Swarthmore College

Works

Psychology Faculty Works

Psychology

1993

On The Creation And Destruction Of Value

Barry Schwartz

Swarthmore College, bschwar1@swarthmore.edu

Follow this and additional works at: <https://works.swarthmore.edu/fac-psychology>

Part of the [Psychology Commons](#)

[Let us know how access to these works benefits you](#)

Recommended Citation

Barry Schwartz. (1993). "On The Creation And Destruction Of Value". *Origin Of Values*. 153-186.
<https://works.swarthmore.edu/fac-psychology/515>

This work is brought to you for free by Swarthmore College Libraries' Works. It has been accepted for inclusion in Psychology Faculty Works by an authorized administrator of Works. For more information, please contact myworks@swarthmore.edu.

8

On the Creation and Destruction of Value

Barry Schwartz

Finally, there came a time when everything that men had considered inalienable became an object of exchange, of traffic, and could be alienated. This is the time when the very things which till then had been communicated, but never exchanged; given, but never sold; acquired, but never bought—virtue, love, conviction, knowledge, conscience—when everything, in short, passed into commerce. It is the time of general corruption, of universal venality.

Karl Marx, *The Poverty of Philosophy*

Can there be a science of values? If so, what might it look like? How should it begin? In this chapter, I will suggest that as the term *science* is ordinarily understood, the most likely result of attempts to construct a science of values will be to reify phenomena that are historically specific and context dependent into timeless generalizations. These generalizations may then take on normative force, becoming a kind of ideology that helps shape social life and social institutions in a way that makes the generalizations self-fulfilling. It may be possible to construct a science of values that is both illuminating and useful, but only if we understand *science* to include the critical examination of history and of culture.

To tell this story, I will first say something about what science, as traditionally understood, is. Then, I will sketch what a traditional science of values might look like. Next, I will attempt to exemplify the shortcomings inherent in a traditional approach by discussing phenomena that illustrate how the values of individuals can be changed—created and destroyed—by certain kinds of experience. These phenomena also illustrate how changes in individual values can contribute to the transformation of social institutions and of the values embodied by these institutions.

Science

The practice of science depends on a set of metaphysical and epistemological commitments that are so commonplace that they rarely rise to the level of explicit consideration or discussion, at least among scientists. However, it is useful to review them, especially in a context in which the extension of a science to a categorically new domain (values) is being considered. Lacey (1988) has recently rendered these commitments in a form that is especially useful for the present discussion.

First, science presupposes that there is an objective causal order that is ontologically independent of human inquiry, perception, and action, an order whose character does not vary with the theoretical commitments, interests, or values of the people investigating it.

Second, science presupposes that underlying this objective causal order are laws that are independent, both ontologically and causally, of human inquiry, perception, and action. These laws capture the state of the world as it is, and their generative (predictive) power defines and circumscribes possible future states of the world.

The principal aim of science is to create theories that are adequate representations of these laws. These theories can be developed with methods and practices that are known to provide adequate representations of the way the world is. And they can be evaluated by appeal to widely accepted criteria for assessing the adequacy of possible theories. The criteria for evaluation depend on data that meet the following conditions:

1. The truth value of any datum can be recognized by anyone with suitable training, simply through making the appropriate observations.
2. All parties to theoretical disputes accept the relevance of a given set of data to the disputes, though they may differ on how those data may best be characterized.
3. Only data in this class are relevant to the resolution of theoretical disputes.
4. Data that represent replicable experimental results have special status in evaluating theoretical representations of a domain.

The data set on which a theory is based can never be exhaustive. Data are always being created, both in experiment, and in the scientist's selection of what to focus on and what to ignore. Nevertheless, the rebuttable presumption is that the data under consideration at any given time are *representative* of the data set as a whole. When theories conflict, theory choice is based on an assessment of which theory has greater explanatory

and predictive power with regard to data that meet the above conditions. And importantly, theory choice is based on nothing else.

Values

Now that we have before us a picture of “canonical” science, what about values? What are the entities that we are to imagine studying and understanding scientifically? I take values to be principles, or criteria, for selecting what is good (or better, or best) among objects, actions, ways of life, and social and political institutions and structures. Values operate at the level of individuals, of institutions, and of entire societies.

A social institution embodies individual values when, in the normal course of its operation, the institution offers people roles that encourage behavior that displays the values, and fosters conditions for the further expression of the values. Thus, for example, elite liberal arts colleges embody the value of intellectual cultivation to a high degree. They embody the value of cooperation and group solidarity to a lesser degree, and the value of service to the poor hardly at all. An entire social order embodies a value to the extent that it provides conditions that nurture social institutions that embody the value. The values that an individual can express are very much constrained by the character of the social institutions and the social order in which that individual lives. Indeed, social stability probably depends on a meshing of personal values and institutional opportunities for their expression.

Science and Values

What, then, might a science of values look like? One possibility is that a science of values would discover, by empirical inquiry, the large and varied set of objects, experiences, and actions that different people value. Such a science would not be especially illuminating. It would not tell us *why* (that is, under what conditions) certain objects, experiences, and actions are valued over others. We could enrich the science by adding to the set of values a set of boundary conditions that specifies when some objects, experiences, or actions will be valued rather than others. Still richer would be a more abstract characterization of what people value, so that deep similarities that underlie surface differences might be detected. For example, the claim that people value whatever makes them feel good, or whatever maximizes utility, or whatever promotes inclusive

reproductive fitness, if true, might permit us to unify quite diverse individual value schemes. The surface diversity could be the result of institutional requirements and constraints that have a heavy hand in determining what will make individuals feel good, or give them utility, or promote their reproductive fitness. Disciplines such as behavior theory, economics, and sociobiology, which purport to discover the hidden universal laws of human motivation by which all societies operate, are presumably guided by this kind of theoretical aspiration.

Problems for a Science of Values

These brief remarks on science and on values are intended to lay the groundwork for a discussion of the problems that arise when the two are combined. My focal concern is the dynamic interplay between human values and social institutions. And the problem I am concerned with is this: the "value scientist," using whatever empirical methods are at his or her disposal, identifies the values that characterize both individuals and institutions at a given time and place. From this empirical work, the scientist makes inferences about underlying causal laws. These are taken to be "laws of nature," generalizations about the way the world must work. But in actual fact, they may not be laws of nature. Instead, they may be facts of history that arise out of the interplay of human beings *making* their institutions and in turn being made by them. Claims that are *contingently* true of people located at a particular time and place become reified into laws of nature.

And to this problem we can add a further one. None of the epistemological criteria of science that I outlined above would ensure that we could detect this reification if it occurred. That is, a careful scientist, following all the rules, could easily confuse observations about what *is* the case for evidence about what *must be* the case (see Schwartz & Lacey, 1982, 1988; Lacey & Schwartz, 1986, 1987). In short, the metaphysical commitments of science make it difficult for the scientist to detect that the phenomena being investigated violate those commitments. Detection of the contingent, historical character of generalizations about values requires a perspective that is both metaphysically and epistemologically more inclusive than science.

I will illustrate these problems by examining a particular example of how values can be created and destroyed. By looking at how these value changes occur, and at how they are reflected in the character of our social institutions, I will try to illustrate what the limitations are to a science of values.

The Destruction of Value and the Creation of New Means—Ends Relations

The part of experimental psychology known as behavior theory or reinforcement theory has focused historically on how instrumental or operant behavior is controlled by its consequences. The study of how behavior is controlled by consequences has had built into it the presumption that means and ends—operant responses and reinforcers—are both conceptually and empirically distinct. The relation between the particular response one requires an organism to make and the reinforcing consequence of that response is *arbitrary*. It does not exist prior to the experimental intervention (e.g., Schwartz, 1989). The various means to reinforcement are essentially interchangeable with one another, and they have no value apart from their relation to the consequences they produce. These kinds of arbitrary response—outcome relations are studied because they are thought to be paradigmatic of means—ends relations that characterize human behavior. The automobile assembly line worker can perform anywhere on the line for his or her weekly wage. Which particular task is required is a matter of indifference, as long as the rate and quantity of reinforcement are held constant.

