

The Behavioral and Social Sciences

Herbert A. Simon

The social and behavioral sciences (1) deal with the whole gamut of behavior of our species, neonate and aged, female and male, wealthy and poor, awake and asleep, as individuals and as members of the numerous collectivities to which we belong. If our ignorance about human in-

of important social science discoveries precede the data whose centenary we are celebrating. The 200th anniversary of the publication of Adam Smith's *An Inquiry into the Nature and Causes of the Wealth of Nations* is past, as is the 100th anniversary of Ebbinghaus's seminal

Summary. In the social sciences, as in other sciences, progress is often paced by advances in observational techniques and instruments. This article reviews some of the recent technical progress in the social sciences and then discusses three substantive frontier areas that are particularly exciting at present: evolutionary theory, especially in relation to sociobiology, the theory of human rational choice, and the newly christened discipline of cognitive science. All three claim to provide explanations for broad areas of human behavior.

dividual and social behavior is still impressive, the body of empirical and theoretical knowledge we have today is substantial—much too voluminous to be reviewed in a single article or by a single person. Even limiting the discussion to “frontiers” does not bring the subject within the competence of any scientist known to me. In undertaking this assignment, I have, therefore, more than an ordinary excuse for being illustrative rather than comprehensive in the topics I address, and in sampling in a highly selective fashion, with a strong bias toward areas in which I have some professional knowledge.

The social sciences are often viewed as newcomers. It is little more than a decade since the programs of the National Science Foundation were broadened to embrace them, and since social scientists were admitted on a basis of equality into the National Academy of Sciences (2). Of course their history goes back much farther, for our curiosity about ourselves was surely awakened not much later than our curiosity about the physical and biological world around us.

But even leaving aside the writings of shrewd practitioners—Machiavelli, or the authors of *The Federalist*—a number

of studies of human memory. Nor was mathematics slow in being applied to human affairs. The first mortality tables were published by John Graunt in 1662, and the first calculations of life annuities based on such tables were carried out by the astronomer Halley in 1693, just 7 years after he assisted Newton with the publication of the *Principia*. Cournot's pioneering work in mathematical economics, containing a deep analysis of the phenomena of imperfect competition, appeared in 1838, a generation before Maxwell wrote out the basic equations governing electricity and magnetism. Newcomers the social sciences are not.

Advances in Empirical Knowledge

In the social sciences, as in all science, the progress of our understanding is paced by the progress in our accumulation of a reliable and systematic body of factual knowledge. Large numbers of facts of human individual and social behavior surround us in our daily lives, just as do large numbers of facts of biology and physics. But in neither case does casual observation provide a satisfactory empirical foundation for general descriptive laws. Without systematic observation, including experimentation where that is possible, our samples will be badly biased, our observations will be se-

verely filtered by our preconceptions, and the phenomena will be altogether too tangled and complex for satisfactory analysis (3).

An important part of the history of the social sciences over the past 100 years, and of their prospects for the future, can be written in terms of advances in the tools for empirical observation and in the growing bodies of data produced by those tools. Even our newspapers today provide us with voluminous and reasonably accurate data about our human world that simply were not available a generation or two ago. The census, beginning as an administrative device for fixing taxes and representation, has become a rich source of demographic, economic, and social data of every variety. Important economic series—price levels, unemployment statistics, gross national product, trade balances—are compiled on quarterly or monthly bases. Shifts in public opinions about issues and political personages are polled systematically by means of carefully designed samples. Vital statistics inform us about demographic trends, and about trends in the specific diseases and accidents that are major causes of death and disability. Normed measurements of intellectual skills monitor the effectiveness of our school systems in imparting basic abilities in reading and arithmetic, and have called to public attention a worrisome decline in the average scores of high school graduates (4).

It is perhaps not important that we have more information than our ancestors; it is vitally important that we have better information. A major part of the effort of trained social scientists has gone into improving our techniques for making the kinds of measurements that I have just enumerated. For example, the quality of the data provided by public opinion polls has been steadily advanced, both by the discovery and adoption of better sampling techniques, and by improved design of the questionnaire items themselves—the last made possible by our growing understanding of the sensitivity of replies to the wording of questions and of the conditions under which replies will be volatile or stable over time. We begin to understand when poll responses reflect deeply held attitudes, and when they are superficial reactions to evanescent events (5).

Our sophistication in carrying out psychological experiments in the laboratory has progressed correspondingly. A large part of this progress stems, as in the natural sciences, from advances in laboratory instrumentation—in this case, especially from the growing availability of

The author is Richard King Mellon Professor of Computer Sciences and Psychology in the Department of Psychology, Carnegie-Mellon University, Pittsburgh, Pennsylvania 15213.

electronic devices for presenting sophisticated stimuli, visual and auditory (for example, laboratory computers with display screens or devices for presenting different messages to the two ears), for measuring response times accurately, for recording verbal responses (tape recorders and videotape), for measuring eye movements, and for recording brain waves. With the advances in instrumentation have come a great enrichment in the varieties of human behaviors that can be studied in the laboratory. Our enhanced precision of measurement gives us access to relatively minute events, of the order of a few milliseconds in duration. This has led, for example, to the discovery of a limited-capacity "iconic" memory of very short duration that we had not previously observed (6). Enhanced ability to deal with verbal behavior enables us to bring into the laboratory complex problem-solving tasks (for example, solving a physics problem, making a medical diagnosis) to complement or replace the simpler experiments of a previous generation, involving memorizing nonsense syllables or running finger mazes (7).

