

INDETERMINACY IN BRAIN AND BEHAVIOR

Paul W. Glimcher

*Center for Neural Science, New York University, New York, New York 10003;
email: glimcher@cns.nyu.edu*

Key Words randomness, decision-making, choice, game theory

■ **Abstract** The central goal of modern science that evolved during the Enlightenment was the empirical reduction of uncertainty by experimental inquiry. Although there have been challenges to this view in the physical sciences, where profoundly indeterminate events have been identified at the quantum level, the presumption that physical phenomena are fundamentally determinate seems to have defined modern behavioral science. Programs like those of the classical behaviorists, for example, were explicitly anchored to a fully deterministic worldview, and this anchoring clearly influenced the experiments that those scientists chose to perform. Recent advances in the psychological, social, and neural sciences, however, have caused a number of scholars to begin to question the assumption that all of behavior can be regarded as fundamentally deterministic in character. Although it is not yet clear whether the generative mechanisms for human and animal behavior will require a philosophically indeterminate approach, it is clear that behavioral scientists of all kinds are beginning to engage the issues of indeterminacy that plagued physics at the beginning of the twentieth century.

CONTENTS

INTRODUCTION	25
Determinism and the Philosophy of Science	26
Do Indeterminacies in the Physical World Matter for Behavioral Scientists?	27
THE RISING TIDE OF APPARENT INDETERMINACY	29
Indeterminacy in the Social Sciences	29
Empirical Measurements of Behavioral Indeterminacy	35
Reducing Uncertainty: Looking for Determinacy with Neurophysiology	40
Indeterminacy at the Cellular and Subcellular Levels	46
THE CHALLENGE OF INDETERMINACY FOR BEHAVIORAL SCIENCE	50

INTRODUCTION

Our modern view that the central function of scientific inquiry is to reduce uncertainty emerged early in the scientific revolution that constituted the Enlightenment; by the time of Galileo's death (cf. Bacon 1620, Descartes 1637, Galilei 1630,

Kepler 1618–1621) it was clear that improving the accuracy with which one could predict future events as determinate processes would be a central feature of the scientific method at both theoretical and empirical levels in the physical sciences. Over the course of the eighteenth and nineteenth centuries, the early social sciences emulated this trend, seeking to develop causal relationships in a testable and determinate fashion (cf. Keynes 1936, Smith 1776). By the twentieth century, the notion that scientific inquiry would reduce animal behavior to deterministic certainty had become a mainstream notion in psychological circles as well. Nowhere is this clearer than in the work of Pavlov. As Pavlov put it in *Conditioned Reflexes*:

The physiologist must thus take his own path, where a trail has already been blazed for him. Three hundred years ago Descartes evolved the idea of the reflex. Starting from the assumption that animals behaved simply as machines, he regarded every activity of the organism as a *necessary* reaction to some external stimulus. . . . A bold attempt to apply the idea of the reflex to the activities of the [cerebral] hemispheres was made by the Russian physiologist I.M. Sechenov, on the basis of the knowledge available in his day of the physiology of the central nervous system. In a pamphlet entitled “Reflexes of the Brain,” published in Russia in 1863, he attempted to represent the activities of the cerebral hemispheres as reflex—that is to say, as *determined*. (Pavlov 1927)

In the period that followed, Skinner and his students (cf. Skinner 1938) strengthened this notion, and psychologists as a group largely embraced the idea that a complete psychological theory would be a determinate one. By studying the causal relationships between environment, organism, and response, these scientists began the process of developing a predictive and testable theory of psychology. The twentieth century witnessed a similar trend in the effective application of the scientific method toward understanding the biological sources of behavior, and as a result, saw the development of a powerful deterministic program for understanding biological systems. Charles Sherrington (1906), for example, applied this programmatic approach to the physiological study of reflexes with great effect.

Determinism and the Philosophy of Science

In philosophical circles, the central role of a determinate worldview in the classical scientific method also became a formalized principle in this period. In the early part of the twentieth century, the philosopher Karl Popper (1934) explicitly defined the goal of modern science as the falsification of extant theories through experimental inquiry. For Popper, theories could never be proven in practice, only subjected to the test of falsification.

If experimental evidence falsifies a theory then it can be discarded; if experimental evidence corroborates a theory then it can be tentatively retained. What is critical about this logic is what it implies about indeterminacy. Consider the theoretical claim that if I flip a certain coin there is a 50% chance it will land

heads-up. As Popper points out, this is not only an unverifiable theoretical claim, but also an untestable one; my assertion predicts all possible empirical outcomes and is thus unfalsifiable. Even more importantly, my theoretical claim predicts as an experimental result all possible finite series of coin flips that could ever be observed. If the coin is equally likely to land heads-up and tails-up, then any specific series of heads and tails is equally likely, whether that be six heads in a row or six flips that alternate between heads and tails. No formal prediction of any particular outcome is ever possible and for this reason Popper argued that probabilistic claims about indeterminate processes were irremediably problematic. Indeed, in his early writings Popper even used this to argue that the notion of a fundamentally indeterminate universe is at base a nonscientific proposition.

In the 1920s and 1930s, however, the emerging discipline of quantum physics raised an important challenge to this notion that had evolved during the Enlightenment and motivated much of Popper's work. Based initially on the work of Heisenberg and his colleagues (Heisenberg 1930, 1952), strong evidence arose suggesting that several phenomena that occur at the atomic and subatomic scales are, in fact, fundamentally indeterminate and thus could be described only probabilistically. This was a critical challenge to the existing philosophy of science as expressed by Popper because that philosophy argued that a theory of physics built upon probability theory was unfalsifiable, perhaps even unscientific. Nevertheless, the empirical evidence gathered during that period seemed to indicate unambiguously that at a small enough scale of analysis, events occur that are fundamentally indeterminate. This indicated that the philosophy of science, rather than the reality of our physical universe, might have to change.

Do Indeterminacies in the Physical World Matter for Behavioral Scientists?

What, if any, are the implications of these issues for the study of behavior? Even if there are fundamental indeterminacies in the physical world, should this matter to behavioral scientists? Many scholars believe that the quantum physicist Edwin Schrodinger provided an answer to that question in his book, *What Is Life* (1944), in which he argued that for any organism to survive it must operate, in principle, in a fully determinate environment. Indeterminacy, he believed, would be lethal to living systems. Schrodinger's own work (cf. 1951) had demonstrated that at the atomic and subatomic scales, matter can be described only in probabilistic terms, but it had also shown that large aggregates of these elementary particles behaved in an effectively determinate manner. His argument was that living cells were large enough objects that they would never interact with single atomic or subatomic particles, but only with these larger determinate aggregates. In essence, he argued that cells were large enough that quantum fluctuations in the properties of individual atoms would have no effect on them. Indeed, he went on to argue that biological cells are the size that they are specifically because quantum indeterminacy prevents them from surviving if they become any smaller. Biologists,

psychologists, and social scientists, he assured us, need not be concerned with fundamental indeterminacy in the universe:

If it were not so, if we were organisms so sensitive that a single atom, or even a few atoms, could make a perceptible impression on our senses—Heavens, what would life be like! To stress one point: an organism of that kind would most certainly not be capable of developing the kind of orderly thought which, after passing through a long sequence of earlier stages, ultimately results in forming, among many other ideas, the idea of an atom. (Schrodinger 1951)

Recently, however, evidence has begun to arise in the social, psychological, and neurobiological domains that suggests that, at larger scales of analysis than the one Schrodinger examined in *What Is Life*, living systems exhibit behavior that is apparently indeterminate (cf. Hastie 2001, Schall 2004, Shafir & LeBoeuf 2002, Staddon & Cerutti 2003). At the largest scale of analysis, social scientists working in areas such as the theory of games have begun to argue that for behavior to be efficient under some circumstances, it must be irreducibly uncertain from the point of view of other organisms and therefore must be studied with the tools of probability theory. In principle, this raises critical problems for game theory. For all of the reasons Popper identified, when game theory makes probabilistic predictions it does so in a manner that is nonfalsifiable. Of course, if Schrodinger was correct, the apparent indeterminacy of game theory presents only a temporary impediment to scientific inquiry. A reductionist approach to human behavior during strategic games would ultimately reveal the mechanisms that give rise to this apparent indeterminacy and thus should ultimately yield a falsifiable determinate theory of human behavior. Although contemporary game theory thus faces indeterminacy, empirical science can hope to resolve this apparent indeterminacy by reduction. Interestingly, however, psychologists working at a lower level of reduction than social scientists have also begun to find evidence of apparent indeterminacy in the systems they study (cf. Staddon & Cerruti 2003). Recently, psychologists have begun the analysis of apparently stochastic patterns of individual responses and have been able to demonstrate classes of individual behavior that appear to be as fully random as can be measured. Indeterminacy, in the hands of these psychologists, seems to be an apparent feature of the behavior of single humans and animals. At a yet deeper level of reduction, neurobiologists have also begun to gather evidence for the existence of apparently indeterminate processes within the architecture of the mammalian brain (cf. Schall 2004). The patterns of action potentials generated by individual neurons, for example, appear highly stochastic for reasons that are not yet well understood.

Growing evidence that apparently indeterminate processes operate at social, psychological, and even neurobiological levels are bringing behavioral scientists face-to-face with the same philosophical and scientific issues faced by Popper, Heisenberg, Schrodinger, and others in the last century. Can such theories be scientific, or is calling a neural signal or a behavior a random process only an excuse for ignorance? It may be that behavioral scientists will choose to assert as an

axiom that all of the physical phenomena we study are fundamentally determinate in order to avoid these issues, but on the other hand such an assertion may force us to neglect a growing body of compelling evidence.

