Mexican Journal of Behavior Analysis
Special Issue on Operant Contingency
Vol. 21, 1995, pp.27-46

Direct Analysis of Contingencies using Working Models

Análisis directo de contingencias utilizando modelos de trabajo

Steven M. Kemp and David A. Eckerman[1]
University of North Carolina at Chapel Hill

Abstract

It is useful to "model" both contingencies and behavior. A reinforcement schedule arranged in the laboratory is a working model of a real-world contingency in the sense that a model of an airplane is a model. It simplifies a natural contingency, but it really will "fly." Thus, a reinforcement schedule is both a model of a contingency and it also IS a contingency, albeit a synthetic one. The authors argue for the use of working models of organisms, computer programs that take in models of stimuli as input and produce models of responses as output, all on a real-time basis. Computer simulations of working models called *in situ* simulations, are more complex than molecular simulations such as Shimp's Associative Learner in that spatio-temporal patterning, not just tempral patterning, is modeled. Working models of organisms, it is claimed, will allow for testing of behavioral theories of many types without the difficulties attendant upon theorizing that concerned Skinner (1950). Advantages to using working models include: (1) Analysis of behavioral theories at levels of observation more microscopic than those readily obtainable elsewhere. (2) Easy integration with neural network models. (3) The possibility of statistical testing in conjunction with the experimental analysis of behavior.

Key words: direct analysis of contingencies, working models, real time in situ simulations.

Resumen

Elaborar modelos tanto de contingencias de reforzamiento como de la conducta, puede ser de utilidad. Un programa de reforzamiento prescrito en condiciones de laboratorio, es un modelo de trabajo de una contingencia de reforzamiento que se

1 Request of reprints and other inquiries should be directed to Steven M. Kemp, Department of Psychology, Davie Hall CB#3270, University of North Carolina, Chapel Hill, NC27599-3270.

presenta en el mundo real en la misma forma en que un avión a escala es considerado un modelo del avión real. El modelo de trabajo permite simplificar la contingencia de reforzamiento que se presenta en el mundo real, además, a diferencia del avión a escala, el modelo sí puede "volar". Un programa de reforzamiento es un modelo de una contingencia y ES, aunque de manera artificial, una contingencia. En este trabajo, los autores presentan argumentos a favor de utilizar programas de computadora para elaborar modelos de trabajo del comportamiento de los organismos. Los programas incorporan modelos de estímulos a manera de "input" y generan modelos de respuestas a manera de "output", todo el proceso se lleva a cabo sobre una base de tiempo real. La simulación de modelos de trabajo, llamadas simulaciones *in situ*, son más complejas que simulaciones moleculares como las del modelo de aprendizaje asociativo de Shimp. La diferencia estriba en que en los primeros se modela no solamente el patrón temporal sino también el patrón espacio-temporal. Se argumenta que modelos de trabajo de los organismos permitirán evaluar diversos tipos de teorías conductuales evitando los escollos teóricos que preocupaban a Skinner (1950). Algunas de las ventajas de utilizar modelos de trabajo son: (1) Permiten evaluar teorías conductuales a un mayor nivel de detalle de lo que permiten otros métodos. (2) Permiten integrar los modelos de redes neuronales. (3) Permiten la utilización de pruebas estadísticas en conjunción con el análisis experimental de la conducta.

Palabras clave: modelos de trabajo, análisis directo.

Contingencies are the environmental relations that Skinner (1969) claimed are involved with learning. Contingencies thus can provide an answer to what Sidman (1960, p. 7-10) would call a what-question. That is when we ask: *What* features of the environment are involved with learning? The radical behaviorist answers: *contingencies*.

Environments are highly complex with many features that might promote learning. We will emphasized three characteristics of contingencies as important to their putative involvement with learning. First, contingencies are aspects of the environment. They are real, observable, measurable, external entities whose relations to behavior can be tested. Next, contingencies are describable by rules. (Call them contingency rules or c-rules for short.) Once contingencies are observed, *c-rules* describing those contingencies can be used to predict behavior. Using c-rules in this way is called *theorizing*. Finally, contingencies and c-rules share a common structure, an if...then... structure. As we will argue below, this if...then... structure is intimately related to Skinner's view of the nature of learning.

The focus of the present paper is on the use of contingencies in theorizing. Skinner (1950) claimed that the answer to the question of what affects learning could be worked out to any degree of detail required without using what he described as "theories of learning." Skinner believed some types of theories to be dangerous. We will have more to say about these dangerous sorts of learning theory later on. For now, it is important to note that several questions can be posed about Skinner's claim.