Laboratory experimentation has been extended to social as well as individual phenomena. Ingenious experimental designs in social psychology, sometimes making use of stooges and deception, have provided important information about human behavior, but have also raised issues about the ethics of experimentation. The best-known example of such work is Milgram's demonstration that many, if not most, human beings are rather easily persuaded to inflict physical pain on others at the request of an experimenter (8, pp. 567-572). I must register my personal view that we learned more than enough about human psychology (that is, about ourselves) from these experiments to justify the possible (and largely conjectural) damage, in the form of guilt feelings, that might have been inflicted on the subjects who obeyed the experimenter's instructions. We learned, for example, that events like the Holocaust are explainable without attributing to the German culture any psychological characteristics that are not widely shared by other cultures, including our own.

Beginning in the 1930's, there have also been some successful efforts at conducting experiments outside the laboratory, in human social groups and organizations. An early example was the famed series of Hawthorne Experiments, studies of work motivation and productivity carried out in the Hawthorne works of the Western Electric Company (9). More

recently, we have had an important series of experiments to test how a negative income tax would affect the job-seeking and work behaviors of low-income persons (10). (Finding: it would probably not greatly reduce work motivation.) There are enormous difficulties in carrying out such experiments, and in interpreting their results, for experimental controls can only be partial and cannot be maintained over very long intervals of time; but it appears that the field experiment will be used more and more frequently, and will provide empirical data of growing importance.

In the past, the social sciences, except history itself, have tended to be ahistorical, drawing largely upon contemporary data and paying little attention to long-run dynamics. This emphasis, or bias, toward the present has been strongly influenced by the kinds of data that are available—or, more accurately, the kinds that are not available. Maintaining comparability of time series over a long period of years can raise difficult technical and conceptual problems. (How do we compare consumer prices in 1980 with consumer prices in 1925, when consumption patterns were quite different?) Tracking a panel of human respondents over a long interval is costly, and accumulating data banks of information about identifiable individuals can raise important questions of privacy and personal freedom (11).

In spite of these difficulties, the time perspectives of the social sciences are lengthening with the gradual extension of reliable and comparable time series. The course of prices and productivity can be traced with reasonable accuracy for a century or more. A panel of gifted children, first tested by Terman 50 years ago, has been followed to retirement age, providing priceless information about human development over almost the entire life cycle (12). Comparable polling data on voting preferences and on levels of work satisfaction extend back over a quarter-century (13). (Finding: there has been no noticeable upward or downward trend in work satisfaction over this period.) We can now also assess empirically how much social mobility there is in our society from data about the relation between occupations of individuals and their parents. We are gradually accumulating the data required to provide the social sciences with an adequate historical perspective, although there is a need for a great deal more such data collection, particularly in noneconomic domains.

At the same time the discipline of history (sometimes under the banner of "cliometrics") has been absorbing the

statistical and quantitative techniques of the other social sciences, applying them, for example, to improve methods for sampling historical data. Social science theories have been tested in historical contexts, as in the still controversial but gradually converging analyses of the economics of slavery (14).

In these ways, the social sciences have been acquiring the ability to deal with phenomena whose dynamics are measured in years and decades, as well as phenomena whose durations are only tens or hundreds of milliseconds. The former are of primary concern to historians, sociologists, political scientists, and economists dealing with broader institutional phenomena and issues; the latter are the concern of experimental psychologists working toward the as-yet-elusive goal of linking our growing understanding of human information processing to the underlying neurological structures and processes that provide the human mind with its biological substrate (8, chaps. 2 and 8).

Some Frontiers

Turning from the methods to the substance of the social sciences, I am going to select (somewhat arbitrarily) three "frontiers" that are almost certain to be important areas of activity and progress over the coming decades. Undoubtedly, other areas might have been selected instead of these three, but the ones I have chosen are focal and will illustrate suitably the flavor of ongoing research activity.

Frontiers are easily identified by the intellectual excitement they engender. The three areas I shall discuss all amply exhibit that excitement. One is evolutionary theory, especially in its relation to recent developments in sociobiology. A second is the theory of human rational choice, especially in the forms given it by the disciplines of economics and statistics. A third is the newly christened discipline (or interdiscipline) of cognitive science. All three, today, attract much interest across almost the whole of the social sciences; and all three show strong imperialistic tendencies: making claims to provide explanations for very broad areas of human behavior.

