12

The Technology of Foolishness

James G. March

Choice and Rationality

The concept of choice as a focus for interpreting and guiding human behavior has rarely had an easy time in the realm of ideas. It is beset by theological disputations over free will, by the dilemmas of absurdism, by the doubts of psychological behaviorism, by the claims of historical, economic, social, and demographic determinism. Nevertheless, the idea that humans make choices has proven robust enough to become a major matter of faith in important segments of contemporary Western civilization. It is a faith that is professed by virtually all theories of social policy-making.

The major tenets of this faith run something like this:

> Human beings make choices. If done properly, choices are made by evaluating alternatives in terms of goals on the basis of information currently available. The alternative that is most attractive in terms of the goals is chosen. The process of making choices can be improved by using the technology of choice. Through the paraphernalia of modern techniques, we can improve the quality of the search for alternatives, the quality of information, and the quality of the analysis used to evaluate alternatives. Although actual choice may fall short of this ideal in various ways, it is an attractive model of how choices should be made by individuals, organizations, and social systems.

These articles of faith have been built upon, and have stimulated, some scripture. It is the scripture of theories of decision-making. The scripture is

partly a codification of received doctrine and partly a source for that doctrine. As a result, our cultural ideas of intelligence and our theories of choice bear some substantial resemblance. In particular, they share three conspicuous interrelated ideas:

The first idea is the *pre-existence of purpose*. We find it natural to base an interpretation of human-choice behavior on a presumption of human purpose. We have, in fact, invented one of the most elaborate terminologies in the professional literature: 'values', 'needs', 'wants', 'goods', 'tastes', 'preferences', 'utility', 'objectives', 'goals', 'aspirations', 'drives'. All of these reflect a strong tendency to believe that a useful interpretation of human behavior involves defining a set of objectives that (a) are prior attributes of the system, and (b) make the observed behavior in some sense intelligent *vis-á-vis* those objectives.

Whether we are talking about individuals or about organizations, purpose is an obvious presumption of the discussion. An organization is often defined in terms of its purpose. It is seen by some as the largest collectivity directed by a purpose. Action within an organization is justified (or criticized) in terms of the purpose. Individuals explain their own behavior, as well as the behavior of others, in terms of a set of value premises that are presumed to be antecedent to the behavior. Normative theories of choice begin with an assumption of a pre-existent preference-ordering defined over the possible outcomes of a choice.

The second idea is the *necessity of consistency*. We have come to recognize consistency both as an important property of human behavior and as a prerequisite for normative models of choice. Dissonance theory, balance theory, theories of congruency in attitudes, statuses, and performances have all served to remind us of the possibilities for interpreting human behavior in terms of the consistency requirements of a limited capacity information-processing system.

At the same time, consistency is a cultural and theoretical virtue. Action should be made consistent with belief. Actions taken by different parts of an organization should be consistent with each other. Individual and organizational activities are seen as connected with each other in terms of their consequences for some consistent set of purposes. In an organization, the structural manifestation of the dictum of consistency is the hierarchy with its obligations of coordination and control. In the individual, the structural manifestation is a set of values that generates a consistent preference ordering.

The third idea is the *primacy of rationality*. By rationality I mean a procedure for deciding what is correct behavior by relating consequences systematically to objectives. By placing primary emphasis on rational techniques, we implicitly have rejected – or seriously impaired two other procedures for choice: (a) the processes of intuition, by means of which people may do things without fully understanding why; (b) the processes of tradition and faith, through which people do things because that is the way they are done.

Both within the theory and within the culture we insist on the ethic of rationality. We justify individual and organizational action in terms of an analysis of means and ends. Impulse, intuition, faith, and tradition are outside that system and viewed as antithetical to it. Faith may be seen as a possible source of values. Intuition may be seen as a possible source of ideas about alternatives. But the analysis and justification of action lie within the context of reason.

These ideas are obviously deeply imbedded in the culture. Their roots extend into ideas that have conditioned much of modern Western history and interpretations of that history. Their general acceptance is probably highly correlated with permeation of rationalism and individualism into the style of thinking within the culture. The ideas are even more obviously imbedded in modern theories of choice. It is fundamental to those theories that thinking should precede action; that action should serve a purpose; that purpose should be defined in terms of a consistent set of pre-existent goals; and that choice should be based on a consistent theory of the relation between action and its consequences.