First, are there questions about learning that do require theories in order to find answers? Second, has behavioral analysis evolved in the last forty-plus years to the point where questions demanding theories need to be asked? Third, contemporary behavioral theories of learning are typically categorized as being either "molar" or "molecular". Are contemporary behavioral theories of learning, either molar or molecular, the kind of theory that concerned Skinner? Fourth, can the kinds of questions asked by molar and molecular theorists be answered using a kind of theory that is not subject to Skinner's (1950) criticism? Fifth, what would be the advantages of this new kind of theory? In particular, what new sorts of question could be answered that cannot be answered otherwise?

Questions about Contingencies

Skinner (1950) argued persuasively that certain sorts of questions about the effect of contingencies on behavior could be answered effectively without the use of "theory" in this dyslogistic sense of that term. Following on Skinner's point, Sidman (1960) distinguished between what-questions and why-questions. He pointed out that the sort of questions that are answerable without the kind of theory considered to be dangerous are what-questions. The first issue for the present paper is to identify the sorts of questions that might require theories (dangerous or not) in order to be answered. Sidman's why-questions provide a starting point. The behaviorist philosopher, Ryle (1949), distinguishes two kinds of why-questions, which we may call why-questions and how-questions.

To use Ryle's example, consider a glass that broke when hit by a stone. What might be said about this event? First, what happened? The glass broke after being hit by the stone. Why did the glass break? It was brittle. How did the glass break? The stone hit it.

What-questions ask about empirical relations between events. For events that occur over and over again, answers to what questions provide important information. For example, what proportion of what types of objects hit by what sorts of projectiles broke immediately thereafter? Do more glasses break before or after being hit by rocks? Etc.

Why-questions[2] ask about general propensities or properties. Brittleness is the classic example of a disposition used as a hypothetical construct. Why-questions demand either hypothetical constructs or their less troublesome cousins, intervening variables (MacCorquodale & Meehl, 1948). Traditional stimulus-response learning theorists, often preferred intervening variables to hypothetical constructs because intervening variables lack "surplus meaning". An example of an intervening variable might be *fragility*. Saying that the glass is fragile is supposedly a better answer to the why-question because "fragile", unlike "brittle", does not suggest hardness and inflexibility, merely break-ability. In either case, however, answering a why-question introduces a dispositional variable.

How-questions ask about mediational processes, that is, intermediate steps connecting the antecedent event and the subsequent event. While such steps intervene, logically or temporally, they need not be intervening variables in MacCorquodale & Meehl's sense of the word. The instantaneous impact of the stone and the glass may be only a small part of the sequence of events that occurred when the stone broke the glass, but it is as real as the flight of the stone, the spreading fractures, and the flying bits of glass. How-questions ask for a more detailed picture of events and raise additional questions (of all three types).

How-questions lead to additional predictions, allowing for more thorough testing. While why-questions multiply explanatory entities, how-questions challenge those entities by addressing their logical entailments. A detailed description of impacts at a microscopic level, for example, may force us to conclude that "brittleness" must mean the possession of a certain sort of crystalline structure. We can now examine the bits of glass to see if they have this sort of structure. Our explanation would no longer stop with the general assertion (assumed true) that all glass is brittle.

Let us now examine the three different kinds of questions as they apply to contingencies:

Since environments are complex, an important initial what-question is: *What* aspects of the environment promote learning? Skinner's answer is contingencies. Additional what-questions involve determining what sorts of contingencies produce what sorts of behavioral patterns and lead to what sorts of learning? When confining themselves to what-questions, experimental analysts of behavior can and do study and explain enormous

2 We realize that in some way "why questions" might be considered the most abstract and therefore should be addressed as the third rather than the second type. As our primary focus is on "how-questions", however, we have placed these last.

amounts of data and come to many important conclusions about behavior and learning. All of this proceeds without recourse to theory of the sort that Skinner deemed problematic and dangerous.

The second sort of question we might have about contingencies are why-questions. Why is it that contingencies, rather than some other feature of the environment, promote learning? In other words, why does describing the environment in terms of if...then...relations make for the most effective prediction of behavior? Skinner (1981) has answered this questions also, though in a more speculative fashion. Organisms learn because environments select certain behavior by providing consequences. Contingency-rules (c-rules) are attempts to describe the environment so as to specify what consequences will be delivered contingent upon what behavior. Thus, contingency-rules tend to pick out the right aspects of the environment to explain learning.

As noted above, what-questions can be answered using the experimental analysis of behavior without problematic theory. On the other hand, why-questions can be, and have been, answered by speculative involve theory of the sort that Skinner described as problematic. This brings to how-questions about contingencies. For example, how do contingencies produce learning? These how-questions are treacherous for behaviorists, because exploring the details of how consequences alter the organism so as to change subsequent behavior involves things going on inside the organism. Just as asking about how a stone impacts a glass leads to questions about the arrangement of molecules inside the glass, so asking about how the selection process works to alter a behavior repertoire leads to questions about biological processes (Think of Skinner's "natural lines of fracture").