It is undeniable that some human activities reflect the kind of means—ends relation that characterizes studies of operant conditioning. However, the relation between means and ends need not have this arbitrary form. For some activities, means and ends are interconnected. To see the point, consider the concrete example of a man who works as an automobile mechanic from nine to five each day, and then goes home to pursue his hobby—restoration of old cars to running order. On his job, fixing cars is an operant. The weekly paycheck is the reinforcer. He would not be fixing cars were it not for the paycheck, and he would just as soon do some other kind of work for an equivalent or greater paycheck. Thus his job, the operant, is a means, and his paycheck, the reinforcer, is an end, and there is no special relation between the means and the end that could not be duplicated by substituting some other job for his current one.

The situation is quite different when he gets home. Now, fixing cars is both means and end, operant and reinforcer. While it is true that he does not tinker with cars just for the sake of tinkering—achieving the goal of a smooth-functioning automobile is an important influence on his activity—it is also true that he would not be satisfied with any old means of achieving that goal. He would not, for example, be satisfied with hiring someone else to restore the old cars for him. The reinforcing consequences of the activity are a part of the activity itself, and other kinds of activity are not interchangeable with it in the service of the same

reinforcer. Indeed, we might even say that "owning old cars that run well" is not even properly a reinforcer, for it will not increase the likelihood of any operant except for "fixing old cars." Similarly, the operant "fixing old cars" is not properly an operant, since it will not be reinforced by any reinforcer except "having old cars that run well."

The distinction between this man's job and his hobby should be familiar. Some people have jobs that are like this man's; they are simply means to an end—pure operants performed solely for the wage that would be given up immediately if a bigger wage came along. Other people are fortunate enough to have jobs that are more like this man's hobby. While the wage is certainly significant, and without it people would not do the job (just as for the hobbyist, having a finished, working automobile is crucial, and without it, he would abandon his hobby), it is not everything. There are aspects of the job itself that make it more than just a means and make people unwilling to substitute other jobs that pay just as well or better. So even though people work at these jobs for the wage, the jobs themselves are both operant and reinforcer.

The relevance of this means—ends distinction to the creation and destruction of values is this: whether activities will be purely instrumental or will possess some intrinsic value or connection to the ends they produce depends on how those activities are organized. And the way in which activities are organized is subject to historical and cultural change. Thus, whether and why activities are valuable is a matter not of natural law, but of cultural contingency. Nowhere is this more clearly in evidence than in the history of the workplace.

Centuries ago, what came to be modern industrial society was feudal. Large portions of land were controlled by lords. The majority of the population worked the lord's land, as serfs. These serfs had no legal alternative to the work they did. In return for his protection, serfs were required to work the lord's land, and to turn over a fixed proportion of their yield to him. They had no choice of the terms they would work under, or of the conditions of their work. They could not hire themselves out to the highest bidder. Nor could the lord sell off his land. The details of the relation between serf and lord were part of a long-standing set of political and social practices that was neither based strictly on economic considerations nor changed on the basis of these considerations. This network of political and social practices is what economic historian Karl Polanyi had in mind when he said, "man's economy, as a rule, is submerged in his social relationships" (1944: 46).

If the factors operative in the choice of work were different in feudal than in modern times, so also was the nature of the work itself. Serfs, and other premodern workers, engaged in a wide variety of different

activities in the course of a day. Their work required flexibility and decision making. The rhythm and pace of their work changed with the seasons. In addition, the work they did for the lord was integrated into the rest of their daily activities. They did not leave home for the shop, work from 9 to 5, then return home to engage in personal pursuits. This pattern of work is in sharp contrast to the modern factory worker, who does the same thing all day, every day, with no flexibility or decision making required.

Over a period of several hundred years after feudalism ended, the descendants of serfs eventually became wage laborers. This change coincided with other changes in work that resulted in the emergence of the factory system. By the end of the eighteenth century in England, many of masses of people were not only working for wages, but were free to hire themselves out to the highest bidder. Moreover, with increasing mechanization and division of labor, work became less and less varied and flexible. When the factory system was fully in place, behavior in the workplace seemed a perfect exemplification of the laws of operant behavior in operation.

Industrialization, as we now know it, did not come all at once (Hobsbawm, 1964). For a time, even when masses of people were working for a wage, the wage they received and the way they did the work were largely determined by social custom, not by the competitive market. That is, workers did not hire themselves out to the highest bidder, and bosses did not try to extract maximal output for minimal cost. Thus, complete control of work by wage rates (reinforcement rates) was not characteristic of early industrialization. This is not to suggest that work was uninfluenced by the reinforcement contingencies. Clearly, if workers received no pay at all, they would not have worked. However, pay rates did not exert the same kind and degree of control over workers as reinforcement rates exert over animals.

As industrialization proceeded, however, wages came completely to dominate the work people chose and the way they performed it. Workers learned to sell their labor to the highest bidder. And the reason that work came to be completely dominated by the wage is that custom, its principal competitor for control, had been systematically and intentionally eliminated. A central component of the final stages of development of the workplace, in its modern form, was a movement explicitly designed to eliminate custom as an influence on behavior. The movement was one of the earliest examples of what is now called "human engineering." It went by the name of "scientific management" and its founder and leader was Frederick Winslow Taylor.

Taylor (1911/1967) argued that custom interfered with efficiency and

productivity. What industry needed was a set of techniques for controlling the behavior of the worker that was as effective as the techniques used for controlling the operation of machines. Accomplishing this control involved two distinct lines of human engineering. First, one would need to discover the rates and schedules of pay that resulted in maximal output (e.g., Gilbreth, 1914). Second, one would need to break up customary ways of doing work, and substitute for them minutely specialized and routinized tasks that could be accomplished mechanically and automatically. The idea was to strip work down to its simplest possible elements, to eliminate the need for judgment and intelligence, and to wrench work free of its customary past. With this done, there would be no possible source of influence on work except for the schedule of pay. And the schedule of pay was something the boss could control.

Thus, work as pure means, as purely operant behavior, is a relatively recent human invention. It is an invention that took all value out of work itself, and located it instead in the wage, the consequence. And it is an invention that was abhorrent to most of the workers subjected to it. As Marglin (1976) has pointed out, bosses had enormous difficulty in harnessing the efforts of their workers. Workers chafed at the confining discipline of the factory. They malingered, they failed to appear, they quit altogether. Harnessing the worker was difficult, and for the successful boss, it was a singular achievement. But eventually, the problem of inducing workers to put up with the conditions of the factory disappeared. Eventually, what for one generation was the wrenching out of a complex network of customs and social relations was for another "only natural." So it was that scientific managers could see themselves as merely increasing the efficiency of work rather than transforming its very character. And so it may be that the reinforcement theorist, looking around at the "natural" order of things, can see his principles as reflecting an eternal necessity of human nature rather than an historical contingency (see Schwartz, Schuldenfrei, & Lacey, 1978, for a more detailed discussion of this process).

Creation of New Means-Ends Relations: Laboratory Evidence

I am suggesting that the historical transformation of the nature of work provides evidence that value can be destroyed. Supporting evidence can be found in several lines of experimental research. One such line can be summarized as showing how rewards can have the effect of "turning play into work." People are given the opportunity to engage in a variety of activities that might be regarded as pleasurable: solving various puzzles, for example. These are activities people would happily

engage in in the absence of any reinforcement. The twist in these demonstrations is that even though no reinforcement is necessary to keep people at the activities, they get it anyway, typically in the form of money.