THE RISING TIDE OF APPARENT INDETERMINACY

Indeterminacy in the Social Sciences

Like scholars in the physical sciences, social scientists in the eighteenth and nineteenth centuries strongly emphasized a determinate scientific approach in their study of human behavior. The classic economic theory of that period, for example, rested on the foundation of a theory of determinate utility developed by Blaise Pascal (1670, Arnauld & Nicole 1662) and Daniel Bernoulli (1738). This utility theory argued that humans act predictably to maximize benefits and to minimize costs, and that the costs and benefits of any action can be reliably computed. Pascal had developed this basic logic in the seventeenth century, arguing that the “expected value” of an action was equal to the product of any possible gain or loss incurred by that action and the likelihood of the gain or loss. Bernoulli had extended this notion with the observation that humans appear at an empirical level to be more averse to risk than Pascal’s formulation predicts. Bernoulli’s conclusion was that humans made decisions based on the product of a subjective estimate of cost or benefit and the likelihood of that gain or loss, rather than based on an objective measure of gains or losses. Because of the precise form of his hypothesis, Bernoulli was able to show that this notion could successfully account for the empirically observed aversion of humans to risk. Thus, the critical idea that utility theory advanced was that one could compute the relative desirabilities of all possible actions to a chooser and, except in the presumably rare case where two actions have identical subjective desirabilities, one could then predict the actions of a chooser with determinate precision. Building on this foundation, Adam Smith (1776) argued that all economic actors could be seen as effectively trading off costs and benefits to maximize gain in a complex marketplace. The prices of goods, for example, were presumed to be set by the determinate interactions of supply and demand curves that represented the aggregate subjective desirabilities and costs of goods to consumers and producers. It was thus a central thesis of eighteenth- and nineteenth-century economic theory that the rational process by which desirability was assessed could be accurately modeled and that these models made deterministic predictions about human behavior.

Importantly, the incorporation of likelihoods into expected utility theory allowed the approach to make determinate predictions even when the environment in which human decision makers operated was unpredictable. Choosers were assumed to consider risk when they determined the desirability of an action, and the theory explicitly and convincingly predicted that no feature of this environmental uncertainty would be presumed to propagate into the behavior of the choosers. The only time that utility theory predicted indeterminacy in behavior was when

two or more mutually exclusive actions had precisely equal subjective desirabilities, and the importance of that particular situation seemed limited to the classical economists.

In the first half of the twentieth century, the theory of games developed by John VonNeumann, Oskar Morgenstern, and John Nash directly challenged this determinate approach (Nash 1950a,b, 1951; VonNeumann & Morgenstern 1944). Game theory represented a deviation from the classical tradition specifically because it proposed that when a rational chooser faces an intelligent and self-interested opponent, rather than a passive economic environment, situations could easily arise in which the subjective desirabilities of two or more actions are driven toward precise equality. The theory went on to make surprising and fundamentally indeterminate predictions about how rational humans would behave under many conditions that could be well described by game theory.

To understand this theoretical insight, consider two opponents repeatedly playing the childhood game of rock-paper-scissors in which the loser pays the winner \$2 on a round won by playing paper and \$1 on a round won by playing scissors or rock. If the behavior of one's opponent is unpredictable, any response can win, in principle. Paper will beat a play of rock for \$2, scissors will beat a play of paper for \$1, and rock will beat a play of scissors, again for \$1. Classical utility theory assumes that humans choose between actions by multiplicatively combining the subjective value and likelihood of each outcome and then selecting the action with the outcome that yields the highest expected utility. Assuming naively that one's opponent is equally likely to play rock, or paper, or scissors, the greater value of winning with paper should lead all players to select paper deterministically on each round. What VonNeumann recognized was that this assumption about the behavior of one's opponent simply could not be correct. A competitor who simply selected scissors could reliably defeat any player who actually behaved in accord with this strategy.

Game theory, as developed by VonNeumann & Morgenstern (1944), addresses this limitation of classical utility theory by making the assumption that both players are aware that they face an intelligent opponent who can anticipate their actions and that both players will shape their behavior accordingly to minimize losses and maximize gains. To accomplish this, players must take into account the potential payoffs associated with each choice, as specified by classical utility theory, but they must also consider how the actions of their opponent will influence those payoffs. Consider again the situation in rock-paper-scissors. Winning with paper yields twice as much money as winning with rock or scissors, but deterministically playing paper leads to certain defeat. What VonNeumann & Morgenstern showed was that under these conditions we can predict that a rational player will titrate risk against gain and play paper two-thirds of the time, scissors one-sixth of the time, and rock one-sixth of the time. Critically, however, he must avoid making his two-thirds, one-sixth, one-sixth selections in a determinate fashion; for example, in a repeated version of the game by playing paper, then scissors, then paper, then rock, then paper, and then paper. Were his opponent to divine the determinate nature

of such a strategy (through observation, for example), then winning would again become trivial for that opponent. He would only have to play scissors, then rock, then scissors, then paper, then scissors, and then scissors to assure a consistent win. The only way to avoid this trap is for a player to incorporate apparent indeterminacy directly into his behavior. He must in essence flip a weighted coin on each round to select between rock and paper and scissors. VonNeumann & Morgenstern were well aware of the implications of this observation. It suggested that under some conditions the study of economic choice would have to become a probabilistic process. As they put it:

Consider now a participant in a social exchange economy. His problem has, of course, many elements in common with a maximum problem. [A problem in which a single economic actor seeks to maximize his gain by classically deterministic processes.] But it also contains some, very essential, elements of an entirely different nature. He too tries to obtain an optimum result. But in order to achieve this, he must enter into relations of exchange with others. If two or more persons exchange goods with each other, then the results for each one will depend in general not merely upon his own actions but on those of the others as well. Thus each participant attempts to maximize a function (his above-mentioned "result") of which he does not control all of the variables. This is certainly no maximization problem, but a peculiar and disconcerting mixture of several conflicting maximum problems. Every participant is guided by another principle and neither determines all of the variables which affect his interest.

This kind of problem is nowhere dealt with in classical mathematics. . . . We hope that the reader will be convinced by the above that they face here and now a really conceptual—and not merely technical—difficulty. And it is this problem which the theory of "games of strategy" is mainly devised to meet. (VonNeumann & Morgenstern 1944)

VonNeumann & Morgenstern's critical insight was that under conditions of this type choosers might not be able to identify a single course of action that is deterministically optimal. Instead, they may be forced to select a course of action in as random a fashion as possible. It is this strategy of random selection, known now as a mixed strategy, that distinguishes VonNeumann & Morgenstern's approach from more classical deterministic approaches to the study of behavior. In sum, VonNeumann & Morgenstern argued that human behavior, under some conditions, must appear indeterminate in order to be efficient. They made this point elegantly when they described, in game theoretic form, a conflict between Sherlock Holmes and his archenemy, Professor Moriarity:

Sherlock Holmes desires to proceed from London to Dover and hence to the continent in order to escape from Professor Moriarity who pursues him. Having boarded the train he observes, as the train pulls out, the appearance of Professor Moriarity on the platform. Sherlock Holmes takes it for granted—and

in this he is assumed to be fully justified—that his adversary, who has seen him, might secure a special train and overtake him. Sherlock Holmes is faced with the alternative of going to Dover or of leaving the train at Canterbury, the only intermediate station. His adversary—whose intelligence is assumed to be fully adequate to visualize these possibilities—has the same choice. Both opponents must choose the place of their detrainment in ignorance of the other's corresponding decision. If, as a result of these measures, they should find themselves, *in fine*, on the same platform, Sherlock Holmes may with certainty expect to be killed by Moriarity. If Holmes reaches Dover unharmed he can make good his escape.

What are the good strategies, particularly for Sherlock Holmes? [Set the value] to Professor Moriarity [of] catching Sherlock Holmes [at], say 100. [Alternatively, consider what happens if] Sherlock Holmes successfully escaped to Dover, while Moriarity stopped at Canterbury. This is Moriarity's defeat as far as the present action is concerned, and should be described by a big negative value of the matrix element [for Moriarity]—in the order of magnitude but smaller than the positive value mentioned above—say, -50 . [Finally, consider what happens if] Sherlock Holmes escapes Moriarity at the intermediate station, but fails to reach the Continent. This is best viewed as a tie, and assigned the matrix element 0.

[From a mathematical analysis of these values [one can conclude that] the good strategies (e for Moriarity, n for Sherlock Holmes) [are]:

$$e = \{3/5, 2/5\}, n = \{2/5, 3/5\}$$

Thus Moriarity should go to Dover with a probability of 60% while Sherlock Holmes should stop at the intermediate station with a probability of 60%, the remaining 40% being left in each case for the other alternative.¹ (VonNeumann & Morgenstern 1944, pp. 177–178)

Of course, this theoretical formulation raises critical questions about the scientific nature of game theory. If game theory predicts that Holmes will get off the train at Canterbury with a 60% probability, any action Holmes takes is compatible with the theory. VonNeumann & Morgenstern recognized this but were adamant that this was still the only rational strategy for Holmes to adopt. Holmes must, they argued, be as indeterminate as possible in selecting a course of action. In essence, he must appear to have flipped a weighted coin (weighted 60% for Canterbury and 40% for Dover) in order to maximize his chance of survival. This was true, VonNeumann & Morgenstern suggested, irrespective of whether the theory of games preserved Popperian falsifiability.

By the early 1950s, John Nash (1950a,b; 1951) had seen an interesting additional level of structure in game theoretic problems that required a mixed strategy,

¹Our result for e, n yields that Sherlock Holmes is as good as 48% dead when his train pulls out from Victoria Station.

or apparently indeterminate, solution. Building on the work of VonNeumann & Morgenstern, he concluded that stable mixed strategies must in principle reflect an equilibrium point at which the subjective desirabilities of the two or more actions being mixed were precisely equivalent. He argued that it was only this equivalence that could produce the indeterminate behavior that VonNeumann & Morgenstern had predicted. Consider again the situation when a single player must, on repeated rounds, select rock or paper or scissors. If any one of these is truly preferable as a choice, then we can assume that the chooser will always select that option. Mixed strategies should thus emerge, Nash reasoned, only when the two or more actions that are being mixed have identical average desirabilities. Working from VonNeumann & Morgenstern's insights, Nash argued that these two equivalent desirabilities emerge when the competitive interactions of the two players drive them toward an equilibrium at which the two or more actions being mixed are of equal desirability. What Nash argued was that mixed-strategy equilibria emerge from dynamic interactions between the players, which yield equal average desirabilities, and thus totally indeterminate patterns of behavioral choice. Classical utility theory had presumed that situations in which two or more actions have precisely equal subjective desirabilities would be encountered only rarely. Nash's insight was that not only are they encountered, but the dynamic interactions that occur during strategic games actively create these situations of equal subjective desirability.

