Published in final edited form as:

Behav Processes. 2001 May 3; 54(1-3): 33–52.

Modeling games from the 20th century

P.R. Killeen*

Department of Psychology, Arizona State University, Tempe, AZ 85287-1104, USA

Abstract

A scientific framework is described in which scientists are cast as problem-solvers, and problems as solved when data are mapped to models. This endeavor is limited by finite attentional capacity which keeps depth of understanding complementary to breadth of vision; and which distinguishes the process of science from its products, scientists from scholars. All four aspects of explanation described by Aristotle trigger, function, substrate, and model are required for comprehension. Various modeling languages are described, ranging from set theory to calculus of variations, along with exemplary applications in behavior analysis.

Keywords

Causality; Complementarity; Models; Theories; Truth

1. Introduction

It was an ideal moment for an aspiring young man to enter the field. Half a century of laboratory research had generated an unparalleled backlog of data that demanded understanding. Very recent experiments had brought to light entirely new kinds of phenomena. The great twenty-[first] century upheavals that were to rock [psychology] to its foundations had barely begun. The era of classical [psychology] had just come to an end. Abraham Pais, *Neils Bohr's Times* 52–53.

Society supports science because it is in society's interest to do so. Every grant application has scientists underline the redeeming social qualities of their work; most students are most interested in applications; and scientists often describe their work to their neighbors in terms of its implications for every man. Outstanding discoveries with practical consequences, such as that of the electron, have 'coat-tails' that support generations of more esoteric inquiries. But application is not the goal of science; it is the goal of its sibling, technology. Technology uses scientific structures to change the world, whereas science uses technology to change its structures. This is an essay on the interplay between scientific structures – the theories and models that constitute knowledge – and their map to the empirical world.

Science does not cumulate; science evolves. Just as telling students is less than teaching them, telling them what we know is less than teaching them to know. Science is the crest of the wave of knowledge: frothy, dangerous, and contemporary. Without an accumulated mass of water beneath a crest, it would be mere foam; without an accumulated mass of knowledge beneath a dissertation, it would be mere foam. But, however important that mass is, it is not science, but its product. It is not the thing that makes science addictive.

*Tel.: +1-602-9652555; fax: +1-602-9658544. E-mail address: E-mail: killeen@asu.edu (P.R. Killeen).

The history of science is cumulative, but its practice is evolutionary. Memory is finite: students cannot learn all that their mentors know and have time for creative work. Those students docile to intense library-work are often refractory to intense laboratory-work. Good scientists are problem-solvers, not pedants. Their's is not the comprehension of the scientific structure of a discipline in toto, but rather the master of a small part that they can perfect. The gift we give our students is not what we have seen, but a better way of looking; not solutions, but problems; not laws, but tools to discover them.

Great tools create problems; lesser tools solve them. There are many uncertainties in the world that are not considered problematic. Great tools transform such nescience into ignorance, reconstruing them as important gaps in knowledge. Method then recasts the ignorance as a series of problems and initiates a complementary research program. The bubble-chamber created problems for generations of physicists. The double-helix was important as a fact but more important as a cornucopia of problems that fed the careers of molecular biologists. Salivating dogs were only a mess until Pavlov recognized the significance of their 'psychic secretions'; his conditioning paradigm unleashed 100 year of problems and associated research programs. Experimental choices were made primarily by humans, until the 2-key experimental chamber made it easy to study the choices of pigeons and rats, which then dominated the operant literature for a generation. Contrast was primarily a confound — until conditions of reinforcement were systematically alternated with techniques such as multiple schedules, yielding an embarrassment of problems largely unsolved today. Constraints on learning were not part of a research program — until Garcia sickened his rats and found that they learned despite response-punishment delays of hours. Fabricating these problem-originating tools is a creative art; the original experiments in which they were deployed, however flawed, were *seminal*. Such paradigm creation eludes the present discussion, which focuses on the nature of the scientific problems they create, and the quantitative techniques that have been deployed to solve them. Discussion starts with the intellectual context of science — its framework. It then reviews the role of theories and some of their products — models — that may be useful in the analysis of behavior.

2. The complementarity, distribution, relativization, and truthfulness of explanations

2.1. The complementarity of explanation

Attention limits our ability to comprehend an explanation/theory/model. The limits can be extended by graphical and mathematical techniques, and by chunking — constructing macros that act as shorthand — but not indefinitely. Neils Bohr promulgated *complementarity theory* as an expression of such constraints. The name comes from the complementary angles created when lines intersect.

Complementarity occurs whenever some quantity is *conserved*. When lines intersect, the 180° measure of a line is conserved, as the angles on either side of the intersection must sum to that value. Bohr noted many scientific complements, such that the more one knows about one aspect, the less one can know about the other. Position and momentum are his classic complements. Precision and clarity, or intelligibility, are others. These are complementary because our ability to comprehend — to hold facts or lines of argument together — is limited. Detailed and precise exposition is a *sine qua non* of science; but if the details do not concern a problem of personal interest, they quickly overwhelm. Conversely, the large picture without the detailed substrate is a gloss. Both are necessary, but the more of one, the less of the other.

The more parameters in an equation, the more precisely it describes a phenomenon. Hundreds of parameters are used to describe the orbit of a satellite around the earth. But the more

parameters, the less certain we can be what each is doing, and the more likely it is that one is doing the work of some of the others. The more parameters, the greater the likelihood that their interactions will generate emergent phenomena. Precision is complementary to comprehension; and both are necessary.

Understanding the principle of complementarity is essential so that students do not discredit models for their complexity, or discredit glosses on them for their superficiality. Complementarity arises from a constraint on our processing abilities, not a shortcoming of a particular theoretical treatment. In a microscope, field of view is conserved; precise visualization of detail must sacrifice a larger view of the structure of the object. In a scientist's life, time is conserved, so that efforts at understanding the relation of one's problem to the larger whole is time away from perfecting technique. One can survey the landscape or drill deeper, but one cannot do both at the same time.

Scientists have yet to develop a set of techniques for changing the field of view of a theory while guaranteeing connectedness through the process: theoretical depth of focus is discrete, not continuous. Ideally, all models should demonstrate that they preserve phenomena one level up and one level down. Non-linear interactions, however, give rise to 'emergent phenomena' not well-handled by tools at a different level. One might show, for instance, that verbal behavior is consistent with conditioning principles. But those principles by themselves are inadequate to describe most of the phenomena of speech.

Constraints on resources exacerbate theoretical distinctions. To provide lebensraum for new approaches, protagonists may deny any relevance to understanding at a different level, much as eucalyptus trees stunt the growth of competing flora. Complementarity of resources — light and moisture in the case of trees, money and student placements in the case of scientists — thus accelerates the differentiation of levels and helps create the universities of divergent inquiries so common today.