Attempting to answer how-questions leads the psychologist toward theory. As the science of behavior has progressed and some what- and why-questions have been answered, how-questions have become more pressing. Each suggested answer to a why-question adds at least one new entity to psychological explanations. As noted above, how-questions are an excellent way of limiting the multiplication of entities. In this age of cognitive science, whose theories are notorious for their many intervening variables, answering how-questions is an important issue for the behaviorist. How-questions are needed to challenge the existence of these postulated entities. Methods for answering how-questions are discussed below.

Answering the How-Questions

Perhaps neurophysiologists will answer the how-questions. But it is unclear how neurophysiologists, unfamiliar with the study of behavior, can know what questions need to be answered. The level of detail of neurophysiological study is orders of magnitude more microscopic than that of behavioral analysis. Who will build the bridge between these levels?

Perhaps an analysis of brittleness will offer an example. Knowing the details of how molecules of glass are arranged is not enough to account for brittleness. First, the fact that the glass breaks when struck by a stone is discovered by macroscopic experiment. Further experimentation determines that many different sort of impacts shatter the glass in many different fashions. These macroscopic investigations are needed to help determine which microscopic details need to be investigated. It is equally important to note that microscopic investigations will help guide which sorts of impacts should be tested. For example, an understanding of the microscopic structure of the glass molecules might suggest whether the hardness of the stone or the momentum of the impact would be more critical to predicting when the glass would break. Cooperation and collaboration between macroscopic and microscopic investigations seems crucial when answering how-questions. Behavior analysts must be ready and able to contribute to a vigorous collaboration with neurophysiologists.

What sort of collaboration with neurophysiologists will help? And, will that collaboration involve theory? If so, what sort of theory would be useful? Cognitivists (and others) argue that suspicions about the perils of theory are old-fashioned, and that enough is now known about what organisms do and why they do it to allow us to safely theorize about how they learn. Baars (1986) describes the cognitive approach as encouraging psychologists to theorize freely. According to Baars, a cognitivist's theorizing should be constrained by two practices: behavioristic experimentation and computational theorizing.

Cognitivists adopted the experimental methods of behaviorists, particularly stimulus-response psychologists such as Woodworth (1938). Early cognitivists used mediational concepts that were closer to being intervening variables than hypothetical constructs. Terms like "memory", "semantic features", and "selective attention", were operationally defined as narrow in application. Only gradually, as experiments confirmed (or, more accurately, failed to disconfirm) their modest hypotheses, did cognitive psychologist begin to introduce broader terms such as "mental representation", "meaning", and even "consciousness" (Baars, 1986, ch. 4).

Computational theorizing (using the digital computer as a metaphor for the mind and/or brain) was adopted by cognitive psychologists for at least two reasons. As Neisser (1967) put it: "Although a (computer) program is nothing but a flow of symbols, it has reality enough to control the operation of very tangible machinery that executes very physical operations". On the one hand, computational mathematics is closely allied with the sort of logic philosophers of the time used to describe hypotheses and dispositions. Symbols were thought to be the stuff of which explanations were made. On the other hand, if a mediational entity could be, in principle, built into a computer program, it could be expected to be able to generate actual physical behavior. Computational constructs were simultaneously hypothetical and physical. This made them reassuring.

Behavior analysts, of course, have not been sanguine about the free theorizing of cognitivists. First, the experimental methods differ. The experimental methods of stimulus-response behaviorism may be empirical, but they are not the methods of behavior analysis. In general, choice of experimental method constrains theorizing in a number of ways. Most obviously, theoretical prediction (and control) involve the prediction of dependent variables from independent variables. The most prominent feature of an experimental method is the choice of variable. The independent and dependent variables of stimulus-response psychology (and cognitive psychology) are stimulus and response, respectively. As Lee (1988, p. 64) points out, the independent and dependent variables of behavior analysis are contingencies and response rates, respectively. In short, cognitive theorizing is constrained to answer questions *other* than those posed by the experimental analysis of behavior.

Second, computational theorizing is no longer as attractive as it once was, either to behavior analysts or to other critics. Researchers far afield from behavior analysis have raised questions about the value of computerized symbol systems both as hypotheses and as physical mechanisms for producing behavior. Phylosophers no longer analyze scientific hypotheses in terms of the logic of dispositions (Salmon, 1989). Further, nearly forty years of research in artificial intelligence have demonstrated clearly that many computational models of mind do not, in fact, produce effective behavior when they are actually programmed into a computer (Dreyfus, 1979).

While objectors to theory are in the majority within behavior analysis, some in that field have taken a middle road and advocated the development of behavioral theory. One way of looking at behavioral theory is as an attempt to begin to build the bridge to neurophysiology. In order to begin to answer questions of how contingencies affect behavior and produce learning, behavioral theorists, both molar and molecular, have attempted to provide answers to intermediate what- and why-questions about contingencies. They ask what aspects of contingencies promote learning and why those particular aspects are efficacious. Molarists look to aspects such as the correlation of behavior with long-term consequences while molecularists look to aspects such as contiguity between specific behavioral "elements" and specific delays of consequence. We will return to this matter shortly.

