The Distinction Between Positive and Negative Reinforcement: Use With Care

Alan Baron

University of Wisconsin–Milwaukee

Mark Galizio

University of North Carolina at Wilmington

It is customary in behavior analysis to distinguish between positive and negative reinforcement in terms of whether the reinforcing event involves onset or offset of a stimulus. In a previous article (Baron & Galizio, 2005), we concluded that a distinction of these terms is not only ambiguous but has little if any functional significance. Here, we respond to commentaries by a group of distinguished behavior analysts about the issues we raised. Although several of the commentators argued for preservation of the distinction, we remain unconvinced that its benefits outweigh its weaknesses. Because this distinction is so deeply embedded in the language of behavior analysis, we hardly expect that it will be abandoned. However, we hope that the terms *positive* and *negative reinforcement* will be used with circumspection and with full knowledge of the confusion they can engender.

Key words: classification of reinforcers, positive reinforcement, negative reinforcement, stimulus onset, stimulus offset

In our previous article about positive and negative reinforcement (Baron & Galizio, 2005), we made three points: (a) The customary distinction between positive and negative reinforcement continues to play an influential role in the analysis of behavior. (b) To judge from discussions in the literature, the distinction is straightforward: The designations *positive* and *negative* pertain to the direction of stimulus change, that is, whether a stimulus is presented (positive) or removed (negative). (c) Despite this apparent consensus, the logical and conceptual questions raised about the distinction 30 years ago by Michael (1975) have not been confronted, let alone refuted.

Our purpose was to promote much-needed discussion, and we are pleased that a number of distinguished behavior analysts have offered their comments in this issue of

The Behavior Analyst (Chase, 2006; Iwata, 2006; Lattal & Lattal, 2006; Marr, 2006; Michael, 2006; Sidman, 2006). We thank them for their thoughtful critiques and hope that others will join the discussion in the future. Each set of comments raised unique and interesting points, and this led us to reply to them on an individual basis. However, when possible, we consider common themes. It seems fitting to begin with a discussion of Michael's contribution, in part because his seminal 1975 paper provided the impetus for our review, and also because he provides additional support for the conclusions we reached. We follow with Marr's comments that raise overarching conceptual issues that run through all the commentaries. We then discuss Chase's article, in which he offers reasons why the distinction has been maintained over the years (particularly its role in education), and Iwata's, which considers the role of the distinction within applied behavior analysis. Lattal and Lattal then expand on these themes by making special reference to ways that historical-cultural forces can perpet-

Please address correspondence to Alan Baron, Department of Psychology, University of Wisconsin–Milwaukee, Milwaukee, Wisconsin 53201 (e-mail: ab@uwm.edu), or Mark Galizio, University of North Carolina at Wilmington, Wilmington, North Carolina 28401 (e-mail: galizio@uncw.edu).

uate theoretical distinctions despite evidence to the contrary. Finally, we review Sidman's comments, which are most strongly at odds with our point of view.

MICHAEL

We are gratified that Jack Michael agreed that we had presented an accurate rendition of his original views (Michael, 1975). He should be given full credit for identifying the limitations of current views of positive and negative reinforcement. In his present remarks, Michael (2006) sharpens his argument with a hypothetical experiment. A food-deprived rat's lever press produces a tone in whose presence food pellets are delivered; when the tone is absent, lever pressing is ineffective. Michael points out that although responding to produce the tone establishes its onset as a positive reinforcer, such a designation is quite arbitrary. The experimenter could just as easily have paired the tone-off condition with pellet delivery; for this reason, the distinction between onset and offset is irrelevant. The reinforcing event is better described as a change in stimulation without reference to the direction of the change.

Michael's concerns extend to other types of reinforcement. An instructive case is research on light reinforcement. Kish and others (see Kish, 1965) found that laboratory animals will respond simply to produce an increase in illumination within the experimental chamber. An obvious interpretation, given that responses turn a light on, is that responding is maintained by positive reinforcement. However, research has also shown that animals will respond to reduce the level of an otherwise present light with similar characteristics, an outcome that technically should be attributed to negative reinforcement. But in the sensory reinforcement literature, this latter behavior customarily is also de-

scribed as under the control of positive rather than negative reinforcement. Thus, for unclear reasons, a definition in terms of direction of change has been abandoned in favor of identifying the reinforcer as the stimulus change per se, regardless of its direction.