2.2. Distribution of explanation

A different complementarity governs what we accept as explanation for a phenomenon. It is often the case that a single kind of explanation satisfies our curiosity, leaving us impatient with attempts at other explanations that then seem redundant. But there are many types of valid explanation, and no one kind by itself can provide comprehension of a phenomenon. Belief that one type suffices creates unrealistic expectations and intellectual chauvinism. Comprehension requires a distribution of explanations, and in particular, the distribution given by Aristotle's four (be)causes:

1. *Efficient causes.* These are events that occur before a change of state and trigger it (sufficient causes). Or they do not occur before an expected change of state, and their absence prevents it (necessary causes). These are what most scholars think of as cause. They include Skinner's 'variables of which behavior is a function'.
2. *Material causes.* These are the substrates, the underlying mechanisms. Schematics of underlying mechanisms contribute to our understanding: the schematic of an electronic circuit helps to troubleshoot it. Neuroscientific explanations of behavior exemplify such material causes. Assertions that they are the best or only kind of explanation is reductionism.
3. *Final causes.* The final cause of an entity or process is the reason it exists — what it does that has justified its existence. Final causes are the consequences that Skinner spoke of when he described *selection by consequences*. Assertion that final causes are time-reversed efficient causes is *teleology*: results cannot bring about their efficient causes. But final causes are a different matter. A history of results, for

instance, may be an agent. A history of conditioning vests in the CS a link to the US; the CS is empowered as an efficient cause by virtue of its (historical) link to a final cause that is important to the organism. Explanations in terms of reinforcement are explanations in terms of final causes. Whenever individuals seek to understand a strange machine and ask “What does that do?”, they are asking for a final cause. Given the schematic of a device (a description of mechanism), we can utilize it best if we are also told the purpose of the device. There are many final causes for a behavior; ultimate causes have to do with evolutionary pressures; more proximate ones may involve a history of reinforcement or intentions.

4. *Formal causes.* These are analogs, metaphors and models. They are the structures with which we represent phenomena, and which permit us to predict and control them. Aristotle’s favorite formal cause was the syllogism. The physicist’s favorite formal cause is a differential equation. The chemists’ is a molecular model. The Skinnerian’s is the three-term contingency. All understanding involves finding an appropriate formal cause — that is, mapping phenomena to explanations having a similar structure to the thing explained. Our sense of familiarity with the structure of the model/explanation is transferred to the phenomenon with which it is put in correspondence. This is what we call understanding.

Why did Aristotle confuse posterity by calling all four of these different kinds of explanation *causes*? He did not. Posterity confused itself (Santayana characterized those translators/interpreters as ‘learned babblers’). To remain consistent with contemporary usage, these may be called causal, reductive, functional and formal explanations, respectively. No one type of explanation can satisfy: comprehension involves getting a handle on all four types. To understand a pigeon’s key-peck, we should know something about the immediate stimulus (Type 1 explanation), the biomechanics of pecking (Type 2), and the history of reinforcement and ecological niche (Type 3). A Type 4 explanation completes our understanding with a theory of conditioning. Type 4 explanations are the focus of this article.

2.3. Relativization of explanation

A formal explanation proceeds by apprehending the event to be explained and placing it in correspondence with a model. The model identifies necessary or sufficient antecedents for the event. If those are found in the empirical realm, the phenomenon is said to be explained. An observer may wonder why a child misbehaves, and suspect that it is due to a history of reinforcement for misbehavior. If she then notices that a parent or peer attends to the child contingent on those behaviors, she may be satisfied with an explanation in terms of conditioning. Effect (misbehavior) + model (law of effect) + map between model and data (reinforcement is observed)=explanation. Explanation is a relation between the models deployed and the phenomena mapped to them.

The above scenario is only the beginning of a scientific explanation. Confounds must be eliminated: although the misbehavior appears to have been reinforced, that may have been coincidence. Even if attention was the reinforcer which maintains the response, we may wish to know what variables established the response, and what variables brought the parents or peers to reinforce it. To understand why a sibling treated the same way does not also misbehave, we must determine the history and moderator variables that would explain the difference. All of this necessary detail work validates the map between the model and the data; but it does not change the intrinsic nature of explanation, which is bringing a model into alignment with data.

Prediction and control also involve the alignment of models and data. In the case of *prediction* a causal variable is observed in the environment, and a model is engaged to foretell an outcome. A falling barometer along with a manual, or model, for how to read it, enables the

sailor to predict storm weather. Observation that students are on a periodic schedule of assignments enables the teacher to predict post-reinforcement pausing. The demonstration of conformity between model and pre-existing data is often called prediction. That is not prediction, however, but rather post-diction. Such alignment signifies important progress and is often the best the field can do; but it is less than prediction. This is because, with outcome in hand, various implicit stimuli other than the ones touted by the scientist may control the alignment; as may various ad hoc responses, such as those involved in aggregation or statistical evaluation of the data. Those stimuli and responses may not be understood or replicable when other scientists attempt to employ the model. Journal editors should, therefore, require that such mappings be spoken of as ‘the model is consistent with/conforms to/ gives an accurate account of/the data.

In the case of *control*, the user of a model introduces a variable known to bring about a certain effect. A model stating that water vapor is more likely to condense in the presence of a nucleus may lead a community to seed the passing clouds to make rain. A model stating that conditioned reinforcers can bridge otherwise disruptive delays of reinforcement may lead a pet owner to institute clicker training to control the behavior of her dog. The operation of a model by instantiating the sufficient conditions for its engagement constitutes control. Incomplete specification or manipulation of the causal variables may make the result probabilistic.

Explanation is thus the provision of a model with (a) the events to be explained as a consequent, and (b) events noticeable in the environment as antecedents; successful explanations make the consequents predictable from the antecedents, given that model. *Prediction* is the use of a model with antecedents being events in the environment and a consequent to be sought in the environment. If the consequents are already known, the relation is more properly called ‘accounting for’ rather than ‘prediction of’. *Control* is the arrangement of antecedents in the context of a model that increases the probability of desired consequences.

2.4. The truth of models

Truth is a state of correspondence between models and data. Models are neither true nor false per se; truth is a relative predicate, one that requires specification of both the model and the data it attempts to map. *He is 40-years-old* has no truth value until it is ascertained to whom the ‘he’ refers. $2 + 2 = 4$ has no truth value. It is an instance of a formal structure that is well-formed. $2 \text{ apples} + 2 \text{ peaches} = 4 \text{ pieces of fruit}$ is true. To make it true, the descriptors/ dimensions of the things added had to be changed as we passed the plus sign, to find a common set within which addition could be aligned. Sometimes this is difficult. What is: $2 \text{ apples} + 2 \text{ artichokes}$? Notice the latency in our search for a superset that would embrace both entities? Finding ways to make models applicable to apparently diverse phenomena is part of the creative action of science. Constraining or reconstruing the data space is as common a tool for improving alignment as is modification of the model.

Not only is it necessary to map the variables carefully to their empirical instantiations, it is equally important to map the operators. They symbol ‘+’ usually stands for some kind of physical concatenation, such as putting things on the same scale of a balance, or putting them into the same vessel. If it is the latter, then $2 \text{ gallons of water} + 2 \text{ gallons of alcohol} = 4 \text{ gallons of liquid}$ is a false statement, because those liquids mix in such a way that they yield less than 4 gallons. *Reinforcement increases the frequency of a response*. This model aligns with many data, but not with all data. It holds for some hamster responses, but not others. Even though you enthusiastically thanked me for giving you a book, I will not give you another copy of the same book. That’s obvious. But why? Finding a formal structure that keeps us from trying to apply the model where it does not work is not always so easy. Presumably here it is ‘A good doesn’t act as a reinforcer if the individual is satiated for it, and having one copy of a book provides indefinite satiation’. Alternatively, one may define reinforcement in terms of effects

rather than operations, so that reinforcement *must* always work, or it is not called *reinforcement*. But that merely shifts the question to why a proven reinforcer (the book) has ceased to be reinforcing. Information is the reduction of uncertainty. If uncertainty appears to be dispelled without information, one can be certain that it has merely been shifted to other, possibly less obvious, maps. Absent information, uncertainty is conserved.