And the reinforcement has two effects. First, predictably, it gains control of the activity, increasing its frequency. Second, and more significant for our purposes, when reinforcement is later withdrawn, people engage in the activity even less than they did before reinforcement was introduced. The withdrawal of reinforcement does not simply reduce responding to its prereinforcement, baseline levels; it eliminates responding almost completely (see Deci, 1975; Greene, Sternberg, & Lepper, 1976; Lepper & Greene, 1978; Lepper, Greene, & Nisbett, 1973).

In one particular demonstration (Lepper et al., 1973), the experimental subjects were nursery school children. They were given the opportunity to draw with felt-tipped drawing pens, an activity that seems to have almost unlimited appeal to young children. After a period of observation, in which experimenters measured the amount of time the children spent playing with the pens, the children were taken into a separate room where they were asked to draw pictures with the pens. Some of the children were told they would receive "Good Player" awards (reinforcement) if they did the drawing; others were not. A week later, back in the regular nursery school setting, the drawing pens were again made available, with no promise of reward. The children who had received awards previously were *less* likely than the others to draw with the pens at all. If they did draw, they spent less time at it than other children, and drew pictures that were judged to be less complex, interesting, and creative. Without the prospect of further awards, their interest in drawing was only perfunctory.

What is important to note about this demonstration is that it is not an example of the failure of the principle of reinforcement. On the contrary, the awards seemed to gain control over the behavior. If they had continued to be available for drawing in the classroom, there is little doubt that high rates of drawing would have been maintained. The point of the demonstration is that prior to the introduction of the reinforcement contingency, something else was influencing the drawing, and that other influence was suppressed or superceded by the reinforcement contingency. Clearly, this laboratory demonstration is an example of how value can be destroyed.

Another series of experiments indicating that reinforcement can usurp control of an activity from other sources has been going on in my laboratory over the past several years (Schwartz, 1982, 1988). In a prototypical procedure, subjects are seated in front of a matrix of light bulbs, five across by five down. Beside them are two push-buttons and a counter to keep track of their score. Periodically, the top left bulb in the matrix

lights up, signalling the start of a trial. "This is a game," subjects are told. "By pushing the two buttons, you can change the position of the illuminated light in the matrix of lights. If you do it right, you get a point. What I want you to do is to figure out the rules of the game; figure out what you have to do to earn a point."

A subject might first push the left button, and observe that the light moves down one position. She might then push the right button, and observe that the light moves across one position. Left, right, left, right. After four alternations, all the lights go out, and the subject gets a point. The next trial begins, and she pushes the buttons in exactly the same order, and again gains a point. She does the same thing, with the same result, on the third trial. Thinking she has the game figured out, she calls over the experimenter. "You have to start on the left and alternate between left and right. Four alternations get you a point." "Wrong" says the experimenter, "try again."

The subject realizes she has made a silly mistake. The experimenter wants to know what one has to do to get a point—what is *necessary*. All that she has discovered is one particular way to do it—what is *sufficient*. The way to find out what is necessary is to vary what one does on each trial, in systematic fashion. For example, to test whether it is necessary to alternate, starting on the left, one starts a trial on the right. The name of this game is "experimental science." There is a phenomenon, the getting of points, and the task is to discover its causes. One goes about this task by formulating guesses or hypotheses, and by doing experiments to test the hypotheses. The process one goes through in attempting to discover the rules of the game is precisely the one that scientists go through as they attempt to discover the rules of whatever "natural game" they are studying. And the process has two essential ingredients: formulation of hypotheses and tests of the hypotheses by systematic variation in experiment.

Moreover, there is nothing about these processes that is unique to science. People engage in them frequently, if somewhat less systematically, in everyday life. Someone interested in learning how to bake bread might find a recipe, and follow it carefully, step-by-step, with the bread turning out delicious. Now if all one cares about is knowing *a* way to make good bread, this recipe will be followed every time the need for bread arises. Similarly, if all our game-player cares about is finding *a* way to get points, once he or she finds that left-right alternation succeeds, he or she will never deviate from it. But suppose one wants to know more than *a* way to make good bread. Suppose one wants to discover the most efficient way to make good bread. Now, what the baker must know is the essentials of good bread baking, what is necessary to make good bread—

the rules, as it were, of bread baking. And as in the case of the button-pushing game, this requires experimentation, the formulating and testing of hypotheses.

In our experiments, college students were asked to try to discover the rules of the game. When they succeeded in discovering a rule, it was changed, and the students did it again. They were given no instructions about how one might most effectively tackle the problem. Sometimes, the students were told that they would get a few cents for each point they earned in the process of discovering the rule. Sometimes they were told they would get a dollar for each rule they discovered. Sometimes, they were able to get a few cents for every point and a dollar bonus for every rule. Finally, sometimes no monetary rewards were available at all.

Think about how these various contingencies of reward might affect students' behavior. Suppose they could earn money for every point. This would put them in conflict. If what they care about is discovering the rule, they will vary their responses from trial to trial. But if what they care about is earning as much money as possible, once they find a successful sequence of responses, they will stick with it. After all, every time they experiment by varying their sequences of responses, they risk failing to earn a point. When one experiments with bread baking technique, one risks producing lousy bread. If, in contrast, they earn money only by discovering the rule, then there is no conflict. Both the contingency of reward and the intrinsic demand of the task itself encourage them to vary their response sequences systematically and intelligently.

When students played the game, these varying conditions of reward made no difference at all. In all cases, students varied their responses from trial to trial with great efficiency. Almost every one of them discovered each of the rules, and they did so quite rapidly. The reinforcement contingencies failed to have any impact. But there is more to the story. Another group of students was exposed to the same set of problems with the same contingencies of reward as the first set. What distinguished the two groups was that this second group had had prior experience playing the game. The prior experience was this: they were brought into the laboratory, shown the game, and told that every point they scored would earn them two cents. They were then given one thousand opportunities (trials) to play the game. Note that there was nothing in the instructions they received urging them to discover the rule. They could have tried if they wished, of course, but they could also just find a sequence of responses that worked, and stick with it, earning as much money as possible. And that is what they did. Each student settled on a particular sequence of responses that occurred on about 90% of all trials. The little game and the contingency of reinforcement had turned the subjects into assembly-line

workers, engaged in the same task, done the same way, over and over again, completely controlled by the reinforcement contingency.

What happened then when these newly formed factory workers were instructed to discover the rules? Compared to the first, inexperienced group, they were much less effective. They discovered fewer of them, and took longer in discovering the rules when they were successful. And unlike the first group, what they did was powerfully influenced by the prevailing contingency of monetary reward. They were especially ineffective at discovering rules if each point they got earned them money. Some of the results of this experiment are presented in Figures 1 and 2. Figure 1 presents the percentage of problems solved by naive and pre-trained subjects in each of the four different reward conditions, and Figure 2 presents the number of trials per solution across the same subjects and conditions.

That we obtained this result is surprising, for several reasons. First, we were dealing with a group of bright college students, people who get their kicks from acting like scientists. Nevertheless, pretraining seemed to overcome this orientation, or at least to place the game outside the problem-solving domain. And it did so though the sums of money up for grabs were rather trivial. Second, we might have expected their previous

Figure 1. Percent of sequence problems solved by pretrained and naive subjects in each of four different incentive conditions identified on the X-axis. Data are from Schwartz (1982).