Let us summarize up to this point. Our answer to the first two questions we posed is: Yes, there are new behavioral questions that were not considered by Skinner that do call for theory. (We will consider a few of these as we discuss molar and molecular theories below). The goal of the present paper is to argue, however, that addressing these questions does not require us to violate Skinner's canons about dangerous theory. In agreement with behavioral theorists and in disagreement with the cognitivists, we contend that theories can be constructed to answer how-questions without the freedom to theorize described by Baars (1986). Our third question concerned contemporary molar and molecular theorists, we do not believe that the way to answer how-questions safely is to re-phase them as what- and why-questions, but to answer them directly using what we will call "working models."

Heteroscopic Theories

Exactly what sort of theory was Skinner worried about? Skinner (1950/1988) defined a problematic theory as "any explanation of an observed fact which appeals to events taking place somewhere else, at some other level of observation, described in different terms, and measured, if at all, in different dimensions." A great deal of exegesis has been done on this definition and we may find time to do more here. The key phase seems to be "at some other level of observation." Therefore, let us dub the sort of theories that concerned Skinner, *heteroscopic* theories. We can now rephrase our questions 4 and 5 as follows: Are contemporary molar and molecular theories heteroscopic? Can the questions asked by molar and molecular theorists about contingencies be answered with non-heteroscopic theories?

The great dilemma about contingencies is that the description of critical contingencies are a crucial part of our answers to what and why questions about learning, but they appear too "abstract" to help answer how-questions. Contingencies themselves, being diachronic environmental relations, are not immediately available to the organism's senses at any point during the actual course of learning. There is no burning bush spelling out the name "FI" when the worker receives her weekly paycheck. No one disputes that contingencies have their effects on organisms via the senses during the ongoing course of the organism's responding, yet it is difficult to identify what sensible aspects of the environment actually gain control over the organism's behavior as the mediators of these contingencies. The dilemma is to find immediate "here and now" events that transmit the impact of contingencies.

This dilemma can be phrased in terms of how-questions: How are sensible and immediately accessible aspects of the environment tied to these more global, perduring —and, hence, less accessible— contingencies?. Further, how are these sensible aspects of the environment involved in the ongoing selection of behavior? This second how-question is riskier, because, as we have mentioned above, it leads to questions about events going on within the skin.

Molarists claim that the efficacious aspects of contingencies are correlations (or similar measures). These correlations can be inferred or calculated from sensible features of the environment by first determining (via measurement and/or calculation) such things as rate of reinforcement. (This is not to say that molarists are speculating as a cognitivist might, as to how such calculations are performed. They are not). The difficulty is that any theory based on rates of events is obviouly heteroscopic because the behavioral events, occurring one at a time, can only be observed one at a time. This single-event level, the level of the behavior protocol, is their "level of observation." To determine a rate, multiple single-events must be observed (a macroscopic level). Thus, "rates" exist at a more macroscopic level of observation than does "behavior."

Molecularists (e.g., as reviewed by Staddon & Ettinger, 1989, p. 402-413) claim that the efficacious aspects of contingencies are contiguities occurring between specific narrowly defined classes of behavior and their consequences. These contiguities occur between sensible features of the environment. While it is less obvious that these molecular theories are also heteroscopic, they are. To determine delay, for example, at least two events are needed - one response event and one reinforcing event. A delay is the temporal gap between these two events. As with rates, delays are therefore measured at a more macroscopic level of observation than are individual behavioral events. The same argument leads to the conclusion that reinforcement probabilities are measured at a macroscopic level.

There is another way of looking at the question of whether theory is or is not heteroscopic. At one single instant in time, an event is either occurring or not. Events can be measured instantaneously. Rates and delays, however, being relations between multiple events, cannot be measured instantaneously. In some cases, instantaneous rates appear to be calculable, yet if the time period chosen as a basis is made longer or shorter, the measured rate may change. That is, rates are determined, in part, by the rules of measurement, not solely by the events observed. Therefore, a more macroscopic level of observation is involved.

By the same logic, delays are also macroscopic. For example, when there are multiple responses followed by a single reinforcer or when a single response is followed by multiple outcomes, multiple delays necessarily overlap. At any single instant in time, the specific delay or delays that are influential will differ depending upon which pair of events is chosen for measurement.

(Though just as heteroscopic, molecular theories do have a "virtue" that molar theories lack. In the physicist's terms, molecular theories are more "local" than molar theories. According to some philosophers of science, locality is a good thing for theories).