The truth of models is relative. A model is true (or *holds*) within the realm where it accurately aligns with data, for those data. A false model may be made true by revising it, or by restricting the domain to which it applies. Just as all probabilities are conditional (upon their universe of discourse), the truth of all models is conditional upon the data set to which they are applied. Life is sacred, except in war; war is bad, except when fought for justice; justice is good, except when untempered by humanity. Assignment of truth value, like the assignment of moral value, or of any label to a phenomenon, is itself thus a modeling enterprise, not a discovery of absolutes.

Truth is the imposition of a binary predicate on a nature that is usually graded; it is relative to the level of precision with which one needs to know, and to competing models. *The earth is a sphere* is in good enough alignment with measurement to be considered true. It accounts for over 99.9% of the variance in the shape of the earth. *Oblate spheroid* is better (truer), and when that model became available, it lessened the truthfulness of *sphere*. *Oblate spheroid with a bump in Nepal and a wrinkle down the western Americas* is better yet (truer), and so on. Holding a correspondent to a higher level of accuracy than is necessary for the purposes of the discussion is called *nitpicking*. Think of the truth operator as $truth(m, x, p, a) \in \{T, F, U\}$; it measures the alignment between a model (m) and a data set (x) in the context of a required level of precision (p) and alternative models (a) to yield a decision from the set *True, False, Undecided*.

A model shown to be false may be more useful than truer ones. False models need not, *pace* Popper, be rejected. Newtonian mechanics is used every day by physicists and engineers; if they had to choose one tool, the vast majority would choose it over relativity theory. They would rather reject Popper and his falsificationism than reject Newton and his force diagrams. It is trivial to show a model false; restricting the domain of the model or modifying the form of the model to make it truer is the real accomplishment.

3. Tools of the trade

A distinction must be made between *modeling tools*, which are sets of formal structures (e.g. the rules of addition, or the calculus of probabilities), and *models*, which are such tools applied to a data domain. Mechanics is a modeling tool. A description of the forces and resultants on a baseball when it is hit is a model. The value of tools derives from being flexible and general, and, therefore, capable of providing models for many different domains. They should be evaluated not on their strength in accounting for a particular phenomenon, but on their ability to generate models of phenomena of interest to the reader. We do not reject mechanics because it cannot deal with combustion, but rather we find different tools.

3.1. Set theory

Behavioral science is a search in the empirical domain for the variables of which behavior is a function; and a search in the theoretical domain for the functions according to which behavior varies. Neither can be done without the other. *The functions according to which behavior varies are models*. Models may be as complex as quantum mechanics. They may be as simple as rules for classification of data into sets: classification of blood types, of entities as reinforcers, of positive versus negative reinforcement. Such categorization proceeds according to lists of criteria, and often entails panels of experts. Consider the criteria of the various juries who

decide whether to categorize a movie as *jejune*, a death as *willful*, or a nation as *favored*. Or those juries who categorize dissertations as *passing*, manuscripts as *accepted*, grants as *funded*.

A model for classifying events as reinforcers is: 'If, upon repeated presentation of an event, the immediately prior response increases in frequency, then call that event a positive reinforcer'. That is not bad; but it is not without problems. If a parent tells a child: "That was very responsible behavior you demonstrated last week at school, and we are very proud of you", and we find an increase in the kind of behavior the parents referred to, can we credit the parents' commendation as a reinforcer? It was delayed several days from the event. Their description 're-minded' the child of the event. Is that as good as contiguity? Even a model as simple as the law of effect requires qualifications. In this case, it requires either demonstrating that such commendations do not act as reinforcers; or generating a different kind of model than reinforcement to deal with them; or discarding the qualifier 'immediate'; or permitting representations, or memories, of events to be reinforced *and* to have that strengthening conferred on the things represented or remembered (and not just on the behavior of remembering). These theoretical steps have yet to be taken.

If our core categorical model requires more work, it is little surprise that stronger models also need elaboration, restriction, refinement, or redeployment. Such development and circumscription is the everyday business of science. It is not the case that the only thing we can do with models is disprove them, as argued by some philosophers. We can improve them. That recognition is an important difference between philosophy and science.

3.1.1. The 'generic' nature of the response—Skinner's early and profound insight was that the reflex must be viewed in set-theoretic terms. Each movement is unique. Reinforcement acts to strengthen movements of a similar *kind*. A *set* is a collection of objects that have something in common. A thing-in-common that movements have that make them responses could be that they look alike to our eyes. Or it could be that they look alike to a microswitch. Or it could be that they occur in response to reinforcement. Skinner's definition of the operant emphasized the last, functional definition. This is represented in Fig. 1. Operant responses are those movements whose occurrence is correlated with prior (discriminative) stimuli, and subsequent (reinforcing) stimuli. After conditioning, much of the behavior has come under control of the stimulus.

Whereas Skinner spoke in terms of operant movements as those selected by reinforcement, he primarily used the second, pragmatic definition in most of his research: responses are movements that trip a switch. Call this subset of movements 'target' responses, because they hit a target the experimenter is monitoring. Target responses are a proper subset of the class of functionally-defined operant responses. Other operant responses that the experimenter is less interested in are called *superstitious* responses or *collateral* responses or *style*. When Pavlov conditioned salivation, he reported a host of other emotional and operant responses, such as restlessness, tail-wagging, and straining toward the food dish, that were largely ignored. The target response was salivation. The collateral responses fell outside the sphere of interest of the experimenter, and are not represented in these diagrams.

Many of the analytic tools of set theory are more abstruse than needed for its applications in the analysis of behavior. But it is important to be mindful of the contingent nature of categorization into sets, and the often-shifting criteria according to which it is accomplished. Representation in these diagrams helps remind us of not only the properties of the responses that are measured, but their relation to other salient events in the environment, as shown in Fig. 2.

The top diagram depicts a discriminated partial-reinforcement schedule: reinforcement is only available when both an appropriate stimulus and movement have occurred, but does not always occur then. In Pavlovian conditioning, the target movement is uncorrelated with reinforcement (even though other movements may be necessary for reinforcement), as in the free-operant paradigm no particular stimulus is correlated with the delivery of reinforcement (other than the contextual ones of experimental chamber, and so on). Under some arrangements, a stimulus is a much better predictor of reinforcement than a movement, and tends to block conditioning of the movement.

3.2. Probability theory

Probabilities, it has been said, are measures of ignorance. As models of phenomena are improved, probabilities are moved closer to certainties. But when many variables interact in the construction of a phenomenon, a probabilistic, or stochastic, account may be the best we can ever do. This is the case for statistical thermodynamics, and may always be the case for meteorology. A stochastic account that provides accurate estimates of probabilities and other statistics may be preferred over a deterministic account that is only sometimes accurate.

Probabilities are long-run relative frequencies. They may be inferred from samples, deduced from parameters, deduced from physical arrangement, or estimated subjectively. Counting 47 heads in 100 tosses of a coin may lead you to infer that the probability of a head is ‘about 0.5’. This means that the inferred population of an indefinitely large number of tosses would show about half heads. I may have performed the same experiment with the same coin 1000 times already, with the result of 492 heads. From this knowledge I can deduce that your sample is likely to show 49 ± 5 heads. I deduce your statistic from my inference of the population parameter. The probabilities may be deduced from principles of fabrication or mixing: Extreme care was taken in milling this almost perfectly symmetric coin. I deduce that there is no reason for it to land on one side more often than another, and that the probability of heads is therefore $1/2$. Finally, I may estimate the probability of rain as 0.5. This may be interpreted as ‘in my experience, half the time the clouds and wind have been this way, it soon rained.’ Obviously, our confidence in these estimates varies with the method of estimation. Associated with the methods are models: models for physically counting, or mixing red and green balls in an urn, or categorizing clouds.