Figure 2. Mean number of trial blocks required per problem by naive and pre-trained subjects in each of four different incentive conditions identified on the X-axis. Data are from Schwartz (1982).

experience to make them better rather than worse at discovering the rules. No doubt they had tested and rejected some hypotheses while developing their stereotyped response sequences. For example, they probably learned that how fast they pushed the buttons, or which hand they used, made no difference. The inexperienced subjects did not know these things when the rule discovery task began. Finally, one of the rules that the pretrained students had to discover was the very same rule that had been in effect during their 1000 trials of pretraining. But even here, they were less effective than students who came to this problem completely naive.

We have followed up these initial findings (Schwartz, 1988) to determine the extent to which the negative effects of rewarded pretraining will generalize to problem-solving situations that are different from the pretraining situation. We have evaluated how pretraining affects subjects' ability to assess accurately the degree of correlation between their actions and environmental events, and how pretraining affects subjects' ability to apprehend and evaluate the truth value of statements of different logical form. I will describe each of these studies briefly.

There is a large literature concerned with assessing how accurately people estimate the degree of contingency or correlation between two

environmental events, or between their action and some outcome. In general, the results of such studies indicate that humans are inaccurate in various systematic ways in assessing contingencies (e.g., Allan & Jenkins, 1980; Alloy & Abramson, 1979; Arkes & Harkness, 1983; Crocker, 1981; Einhorn & Hogarth, 1978; Jenkins & Ward, 1965), an outcome that is somewhat surprising in light of the apparent accuracy of pigeons and rats when faced with similar tasks (e.g., Alloy & Tabachnik, 1984; Rescorla, 1972; see Schwartz, 1989 for a textbook review of this literature).

There have been a few attempts to reconcile the two literatures, the most sweeping of which was offered by Alloy and Tabachnik (1984), who suggested that accuracy in contingency assessment is largely determined by whether assessment is based on expectations derived from past experience or upon analysis of current situational data. Humans are perhaps especially susceptible to having theories based on past experience shape expectations that in turn influence the perception of current events. These theory-based perceptions are the source of various biases that have been well documented in the psychological laboratory, in applied clinical or industrial settings, and even in the history of science (e.g., Alloy & Tabachnik, 1984; Einhorn & Hogarth, 1978; Nisbett & Ross, 1980; Platt, 1964; Tweney, Doherty, & Mynatt, 1981).

One particular source of error in contingency judgment has often been called a "confirmation bias." It refers both to people's tendency to seek evidence that can only confirm but not falsify hypotheses, and to overvalue confirming evidence and undervalue disconfirming evidence when it does appear (e.g., Einhorn & Hogarth, 1978; Schwartz, 1982; Wason & Johnson-Laird, 1972; but see Baron, 1985; and Klayman & Ha, 1987, for a critical discussion of confirmation bias). A suggestive account of where this particular bias might come from was offered by Einhorn and Hogarth (1978). They suggested that in natural settings, people rarely have the opportunity to falsify hypotheses. A variety of practical constraints operate to make only partial tests of hypotheses possible, and these partial tests are often of the sort that make confirmation easy and disconfirmation difficult. These practical constraints can lead to habits of inference that prevent people from either collecting the appropriate data, or processing appropriate data in an appropriate fashion, even in circumstances in which practical constraints are absent. It appears from our research that a history of reinforcement for successful responses may create just such habits of inference. When a subject evolves a particular sequence of responses to produce payoffs, his or her orientation is toward producing desirable outcomes rather than true generalizations,

and this tendency persists even in the face of explicit instructions to do otherwise.

We therefore examined whether a history of reinforcement for correct sequences of responses in fact reduces accuracy in a contingency detection task conducted in a different context. The study was a replication of contingency detection experiments reported by Alloy and Abramson (1979). They found that judgments of contingency were inaccurate in several ways. When the environmental events in question were frequent, people judged that they controlled their occurrence, even when in fact they did not. Also, when the environmental events in question were hedonic in nature (wins or losses of money), contingency estimates were affected by the hedonic nature of the events. People gave higher estimates of their ability to control good outcomes than of their ability to control bad ones though the actual degree of control that they had was the same.

More specifically, Alloy and Abramson (1979) gave subjects a series of trials in which they could either push a button or not push it in any 3-second period. At the end of the 3 seconds, either a light would come on or it would not. This was the environmental event the control of which subjects had to estimate. The actual degree of control that subjects had was manipulated by manipulating two conditional probabilities: the probability of the light given a response; and the probability of the light given no response. When these probabilities were equal, subjects had no control. When they were unequal, subjects had control to a degree that could be quantified as the difference between the two conditional probabilities. After a series of such trials, subjects were asked to estimate the degree of control they had, on a 100-point scale.

In our experiment, there were three major groups of subjects. Two of them had been previously exposed to the button-pushing, light matrix task described above (see Figures 1 and 2). Group Sequence experienced a 300-trial session of the task in which correct sequences (four pushes, in any order, on each of the two buttons) earned \$.02, 1 or 2 days prior to the contingency estimation task. Group Rule experienced three sequence problems in which correct sequences had to include four responses on each button and satisfy additional constraints (for example, beginning with two left-button pushes), and they were instructed to discover the rule that determined whether or not they earned points. These subjects were paid \$5 for participating in the session, which occurred a day or two prior to the contingency estimation task. Finally, the third group, Group Naive, had no pretraining.

Subjects were seated at a console that contained a push button and a

red light. They were then read instructions that described the task (see Schwartz, 1988, for verbatim instructions). They were then given a series of five blocks of 40 trials. In each trial, they could either push the button or refrain from pushing it, and either the red light would go after the trial or it would not. At the end of the forty trials, they had to estimate, on a 100-point scale, the degree to which their behavior controlled the illumination of the red light. An estimate of 0 indicated that they thought they had no control, and an estimate of 100 indicated that they thought they had complete control. The conditional probabilities of a red light given a response or not in the five blocks of trials were .75-.75, .75-.50, .75-.25, .50-.50, and .25-.25. Subjects received these blocks in random order.

The results of our experiment are presented in Figures 3-5. Conditional probabilities (of the light coming on after a button press or no button press) are on the X-axis and estimates of control are on the Y-axis. Figure 3 presents control estimates across three series where degree of control actually varied. The data are presented separately for the different pretraining groups. It is clear from the figure that subjects in Group Rule were quite sensitive to the degree of actual contingency, while subjects in Group Sequence were insensitive. The naive subjects were in the middle; their estimate of control in the .75-.25 series was significantly higher than in the other two.

Figure 3. Estimates of control in the conditions in which the actual control that subjects had varied. Data are presented separately for each pretraining group. Data are from Schwartz (1988).

Figure 4 presents similar data from the problem series in which subjects had no control. The data for naive subjects were very much like those reported by Alloy and Abramson; the degree of control was consistently overestimated (since actual control was zero), and the magnitude of the overestimation varied directly with the frequency of the outcome. For Group Sequence, the effect of frequency was diminished as estimates of control were high at all frequencies. For Group Rule, the effect of frequency of outcome was also diminished as all estimates of control were quite modest.

Finally, Figure 5 presents data that compare estimates of control when outcomes are positive with estimates of control when outcomes are negative. Some subjects (Win) started with no money, and won a quarter whenever the light came on. Others (Lose) started with \$10, and lost a quarter whenever the light came on. The Y-axis in Figure 5 represents the difference between Win and Lose estimates of control under the three different degrees of actual contingency. Since contingency estimation should be independent of whether outcomes are positive or negative, accurate performance should yield difference scores close to zero. The scores of Group Rule are close to zero. Naive subjects give substantially higher estimates of control when they are winning than when they are

Figure 4. Estimates of control in the conditions in which subjects had no control but the overall frequency of the outcome varied. Data are presented separately for each pretraining group. Data are from Schwartz (1988).