Every tool of management decision that is currently a part of management science, operations research, or decision theory assumes the prior existence of a set of consistent goals. Almost the entire structure of microeconomic theory builds on the assumption that there exists a well-defined, stable, and consistent preference-ordering. Most theories of individual or organizational choice behavior accept the idea that goals exist and that (in some sense) an individual or organization acts on those goals, choosing from among some alternatives on the basis of available information. Discussions of educational policy, for example, with the emphasis on goal-setting, evaluation, and accountability, are directly in this tradition.

From the perspective of all of man's history, the ideas of purpose, consistency, and rationality are relatively new. Much of the technology currently available to implement them is extremely new. Over the past few centuries, and conspicuously over the past few decades, we have substantially improved man's capability for acting purposively, consistently, and rationally. We have substantially increased his propensity to think of himself as doing so. It is an impressive victory, won – where it has been won – by a happy combination of timing, performance, ideology, and persistence. It is a battle yet to be concluded, or even engaged, in many cultures of the world; but within most of the Western world, individuals and organizations see themselves as making choices.

The Problem of Goals

The tools of intelligence as they are fashioned in modern theories of choice are necessary to any reasonable behavior in contemporary society. It is difficult to see how we could, and inconceivable that we would, fail to continue their development, refinement, and extension. As might be expected, however, a theory and ideology of choice built on the ideas outlined above is deficient in some obvious, elementary ways, most conspicuously in the treatment of human goals.

Goals are thrust upon the intelligent man. We ask that he act in the name of goals. We ask that he keep his goals consistent. We ask that his actions be oriented to his goals. We ask that a social system amalgamate individual goals into a collective goal. But we do not concern ourselves with the origin of goals. Theories of individual organizational and social choice assume actors with pre-existent values.

Since it is obvious that goals change over time and that the character of those changes affects both the richness of personal and social development and the outcome of choice behavior, a theory of choice must somehow justify ignoring the phenomena. Although it is unreasonable to ask a theory of choice to solve all of the problems of man and his development, it is reasonable to ask how something as conspicuous as the fluidity and ambiguity of objectives can plausibly be ignored in a theory that is offered as a guide to human choice behavior.

There are three classic justifications. The first is that goal development and choice are independent processes, conceptually and behaviorally. The second is that the model of choice is never satisfied in fact that deviations from the model accommodate the problems of introducing change. The third is that the idea of changing goals is so intractable in a normative theory of choice that nothing can be said about it. Since I am unpersuaded of the first and second justifications, my optimism with respect to the third is somewhat greater than most of my fellows.

The argument that goal development and choice are independent behaviorally seems clearly false. It seems to me perfectly obvious that a description that assumes goals come first and action comes later is frequently radically wrong. Human choice behavior is at least as much a process for discovering goals as for acting on them. Although it is true enough that goals and decisions are 'conceptually' distinct, that is simply a statement of the theory. It is not defense of it. They are conceptually distinct if we choose to make them so.

The argument that the model is incomplete is more persuasive. There do appear to be some critical 'holes' in the system of intelligence as described by standard theories of choice. There is incomplete information, incomplete goal consistency, and a variety of external processes impinging on goal development – including tuition and tradition. What is somewhat disconcerting about the argument, however, is that it makes the efficacy of the concepts of intelligent choice dependent on their inadequacy. As we become more competent in the techniques of the model, and more committed to it, the 'holes' become smaller. As the model becomes more accepted, our obligation to modify it increases.

The final argument seems to me sensible as a general principle, but misleading here. Why are we more reluctant to ask how human beings might find 'good' goals than we are to ask how they might make 'good' decisions? The second question appears to be a relatively technical problem. The first seems more pretentious. It claims to say something about alternative virtues. The appearance of pretense, however, stems directly from the theory and the ideology associated with it.

In fact, the conscious introduction of goal discovery as a consideration in theories of human choice is not unknown to modern man. For example, we have two kinds of theories of choice behavior in human beings. One is a theory of children. The other is a theory of adults. In the theory of childhood, we emphasize choices as leading to experiences that develop the child's scope, his complexity, his awareness of the world. As parents, or psychologists, we try to lead the child to do things that are inconsistent with his present goals because we know (or believe) that he can only develop into an interesting person by coming to appreciate aspects of experience that he initially rejects.