Of course, the model may not hold; I may forget to replace the items I have sampled, or the urn in which the balls are contained may not have been thoroughly mixed, or I may prefer the way red balls feel, and select them differentially, and so on. Like all models, probability models stipulate many of the conditions under which they are useful. They may be useful beyond those conditions, but then *caveat emptor*. If a statistic is not normally distributed, conventional statistical tests may still be useful; but they may not be telling the user exactly what he thinks, and may lead him to false conclusions.

Appreciation of *recherché* statistical models may be deferred until confrontation by a problem that demands those tools. But basic probability theory is important in all behavioral analyses. It is thumbnailed here.

All probabilities are conditional on a universe of discourse. The rectangle in Fig. 3 represents the universe of discourse: All probabilities are measured given this universe. Other universes hold other probability relations. The probability of the Set *A* is properly the probability of Set *A* given the universe, or $p(A/U)$. It is the measure of *A* divided by the measure of the universe. Here those measures are symbolized as areas. Consider the disk *A* to be laid atop the universe, and throw darts at the universe from a distance such that they land randomly. If they hit *A*, they will also go through to hit *U*. Then we may estimate the probability of *A* as the number of darts that pierce *A* divided by the number that pierce *U*. This is a sample. As the number of darts

thrown increase, the estimates become increasingly accurate. That is, they converge on a limiting value.

The term *given* redefines the universe that is operative. The probability of *A* given *B* is zero in the universe of Fig. 3, because if *B* is given, it means that *it* becomes the universe, and none of the area outside it is relevant to this redefined universe. No darts that hit *B* can also hit *A*. The probability of *A* given $C=1$: all darts that go through *C* must also go through *A*, and because *C* is given, none that go elsewhere are counted.

The area of *D* is about 1/50 that of the universe, and about 1/20 that of *A*. Therefore, the probability of *D* given *A* is greater than the probability of *D* (given *U*). *A* and *D* are said to be positively correlated. When probabilities are stipulated without a conditional given, then they are called base rates, and the given is an implicit universe of discourse. It is usually helpful to make that universe explicit. The probability that what goes up will come down is 1.0 (*given* that it is heavier than air, is not thrown too hard, is thrown from a massive body such as the earth, has no source of propulsion, etc.).

The probability of *E* given *A*, $p(E/A)$, is less than the base rate for *E*, $p(E/U)$, so *E* and *A* are said to be negatively correlated. The $p(F/A)$ equals the $p(F/U)$, so the events *F* and *A* are said to be independent.

Instrumental conditioning occurs when the probability of a reinforcer given a response is greater than the base rate for reinforcement; behavior is suppressed if the probability is less than the base rate. What is the universe for measuring the base rate? Time in the chamber? When a trial paradigm is used, stipulation of the universe is relatively straightforward. When behavior occurs in real-time, the given could entail a universe of the last 5, 50 or 500 s. But a response 5 min remote from a reinforcer will not be conditioned as well as one occurring 5 s before it. Our goal is to define the relevant universe (context) the same way as the animal does, so that our model predicts its behavior. Exactly how time should be partitioned to support a probability-based (i.e. correlation-based) law of effect is as yet an unsolved problem. Just as all probabilities are conditional (on some universe), all conditioning is, as its name implies, conditional: If no stimuli are supplied by the experimenter, some will nonetheless be found by the organism.

Table 1 gives the names of experimental paradigms or the resulting behavior that is associated with various correlations. The second row assumes a (positive) reinforcer. The bottom row indicates that movements that co-occur tend to become part of the same operant class; this may be because they are physically constrained to do so (topographic effects), because they are strongly co-selected by reinforcement (operant movements), or because many different constellations of movements suffice, and whatever eventuates is selected by reinforcement (style). Responses that are independent can occur in parallel with no loss of control. Responses that compete for resources such as time or energy often encourage an alternation or exclusive choice. Responses that appear independent may be shown to be competing when resources are restricted, a phenomenon known in politics as the Ford effect.

3.2.1. Bayes—One stochastic tool of general importance is the Bayesian paradigm. Consider the probability spaces shown in Fig. 3. Whereas $p(A/C) = 1$, $p(C/A) < 1$. Here, the presence of *C* implies *A*, but the presence of *A* does not imply *C*. These kinds of conditional probabilities are, however, often erroneously thought to be the same. If we know the base rates, we can calculate one given the other, with the help of a chain rule for probabilities. Consider the area that is enclosed by both *A* and *D*. It is called the intersection of *A* and *D*, and can be calculated in two ways: First, $p(A.D) = p(A/D)p(D)$. The probability of both *A* and *D* occurring (given the universe) is the probability of *A* given *D*, times the probability of *D* (given the universe).

Second, $p(A.D) = p(D/A)p(A)$. From this we can conclude $p(A/D)p(D) = p(D/A)p(A)$, or $p(A/D) = p(D/A)p(A)/p(D)$. This last relation is Bayes' rule. The probability of D given A equals the probability of A given D times the ratio of base rates of A and D . If and only if the base rates are equal will the conditional probabilities be equal. The probability that you are a behaviorist given that you read The Behavior Analyst (TBA) equals the probability that you read TBA given that you are a behaviorist, multiplied by the number of behaviorists and divided by the number of TBA readers (in the relevant universe). Because there are many more behaviorists than readers of TBA, these conditional probabilities are not equal. In the context of Bayesian inference, the base rates are often called prior probabilities (prior to evaluating the conditionals), or simply priors.

3.2.2. Foraging problems involve Bayesian inference—The probability that a patch still contains food given that none has been found during the last minute equals the probability none would be found in the last minute given that it still contains food, times the relevant base rates. If the operative schedule is a depleting VI 300 s schedule, the conditional probability is good; if it is a depleting 15 s schedule, then it is bad. This is because 1 min of dearth on a VI 300 is typical and gives us little information; the same epoch on a VI 15 is atypical, and informs us that it is likely that the source has gone dry. This is the logic; Bayes gives the predictions.

Learning in the context of probabilistic reinforcement involves Bayesian inference. It is sometimes called the *credit allocation problem*. A pellet (R) is delivered after a lever-press (M) with a probability of $p(R/M)=1/5$. It is not enough to know that to predict the probability of L given P , which is $p(M/R) = p(R/M) \times p(M)/p(R)$. If the base-rates for reinforcers, $p(R)$, is high, conditioning is unlikely. In terms of mechanisms, it is said that the background is being conditioned, and little credit is left to be allocated to M . Thus, $p(M/R)$ is a better model for the conditional most relevant to conditioning than is $p(R/M)$.

Reinforcers or threats may elicit certain behavioral states, in the context of which specific responses (e.g. *prepared* or *species-specific* appetitive or defensive responses, such as pecking or flight for pigeons), may be vested by evolution with high priors, and other responses (e.g. *contraprepared* responses) with low priors.

Scientific inference involves Bayesian inference. What is the probability that a model, or hypothesis, is correct, given that it predicted some data, $p(H/D)$? Experiments and their statistics give us the probability that those data would have been observed, given the model, $p(D/H)$. These are not the same conditionals. Inferential statistics as commonly used tell us little about the probability of a model being true or false, the very question we usually invoke statistics to help us answer. Null hypothesis statistical testing confounds the problem by setting up a know-nothing model that we hope to reject. But we cannot reject models, unless we do so by using Bayes theorem. A p -level of less than 0.05 means that the probability of the data being observed given the null hypothesis is less than 5%. It says nothing about the probability of the null hypothesis given those data, and in particular it does not mean that the probability of the null hypothesis is less than 5%. To calculate the probability of the hypothesis, we need to multiply the p -level by the prior probability of the model, and divide by the prior probability of the data: $p(H/D) = p(D/H)p(H)/p(D)$. If the priors for the data are high — say, if my model predicts that the sun will rise tomorrow [$p(S/H) = 1$] — the probability of the model given a sunrise is not enhanced. This is because $p(S) \approx 1$, so that, $p(H/S) \approx 1 \times p(H)/1$. The experiment of rising to catch the dawn garners data that do little or nothing to improve the credibility of the model. If the priors are low, however — if the model predicts that at 81° ambient temperature, the sun will flash green as it rises, and it does — the model gains appreciably.