Figure 5. Difference scores between estimates of control when subjects could only win money and estimates of control when subjects could only lose money. Data are presented separately for different pretraining groups across different contingency conditions. Data are from Schwartz (1988).

losing. For Sequence subjects, the disparity between winning and losing is even higher.

The bias that reveals itself in inaccurate estimates of contingency relations has also been investigated in other contexts. Perhaps the best known of these is the so-called "selection task" (Wason, 1964; Wason & Johnson-Laird, 1972). Subjects are shown four cards, two of which have letters (A or B), and two of which have numbers (1 or 2), and are made to understand that each card has a letter on one side and a number on the other. They are then read a proposition and asked to select the card or cards they would have to inspect to assess the validity of the proposition. The proposition is typically of the conditional form, if p , then q . For example, the proposition might be: "If there is an A on one side, there is a 1 on the other."

Propositions of this form are false only when the antecedent (p) is true and the consequent (q) is false. So faced with four cards displaying A, B, 1, and 2, and the proposition "If there is an A on one side, there is a 1 on the other," the schooled logician would know that the A card (p) and the 2 card (not- q) would have to be examined. The other two cards, while relevant to evaluating some kinds of propositions, are not relevant to this one.

Schooled logicians may know this, but typical college students do not. Very few subjects choose all and only the relevant cards. The most common choices are of card A alone, or of card A and card 1. The evidence on card A is relevant to efforts at either confirmation or falsification of the proposition (a 1 on the hidden side confirms and a 2 falsifies). However, the evidence on card 1 can only confirm (an A on the hidden side confirms and a B is irrelevant). The fact that subjects choose card 1, which can only confirm, and ignore card 2, which can only falsify, is what leads to the conclusion that subjects come to tasks like this with a confirmation bias.

Does pretraining of various kinds enhance or diminish people's tendencies to appreciate the value of some kinds of evidence and ignore the value of other kinds? Does it enhance or impair their ability to apprehend the logical form of various propositions? We asked these questions experimentally by assessing the effects of pretraining on people's behavior when they were required to test propositions that were presented in varying linguistic but equivalent logical forms.

There were three groups of subjects in the experiment. Group Sequence subjects experienced a 300-trial session in which correct sequences earned \$.02. Group Rule subjects were paid \$5 for the session, experienced three sequence problems, and were instructed to discover the rule that determined whether or not they earned points. Group Naive had no pretraining. Within 2 days of their pretraining, all subjects then experienced an identical test phase of the experiment. In the test phase, they were seated in front of the sequence apparatus, and given a series of problems. For each problem, they were given a diagram of the light matrix with one of the 25 squares shaded. Beneath the diagram was a statement, like, for example, "if you go through the shaded square, you get a point," or "to get a point, it is necessary to go through the shaded square." Their task was to test the statement by doing "experiments" with the light matrix. That is, they had to determine whether the statement was true or false by giving themselves trials with the matrix.

Subjects were then given a series of eight problems, each including a matrix diagram with one shaded square and a statement. These materials occupied about half of an 8.5 by 11 inch piece of white paper. The remainder of the sheet was available for subjects to use for taking notes, framing hypotheses, or whatever they thought useful. Four of the eight problems contained statements that were logically equivalent forms of the claim "if p , then q ," with p standing for the shaded square and q standing for the getting of a point. Each of these statements was a claim about *sufficiency*. The other four problems contained statements that

were claims about *necessity*, logically equivalent to “if q , then p .” Specifically, the eight statements were these:

Sufficiency

1. If the light goes through the shaded square I get a point.
2. To get a point, it is sufficient for the light to go through the shaded square.
3. Either the light does not go through the shaded square, or I get a point.
4. The light went through the shaded square only if I got a point.

Necessity

5. If the light does not go through the shaded square, I don't get a point.
6. To get a point, it is necessary for the light to go through the shaded square.
7. Either I didn't get a point, or the light went through the shaded square.
8. I get a point only if the light goes through the shaded square.

Subjects received these eight problems in random order, and for each subject, a random half of both the necessity and the sufficiency statements were true. Because we expected that pretraining might enhance an already existing tendency in subjects to treat questions about necessity as if they were questions about sufficiency, the results of the experiment were analyzed separately for the two types of statements. Figure 6 presents data on accuracy; the percentage of correct conclusions for each group on each subset of the problems. On the problems asking about sufficiency, there were differences between the groups, but they were not statistically significant. On the necessity problems, differences were significant; rule pretrained subjects were significantly more accurate than naive ones, who in turn were significantly more accurate than sequence pretrained subjects.

A second measure of performance was the number of trials subjects took in evaluating each problem, and these data are presented in Figure 7. First, evaluations of necessity took more trials than evaluations of sufficiency. Second, within problem types, there was a significant effect of pretraining. In the case of both types of problems, Group Sequence took significantly more trials than either of the other groups, which did not in turn differ from each other.

When hypotheses make claims about sufficiency, they can be falsified only by trials in which the shaded square is in fact illuminated. The statement that, for example, “if the middle square in the 5×5 matrix is

Figure 6. Percentage correct guesses about the validity of the statements to be tested. Data for statements about sufficiency are on the left, and data for statements about necessity are on the right. Data are presented separately for each pretraining group. Data are from Schwartz (1988).

Figure 7. Mean number of trials per statement taken by subjects prior to making their guesses about statement validity. Data for statements about sufficiency are on the left and data for statements about necessity are on the right. Data are presented separately for each pretraining group. Data are from Schwartz (1988).

illuminated you get a point" is tested only by trials in which that square is illuminated. Conversely, statements about necessity can be falsified only by trials that *do not* illuminate the square in question. To assess the claim that "you get a point only if the middle square is illuminated," one must generate trials that avoid illuminating the middle square.

We evaluated, for each subject, the proportion of trials generated that actually were appropriate tests of the proposition in question. In other words what proportion of trials that evaluated statements about sufficiency were actually tests of sufficiency, and what proportion of trials that evaluated statements of necessity were actually tests of necessity? The data are presented in Figure 8. On sufficiency problems, both pre-trained groups generated a significantly higher proportion of trials that actually tested for sufficiency than did the naive subjects. This might lead to the conclusion that both kinds of pretraining improved performance. However, the data from necessity problems suggest quite a different conclusion. Here, the Rule group was significantly better than the Naive group while the Sequence group was significantly worse. The

Figure 8. Percentage of trials in which the tests subjects generated were appropriate to the claims in the statements being tested. Thus data on the left present the percent of tests of statements about sufficiency that were actually tests of sufficiency, while data on the right present the percentage of tests of statements about necessity that actually were tests of necessity. Data are presented separately for each pretraining group. Data are from Schwartz (1988).

appropriate conclusion from these data seems to be that rule pretraining indeed helps subjects in their analysis of the logical form of various propositions so that they understand what claim is being made and test it correctly. In contrast, sequence pretraining induces subjects to treat all logical claims as claims about sufficiency, and to test them as such. Or alternatively, perhaps it induces them to attempt to produce positive results and ignore logic all together. This will improve the accuracy and appropriateness of subjects' behavior only as long as the claims in question actually *are* about sufficiency. It seems, therefore, that the kind of bias identified and studied by Wason and Johnson-Laird (1972) and by Einhorn and Hogarth (1978) is ameliorated by experience with systematic hypothesis testing and rule discovery and exacerbated by experience generating particular behavior patterns to produce particular outcomes.