From the point of view of indeterminacy, the critical insight that VonNeumann, Morgenstern, and Nash offered was that indeterminacy is a requisite feature of efficient behavior in a competitive world. That insight means either that humans and animals appear indeterminate to each other under some conditions or they behave inefficiently.

In 1982, the evolutionary biologist John Maynard Smith also engaged indeterminacy during strategic games, but from an evolutionary perspective. He argued that any species involved in an internal competition for resources could be described using game theory and that this mathematical formalism predicted that organisms capable of producing apparently indeterminate behavior would be favored by natural selection.

Imagine, Maynard Smith proposed, a species of animals in which individuals compete for access to territories that increase the number of young an individual can produce. Individuals without territories produce a small number of young, while individuals with territories produce a large number of young. Obviously, under these conditions it is in the interest of individuals to obtain territories. Now consider a situation in which there are more individuals than territories. In this hypothetical species, territories change hands when an animal without a territory "displays" to an animal with a territory, essentially threatening that individual for control of the territory. In the hawk-dove game, as this competition has come to be known, after such a display each animal must make a decision: whether to escalate the conflict (to fight for the territory) or whether to retreat (give up the territory without a fight). If one of the animals elects to escalate, behaving as a hawk, and

TABLE 1 Payoffs for challenger in the hawk-dove game

	Challenger chooses hawk	Challenger chooses dove
Defender chooses hawk	50% chance of gaining territory 50% chance of injury	Nothing gained
Defender chooses dove	Value of territory gained	50% chance of gaining territory

one decides to retreat, behaving as a dove, then the hawk takes the territory. If both animals elect to be doves, then one of them at random takes the territory. Finally, if both animals elect to be hawks, then they fight, one sustains injuries that reduce the number of young that individual can produce, and the other gains the territory. Table 1 illustrates this simple game as a two-by-two matrix that specifies the costs and benefits of all possible actions to each player.

What Maynard Smith realized at a mechanistic level was that each of these values could be expressed in terms of evolutionary fitness, the gain in reproductive success, that an individual achieves with each outcome. Gaining a territory confers an increase in fitness, whereas sustaining an injury confers a decrease. Thus, if the value of a territory is high and the magnitude of injury in a hawk versus hawk fight is low, then animals that behave as hawks are more fit than those that behave as doves. Under these conditions, Maynard Smith reasoned, the population will evolve toward a single pure strategy equilibrium: All animals in the population will always be hawks. Similarly, if the value of a territory is low and the magnitude of injury is high, then all animals that behave as doves will produce more offspring, be more fit, than animals that act as hawks. Under these conditions, the population should converge on a pure strategy equilibrium of dove. If, however, the value of a territory is high and the cost of an injury also is relatively high, then an interesting thing happens. The only reproductively stable strategy for the animal and its offspring is to behave sometimes as a hawk and sometimes as a dove. To be more specific, a single dominant and unbeatable strategy emerges in a population playing the hawk-dove game. The probability that on any given encounter an individual will choose to behave as a hawk must be equal to the value of a territory divided by the magnitude of the injury sustained in a hawk versus hawk conflict. Critically, on each encounter individuals have to behave in an unpredictable fashion, never allowing their opponent to know whether they will be a hawk or dove². But across many such encounters the only stable and unbeatable solution for the population is for the probability of being a hawk to be equal to the value of a territory divided by the cost of injury.

This theoretical analysis suggests that, at least from the point of view of other individuals in this same species, evolution would drive behavior toward

²Maynard Smith showed mathematically that a population of unpredictable individuals would dominate a population in which separate individuals were committed at birth to playing hawk or dove. For details of that proof, see (Smith 1982).

unpredictability. As in the game theoretic work of VonNeumann, Morgenstern, and Nash, the ability to generate apparently unpredictable behavior seems advantageous. One interesting feature of Maynard Smith's argument, however, is the mechanism by which this uncertain behavior would be presumed to arise. We have strong reasons to believe that completely novel behaviors arise, at least in part, from genetic mutations. Random changes occur in the genomes of these animals and then selection operates to preserve useful variations in behavioral traits. Atomic-level fluctuations in DNA molecules, induced by quantum-level forces like cosmic radiation, produce unpredictable changes in the genetic make-up of a species. These random changes then influence behavior. We have every reason to believe that the mechanism by which apparently indeterminate behaviors would arise would itself be truly indeterminate.

Game theory, whether directed toward the actions of an individual or the evolution of a species, predicts that under some conditions behavior must appear indeterminate in order for it to be efficient. What implications, if any, does this have for the determinate scientific method? Does this mean that social scientists have to abandon Popperian falsifiability? Probably not, for at least two important reasons. First, the theoretical observation that behavior *should* appear indeterminate does not mean that behavior *does* appear indeterminate. Physical constraints may make it impossible, or unlikely, for mutations to generate behavior that even appears indeterminate. If this is the case, then at an empirical level we may simply find that apparently indeterminate behavior does not occur. Second, even if apparently indeterminate behavior were to be observed, this would not require that the physical generative process for the behavior itself be indeterminate. The psychological system responsible for decision making during strategic games might operate on totally deterministic grounds. Like a modern digital computer, it may simply generate an appearance of indeterminacy sufficient to defy prediction by the opponent. In summary, there may be reasons why fundamental indeterminacies like those that arise at the quantum level cannot influence the systems that generate behavior. Either of these two observations would rescue Popperian falsifiability in its strongest form.

Empirical Measurements of Behavioral Indeterminacy

The theory of games makes it clear that an organism with the ability to produce apparently indeterminate patterns of behavior would have a selective advantage over an animal that lacked this ability. Were apparently indeterminate behavior to have arisen in the evolutionary history of vertebrates, there seems every reason to believe that this behavioral phenotype would be preserved. Do humans have this ability? A common answer to this question, based on studies of humans, is no.

Over the course of the last 40 years, a number of psychological studies have suggested that, perhaps because of some fundamental constraint in the human nervous system, humans cannot generate behavior that appears indeterminate (for a review of this literature see Wagenaar 1972). For example, in one of the first of

these studies, Bakan (1960) asked humans to simulate the action of a random coin flip: subjects were asked to make up a sequence of heads and tails that was fully random in order. When Bakan analyzed the sequences generated by these subjects they were found to be highly nonrandom despite the instructions that the subjects received. Bakan found that the subjects tended to overproduce alternations between heads and tails and to underproduce the occasional long runs of heads or tails that would be predicted from a truly random process. In sum, the humans behaved in a fairly determinate fashion, despite their instructions to do otherwise. Since 1960, dozens of studies have replicated this basic result. When human subjects receive a verbal instruction to produce a random sequence, they reliably fail. On the basis of this evidence, many psychologists have concluded that humans lack the ability, in principle, to generate patterns of behavior that appear indeterminate.

According to game theory, however, environmental conditions should arise in which apparently indeterminate behavior would be truly beneficial. Organisms in their natural environment would be reinforced for producing apparently indeterminate behavioral patterns under some conditions. Regardless of these human data, then, can nonhuman animals produce apparently indeterminate behaviors if they are reinforced for doing so? Blough (1966) was one of the first to ask this question directly by specifically reinforcing pigeons for producing behavior that approximated a random process. In that experiment, pigeons were trained to peck a key in a Skinner box, and the amount of grain that they received after each peck was contingent upon the length of time that had intervened since the last peck. The more closely the set of interresponse intervals produced by the pigeon approximated the output of a random Poisson-like process, the more grain the bird earned. Blough found that under these conditions the birds quickly adopted a response strategy that was virtually indistinguishable from the output of a truly random operator. Figure 1 shows the frequency distribution of interresponse intervals Blough obtained from a single pigeon and, plotted as a solid line, the pattern of intervals that would be expected from a fully random process. While Blough's analysis did not show that the behavior of the pigeons was random by all possible measures, it did demonstrate that when an apparently indeterminate behavior was reinforced, pigeons could produce a behavior of this general type. This study was critical because it provided the first evidence that the ability to produce apparently indeterminate behavior had arisen in the vertebrate line.

Since that original study, a voluminous literature has examined the ability of several species of animals to generate apparently indeterminate behavioral sequences when they are specifically reinforced for doing so, and tasks more closely approximating the conditions described by game theorists have also been examined (see Neuringer 2002 for a review of this literature). Shimp (1967) introduced one paradigm that has been particularly widely studied. In that paradigm, pigeons were trained to choose sequentially between left and right response keys for four responses during each of thousands of trials. The behavior of the pigeons on each trial thus produced one of 16 possible patterns, for example, left-right-left-left. The animals then were reinforced for producing the 16 possible patterns with an




Figure 1 Mean interresponse times (IRTs) from two replications of three experimental conditions for a single pigeon (Blough 1966). The graph plots the frequency of each IRT in half-second bins. A fully indeterminate process would produce points that fall along straight lines in this space. (Reproduced with permission from *Journal of the Experimental Analysis of Behavior*, copyright 1966, Society for the Experimental Analysis of Behavior.)

apparently random frequency distribution. In one important and well-controlled study, Page & Neuringer (1985) employed a strategy of this type to examine indeterminacy in behavior and to see whether the ability of pigeons to produce random sequences depended specifically upon whether or not they were reinforced for apparent randomness. In that experiment, pigeons produced long sets of left and right responses, but under two reinforcement contingencies. Under the first contingency, the animals were specifically reinforced for producing patterns of left and right responses that had a random-like frequency distribution. Under the second contingency, the randomness of the emitted frequency distribution was irrelevant to the reward received. Page & Neuringer found that when reinforcement was contingent on variability, the variability of the pigeons' responses increased, but when the level of variability was not reinforced directly, the pigeons adopted much more determinate response patterns. More specifically, they found that an information theoretic analysis of the response patterns showed nearly perfect indeterminacy when, but only when, indeterminacy was reinforced. These results suggest two interesting conclusions. First, they suggest that the degree of apparent indeterminacy included in behavior is variable. Animals can be more or less indeterminate based on the requirements of their environment. Second, they suggest that when

indeterminacy is irrelevant, this species of animal prefers to adopt a fairly determinate response strategy.