The difficulty in applying Bayesian inference is the difficulty in specifying the priors for the model that one is testing. Indeed, the very notion of assigning a probability to a model is alien

to some, who insist that models must be either true or false. But all models are approximations, and thus their truth value must be graded and never 1.0; burdening models with the expectation that some must be absolutely true is to undermine the utility of all models. So the question is, how to construe models so that the probability calculus applies to them. One way to do this is to think of a universe of models — let us say, all well-formed statements in the modeling language we are using. Then to rephrase the question as: ‘What is the probability that the model in question accounts for more of the variance in the empirical data than does some other model of equal complexity?’ Techniques for measuring computational complexity are now evolving.

Some models have greater priors than others, given the current state of scientific knowledge, and all we may need to know is order-of-magnitude likelihood. It is more likely that a green worm has put holes in our tomatoes than that an alien has landed to suck out their vital fluids. Finding holes in our tomatoes, you go for the bug spray, not the bazooka. Are the priors higher that a behavioral theory of timing or a dynamic theory of timing provides the correct model of Fixed Interval performance? This is more difficult to say, a priori. Because we may not be able to say, does not entail that we should not take the Bayesian perspective; for the problem is not a by-product of Bayesian inference. Bayes theorem merely gives voice to the problem of inverse inference most clearly. There are ways of dealing with these imponderables that is superior to ignoring them. Calculating the ratio of conditional probabilities for two competing hypotheses based on a common data set eliminates the need to estimate the priors on the data. Other techniques (e.g. entropy maximization) help determine the priors on the models.

Another approach is simply to feign indifference to Bayesian and all other statistical inference. This is more of a retreat than an approach; yet there remain corners of the scientific world where it is still the mode. But science inherently concerns making and testing general statements in light of data, and is thus intrinsically Bayesian. Whether or not one employs statistical-inference models, it is important to understand the probability that a general statement is true, based on a particular data set is, $p(H/D) = p(D/H)p(H)/p(D)$. This provides a qualitative guide to inference, even if no numbers are assigned.

3.3. Algebra

3.3.1. A question of balance—We use algebra regularly in calculating the predictions of simple models of behavior. Think of the equals-sign in an equation as the fulcrum of a balance. Algebra is a technique for keeping the beam horizontal when quantities are moved from one side to another. Angles of the beam other than 0° constitute error. A model is fit to data by putting empirical measurements on the left-hand side of the beam, and the equation without the y-value on the right-hand side of the beam. If the left side moves higher or lower than the right, the model mispredicts the data. The speed with which the beam deviates from horizontal as we repeatedly place pairs of x and y measurements in the right and left sides indicates the residual error in the model. Descartes’ analytic geometry gives us a means to plot these quantities as graphs.

3.3.2. Inducing algebraic models—Assume that the preference for a reinforcer increases as some function of its magnitude, and a different function of its delay. Imagine a procedure that lets us pair many values of amount with delay, and contrast them with an alternative of amount 1 at delay 1. The subjects generate a data set (Table 2), where the entries are the ratios of responses for the variable alternative to those for the constant alternative. Can we reconstitute the simple laws governing delay and amount from this matrix? In this case, there is no noise, and so the solution is straightforward (hint: focus on the first row and first column; graph their relation to the independent variables; generate hypothetical relations, and see if they hold for other cells). When there is noise in the data, or the functions that generate them are more

complex, other techniques, called conjoint measurement, or functional analysis, are available. They constitute the essence of Baconian induction.

For more than two dimensions, or where the dimensions can not be identified before-hand, or in cases where there is an additional transformation before an underlying construct — preference, strength, momentum — is manifest as behavior, other techniques may be used. In one of the most creative methodological inventions of the 20th century — non-metric multidimensional scaling — Roger Shepard showed how the mutual constraints in a larger version of such a table could be automatically parsed into relevant dimensions and transformations on them.

3.3.3. A new twist—Traditional algebras are so well-ingrained that we take their reality for granted. Yet, other algebras provide better modeling tools for some operations. Pick up this journal and rotate it counterclockwise 90° , and then flip it 180° from left to right. Note its position. Restore it to normal. Now perform those operations in reverse order. Its terminal position/orientation is different. These spatial operations do not commute: their order matters. This is also the case with some operations in matrix algebra. Conduct another experiment with two conditions: present a stimulus that can serve as a reinforcer to an organism (US), and then present a discriminative stimulus (CS). Repeat several times. Next find another pair of stimuli, and present them in the order CS, US. The behavior of the organism to these two sets of operations will be quite different, showing associative learning in the second case, but not in the first; or possibly inhibitor conditioning in the first. Do other algebras offer superior models for the operations of conditioning?

3.4. Calculus

Calculus is a generic term, meaning any modeling system. The set of laws governing probabilities is the probability calculus. But by itself, the term is most closely associated with techniques for finding the tangent to a curve (differential calculus) or the area under a curve (integral calculus). It was Newton who recognized that these systems were inverses: that the inverse of the differentiation operation, anti-differentiation — is equivalent to integration.

3.4.1. Differential—Why should such esoterica as tangents to curves and areas under them have proved such a powerful tool for science? Because change is fundamental, as noted by Parmenides, and calculus is the language of change. Behavior is change in stance over time. In Skinner's words, "it is the movement of an organism within a frame of reference". If we plot position on the y -axis and time on the x -axis, then the slope of the curve at any point measures the speed of motion at that point in time. This slope is the derivative of the function with respect to time. The slope of a cumulative record is the rate of responding at that point.

If we plot the proportion of choices of a sucrose solution as a function of its sweetness, the curve will go through a maximum and then decline as the solution becomes too sweet. At the point of the maximum the slope goes from positive — increases in sucrose increase preference — to negative — further increases in sucrose decrease preference. At the balance point the sweetness is *optimum*. We can define that optimum as the point at which the derivative of choice with respect to sucrose concentration goes from positive to zero. If the function governing preference were $p = ax - bx^2$, then this point would occur when $0 = a - 2bx$; that is, at $x = a/(2b)$. All talk of optimal foraging, or of behavior in a concurrent experiment as optimizing rate of reinforcement, or optimal arrangements of contingencies, can be thought of in terms of derivatives. Think of a multidimensional space, with one dimension being the dependent variable, and all the other dimensions being the various experimental conditions. The point at which the derivative of the dependent variable with respect to all the dependent variables goes to zero (if one exists) is the optimum. Finding the right balance in life — the

optimal trade-off between work, study, leisure, exercise, and community activities — is finding the point at which life-satisfaction has zero-derivatives. Be careful though: a slope of zero also occurs at the minimum, where the derivatives change from negative to zero. If you want to determine if you are as bad off as you could get, once again, check the derivative!

3.4.2. Integral—Behaviorists are seldom interested in calculating the area under a curve; yet it is something we do all the time, at least implicitly. The total number of responses in extinction is the area under the rate versus time graph. If reinforcers can increase the probability of more than one prior response, then the strengthening effect of reinforcement is the sum of the decaying trace of reinforcement multiplied by the probability of a response at all distances from the reinforcer.