In the theory of adulthood, we emphasize choices as a consequence of our intentions. As adults, or economists, we try to take actions that (within the limits of scarce resources) come as close as possible to achieving our goals. We try to find improved ways of making decisions consistent with our perceptions of what is valuable in the world.

The asymmetry in these models is conspicuous. Adults have constructed a model world in which adults know what is good for themselves, but children do not. It is hard to react positively to the conceit. The asymmetry has, in fact, stimulated a rather large number of ideologies and reforms designed to allow children the same moral prerogative granted to adults – the right to imagine that they know what they want. The efforts have cut deeply into traditional child-bearing, traditional educational policies, traditional politics, and traditional consumer economics.

In my judgment, the asymmetry between models of choice for adults and models of choice for children is awkward; but the solution we have adopted is precisely wrong-headed. Instead of trying to adapt the model of adults to children, we might better adapt the model of children to adults. For many purposes, our model of children is better. Of course, children know what they want. Everyone does. The critical question is whether they

are encouraged to develop more interesting 'wants'. Values change. People become more interesting as those values and the interconnections made among them change.

One of the most obvious things in the world turns out to be hard for us to accommodate in our theory of choice: A child of two will almost always have a less interesting set of values (yes, indeed, a *worse* set of values) than a child of twelve. The same is true of adults. Values develop through experience. Although one of the main natural arenas for the modification of human values is the area of choice, our theories of adult and organizational decision-making ignore the phenomenon entirely.

Introducing ambiguity and fluidity to the interpretation of individual, organizational, and societal goals, obviously has implications for behavioral theories of decision-making. The main point here, however, is not to consider how we might describe the behavior of systems that are discovering goals as they act. Rather it is to examine how we might improve the quality of that behavior, how we might aid the development of interesting goals.

We know how to advise a society, an organization, or an individual if we are first given a consistent set of preferences. Under some conditions, we can suggest how to make decisions if the preferences are only consistent up to the point of specifying a series of independent constraints on the choice. But what about a normative theory of goal-finding behavior? What do we say when our client tells us that he is not sure his present set of values is the set of values in terms of which he wants to act?

It is a question familiar to many aspects of ordinary life. It is a question that friends, associates, students, college presidents, business managers, voters, and children ask at least as frequently as they ask how they should act within a set consistent and stable values.

Within the context of the normative theory of choice as it exists, the answer we give is: First determine the values, then act. The advice is frequently useful. Moreover, we have developed ways in which we can use conventional techniques for decision analysis to help discover value premises and to expose value inconsistencies. These techniques involve testing the decision implications of some successive approximations to a set of preferences. The object is to find a consistent set of preferences with implications that are acceptable to the person or organization making the decisions. Variations on such techniques are used routinely in operations research, as well as in personal counseling and analysis.

The utility of such techniques, however, apparently depends on the assumption that a primary problem is the amalgamation or excavation of pre-existent values. The metaphors – 'finding oneself', 'goal clarification', 'self-discovery', 'social welfare function', 'revealed pre-reference' – are metaphors of search. If our value premises are to be 'constructed' rather than 'discovered', our standard procedures may be useful; but we have no *a priori* reason for assuming they will.

Perhaps we should explore a somewhat different approach to the normative question of how we ought to behave when our value premises are not yet (and never will be) fully determined. Suppose we treat action as a way of creating interesting goals at the same time as we treat goals as a way of justifying action. It is an intuitively plausible and simple idea, but one that is not immediately within the domain of standard normative theories of intelligent choice.

Interesting people and interesting organizations construct complicated theories of themselves. In order to do this, they need to supplement the technology of reason with a technology of foolishness. Individuals and organizations need ways of doing things for which they have no good reason. Not always. Not usually. But sometimes. They need to act before they think.

Sensible Foolishness

In order to use the act of intelligent choice as a planned occasion for discovering new goals, we apparently require some idea of sensible foolishness. Which of the many foolish things that we might do now will lead to attractive value consequences? The question is almost inconceivable. Not only does it ask us to predict the value consequences of action, it asks us to evaluate them. In what terms can we talk about 'good' changes in goals?

