

Behavior Analysis, Relational Frame Theory, and the Challenge of Human Language and Cognition: A Reply to the Commentaries on *Relational Frame Theory: A Post-Skinnerian Account of Human Language and Cognition*

**Steven C. Hayes, University of Nevada, Reno, and
Dermot Barnes-Holmes and Bryan Roche,
National University of Ireland, Maynooth**

Answers to a series of commentaries are presented and the challenge Relational Frame Theory (RFT) presents to behavior analysis is explicated. RFT is a behavior analytic theory, based on extensive behavior analytic data, which appeals only to known principles to explain arbitrarily applicable relational responding. The claim that such responding is operant must be answerable within behavior analysis. RFT has too much empirical support for the field to avoid this challenge. If the answer is “yes,” behavior analysis seems destined to enter a new era.

Any disciplinary approach is composed of multiple levels of scientific activity: assumptions, methods, techniques, concepts, research questions, research answers, and the like. Approaches that are long lived, and particularly those associated with significant intellectual figures, also have substantial histories and traditions, complete with heroes and tales of former battles won. As traditions become more extended, new directions become harder to assimilate, especially if these new directions challenge previously dominant methods or questions, or if the views of historical figures are challenged.

Behavior analysis is one of the longer-lived traditions in psychology. Behavior analysis sought nothing less than a comprehensive account of the psychological activity of organisms, particularly and ultimately, the complex behavior of human beings, based on a functional, monistic, scientific approach. Behavior analysis adopted some unusual tactics in pursuing this goal, including the emphasis on rela-

tively simple organism–environment interactions studied largely in non-human organisms. Behavior analysis has been enormously successful, but that very success may make it more difficult to take new directions in areas that have not been as successfully addressed.

Relational Frame Theory seems to be a test case. RFT is making explicitly behavior analytic claims based on behavior analytically sensible research data. Nevertheless, it suggests significant changes in some of the methods, techniques, concepts, research questions, and research answers of the field. For example, if RFT is correct, we need many more basic human laboratories; we need new research methods; we need much more experimental research on language processes; and many of Skinner’s specific ideas (e.g., about the role of human emotion; about the nature and function of self-knowledge; about the need to understand private events in order to understand overt behavior; and so on) are challenged in whole or in part. This is a difficult possibility for behavior analysis to face head on, and it helps explain a sense of discomfort that is palpable throughout the commentaries. We understand the discomfort, and have felt it ourselves, but we do not think behavior analysis can avoid that discomfort without harm. The challenge Relational Frame Theory presents to behavior analysis as a whole must be faced and resolved, one way or the other.

Please address correspondence concerning this article to Steven C. Hayes, Department of Psychology / 296, University of Nevada, Reno, Reno, NV 89557-0062 (email: hayes@unr.nevada.edu) or to Dermot Barnes-Holmes, Department of Psychology, National University of Ireland, Maynooth, Maynooth, Co. Kildare, Ireland (email: Dermot.Barnes-Holmes@may.ie).

The Challenge of Relational Operants

It is precisely because RFT is a behavior analytic theory that the challenge is so great. RFT claims that verbal behavior is a particular kind of operant behavior: arbitrarily applicable relational responding. Considerable detail is devoted to delineating the kind of operant we are speaking about. RFT researchers have provided several lines of evidence that support this claim and year after year more data come in. To our knowledge, there has not yet been a sour note in the data that have been generated on this point so far.

It should not be easy for behavior analysts to dismiss a credible claim that a particular type of operant exists. Behavior analysis as a field must have concepts to decide if something can or cannot be an operant, and if a particular operant is conceptually possible, methods and techniques to determine empirically if it is in fact. If such is not the case, behavior analysis is not a field worthy of the name "operant psychology."

RFT is now 18 years old. It has spawned more basic human operant work than almost any theory put forward during that time. Sooner or later the field must reach a conclusion about the central claim of RFT. Arbitrarily applicable relational responding either is an operant or it is not. That conclusion does not have to be reached tomorrow, but it cannot be put off indefinitely. If critics believe that relational operants are conceptually impossible, they must say why. If they believe it is an empirical question and the question has not yet been answered, they should be able to say precisely why the existing data are insufficient and precisely what additional data are still needed.

A New Behavioral Principle?

One basis for avoiding this challenge is the suggestion that new principles are implicated by RFT, but closer examination shows that this cannot be used to delay a decision about the central claim of RFT. In several places in the current reviews it is said that RFT theorists appeal to new principles to explain verbal phenomena. For example, Malott says that we seem to "feel it more parsimonious to explain a result in terms of one new principle of behavior than two existing principles" adding "My concern about this new principle is that it seems to

be merely a molar *description* of that which we are trying to understand rather than an *explanation* of what we are trying to understand." (p. 17)

If RFT used a new principle to explain relational operants, it might be appropriate for behavior analysts to resist these claims until "more conservative" alternatives based on existing principles could receive more attention. That is not the case, however.

The principles and procedures used to explain relational operants in RFT are identical to those used to explain any operant: contacted consistencies in contingencies across multiple exemplars. The *explanation* of arbitrarily applicable relational operants *appeals to no new principles at all*. It is only after assuming that the operant described by RFT exists that we then claim that a new principle exists. The concept of a new principle of behavior is not used to *explain the process* that gives rise to relational operants, rather it is used to *explain the implications* of such an operant. This distinction fundamentally changes the nature of the challenge RFT presents to behavior analysis as a whole. It may be scientifically conservative to avoid using new explanatory principles, but it is not scientifically conservative to use old labels for phenomena that fail to fit the definitional limits for those labels.

What is new about RFT? We point out that "the RFT approach invokes a purely functional concept of an operant" (p. 146) rather than one that relies on topography, but this is not anything new. We use terms like "generalized operants" or "overarching operant class" to emphasize the functional nature of the operant but we are also careful to note that "The concept of a response class with an infinite range of topographies is a defining property of operant behavior, and has been from the very beginning (e.g., see Skinner, 1938, p. 33-41)." (p. 147), and that "No new type of operant is supposed by these terms—the qualifiers are merely to avoid confusion." (p. 147). Several of the commentators note that we appeal to a special kind of operant, which is true, but only in the sense that we define the relational operant. That operant is not special because it is "overarching"—as Spradlin correctly notes, this qualifier is just a way to emphasize Skinner's functional definition of the operant—rather is it special because of its particular function altering features. That, however, is a re-

sult—the process that is supposed to establish those features is not special in any way.

For behavior analysts, claiming that arbitrarily applicable relational responding is operant behavior *is* an explanation. The explanation could be empirically wrong, of course, but to criticize such a claim on the grounds that it is “merely a description” (Malott) should be rejected by behavior analysts on metatheoretical grounds. Behavior analysts describe contingencies and their results. This is explanation for behavior analysts, and claiming that something is an operant is a testable claim about process.