Machado (1989) employed a similar approach in another important study. In that experiment, pigeons once again emitted four left or right responses in each of thousands of sequential trials, and the variability of the response pattern they produced was assessed statistically to determine the amount of reward that the pigeon would receive. This was accomplished simply by counting the number of trials since that sequence had last been produced and assigning this number, a variability score, to that trial. To determine whether a reinforcement was delivered, the variability scores for the last 20 trials were cumulated and the variability score for the current trial was compared to the variability scores for those last 20 trials. If the percentile rank of the variability score for the current trial exceeded some fixed threshold, for example 50%, then a reinforcement was delivered. If pigeons responded truly randomly, then the probability of emitting all 16 possible sequences within 25 trials would be less than 1%. Accordingly, Machado adjusted the threshold requirement so that it never reinforced patterns of sequences shorter than 25 trials in length. Machado found that under these conditions the frequency distribution of variability scores actually produced by the birds was nearly identical to the frequency distribution that would be expected from a truly random process. Figure 2 plots this relationship for one of Machado's animals. In this figure, the points plot the frequency with which each possible variability score was observed




Figure 2 Dots are used to plot the distribution of variability scores obtained from a single animal in the Machado (1989) paper. The solid line indicates the frequency distribution that would be expected from a perfectly indeterminate process. (Reproduced with permission from *Journal of the Experimental Analysis of Behavior*, copyright 1989, Society for the Experimental Analysis of Behavior.)

during a reinforcement contingency that maximized indeterminate behavior. Low variabilities were observed often and high variability scores occurred more rarely. The solid line plots, for comparison, the pattern of variability scores that would be expected from a fully random process measured in this way. Machado's critical observation is that this particular reinforcement protocol yields a fourth order pattern of responses indistinguishable from the pattern that would be expected from a fully indeterminate process. In sum, many studies (Neuringer 2002) that have yielded data similar to Blough's, Page & Neuringer's, and Machado's suggest that when nonhuman animals are reinforced for producing apparently indeterminate patterns of behavior, they can produce behavior of this type.

This set of observations thus led Neuringer (1986) to test the hypothesis that previous studies with human subjects had failed to yield apparently indeterminate behavior because human subjects had not been specifically reinforced for producing apparently indeterminate behaviors. In this study, Neuringer instructed human subjects to produce a random sequence of ones and twos on a computer keyboard. He then analyzed the resulting sequences for nonuniform distributions of ones and twos, and first and second order patterns in the ones and twos. He also analyzed the sequences with a set of statistics related to autocorrelation functions. What he found under these conditions was that the human subjects produced highly non-random sequences exactly as had been observed in previous studies. Neuringer provided feedback to these subjects by showing them, after a run of 100 trials, how the distribution they had produced deviated from the distribution that would be expected from a random sequence according to one of the statistical measures of randomness that he employed. In sequence, Neuringer then presented the subjects with each of the additional statistical metrics until they were receiving feedback according to all five metrics at the end of each 100-trial run. Finally, the feedback was terminated and the subjects were told that if they could produce a sequence that could not be discriminated from the product of a computer pseudorandom number generator, they would receive a cash bonus. Under these conditions, Neuringer found that the human subjects essentially all produced sequences that could not be discriminated from random sequences by any of the metrics he employed. From these data, he concluded that like pigeons and other vertebrates, human subjects could produce apparently indeterminate sequences under some conditions.

In a similar experiment, Rapoport & Budescu (1992) examined the behavior of humans playing two-person games of the type VonNeumann, Morgenstern, and Nash studied. In Rapoport & Budescu's experiments, random-like behaviors were reinforced monetarily, and they found that humans could produce behavioral sequences that appeared indeterminate. Under conditions in which game theory predicts that indeterminacy will be reinforced, apparently indeterminate behavior can be produced.

Of course, these data do not demonstrate that humans can produce fully indeterminate behavior. All of these data suggest that humans and animals can produce behavior that appears indeterminate, but it seems probable that, like a random-number generator in a computer, the generative process for this behavior is likely

determinate at a lower level of reduction. To test that hypothesis, however, one would have to turn to a neurophysiological level of analysis.

Reducing Uncertainty: Looking for Determinacy with Neurophysiology

Perhaps the most influential study of choice behavior at the level of interacting neurons has been the work of William Newsome and his colleagues at Stanford University (for a review of Newsome's work, see Batista & Newsome 2000; for a review of neurobiological choice literature, see Glimcher 2003). Newsome and his colleagues trained rhesus monkeys to monitor a visual display that presents a circular patch of chaotically moving spots of light (Figure 3). Upon viewing a display of this type, human observers report a chaotic blizzard of randomly moving white spots. However, when 15% of the spots move coherently in a single direction, humans subjects report a strong sense that the spots are moving, overall, in that direction. If the fraction of spots moving coherently is reduced, the strength of this perceived motion is reduced. By systematically varying the fraction of spots moving coherently, Newsome and his colleagues could therefore systematically manipulate how difficult it was for observers to determine the average direction in which the spots were moving.

In their original experiment (Newsome et al. 1989), monkeys were presented with a display of this type for 2 seconds, after which they had to decide in which of two possible directions the spots, on average, were moving. The animals indicated their decision with an eye movement that shifted the animal's point-of-gaze in the direction of perceived average motion. If the animals had judged the direction of spot motion correctly, they received a fruit juice reward.

While animals made these decisions, the activity of single motion-sensitive neurons in the middle temporal visual cortex (area MT) was monitored. Under these conditions, Newsome and his colleagues found that if 15% of the dots in the display moved to the right, the monkeys always reported that they saw rightward motion, and cells in area MT activated by rightward motion rapidly generated action potentials. As the percentage of rightward dots was systematically decreased, both the probability that the monkey would report that he had seen rightward motion and the probability that the neurons would show an increase in firing rate decreased at roughly the same rate.

What can we learn from data of this type about the mechanisms that give rise to apparently indeterminate behaviors? In one interesting study, Britten and colleagues (1996) examined the activity of MT neurons while monkeys viewed a display in which either all of the spots moved in random directions (there was no coherent direction of spot motion) or only a small fraction of the dots moved in a coherent direction. These conditions were selected to examine the relationship between neural activity in area MT and the decisions that an animal made when the visual stimulus was ambiguous. Interestingly, on a subset of the trials in which there was no coherent motion of the spots, the exact same pattern of randomly

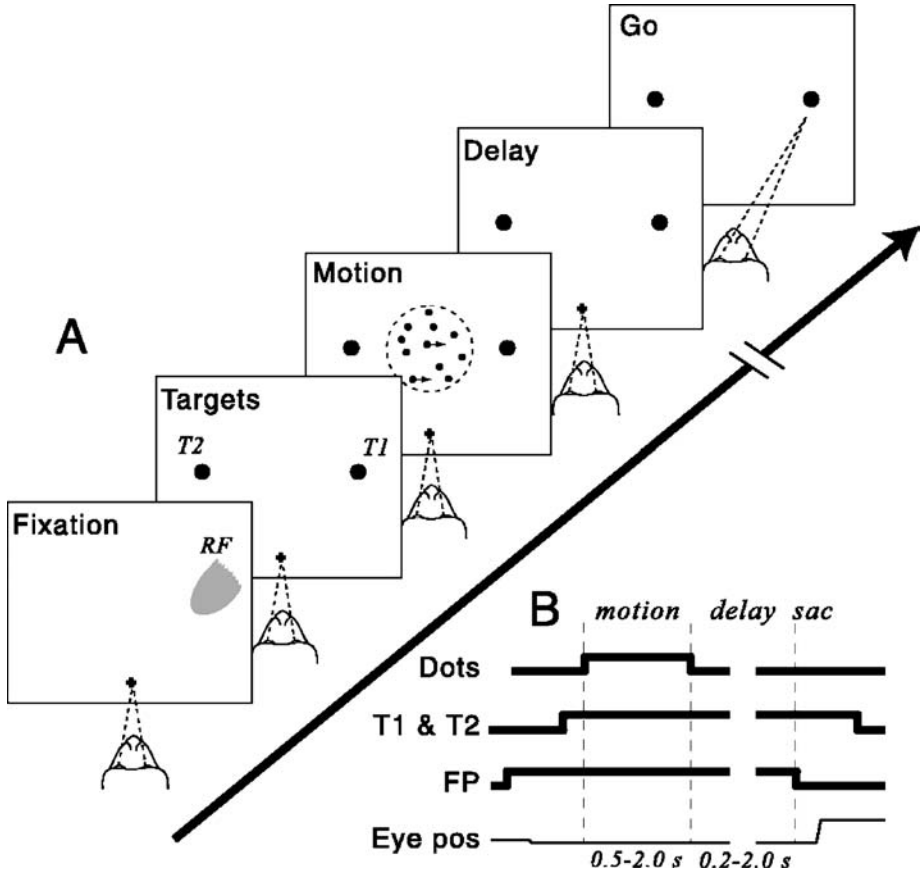


Figure 3 The Moving Spot Task. Monkeys fixate a central point while chaotically moving spots of light are presented within a circular aperture. On any given trial, a small fraction of the spots moves in a coherent manner in one of two possible directions. Across trials, the fraction of dots moving in this coherent fashion can be varied systematically to increase or decrease the strength of the perceived motion signal in the two possible directions. After viewing the display for 2 seconds, monkeys indicate the direction of perceived motion with a saccadic eye movement. Correct responses are reinforced with water or fruit juice. (From Shadlen & Newsome 2001. Reproduced with permission from *Journal of Neuroscience*.)

moving spots was presented. Under these conditions, the animals viewed the exact same stimulus while the activity of MT neurons was monitored. Perhaps surprisingly, even under these conditions the activity of the neurons varied from trial to trial. The precise number of action potentials generated and the precise pattern of action potential generation differed in an apparently random manner from trial to trial, even when the visual stimulus that the animal was evaluating was identical.

Britten and colleagues also found that the perceptual judgments of the animals were unpredictable on these trials. Like the neurons, the behavior of the animals was variable. Finally, these authors found that the judgments of the animals were always correlated, although only weakly, with the activity of the neurons. In sum, the neurons appeared to be indeterminate with regard to the stimulus, and the decisions that the animals made were correlated with these apparently indeterminate neural events.