In effect, we are asked either to specify a set of super-goals in terms of which alternative goals are evaluated, or to choose among alternatives *now* in terms of the unknown set of values we will have at some future time (or the distribution over time of that unknown set of future values). The former alternative moves us back to the original situation of a fixed set of values – now called 'super-goals' – and hardly seems an important step in the direction of inventing procedures for discovering new goals. The latter alternative seems fundamental enough, but it violates severely our sense of temporal order. To say that we make decisions now in terms of goals that will only be knowable later is nonsensical – as long as we accept the basic framework of the theory of choice and its presumptions of pre-existent goals.

I do not know in detail what is required, but I think it will be substantial. As we challenge the dogma of pre-existent goals, we will be forced to re-examine some of our most precious prejudices: the strictures against imitation, coercion, and rationalization. Each of those honorable prohibitions depends on the view of man and human choice imposed on us by conventional theories of choice.

Imitation is not necessarily a sign of moral weakness. It is a prediction. It is a prediction that if we duplicate the behavior or attitudes of someone else, the chances of our discovering attractive new goals for ourselves are relatively high. In order for imitation to be normatively attractive we need a better theory of who should be imitated. Such a theory seems to be eminently feasible. For example, what are the conditions for effectiveness of a rule that you should imitate another person whose values are in a close neighborhood of yours? How do the chances of discovering interesting goals through imitation change as the number of other people exhibiting the behavior to be imitated increases?

Coercion is not necessarily an assault on individual autonomy. It can be a device for stimulating individuality. We recognize this when we talk about parents and children (at least sometimes). What has always been difficult with coercion is the possibility for perversion that it involves, not its obvious capability for stimulating change. What we require is a theory of the circumstances under which entry into a coercive system produces behavior that leads to the discovery of interesting goals. We are all familiar with the tactic. We use it in imposing deadlines, entering contracts, making commitments. What are the conditions for its effective use? In particular, what are the conditions for coercion in social systems?

Rationalization is not necessarily a way of evading morality. It can be a test for the feasibility of a goal change. When deciding among alternative actions for which we have no good reason, it may be sensible to develop some definition of how 'near' to intelligence alternative 'unintelligent' actions lie. Effective rationalization permits this kind of incremental approach to changes in values. To use it effectively, however, we require a better idea of the kinds of metrics that might be possible in measuring value distances. At the same time, rationalization is the major procedure for integrating newly discovered goals into an existing structure of values. It provides the organization of complexity without which complexity itself becomes indistinguishable from randomness.

There are dangers in imitation, coercion, and rationalization. The risks are too familiar to elaborate. We should, indeed, be able to develop better techniques. Whatever those techniques may be, however, they will almost certainly undermine the superstructure of biases erected on purpose, consistency, and rationality. They will involve some way of thinking about action now as occurring in terms of a set of unknown future values.

Play and Reason

A second requirement for a technology of foolishness is some strategy for suspending rational imperatives toward consistency. Even if we know which of several foolish things we want to do, we still need a mechanism for allowing us to do it. How do we escape the logic of our reason?

Here, I think, we are closer to understanding what we need. It is playfulness. Playfulness is the deliberate, temporary relaxation of rules in order to explore the possibilities of alternative rules. When we are playful, we challenge the necessity of consistency. In effect, we announce – in advance – our rejection of the usual objections to behavior that does not fit the standard model of intelligence.

Playfulness allows experimentation. At the same time, it acknowledges reason. It accepts an obligation that at some point either the playful behavior will be stopped or it will be integrated into the structure of intelligence in some way that makes sense. The suspension of the rules is temporary.

The idea of play may suggest three things that are, in my mind, quite erroneous in the present context. First, play may be seen as a kind of Mardi Gras for reason, a release of emotional tensions of virtue. Although it is possible that play performs some such function, that is not the function with which I am concerned. Second, play may be seen as part of some mystical balance of spiritual principles: Fire and water, hot and cold, weak and strong. The intention here is much narrower than a general mystique of balance. Third, play may be seen as an antithesis of intelligence, so that the emphasis on the importance of play becomes a support for simple self-indulgence. My present intent is to propose play as an instrument of intelligence, not a substitute.