So Michael's example is not just hypothetical, and, in fact, arrangements such as these have generated confusion about negative as well as positive reinforcement. A case in point comes from experiments on time-out from avoidance. Verhave (1962) was the first to demonstrate that response-contingent periods during which a concurrently operating free-operant avoidance schedule was suspended would maintain behavior. Time-outs in Verhave's study were signaled by tone onset, scheduled in much the same way as positive reinforcers, and he referred to the effect as an example of positive reinforcement derived from an avoidance schedule. One of Verhave's noteworthy findings was that responding could be maintained by continuous schedules of time-out but not by intermittent schedules. Years later, we also conducted research on time-out from avoidance (e.g., Perone & Galizio, 1987) that showed, among other things, that responding could be controlled by variable-interval schedules of time-out as well as by continuous schedules. However, in our research time-outs were signaled by *termination* of white noise and house-light, and we described the time-outs as negative reinforcers. To be sure, our description was not solely based on the onset or offset of the stimuli signaling time-out, a distinction that we had demonstrated to be functionally unimportant, but rather was based on the fact that responses removed the avoidance schedule. Nonetheless, this case shows the pertinence of Michael's example.

Michael's comments about what has been called positive psychology provide an example of the continuing

confusion associated with the positive-negative distinction in applied behavior analysis. No doubt, there are advantages in the use of “positive” techniques of behavior change, not only because they are “pleasant,” but also because of the well-documented side effects of procedures that are variously described as “aversive,” “coercive,” or “unpleasant.” However, we would underscore Michael’s reminder that favoring pleasant over unpleasant interventions is not the same as favoring positive over negative reinforcement (see Perone, 2003, for a cogent review of this issue).

MARR

Jack Marr is well known for his efforts to integrate behavior-analytic concepts with scientific advances in other disciplines, as well as for bringing conceptual clarity to complex issues. His comments (Marr, 2006) provide several welcome illustrations of this. In particular, he underscores the value of symmetry in scientific analysis. Principles that possess symmetry can be described with fewer terms (they are more parsimonious) and they “confer an internal and consistent unity to the field” (p. 126). With respect to reinforcement, he notes that if there are “no functional distinctions to be made between positive and negative reinforcement, then reinforcer effectiveness (by various measures) is invariant under a simple inversion of procedure, that is, arranging, say, onset as opposed to offset of pertinent events” (p. 125). Indeed, this is the crux of our case, but although Marr seems to agree that our argument succeeds on the whole, he also sees problems in the way we developed it.

Marr states that in our discussion of positive and negative reinforcement we failed to identify a “conceptual knot” that entangles the specific consequent operations, their effects, and the reasons for the effects seen.

This knot permeates the controversy and crops up in one way or another in each of the other critiques. No doubt, it is useful to use Marr’s framework to reconsider some of our examples. When we did this, we found that they fell into two categories (Chase reaches a similar conclusion; see below). In some instances, the ambiguity is intrinsic to the specific operations. For example, with the operation of temperature change it cannot be determined whether the reinforcer is the offset of cold or the onset of heat because a change from one level on the temperature scale to another requires the removal of the initial level and the onset of the following level. However, other examples involve two continua, one pertaining to the specific operation (such as aspirin delivery) and the other to the reason for the operation’s reinforcing effects (termination of the headache). Marr is particularly concerned about descriptions that rely on the latter, more remote, effects of reinforcing operations (as is Sidman; see below). He notes, “Traditionally, reinforcers were said to act through the onset of pleasure or the offset of pain. But the Tweedledee–Tweedledum characterization of reinforcement often seems contrived, if not outright implausible” (p. 127).