Evolutionary Theory

The idea of social evolution was popularized by Herbert Spencer and others a century ago as a sort of corollary to biological evolution. Later it became some-

what unfashionable, both because the social analogs to the biological mechanisms of mutation and selection were never clearly specified, and because some of the early social evolutionists introduced evaluative overtones by identifying "evolution" with "progress."

Recently, however, theories of social evolution have taken on new life in what is to be hoped are new and improved forms. The revival appears to have several independent points of origin. As one source, some economists introduced a "survival of the fittest" argument to buttress the claim that business firms maximize their profits—more exactly, that only those business firms that happen to behave as though they are profit maximizers will survive competition. More recently, widespread attention was drawn to evolutionary theories and their implications for the social sciences by E. O. Wilson's *Sociobiology* (15) and especially by its final chapter, which extrapolated from animal evidence to broad and controversial hypotheses about human social behavior.

Three topics will illustrate current applications of evolutionary theory in the social sciences: the "survival of profit maximizers" argument, evolutionary models of the dynamics of business firm growth, and the current debate about evolutionary selection of traits of "egoism" and "altruism" (16).

Survival of profit maximizers. The theory of the business firm provides the foundation for economic theories of markets and economic equilibrium. According to the classical theory, the firm behaves in such a way as to maximize its profits in the environment of prices, production costs, and demand schedules in which it finds itself.

The postulate of profit maximization is a powerful assumption, having strong implications for the theory of markets that can not be derived without it. Thus a great many of the classical and neo-classical theorems showing the existence and stability of market equilibrium and the social (Pareto) optimality of the equilibrium positions depend on this postulate, as do predictions of rates of investment and estimates of the respective contributions of capital and technological innovations to economic development (17). Much is at stake for theoretical economics and its application to public policy in determining whether the assumption of profit maximization does describe human behavior accurately, or if it does not, whether better assumptions can be found to replace it.

Empirical work in economics is mainly based on public data gathered for non-

research (governmental and business) purposes—data that are mostly of a highly aggregated sort. Typical examples are census data and information published by taxing bodies. Inquiries that carry economists inside business firms to observe how decisions are actually made are quite rare, with the result that our evidence about profit maximization is mostly of a very indirect kind. Moreover, such direct empirical evidence as we do have on the decision-making processes of business executives largely contradicts the hypothesis that business firms behave explicitly as profit maximizers (18).

In the face of this evidence, it has been proposed that the forces of competition, including competition for the investment capital that permits growth, ensure that only the most efficient businesses—those that are de facto or unwitting profit maximizers—will survive and grow. They can not and do not carry out the profit-maximizing calculations, but they survive because they happen to behave as if those calculations had been made (19).

Today, this argument is surrounded by controversy. We have not yet clearly identified the economic analogs of mutation and inheritance that permit business firms to find profitable innovations and that select improved practices for survival. A more serious difficulty is that the uses often made of the evolutionary theory (that is, to support classical theories of market equilibrium) require that the system attain global maxima and not merely local maxima. But most forms of natural selection work by myopic, hill-climbing improvement: they can drive a system to the top of the nearest hill, but cannot guarantee that it will ever scale the highest hill.

When economists have examined the empirical characteristics of business firms, and especially their distribution by size, what they observe does not seem to correspond to the static equilibrium of profit maximization, but to a dynamic steady state. Contrary to what would be inferred from classical theory, firms of the most diverse sizes compete in any given industry, with very little difference in either profitability or average rates of growth between small firms and large, but with a high degree of stability in the distribution of relative sizes (20).

Thus it remains an open question whether evolutionary theory can provide a satisfactory foundation for classical theories of the firm, and this domain is likely to remain for some time an active area of theoretical and empirical inquiry.

Economic growth. An alternative way of applying evolutionary theory to the

behavior of business firms, without commitment to the validity of the classical theory, is to model growth (which may, of course, incorporate competitive pressures) as an explicit dynamic process. The potential fruitfulness of this approach is suggested by the fact that firm-size distributions very similar to those that have been observed empirically can be derived as the steady-state distributions of stochastic processes incorporating simple, plausible assumptions about growth. The key assumption in such derivations, which is consistent with the empirical data, is that (percentage) rate of growth is uncorrelated with present firm size (21).

Winter and Nelson (22) have experimented with evolutionary models, incorporating specific assumptions about processes for generating technological improvement and processes to determine the rate of new investment, to simulate by computer the growth of firms in an industry and to observe the effects of different assumptions about technological innovation upon the growth patterns. The important processes that they model include the effects of the rate of profit on the search for improved technology and on new investment. The dynamic processes the Winter-Nelson models postulate produce patterns of growth and distributions of business size which, like the more aggregative models mentioned above, fit the empirical data well.