These data led to the generation of a model (Figure 4) designed to simulate the brain circuits for making this perceptual decision about the direction of spot motion (Shadlen et al. 1996). The model proposed that a group of, for example, rightward motion-sensitive neurons in area MT pooled data according to a fully defined algorithm to yield an instantaneous estimate of the current strength of rightward motion in the moving spot display. In a similar way, a group of leftward motion-sensitive neurons was hypothesized to extract an estimate of the

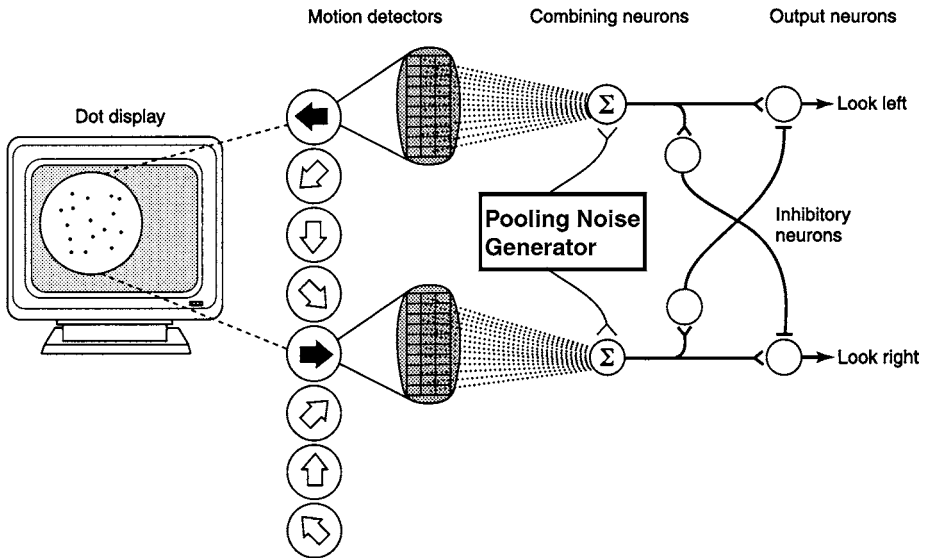


Figure 4 Shadlen et al.'s (1996) model of a perceptual decision circuit. Pools of neurons in area MT extract the instantaneous strength of visual motion occurring in the display, for motion in all possible directions. The instantaneous pooled estimates of motion strength in each of the two possible directions are passed to elements that compute the time integral of that signal to derive an estimate of the average motion signal over a 2-second display interval. The process of pooling is presumed to involve the addition of a fundamental indeterminacy called the “pooling noise” in the original model and labeled here as the pooling noise generator. These integrative elements project, in turn, to eye movement-producing neurons. The integrative elements are postulated to be mutually inhibitory, assuring that only one eye movement is triggered at a time.

instantaneous strength of leftward motion. Because the monkeys were allowed to view the motion stimulus for 2 seconds, this had allowed them, in principle, to sum 2 seconds of instantaneous motion information before making a choice. Accordingly, Shadlen and his colleagues proposed that the output of each of the neuronal pools of direction-sensitive neurons was summed, or integrated, over the 2-second period, to develop an estimate of the average direction of motion throughout the stimulus interval. They next proposed that the neurons that integrated rightward activity should be able to directly activate circuits that produced rightward eye movements and that leftward integrating neurons should be able to activate circuits for producing leftward eye movements. To make the system capable of decision making, in the sense of making choices, the model employed two inhibitory linkages that allowed the output of each integrator to inhibit the other integrator's access to the eye movement control circuits.

A quantitative analysis of the behavior of the model, however, revealed a surprising result. The behavior of the actual monkeys appeared much more random than would be predicted from the neurobiologically derived model. Moreover, interestingly, this apparent randomness could be accounted for only by assuming neural circuitry that specifically incorporated a degree of intrinsic randomness that they referred to as a neuronal pooling noise. Shadlen and colleagues were forced to incorporate into the model a fully random element in order to account for their results.

The Shadlen model was intended to link the activity of neurons in area MT to behavior in as determinate a fashion as possible, but Shadlen and his colleagues concluded that this linkage could only be accomplished if it was presumed that the nervous system incorporated an indeterminate element. Of course, there was no specific claim about the mechanistic nature of this apparently indeterminate neural element. The pooling noise generator could be a determinate device that yields an apparently indeterminate signal, but it is interesting that even at this level of analysis an indeterminate process seemed to operate.

Motivated in part by those findings, Dorris & Glimcher (2004) elected to examine the behavior and brains of monkeys employed in a game theoretic conflict that actually required an apparently indeterminate type of behavior, a mixed strategy game of the type VonNeumann and Nash had modeled. In the human version of their inspection game, two opponents face each other, an employer and an employee (Figure 5). In each round of the game, the employee must decide whether to go to work, in which case he earns a fixed wage, or to shirk, in hopes of earning his wage plus a bonus. The goal of the employee is simply to maximize his gains in terms of salary and bonus. The employer, on the other hand, must decide between trusting his employee to arrive for work or spending money to hire an inspector who can check and see whether the employee arrived for work that day. The goal of the employer is to spend as little money as possible on inspections while maximizing the employee's incentive to work.

In this game, both human and monkey contestants played the role of the employee against a standardized, and strategically sophisticated, computer employer. Each round began with the illumination of two lights, one for working and one for




Figure 5 General form of the payoff matrix for the inspection game for both the experimental subjects (employees) and their opponents (employers). The variables in the *bottom left* of each cell determine the employee's payoffs and the variables in the *top right* of each cell determine the employer's payoffs for each combination of player's responses. V = value of hypothetical product to the employer, fixed at 4; W = wage paid by employer to employee, fixed at 2; C = cost of working to employee, fixed at 1; I = cost of inspection to the employer, varied from 0.1 to 0.9 in steps of 0.2. *Middle* and *right* panels show payoff matrices for 70% and 30% employee-shirk rates. The predicted equilibrium strategy for the employer remains constant at a 50% inspect for all blocks of trials. One unit of payoff is 0.25 mL of water for monkey or 5 cents for human.

shirking. At the end of each round, players selected one light and the computer employer simultaneously decided whether to pay for an inspection on that round. These responses were then compared by a computer arbiter that paid off players according to a fixed payoff matrix (paying off in juice for monkeys, real currency for humans, and virtual currency for the computer employer).

Dorris & Glimcher (2004) found that the overall probability that a human playing the inspection game for money would chose to shirk was reasonably well predicted by the Nash equilibrium computations, but more importantly, they found that human subjects behaved almost perfectly randomly from trial to trial. An analysis of the human data revealed when the Nash solution in the game was for a player to shirk 50% of the time, not only did the players shirk about 50% of the time, but they also showed essentially no patterns in their behavior out to a third-order statistical analysis. As in the experiments of Neuringer (1986) and Rapoport & Budescu (1992), subjects appeared capable of producing largely random patterns of behavior when they were reinforced for doing so. When Dorris & Glimcher analyzed the behavior of their monkeys, they found that the behavior of the monkeys was surprisingly similar, even essentially identical, to the behavior of their human employees. Just like humans, the monkeys seemed to precisely track the Nash equilibrium solutions and to produce those average solutions using largely random sequences of working and shirking.

When Dorris & Glimcher (2004) examined the activity of neurons in the posterior parietal cortex while monkeys played the inspection game, they found that the posterior parietal cortex carried a signal essentially identical to one predicted by game theory. The neural activity was correlated with the theoretical quantity economists refer to as expected utility. Importantly, however, this neural encoding of an economic choice variable was not accomplished in a totally deterministic fashion. The cortical neurons responded with an average rate that was correlated with expected utility, but on a moment-by-moment basis, the neurons behaved unpredictably. At a formal level, the neurons behaved roughly like Poisson devices, producing action potentials with random interspike intervals much like the interpeck intervals Blough's (1966) pigeons produced.

So what do we know of the mechanism that generates choice behavior under these conditions? Shadlen's computational model of the choice process (Shadlen et al. 1996) seems to suggest that at the level of the neural computation we still can see evidence of apparent indeterminacy, and other models loosely related to the original Shadlen model seem to make a similar point (Barraclough et al. 2004, Corrado et al. 2003, Glimcher & Dorris 2005, Lau & Glimcher 2003). The absolute variability of primate behavior seems to be adjustable, and neural models of the machinery that generates this behavior at a neuronal level seem to include apparently random elements. What, then, is the implementation of this random element? One hopeful possibility is that the cellular-level mechanisms that implement this randomness may be, in fact, fully determinate processes. It is at least possible that if we better understood the mechanisms by which cells (for example, cortical neurons) generate action potentials, it would still be possible

to reduce this apparent indeterminacy to a determinate process at a subcellular level.

Indeterminacy at the Cellular and Subcellular Levels

Among the first scientists to examine the pattern of cortical neuronal firing rates with regard to indeterminacy were Tolhurst et al. (1981) and Dean (1981), who were extending studies of neuronal variability pioneered by Barlow & Levick (1969; see also Heggelund & Albus 1978). In two landmark papers, Tolhurst et al. (1981) and Dean (1981) examined the firing patterns of neurons in the visual cortices of anesthetized cats viewing visual displays that presented moving bars of light. When a given visual stimulus was presented to the animals, cortical neurons always responded with a fixed average rate of firing. A vertically oriented bar, for example, always produced a fixed rate of average action potential generation. As the bar was rotated toward a horizontal orientation, for example, the cell responded with a different, but also consistent, average rate of action potential production. They found, however, that the exact pattern of firing that gave rise to this average rate seemed to be almost completely unpredictable. Indeed, as the average firing rate increased, the moment-by-moment variability of the spike rate also increased, almost exactly in proportion to this mean rate. Put more formally, Tolhurst et al. and Dean found that the average firing rate was proportional to the square of the variance across a broad range of rates (Figure 6). This was a statistical distribution that would occur if the process of generating an action potential could be described in the following way: Immediately after an action potential is generated there is a 0% chance of generating an action potential for some largely fixed interval. After that interval has elapsed, the probability of generating an action potential in any given instant becomes fixed at a low level until an action potential occurs, after which the probability of action potential generation is again zero and the process repeats. Of course, during the interval when the probability of action potential generation was fixed at this low level, the spike generation could be characterized as, in principle, fully random. The time at which a spike occurred could be described as a fully random process that had all of the hallmarks of a truly stochastic Poisson operator.