Marr illustrates his concerns with our textbook example of a child whose behavior is reinforced by viewing a cartoon. We pointed out that one may say that the child’s behavior is positively reinforced by the onset of the cartoon, or, alternatively, that the behavior is negatively reinforced by escape from boredom or from cartoon deprivation. With respect to an analysis that focuses on the reasons for the reinforcing effects of a stimulus change, Marr notes, “Many behavior analysts have, I think wisely, tended to put that question aside with respect to *any* putative reinforcer, positive or negative, while exploring the conditions under which some contingencies

serve to modify behavior”(p. 127). We would certainly agree that it is both customary and usually preferable to focus on the specific operation, in this case, cartoon delivery. But such focus is often blurred in an effort to make the positive–negative distinction something more than a purely procedural distinction. For example, in his mention of our headache scenario, when there is also a specific consequent operation (aspirin delivery), Marr selects the effect of the drug as the reinforcer (termination of the headache). How, then, do we decide when to focus on the specific operation and when to shift to the reasons for its effects? If aspirin delivery is to be viewed as an example of negative reinforcement because the headache is reduced, why is food delivery to a hungry rat to be viewed as positive reinforcement despite the reduction of hunger? What is the basis for distinguishing between the headache and hunger? Of course, this was just our point. Even when a consensus has developed about the way to specify a particular reinforcing event (e.g., that the lever press is reinforced by the food pellet, not by hunger reduction), the rule for choosing between the operation or its effects to define the contingency remains obscure. The interpretative examples that may be found in many textbooks (what Marr terms “just-so” stories) are a product of this conceptual confusion, and to us, this suggests the need to reevaluate the terminology we use to talk about reinforcement.

A related issue posed by the so-called “Tweedledee–Tweedledum” characterization of reinforcement pertains to the contribution of the deprivation state to the outcome. As a way of illustrating this second problem, Marr reports that although he enjoys listening to Beethoven, he doesn’t sense that in the absence of hearing such music he is in an active state of being deprived of 19th century romanticism, a state that

might, for example, be relieved by listening to the Egmont Overture. To be sure, discussions of questions such as these have sometimes included in references to private events (see Baron & Galizio, 2005, pp. 92–93; the problem is that other individuals might report that deprivation of music of a particular genre is accompanied by feeling states that are relieved by listening to music of that type). Regardless of one’s view of the place of such approaches in the analysis of reinforcement, the specific operations are enough to carry our argument. The contingent presentation of the Egmont Overture implies its prior absence and, therefore, entails a period of deprivation that is terminated by the music. Having said this, we should reiterate that we do not dispute the virtues of focusing on the specific consequent operations in preference to less accessible events. This approach has stood behavior analysis in good stead, and it forms an important part of the conceptual framework within which we work. The unresolved dilemma posed by our article resides in those instances in which this familiar approach is apparently abandoned for the sake of the positive–negative distinction.

Marr’s identification of the distinction between specific operations and the reasons for their effects is an important first step in the direction of untying the conceptual knot that he identifies. There are other tangles as well. In addition to the question of whether we should focus on the conditioned or the backup reinforcer (see our discussion of Michael above), what about operations that cannot easily be characterized in terms of onset or offset? And what of the difficulties associated with treating consequences as discrete events, when, more often, they involve transitions from one complex set of environmental conditions to another? As Marr also notes, these sorts of issues relate to questions

about how to deal with reinforcement in general. It was these unanswered questions that led us to ask whether behavior analysis is aided or hindered by distinguishing the two forms of reinforcement. At the end of the day, Marr appears to agree that until a stronger case for the distinction can be made, we probably are better off simply assuming a single, symmetric reinforcement process.

CHASE

We do not have much reason to quarrel with Phil Chase's views of positive and negative reinforcement (Chase, 2006). He concurs with our conclusion that the distinction cannot be sustained in the grounds of logic or the assumption that different processes underlie the two. As a recognized leader in the study of behavioral education and human operant behavior, he is in a good position to explore the reasons why we continue to teach this admittedly ambiguous distinction to students of behavior analysis. According to Chase, consideration of the distinction can serve as an object lesson for the problems that ensue when behavior is not carefully analyzed and precisely measured. Indeed, we have adopted the same strategy in our own teaching, not only to dissuade students from using loose terminology when they talk about reinforcers but also to acquaint them with a concept that has played a central role in the history of behavior analysis.

