, though, even this modest claim couldn't be defended. It turned out that the programs that seemed most successful often had been combined with other, more substantive reforms (including, in the case of schools, reducing the size of classes, getting parents more involved, and so forth) These other changes, of course, might well have been responsible for any beneficial effects. "As a general rule," Kazdin wrote, with an almost audible sigh, "it is still prudent to assume that behavioral gains are likely to be lost in varying degrees once the client leaves the program."

One study conducted in a classroom should convey a feel for the kind of research he reviewed. Over the course of twelve days, fourth and fifth graders were rewarded for playing with certain math-related games and were not rewarded for playing with others. (None of these activities was inherently more interesting than any other.) When the rewards started, the kids promptly gravitated to the games that led to a payoff. 'When the rewards disappeared, their interest in those games dropped significantly, to the point that many were now less interested in them than were children who had never been rewarded in the first place. The researchers concluded that: the use of powerful systematic reward procedures to promote increased engagement in target activities may also produce concomitant decreases in task engagement, in situations where neither tangible nor social extrinsic rewards are perceived to be available.

Other reviewers surveying the landscape, including committed behaviorists, have reached a conclusion similar to Kazdin's, finding plenty of reason to doubt the long-term effectiveness of token economy programs or, at best, claiming that there still isn't enough research to know for sure that they work. Perhaps most telling is the finding that in those cases where behavior change did continue after the initial rewards were withdrawn, it was only because new rewards were substituted for the original ones.

Is this failure to promote enduring change restricted to token economies? These programs are rarely used today, but behaviorists are still called in to help people develop good habits or break bad habits. Three areas where there is enough evidence to permit at least provisional judgments about their success are losing weight, quitting smoking, and using seatbelts. Here's what the data show:

Losing weight In one dieting study, some subjects were promised a twice-a-week reward of five dollars each time the scale showed good news, while others got nothing. Those who were paid did make more progress at the beginning, but then gained back the weight — and then some — over the next five months. By contrast, those who had not been rewarded kept getting slimmer.' This study was quite small, and a lot of the subjects were unavailable for measurement at the end, so we probably shouldn't give it too much weight. But a similar study published ten years later offered little solace for behaviorists: after a year, no difference was found between the payment and nonpayment groups. (Actually, there was one difference: many of those who had been promised money for shedding pounds failed to show up for the final weigh-in.)

Quitting smoking Losing weight and keeping it off are inordinately difficult, so it may be unfair to reject pop behaviorism just because it hasn't worked miracles here. The trouble is that it hasn't done much better elsewhere, assuming we are looking for long-term gains. Take smoking cessation. A very large study, published in 1991, recruited

subjects for a self-help program designed to help people kick the habit. Some were offered a prize for turning in weekly progress reports; some got feedback designed to enhance their motivation to quit; everybody else (the control group) got nothing. What happened? Prize recipients were twice as likely as the others to return the first week's report. But three months later, they were lighting up again more often than those who received the other treatment and even more than those in the control group! Saliva samples revealed that subjects who had been promised prizes were twice as likely to lie about having quit. In fact, for those who received both treatments, "the financial incentive somehow diminished the positive impact of the personalized feedback." Not only were rewards unhelpful; they actually did harm.*

[* Recall that rewards were also positively harmful in the program that promised goodies to children for playing with certain math games. There, interest in the games dropped below what was found at the beginning; here, getting a prize for quitting; smoking was worse than getting nothing at all. This sort of thing happens often enough that behaviorists have had to invent a neutral-sounding name for it: it is technically known as the "contrast effect."]

Using seat belts Even more research has been done on applying behaviorism to the promotion of seat belt use. In fact, an enthusiastic partisan of behaviorism and his colleagues reviewed the effects of twenty-eight programs used by nine different companies to get their employees to buckle up; nearly half a million vehicle observations were made over six years in this research. The result: programs that rewarded people for wearing seat belts were the *least* effective over the long haul. In follow-up measures ranging from a month to more,, than a year later, programs that offered prizes or cash for buckling up found changes in seat belt use ranging from a 62 percent increase to a 4 percent decrease. Programs without rewards averaged a 152 percent increase. The authors, who clearly did not expect this result, had to confess that "the greater impact of the no-reward strategies from both an immediate and long-term perspective... [was] not predicted and [is] inconsistent with basic reinforcement theory."