In these experiments we have experimental simulations of the historical process I argued for above. Exposure to contingencies of reinforcement creates an efficient, stereotyped pattern of behavior that can be executed with effortless and mechanical precision, just like work on the assembly line. The contingency makes it possible for people to do the right thing, over and over again. Lapses of attention have no cost, because attention is not required. Lack of intelligence has no cost, because intelligence is not required. The people become an extension of the machinery, a realization of Taylor's Scientific Management vision.

That this automatization is achieved at the expense of another potential influence on the nature of the activity becomes apparent when these students are later asked to discover rules, detect contingencies, or test logical claims. Students without pretraining know what rule discovery means. They have been participating in a tradition of rule discovery or problem solving for years, and it is a relatively simple matter to plug this new challenge into one's traditional wisdom from previous ones. The pretrained students are a part of this same tradition. But they have been induced, by their pretraining, to place this particular task outside it. The debilitating effect of this pretraining is modest. But one can imagine that if they were required to engage in this pretraining task for 8 or more hours a day, day after day, week after week, year after year, the effect might be considerably more dramatic. And if everyone around them was engaged in a similarly repetitive activity, the effect might be more dramatic still. For instead of simply failing to locate this task in the problem-solving tradition, that tradition might erode and disappear all together. Adam Smith (1776/1937: 734–735) of all people, captured this possibility most forcefully. He had this concern about the side effects of factory work:

The man whose life is spent in performing a few simple operations . . . has no occasion to exert his understanding, or to exercise his invention in finding out expedients for difficulties which never occur. He naturally loses, therefore, the habit of such exertion and generally becomes as stupid and ignorant as it is possible for a human creature to become.

Exactly so.

Practices and Their Contamination

The research just described explores an example of how activities that are valuable in themselves can be transformed into activities that are simply means to ends that could be attained in other ways. To appreciate better the general character and significance the transformation of value reflected in this research, we need a clearer idea of what it means for a domain of activity to be valuable in itself. What is it that participants strive for when they engage in activities that are not purely instrumental?

This question has been illuminated by Alasdair MacIntyre, in his book *After Virtue* (1981). MacIntyre attempts in that book to reconstruct a moral philosophy, and central to that attempt is the concept of a *practice*. Practices are certain forms of complex and coherent, socially based, cooperative human activities. Among their characteristics are these:

1. They establish their own standards of excellence, and indeed, are partly defined by those standards.
2. They are teleological, that is, goal directed. Each practice establishes a set of "goods" or ends that is internal or specific to it, and inextricably connected to engaging in the practice itself. In other words, to be engaging in the practice is to be pursuing these internal goods.
3. They are organic. In the course of engaging in a practice, people change it, systematically extending both their own powers to achieve its goods, and their conception of what its goods are.

Thus practices are established and developing social traditions that are kept on course by a conception of their purpose that is shared by the practitioners. And most importantly, the goals or purposes of practices are specific or peculiar to them. There is no common denominator of what is good, like utility maximization, by which all practices can be assessed.

We can illustrate the concept of a practice and its significance by considering an example of a practice in some detail. The collection of activ-

ities referred to as “science” is a practice. Sciences are certainly complex, social activities. They establish their own standards of excellence. They have a set of “goods,” the pursuit of which partly defines them. And they develop. The goal of science is to discover generalizations that describe and explain the phenomena of nature.

Different scientific disciplines develop traditions that provide guidance as to which generalizations are worth going after, which methods are best suited for going after them, and which standards should be used for determining whether one has succeeded. Now not all people who do what looks like scientific work are engaged in the practice of science. People who do experiments to achieve impressive publication records are not engaged in the practice. The goods they seek—fame, wealth, status, promotion—are not internal to science. Science is just one means to those goods among many. It is certainly true that people who are pursuing such external goods may do good science; that is, they may contribute to the development of the practice. But they are not themselves practitioners. And if everyone engaged in science were to start pursuing these external goals, the practice of science would cease to exist. The core of the practice of science—the thread that keeps it going as a coherent and developing activity—lies in the actions of those whose goals are internal to the practice.

And these internal goals are not commensurable with other kinds of goals. The scientist does not choose from among a variety of market baskets, each containing some amount of truth and some amount of status and money, the one market basket that maximizes his preferences. One does not bargain away portions of truth for portions of something else, at least not if one is working within the practice of science. But in the experiments I just described, some of the experimental subjects did precisely this. They bargained away truth, or more accurately, the best techniques for discovering truth, in return for money. For subjects who were pretrained to perform a particular task for monetary reward, the problem solving task was not “pure” science. It was an amalgam of truth seeking and money seeking, of doing what will yield a general principle and doing what works. These students struck a compromise between two competing masters when they faced the problem-solving task. Their compromise was not necessarily deliberate, but it was there nevertheless. They forsook traditional methods of doing good science that their untrained colleagues followed, so that they could earn more money.

One does not need our laboratory demonstrations to see this compromise between scientific and economic objectives. Economic considerations have been affecting the behavior of real scientists, doing real science, for years, and continue increasingly to do so. It is not good science to do the same experiment again and again—to repeat what

works. Yet with research success, promotion and the granting of tenure largely determined by rate of publication, many scientists do so. Each experiment is a minor variant on the preceding one, because such mechanical and unimaginative variation is the quickest road to print. It is not good science to decide what to study on the basis of what people are willing to pay for. Yet government agencies are able to manipulate fields of inquiry by shifting funding from one domain to another. It is not good science to keep one's results a secret, keeping others in the dark, or even intentionally misleading them. Yet in areas that are hot, scientists often do this, as a way of protecting claims to priority, even at the cost of scientific progress. Above all, it is not good science to lie—to misrepresent results willfully, or to invent results of experiments that were never conducted. Yet, in the last few years, several examples of blatant falsification have been uncovered at major research institutions. This last perversion of science, presumably in the interest of self-aggrandizement, is especially crippling. Science must proceed on the presumption that its practitioners always tell the truth, even if they are not always successful at finding it. Were this presumption seriously undermined, science would grind to a halt. All experiments would have to be repeated by all interested parties, to make sure of the veracity of published reports.

Spurred by the growing number of research partnerships between universities and industry, concerned scientists worry that it is just a matter of time before corporate concerns start taking control of the university laboratory, dictating which problems are to be studied, who is to be hired to study them, what information is allowed to be publicly disseminated when, and the like. At the very least, even if they do not dictate the direction of research, or encourage distortion, corporate sponsors can be expected to exert substantial control on communication. It does them little good to foot the bill for research if everyone can gain access to its products, through the scientific journals, at the same time they do. If the practice of delaying or withholding publication becomes widespread enough, it will create substantial doubt among scientists about just how accurate and up to date the reports they see in their journals really are (see Schwartz, 1986, for further discussion of this and other examples).

Economic Imperialism and the Destruction of Value

The distinction between internal and external goods, and the concept of a practice more generally, is illuminated by the notion of "economic imperialism" (see Hirsch, 1976; Schwartz, 1986). Economic imperialism is the spread of economic calculations of "interest" to domains that were once regarded as noneconomic. It is the infusion of a practice with the

pursuit of external goods. This pursuit pushes a practice in directions it would not otherwise take, and in so doing, undercuts the traditions that comprise it. Economic imperialism evaluates practices by a common, economic denominator, abandoning the ones that fall short and encouraging the ones that do not, without regard to the internal goods that each practice possesses uniquely. And whenever internal and external goods conflict, economic imperialism moves in the direction of maximizing the latter, sometimes at the cost of eliminating the former.