Salzinger approaches this same type of criticism in the following quote:

What is reinforced then is the relational response. Sidman’s (2000) analogy is at least equally attractive, viz. p. 144). “An equivalence relation can be thought of as a bag that contains ordered pairs of all events that the contingency specifies; the bag can be shaken and the elements mixed without regard to any spatial or temporal relations among them.” That leaves unspecified what the mechanism is that seems to transcend spatial and temporal relations. (p. 8)

It is not clear why Salzinger believes that Sidman’s “bag theory” is equally attractive. Sidman’s bag theory could apply only to equivalence, not to all of human language, and it is at best a behavioral result, not a process. He is right to note that Sidman’s idea leaves unspecified what transcends spatial and temporal relations, but RFT is clear on this point. RFT claims that what transcends spatial and temporal relations is the relational operant itself. This is a process claim and it points directly to what is new in an RFT account.

RFT claims that:

the instrumental behavior of relational framing alters the functions of behavioral processes ... If Relational Frame Theory is correct, the alteration of these behavioral processes was itself a learned process. Said another way, relational framing is operant behavior that affects the process of operant learning itself. We know of no term for such an effect. (p. 45)

This quote shows clearly that we use a “new principle” only as a label for a new effect, not for the processes that created that effect. This

claim is both empirical and conceptual.

In the book an example is offered:

A discriminative stimulus is a stimulus in the presence of which there has been a greater probability of reinforcement for a given behavior than in its absence. Suppose a child is rewarded for waving when the word “dog” is heard. The word “dog” is a discriminative stimulus. Suppose, however, that the child is now taught to say “dog” given the word D-O-G, and to point at actual dogs given D-O-G. Suppose that as a result of this training the child now waves upon seeing a dog. Such an outcome has repeatedly been seen in the literature (e.g., Hayes et al., 1987). The dog cannot be a discriminative stimulus because the child has no history of greater reinforcement for waving in the presence of dogs than in the absence of dogs. The effects cannot be stimulus generalization because there are no formal properties that are shared between the word and actual dogs. The effect cannot be due to classical conditioning because it would require an appeal to backward conditioning. The effect cannot be due to compounding because “dog” and dogs have not even occurred together.

Relational Frame Theory suggests that the performance is due to a learned process that transformed these discriminative functions. In normal discriminative control, the stimulus function is learned, but not the process itself. In contrast, the derived performance is discriminative-like, but it is not discriminative. These discriminative-like effects seem to depend on a learned process of altering behavioral processes, and that is something that is not covered by an existing technical term. Despite the conservatism of an RFT approach, therefore, ... in our analysis, verbal events (and relational frames) instantiate a newly identified behavioral process (Hayes and Hayes, 1992). (p. 46)

If operants exist that are focused on and alter how behavioral processes themselves function, then we need a new term because no existing term will do. This is a simple conceptual claim, and we see no logical flaw in it. But the “if ... then” nature of the claim clearly separates the issue of implications (the “then” part of the statement) from the issue of process (the “if” part of the statement).

Let’s dream up a new operant to show the point in another area. Suppose it were possible to combine operant and classical conditioning in the following way. Imagine that an animal

is presented with several CS–UCS pairings. Subsequently, the UCS is presented alone and reinforcement is delivered contingent upon whether the animal responds in a way that was originally evoked by the CS. For example, in our imaginary experiment, reinforcers would be delivered to Pavlov’s dogs if they showed some minimal auditory response when presented with food powder. Over scores of trials, varying in the magnitude of the response required and varying across different specific kinds of CS and UCS events, imagine that this operant contingency trained our subjects in some contexts to show strong CS related responses to the UCS following any new set of CS–UCS pairings. Now imagine that those animals who have been through our training are, unbeknownst to an experimenter, mixed in with naive subjects in a classical conditioning study examining backward conditioning. Two very different response patterns result. Some subjects (the naive ones) show little or no CS-related responses to the UCS. Some subjects (our trained subjects) show strong and reliable responses of this kind. If no one knew of the different histories in our two sets of subjects it might be reasonable to conclude that “sometimes classical conditioning leads to strong CS-related responses to the UCS but we do not know why it sometimes does and sometimes does not.” It would be scientifically irresponsible to make this same claim if the different histories were known, however.

Suppose that someone suspected the true source of the effect, and conducted a series of experiments that showed that operant histories could indeed have an effect of this kind. It would be perfectly proper and scientifically conservative to ask for evidence on this point. But suppose that the scientist pointed out, very reasonably, that if operant contingencies could have this effect on classical conditioning we would need a new term to talk about it. Purely descriptive terms like “inverse CS–UCS relations” would not do because they merely described the result, not the source of the result. Classical conditioning terms such as “backward classical conditioning” would not do because they would imply that the process involved was classical conditioning *per se*, whereas the key process was operant and only the preparation was classical. If a new term were used, however, it would not be “conservative” for scientists to reject the idea of such an operant im-

pacting on classical conditioning in this way because the scientist is “making up new principles.” The account did no such thing. The “new principles” were needed to describe the previously unknown impact of *known* principles, thus making the issue more empirical than theoretical. It is certainly not a mere philosophical matter, as if the data on this new operant could be rejected on the basis of broad philosophical or terminological preferences.

This parable parallels rather precisely the current situation with regard to RFT. It helps explain why we believe that behavior analysis cannot walk away from this challenge unscathed. If relational operants of the sort RFT imagines exist then behavior analysis is in a new era. For that not to occur, relational operants must either be shown to be incoherent in principle, or not to exist empirically, or for the implications to be incoherent in principle, or for the implications not to exist empirically. Any empirical or conceptual resolution will be positive for behavior analysis if it is true to this own approach. If, however, the field simply ducks the empirical and conceptual challenge presented by a major theory within its ranks by (falsely, we believe) treating this challenge as if it is mere a matter of preference, terminology, or philosophy, the field will be risking its soul in the name of keeping its identity. RFT is simply too well established as a behavior analytic theory to walk away from in that way.

Relational Responding as an Active Process

Another putative reason for avoiding the challenge that Relational Frame Theory presents to the field of behavior analysis is the suggestion that the concept of arbitrarily applicable relational responding, *as an active process*, does not fit with traditional behavior-analytic thinking. McIlvane summarizes this argument as follows:

Another major departure from the abstraction analysis is that RFT views relational learning as behavior. One does not merely exhibit relational stimulus control. One “relates.” By contrast, the traditional behavioral analysis of abstraction has been the narrowing of stimulus control. Thus, relational responding in RFT seems to be an active rather than a passive process. This characteristic has set many traditionally trained behavior analysts to head scratch-

ing as they tried to fit the theory within a behavior analytic framework.” (p. 30).

The main difficulty here is that in behavior analysis we do not normally separate the discrimination from the response. When a rat presses a lever in the presence of a green light this is the discriminative response—we don’t say that the light was discriminated and then this caused the rat to press the lever.

Behavior analysts are suitably wary of using behaviors as causes of other behaviors. Focusing on behavior–behavior relations can easily work against more complete experimental and applied analyses in which current and historical contextual variables are identified and manipulated to achieve the scientific goals of prediction-and-influence (Hayes & Brownstein, 1986). In the context of arbitrarily applicable relational responding, however, active terminology highlights rather than obscures the historical and current contextual variables responsible for this type of relational control. This is a subtle but fundamentally important issue for RFT and warrants further elaboration.