To show how integration may be employed, consider the following model for the strength of a continuous CS: associative strength is the average association of a stimulus with a reinforcer. This may be calculated as the sum of each instant of the CS weighted by an exponentially decaying trace of the primary reinforcer, divided by the duration of the CS:

$S(CS) = 1/t \int_0^t e^{-kx} dx$. The solution of this equation is $S(CS) = (1 - e^{-kt})/(kt)$ which provides a plausible model of associative conditioning as a function of CS duration (see Fig. 4). The curve is similar to a hyperbolic function over the range of interest. If choice behavior is under the control of conditioned reinforcers bridging the delay until a primary reinforcer, this model predicts an almost-hyperbolic decay in preference as a function of delay.

3.4.3. The calculus of variations—One of the most elegant generalizations of the calculus is to use it not to find a point on a function that minimizes some value, but rather to find a *function* that minimizes some value. The most famous problem in the calculus of variations, and one of the oldest, is the brachistochrone problem: Given a starting point S , x ft above the floor, and an ending point E , on the floor y ft distant, what is the shape of a surface that will cause a ball rolling down the surface to arrive at E in the minimum time? There are many candidates, an inclined plane being the simplest — but not the fastest. What is needed is a statement of constraints, and a statement of the thing to be optimized. Here, the constraints are the coordinates of S and E , the force of gravity, and the system of mechanics that tells us how that force accelerates bodies. The thing to be optimized is transit time. There are an infinity of surfaces that will give long or maximum times; only one, the cycloid, that gives a minimum time. Another problem requiring the calculus of variations is that of determining the shape of a hanging clothesline, which distributes the force of gravity evenly over its length.

The brachistochrone problem was posed by Johann Bernoulli as a open problem; he received five solutions — an anonymous one within several days. Looking at the elegant style, he recognized the author, revealed ‘as is the lion by the claw print’; the print was Newton’s. (The other solutions came from Johann himself, his brother Jakob, one of their students, L’Hospital, and Leibniz). The calculus of variations is one step beyond differential calculus. As with any skill, some time spent working problems will confer ability. In the analysis of behavior, we might ask questions of the following sort:

- Brachistochrone/shaping: given a continuum of responses, a hyperbolic delay of reinforcement gradient, and a logarithmic satiation function, what is the allocation of reinforcers to behavior that will move an organism from an initial state S to a target state E as quickly as possible? What else needs to be specified?
- Clothesline/foraging: given that an organism must transit from point S to point E under a threat of predation that decreases as an inverse square function from some third point, what trajectory will minimize exposure to the threat?

- Surprise/conditioning: assume that an organism is conditioned by the appearance of unexpected events, and that conditioning acts to minimize future surprise. It does this because responses or stimuli that predict the US come to act as CSs. We may operationalize uncertainty as the temporal variance in predicting the US. Assume that temporal epochs (CSs) occur continuously before a US, that the organism has a limited amount of attention λ that it can deploy among the temporal stimuli, and that the standard deviation of prediction increases linearly with the time between a CS and the US. What is the form of the delay of reinforcement gradient for conditioning that minimizes surprise over a standard interval? Will, for instance, it be the exponential decay function used in the example above? Hyperbolic? Something else?

3.5. Decision theory

To be, or not to be; that is a question that decision theory can help answer. Decision theory is a way of combining measures of stimuli and reinforcers to predict which response will be most strongly reinforced. In the avatar of signal detection theory (SDT), decision theory focused on stimulus presence and predictiveness, and asked how to maximize reinforcement given payoffs of stipulated value for errors and corrects. More recently it has been generalized to incorporate a behavioral theory of reinforcement value, and also responses that are not perfectly discriminable. Although possibly strange to our ears, the notion of response discriminability is a central issue in our field: Whenever a new skill, such as golf, is being learned, it is often unclear what was done that resulted in a great/terrible shot. Response discrimination is a natural complement to stimulus discrimination.

Decision theory builds on probability theory, on Bayesian inference, and on optimality. Its icon is that of a *t*-test: overlapping Gaussian densities (Fig. 5), one representing the distribution of percepts (P_2) that might have arisen from a stimulus (S_2), the other the distribution of percepts arising from an alternate stimulus (S_1) or irrelevant stimuli (*noise*). It is the organism's challenge to determine the probability of S_2 given the percept P_2 ; but all he has a priori is a history that tells him the probabilities of the percepts given the stimuli, depicted by the densities. Bayes theorem is required: take the ratio of the evidence for S_2 over that for S_1 to determine a likelihood ratio

$$\frac{p(S_2|P_2)}{p(S_1|P_2)} = \frac{p(P_2|S_2)p(S_2)}{p(P_2|S_1)p(S_1)}$$

The optimal strategy is to select a criterial value for this ratio, determined by the payoffs for correctly identifying S_1 and S_2 and the penalties for mistakes, and to respond '2' if the percept is above the criterion, and '1' otherwise.

There are many variant models of SDT available. The core idea of the theory is to take into account the probabilistic nature of perception, assume that organisms are acting to optimize reward, and to create models that combine inferred knowledge of stimuli and reinforcement to predict responding.

Organisms are afloat on a raft of responses in a sea of stimuli, reinforcers and punishers (Fig. 6). Survival entails finding correlations between stimuli, responses and reinforcers that will maximize aspects of a life-trajectory. The particular tools, such as SDT, that are evolving to understand these relations are simple one- and two-dimensional slices of an evolving multidimensional space. It is the challenge of the next century to evolve models that more closely map the complexities we hope to understand.

3.6. Game theory

Another dimension is added to complexity when organisms operate in a *closed-loop* environment. These are situations in which their actions change the environment, which in turn changes future actions. Many experimental paradigms open these loops to establish relatively constant environments. A variable-interval schedule keeps rate of reinforcement relatively invariant over substantial changes in rate of responding. If all environments were open loop, there would be little point in learning. Plants largely inhabit open-loop environments. To the extent that an organism can benefit from sensitivity to the effects of its behavior — that is, to the extent that important features of its interaction with the world are closed-loop — learning will improve the organism's fitness. The difficulty in analyzing such systems is that small variations from expected reactions at one step quickly become magnified by succeeding steps.

The non-linear effects found in closed-loop systems are especially important when individuals interact. Consider three societies in which both kindness and meanness are reciprocated; In Society *E* each gives as good as he gets; In Society *A*, each gives 5% more than he gets; and in Society *D*, each gives 95% of what he gets. Which societies will be stable; which polarized? Can you generate a spreadsheet model of these societies?

If interactions are on a trials basis, with outcomes of each interaction known only after each individual makes a move, game theory provides a modeling system. Classic games include the *prisoner's dilemma*. Depending on the reinforcement contingencies, the most effective strategies (ones that optimize reward) may be mixed, with alternate responses chosen probabilistically.

Most models of conditioning assume a power differential — the experimenter sets the rules and motivates the organism. In game theory, however, the players have equal power, and the optimum sequence of responses are ones that not only provide the best short-term payoff, but also condition the opponent/cooperator to behave in ways that sustain payoffs. Since this may mean foregoing the largest short-term payoff, optimal strategies entail self-control. Conversely, self-control may be thought of as a game played against one's future self, an entity whose goals are similar to, but not the same as, those of the present self.

If the games are real-time, such as the game of chicken, dynamic system models are necessary. If signaling is possible, players will seek signs of character — that is, predictors of future behavior — in posture, couture, and coiffure; strategies of bluff, deception, and seduction are engaged; signal detection becomes a survival skill. Due to the potential instability of such interacting systems, strong reinforcement contingencies are necessary to avoid chaos. These are provided by charismatic leaders, repressive regimes, or elaborate legislative and legal systems reinforced by ritual and myth.