The key point in all of this discussion is that the original "level of observation" is the observation of events as they happen to the subject organism. We can call this a first-person status perspective on the events (Danziger, 1990). Twentieth Century psychological methods, behaviorist and cognitivist alike, are conceptualized from a third-person perspective. Behavior is described from the experimenter's point of view, not the subject's. As one might expect, a third-person perspective is used in constructing behaviorist theories as well as in describing behavioral data. When this use of the third-person perspective requires the theorist to "look beyond" what the subject actually experiences at any moment, the result is hetersocopic theory.

It is extremely important to note that molar or molecular theories status as heteroscopic (and therefore dangerous according to Skinner) has nothing to do with whether these theories are right or wrong. Molarists contend that the feature of contingencies that produce learning is the rate of reinforcement. Molecularists contend that the crucial feature is delay to reinforcement. Either (or neither or both) could be right. Instead of declaring them to be right or wrong, our question is methodological, namely: Is there a way to test these alternative contentions without engaging in heteroscopic theorizing?

Working Models

This paper proposes a novel technique for explaining behavior that operates solely at the level of observation of the original events to be explained. We call the technique a working model. Unlike cognitivist mathematical models, working models are models in the same way that a radio-controlled model airplane is a model of an airplane. A model airplane is both a model of an airplane and an airplane. It works. It flies.

The experimental analysis of behavior already uses working models. In the laboratory, reinforcement schedules are working models of contingencies. Reinforcement, schedules are models of natural contingencies. But they are also contingencies themselves. They work. They condition responding and promote learning. We propose here an extension of the working model approach: a working model of an organism. A working model of an organism would read in stimuli and reinforcers sequentially in real time and output responses in the same way. A working model is necessarily homoscopic because it requires actual contingencies as input and produces actual response sequences as output on a sequential real-time basis. As the reader might guess, such working models would be constructed using computer programs.

In evaluating how a working model is influenced by a reinforcement schedule, the reinforcement schedule could be programmed into the same computer as the model of the organism. Stimuli and reinforcers output from the scheduling program would be input directly into the model organism program. Responses output from the model organism program would be input directly into the scheduling program. (This design is a modification of Klopf & Morgan's, 1990, spatio-temporal simulation technique). We will call this an *in situ* simulation.

Certain behavior analytic simulations, particularly some simulations of Shimp's *Associative Learner* (e.g., Shimp, Childers, & Hightower, 1990), closely resemble *in situ* simulations. In order to clarify what makes a working model a distinctly novel approach for simulations, we will contrast the working model approach to that of Shimp et al. (1990).

In behavior analysis, the particular topography of responding is usually neither measured nor modeled. So long as the criterion for responding (sufficient to register on the cumulative recorder) is met., the only issue is the timing of the responses. Of course, the topography of responding and the structural details of the operant chamber may impact the timing of responses. For instance, the distance between the key and the feeder may impact the rate of key-pecks when the pigeon accesses the feeder. Likewise, the distance between two keys may alter the maximum rate at which a pigeon could exhibit a pattern of pecking in alternation.

Shimp et al. (1990) are concerned with what they call the "temporal patterning" of behavior. As the above discussion shows, temporal patterning cannot always be disentangled from spatial relations. In order to guarantee correct simulation of the temporal sequence of responding, *spatio-temporal patterning*, not just temporal patterning, must be considered. Those aspects of the topography of responding that may impact temporal patterning must be included in a simulation in order to make it a full-fledged working model. Aside from a few rare exceptions, such as Klopf & Morgan's (1990) spatio-temporal simulations, computer simulations, including simulations of Shimp's Associative Learner, neglect those critical features of response topography and are thus not full working models.

Of course, this does not mean that a *in situ* simulation of a working model must include all, or even most, of the topographic details of conditioning. A simplified computer model of the operant chamber might consist of a matrix with two positions, one above for the key and one below for the feeder. A simplified computer model of the pigeon's kinematics might consist of three topographies, move head up, move head down, and thrust head forward (to peck or to feed). the result of testing a computational theory in such a simulator would be to guarantee that the theory explains how the behavior of the organism, constrained by the structure of the operant chamber and the limitations of its own body, produces the response protocol measured by the cumulative recorder.

In sum, a reinforcement schedule is a synthetic contigency that successfully promotes learning in natural organisms. A working model of an organism would have to be a synthetic organism susceptible to conditioning by environmental contingencies. To qualify as a working model, the inputs and outputs must be constrained by the *in situ* model of the contingencies in a manner analogous to the way natural organisms are constrained by real contingencies, natural or synthetic. At the limit, a computer program that qualified as a working model could serve as the automated program inside a robot, allowing that robot to function in a natural environment. Such a working model, of course, needs to be far more complex than the simple working model of pigeon pecking proposed above.

How are Working Models Useful?