What Tolhurst et al. (1981) and Dean (1981) found, therefore, was that at the level of action potential generation, cortical neurons could be described as essentially stochastic. This was a surprising result at the time, and it has been widely confirmed (Rieke et al. 1997, Shadlen & Newsome 1998). What then is the source of this apparent stochasticity, and would a more detailed biophysical analysis of the spike generation mechanism reveal an underlying deterministic process that would yield this apparent indeterminacy?

To examine one possible answer to that question, Mainen & Sejnowski (1995) sought to determine whether the biophysical process that actually generates action potentials in response to changes in membrane voltage was determinate. They performed intracellular manipulations of single cortical neurons in cortical networks by employing a brain slice preparation, inserting a microelectrode inside a single




Figure 6 Tolhurst et al.'s (1981) plot of variance as a function of the mean firing rates for a cat's visual-cortical neurons. The different symbols represent different averages of stimulus conditions; the straight lines plot regressions. The graph indicates that the square of variance and mean rate are related by a relatively fixed constant of proportionality. (Reproduced with permission from *Vision Research*. Copyright Pergamon Press.)

neuron and recording the pattern of membrane voltage produced in the cell by the network in which it was embedded. While membrane voltage was monitored, they also recorded the precise times at which the cell generated action potentials. This allowed them to determine the relationship between membrane voltage and action potential generation under reasonably normal conditions. Next, they disconnected this cell from the rest of the network in which it was embedded and used the microelectrode to reinject exactly the same pattern of membrane voltages that had occurred originally. They found that under these conditions the cell fired action potentials at exactly the same time, with regard to the membrane voltage signal, as it did previously. They found that the spike-generating mechanism was fully

deterministic. A given pattern of membrane voltage gave rise to exactly the same pattern of action potentials no matter how many times it was injected into the cell.

On the one hand, this was a reassuring result. At base, the pattern of action potential generation was found to be governed by a determinate device. However, on the other hand, it was puzzling. Spike rates are not determinate in this sense. Tolhurst et al. and Dean's work indicates that spike rates are distributed in a Poisson-like fashion, and there clearly is nothing about the spike generator within each cell that produces this pattern. The Mainen & Sejnowski (1995) data indicate that the apparent randomness in spike patterns must be a function of apparent randomness in the underlying membrane voltages. What then are the sources of these Poisson-like fluctuations in membrane voltage?

We know that membrane voltages are governed, ultimately, by the pattern of synaptic activations that a cell receives from the neurons that impinge upon it. Each cortical neuron receives about 10,000 synapses from the tissue that surrounds it. The fact that about half of these synapses are excitatory and half are inhibitory is also important. It means that net excitation and inhibition are largely balanced in an active neuron and small shifts in this balance cause the membrane voltage to rise and fall, and thus cause action potentials to be generated. Together, these observations make a clear suggestion. The source of the apparent stochasticity in the membrane voltage either is a determinate pattern of synaptic activations that carefully sculpts the membrane voltage to yield an apparently indeterminate pattern of action potentials for reasons we do not yet understand or the process of synaptic activation is itself apparently indeterminate.

A number of groups have investigated this latter possibility by studying the activity of single synapses (see Auger & Marty 2000, Stevens 2003 for reviews of this literature). The basic approach taken by these groups has been to activate a neuron and then monitor the rate at which individual synaptic vesicles are released into the synaptic cleft. Before these experiments were undertaken one could have speculated that synapses were simple determinate mechanisms: When an action potential invades the presynaptic region, it might be presumed that synaptic vesicles of neurotransmitter were deterministically released into the synaptic cleft. Modern studies of this process seem to contradict this view, however. Current evidence indicates that when an action potential invades the presynaptic terminal, the chance that a single synaptic vesicle will be released can be as low as 20%. Examinations of the precise patterns of vesicular release suggest that the likelihood that a vesicle of neurotransmitter will be released in response to a single action potential can be described as a random Poisson-like process. Vesicular release seems to be an apparently indeterminate process.

Careful study of other elements in the synapse seems to yield a set of similar, and highly stochastic, results. Postsynaptic membranes, for example, seem to possess only a tiny number of neurotransmitter receptors (cf. Takumi et al. 1999), and during synaptic transmission as few as one or two of a given type of receptor molecules may be activated (Nimchinski et al. 2004). Under these conditions, a single open ion channel may allow a countable number of calcium or sodium ions

to enter the neuron, and there is evidence that the actions of a single receptor and the few ions that it channels into the cell may influence the postsynaptic membrane. Together, all of these data suggest that membrane voltage is the product of interactions at the atomic level, many of which are governed by quantum physics and thus are truly indeterminate events. Because of the tiny scale at which these processes operate, interactions between action potentials and transmitter release as well as interactions between transmitter molecules and postsynaptic receptors may be, and indeed seem likely to be, fundamentally indeterminate.

In 1944, Schrodinger argued that the fundamental indeterminacy of the physical universe would have no effect on living systems. He argued that were biological systems to become so small that the actions of single atoms or molecules could influence cells, the resulting organisms would surely perish from the evolutionary landscape. Studies of the mammalian synapse, however, seem to indicate that Schrodinger (1944) was simply wrong in this regard. Single synapses appear to be indeterminate devices; not apparently indeterminate, but fundamentally indeterminate. At base, physical indeterminacy seems to be a fundamental property of the brain. But how sure can we be that this fundamental indeterminacy at the level of the synapse has anything to do with indeterminacy at the level of a single cortical neuron, at the level of a cortical network, at the level of behavior, or at the level of a social theory of behavior?

The evidence that we have today suggests that membrane voltage can be influenced by quantum level events, like the random movement of individual calcium ions. So there is every reason to believe that membrane voltage can be viewed, at least under some circumstances, as a formally indeterminate process of the type that precludes Popperian falsifiability. How does this membrane voltage influence action potential generation? Recall that cells receive a mixture of excitation and inhibition from thousands of synapses and that the ratio of this mixture is variable. Imagine that the correlations between the activity of the individual synapses impinging on a given cell were variable. Under conditions in which the activity of many synapses is correlated and the membrane voltage is driven either way above or way below its threshold for action potential generation, the network of neurons itself would maintain a largely determinate characteristic even though the synapses themselves might appear stochastic. Alternatively, when the synaptic activity is uncorrelated and the forces of excitation and inhibition are balanced, small uncorrelated fluctuations in synaptic probabilities drive cells above or below threshold. Under these conditions, indeterminacy in the synapses propagates to the membrane voltage and thence to the pattern of action potential generation. Indeterminacy in the pattern of action potential generation, although variable, would reflect a fundamental indeterminacy in the nervous system.

At the level of behavior, apparent indeterminacy is reinforced by the environment and has been observed. Animals can produce behavior that appears to scientists to be indeterminate. How does this apparent indeterminacy arise? Given what we know about the behavior of synapses and action potentials, two possibilities present themselves. The fundamental indeterminacy observed at the cellular

level could be prevented from influencing higher-level phenomena in the nervous system, rendering these higher-level phenomena determinate. These determinate processes could then instantiate pseudorandom computations that emulate the underlying cellular indeterminacy and yield apparently indeterminate behavior. Alternatively, we can propose the hypothesis that indeterminacy observed at the cellular level could propagate to behavior under some circumstances, yielding truly indeterminate behavior under some conditions and more determinate behaviors under others.

THE CHALLENGE OF INDETERMINACY FOR BEHAVIORAL SCIENCE

The traditional scientific method, or at least one interpretation of that method, suggests that the goal of an investigator should be to reduce uncertainty. We make predictions from our data about future states of the world, Ψ_i , which have some error, ε . One goal of science is to reduce ε to the smallest possible value and then to use the Ψ_i measured under these conditions to falsify incorrect theories. Formally, situations in which behavior *appears* highly indeterminate are those in which ε is large with regard to Ψ_i . The argument that aspects of the world are, however, *truly* indeterminate necessitates a change in the way measurements of Ψ_i are approached. In a fundamentally indeterminate world, ε would have a fixed minimum value beyond which the reduction of uncertainty would be impossible. If this is the case, and ε does have a fixed minimal value, then two critical problems arise for the scientific method. First, a measurement problem arises. If variability is observed during a scientific measurement, does that represent an error on the part of the scientist or variability in the world? Without a reliable technique for specifying the minimum value of ε under a given set of circumstances, there is no way to know if a measurement is accurate. This promotes anarchy in the method by permitting a confusion between error and observation. Second, a falsification problem arises. The existence of a lower limit on ε precludes hard falsification of the type Popper advocated. If a given set of scientific predictions must be couched in probabilistic terms, then—for all the reasons Popper outlined—rigorous falsification is impossible.

Good examples of these measurement and falsification problems arise in the contemporary debate about what information is carried in the Poisson-like patterns of action potentials produced by cortical neurons. Cortical neurons produce variable patterns of interspike intervals. All efforts to reduce that variability to a determinate pattern have essentially failed. Some scientists conclude from this failure that spike trains are, at root, indeterminate and that the only information carried by these patterns of action potentials is encoded by the mean rate at which they occur (Shadlen & Newsome 1998). Others propose theories that would yield Poisson-like patterns of interspike intervals, but from underlying determinate processes. Is the first of these hypotheses testable, falsifiable, and scientific? The

answer to that question seems far from clear. What is clear is that two kinds of indeterminacy in principle could occur: a fundamental indeterminacy for which ε cannot be reduced and an apparent indeterminacy for which ε can be reduced. Fundamental indeterminacy challenges the scientific method. Apparent indeterminacy only serves to challenge scientists. In which category does the Poisson-like variability of cortical neurons belong?

Popper (1934) argued that science proceeds by falsification. A hypothesis is never proven; it is just discarded when it becomes clearly false. For this reason, Popper was deeply troubled by scientific theories that were fundamentally probabilistic in nature. If a theory proposes that a given neuron will fire an action potential with a probability of 0.2 in the next millisecond, any observation made during the next millisecond is commensurate with the theory. Of course, the longer the neuron is observed, the more robustly the frequency of action potential generation can be described in the past, but the ability of the theory to predict the future remains untestable and perhaps even unscientific, Popper argued. In the end, the theory is untestable because it predicts that given an infinitely long period of observation all possible patterns of action potentials will occur and thus no given pattern can be used to formally falsify the hypothesis.