Other uses Some psychotherapists and couples counselors also use rewards to change behavior. One behaviorist, for example, actually suggests that spouses "use tokens with each other to encourage conversation [or] to control excessive talk." An attempt to assess the effectiveness of rewards here would pull us into a long and complicated discussion about how therapy works and how its success can be measured. What stands out, though, in any evaluation of the topic is that plans to solve problems by rewarding certain behaviors can only be expected to work for as long as these rewards are still in effect. Moreover, "eliciting desired behaviors is not the only, indeed, not even necessarily the most important outcome of psychotherapy"; what matters more are "the underlying psychological processes," which behavioral approaches refrain from addressing.

3. ***At what, exactly, are rewards effective?*** To ask how long rewards last, and to learn that they rarely produce effects that survive the rewards themselves, is to invite curiosity about just what it is that rewards are doing. Why don't people keep acting the way they were initially reinforced for acting? The answer is that reinforcements do not generally alter the attitudes and emotional commitments that underlie our behaviors. They do not make deep, lasting changes because they are aimed at affecting only what we *do*. If, like Skinner, you think there is nothing to human beings other than what we do — that we are

only repertoires of behavior — then this criticism will not trouble you; it may even seem meaningless. If, on the other hand, you think that actions reflect and emerge from who a person is (what she thinks and feels, expects and wills), then interventions that just control actions wouldn't be expected to help a child grow into a generous person or even help an adult decide to lose weight.

What rewards and punishments do is induce compliance, and this they do very well indeed. If your objective is to get people to obey an order, to show up on time and do what they're told, then bribing or threatening them may be sensible strategies. But if your objective is to get long-term quality in the workplace, to help students become careful thinkers and self-directed learners, or to support children in developing good values, then rewards, like punishments, are absolutely useless. In fact, as we are beginning to see, they are worse than useless — they are actually counterproductive.

Do Rewards Improve Performance?

In 1961, a graduate student at the University of Kentucky found something she didn't expect. For her dissertation, Louise Brightwell Miller arranged a series of simple drawings of faces so that pairs of nearly identical images would be flashed on a screen. Then she brought 72 nine-year-old boys into her laboratory one at a time and challenged them to tell the two faces apart. Some of the boys were paid when they succeeded; others were simply told each time whether or not they were correct.

Miller expected that the boys would do a better job when there was money at stake. Instead, she found that those who were trying to earn the reward made a lot more mistakes than those who weren't. It didn't matter how much they were paid (one cent or fifty cents) or whether they were highly motivated achievers (as measured by a personality test). The discovery left her scratching her head: "The clear inferiority of the reward groups was an unexpected result, unaccountable for by theory or previous empirical evidence," she and her adviser confessed.

The following year, another graduate student, Sam Glucksberg, published the results of his own dissertation research at New York University in the same journal. This time it was undergraduates, 128 of them in all, who were brought into a lab individually. Each was given matches, thumbtacks, and the boxes they came in and told to mount a candle on a wall using only these materials. (They were supposed to figure out that an empty box could be tacked to the wall and the candle placed on top of it.) Some of the students were given; empty boxes, with matches and tacks on the side; others got full boxes, which made the solution much less obvious

As with Miller's experiment, some of the students were informed that they could earn anywhere from $5 to $20 — quite a lot of money in 1962 if they succeeded; others weren't promised anything. Even though the subjects were older and the assignment quite different, Glucksberg's results echoed Miller's: when the task was more challenging, those who were working for the financial incentive took nearly 50 percent longer to solve the problem.

Nobody paid much attention to these studies at the time — or since, for that matter. But because most of us assume that better work is done by people who know they are going to be rewarded for doing it, this early research takes on a certain retrospective

significance. Those graduate students may not have realized it, but they had stumbled onto something enormously important.

In the early 1970s, a batch of new reports came out that showed the early results were no flukes. Janet Spence, a psychologist at the University of Texas who later became the president of the American Psychological Association, published two studies in which children were asked to remember which of two words was "right" (as arbitrarily determined by the experimenter) and then choose that one over others when it came up again later. Some of the children simply saw a light come on (or a bean drop down a chute) when they chose correctly; others got either an M&M or a token that could be exchanged for an M&M when they were finished. It turned out that the children who received candy or the promise of candy got fewer right than those who received nothing more than information about how well they were doing — a result that led her to make the comment that appears at the beginning of this chapter.