Thinking about social growth processes in evolutionary terms does not imply imitating slavishly the mechanisms of Darwinian evolution. Social evolution can obviously be Lamarckian—acquired improvements can persist and can be copied. Nor is there anything in social structures that quite corresponds to the genotype-phenotype distinction. Hence, while the Darwinian metaphor has been stimulating and useful, the current interest in evolutionary models of social phenomena demands a careful building of theory from the ground up, with the closest attention to the initial assumptions about mechanism. At least a generation's work lies ahead to explore the potential of this direction of inquiry.

Altruism and egoism. A key issue raised by recent writings in sociobiology is the respective role played in human behavior by altruism and egoism (23). In the simplest form of Darwinian theory, there is little room for the survival of altruism as a motive. Those who look out most effectively for themselves and their offspring will have the highest reproductive ratios. Of course the kind of altruism known as "enlightened self-interest" is not excluded, and has in fact been exam-

ined extensively by sociologists and anthropologists under the rubric of exchange theory. Reciprocity, at least over extended time intervals, is seen to be fundamental to most forms of social interaction, and is not limited to economic transactions.

Darwinism leaves open a small cranny for another form of altruism. Since genes and not organisms are the units of evolution, willingness to sacrifice one's life for the lives of two or more siblings or eight or more first cousins is adaptive. Hence, organisms may be selected for a modest measure of altruism, but even here, blood is thicker than water.

Uncritical transfer of these Darwinian models to social evolution can lead to a dog-eat-dog picture of the human condition. However, a number of social scientists have pointed out that matters are not nearly this simple. A social system may influence, in many ways, the relative reproduction rates of its participants. In particular, it can reward altruistic behavior and punish selfish behavior. Hence it is entirely possible to have positive selection for cooperative and altruistic behavior in such a system. Moreover, if we take the social system itself as the unit of competition for survival, systems that have selected for altruism (that is, altruism within the group) may grow more rapidly than systems whose members are purely selfish.

Only the first halting steps have been taken toward modeling natural selection in complex systems with reproductive units at several distinct levels—individuals, groups, societies—and further progress, both theoretical and empirical, along this path appears to be one of the most intellectually challenging and socially important goals for social science research today.

Human Rationality

One of the crowning achievements of the social sciences in the past two or three generations has been to elucidate the concept of rationality, a concept that is central to understanding human behavior. This is not to say that actual behavior is always or even usually rational in all of its aspects. But human behavior is manifestly directed toward goals and the satisfaction of wants and needs, hence rational insofar as its direction is appropriate. In fact, in line with our previous discussion of evolution, we would expect natural selection to favor mechanisms that enhance an organism's capacity for rationality. Even madness has its method, as Freud showed us when he

analyzed the myopic adaptiveness of many neurotic and psychotic manifestations.

Research on rationality has two aspects, one normative and the other descriptive. On the one hand, a theory of rationality can be used to infer and prescribe right (that is, rational) courses of action. The theory of the firm we discussed in a previous section might be taught to business school students in this vein: instructing them how to make profit-maximizing decisions. On the other hand, a theory of rationality might claim to describe actual human behavior, as has sometimes been claimed of the classical theory of the business firm. Contemporary developments in the theory of rationality are relevant to both the normative and descriptive applications of the concept.

SEU theory. The core of the formal treatment of rationality is the so-called subjective expected utility (SEU) theory, whose formal axiomatization was achieved shortly after World War II (24). In SEU theory, the rational actor is assumed to have a consistent preference ordering of all possible states of the world, and a prior probability distribution of exogenous events. This probability distribution, together with the actor's strategy, determines, in turn, the probabilities of various states of the world occurring. Under these assumptions, it can be shown that if the actor chooses consistently, in a plausible sense, among all possible contingencies, then one can assign to each possible state of the world a cardinal number (utility) in such a way that the actor will choose as if he or she were maximizing the expected value of this utility.

Stated less precisely, the SEU axioms imply that the actor maximizes his or her expected utility in the light of subjective estimates of the probabilities of events. The purpose of the more convoluted statement of the formal theory is to preserve the operationality and empirical testability of the axioms. Given a sufficient number of observations of the choices of a rational actor, it is possible (in principle) to infer from these choices alone whether the actor is behaving consistently and, if he or she is, it is possible to estimate simultaneously the actor's subjective probability distribution and utility function.

The importance of the formal SEU theory is at least threefold. First, it provides a precise (and intuitively plausible) definition of the concept of rationality, and thereby also a clear standard with which alternative definitions can be compared. Second, it clarifies the extent to

which the assumption of human rationality is an empirical assumption that can be tested with data about choices. Third, it provides a conceptual framework for normative methods of real-world rational decision-making.

The axiomatization of the SEU theory has stimulated numerous experiments aimed at testing it empirically. A fair summary of the findings of these experiments is that actual human choices depart radically from those implied by the axioms except in the simplest and most transparent of situations (25). Humans are unable to choose consistently in the face of even moderate complexity or uncertainty (18). In spite of its failure as a description of human choice, the SEU theory provides a useful framework for categorizing and diagnosing the nature of the departures of actual choice from the requirements of rationality. Kahneman and Tversky (25) have been particularly active and successful in recent years in developing a taxonomy of deviations from rationality and proposing mechanisms of human choice to account for the deviations. Thus, there is emerging a theory of bounded rationality or procedural rationality that describes man as an organism of limited computational ability and possessing limited information and limited imagination, seeking to survive in a world rich in complexity.