This is arguably what happened in the factory, as I indicated above. Tailoring, for example, is a practice. There are goods internal to it that help to shape it even though people do it to earn their livelihood. The pursuit of livelihood does not completely determine the character of weaving just because of these other, internal goods. But when tailoring is organized by the assembly line, the goods internal to it disappear. Assembly line work is not a practice. Good assembly line workers show up for work, and work just as hard and as fast as they are told. What they are doing on the line does not matter. There is no developing tradition of which they are a part. The goods of assembly line work are entirely external. This was part of the point of creating the assembly line to begin with. And economic imperialism threatens to turn other practices into the essential equivalent of assembly line work, by moving them always in the direction of maximization of marginal productivity, utility, or profitability, at the expense of their traditional, internal goods.

For MacIntyre, the concept of a practice has a central place in a theory of what it is to be a good person. Good people possess just those characteristics, or virtues, that permit them to engage successfully in practices. The list of these virtues is fairly traditional—justice, honesty, courage, wisdom, respect, constancy, determination, and so on. And the continued existence and development of practices depend on the continued existence of people who possess these virtues. Thus, our judgment of moral worth is bound up with our judgment of the set of practices to which that worth contributes. What then happens to moral worth if the practices disappear, if economic imperialism transforms them into simply means to external goods? If that happens, there is only one practice—the practice of utility maximization. And the good person will be the good utility maximizer. Virtuous, moral people will be indistinguishable from rational, economic agents. By penetrating and transforming the set of practices that comprise human, social life, economic imperialism will have created the conditions under which its conception of human nature is true.

Hirsch (1976) identifies many aspects of modern life that once were largely independent of economic considerations but are now becoming increasingly pervaded by them. One example he discusses is education.

With increasing competition among members of society for good jobs, employers keep erecting new hurdles that must be jumped before job entry is possible. First high school degrees were required, then college degrees, then special training programs, then masters degrees, then doctoral degrees, then doctoral degrees at only a handful of select, certified institutions. The training one receives in these various programs may or may not be relevant to the requirements of the job; relevant training is really not the essential point. What is critical is to make the path to the job arduous enough so that only a few dedicated souls will embark on it.

These hurdles have a profound effect on the way people view education. With education so closely tied to job entry, and job training, it becomes an "investment" in one's future. The money spent on school is expected to be returned, *in kind*, and with interest, later on. One can put a dollar value on a college degree by surveying the wages paid on the jobs to which it gives access. It is easy to imagine people engaged in the following kind of calculations: a degree from Harvard will cost \$80,000. If one took that money and invested it, and entered the job market 4 years earlier than one otherwise would, would the interest on investment coupled with the extra years of earning power compensate for the high-paying jobs forgone? Or perhaps the calculations might go like this: Harvard will cost \$80,000, while the state university will cost \$30,000. Will the job opportunities provided by the Harvard degree pay back the extra \$50,000 invested?

It is easy to see how thinking about education in these terms—as an economic investment—can affect what people want out of education, and thus how they evaluate what they get. If enough people assessed their education in these terms, what actually went on in the college classroom would surely change. Colleges and universities would have to be sensitive to market demand; they would have to provide what students wanted, or the students would go elsewhere. The goal of education would shift from creating well-informed, sensitive, and enlightened citizens to creating skilled workers. This is an example of economic imperialism. To the claim that one cannot put a dollar value on having an educated citizenry comes the reply, of course one can. One simply looks at how much extra salary the education makes possible. Extra salary becomes the yardstick for evaluating the effectiveness of an educational institution. Before long, the institution changes what it does, so that the creation of extra salary potential becomes the goal itself, instead of just a measuring stick.

Another example of economic imperialism is that our everyday social relations—as friends, neighbors, spouses, and parents—are taking on an economic component. In part, this comes from the fact that con-

sumption takes time. If one has money only for the essentials of life, finding the time in which to consume them is not an issue. But if there is money for stereos, video recorders, dinners in nice restaurants, the theater, and vacations, one must find the time to decide which stereo, restaurant, play, or resort to partake of. In addition, one must find the time actually to partake of it. Dinner at home is an hour; dinner out is an evening. No matter how rich a person is, time is a resource that cannot be increased; there are only 24 hours in a day.

The pressure for time to consume has real costs. It produces what Hirsch calls "the economics of bad neighbors." Time spent being sociable is time taken away from consumption. Chatting over the backyard fence or helping a neighbor cut down a tree are actions taken at a cost of using the video recorder, or going into town for a nice dinner. Whether we like it or not, the decision to be sociable becomes an economic decision, another example of the spread of economic considerations to traditionally noneconomic domains. Many people have experienced how much harder it has become to find the time to spend a quiet evening sipping beer and chatting with a few friends. It is becoming increasingly rare for such occasions to develop spontaneously; they must be planned days, or even weeks, in advance. And of course it seems ludicrous to "plan" an evening of casual conversation. So instead it becomes a dinner party. This, in turn only adds to the time pressure, since now food must be purchased, and an impressive meal must be prepared.

In the economic world, people get what they pay for. Certainly, they get nothing more, and vigilance is required to see that they do not get less. People are not in business for their health, after all. So what happens when social "goods" become economic? Presumably, people start getting only what they pay for in social relations as well as economic ones. In the economic world, people are prepared to operate on this assumption. Products come with explicit guarantees, services are provided in accordance with detailed and specific contracts. People enter into exchanges with their eyes open, expecting, and guarding against, the worst. They are not so prepared in the social world. People assume that friends, lovers, families, doctors, and teachers will act with good will, doing, insofar as is possible, what is best for them. As a result, they ask no guarantees, and write no contracts. People trust that part of what it means to be a spouse, lover, parent, doctor, or teacher ensures that people close to them will behave honorably, truthfully, courageously, and dutifully in social interactions.

As social relations become commercialized, however, this assumption grows more and more suspect. Increasingly, people feel the need to have things written down in contracts. Increasingly, they feel the need to be able to hold others legally accountable—whether doctors, lawyers, teach-

ers, or even friends or lovers—to have a club to wield to ensure that they are getting what they pay for out of their social relations.

One might argue that this shift from a dependency on what is implicit in various social relations to what is explicit and contractual is merely a recognition of cold, hard reality. But what the person who makes this argument fails to realize is that the process of commercialization of social relations affects the product. By treating the services of doctors and teachers as commodities being offered to the wary consumer, we change the way doctors doctor and teachers teach. Doctors practice defensively, doing not what they regard as the best medicine, but what they regard as the best hedge against malpractice suits. Medical costs soar, but medical care does not improve. Teachers teach defensively, making sure their students will perform well on whatever tests will be used to evaluate their progress, at the expense of genuine education. Test scores go up, but students are no wiser than before.

There is, in short, a self-fulfilling character to the commercialization of social relations. The more we treat such relations as economic goods, to be purchased with care, the more they become economic goods about which we must be careful. The more that an assumption of self-interest on the part of others governs social relations, the truer that assumption becomes. As Hirsch (1976: 88) has said, “the more that is in the contracts, the less can be expected without them; the more you write it down, the less is taken—or expected—on trust.”

Probably, not even the most committed economist is sanguine about this vision of the world. If it is true that moral traditions depend on practices, and practices can be corrupted by the pursuit of external goods, and the pursuit of external goods is encouraged by economic imperialism, then all we have to do is be vigilant, and keep economic considerations from penetrating into all our practices. By keeping practices relatively pure, we can preserve a proper place for morality in a highly industrialized, productive, and affluent culture.

But choosing to keep practices pure, and economic and noneconomic goods distinct, is not a decision that can be easily made individually. If few others make that choice, and instead enter practices with external orientations, the practices themselves will change so that the pursuit of previously internal goods will no longer be possible. Return for a moment to the experiments in which intelligent college students treated a scientific, problem-solving situation as if it were an assembly line job. All that the experiments really did was trick them. They knew what problem solving was, and what it required, and they knew what mechanical, repetitive activity was, and what it required, and the experiments merely induced them to treat an instance of the former as if it were the latter. It

was a trick that could easily be corrected. But suppose there was no practice of science left as we know it. Suppose science had been completely penetrated by economic imperialism. What help would it be to the students then to be told that they were supposed to be doing science? What would "doing science" mean? The economist may expect that economic imperialism can be exercised in moderation, but the problem is that there may be nothing left to provide that moderation.