Chase adds three additional reasons why the distinction may be playing a useful role, namely that one of the two directions of change may be more salient to the observer; that one direction may be the easier to measure; and that one direction may allow a description that employs "less awkward" (p. 114) terminology. For example, in the case of a student who is paid for performance, delivery of a monetary reinforcer is more obvious than a reduction of a presumed state

of monetary deprivation, delivery of the reinforcer can easily be recorded, and it is simpler to talk about delivery of money than termination of a moneyless period. In one way or another, these same points come up in the other commentaries as well. No doubt the features identified by Chase have helped to maintain the traditional distinction, and we certainly do not wish to advocate that behavior analysts focus on less salient consequent events or engage in more awkward language. Indeed, it is easier to talk about contingent monetary payment than contingent termination of moneyless periods, but what do we gain by added the adjectives *positive* or *negative* to our specification of these events? If the distinction cannot be justified on logical or conceptual grounds (Chase's view as well as our own), then how are accounts of behavior furthered by incorporating the distinction? Moreover, as Michael and we have noted, there are pitfalls for the unwary in relying on the addition-subtraction distinction, particularly with respect to more complex contingencies and when the distinction enters into ethical and social decisions.

In his conclusions, Chase makes an important point when he emphasizes that scientific concepts are instances of verbal behavior with multiple determinants that go beyond the sorts of logical arguments that we developed in our article. A scientific analysis also is rule-governed behavior, and we hope that the present discussions will encourage the behavior-analytic community to review the functions of our verbal behavior when we talk about reinforcement, and, if appropriate, to consider alternative modes of expression.

IWATA

As one of the major architects of contemporary functional behavior analysis, Brian Iwata's focus on the

value of the distinction for inquiry into applied issues is particularly valuable. Iwata (2006) starts by reporting that he has always found Michael's arguments compelling. Yet, he also observes that in the applied literature, at least, the distinction has not only remained in use but actually has become more pronounced in recent years (something he documents through a count of the indexing of the term *negative reinforcement* in the *Journal of Applied Behavior Analysis*). Most of Iwata's subsequent comments are directed toward ways that the distinction may be useful for research or practice.

In particular, Iwata attempts to remove the ambiguity of the distinction by characterizing a stimulus change in terms of events that are directly controlled by "the actions of the experimenter" (p. 122; in this regard, his approach is not unlike that of Chase). In Iwata's view, the stimuli provided by delivery of a food pellet should be viewed as the focal reinforcing event. By comparison, the reduction in pellet deprivation that accompanies pellet delivery should not be regarded as "delivery" but rather as the "reversal of a condition" that has previously been arranged by the experimenter (i.e., withholding food). Similarly, in the case of shock termination what is directly delivered to the subject is the termination of shock, not a period of safety. So on this basis, Iwata concludes that "the transition between conditions can easily be described as one involving presentation or removal by the experimenter and classified as positive and negative reinforcement, respectively" (p. 122).

It seems to us that differences between Iwata's interpretation and ours hinge largely on the import of the phrase "actions of the experimenter." Behavior, as we know, can have complex consequences, both local and remote. To reiterate the example of the Egmont Overture

provided by Marr (see above), the action of delivering this musical event to an individual who has been deprived of it is just as surely the agent that terminates the period of deprivation. It is difficult, if not impossible, to speak of one part of such transitions without implying the other. Further, even when it seems quite clear that the experimenter's operation is termination or addition, it is not always clear that characterizing the stimulus change in this way is helpful in a behavior analysis. With regard to Iwata's example of the reinforcing effects of access to free time, should we characterize the event differently if it were signaled by an increase in illumination than we would if it were signaled by dimming the lights?

Iwata also points to the value in applied situations of being able to identify the stimulus changes that serve as reinforcers. For example, the reinforcing effects of free time contingent on completing a task may accrue simply from completion of the work requirement. Alternatively, the behavior may be positively reinforced by activities that are available during the free-time period. Identifying which of these alternative activities will serve as reinforcers in the context of work termination and which will not is precisely the proper goal of a behavior analysis. However, the detailed functional analysis of contingent transitions from one set of conditions to another may be quite complex, and, in our view, there are distinct benefits to be gained by focusing more on the full context of stimulus change and less on addition versus subtraction. At the least, it is not clear to us how the positive-negative distinction adds to the analysis. As Michael stated in his original article, references to removal or presentation may sometimes "stand in place of a more complete description of both the prechange and post-change conditions" (Michael, 1975, p. 41). In this regard, Iwata's com-

ments pinpoint what we regard as the essential problem with the positive–negative distinction. Despite its apparent utility, it may result in misleading or, at the least, premature characterizations of the reinforcement process.