3.7. Automata theory

A toggle switch is a simple automaton. It alternates its state with each input. Each time it is switched on it may send a signal to another switch. Since it takes two operations — on and off to send a signal — the second switch will be activated at half the frequency of the former. A bank of such switches constitutes a *binary counter*. Another switch requires two simultaneous inputs to change its state; this constitutes an *and* gate. Another switch assumes a state complementary to its input; this is a *not* gate. Wired together properly these elements constitute a computer. As described, it is a *finite-state* automaton because it has a finite amount of memory. If an endless tape were attached so it had unlimited memory, a Turing machine could be constructed. Turing machines are universal, in that they can solve any problem that is, in theory, solvable.

Organisms such as rats and humans may be viewed as finite-state automata, differing primarily in the amount of memory that is available to them. This statement does not mean that they are nothing but automata. All models abstract from real phenomena to provide a more comprehensible picture. City maps do not show the trees or litter on the streets they describe, and are of reduced scale. Viewed as automata, the primary difference between rats and humans is the amount of memory available for computations. The more memory that is available, the more information about the stimulating context may be maintained to provide probability estimates. More memory means more capacity to retain and relate conditional probabilities. Enhanced ability to conditionalize permits nuanced reactions.

Automata have been used as metaphors for human computational ability for centuries, with analog computers favored 50-years ago, digital computers 30-years ago, parallel architectures 20-years ago, and genetic algorithms/Darwin machines 10-years ago. Within the learning community, they have seldom been used to generate models, with the implications of architecture and memory capacity taken seriously. Is this because they are intrinsically poor models of organisms, or because no one has tried?

4. Last words

The mind of science may be claimed by philosophy, but its heart belongs to tinkerers and problem solvers. Good scientists love puzzles, and the tools that help unravel them. The difference between scientists and anagram fans is the idea that scientific problems are part of a larger puzzle set; that one puzzle solved may make other pieces fall into place. But, then, jigsaw puzzles have that feature too.

Another difference is that the Nature, not the New York Times, created the scientists problems. But, in fact, all Nature did was exist; scientists themselves transformed aspects of it into problems to be solved. Skinner and his students, not Nature, put two keys in an ice-chest. Another difference is that the puzzles solved by scientists resonate with aspects of the world outside their problem set: Once a puzzle is solved, the scientist can turn back to the world from which it was abstracted and find evidence for the mechanisms discovered in the laboratory in the world at large. A model of speed on inclined planes speaks volumes about the velocity of all motions forced by gravity. But such full-cycle research — reinvestment of the solution in the natural world from which the problem was abstracted — is preached more often than practised. Most scientists are happy to give science away; putting it to work is not in their job description. Because an infinity of puzzles may be posed, society's selection of which puzzle solvers to support is biased by their perception of communal benefits. It follows that scientific societies must pay their way through applications, not assertions of eternal verities that are often beyond the ken of the community. Scientists must work hand-in-hand with technologists to survive.

One of the practical benefits of science is understanding, but understanding is itself only poorly understood. Physicists from Aristotle through Bohr to Einstein help us understand understanding. Aristotle taught us about its multidimensionality; Bohr of complementary limits on its depth and breadth. For Einstein, understanding was the reduction of a problem to another that we already think we understand — essentially, a successful map to a model.

Truth, we learn, is a relation, not a thing, even when proceeded by *the* and capitalized. It is a relation between a statement/model and a fact. Assignment of a truth value works best for binary statements concerning binary facts; but most data must be forced into such polarity; hedges and other conditionals are, therefore, common ('Innocent by reason of insanity'). A generalization of the operator *truth* for continuous variables is provided by the coefficient of determination, which measures the proportion of variance in the data field that is accounted for by the model. This is a mainstay in judging the accuracy of models. A complementary

index, proportion of the variance available from a model that is relevant to the data, provides a measure of specificity or parsimony; it is often crudely approximated by the number of free parameters in a model. Accuracy and parsimony measured in such fashions are the scientist's meta-models of truth and relevance.

The utility of models depends on their accuracy relative to others that are available. A model with mispredictions that still accounts for a significant and useful amount of variance in the data should not necessarily be spurned. It is foolish for a philosopher to deprive a laborer of a shovel because it is dull — unless he offers a better one in its place. The goal of science is not perfect models, because the only perfect renditions are the phenomena *sui generis*; the goal is better models.

Modeling languages are not models. Algebra and calculus and automata theory provide tools to craft the special purpose models, and it is those that are the cornerstones of scientific progress. The distinction between models and modeling languages is that of relevance; modeling languages can say too much that is not relevant to the data field under study. Models are proper subsets of all that may be said in their language. The more laconic a model, the more likely we can extrapolate its predictions to new situations without substantial tinkering. The most succinct models are called *elegant*.

There are many modeling tools available for scientists, a small set of which was sketched to give the flavor of their applications. The community of behavioral scientists has been conservative in exploiting modeling systems that might enrich their practice. Just as the microscope opened a new field of science, so also did tools such as the calculus and the probability calculus. This next century is a fertile one for the deployment of all available modeling tools on the problems of behavioral science; our questions are complex enough to support them. The trick will be to reformulate our questions in ways that make them amenable to using such tools. To do this, all we need is practice; and we all need practice.

Acknowledgement

I thank Tom Zentall for a close and helpful reading of the ms, and the National Institutes of Health for giving me time to write it under the aegis of Grant K05 MH01293.

Science as problem solving

1. Laudan, L. Progress and its problems. Berkeley, CA: University of California Press; 1977.
2. Hestenes D. Secrets of genius. *New Ideas Psych* 1990;8:231–246.
3. Hestenes D. Modeling games in the Newtonian world. *Am. J. Phys* 1992;60:732–748.
4. Pais, A. The original read *twentieth* century, and *physics* in place of *psychology*. New York: Oxford University Press; 1991. Neils Bohr 's Times.
5. Stove, D. Anything Goes: Origins of the Cult of Scientific Irrationalism. Paddington, NSW, Australia: Maclea Press; 1998.

Complementarity

6. French, AP.; Kennedy, PJ. Niels Bohr: A Centenary Volume. Harvard: Harvard University Press; 1985.

Aristotle on comprehension

7. Hocutt M. Aristotle 's four because. *Philos* 1974;49:385–399.

Probability theory

8. Estes, WK. Statistical Models in Behavioral Research. Mahwah, NJ: Lawrence Erlbaum Associates; 1991.
9. Wickens, TD. Models for Behavior: Stochastic Processes in Psychology. San Francisco: Freeman; 1982. 353 p.