Relational stimulus control in behavior analysis is typically thought of as a process of

abstraction [that] may develop when a set of otherwise physically different discriminative stimuli *have a physical property in common*. For example, having learned via discrimination training to respond differentially to a red flower, a red car, a red ball, and other red items, the learner may subsequently respond differentially to other red items without the need for explicit discrimination training. (McIlvane, p. 29, emphasis added)

In effect, the traditional definition of relational stimulus control in behavior analysis draws heavily upon the abstraction of common physical properties, and in one sense the organism that learns to abstract such properties is responding directly to them. Arbitrarily applicable relational responding, however, does not involve direct stimulus control by non-arbitrary physical properties; instead the relational response is arbitrarily applicable and under the control of a contextual cue or cues. In the book, we say it this way:

In order to abstract the behavior of relating, the organism must be exposed to training that

allows it to discriminate between the relevant features of the task (responding to one event in terms of another based on a contextual cue) and the irrelevant features (the actual physical properties of the related objects).” (p. 26)

The final clause, in parentheses, serves to highlight that the traditional definition of relational stimulus control and abstraction does not apply to RFT—in fact, the abstraction in RFT involves learning in some contexts *to ignore* the physical properties of the stimuli as the basis for a relational response.

An arbitrarily applicable relational response is not, therefore, a direct response to some common physical property across two or more stimuli. Rather, it involves the discrimination of a contextual cue that brings to bear a relevant history of non-arbitrary and arbitrary relational responding. In this sense, the cue causes an *active* relational response. McIlvane is correct when he states, “Relational frame theory might be briefly summarized as an extension of the traditional behavior analytic account of abstraction” (p. 29). However, this extension is substantive and RFT should not, therefore, be confused with relational stimulus control as traditionally defined. Interestingly, Osborne appears to agree that there has been some confusion in this regard; “Behavior analysts have probably used the phrase *relational responding* to denote the act of *relating* but didn’t realize at the same time that they were speaking of *relational framing*” (p. 23). From our perspective, therefore, the overarching and extended history of contextually controlled non-arbitrary and arbitrary relational responding that is required for relational framing to develop is so important that it seems wise to highlight this feature of RFT by referring to relational framing as an active or behavioral process.

It is worth noting that this way of talking could also be applied to traditional discrimination learning. The classic perception study by Held and Hine (1963), using new-born kittens, demonstrated that even seeing vertical lines can be conceptualized as operant behavior (see Catania, 1998). However, because the operant of seeing is established so early and so thoroughly in the behavioral history of the organism, it is easy to forget that the first part of any visually based discrimination response involves the operant of seeing the stimulus. On

balance, there are no good pragmatic reasons for remembering this fact (at least in behavior analysis), and thus we rarely say that the history of reinforcement for seeing stimuli leads to an active discriminative response. For arbitrarily applicable relational responding, however, the slide into traditional ways of thinking about relational stimulus control (i.e., in terms of common physical properties) is so steep that, at least for now, it seems wise to differentiate it in this way—as an active behavioral process involving an overarching operant history that in principle could stretch back for decades.

The Issue of Continuity

Another issue raised by some of the commentators that seems to have caused some consternation is the relationship between Relational Frame Theory and the continuity assumption. For example, McIlvane asks “Why do relational frame theorists not more directly consider arguments that behavior may be the product of phylogenetic contingencies” (p. 35); and Osborne states, “It is at this level, that of continuity or discontinuity between homo sapiens and all other animals, that RFT wishes to divide from Skinner, and this is potentially a huge divide” (p. 23); and Spradlin notes “Their position implies discontinuity between the behavior of humans and nonhumans” (p. 3). Implicit in McIlvane’s comment, and explicit in both Osborne’s and Spradlin’s, is the view that RFT is self-consciously opposed to the continuity assumption.

The role of this assumption in behavior analysis, and its relationship to RFT is a complex issue (Hayes, 1987a, 1987b; see Dymond, Roche, & Barnes-Holmes, in press, for a detailed discussion), but the RFT position was clearly stated in the book—the theory remains agnostic:

Arbitrarily applicable relational responding occurs readily, even with human infants and with difficulty or not at all with nonhumans. After thirty years of behavior analytic research on derived stimulus relations, that statement is still true. We do not need to take the stand that nonhumans will never show derived stimulus relations in order to begin to launch an extensive and coherent program of basic research into these processes in human beings. The findings will be no less useful and no less scien-

tific simply because they will not necessarily generalize across tips of evolutionary branches, any more than, say, the finding that operant conditioning does not apply to bacteria would limit the validity of such findings in birds and mammals. All that is needed is that behavioral researchers must not be so wedded to this strategic assumption that alternative strategies cannot be pursued.

Relational frame theory is oriented toward human language and cognition. Whether the richness and complexity of human and language cognition may yet be shown to be in the same functional class as behavior studied with other organisms is yet to be determined. Only empirical research, not assumptions, will resolve this issue. This is precisely the view adopted by RFT (p. 145).

By adopting an agnostic position with respect to the continuity assumption, RFT, contrary to Osborne’s claim, is consistent with Skinner’s early thinking on this issue. Within behavior analysis, we point out that

The traditional focus on nonhumans was based on the idea that the principles of behavior identified with such populations would be generally applicable to humans.... This form of the continuity assumption was a strategic assumption: it was a means to an end. It was not a categorical assumption—that is, one that is fundamental to the conceptual coherence of the field. This is why Skinner warned that “We can neither assert nor deny discontinuity between the human and the subhuman fields so long as we know so little about either” (p. 442, 1938), and that “It is possible that there are properties of human behavior which will require a different kind of treatment” (p. 442, 1938). (p. 145)

It should be emphasized that the “continuity assumption” we are speaking about is the strategic research assumption made by early behavior analysts: it is not the one that built in evolutionary theory:

The evolutionarily sensible form of continuity assumes that new contains old. There is no guarantee or assumption that old contains new. Biological evolution itself would be turned into nonsense by such an assumption. When we look across tips of evolutionary branches we are not looking back in time: we are always looking both back in time (to the point at which specific species differentiated) and forward in time to the present. Thus, discontinuity across

present day species would in no way contradict a biologically sensible form of the continuity assumption. (p. 145).

Adopting an agnostic position with regard to the continuity assumption does not imply that RFT underestimates the importance of phylogenetic contingencies. Relational frame theory, like any behavior analytic account, assumes that all learned behaviors have their ultimate origins in phylogenetic contingencies. Furthermore, RFT assumes that these contingencies may be partly responsible for derived relational responding:

If ... language and cognition will not yield entirely to principles derived from nonhuman research, it means simply that processes emerged in one evolutionary branch and not another. What was new could be extremely small, and yet produce huge differences in behavioral outcomes. Metaphorically, a person standing at the edge of a cliff may step forward an inch and fall hundreds of feet. The step was not large: only the outcome was large.