LATTAL AND LATTAL

Andy Lattal is well known for his important theoretical and research contributions in the area of reinforcement, and his coauthor, Alice Darnell Lattal, has helped to extend applied behavior analysis to management settings. Their joint consideration of our article sheds new light on the issues that we raised by extending our concerns to neglected areas (e.g., stimulus generalization, resistance to change) and elaborating issues that we touched on only briefly (e.g., the role of positive and negative reinforcement within social contexts).

Lattal and Lattal (2006) start by making an important point about the difference between positive and negative reinforcement. They note that the logic of scientific research calls for the assumption that a variable has no effect unless there is clear evidence to the contrary. In other words, the accepted practice is that one must seek evidence that might overturn the so-called null hypothesis. By comparison, interpretations of reinforcement have started with the curious assumption that there *is* a difference between the positive and negative varieties. Lattal and Lattal phrase the essential question this way: “If the formal differences between positive and negative reinforcement are not supportable and the functional differences are at least questionable, it remains to be answered why the distinction persists” (pp. 130–131). Why has this conceptual distinction led such a charmed life?

In Lattal and Lattal’s view, the answer lies within the contingencies of our culture. Their point is that principles persist because they work within the culture. Although the

origins of the practice of categorizing events as positive and negative, together with the complementary notions of additive and subtractive processes, are obscure (they trace these notions back to the Book of Job in the Old Testament), the fact that they are still with us attests to their utility. In their discussion, Lattal and Lattal enumerate some of the favorable consequences. Insofar as the distinction aids communication, once it has been established it perpetuates itself. For example, one cannot appreciate the literature on conditioning and learning without being conversant with an issue that has drawn so much attention over the years. Special advantages accrue on the applied side. Employing terms that differentiate between positive and negative reinforcement allow professionals to represent their procedures to the community at large as positive rather than negative. Within this context, the term *positive* is something “good,” to be valued and supported, whereas *negative* is “bad.” Thus, “the language of positive reinforcement” meshes closely with a culture that “places moral value on circumstances construed as positive as opposed to those that are negative” (p. 132).

Lattal and Lattal then ask how we should address the “inconsistency between what is known and what is said” (p. 132). If one recognizes that the distinction between positive and negative reinforcement lacks support (as they do), should one continue to employ it? Here, we are a bit uncomfortable with their answer. In effect, they recommend that behavior analysts lead a double life, speaking in one way to the lay public and a different way to members of our field. They point out that an emphasis on “positive” procedures is virtually required to implement behavioral techniques, given current legal and institutional considerations. As they put it, “communication in practical settings often means putting precision

on the back burner in favor of more user-friendly descriptions than those employed with colleagues" (p. 133). It is not that they do not see dangers in this approach. They add that they regard it as essential that students and professional practitioners of behavior analysis understand the logical and empirical limitations of the distinction. Without such understanding, the distilled (sanitized) version of the concept that has been translated for the sake of communication will serve as "a bastardization and misrepresentation" (p. 133).

Although Lattal and Lattal recognize the dangers as well as the advantages of maintaining the distinction, we would put more weight on the dangers. If behavior analysts speak one way to each other and a different way to a public that is becoming increasingly aware of behavioral principles, this discrepancy will invite the accusation that the behavior-analytic approach is a cynical one. The legal decisions that have surrounded use of token economies in psychiatric hospitals and other custodial institutions are instructive here. The behavioral procedure of providing patients with tokens that could be exchanged for various privileges had been represented as a system based on positive reinforcers. But in the view of the courts, such procedures cannot be employed if they are based on denying patients goods and services that are their right, and, as a consequence, use of such procedures to manage patients' behavior has become uncommon. Thus, although describing a procedure as negative may create bias against it, emphasizing that the procedure involves positive reinforcement is not always sufficient to make it socially acceptable. This leads us to wonder whether the positive-negative distinction is really all that helpful in public dissemination of behavioral techniques. If we spoke only of reinforcement without the positive and negative modifiers, the applied

behavior analyst would still be in a position to emphasize and defend the use of specific events and contingencies as they relate to treatment goals. Importantly, they could do so without distortion or misuse of the scientific terminology.