In order to illustrate the use of this novel contraption, let us suppose we wished to test the molarist claim that correlations effectively control behavior. Initially, the molarist must lay out the claim in general terms. This bit of verbal behavior would consist of a statement such as *the amount of time spent in an activity is proportional to the rate of reinforcement for that activity* (Baum, 1973). It is perfectly obvious that this bit of verbal behavior is theoretical in nature. Importantly, it is a type of theory that Skinner (1950) found acceptable and specifically *excluded* from his criticism: "Certain statements are also theories simply to the extent that they are not yet facts. A scientist may guess at the result of an experiment before the experiment is carried out". So long as the prediction can be stated in "the same terms in the same syntactic arrangement" as the result, Skinner has no complaint.

Clearly, the molarist claim above could be a description of the result of an experiment, or many experiments. Rates may be at a different level of observation than homoscopic measures such as cumulative records, but an experiment certainly may be designed in order to test hypotheses about rates. The independent and dependent variables of the experiment may be observable rates.

A problem arises when we consider experiments designed to investigate aspects of behavior in addition to mere rates. Consider the experiments of Ferster & Skinner (1957). For the most part these experiments were designed to answer what-questions about the detailed patterns of responding in behavioral protocols that varied systematically with the contingencies supplied. The level of observation of such experiments is the level of the individual behavioral protocol, measured in a cumulative record. (In general, the design of the experiment reflects the scientific question posed and the level of observation follows the design of the experiment). Just as determining a rate of responding requires the observation of more than one response, many different behavioral protocols can correspond to one and the same rate. It should be clear that predictions about rates made by molar theories are insufficient to predict the specific response protocol generated in answer to the experimenter's original question. (This problem is a reflection of a far more general problem. The predictions made by molar theories are appropiated to experiments specifically designed to test hypotheses about molar rates. Those predictions are inappropriate elsewhere. "Elsewhere" includes experiments designed to investigate characteristics of behavior more microscopic than rates. But "elsewhere" also includes behavior outside the laboratory. According to the above argument, an important difficulty with heteroscopic theories is that they cannot be extrapolated effectively. In science, however, the ultimate power of any theory rests with the ability to extrapolate its predictions to the real world.)

Suppose, however, that a real-time computer program incorporating the molarist claim was written and loaded into a computer. As each reinforcer is input, reinforcement rate would be recalculated "on the fly." At every opportunity to respond, the program would calculate whether or not to emit a response. This calculation would be performed moment by moment at the same level as the behavioral observations. The overall rate required by the molarist's theoretical claim would be arrived at by the ongoing calculations and re-calculations. Obviously, the algorithm for this computer program would involve a good deal more than would be found in the initial English-language statement of the claim. This additional programming is what would tie the molarist rule about proportionality to the requisite moment by moment responding that working models, like natural organisms, must produce in order to "fly." Unexpected results from the working model would provide additional, newly testable predictions of the molarist's theoretical claims.

A molecular theory could also be evaluated using a working model such as we are proposing. Let us next assume that a statement of a molecularist claim about the effectiveness of delay to reinforcement was used as the basis for a second computer program. A series of computer runs can then be performed in order to compare the two programs. Both programs can be run on the computer, each separately, each with a variety of reinforcement schedules. Suppose one program produces cumulative records more similar to the actual records produced by real organisms across a wide variety of reinforcement schedules. Such a result would appear to be a strong confirmation of the theoretical claim underlying the better performing computer program. In our testing scenario, the variety of reinforcement schedules are used to allow comparison of the behavioral protocols emitted by natural organisms exposed to those same reinforcement schedules. As pointed out above, such behavioral protocols are not predictable by heteroscopic theories. Translated in molar or molecular theories, the initial theoretical claims cannot be appropriately measured against detailed protocol-level data. However, the very same theoretical claims, instantiated in practical working computer models, are homoscopic, safe and able to be used for the prediction of protocol-level data. Events are input; events are output.

Each working model is an explanation for the behavior it simulates. Each explanation appeals indirectly to the broader theoretical claim of which it is a model. If no working model can be constructed from the broader theory that produces behavior comparable to natural behaviors, then the broader theoretical claim is disconfirmed or currently untestable. Over multiple simulation tests with multiple working models, a given theoretical claim can ultimately be accepted or rejected. Thus, broad, informal heteroscopic claims can be tested using a collection of narrow, computational, homoscopic working models.

Advantages to Working Models

There are a number of potential advantages to working models. Space limitations preclude mention of more than three. First, *in situ* simulations allow analysis of behavioral theories at levels of observation more microscopic than those readily obtainable in the usual operant settings. Second, working models are easily integrable with neural network models. Finally, working models make possible the use of statistical testing in conjunction with the experimental analysis of behavior. Mechner (1992) has championed a microscopic analysis of operant behavior. He discusses issues surrounding the spatio-temporal patterning of behavior in terms of "suboperants," which he defines as "physically necessary antecedents of the final effect" of the operant. (In the experimental analysis of behavior, the final effect of the operant is the change in the *operandum* that is recorded on the cumulative record, typically the closure of an electrical switch.) Examples would be the rat's placing a paw on the bar prior to pressing or the pigeon's head movements that position its head in front of the key.