If this is the conclusion we are eventually driven to, what kind of processes could produce such an effect? The ability of a listener to derive a bidirectional relation readily from multiple exemplars would be a ready nominee. (p. 146)

RFT does *not* place phylogeny center stage in the theory, however. Focusing on phylogeny, and more specifically on whether or not nonhumans are capable of demonstrating the simplest forms of derived stimulus control, seems counterproductive to the analytic goal of the theory. McIlvane may well be correct in arguing that such a focus could lead to “debate” like that “paralleling similar discussions in psycholinguistics,” but RFT is chasing what we see as a far more important prize than mere discussion. RFT aims to provide a naturalistic, monistic, and purely functional-analytic account of human language and cognition that readily facilitates meaningful change within these two domains in both basic and applied research settings. This is something that psycholinguistics, and cognitive psychology more generally, has not achieved, and probably never will given their analytic goals and research strategies.

If RFT is to make a contribution that other psychological traditions have not, we need to stay focused on the analytic goal of prediction-

and-influence and avoid getting sidetracked into well-worn and long since barren furrows of scientific debate concerning the nature–nature controversy. How many animal studies will have to be conducted before negative evidence is taken as final? (cf. Lionello-De-Nolf & Urcuioli, 2002). And if evidence is finally obtained that proves unequivocally that a non-human can demonstrate the basic equivalence relation, what then? How important will this research be for mainstream psychology, or science more generally? As behavior analysts, we are at a crucial stage in our history—we need to choose our research agendas very carefully and not waste time on wild goose chases (see Malott, 1991).

Parsimony

Another issue that appeared to be of concern to some of the commentators was an apparent lack of parsimony in RFT. Malott, for example, suggests that his own chaining account of symmetry and transitivity is more parsimonious than RFT because the former is cast “in terms of the basic principles of behavior” (p. 17) and McIlvane remarked “One problem that I had with *RFT:ALC*’s discussions of transformation of functions was the general lack of an obvious relationship to the basic processes of behavior analysis” (p. 33).

One can only appeal to parsimony as a measure of theoretical elegance and viability when two relatively adequate explanations co-exist. RFT has made significant progress by using particular behavioral concepts applied to an operant conception of relational responding. If we are correct, relational operants require that we rethink and some of our units of analysis. The “natural lines of behavioral fracture” for verbal humans may be fundamentally different from those that we find with non-verbal organisms. Without a clear focus on the alternatives, behavior analysis could well cut itself on Occam’s razor.

Several of the commentators suggested alternatives—either their own or others. In our view, however, none of the currently available accounts can readily describe, let alone explain, all of the data that has been generated in RFT studies (see Barnes, 1994; Barnes & Roche 1996; Barnes-Holmes, Healy, & Hayes, 2000; Barnes-Holmes, et al., 2001; Dymond & Barnes, 1995; Hayes & Barnes, 1997; Hayes

& Wilson, 1996; O'Hora, Roche, Barnes-Holmes, & Smeets, 2001). It is to these alternatives that we now turn.

Alternative Accounts

One way of responding to the challenge presented by Relational Frame Theory is to explore alternative accounts, and a number of the commentators have indeed suggested alternative views. We will list each of these here and respond briefly to what has been offered.

Sidman's Account. One of the commentators (McIlvane) suggests that there is little difference between Sidman's treatment of equivalence relations and RFT:

The main difference between Sidman's thinking and that articulated in *RFT:ALC* is whether complex forms of stimulus control are *selected by or constructed* via reinforcement contingencies. Once the selectionist vs. constructivist issue is ultimately settled, I anticipate fairly few points of major disagreement between Sidman's followers and the thinking represented in *RFT:ALC* (p. 33).

Furthermore, as noted earlier, Salzinger argued that Sidman's "bag analogy" is at least equally attractive as the concept of a relational frame. The key difference, then, seems to be that for Sidman, the contingencies select a phylogenically established basic stimulus function, but for RFT the contingencies construct learned or overarching operant classes.

Theoretically, we don't think this difference makes a difference—as argued in the previous section, all learned behaviors have their origins in phylogenetic contingencies. We also pointed out, however, that in practical terms the different foci on phylogeny versus ontogeny may have massive implications for the research programs that develop from the two theoretical perspectives. Interestingly, Spradlin, who devotes much of his commentary to comparing and contrasting Sidman's approach with RFT, seems to agree with our view when he writes, "... the research agendas implied by the two models are very different" (p. 3).

A related problem in brushing over the differences between the two accounts is that Sidman has focused almost exclusively on stimulus equivalence, which he argues is a basic stimulus function or biological given. But what of other derived relations? Are these, too,

basic stimulus functions? Surely not, since we can make up new relations at a whim (see page 40 in the book). At the present time, it is difficult to judge Sidman's position on this, because the work on non-equivalence relations within that tradition is so sparse and theoretically confused. On the one hand, for example, Carrigan and Sidman (1992) stated "Difference, opposition, and negation ... are never equivalence relations; they do not exhibit reflexivity, and need not exhibit transitivity" (p. 185), but Sidman has also argued that these other relations may be interpreted as equivalence relations under contextual control. In the book, we quoted Sidman on this very point: "The fact that a stimulus pair can be brought via contextual control into such differing relations as same, opposite, different, and so forth, can be handled by any formulation of equivalence that recognizes the role of context" (1994, p. 561)." We then went on to point out that "This sentence appears to be Sidman's only treatment of multiple stimulus relations."

McIlvane, in his commentary, referred to research on ordinal classes as an example of Sidman's approach to non-equivalence relations, but this work begs the same question: are ordinal relations another biological given or contextually controlled equivalence classes? Furthermore, in the book we point out that RFT has generated data that does not appear to yield to equivalence or ordinal-class based analyses (pp. 59-62; see also Dymond & Barnes, 1995, pp. 182-183). At the present time, therefore, it is extremely difficult to compare directly Sidman's account with that of RFT. The two models, as Spradlin points out, appear to call for very different research programs, and Sidman's treatment of non-equivalence relations, which are the *sin qua non* of RFT, remain at best underdeveloped and at worst confused and contradictory.

Malott's Chaining Account

In the commentary provided by Malott an explanation of symmetry and transitivity is presented in terms of a behavioral chain. The explanation rests critically upon a vocal repertoire that is accidentally reinforced as a tact during the symbolic matching training (i.e., the echoic "Mark" is reinforced in the presence of the photograph). Like Horne and Lowe's (1996) naming hypothesis, therefore, Malott's

account seems to require a vocal response or tact, which mediates the derived relational response. Even the relatively limited data that is currently available suggests that this type of vocal chaining explanation is inadequate. Shusterman and Kastak (1993), for example, provided some evidence to suggest that exemplar training with a sea lion (without any differential name training) can produce symmetry and transitivity in a matching-to-sample task. Furthermore, Lipkens, Hayes, and Hayes (1993) reported a longitudinal study in which symmetrical responses between a heard name and a seen object were observed in an infant who failed tests for differential naming. More recently, Carr, Wilkinson, Blackman, and McIlvane (2000) have reported equivalence in learning disabled individuals with very limited language abilities (but years of experience with MTS procedures). Finally, a current research project in the second author's lab has obtained good evidence for equivalence responding in a mute autistic child (but with an extensive history of matching-to-sample training). These types of findings clearly suggest that derived relational responding (at least simple forms) may occur in the absence of an expressive vocal repertoire. Insofar as Malott's chaining account requires such a repertoire, his mediational explanation for symmetry and transitivity seems untenable.