SIDMAN

We are not comfortable finding ourselves in such disagreement with Murray Sidman. His intellectual example has been an inspiration to both of us throughout our careers. We would like to believe that at least some of his discontent may reflect our own failing to express ourselves clearly enough, and we will do our best here to straighten things out. There is no doubt, however, that he also expresses more fundamental disagreements. Although we may be unable to resolve all of our differences, we will try to define their nature and extent.

Sidman (2006) attributes a more ambitious effort to us than the one we actually undertook. He wondered whether we may have intended to propose a new terminological convention or, perhaps, "a change in basic principles" (p. 135). To the contrary, we simply noted some of the difficulties inherent in the convention of distinguishing between positive and negative reinforcers strictly in terms of the onset or offset of stimulus energy. Such a distinction seemed arbitrary to us insofar as the direction of change often is more a procedural matter than anything else, and the direction of change does not appear to make a fundamental contribution (see Michael's comments). In our article, we explored the literature bearing on whether there are functional differences that are correlated with the onset–offset difference. For example, evidence for different psychophysiological processes would help buttress the distinction. Conversely, in the absence of convincing evidence, we advised

caution in the ways the terms *positive* and *negative* reinforcement are used (see the final paragraph of Baron & Galizio, 2005), and with Michael (1975, 2006) we questioned the merits of preserving the traditional distinction.

Like several of the other commentators, Sidman is at odds with our views about the ambiguity contained within statements about positive and negative reinforcement, and he also expressed particular concern that we found it necessary to appeal to changes in internal states to make our case. We did present several such examples, but most were drawn from accounts of others to illustrate some of the historical approaches to the problem (e.g., food reinforcement as drive reduction). In each example, the ambiguity can be equally well described without reference to physiological or emotional processes. So in the example of contingent attention (an apparent case of positive reinforcement), one could speak of negative reinforcement that accrues from termination of a condition in which attention is lacking rather than saying, "relief from loneliness." (As an aside, it is worth remembering that the primary dictionary definition of *lonely* does not refer to an internal emotional state, as implied in several of the commentaries, but rather to the environment, i.e., "alone, solitary, without company.") Similarly, food delivery can be described as reducing a state of hunger or, alternatively, as terminating a period without food.

Sidman's second and closely related concern is that we too easily accept that it is impossible to determine whether the reinforcer is a consequence of presentation or of removal. He notes that empirical analysis can often identify the specific events or activities that serve to maintain the behavior. As we have already indicated, empirical analysis of reinforcing events can be valuable, but emphasizing addition versus sub-

traction seems to add little to the analysis. That said, there are surely cases, including some common examples from the laboratory, that appear to be fairly clear-cut instances of positive or negative reinforcement (Sidman's example of shock escape is one) or, at least, in which reversing the emphasis may seem contrived. But many examples, including several mentioned by Sidman (divorce, aspirin consumption, TV watching), seem far less clear. Concerning the exceptions, Sidman pleads tolerance in telling us that "areas of definitional confusion exist at the edge of many, if not most, concepts" (p. 136). Obviously, there can be legitimate differences of opinion about when the frequency of exceptions to a rule tip the balance in favor of abandoning, or at least reconsidering, the rule. It is not difficult to find examples of classifications within science that have failed, or at least been seriously questioned, because of exceptions (in the behavioral area, Lattal and Lattal cite the operant–respondent distinction in this regard).

A thread that runs through Sidman's comments pertains to the relation between negative reinforcement and punishment. In our article, we did no more than follow contemporary usage, one with which Sidman is strongly at odds. He vigorously supports an early view of punishment that makes punishment secondary to negative reinforcement, more specifically the view that a stimulus can serve a punishing function only if it also has the properties of a negative reinforcer. Over the years, this view has fallen out of favor largely because of Azrin and Holz's (1966) cogent analysis that placed punishment on a par with reinforcement. In their terms, the distinguishing feature of punishment is that the behavior is weakened, thus making punishment the opposite of reinforcement, which strengthens behavior. The labels *positive* and *negative* are reserved as qualifiers of reinforce-

ment and punishment, that is, to designate strengthening and weakening effects that accompany the presentation and removal of reinforcing and punishing stimuli. Sidman observes that this way of looking at things is rarely questioned these days, and we appreciate his concern that conclusions about negative reinforcement and punishment may have been too hastily drawn. At the least, a careful review of their relation seems in order if we are to take seriously Michael's suggestion that we refer to "good" things as reinforcers (changes that strengthen behavior) and "bad" things as punishers (changes that weaken behavior).