The spatio-temporal patterning of the suboperants may vary widely. Typically, the experimenter does not specify that patterning. Reinforcement is determined solely on the basis of whether or not the switch is closed, irrespective of which of the possible sequences of suboperants precedes that closure. However, as we have noted above, the construction of the operant chamber and the details of the reinforcement schedule may constrain the spatio-temporal patterning and thereby indirectly impact the behavioral protocol produced. In principle, the behavioral theories should be able to predict the indirect effects of contingencies upon suboperants as well as the effects of constraints of suboperants upon operants. In practice, however, since suboperants are rarely indentified and recorded,behavioral theories are used solely for the prediction of operants, not suboperants.

Mechener (1992) recommends designing *operanda* so as to record important features of suboperants, thus exposing spatio-temporal patterning to the experimental analysis of behavior. Working models make possible concomitant theoretical analysis to support such an experimental-analysis. As discussed above, *in situ* simluations can be constructed to model contingencies to any degree of detail, up to and including the point where a virtual robot pigeon would be simulated inside a virtual operant chamber, all inside the computer. A working model must function in the *in situ* simulation, no matter what its level of detail. If, as both Mechner and we suggest, spatio-temporal patterning is an important issue in the prediction and control of operant behavior, then working models provide an additional tool for investigating that issue.

A second advantage to working models is that they are readily integrable with neural network models. As currently implemented, most neural networks are not working models by the present definition. Yet, like working models, most neural network models are computer programs, with a sequential series of inputs and outputs. Unlike working models, however, the timing of those inputs and outputs is neither proportional to real-time events in the laboratory, nor is it controlled according to any reinforcement schedule.

Further, in the design of neural network models, different kinds of inputs are *arbitrarily* assigned to different kinds of stimuli and different kinds of outputs are arbitrarily assigned to different kinds of responses. In working models, the inputs and outputs must be assigned so as to correspond to the corresponding assignment of outputs and inputs of the *in situ* simulation. In this way, the behavior analyst determines the how-questions to be answered. The how-questions determine the level of detail and intricacy with which the spatio-temporal patterning is modeled *in situ*. The level of detail of the simulation determines the assignment of the inputs and outputs of the working model.

Because neural networks models are computer models, however, the timing of inputs *could be* controlled by a reinforcement schedule and the inputs and outputs *could* be coded to correspond to outputs and inputs of a *in situ* simulation. Such a design would convert a neural network into a working model. An advantage of working models of neural networks is that working models of all kinds can be compared against the same data as was discussed above. At present, without working models, comparing neural networks to behavioral theory involves more speculative assumptions than does the construction of the theory.

Consider how the theoretical claims of a neural network modeler could be evaluated using working models. Suppose that a theorist, in this case, neither a molarist nor a molecularist, wishes to explore how-questions about contingencies and their effect on learning. This theorist does not believe in the effectiveness of specific features of contingencies said by others to be responsible for conditioning. Rather, she believes that all of the sensible effects of contingencies collaboratively establish stimulus control over behavior. This collaborative control is integrated by neurophysiological changes analogous to computations. The theorist authors a neural network model as a working model and can now test it head-to head against the molar and molecular models using the same sets of schedules.

A third advantage to working models is testability. Traditional learning theories, including behavioral theories, either molar or molecular, are attempts to answer why-questions. Dispositional variables, either intervening variables or hypothetical constructs (MacCorquodale y Meehl, 1948), are thus always introduced. The virtue of modeling using dispositional variables was testability. If the glass broke because it was brittle, then hitting other brittle, non-glass objects will test the new theory that all brittle objects hit with stones break. The vice of modeling using dispositional variables was simplicity. A great number of facts besides the brittleness of the glass might have contributed to its breaking on that particular occasion. (The size and speed of the stone, the absence of something to cushion the blow, etc.)

Answering why-questions must always begin with simple models for technical reasons. The more facts about the situation that are described using additional dispositional variables, the harder it is to establish which combination of variables are actually responsible for the effects observed. If we allow the theorist to add variables to the model, it becomes harder to design an experiment that will falsify the theory. Each time we find a brittle object that does not break when hit by a stone, the proponent of the theory just adds a new variable. The stone wasn't big enough. The stone was big enough but wasn't moving fast enough, etc., etc., *ad infinitum*. A theory that cannot be falsified does not explain.

Skinner (1950) believed that the cost of simplification was not worth the benefits of testability. This is why he doubted the value of theories of learning that used dispositional variables. Unlike why-questions,however, how-questions can be answered without dispositional variables. The mathematical "variables" inside a computer program are not necessarily dispositional variables.