It is also worth noting that mediational accounts of stimulus equivalence have also been around since the 1930's, and they failed. Sidman (1994), himself, abandoned them early on in the equivalence research program. Recently, Tonneau (2001) has argued that we should give the mediational concepts of Pavlovian conditioning another go, but he also points out that his Pavlovian approach to anything beyond the most simple transfer performances will be too complex to yield to experimental analysis (p. 123).

It may be possible to construct a mediational account of the human language and cognitive abilities that are covered in RFT. However, no one has done it yet in over a century of mediational theorizing. If someone does, we will see if it moves behavior analysis significantly closer to the goals of prediction-and-influence in the domains of human language and cognition in both basic and applied settings. Only this outcome will decide which is the best option to take—the RFT approach or some me-

diational or Pavlovian account? We invite Malott and others to build out their theoretical and empirical analyses to a level comparable to that of RFT to see if the alternatives can achieve something that our theory cannot. This would certainly invigorate the discipline of behavior analysis.

Attributes, Rules, and Logic

In the commentary provided by Salzinger, seven formal attributes of verbal behavior were presented, with a view to seeking out the various characteristics that distinguish it from all other behavior. In doing so, Salzinger implies that this strategy provides a viable alternative to RFT, which has involved identifying one characteristic. Later in his commentary Salzinger also suggests that rule-governed behavior and formal logic may provide an alternative explanation for derived relational responding. In our view, these cannot be considered viable alternatives to RFT.

Listing the formal attributes of verbal behavior may be a useful exercise early in a research program, but such a list will never provide the technical or functional-analytic definition of verbal behavior that we have sought to develop in RFT. If we are to use the term verbal in a technical sense in behavior analysis, we need to delineate it *functionally* from all other classes of non-verbal behavior. The list Salzinger provides, like the definition by Hockett, may help orient us towards important functional properties, but it cannot substitute for a genuine functional-analytic definition or analysis. In the book we are precise in our definition of verbal behavior:

Relational frame theory takes the position that derived stimulus relations constitute the core of verbal behavior. *Verbal behavior is the action of framing events relationally.* Both speakers and listeners engage in verbal behavior. When a speaker frames events relationally and produces sequences of stimuli as a result, the speaker is engaging in verbal behavior. In more lay terms, we say that the speaker is speaking with meaning. If the same formal stimuli are produced but not because the speaker has framed events relationally (e.g., when a parrot repeats what is said), then no verbal behavior is involved. Verbal meaning, in this approach, is not a mental event, nor an inference, nor a simple effect. It is a highly specified behav-

ioral process (see the earlier section on the nature of relational frames for that specification) (pp. 43–44).

Salzinger and others may not agree with the RFT functional-analytic definition of verbal behavior, but in so doing it is incumbent upon them to provide an alternative functional definition—providing a list of formal characteristics will not suffice.

Appealing to rule-governed behavior or formal logic as possible alternatives to RFT falls victim to the same problem. Both of these concepts do *not* have clear functional-analytic definitions, and thus they cannot be used as the basis for a behavior-analytic explanation. In short, rule-governance and formal logic themselves require explanation, and it was for this reason that we subjected them to RFT analyses in the book. Once again, we were relatively precise in our functional-analytic treatments of these concepts. With respect to rule-governed behavior we suggested the following:

In RFT there is a clear difference between non-verbal and verbal regulation. Rule-governed behavior is a subset of verbal regulation. The term becomes more likely to be used when the verbal antecedent is a relational network or a comparison of such networks, and especially when comparison between a verbal antecedent and the verbal construction of ongoing events is part of the source of control over behavior regulation. Rule-governance is also more likely to be used when the non-arbitrary features of the environment are abstracted and transformed, and when the verbal network is generally applicable (p. 108).

And with respect to logic, we argued as follows:

Logic, from an RFT perspective, is in essence a relational activity that involves the derived transformation of function in accordance with multiple stimulus relations. Although specific examples of relational framing may sometimes appear logical (e.g., if A is the same as B and B is the same as C, then A is the same as C), logic does not provide an explanation for relational framing. Instead, it is RFT that provides the basis for a behavioral explanation of logical reasoning, including instances in which individuals fail to reason logically (p. 191).

We should emphasize that we think it entirely reasonable, and indeed intellectually

healthy, to disagree with the RFT explanations that we offer of verbal behavior, rules, and logic. Nevertheless, we also believe that it is behaviorally nonsensical to argue that non-functionally defined concepts, such as rules and logic, can be used to explain the functionally defined concepts of RFT. Doing so is like arguing that the concept of the operant should be explained using the lay concepts of intention and purpose.

SUMMARY

As we looked over the various alternative accounts it became apparent that none of them attempt to cover the entire domain of human language and cognition. Even Sidman's treatment of equivalence relations, which is now over thirty years old, is still focused almost exclusively on one particular derived relation and seems to have moved progressively further away from where it started—a functional analysis of human symbolic control (Barnes-Holmes, Hayes, & Roche, 2001). Other possible alternatives, such as Malott's (and Tonneau's) chaining and mediational accounts, are likewise limited in their foci and have their historical roots in analyses that have not worked very well empirically or conceptually. Furthermore, none of the accounts, apart from Sidman's, have a robust research program linked to them. And the remaining alternatives, which appeal to formal definitions of verbal behavior and to the concepts of rules and logic, ipso facto, fail to provide genuine functional-analytic explanations. In contrast, RFT has made significant progress, both conceptually and empirically, in tackling human language and cognition from a purely functional-analytic perspective.

Much more remains to be done. One direction for future research could involve developing another viable and data-driven functional account that attempts to grapple with the broad range of issues covered by RFT. We would welcome any genuine attempt to step up to this challenge, which we believe would be an extremely healthy sign for our field.

More Minor Points and Misunderstandings

Before moving forward with the closing sections of our reply, we feel that it will be useful

to clarify a number of minor points and address possible misunderstandings that we identified in the commentaries. Because these are mostly minor points of clarification, we will make them relatively brief and refer the reader to the relevant pages of the book for further elaboration and detail. Furthermore, in the interests of brevity, in general we will not provide direct quotations from the commentators, but simply identify the relevant individual(s) and page number(s) from which the point was derived.

The Distinction between a Contextual Cue and a Discriminative Stimulus (McIlvane & Salzinger). Relational frame theory reserves the term contextual cue for those stimuli that control (bring to bear) a history of arbitrarily applicable relational responding. We use the term contextual cue, rather than discriminative stimulus because; (1) it brings to bear a specific history of relational responding, and thus the relational response cannot be traced directly to the physical properties of the stimuli, and (2) a contextual cue can itself acquire its controlling properties via relational framing, and thus it does not fit the standard definition of a discriminative stimulus (see pp. 25–27, 32–33, 62–64).