Four other studies, each conducted by a different experimenter and published in a different journal, were reported that same year:

• One researcher asked undergraduates to "select the pattern on each page that was least like the other two patterns on that page." To his surprise, he found that students "who were not offered money performed significantly better than those who were paid." So he doubled the amount of the reward — and got exactly the same result.

• In an experiment that ventured outside the psychology laboratory, college students who worked on the school newspaper were observed as they learned to "writ headlines according to prescribed rules." As they got better over time, they were able to work more quickly. For a while some students were paid for each headline they turned out — with the result that their performance stopped improving. Those who received no money kept getting better.

• Fourth graders performed more poorly on a task when they were offered the very reward (some sort of toy or candy) that they had earlier indicated they especially liked. The experimenter pronounced the results "puzzling."

• High school students were given five different tasks, some testing their memory and some requiring true creativity. Once again, some were promised a reward while others were not. And once again, regardless of the task, the rewarded subjects didn't do nearly as well.

As the 1970s wore on, still more evidence accumulated. Preschoolers who expected an award for drawing with felt-tip pens drew at least as many pictures as those who didn't expect an award, but the quality of their drawings was judged to be appreciably lower. (That rewards can have one effect on quantity and another on quality has been noticed by other researchers, too.)

Another group of college students took longer to solve a problem requiring creativity when they were rewarded for doing so. And in a particularly intriguing experiment, sixth-grade girls who were promised free movie tickets for successfully teaching younger girls to play a new game wound up doing a lousy job as tutors: they got frustrated more easily, took longer to communicate ideas, and ended up with pupils who didn't

understand the game as well as those who learned from tutors who weren't promised anything.

By the 1980s, anyone who kept up with this sort of research would have found it impossible to claim that the best way to get people to perform well is to dangle a reward in front of them. As the studies became more sophisticated, the same basic conclusion was repeatedly confirmed. College students exhibited "a lower level of intellectual functioning" when they were rewarded for their scores on the more creative portions of an intelligence test. (Their scores on the portions of the test requiring less insight and discovery were neither hindered nor helped by rewards.) In a separate study, third graders who were told they would get a toy for working on some "games" (which, in reality, were also IQ tests) didn't do as well as those who expected nothing. In research by Barry Schwartz, adults who had to figure out the rules to another sort of game, in effect trying to think like scientists, were less successful if they had been trained at the task earlier and promised a monetary reward for doing well. (They had trouble breaking out of the fixed pattern of behavior that had succeeded in producing rewards for them before.)

A few years later, Teresa Amabile, a leading student of creativity, published two reports that clinched the case against the use of rewards. In the first, young creative writers who merely spent five minutes *thinking* about the rewards their work could bring (such as money and public recognition) wrote less creative poetry than others who hadn't been reflecting on these reasons for pursuing their craft.

The quality of their writing was also lower than the work that they themselves had done a little while earlier. Then Amabile conducted a series of studies with children and adults that involved such tasks as making collages and inventing stories. Some subjects were promised rewards — real ones this time — and others weren't. Again, rewards killed creativity, and this was true regardless of the type of task, the type of reward, the timing of the reward, or the age of the people involved.

As recently as 1992, researchers were still finding that rewards undermine different kinds of performance. Amabile and her associates discovered that professional artists do less creative work when that work is commissioned — that is, when they have contracted in advance for a reward. And Mark Lepper, who was once Amabile's adviser, conducted a study with a graduate student in which fourth and fifth graders were given a problem-solving task similar to the board game Clue. Those who were promised a toy for doing well "formulated hypotheses in a much less systematic fashion" and took longer to get the solution than those who weren't promised anything. Even more disturbing, those anticipating a reward also did a poorer job on an entirely different task a week later.

Still other researchers, approaching the topic from different angles, have found additional reasons to question the wisdom of pop behaviorism. One series of investigations considered the basis on which people are rewarded. After conducting six separate studies, Morton Deutsch concluded that "there is no evidence to indicate that people work more productively when they are expecting to be rewarded in proportion to their performance than when they are expecting to be rewarded equally or on the basis of need." (In later chapters I will cite other evidence suggesting that pay-for-performance in the workplace

and an emphasis on grades in the classroom are both counterproductive, exactly as Deutsch's data would lead us to predict.)