Normative applications of SEU theory have shown vigorous growth in recent years. Within mathematical statistics and its applications, the Bayesian approach (the prior probabilities of SEU theory are usually attributed to the Reverend Thomas Bayes) has gained ground, although in the face of formidable mathematical difficulties in translating it into practice (26). At a slightly more informal level, Raiffa and his colleagues have made many applications to practical problems of a decision-tree approach that begins by assigning prior probabilities to events, and then proceeds to estimate expected values for alternative strategies (27). These techniques have been applied to such diverse problems as determining oil-well prospecting strategy and evaluating the desirability of cloud seeding as a means for decreasing the violence of hurricanes.

Beyond these specific applications of SEU theory, there has been since World War II an enormous development of mathematical tools to aid decision-making, which has spawned the new discipline of operations research or management science (the two terms are essentially synonymous). Among these tools are linear programming, which provides algorithms for maximizing a linear

form subject to linear constraints; integer programming, which is similar, but restricts solutions to be integers; dynamic programming; queuing theory; heuristic search; and a number of others (28). While not all these methods fall strictly within the specifications of SEU theory, they are much in the same spirit, seeking "optimal" solutions to decision problems. They are widely applied today in business and government decision-making, especially at middle management levels, and continue to be the objects of intensive study directed toward developing more powerful computational algorithms so that more complex problem situations and wider domains of application can be handled.

Game theory. The SEU model of rationality breaks down in situations where there are two or more actors, each trying to maximize a different utility function. In 1944, in a celebrated work entitled *The Theory of Games and Economic Behavior*, von Neumann and Morgenstern developed a theory of competitive games which, though it fell far short of solving all of the problems, has provided the same kind of authoritative framework for subsequent discussion and analysis that SEU theory has provided for one-person choice situations (29).

For two-person zero-sum ("what I gain, you lose") games, the solution is straightforward enough: each player should adopt a particular mixed (probabilistic) strategy that will prevent his opponent from gaining by outguessing him. Such strategies generally involve bluffing. If the calculations required to find the optimum strategy were not too complex for humans, the game of poker would be an ideal environment for application of this theory.

Difficulties arose in trying to extend the theory to non-zero-sum games, and to games involving more than two players, which have continued to be very active areas of theoretical, and to some extent empirical, research. The Prisoners' Dilemma is an important and instructive example of the two-person non-zero-sum game (30). In this kind of game, each player chooses between two strategies, A and B. If both choose A (for example, not confessing to a crime), they get off with light penalties. If both choose B (confessing), they receive identical rather heavy punishments. If one chooses A and the other B, the one who rats receives a reward, while the other receives a very heavy punishment. What is the rational strategy?

The Prisoners' Dilemma game has obvious applications to competitive situa-

tions in international relations and in business, and more generally, to the problem of public goods not priced in the market. What we should like is a plausible definition of rationality that would recommend the cooperative, non-aggressive strategy A to both players. But it is extremely difficult to explain how such an equilibrium can be stable, even if the game is played repeatedly between the players. Each player has a strong incentive to betray his opponent before he is betrayed. Recently Radner (31) has shown that if the players aim at a satisfactory solution rather than an optimal one, the cooperative strategy will be stable; but there is still no consensus that this way of looking at the situation, or some other, is the correct one—if, indeed, "correct" has any meaning in this context.

In the case of multiperson games, additional difficulties arise out of the multitudinous possibilities for forming coalitions among the players. For most games, no single coalition-forming strategy appears to dominate, unequivocally, alternative strategies. Hence, the definitions of rationality that have been proposed have been almost as numerous as the researchers in this field (32).

Finally, the results in game theory and SEU theory are closely related to the celebrated Impossibility Theorem of Arrow (33). Arrow showed that it is impossible to define a welfare function for a democratic society satisfying all of a small list of conditions that would appear to be natural conditions to impose on such a function. In particular, there is no obvious way to convert to consistency the utilitarian slogan of "the greatest good for the greatest number."

Cognitive Science

Over the past quarter-century, no development in the social sciences has been more radical than the revolution—often referred to as the information processing revolution—in our way of understanding the processes of human thinking. At mid-century, behaviorism dominated experimental psychology. Behaviorism was suited to the predominately positivist and operationalist views of the methodology and philosophy of science, and seemed to provide some guarantee against metaphysical "mentalist" explanations of human behavior. The price paid for these qualities was to confine experimental psychology to relatively simple memory and learning experiments, and to a preoccupation with labo-

ratory rats rather than humans engaged in complex thinking and problem-solving tasks.