Correlation in behavior

10. Baum WM. The correlation-based law of effect. *J. Exp. Anal. Behav* 1973;20:137–153. [PubMed: 16811687]
11. Galbicka G, Platt JR. Response-reinforcer contingency and spatially defined operants: testing an invariance property of ϕ . *J. Exp. Anal. Behav* 1989;51:145–162. [PubMed: 2921584]

Bayes and priors on behavior

12. Green RF. Bayesian birds: a simple example of Oaten 's stochastic model of optimal foraging. *Theor. Pop. Biol* 1980;18:244–256.
13. Jaynes, ET. Bayesian methods: an introductory tutorial. In: Justice, JH., editor. Maximum entropy and Bayesian methods in applied statistics. Cambridge, UK: Cambridge University Press; 1986.
14. Keren, G.; Lewis, C., editors. A Handbook for Data Analysis in the Behavioral Sciences. Mahwah, NJ: Lawrence Erlbaum Associates; 1993.
15. Killeen PR, Palombo G-M, Gottlob LR, Beam JJ. Bayesian analysis of foraging by pigeons (*Columba livia*). *J. Exp. Psych.: Anim. Behav. Process* 1996;22:480–496.
16. Myung IJ, Pitt MA. Applying Occam's razor in modeling cognition: a Bayesian approach. *Psych. Bull. Rev* 1997;4:79–95.
17. Shettleworth SJ. Reinforcement and the organization of behavior in golden hamsters: hunger, environment, and food reinforcement. *J. Exp. Psych: Anim. Behav. Process* 1975;1:56–87.
18. Skilling, J. The axioms of maximum entropy. In: Erickson, GJ.; Smith, CR., editors. Maximum-Entropy and Bayesian Methods in Science and Engineering. Dordrecht: Kluwer; 1988. p. 173-187.
19. Skilling, J. Maximum Entropy and Bayesian Methods. Dordrecht, The Netherlands: Kluwer; 1989.
20. Staddon, JER. Learning as inference. In: Bolles, RC.; Beecher, MD., editors. Evolution and Learning. Hillsdale, NJ: Lawrence Erlbaum; 1988. p. 59-77.

Induction/retroduction

21. Cialdini, RB. Full-cycle social psychological research. In: Beckman, L., editor. Applied Social Psychology Annual. Vol. vol. 1. Beverly Hills, CA: Sage; 1980.
22. Rescher, N. Peirce 's Philosophy of Science. Notre Dame, IN: University of Notre Dame Press; 1978.

Decision theory

23. Alsop, B. Behavioral models of signal detection and detection models of choice. In: Commons, ML.; Nevin, JA.; Davison, MC., editors. Signal Detection: Mechanisms, Models, and Applications. Hillsdale NJ: Erlbaum; 1991. p. 39-55.
24. Davison M, Nevin JA. Stimuli, reinforcers, and behavior: an integration. *J. Exp. Anal. Behav* 1998;71:439–482. [PubMed: 16812904]
25. Levitt, H. Decision theory, signal detection theory, and psychophysics. In: David, EE., Jr; Denes, PB., editors. Human Communication: A Unified View. NY: McGraw-Hill; 1972. p. 114-174.
26. Macmillan, NA. Signal detection theory as data analysis method and psychological decision model. In: Keren, G.; Lewis, C., editors. A Handbook for Data Analysis in the Behavioral Sciences: Methodological Issues. Lawrence Erlbaum Associates: Mahwah, NJ; 1993. p. 21-57.
27. Macmillan, NA.; Creelman, CD. Detection Theory: A User 's Guide. Cambridge, England: Cambridge University Press; 1991. 407 p.

28. Marston M. Analysis of cognitive function in animals, the value of SDT. *Cog. Brain Res* 1996;3:269–277.
29. McNamara JM, Houston AI. The application of statistical decision theory to animal behavior. *J. Theor. Biol* 1980;85:673–690. [PubMed: 7442286]
30. McNicol, D. A primer of signal detection theory. London: George Allen & Unwin, Ltd.; 1972. 242 p.
31. Swets JA, Dawes RM, Monahan J. Psychological science can improve diagnostic decisions. *Psych. Sci. Pub. Interest* 2000;1:1–26.
32. Wixted JT. A signal detection analysis of memory for non-occurrence in pigeons. *J. Exp. Psych.: Anim. Behav. Process* 1993;19:400–411.

Game theory

33. Bonta BD. Cooperation and competition in peaceful societies. *Psych. Bull* 1997;121:299–320.
34. Clements KC, Stephens DW. Testing models of nonkin cooperation — mutualism and the prisoners-dilemma. *Anim. Behav* 1995;50:527–535.
35. Crowle P, Sargent R. Whence tit-for-tat? *Evol. Ecol* 1996;10:499–516.
36. Green L, Price PC, Hamburger ME. Prisoners-dilemma and the pigeon: control by immediate consequences. *J. Exp. Anal. Behav* 1995;64:1–17. [PubMed: 16812760]
37. Mérö, L. Moral calculations: game theory, logic, and human frailty. Copernicus/Springer: New York; 1998.
38. Nowak MA, Ma R, Sigmund K. The Arithmetics of Mutual Help. *Sci. Am* 1995;272:76–81. [PubMed: 7761818]
39. Stephens DW, Anderson JP, Benson KE. On the spurious occurrence of tit for tat in pairs of predator-approaching fish. *Anim. Behav* 1997;53:113–131.

Automata

40. Hopkins, D.; Moss, B. Automata. North Holland, New York: 1976.
41. Minsk, M. Computation: Finite and Infintite Machines. Englewood-Cliffs, NJ: Prentice-Hall; 1967.

Functional/conjoint measurement

42. Anderson N. Algebraic rules in psychological measurement. *Am. Sci* 1979;67:555–563.
43. Shepard, RN. Approximation to uniform gradients of generalization by monotone transformations of scale. In: Mostofsky, DI., editor. *Stimulus Generalization*. Stanford, CA: Stanford University Press; 1965. p. 94-110.

General

44. Boas, ML. Mathematical Methods in the Physical Sciences. New York: Wiley; 1966. 778 p.
45. Maor, E. The Story of a Number. Princeton, NJ: Princeton University Press; 1994.
46. Mesterton-Gibbons, M. A Concrete Approach to Mathematical Modeling. New York: Addison-Wesley; 1989.
47. Shull, RL. Mathematical description of operant behavior: an introduction. In: Iversen, IH.; Lattal, KA., editors. *Experimental Analysis of Behavior*. Vol. vol. 2. New York: Elsevier; 1991. p. 243-282.
48. Timberlake W. Behavior systems and reinforcement: an integrative approach. *J. Exp. Anal. Behav* 1993;60:105–128. [PubMed: 8354963]

S = Stimulus
M = Movement
R = Reinforcer

Before Conditioning

After Conditioning

Fig. 1.

The process of conditioning increases the frequency of target movements (*M*) that are emitted in the presence of a discriminative stimulus (*S*). The correlation of these with reinforcement (*R*) is determined by the contingencies of reinforcement.

Fig. 2.
Different arrangements of sets of stimuli, responses and reinforcers correspond to traditional conditioning paradigms.

Fig. 3.
Probability as a dart game.

Fig. 4.

A model of response strength: an exponentially decaying trace with rate constant $k = 1/8$ averaged over CS durations given by the abscissae (symbols). The curve is the hyperbolic $y = 1/(1 + 0.08t)$.

Fig. 5.
Distributions of effects arising from presentation of stimuli.

Fig. 6.
The multidimensional signal-detection/optimization problem faced by real organisms: how to access a particular reinforcer.

Table 1

Conditions associated with correlations among parts of an act

Correlation	+	0	–
<i>S–R</i>	Classical conditioning	Rescorla control	Inhibitory conditioning
<i>M–R</i>	Positive reinforcement	Superstitious conditioning	Negative punishment
<i>S–S</i>	Compound	Noise	Contrast
<i>M–M</i>	Operant/style	Orthogonal/parallel	Competition: Ford effect

Table 2
A hypothetical set of preferences based on simple non-interacting functions on amount and delay of reinforcement

Delay: amount	1	2	3	4	5
1	1.00	0.67	0.50	0.40	0.33
2	1.41	0.94	0.71	0.57	0.47
4	2.00	1.33	1.00	0.80	0.67
6	2.45	1.63	1.22	0.98	0.82
8	2.83	1.89	1.41	1.13	0.94