Dispositional variables, such as brittle and fragile, predict individual experimental outcomes: one variable, one outcome. Statisticians measure dispositional variables in terms of degrees of freedom. It may take many computer program variables to make one dispositional variable worth one degree of freedom. The difficulty is that there is no consensus amongst statisticians as to how to test theories instantiated as computer programs.

Because working models can generate data that are formally identical to experimental data, the issue of statistical analysis is sidestepped. Whatever statistical tecnique is appropriate for evaluating the experimental data is also appropriate for evaluating any simulated data. In the case of experimental analysis of behavior, where statistics are often not used, the simulated data would be as accessible to non-statistical analyses as are the experimental data. In short, working models allow different theories to be compared directly to one another with or without statistics. This is what we mean by "direct analysis."

Skinner's concern about over-simplification is also addressed. Neural network models and other computational models need not begin by postulating dispositional variables at all. They do not need to be simplified in order to be tested. Working models, together with *in situ* simulations, allow computational models, whatever their conceptual origins, to be tested directly and compared to one another.

Conclusion

Heteroscopic theories are problematic because they cannot be directly tested using behavioral experiments. Their apparent elegant simplicity bellies the complexities of the environmental contingencies whose effects they are supposed to explain. Conceptual analysis of heteroscopic theories may miss subtle aspects of environment-organism interactions with which those theories may not be able to deal. Working models on the other hand, allow behavior analysts to test each theoretical explanations against actual behavior. Speculative learning theories become testable when instantiated in a working model. These theories then can be tested far more thoroughly, and examined for implications outside the scope of tecniques currently available in the experimental-analysis of behavior. Given the possibility of real time computerized working models of behavior, (heteroscopic) theories of learning are truly unnecessary.

References

Baars, B. J. (1986). *The cognitive revolution in psychology*. New York: Guilford Press.

Baum, W. M. (1973). The correlation-based law of effect. *Journal of the Experimental Analysis of Behavior*, 20, 137-153.

Danziger, K. (1990). *Constructing the subject: Historical origins of psychological research*. Cambridge: Cambridge University Press.

Dreyfus, H. L. (1979). *What computers can't do: The limits of artificial intelligence* (revised ed.). New York: Harper and Row.

Ferster, C. B. y Skinner, B. F. (1957). *Schedules of reinforcement*. New York: Appleton-Century-Crofts.

Klopf, A. H. y Morgan, J. S. (1990). The role of time in natural intelligence: Implications of classical and instrumental conditioning for neuronal and neural network modeling. In M. Gabriel and J. Moore (Eds.), *Learning and computational neuroscience: Foundations of adaptive networks*. Cambridge, MA: MIT Press.

Lee, V. L. (1988). *Beyond behaviorism*. Hillsdale, NJ: Lawrence Erlbaum Associates.

MacCorquodale, K. y Meehl, P. E. (1948). On a distinction between hypothetical constructs and intervening variables. *Psychological Review*, 55, 95-107.

Mechner, F. (1992). *The revealed operant: A way to study the characteristics of individual occurrences of operant responses*. Cambridge, MA: Cambridge Center for the Behavioral Studies.

Neisser, U. (1967). *Cognitive Psychology*. Englewood Cliffs, NJ: Prentice-Hall.

Ryle, G. (1949). *The concept of the mind*. New York: Harper and Row.

Salmon, M. H. (1989). *Explanations in the social sciences*. In P. Kitcher and W. C. Salmon (Eds.), *Minnesota studies in the philosophy of science: Vol. 13*. Scientific explanation. Minneapolis, MN: University of Minnesota Press.

Shimp, C. P. , Childers, L. J., y Hightower, F. A. (1990). Local patterns in human operant behavior and a behaving model to interrelate animal and human performances. *Journal of Experimental Psychology: Animal Behavior Processes*, 16, 200- 212.

Sidman, M. (1960). *Tactics of scientific research: Evaluating experimental data in psychology*. New York: Basic Books.

Skinner, B. F. (1950). Are theories of learning necessary? *Psychological review*, 54, 193-216.

Skinner, B. F. (1969). *Contingencies of Reinforcement: A theoretical analysis*. New York: Appleton-Century-Crofts.

Skinner, B. F. (1981). Selection by consequences. *Science, 213*, 501-504. (Reprinted with commentary, in *Behavioral and Brain Sciences*, 1984, 7, 477-481. Also, in A. C. Catania and S. Harnad (Eds.), The selection of behavior: The operant behaviorism of B. F. Skinner: Comments and consequences, 1988, 11-76.)

Staddon, J. E. R. and Ettinger, R. H. (1989). *Learning: An introduction to the principles of adaptive behavior*. New York: Harcourt-Brace-Jovanovich.

Woodworth, R. S. (1938). *Experimental psychology*. New York: Holt.