A final area of potential misunderstanding pertains to views of the societal implications of the results of basic science. One of the reasons for retaining the distinction that was considered and rejected by Michael (1975) was that the distinction could be used to make applied behavior analysts more aware of the "undesirable" aspects of negative reinforcement. Michael noted (as we have) that the utility of the distinction is restricted by its lack of clarity. A further limiting feature is that procedures conventionally referred to as involving positive reinforcement may possess many of the same features as those that make negative reinforcement undesirable (see Perone, 2003). Finally, Michael noted that "to maintain a distinction at the level of basic science because of its possible social implications seems a risky practice, and one that is usually avoided in other sciences where possible" (p. 42), a point of view with which we concur. However, Sidman concludes that Michael, and perhaps the present authors, have asserted that societal relevance does not have a role to play in scientific work. To the contrary, we agree that the scientist as a member of society should be sensitive to the social implications and importance of his or her work. For example, the

temptation must be resisted to shy away from research on aversive control simply because of the social problems such techniques may have produced. As Sidman points out, the widespread interest in and use of aversive events as a means of social control demands that the behavioral scientist investigate their properties. As far as we can tell, Sidman, Michael, and the present authors are in agreement on this point.

SUMMARY

Each of the commentators provided valuable insights, and we are grateful to them for identifying areas that call for further discussion. Although several argue that the distinction between positive and negative reinforcement should be preserved, their arguments leave us unconvinced that its benefits outweigh its limitations. We certainly agree with the commentators that there are more avenues to explore and much work to be done to improve the way behavior analysts talk about good and bad things. We hope this exchange will stimulate a renewed interest in these issues. Perhaps further analysis will lend support to a continued distinction between positive and negative reinforcement, but in the meantime we continue to wonder whether the distinction does more harm than good. After all, despite popular belief, the contingencies that we call positive are not always kind and those we call negative can sometimes be more humane.

REFERENCES

- Azrin, N. H., & Holz, W. C. (1965). Punishment. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application* (pp. 380-447). New York: Appleton-Century-Crofts.
- Baron, A., & Galizio, M. (2005). Positive and negative reinforcement: Should the distinction be preserved? *The Behavior Analyst*, 28, 85-98.
- Chase, P. N. (2006). Teaching the distinction between positive and negative rein-

- forcement. *The Behavior Analyst*, 29, 113–115.
- Iwata, B. A. (2006). On the distinction between positive and negative reinforcement. *The Behavior Analyst*, 29, 121–123.
- Kish, G. B. (1965). Studies of sensory reinforcement. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application* (pp. 109–159). New York: Appleton-Century-Crofts.
- Lattal, K. A., & Lattal, A. D. (2006). And yet ...: Further comments on distinguishing positive and negative reinforcement. *The Behavior Analyst*, 29, 129–134.
- Marr, M. J. (2006). Through the looking glass: Symmetry in behavioral principles? *The Behavior Analyst*, 29, 125–128.
- Michael, J. (1975). Positive and negative reinforcement: A distinction that is no longer necessary, or a better way to talk about bad things. *Behaviorism*, 3, 33–44.
- Michael, J. (2006). Comment on Baron and Galizio (2005). *The Behavior Analyst*, 29, 117–119.
- Perone, M. (2003). Negative effects of positive reinforcement. *The Behavior Analyst*, 26, 1–14.
- Perone, M., & Galizio, M. (1987). Variable-interval schedules of timeout from avoidance. *Journal of the Experimental Analysis of Behavior*, 47, 97–113.
- Sidman, M. (2006). The distinction between positive and negative reinforcement: Some additional considerations. *The Behavior Analyst*, 29, 135–139.
- Verhave, T. (1962). Functional properties of a time out from avoidance schedule. *Journal of the Experimental Analysis of Behavior*, 5, 391–422.