Playfulness is a natural outgrowth of our standard view of reason. A strict insistence on purpose, consistency, and rationality limits our ability to find new purposes. Play relaxes that insistence to allow us to act 'unintelligently' or 'irrationally', or 'foolishly' to explore alternative ideas of possible purposes and alternative concepts of behavioral consistency. And it does this while maintaining our basic commitment to the necessity of intelligence.

Although play and reason are in this way functional complements, they are often behavioral competitors. They are alternative styles and alternative orientations to the same situation. There is no guarantee that the styles will be equally well-developed. There is no guarantee that all individuals, all organizations, or all societies will be equally adept in both styles. There is no guarantee that all cultures will be equally encouraging to both.

Our design problem is either to specify the best mix of styles or, failing that, to assure that most people and most organizations most of the time use an alternation of strategies rather than perseverate in either one. It is a difficult problem. The optimization problem looks extremely difficult on the face of it, and the learning situations that will produce alternation in behavior appear to be somewhat less common than those that produce perseveration.

Consider, for example, the difficulty of sustaining playfulness as a style within contemporary American society. Individuals who are good at consistent rationality are rewarded early and heavily. We define it as intelligence, and the educational rewards of society are associated strongly with it. Social norms press in the same direction, particularly for men. Many of the demands of modern organizational life reinforce the same abilities and style preferences.

The result is that many of the most influential, best-educated, and best-placed citizens have experienced a powerful overlearning with respect to rationality. They are exceptionally good at maintaining consistent pictures of themselves, of relating action to purposes. They are exceptionally poor at a playful attitude toward their own beliefs, toward the logic of consistency, or toward the way they see things as being connected in the world. The dictates of manliness, forcefulness, independence, and intelligence are intolerant of playful urges if they arise. The playful urges that arise are weak ones.

The picture is probably overdrawn, but not, I believe, the implications. For societies, for organizations, and for individuals, reason and intelligence have had the unnecessary consequence of inhibiting the development of purpose into more complicated forms of consistency. In order to move away from that position, we need to find some ways of helping individuals and organizations to experiment with doing things for which they have no good reason, to be playful with their conception of themselves. It is a facility that requires more careful attention than I can give it, but I would suggest five things as a small beginning:

1 We can treat *goals as hypotheses*. Conventional decision theory allows us to entertain doubts about almost everything except the thing about which we frequently have the greatest doubt – our objectives. Suppose we define the decision process as a time for the sequential testing of hypotheses about goals. If we can experiment with alternative goals, we stand some chance of discovering complicated and interesting combinations of good values that none of us previously imagined.

2 We can treat *intuition as real*. I do not know what intuition is, or even if it is any one thing. Perhaps it is simply an excuse for doing something we cannot justify in terms of present values or for refusing to follow the logic of our own beliefs. Perhaps it is an inexplicable way of consulting that part of our intelligence that is not organized in a way anticipated by standard theories of choice. In either case, intuition permits us to see some possible actions that are outside our present scheme for justifying behavior.

3 We can treat *hypocrisy as a transition*. Hypocrisy is an inconsistency between expressed values and behavior. Negative attitudes about hypocrisy stem from two major things. The first is a general onus against inconsistency. The second is a sentiment against combining the pleasures of vice with the appearance of virtue. Apparently, that is an unfair way of allowing evil to escape temporal punishment. Whatever the merits of such a position as ethics, it seems to me distinctly inhibiting toward change. A bad man with good intentions may be a man experimenting with the possibility of becoming good. Somehow it seems to me more sensible to encourage the experimentation than to insult it.

4 We can treat *memory as an enemy*. The rules of consistency and rationality require a technology of memory. For most purposes, good memories make good choices. But the ability to forget, or overlook, is also useful. If I do not know what I did yesterday or what other people in the organization are doing today, I can act within the system of reason and still do things that are foolish.

5 We can treat *experience as a theory*. Learning can be viewed as a series of conclusions based on concepts of action and consequences that we have invented. Experience can be changed retrospectively. By changing our interpretive concepts now, we modify what we learned earlier. Thus, we expose the possibility of experimenting with alternative histories. The usual strictures against 'self-deception' in experience need occasionally to be tempered with an awareness of the extent to which all experience is an interpretation subject to conscious revision. Personal histories, and national histories, need to be rewritten rather continuously as a base for the retrospective learning of new self-conceptions.