In the behavioral sciences, however, even determinate theories rarely proceed through a process of unambiguous falsification. Nearly all of the measurements made by behavioral scientists are clouded by variability. Variability results from measurement error, from uncontrolled factors that influence the outcome of the experiment, and perhaps even from variability intrinsic to the system under study. As a result, behavioral hypotheses typically are falsified not with unambiguous observations but with statistical generalizations. Further, falsification tends to be iterative. Instead of demonstrating that a single observation is incompatible with a given hypothesis, behavioral scientists gather a distribution of observations and use this distribution to assess the accuracy of the theory. An existing theory is replaced when a new theory can account for a portion of the residual variance unexplained by the old theory. Behavioral scientists accept that measurements are clouded by variance, ε ; they work to minimize the magnitude of ε and they make statistical arguments that accommodate ε . However, at a fundamental level the goal of the scientific method remains a reduction in ε . Bacon argued that science must reduce uncertainty, and for working scientists, this usually means reducing ε .

For this reason it is difficult to use the formal logical approach embodied by the Popperian scientific method to argue that variance itself, intrinsic indeterminacy, is a fundamental property of a behavioral system. Accepting the level of variance associated with our best theories as the lowest possible variance necessarily forces an abandonment of further inquiry. If some arbitrarily observed variation in a set of measurements is presumed, a priori, to reflect an irreducible feature of the system under study, then there is no reason to engage in further scientific examination. The search for new theories is, in essence, a technique for reducing ε .

Over the last century, scholars seeking to understand behavior have struggled with this problem because they have again and again identified systems in which

variability, ε , seems irreducible. Neuringer (1986), for example, demonstrated that under some conditions human behavior is indistinguishable from a fundamentally indeterminate process. Tolhurst et al. (1981) and Dean (1981) made similar observations in their studies of cortical neurons. How can we ever hope to rigorously test hypotheses that include irreducible variation if the scientific method always seeks to reduce variance to zero?

One answer would be to employ a strategy first used by the quantum physicists who encountered evidence of fundamental uncertainty in the physical world. Consider as they did a fundamental process, like their quantum events, that impose a known amount of uncertainty, ω , on a set of scientific measurements. Under such conditions scientists would still make measurements and those measurements would still include an uncertain component, $\varepsilon_{\text{total}}$, but under those conditions $\varepsilon_{\text{total}}$ would be the sum of ω and the variances due to factors like measurement error, $\varepsilon_{\text{error}}$. As theories were replaced iteratively by increasingly accurate theories, $\varepsilon_{\text{total}}$ would begin to approach ω . Under these conditions, knowledge of the value of ω would solve the measurement problem posed by the existence of fundamental indeterminacy. Knowledge of the minimum possible level of indeterminacy would allow one to discriminate between an error on the part of the scientist and variability in the world. One of the two problems posed by uncertainty would become tractable.

The existence of a known nonzero ω , however, would do nothing to resolve the falsification problem. Under conditions in which ω has a nonzero value, scientific predictions must always be couched in probabilistic terms and thus rigorous falsification would remain impossible. The two critical issues that would arise were the behavioral world to be indeterminate would therefore be whether ω could be determined, and how one could proceed without rigorous falsification as a scientific goal.

Unfortunately, behavioral scientists do not yet have a theory that would allow them to specify the magnitude of ω , and it seems unlikely that such a theory is imminent. This is probably a very important problem, and one with which behavioral scientists are beginning to grapple. The most promising strategy for defining ω today may be to develop converging evidence, from several levels of analysis, for a specific value of ω under a specific set of behavioral conditions. For example, if game theoretic, behavioral, and neurobiological studies all suggested a specific value for ω under some set of conditions, then one could be much more confident that the traditional scientific method could be pursued. The data presented here suggest that the rudiments of just such an approach may be evolving, although it is far too early to suggest that estimates of the behavioral uncertainty intrinsic to any given situation can be made accurately. Research like that of VonNeumann, Neuringer, and Shadlen points to the existence of indeterminate elements that participate in the generation of behavior, and each provides quantitative estimates of that indeterminacy. One goal of these approaches, in the long run, will have to be quantitative convergence around specific predictions for ω .

The loss of rigorous falsification may be a more difficult philosophical problem, but may pose fewer difficulties to us as working scientists, especially since quantum

physicists have already begun to engage that problem. In practice, scientists rarely proceed through a process of unambiguous falsification. Instead, we often test theories against each other. We ask which theory provides more explanatory power, which yields a smaller ε , and we then discard the less-efficient theory. Theories are used to falsify each other in an iterative process, and there is nothing about this sequence of events that requires determinacy in the real world. This process of iterative falsification does work, although it is less elegant than the strong falsification Popper advocated. It will probably have to form the philosophical basis on which the study of indeterminate behavior is based, and it likely will form an entirely adequate base.

Indeterminacy becomes particularly problematic, however, when we try to ask whether the best currently available theory is a good theory. We traditionally consider a theory good when the predictions it makes are accurate. However, under conditions in which ω is large, good theories—even perfect theories—may not be accurate. An interesting example of this may be contemporary models developed to explain the choice behavior of humans during economic games (Camerer 2003, Erev & Roth 1998). These models seek to explain the play-by-play actions of individuals engaged in repeated rounds of games like rock-paper-scissors. The models seek to explain how human players learn from their experiences, and then use what they have learned to generate actions. Under these game theoretic conditions, however, there may be good reasons to believe that behavior is at least partially indeterminate. If behavior is truly uncertain on a choice-by-choice level, then how much of the behavior should a good theory explain? We can only assess the overall quality of theories like these if we can discover the fraction of the behavior that they seek to model, ω , which is truly indeterminate.

The critical point that all of these observations make is that if human behavior is at root indeterminate, we do not need to abandon the scientific method as it is practiced today. The existence of indeterminacy does raise measurement and falsification problems. The measurement problem can be addressed by efforts to bound ω . The falsification problem has to be addressed in a different way. If the behavioral world is indeterminate, we will have to abandon rigorous falsification. That would be a shame, but it is important to remember that as behavioral scientists, we typically rely on an iterative process of theory-by-theory falsification, and there is no compelling reason to believe that this iterative method is challenged by the existence of fundamental indeterminacy in behavior. These considerations suggest that behavioral indeterminacy may be a good deal less threatening to scientists and the scientific method than Popper may have feared originally. At the same time, the empirical observations presented in this review hint that behavioral indeterminacy may be much more likely to occur than Schrodinger imagined. He argued that fundamental indeterminacy would never arise in the living world because

If it were not so, if we were organisms so sensitive that a single atom, or even a few atoms, could make a perceptible impression on our senses—Heavens, what would life be like! To stress one point: an organism of that kind would

most certainly not be capable of developing the kind of orderly thought which, after passing through a long sequence of earlier stages, ultimately results in forming, among many other ideas, the idea of an atom. (Schrodinger 1944)

Our existing data, although ambiguous, clearly challenge Schrodinger's conclusion. The vertebrate nervous system is sensitive to the actions of single quantum particles. At the lowest levels of perceptual threshold, the quantum dynamics of photons, more than anything else, governs whether or not a human observer sees a light (Rieke & Baylor 1998). Synapses and neurotransmission also seem to violate this assumption of Schrodinger's, and these are the building blocks from which neurocomputation is achieved. In the end, Schrodinger may be right, behavior may be fundamentally determinate, but it would be premature to draw that conclusion now. Behavioral scientists will have to continue to explore apparent indeterminacy in behavior and will have to develop the methodological tools for determining whether this apparent indeterminacy is fundamental.

ACKNOWLEDGMENTS

The author expresses his profound gratitude to B. Lau and H. Bayer for hours spent in discussion. This work was supported by the James S. McDonnell Foundation and the National Eye Institute.

The *Annual Review of Psychology* is online at <http://psych.annualreviews.org>

LITERATURE CITED

- Arnauld A, Nicole P. 1662/1996. *Logic or the Art of Thinking*. Transl. JV Buroker, ed. London: Cambridge Univ. Press
- Auger C, Marty A. 2000. Quantal currents at single-site central synapses. *J. Physiol. (London)* 526:3–11
- Bacon F. 1620/1994. *Novum Organum*. Transl. P Urbach, J Gibson. Chicago: Open Court
- Bakan P. 1960. Response-tendencies in attempts to generate binary series. *Am. J. Psychiatry* 73:127–31
- Barlow HB, Levick WR. 1969. Three factors limiting the reliable detection of light by retinal ganglion cells of the cat. *J. Physiol. (London)* 200:1–24
- Barracough DJ, Conroy ML, Lee D. 2004. Prefrontal cortex and decision making in a mixed-strategy game. *Nat. Neurosci.* 7:404–10
- Batista AP, Newsome WT. 2000. Visuo-motor control: giving the brain a hand. *Curr. Biol.* 10:R145–48
- Bernoulli D. 1738/1954. Exposition of a new theory on the measurement of risk. Transl. R Sommer. *Econometrica* 22:23–36
- Blough DS. 1966. The reinforcement of least frequent interresponse times. *J. Exp. Anal. Behav.* 9:581–91
- Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA. 1996. A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis. Neurosci.* 13:87–100
- Camerer C. 2003. *Behavioral Game Theory*. Princeton, NJ: Princeton Univ. Press
- Corrado GS, Sugrue LP, Newsome WT. 2003. The efficacy of linear integration for reward gathering in a nonstationary foraging task. *Soc. Neurosci. Abstr.* 385.2
- Dean AF. 1981. The variability of discharge