A quarter-century later, the picture had changed radically. Experimental psychology had achieved a new sophistication and a new confidence both in studying with precision simple, fundamental mental processes (for example, reaction times and short-term memory capacities), and in bringing into the laboratory professional-level cognitive tasks like chess playing, solving mathematics or physics problems, understanding natural language, or making medical diagnoses. Moreover, the analysis and explanation of all of these diverse sorts of processes had been brought within a general paradigm, the information processing paradigm—without loss of operability, and with a great gain in precision and rigor.

The digital computer and related mathematical and engineering developments in information theory and control theory played a key role in triggering the new paradigm as psychologists and linguists came into contact with these new technologies during and after World War II (7). What these devices and theories demonstrated was that a physical system whose basic currency is information embodied in symbols, rather than energy, can be used to make decisions and to carry on most or all of the processes that, in humans, we call thinking.

By symbols I mean patterns, whether of electromagnetism or ink, that can be stored, copied, compared, read from external sources, and written on external records. A computer is a symbol system whose underlying substrate of symbol-manipulating processes (its instruction code) is specified exactly. Hence, programming a computer to play chess, to diagnose diseases, to solve puzzles and problems, to learn concepts, or to interpret language provides a proof of the sufficiency of such a symbol system and such processes for carrying on thought-like activity.

Almost as soon as computers were invented, they were seen to provide a new metaphor for the brain, largely supplanting the earlier "switchboard" metaphor. But by 1956 they had acquired more than a metaphoric significance as computer programs were written to compose music, to play chess and checkers, to parse sentences, and to discover proofs for theorems. Still, these achievements left open the question of whether the processes used by computers to perform such tasks bore any resemblance to human processes, or whether they were quite

different. After all, the mechanisms that enable airplanes to fly are quite different, in most regards, from the mechanisms used by birds.

The effort since the 1950's to understand the relation between complex information processing in computers and human thinking has created an entirely new scientific discipline—artificial intelligence—and has brought about a complete reformulation of theory in cognitive psychology and revealed basic communalities among several social sciences (notably between linguistics and experimental psychology) that had been following parallel but largely independent paths. It has had a substantial, and growing, impact on philosophy also, which is increasingly discovering that the computer “metaphor” casts significant light on classical issues of epistemology. The new field of cognitive science (it is perhaps premature to call it a “discipline,” although it has already established a professional organization and its first journal) is supplying a meeting place and communication channels for these and other interests that are converging around the information processing paradigm (34).

Artificial intelligence, associated in academic organization with computer science departments rather than with the social sciences, is concerned with learning how computers can be programmed to do “smart” things. Many artificial intelligence programs (for example, the most powerful contemporary chess-playing programs) make intensive use of the brute strength and speed of computers, and do not imitate at all closely the human processes for performing the same tasks. Such programs belong to the class of airplanes, rather than birds—imitations of function but not of process. But in many cases it turns out to be effective to incorporate in artificial intelligence programs many of the same selective heuristics that humans use in order to avoid interminable searches through immense problem spaces. For this reason, research in artificial intelligence continues to be closely associated with work on computer simulation of human cognitive processes, although it is possible that in the future the two fields may diverge.

Computer programs that simulate human cognitive processes are theories of those processes in the same way that appropriate systems of differential equations are theories of physical processes. To test these cognitive theories, a wide range of experiments are conducted with human subjects. One technique now

used extensively is to have human subjects perform difficult tasks while thinking aloud, and to compare data from the thinking-aloud protocol with the output of the computer simulation programs performing the same tasks (7). More traditional experimental techniques can also be used to compare computer simulation theories with human behavior. As this work has cumulated, we are now able to give a rather extensive account and explanation of human thinking in a significant range of problem-solving, pattern-finding, and concept-attaining situations (35).

The long-run strategy of research in information processing psychology includes both gradually extending the range of human complex performance that can be accounted for in this way, and at the same time linking these theories of complex behavior to theories and experimental data that refer directly to the basic elementary information processes and their organization. Thus, the complex theories increasingly incorporate assumptions about short-term memory and the organization of long-term memory that can be tested directly by experiments, using simpler tasks, on some of the component processes. This two-prong strategy is bringing about a gradual convergence, which had been absent from the earlier history of psychology, between the mainstream of experimental psychology, focusing on simpler tasks, and research on so-called “higher mental processes.” Both fit comfortably in the same information processing framework.

The long-range strategy of psychological research calls for a second bridge-building exercise that is less far along, and whose immediate prospects are still doubtful. Since the human mind resides in the brain, we cannot be satisfied with our explanations of human thinking until we can specify the neural substrates for the elementary information processes of the human symbol system. Of these connections we know next to nothing. While an enormous amount of knowledge has been gathered about brain structures and functions at chemical and neurological levels, we still do not even know the physiological basis for long-term or short-term memory—whether it involves macromolecules, neuronal circuits, some combination of these, or something entirely different (8). We are in a position similar to that of 19th-century chemistry, which had developed an extensive theory of chemical combination long before that theory could be linked to the physics of atoms. My crystal ball

does not reveal to me when we are going to be able to begin the design and construction of the vital bridge between information processing psychology and neurophysiology.