Transformation of Function (McIlvane & Osborne). In RFT, the transformation of function is the result of an overarching history of arbitrarily applicable relational responding, and this effect does not seem to fit with any of the traditional technical concepts in behavior analysis. If three stimuli, A, B, C, participate in an equivalence relation, and A is established as a discriminative stimulus through a history of differential reinforcement, given appropriate contextual cues both B and C may acquire similar discriminative-like properties in the absence of any direct reinforcement. When this occurs, we say that the functions of the B and C stimuli have been *transformed* in accordance with the mutually and combinatorially entailed relations. The term *transformation* denotes that the previously “neutral” functions of the B and C stimuli have been changed or altered in a specific way based on an extended history of arbitrarily applicable relational responding, rather than through direct training, primary stimulus generalization, respondent conditioning, and the like. Stimulus functions can also be transformed in accordance with non-equivalence relations; for example if A and B partici-

pate in a frame of opposite, and A is established as a punishing stimulus, B may acquire a reinforcing function (i.e., if A and B are opposite and A is bad, then B must be good). (pp. 31–33, 150).

Resistance to the idea of a new principle is understandable, particularly the transformation of function, because it seems to impact upon every other known behavioral principle, and thus RFT is revolutionary in this sense (Spradlin). For example, a respondent conditioned appetitive CS can be transformed into an aversive CS if it comes to participate in the appropriate relational frames. It is for this reason, we suspect, that McIlvane, following a similar line of thinking to that of Malott’s chaining example, suggests that transformation of function may parallel “the concepts of behavioral resurgence, spontaneous interconnection of repertoires, and other similar phenomena” (p. 33). From this perspective, the transformation of function, to be readily understood, should be broken down into smaller or more traditional analytic units. Of course, any behavioral researcher is free to engage in this interpretive exercise (although we note that no behavior analyst has attempted to undertake this task in any serious way). As stated previously, however, it is our position that the power and utility of RFT is derived, in large part, from the way in which it reconceptualises the analytic units of operant psychology for the treatment of human language and cognition. If others feel that the concept of transformation of function is too molar or descriptive (e.g., Malott), it is incumbent upon them to develop alternative, coherent treatments that prove to be equally, if not more productive than RFT, in the empirical and conceptual domains covered by RFT.

Stimulus Generalization as an Agreed upon Process (McIlvane). Many of the core concepts in behavior analysis have extended histories of conceptual and empirical research attached to them with certain key issues remaining unresolved, and as McIlvane correctly points out stimulus generalization is one of these. However, the same could be said for so many other concepts, too, including the very definition of operant versus respondent behaviors (e.g., Donahoe, Burgos, & Palmer, 1993). To engage at this level in the RFT book, however, would have undermined one of the primary purposes of the text—to broaden the behavior-analytic

research agenda and to attempt wider communication with those who have virtually forgotten that behavior analysis even exists (p. 1). We do not suggest, however, that the study of stimulus generalization (and numerous other issues, too) are inherently unimportant and resolved to everyone's satisfaction—merely that the broadly agreed-upon usage of such terms in behavior analysis is sufficiently well established to allow us to move forward with the RFT research agenda.

Common Sense Examples of Relational Networks of Familiar, Meaningful Stimuli Are Less Arbitrary (McIlvane). In the book we provide a number of common-sense examples of relational networks (i.e., based on the types of semantic relations found in natural languages).

The purpose of these examples was to illustrate to the non-behavioral reader that the networks typically found in derived relations studies could well provide a valid model of the types of relational networks found in natural language (e.g., pp. 52-57). To the average derived stimulus relations researcher these examples may well seem course, but in the service of communicating with a wider audience we felt it was a useful rhetorical device. However, and this is extremely important, in no way did we imply that natural language relational networks are somehow less arbitrarily applicable than laboratory-induced networks composed of nonsense syllables and/or abstract stimuli. Indeed, we gave explicit examples to the contrary (p. 55). To make such a claim would undermine the RFT definition of verbal behavior itself (p. 43-44).

Arbitrarily Applicable Relational Responding Is All that Humans Do (Osborne). From the RFT perspective, it is almost certain that as a history of relational framing increases, it will predominate in the human behavioral repertoire. In fact, many religions, some therapies, and the Acceptance and Commitment Therapy (Hayes, Strosahl, & Wilson, 1999) work in particular, are devoted, in part, to undermining excessive verbal or arbitrary relational control in those domains in which it seems to create human suffering (pp. 214-219, 244-251). For RFT, therefore, humans can engage in both verbal and non-verbal behavior, although excessive levels of the former may be the basis for a wide range of psychological problems. It should also be noted, however, that there are subtleties in distinguishing between verbal and

non-verbal behavior, which none of the commentators picked up on. For example, in the book we distinguish between verbal and non-verbal contextual control over arbitrarily applicable relational responding; in the former case, the contextual cue itself participates in a relational frame, but in the latter case it does not (this distinction may have some bearing on conditioning without awareness, the development of well-practiced skills, and the like; see pp. 88-91). In short, we need to distinguish between verbal and non-verbal contextual control over relational framing, and nonverbal behavior (in which relational framing is completely absent). Where the boundaries lie and how one crosses from one domain to the other is largely an empirical matter. In fact, the talk of boundaries might be misleading. It may be more useful and accurate to talk of degrees of "verbalness" in any given human action, rather than discrete categories of verbal versus non-verbal behavior. This would certainly be in keeping with RFT's contextualistic and developmental focus (p. 151-154).

The Goal: An Account of Language

One generic "misunderstanding" that seemed to emerge from time to time across some of the commentaries is particularly troublesome—the idea that RFT is a theory of stimulus equivalence rather than a theory of human language and cognition. For example, McIlvane states:

Given the voluminous body of behavior analytic work on relational learning summarized in *RFT:ALC*, Horne and Lowe's article, and Sidman's (1994) book, one would expect there to have already published a comprehensive, detailed treatment of *theories of stimulus equivalence* and related phenomenon. However, such a treatment has not yet appeared, and the authors chose not to attempt one in *RFT:ALC*. (p. 32, emphasis added)

Furthermore, as pointed out previously, Spradlin devoted much of his review to comparing and contrasting Sidman's (1994) book with the current volume (although he clearly recognizes the different implications arising from the two works), and Salzinger compared Sidman's bag analogy for equivalence classes with the RFT concept of the relational response.

This “equivalence-centric” approach to RFT, although inappropriate, is understandable given the history of the theory and the research histories of some of the commentators (two of whom are leaders in the field of stimulus equivalence). Nevertheless, some of the criticisms voiced by the commentators can be traced to their thinking of RFT as a theory of stimulus equivalence, and thus in our view these comments seem off target. For example, both McIlvane and Osborne bemoan the lack of engagement with, and references pertaining to, other behavior-analytic thinking on relational learning. In the Preface to the book, however, we pointed out that the “scholarly engagement” material could be found in various other RFT articles, some of which contain detailed comparisons of RFT with alternative accounts (e.g., Barnes, 1994; Barnes & Roche 1996; Barnes-Holmes, Healy, & Hayes, 2000; Barnes-Holmes, et al., 2001; Dymond & Barnes, 1995; Hayes & Barnes, 1997; Hayes & Wilson, 1996; O’Hora, et al., 2001) and that the broader purposes of the book would not be served by repeating these arguments in detail. Many of the arcane and abstruse “in-house” issues on relational learning and equivalence seem unimportant as measured against the goal of an account of human language and cognition as opposed to a theory of stimulus equivalence as such.