Other investigators, meanwhile, have been looking at people's attitudes toward rewards. Ann Boggiano and Marty Barrett found that children who are extrinsically motivated — that is, concerned about things like the rewards and approval they can get as a result of what they do in school — use less sophisticated learning strategies and score lower on standardized achievement tests than children who are interested in learning for its own sake. The reward-driven children do more poorly even when they are compared with children whose scores the previous year were identical to their own.

I have described these studies individually rather than just summarizing the basic finding because without the supporting details of the research the conclusion might be hard to accept. After all, it was hard for the researchers themselves to accept — at least, until the results appeared so consistently that they had no choice. But before going on to examine the reasons for these results, let us take a moment to sort them out and think about what they imply and why they seem so startling.

Recall the three questions posed at the beginning of this chapter: For whom are rewards effective, for how long, and at what? We know that some people will do a better job at some things when there's a goody at stake, but few of us have stopped to consider just how limited the circumstances are in which this is true. For whom do rewards work best? For those who are "alienated from their work," according to Deutsch. If what you've been asked to do seems silly or simple, you might decide to make a real effort only when there is something else, something outside the task itself, to be gained. (It shouldn't be surprising, then, that researchers find rewards are least effective — in fact, positively counterproductive — when people get these rewards for doing things that are optimally challenging for them, neither too hard nor too easy.) This, of course, says something about the task as much as about the individuals involved; more accurately, it speaks to the relation between the two.

For how long do rewards work? Most of the research on this question concerns behavior change, the sort of effects discussed in the preceding section. Virtually all of the studies concerned with performance look at how well people do at a task immediately after getting, or being promised, a reward. In order for rewards to have any hope of boosting performance over a long period of time, we typically have to continue giving them out, or at least holding out the possibility that more will follow.

We come, finally, to the key questions: At what sort of tasks do people do a better job when they are rewarded? And "better" in what sense? By now we have already seen enough evidence to guess the answers. *Rewards usually improve performance only at extremely simple — indeed, mindless — tasks, and even then they improve only quantitative performance.* The unexpected results from those first dissertation studies by Miller and Glucksman appeared against a backdrop of research by behaviorists that counted things like an increase in the number of times an organism pressed a lever as proof that reinforcement improves performance. If you were given an enormous pile of envelopes to seal, you would probably lick them a lot faster if you were paid to do so. The trouble is that we have incorrectly inferred from this fact a general law of human

nature — Reward people and they'll do a better job — and applied it in our workplaces and schools. This faulty application (which a behaviorist might refer to as "response overgeneralization") goes a long way toward explaining why so many of our workplaces and schools are now in trouble.

One of the most influential papers on the topic of rewards (influential, that is, for the very few social psychologists who are specialists in the field) reached the following conclusion based on research conducted up until the mid-1970s:

> Incentives will have a detrimental effect on performance when two conditions are met: first, when the task is interesting enough for subjects that the offer of incentives is a superfluous source of motivation; second, when the solution to the task is open-ended enough that the steps leading to a solution are not immediately obvious.

This analysis by Kenneth McGraw provides us with a good point of departure from which to figure out when rewards are likely to fail. Subsequent investigations, for example, have confirmed that a Skinnerian approach is particularly unlikely to prove useful when it is creativity that we are trying to promote.

But McGraw's rule may understate the failure of rewards by suggesting that they will miscarry only when used with interesting and creative tasks. I think it is more accurate to say that they are most *likely* to have a detrimental effect, or to have the *most pronounced* detrimental effect, with these tasks. It is true that some studies have found that people's performance at very basic things, such as multiplication, may improve when they are expecting to receive a reward. But the research I have described in this section includes enough examples of impaired performance at rather straightforward tasks — or at least a failure to enhance performance at these tasks — that we cannot casually assume it makes sense to reach for the reinforcements for everything that doesn't demand creativity.

"Do this and you'll get that" turns out to be bad news whether our goal is to change behavior or to improve performance, whether we are dealing with children or adults, and regardless of whether the reward is a grade, a dollar, a gold star, a candy bar, or any of the other bribes on which we routinely rely. Even assuming we have no ethical reservations about manipulating other people's behavior to get them to do what we want, the plain truth is that this strategy is likely to backfire.

As one psychologist read the available research, people who are offered rewards tend to choose easier tasks, are less efficient in using the information available to solve novel problems, and tend to be answer oriented and more illogical in their problem-solving strategies. They seem to work harder and produce more activity, but the activity is of a lower quality, contains more errors, and is more stereotyped and less creative than the work of comparable non-rewarded subjects working on the same problems.