Conclusion

I must plead limitations of space and knowledge as my excuses for ending my account at this point. Many of my social science colleagues will find their favorite subjects missing here, but these pages are not written primarily for them, and in any event, I did not promise more than samples. I have tried to choose those samples both for their intrinsic significance and for my knowledge of them. The latter criterion has undoubtedly caused a bias toward the “behavioral”—that is, toward individual human behavior—and away from the “social”—the behavior of such aggregates as families, organizations, and nations. The agenda of ongoing work in the social and behavioral sciences has been revealed here in only the most fragmentary way. But I hope these fragments will provide some glimpse into the excitement and significance of the whole.

References and Notes

1. From now on, they will simply be called “social sciences,” their traditional name.
2. Committee on the Social Sciences in the National Science Foundation, *Social and Behavioral Science Programs in the National Science Foundation* (National Academy of Sciences, Washington, D.C., 1976).
3. Two of the social sciences, statistics and social psychology, have made especially important contributions in recent years to our understanding of the requirements for objective, unbiased sampling of social phenomena, and of the nature of the perceptual and inferential biases that make our naive interpretations of experience so unreliable. For the former, see, for example, W. H. Kruskal and J. M. Tanur, Eds., *International Encyclopedia of Statistics* (Free Press, New York, 1977), vol. 2, pp. 1071–1093; and for the latter, D. Kahneman and A. Tversky, *Psychol. Rev.* **80**, 237 (1973).
4. W. Wintz, chairman, *Report of the Advisory Panel on the Scholastic Aptitude Test Score Decline* (College Entrance Examination Board, New York, 1977).
5. S. Oskamp, *Attitudes and Opinions* (Prentice-Hall, Englewood Cliffs, N.J., 1977); J. A. Sonquist and W. C. Dunkelberg, *Survey and Opinion Research: Procedures for Processing and Analysis* (Prentice-Hall, Englewood Cliffs, N.J., 1977).
6. G. Sperling, *Psychol. Monogr.* **74**, whole No. 11 (1960).
7. A. Newell and H. A. Simon, *Human Problem Solving* (Prentice-Hall, Englewood Cliffs, N.J., 1972); J. Larkin, J. McDermott, D. P. Simon, H. A. Simon, *Science* **208**, 1335 (1980).
8. See the discussion of Milgram's work in P. H. Lindsay and D. A. Norman, *Human Information Processing* (Academic Press, New York, 1972).
9. F. J. Roethlisberger and W. J. Dickson, *Management and the Worker* (Harvard Univ. Press, Cambridge, Mass., 1939). In recent years there has been some reinterpretation of the experiments, but though flawed, they were pioneering efforts.
10. *The New Jersey Income-Maintenance Experiment*, vol. 1, D. Kershaw and J. Fair, Eds., vols. 2 and 3, H. W. Watts and A. Rees, Eds. (Academic Press, New York, 1976–1977).