Conclusion

Systematic inquiry into the character of human values should help us answer questions like these:

1. Where do values come from?
2. How do values change?
3. How do values influence the actions and life plans of individuals?
4. How are values embodied in social institutions?
5. How does the shape of social institutions influence or constrain the formation of values in individuals?

I have tried to suggest in this chapter that science, as ordinarily understood, is probably not the best way to develop answers to questions such as these. This is because the real answers to questions such as these lie in the particulars. People are situated in particular times and places, are influenced by particular social norms and institutions, and contribute to the development of new social norms and institutions. Attempts to extract what is timeless and universal about the formation of human values from the particulars that are important at any given time or place are likely to lead to distortion. This distortion can be especially serious if historical contingencies are mistakenly identified as natural laws.

Gould (1987) has recently written about the difference between the universal and the historical in the study of geology. Gould refers to the universal as "time's cycle," to highlight its repetition or recurrence. He refers to the historical as "time's arrow," to highlight its uniqueness and directionality. Gould (1987: 15–16) says of the distinction that

Time's arrow and time's cycle is, if you will, a "great" dichotomy because each of its poles captures, by its essence, a theme so central to intellectual (and practical) life that Western people who hope to understand history must wrestle intimately with both—for time's arrow is the intelligibility of distinct and irreversible events, while time's cycle is the intelligibility of timeless order and lawlike structure. We must have both.

We must have both. But what this chapter has tried to suggest and to illustrate is that in the domain of human values, it is time's arrow that should properly do most of the explanatory work.

Acknowledgments

The research and analysis reported in this chapter was supported by a Eugene M. Lang Faculty Fellowship and by grants from the National Science Foundation. Alan Heubert and Heidi McBride provided invaluable assistance in the design, execution, and analysis of the experiments.

References

- Allan, L. G., & Jenkins, H. M. (1980). The judgment of contingency and the nature of the response alternatives. *Canadian Journal of Psychology*, 34, 1-11.
- Alloy, L. B., & Abramson, L. Y. (1979). Judgment of contingency in depressed and nondepressed students: Sadder but wiser? *Journal of Experimental Psychology: General*, 108, 441-485.
- Alloy, L. B., & Tabachnik, N. (1984). Assessment of covariation by humans and animals: The joint influence of prior experience and current information. *Psychological Review*, 91, 112-149.
- Arkes, H. R., & Harkness, A. R. (1983). Estimates of contingency between two dichotomous variables. *Journal of Experimental Psychology: General*, 112, 117-135.
- Baron, J. (1985). *Rationality and intelligence*. New York: Cambridge University Press.
- Crocker, J. (1981). Judgment of covariation by social perceivers. *Psychological Bulletin*, 90, 272-292.
- Deci, E. L. (1975). *Intrinsic motivation*. New York: Plenum.
- Einhorn, H. J., & Hogarth, R. M. (1978). Confidence in judgment: Persistence of the illusion of validity. *Psychological Review*, 85, 395-416.
- Ferster, C. B., & Skinner, B. F. (1957). *Schedules of reinforcement*. New York: Appleton-Century-Crofts.
- Gilbreth, L. (1914). *The psychology of management*. New York: Sturgis and Walton.
- Gould, S. J. (1987). *Time's arrow, time's cycle*. Cambridge: Harvard University Press.
- Greene, D., Sternberg, B., & Lepper, M. R. (1976). Overjustification in a token economy. *Journal of Personality and Social Psychology*, 34, 1219-1234.
- Hirsch, F. (1976). *Social limits to growth*. Cambridge: Harvard University Press.
- Hobsbawm, E. J. (1964). *Labouring men*. London: Weidenfeld and Nicholson.

- Jenkins, H. M., & Ward, W. C. (1965). Judgment of contingency between responses and outcomes. *Psychological Monographs*, 79 (1, Whole No. 594).
- Klayman, J., & Ha, Y. (1987). Confirmation, disconfirmation and information. *Psychological Review*, 94, 211–228.
- Lacey, H. (1988). Realism and the value freedom of science. Unpublished manuscript.
- Lacey, H., & Schwartz, B. (1986). Behaviorism, intentionality and sociohistorical structure. *Behaviorism*, 14, 193–210.
- Lacey, H., & Schwartz, B. (1987). The explanatory power of radical behaviorism. In S. Modgil & C. Modgil (Eds.), *B. F. Skinner: Consensus and controversy* (pp. 165–176). New York: Falmer Press.
- Lepper, M. R., & Greene, D. (Eds.) (1978). *The hidden costs of reward*. Hillsdale, NJ: Erlbaum.
- Lepper, M. R., Greene, D., & Nisbett, R. E. (1973). Undermining childrens' intrinsic interest with extrinsic rewards: A test of the "overjustification" hypothesis. *Journal of Personality and Social Psychology*, 28, 129–137.
- MacIntyre, A. (1981). *After virtue*. South Bend: University of Notre Dame Press.
- Marglin, S. (1976). What do bosses do? In A. Gorz (Ed.), *The division of labour* (pp. 13–54). London: Harvester Press.
- Marx, K. (1955). *The poverty of philosophy*. London: Lawrence and Wishart.
- Nisbett, R., & Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgment*. Englewood Cliffs, NJ: Prentice-Hall.
- Platt, J. R. (1964). Strong inference. *Science*, 146, 347–353.
- Polanyi, K. (1944). *The great transformation*. New York: Rinehart.
- Rescorla, R. A. (1972). Informational variables in Pavlovian conditioning. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 6, pp. 216–282). New York: Academic Press.
- Schwartz, B. (1982). Reinforcement induced behavioral stereotypy: How not to teach people to discover rules. *Journal of Experimental Psychology: General*, 111, 23–59.
- Schwartz, B. (1986). *The battle for human nature*. New York: W. W. Norton.
- Schwartz, B. (1988). The experimental synthesis of behavior: Reinforcement, behavioral stereotypy, and problem solving. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 22, pp. 93–135). New York: Academic Press.
- Schwartz, B. (1989). *The psychology of learning and behavior*. New York: W. W. Norton.
- Schwartz, B., & Lacey, H. (1982). *Behaviorism, science and human nature*. New York: W. W. Norton.
- Schwartz, B., & Lacey, H. (1988). What applied studies of human operant conditioning tell us about humans and about operant conditioning. In G. Davey & C. Cullen (Eds.), *Human operant conditioning and behavior modification* (pp. 27–42). New York: Wiley.
- Schwartz, B., Schuldenfrei, R., & Lacey, H. (1978). Operant psychology as factory psychology. *Behaviorism*, 6, 229–254.
- Smith, A. (1753/1976). *The theory of moral sentiments*. Oxford: Clarendon Press.

- Smith, A. (1776/1937). *The wealth of nations*. New York: Random House.
- Taylor, F. W. (1911/1967). *Principles of scientific management*. New York: W. W. Norton.
- Tweney, R. D., Doherty, M. E., & Mynatt, C. R., Eds. (1981). *On scientific thinking*. New York: Columbia University Press.
- Wason, P. C. (1964). The effect of self-contradiction on fallacious reasoning. *Quarterly Journal of Experimental Psychology*, 16, 30-34.
- Wason, P. C., & Johnson-Laird, P. N. (1972). *The psychology of reasoning*. Cambridge: Harvard University Press.