Each of these procedures represents a way in which we temporarily suspend the operation of the system of reasoned intelligence. They are playful. They make greatest sense in situations in which there has been an overlearning of virtues of conventional rationality. They are possibly dangerous applications of powerful devices more familiar to the study of behavioral pathology than to the investigation of human development. But they offer a few techniques for introducing change within current concepts of choice.

The argument extends easily to the problems of social organization. If we knew more about the normative theory of acting before you think, we could say more intelligent things about the functions of management and leadership when organizations or societies do not know what they are doing. Consider, for example, the following general implications.

First, we need to re-examine the functions of management decision. One of the primary ways in which the goals of an organization are developed is by interpreting the decisions it makes, and one feature of good managerial decisions is that they lead to the development of more interesting value-premises for the organization. As a result, decisions

should not be seen as flowing directly or strictly from a pre-existent set of objectives. Managers who make decisions might well view that function somewhat less as a process of deduction or a process of political negotiation, and somewhat more as a process of gently upsetting preconceptions of what the organization is doing.

Second, we need a modified view of planning. Planning in organizations has many virtues, but a plan can often be more effective as an interpretation of past decisions than as a program for future ones. It can be used as a part of the efforts of the organization to develop a new consistent theory of itself that incorporates the mix of recent actions into a moderately comprehensive structure of goals. Procedures for interpreting the meaning of most past events are familiar to the memoirs of retired generals, prime ministers, business leaders, and movie stars. They suffer from the company they keep. In an organization that wants to continue to develop new objectives, a manager needs to be relatively tolerant of the idea that he will discover the meaning of yesterday's action in the experiences and interpretations of today.

Third, we need to reconsider evaluation. As nearly as I can determine, there is nothing in a formal theory of evaluation that requires that the criterion function for evaluation be specified in advance. In particular, the evaluation of social experiments need not be in terms of the degree to which they have fulfilled our *a priori* expectations. Rather we can examine what they did in terms of what we now believe to be important. The prior specification of criteria and the prior specification of evaluational procedures that depend on such criteria are common presumptions in contemporary social policy-making. They are presumptions that inhibit the serendipitous discovery of new criteria. Experience should be used explicitly as an occasion for evaluating our values as well as our actions.

Fourth, we need a reconsideration of social accountability. Individual preferences and social action need to be consistent in some way. But the process of pursuing consistency is one in which both the preferences and the actions change over time. Imagination in social policy formation involves systematically adapting to and influencing preferences. It would be unfortunate if our theories of social action encouraged leaders to ignore their responsibilities for anticipating public preferences through action and for providing social experiences that modify individual expectations.

Fifth, we need to accept playfulness in social organizations. The design of organizations should attend to the problems of maintaining both playfulness and reason as aspects of intelligent choice. Since much of the literature on social design is concerned with strengthening the rationality of decision, managers are likely to overlook the importance of play. This is partly a matter of making the individuals within an organization more playful by encouraging the attitudes and skills of inconsistency. It is also a matter of making organizational structure and organizational procedure more playful. Organizations can be playful even when the participants in them are not. The managerial devices for maintaining consistency can be varied. We encourage organizational play by permitting (and insisting on) some temporary relief from control, coordination, and communication.

Intelligence and Foolishness

Contemporary theories of decision-making and the technology of reason have considerably strengthened our capabilities for effective social action. The conversion of the simple ideas of choice into an extensive technology is a major achievement. It is, however, an achievement that has reinforced some biases in the underlying models of choice in individuals and groups. In particular, it has reinforced the uncritical acceptance of a static interpretation of human goals.

There is little magic in the world, and foolishness in people and organizations is one of the many things that fail to produce miracles. Under certain conditions, it is one of several ways in which some of the problems of our current theories of intelligence can be overcome. It may be a good way. It preserves the virtues of consistency while stimulating change. If we had a good technology of foolishness, it might (in combination with the technology of reason) help in a small way to develop the unusual combinations of attitudes and behaviors that describe the interesting people, interesting organizations, and interesting societies of the world.