The same general response applies to the criticism that RFT has focused too much attention on normal adults, without dealing with the relevant pre-experimental histories (McIlvane; see also Spradlin). Because RFT is concerned with all areas of human language and cognition, all humans, young and old, are relevant and important for the research agenda. The origin of stimulus equivalence is but one very small part of RFT. Even if the equivalence research community manages one day to work out whether or not equivalence is a behavioral primitive, we will need to understand how this single behavioral repertoire feeds into the vast universe of human language and cognition. It is our view that science does not always do well by waiting until a specific and perhaps not very important question has been answered before moving onto important areas of inquiry.

A similar argument applies to the question concerning the early RFT correlational studies using language-disabled children. For example, McIlvane asks why the Devany, Hayes, and

Nelson (1986) study has not been replicated, except for the ambiguous results of Barnes, McCullagh, and Keenan (1990). First, we don’t agree that the Barnes, et al. results were ambiguous. The one child who failed to demonstrate equivalence also failed to demonstrate symmetry between her expressive and receptive repertoires—the other children who passed the equivalence test possessed both of these repertoires. Second, we are not sure if this question is an extension of comments McIlvane and co-authors previously made (Carr et al., 2000) that the *training* results for Devany et al. (1986) were unusual, or if McIlvane is speaking about the RFT-relevant testing findings from that study. If the former, we would note only what Carr et al. (2000) agree to: every component of the training results have been shown elsewhere. If the latter, we would note that correlational studies of this kind are only one method useful in confirming a general relationship between derived stimulus relations and language. Subsequently, experimental studies have provided stronger systematic replications of the RFT approach to this area. Some of this work is described briefly in the book (p. 183–191), and other full research articles are either currently in preparation or under submission to relevant journals.

We agree with McIlvane that longitudinal studies with normally developing infants are required, but that applies with equal force to the theoretical ideas of other researchers in this broad area. The only currently published study of this type was conducted in the RFT lab at Reno; Lipkens, et al., 1993, and other RFT research is currently underway, including studies using event-related potentials (ERPs) to get around the experimental problems associated with behavioral procedures using infants. In our view, however, the true test of the approach is whether or not it leads to greater prediction-and-influence across the entire domain of human language and cognition. Human infant studies are only a small part of the research agenda being pursued.

RFT, as the book title declares, is a theory of human language and cognition. In fact, from the language-and-cognition focused perspective of RFT, one could turn the infant-studies argument on its head and ask why has there been no published research conducted with senior citizens. Our response is that we have plans in that direction too, but the field clearly needs

many more human operant laboratories to mount the entire RFT research program.

Post-Skinnerian Behavior Analysis?

Some of the comments seem more concerned with past identity than in assessing the challenge of RFT per se, and particular commentators react strongly to the “post-Skinnerian” label we attach to Relational Frame Theory. For example, Osbourne says “As much as the authors of *RFT* appear to want to separate themselves from Skinner—witness their subtitle—there are interesting parallels that they cannot avoid” (p. 20). Osbourne and others spend time arguing that RFT owes much to Skinner’s work.

This is undoubtedly true, and hardly surprising. The book was written by behavior analysts. Indeed, the book is dedicated to Skinner and Kantor. Fortunately, we were very precise about what we meant by the term “post-Skinnerian”:

This book is behavior analytic but it is also post-Skinnerian. Calling it that is meant both to do honor to a great intellect of the twentieth century, and to acknowledge that this volume represents a form of that approach that is quite different in many of the specific topical domains of interest to human psychology. If a behavioral psychologist adopts a Relational Frame Theory perspective, the world of human events suddenly changes. A whole host of topics that have been written about and studied within the behavioral tradition now are viewed differently.... [following a list of 19 items] When a list of this kind becomes long enough it is time to stop and note that something has happened. For us, this book is such an event. (pp. 253-254)

Based on what we actually say, there should be no confusion about RFT as a “post-Skinnerian” account. It is not an attempt to separate from Skinner, nor merely to note, as Osbourne also suggests, that we are writing after Skinner is dead. It is not “anti-Skinnerian” or “anti-behavioral” (as one track of Salzinger’s work associations lead to). Spradlin has it right: the sub-title “suggests that while the account has its roots in Skinner’s work, that it progresses beyond Skinner’s work” (p. 5). We use the term “post-Skinnerian” a) to note that if RFT is correct, many specific Skinnerian analyses related to complex human behavior need to be redone,

and b) to affirm that we believe that this reanalysis can be done from within the behavior analytic tradition that Skinner established, despite the fact that the reanalysis leaves many of Skinner’s specific ideas behind.

Interestingly, none of the commentators try to show that the implications we see for specific Skinnerian concepts are incorrect. To the extent that the challenge is met, it takes the form of suggesting that other concepts might account for some small part of relational learning (e.g., Malott, McIlvane, Salzinger). However, none of the authors argue that these alternatives could account for the range of RFT data already produced. Furthermore, the commentators sometimes seem to use traditional Skinnerian analyses in ways that themselves would fundamentally challenge Skinner’s ideas if RFT data were explained by them. Salzinger argues that RFT might be accounted for by rule-governed behavior, for example, but since human infants show derived relational responding (e.g., Lipkens, et al., 1993) this account would have the effect of claiming that human infants with minimal verbal abilities are following rules—not a happy result for Skinner’s account of rule governance.

CONCLUSION

In considering the five commentaries together, the following themes seem quite strong. There is general agreement about the need to deal with human language and cognition from a behavioral perspective. Furthermore, there appears to be general agreement that RFT may well have something of value to offer in this regard. The reviewers also recognise that RFT is a behavioral account, and all of them pointed to links with other areas of behavior analysis. McIlvane and Spradlin see considerable overlap with Sidman’s work on equivalence classes, Salzinger with rule-governed behavior, Malott with his work on linguistic productivity, and Osbourne sees RFT as a natural extension of Skinner’s work.

However, the main criticism of *RFT* by three of the commentators seems to be focused on the way in which the theory draws on the concept of the operant. Both Malott and McIlvane, for instance, call for an account in terms of basic behavioral processes. Malott attempts to

provide one example, in terms of stimulus-response chains, and McIlvane mentions “behavioral resurgence, spontaneous interconnection of repertoires, and other similar phenomena....” Salzinger points to rule-governed behavior and logic but does not attempt to explain either at a technical level. In contrast, Osborne seems satisfied with RFT as an operant account to such an extent that he argues that there really is nothing very new here. *In every case, however, there is no attempt to deal directly with multiple stimulus relations and the transformation of function that have been the central focus of RFT studies for the past 13 years.* When examples are given, such as that provided by Malott, or the stimulus class definition provided by Osborne (p. 18), they concern only frames of coordination and thus can yield (in terms of description if not explanation) to class-based analyses.