- of simple cells in the cat striate cortex. *Exp. Brain Res.* 44:437–40
- Descartes R. 1637/2000. *Discourse on Method*. Transl. DM Clarke. Harmondsworth, UK: Penguin
- Dorris MC, Glimcher PW. 2004. Lateral intraparietal area activity during manipulations of the expected value and utility of choices. *Soc. Neurosci. Abstr.* 767.1
- Erev I, Roth AE. 1998. Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria. *Am. Econ. Rev.* 88:848–81
- Galilei G. 1630/2001. *Dialogo Sopra i due Massimi Sistemi del Mondo, Tolemaico e Copernicano (Dialogue Concerning the Two Chief World Systems, Ptolemaic and Copernican)*. Transl. S Drake. New York: Modern Library
- Glimcher PW. 2003. The neurobiology of visual-saccadic decision making. *Annu. Rev. Neurosci.* 25:133–79
- Glimcher PW, Dorris MC. 2005. Neuronal studies of decision making in the visual-saccadic system. In *The Cognitive Neurosciences*, ed. MA Gazzaniga. Cambridge, MA: MIT Press. In press
- Hastie R. 2001. Problems for judgment and decision making. *Annu. Rev. Psychol.* 52:653–83
- Hegglund P, Albus K. 1978. Response variability and orientation discrimination of single cells in striate cortex of cat. *Exp. Brain Res.* 32:197–211
- Heisenberg W. 1930. *Physical Principles of Quantum Theory*. New York: Dover
- Heisenberg W. 1952/1979. *Philosophic Problems of Quantum Physics*. Woodbridge, CT: Ox Bow
- Kepler J. 1618–21/1995. *Epitome of Copernican Astronomy*. Transl. CG Wallis. New York: Prometheus
- Keynes JM. 1936/1997. *The General Theory of Employment, Interest, and Money*. New York: Prometheus
- Lau B, Glimcher PW. 2003. How monkeys learn from their decisions: behavior and models for physiology. *Soc. Neurosci. Abstr.* 518.16
- Machado A. 1989. Operant conditioning of behavioral variability using a percentile reinforcement schedule. *J. Exp. Anal. Behav.* 52:155–66
- Mainen ZF, Sejnowski TJ. 1995. Reliability of spike timing in neocortical neurons. *Science* 268:1503–6
- Maynard Smith J. 1982. *Evolution and the Theory of Games*. London: Cambridge Univ. Press
- Nash JF. 1950a. Equilibrium points in n-person games. *Proc. Natl. Acad. Sci. USA* 36:48–49
- Nash JF. 1950b. The bargaining problem. *Econometrica* 18:155–62
- Nash JF. 1951. Non-cooperative games. *Ann. Math.* 54:286–95
- Neuringer A. 1986. Can people behave “randomly?” The role of feedback. *J. Exp. Psychol.: Gen.* 115:62–75
- Neuringer A. 2002. Operant variability: evidence, functions, theory. *Psychon. Bull. Rev.* 9:672–705
- Newsome WT, Britten KH, Movshon JA. 1989. Neuronal correlates of a perceptual decision. *Nature* 341:52–54
- Nimchinski EA, Yasuda R, Oertner TG, Svoboda K. 2004. The number of glutamate receptors opened by synaptic stimulation in single hippocampal spines. *J. Neurosci.* 24:2054–64
- Page S, Neuringer A. 1985. Variability is an operant. *J. Exp. Psychol.: Anim. Behav. Process.* 11:429–52
- Pascal B. 1670/1966. *Pensees*. Transl. AJ Krailsheimer. Harmondsworth, UK: Penguin
- Pavlov IP. 1927. *Conditioned Reflexes: An Investigation of the Physiological Activity of the Cerebral Cortex*. London: Oxford Univ. Press
- Popper K. 1934/1992. *The Logic of Scientific Discovery*. Transl. K Popper. London: Routledge
- Rapoport A, Budescu DV. 1992. Generation of random series in two-person strictly competitive games. *J. Exp. Psychol.: Gen.* 121:352–62
- Rieke F, Baylor D. 1998. Origin of reproducibility in the responses of retinal rods to single photons. *Biophys. J.* 75:1836–57

- Rieke F, Warland D, de Reuter van Steveninck R, Bialek W. 1997. *Spikes*. Cambridge, MA: MIT Press
- Schall JD. 2004. On building a bridge between brain and behavior. *Annu. Rev. Psychol.* 55:23–50
- Schrodinger E. 1944/1967. *What is Life?* London: Cambridge Univ. Press
- Schrodinger E. 1951. *Science and Humanism: Physics in Our Time*. London: Cambridge Univ. Press
- Shadlen MN, Britten KH, Newsome WT, Movshon JA. 1996. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J. Neurosci.* 16:1486–510
- Shadlen MN, Newsome WT. 1998. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* 18:3870–96
- Shadlen MN, Newsome WT. 2001. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* 86:1916–36
- Shafir E, LeBoeuf RA. 2002. Rationality. *Annu. Rev. Psychol.* 53:491–517
- Sherrington CS. 1906. *The Integrative Action of the Nervous System*. New York: Scribner's
- Shimp CP. 1967. Reinforcement of least-frequent sequences of choices. *J. Exp. Anal. Behav.* 10:57–65
- Skinner BF. 1938. *The Behavior of Organisms*. East Norwalk, CT: Appleton-Century-Crofts
- Smith A. 1776/1991. *Wealth of Nations*. New York: Prometheus
- Staddon JER, Cerutti DT. 2003. Operant conditioning. *Annu. Rev. Psychol.* 54:115–44
- Stevens CF. 2003. Neurotransmitter release at central synapses. *Neuron* 40:381–88
- Takumi Y, Ramirez-Leon V, Laake P, Rinivik E, Ottersen OP. 1999. Differential modes of expression for AMPA and NMDA receptors in hippocampal synapses. *Nat. Neurosci.* 2:618–24
- Tolhurst DJ, Movshon JA, Dean AF. 1981. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision. Res.* 23:775–85
- VonNeumann JV, Morgenstern O. 1944. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton Univ. Press
- Wagenaar WA. 1972. Generation of random sequences by human subjects: a critical survey of literature. *Psychol. Bull.* 77:65–72

CONTENTS

Frontispiece— <i>Richard F. Thompson</i>	xviii
PREFATORY	
In Search of Memory Traces, <i>Richard F. Thompson</i>	1
DECISION MAKING	
Indeterminacy in Brain and Behavior, <i>Paul W. Glimcher</i>	25
BRAIN IMAGING/COGNITIVE NEUROSCIENCE	
Models of Brain Function in Neuroimaging, <i>Karl J. Friston</i>	57
MUSIC PERCEPTION	
Brain Organization for Music Processing, <i>Isabelle Peretz</i> and <i>Robert J. Zatorre</i>	89
SOMESTHETIC AND VESTIBULAR SENSES	
Vestibular, Proprioceptive, and Haptic Contributions to Spatial Orientation, <i>James R. Lackner and Paul DiZio</i>	115
CONCEPTS AND CATEGORIES	
Human Category Learning, <i>F. Gregory Ashby and W. Todd Maddox</i>	149
ANIMAL LEARNING AND BEHAVIOR: CLASSICAL	
Pavlovian Conditioning: A Functional Perspective, <i>Michael Domjan</i>	179
NEUROSCIENCE OF LEARNING	
The Neuroscience of Mammalian Associative Learning, <i>Michael S. Fanselow and Andrew M. Poulos</i>	207
HUMAN DEVELOPMENT: EMOTIONAL, SOCIAL, AND PERSONALITY	
Behavioral Inhibition: Linking Biology and Behavior Within a Developmental Framework, <i>Nathan A. Fox, Heather A. Henderson,</i> <i>Peter J. Marshall, Kate E. Nichols, and Melissa A. Ghera</i>	235
BIOLOGICAL AND GENETIC PROCESSES IN DEVELOPMENT	
Human Development: Biological and Genetic Processes, <i>Irving I. Gottesman and Daniel R. Hanson</i>	263

SPECIAL TOPICS IN PSYCHOPATHOLOGY

- The Psychology and Neurobiology of Suicidal Behavior,
Thomas E. Joiner Jr., Jessica S. Brown, and LaRicka R. Wingate 287

DISORDERS OF CHILDHOOD

- Autism in Infancy and Early Childhood, *Fred Volkmar,
Kasia Chawarska, and Ami Klin* 315

CHILD/FAMILY THERAPY

- Youth Psychotherapy Outcome Research: A Review and Critique
of the Evidence Base, *John R. Weisz, Amanda Jensen Doss,
and Kristin M. Hawley* 337

ALTRUISM AND AGGRESSION

- Prosocial Behavior: Multilevel Perspectives, *Louis A. Penner,
John F. Dovidio, Jane A. Piliavin, and David A. Schroeder* 365

**INTERGROUP RELATIONS, STIGMA, STEREOTYPING,
PREJUDICE, DISCRIMINATION**

- The Social Psychology of Stigma, *Brenda Major
and Laurie T. O'Brien* 393

PERSONALITY PROCESSES

- Personality Architecture: Within-Person Structures and Processes,
Daniel Cervone 423

PERSONALITY DEVELOPMENT: STABILITY AND CHANGE

- Personality Development: Stability and Change, *Avshalom Caspi,
Brent W. Roberts, and Rebecca L. Shiner* 453

WORK MOTIVATION

- Work Motivation Theory and Research at the Dawn of the Twenty-First
Century, *Gary P. Latham and Craig C. Pinder* 485

GROUPS AND TEAMS

- Teams in Organizations: From Input-Process-Output Models to IMOI
Models, *Daniel R. Ilgen, John R. Hollenbeck, Michael Johnson,
and Dustin Jundt* 517

LEADERSHIP

- Presidential Leadership, *George R. Goethals* 545

PERSONNEL EVALUATION AND COMPENSATION

- Personnel Psychology: Performance Evaluation and Pay for Performance,
Sara L. Rynes, Barry Gerhart, and Laura Parks 571

**PSYCHOPHYSIOLOGICAL DISORDERS AND PSYCHOLOGICAL EFFECTS
ON MEDICAL DISORDERS**

- Psychological Approaches to Understanding and Treating Disease-Related
Pain, *Francis J. Keefe, Amy P. Abernethy, and Lisa C. Campbell* 601

TIMELY TOPIC

- Psychological Evidence at the Dawn of the Law's Scientific Age,
David L. Faigman and John Monahan 631

INDEXES

- Subject Index 661
Cumulative Index of Contributing Authors, Volumes 46–56 695
Cumulative Index of Chapter Titles, Volumes 46–56 700

ERRATA

An online log of corrections to *Annual Review of Psychology* chapters
may be found at <http://psych.annualreviews.org/errata.shtml>