11. See, for example, B. R. Clarridge, L. L. Sheeky, T. S. Hauser, in *Sociological Methodology*, 1978, K. F. Schuessler, Ed. (Jossey-Bass, San Francisco, 1977), pp. 185-203.
12. L. M. Terman, Ed., *Genetic Studies of Genius* (Stanford Univ. Press, Stanford, Calif., 1925-1959), vols. 1 to 4.
13. U.S. Department of Labor, *Manpower Research Monograph No. 30* (Government Printing Office, Washington, D.C., 1974).
14. R. W. Fogel and S. L. Engerman, *Time on the Cross* (Little, Brown, Boston, 1974); H. G. Gutman, *Slavery and the Numbers Game* (Univ. of Illinois Press, Urbana, 1975).
15. E. O. Wilson, *Sociobiology* (Harvard Univ. Press, Cambridge, Mass., 1975).
16. J. Hirshleifer, *J. Law Econ.* **20**, 1 (1977).
17. K. J. Arrow and F. H. Hahn, *General Competitive Analysis* (Holden-Day, San Francisco, 1971).
18. H. A. Simon, *Am. Econ. Rev.* **69**, 493 (1979).
19. M. Friedman, *Essays in Positive Economics* (Univ. of Chicago Press, Chicago, 1953), pp. 16-30.
20. To a very close approximation the logarithms of the sizes of business firms vary linearly with the logarithms of their rank in the industry or economy—a regularity usually known in economics as the Pareto Law, and observed also in a wide variety of other social phenomena (for example, city size distribution). See Y. Ijiri and H. A. Simon, *Skew Distribution and the Sizes of Business Firms* (North-Holland, Amsterdam, 1977).
21. These phenomena provide an excellent example of the importance of having hard data against which to check widely held popular beliefs. Contrary to popular belief, the data show that large companies do not grow, on average, more rapidly than small companies. A. Singh and G. Whittington, *Growth, Profitability and Valuation* (Cambridge Univ. Press, Cambridge, 1968); J. Steindl, *Random Processes and the Growth of Firms* (Hafner, New York, 1965).
22. R. R. Nelson and S. G. Winter, *Bell J. Econ.* **9**, 524 (1978).
23. For a wide-ranging discussion of the issues of this section see *Behav. Sci.* **24**, whole No. 1 (1979); also see Wilson (15).
24. The definitive treatment is L. J. Savage, *The Foundations of Statistics* (Wiley, New York, 1954).
25. D. Kahneman and A. Tversky, *Psychol. Rev.* **80**, 237 (1973).
26. G. E. Box, *Bayesian Inference in Statistical Analysis* (Addison-Wesley, Reading, Mass., 1973).
27. H. Raiffa, *Decision Analysis* (Addison-Wesley, Reading, Mass., 1968).
28. J. J. Moder and S. E. Elmaghraby, Eds., *Handbook of Operations Research* (Van Nostrand-Reinhold, New York, 1978), vols. 1 and 2; H. M. Wagner, *Principles of Operations Research* (Prentice-Hall, Englewood Cliffs, N.J., ed. 2, (1975); see also H. J. Miser, *Science* **209**, 139 (1980).
29. J. von Neumann and O. Morgenstern, *The Theory of Games and Economic Behavior* (Princeton Univ. Press, Princeton, N.J., 1944).
30. A. Rapoport and A. M. Chammak, *Prisoner's Dilemma* (Univ. of Michigan Press, Ann Arbor, 1965).
31. R. Radner, personal communication.
32. An excellent example of recent research in this domain is J. C. Harsanyi, *Rational Behavior and Bargaining Equilibrium in Games and Social Situations* (Cambridge Univ. Press, Cambridge, 1977).
33. K. J. Arrow, *Social Choice and Individual Values* (Wiley, New York, 1951).
34. The first volume of *Cognitive Science* was published in 1977. Papers of the first National Conference on Cognitive Science will be published in the 1980 volume.
35. J. R. Anderson, *Language, Memory and Thought* (Erlbaum, Hillsdale, N.J., 1976); H. A. Simon, *Models of Thought* (Yale Univ. Press, New Haven, Conn., 1979).

Frontiers of the Biological Sciences

Bernard D. Davis

Those of us who were entering biology in the 1930's were very much encouraged by the essays of J. D. Bernal and J. B. S. Haldane, who predicted that the age of biology would soon emerge. Equally confidently, these authors pre-

Having had such a dismal experience with the crystal ball, I shall concentrate in this article not on predicting the future but on characterizing the recent revolution in biology. In its vanguard was a group of bold investigators, in whose

Summary. The history of the molecular revolution in biology is described, emphasizing its dependence on the emergence of bacterial genetics, the fusion of genetics and biochemistry, and the development of greatly improved techniques for studying macromolecules. Central concepts have included molecular information transfer, both by nucleic acids and by allosteric proteins; the spontaneous conversion of one-dimensional information into three-dimensional structures; and the extraordinary unity in the molecular mechanisms underlying the rich diversity of biology. The merging of molecular and morphological studies, to yield the very broad field of cell biology, is described more briefly, as are also some present frontiers in several areas of biology that present challenges at other levels of organization.

dicted a similar success of scientific planning in solving the problems of economic and political organization. Little could the students of my generation foresee that biology would mature so rapidly, while the predicted social utopia would become more distant than ever.

The author is Adele Lehman Professor of Bacterial Physiology, Harvard Medical School, Boston, Massachusetts 02115. Please do not request reprints.

hands the formal entities and processes of genetics became concrete substances and reactions. The resulting field, generally called molecular biology, has focused on the detailed analysis of structures and interactions in a range of dimensions that had long been terra incognita, lying between the small molecules of the biochemist and the visible structures of the morphologist. These ad-

vances depended heavily on the development of techniques that greatly simplified and refined the analysis of macromolecules, and on the development of an industry to supply the required instruments and materials.

One consequence has been a heightened sense of the unity of biology, in terms of the universal features developed early in evolution and also in terms of the continuity between molecular and morphological structure. An even more profound consequence has been a deep understanding of the mechanisms that give living systems their unique capacity to function and to grow, by accumulating negative entropy at the expense of increased entropy in the surrounding environment.

Although physicists had expected to find novel physical laws underlying this property, only the same old physical forces were encountered—but they were found to be organized in novel ways that yield molecular information storage and a flow of information. Thus DNA stores a program, like a computer, on a one-dimensional tape; and cells use this information to specify the structures of their working machinery. Moreover, this machinery includes regulatory proteins, whose flow of information about the state of the cell and its surroundings directs the flow of material and energy. The foundation of these insights was already implicit in the early term organism, with its emphasis on organization. The idea of biological information also parallels modern advances in electronics and information theory, but the interactions have not been substantial.

With these developments the aim of