We believe that it is telling that none of the reviewers attempted to provide a description, let alone an account, of multiple stimulus relations—the very heart of RFT research. We have been asking this question of our colleagues for some time now (e.g., Hayes & Barnes, 1997), and few seem willing even to try, and then only in the most minimalist fashion. For example, Tonneau (2001) mentions the issue in relation to his account and then admits that he believe that the domain is too complex to yield to an experimental analysis cast in his terms (see Barnes-Holmes & Hayes, 2002).

If behavior analysis, as a field, is to face the challenge that RFT presents, the following questions will need to be answered.

1. Are we satisfied that an operant is a basic behavioral process?
2. If so, can we define and empirically identify operant behavior?
3. If we can, can we define and empirically identify traditional relational responding (based on formal properties of relata) as operant behavior?
4. If so, can we define and empirically identify arbitrarily applicable relational responding (as defined in our book) as operant behavior?
5. If so, is it the case empirically that this relational operant modifies stimulus functions established by other learning processes?

If the answer is “Yes” to each of these questions, then the field as a whole will have to deal with the wide-ranging and revolutionary implications that arise from this relational operant. Behavior analysis will have unquestionably entered the “post-Skinnerian” era because, in effect, behavioral psychology will have to re-examine the impact of a wide variety of behavioral processes in verbal organisms that have hitherto provided the bedrock upon which our science is built. Stepping up to this challenge is exactly what RFT attempts to do, but stepping up to the challenge of these five questions is something that the entire field of behavior analysis can no longer comfortably avoid.

REFERENCES

- Barnes, D. (1994). Stimulus equivalence and relational frame theory. *The Psychological Record, 44*, 91–124.
- Barnes, D., McCullagh, P. D., & Keenan, M. (1990). Equivalence class formation in non-hearing impaired and hearing impaired children. *The Analysis of Verbal Behavior, 8*, 19–30.
- Barnes, D. & Roche, B. (1996). Relational frame theory and stimulus equivalence are fundamentally different: A reply to Saunders. *The Psychological Record, 46*, 489–508.
- Barnes-Holmes, D., Healy, O., & Hayes, S. C. (2000). Relational frame theory and the relational evaluation procedure: Approaching human language as derived relational responding. In J. C. Leslie and D. E. Blackman (Eds.), *Experimental and applied analysis of human behavior* (pp. 149–180). Reno, NV: Context Press.
- Barnes-Holmes, D., Hayes, S. C., & Roche, B. (2002). The (not so) strange death of stimulus equivalence. *European Journal of Behavior Analysis, 2*, 35–41.
- Carr, D., Wilkinson, M., Blackman, D. E., & McIlvane, W. J. (2000). Equivalence classes in individuals with minimal verbal repertoires. *Journal of the Experimental Analysis of Behavior, 74*, 101–114.
- Carrigan, P., & Sidman, M. (1992). Conditional discrimination and equivalence relations: A theoretical analysis of control by negative stimuli. *Journal of the Experimental Analysis of Behavior, 58*, 183–204.

- Catania, A. C. (1998). *Learning* (4th ed.). Upper Saddle River, NJ: Prentice-Hall.
- Devany, J. M., Hayes, S. C., & Nelson, R. O. (1986). Equivalence class formation in language-able and language-disabled children. *Journal of the Experimental Analysis of Behavior*, *46*, 243–257.
- Donahoe, J. W., Burgos, J. E., & Palmer, D. C. (1993). A selectionist approach to reinforcement. *Journal of the Experimental Analysis of Behavior*, *60*, 17–40.
- Dymond, S., & Barnes, D. (1995). A transformation of self-discrimination response functions in accordance with the arbitrarily applicable relations of sameness, more-than, and less-than. *Journal of the Experimental Analysis of Behavior*, *64*, 163–184.
- Hayes, S. C. (1987a). Upward and downward continuity: It's time to change our strategic assumptions. *Behavior Analysis*, *22*, 3–6.
- Hayes, S. C. (1987b). Language and the incompatibility of evolutionary and psychological continuity. *Behavior Analysis*, *22*, 49–54.
- Hayes, S. C., & Barnes, D. (1997). Analyzing derived stimulus relations requires more than the concept of stimulus class. *Journal of the Experimental Analysis of Behavior*, *68*, 265–270.
- Hayes, S. C., & Brownstein, A. J. (1986). Mentalism, behavior-behavior relations and a behavior analytic view of the purposes of science. *The Behavior Analyst*, *9*, 175–190.
- Hayes, S. C., & Hayes, L. J. (1992). Verbal relations and the evolution of behavior analysis. *American Psychologist*, *47*, 1383–1395.
- Hayes, S. C., Strosahl, K., & Wilson, K. G. (1999). *Acceptance and Commitment Therapy: An experiential approach to behavior change*. New York: Guilford Press.
- Hayes, S. C., & Wilson, K. (1996). Criticisms of relational frame theory: Implications for a behavior analytic account of derived stimulus relations. *The Psychological Record*, *46*, 221–236.
- Held, R., & Hein, A. (1963). Movement-produced stimulation in the development of visually guided behavior. *Journal of Comparative and Physiological Psychology*, *56*, 872–876.
- Horne, P. J., & Lowe, C. F. (1996). On the origins of naming and other symbolic control. *Journal of the Experimental Analysis of Behavior*, *65*, 185–241.
- Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In I. Lakatos & A. Musgrave (Eds.), *Criticism and the growth of knowledge* (pp. 91–196). Cambridge: Cambridge University Press.
- Lionello-De-Nolf, K. M., & Urcuioli, P. J. (2002). Stimulus control topographies and tests of symmetry in pigeons. *Journal of the Experimental Analysis of Behavior*, *78*, 467–495.
- Lipkens, G., Hayes, S. C., & Hayes, L. J. (1993). Longitudinal study of derived stimulus relations in an infant. *Journal of Experimental Child Psychology*, *56*, 201–239.
- Malott, R. W. (1991). Equivalence and relational frames. In L. J. Hayes & P. N. Chase (Eds.), *Dialogues on verbal behavior* (pp. 41–44). Reno, NV: Context Press.
- O'Hara, D., Roche, B., Barnes-Holmes, D., & Smeets, P. M. (2001). Response latencies to derived stimulus relations: Testing two predictions of relational frame theory. *The Psychological Record*, *52*, 51–75.
- Schusterman, R. J., & Kastak, D. (1993). A California sea lion (*Zalophus californius*) is capable of forming equivalence relations. *The Psychological Record*, *43*, 823–839.
- Sidman, M. (1994). *Equivalence relations: A research story*. Boston, MA: Authors Cooperative.
- Sidman, M. (2000). Equivalence relations and the reinforcement contingency. *Journal of the Experimental Analysis of Behavior*, *74*, 127–146.
- Skinner, B. F. (1957). *Verbal behavior*. New York: Appleton-Century-Crofts.
- Tonneau, F. (2001). Equivalence relations: A critical analysis. *European Journal of Behavior Analysis*, *2*, 1–128. (Includes commentary.)