

# Moving Beyond Reinforcement and Response Strength

Timothy A. Shahan<sup>1</sup>

Published online: 19 June 2017

© Association for Behavior Analysis International 2017

**Abstract** Behavior analysis has often simultaneously depended upon and denied an implicit, hypothetical process of reinforcement as response strengthening. I discuss what I see as problematic about the use of such an implicit, possibly inaccurate, and likely unfalsifiable theory and describe issues to consider with respect to an alternative view without response strengthening. In my take on such an approach, important events (i.e., “reinforcers”) provide a means to measure learning about predictive relations in the environment by modulating (i.e., inducing) performance dependent upon what is predicted and the relevant motivational mode or behavioral system active at that time (i.e., organismic state). Important events might be phylogenetically important, or they might acquire importance by being useful as signals for guiding an organism to where, when, or how currently relevant events might be obtained (or avoided). Given the role of learning predictive relations in such an approach, it is suggested that a potentially useful first step is to work toward formal descriptions of the structure of the predictive relations embodied in common facets of operant behavior (e.g., response-reinforcer contingencies, conditioned reinforcement, and stimulus control). Ultimately, the success of such an approach will depend upon how well it integrates formal characterizations of predictive relations (and how they are learned without response strengthening) and the relevant concomitant changes in organismic state across time. I also consider how thinking about the relevant processes in such a way might improve both our basic science and our technology of behavior.

**Keywords** Reinforcement · Response strength · Motivation · Information

---

✉ Timothy A. Shahan  
tim.shahan@usu.edu

<sup>1</sup> Department of Psychology, Utah State University, Logan, UT 84322, USA

Reinforcement, *n.*

Something that reinforces or strengthens a material or structure; *esp.* the strengthening structure or material employed in reinforced concrete or plastic (Oxford English Dictionary, 2016).

Although this is clearly not the relevant definition of reinforcement for behavior analysis, it does seem to capture the way we have thought about reinforcement for a long time. Not exactly though, because even most students know better than to say that we have committed the grisly act of reinforcing a living organism. We do, however, reinforce behavior all the time—not with rebar, but with reinforcers. But what is a reinforcer? It is something that reinforces, of course. Unfortunately, if you want to understand the *process* of reinforcement as a basic scientist, this can be a discomfoting state of affairs. Killeen and Jacobs (2016 and this issue) suggest one way this state of affairs might be avoided. They argue that whether or not something serves as a “reinforcer” depends on the motivational state of the organism and the possibilities for perception and action. Both motivation and stimuli predictive of biologically important events change the state of the organism—they change its disposition to make the desired response. When an organism is so disposed, a relevant consummatory activity

will be a satisfier (Thorndike)—an important event (Baum). When that is the case, if an additional contingency is imposed, the organism will make  $R_I$ —an instrumental response—to approach and begin that consummation; the organism is literally moved—motivated—in that context.

So, although this is not exactly reinforcing an organism, the organism is certainly back in the mix—reinforcers are something an organism will work to approach when it is in the relevant state. I find myself moving toward such an argument. It seems important and, overall, pretty satisfying. Killeen and Jacobs ask us to “join the conversation,” so I will.

In what follows, I will expand upon what I see as problematic about the way our field typically thinks about reinforcement—an often simultaneous dependence upon and denial of an implicit response-strength-based theory of reinforcement. I will then briefly review examples of the problem with that implicit theory and describe what I see as important issues to be considered in an alternative view of operant *learning* and *performance* without response strengthening.

## Response Strength as Our Default Theory of Reinforcement

Although the dictionary definition of reinforcement above is not our definition, it seems clear that the root metaphor it provides still serves as the basis of a common implicit theory of reinforcement for behavior analysis. Granted, our typical formal definition of reinforcement does not necessarily refer to the strengthening of behavior. When pushed, most behavior analysts will fall back to a descriptive definition based on only the future frequency or likelihood of a behavior. For example, Cooper, Heron, and Heward (2007) define it thus: “*Reinforcement*. Occurs when a stimulus change immediately follows a response and increases the future frequency of that type of behavior in similar conditions” (p. 702). Usually, we are cautioned to avoid asserting that the behavior increased *because* of reinforcement, as doing so leads one to a discomfoting

circularity. For example, Catania (2013) notes, “Reinforcement is not a theory. It is something that happens in behavior, and we must learn to spot it when it happens...” (p. 66). But, Catania further notes that the term serves both as the name for the outcome and for the *process*, and that to avoid circularity, we must recognize that the term is descriptive, not explanatory. If reinforcement is both a name for the outcome and for the process, what is that process? The standard answer is that the process *is* just the description of the outcome. Commenting on the fields’ descriptive behaviorism, Killeen and Jacobs (2016) suggest that “No, it’s not broke. But these authors believe that it is starting to run on empty”. In my view, it very well might be broke, partially because it is implicitly running on *reserves*.

As Timberlake (1988) cogently described, Skinner’s descriptive framework for behavior has always been based implicitly on the notion of reinforcement as a process of response strengthening. Of course, it was not always implicit. Skinner’s notion of reinforcement was based on a fairly clearly stated theory of response strength as characterized by changes in a hypothetical operant reserve in the *Behavior of Organisms* (Skinner, 1938). Indeed, Killeen (1988) showed how to more formally describe many of the basic properties of this hypothetical system (cf. Catania, 2005). In short, under the appropriate motivational conditions, delivery of reinforcement increases a reserve of responses, and emission of responses drains the reserve. In this theory, response strength is proportional to the size of the reserve. However, as Killeen (1988) made clear, it is not that simple. In attempting to account for the data presented within *The Behavior of Organisms* with this theory, Skinner made use of multiple reserves and a number of complex interactions between them. Ultimately, because of the theory’s shortcomings with data, and perhaps because of its increasing complexity, Skinner abandoned it. But, as Timberlake (1988) notes, although Skinner abandoned trying to *formalize* the theory of operant reserve and response strength, his general approach to operant behavior nevertheless appears to have been heavily influenced implicitly by the process invoked by the theory. To make matters worse, while using this implicit theory, Skinner subsequently directed the rest of us to avoid trying to formalize our thinking about reinforcement and operant behavior. What we were left with was a closed, impenetrable system based entirely on the interdependent definitions of the terms of the three-term contingency. These interdependent definitions are wholly dependent upon unspecified “orderly” changes in behavior (i.e., “future frequency”), and within the implicit theory of response strength that subsumes them, it is difficult to identify the underlying assumptions and to subject them to formal testing that might promote conceptual revision (i.e., to subject them to a scientific approach). We are taught that reinforcers just reinforce, and our job is to spot them and make use of them. We are also taught that the how and/or why of this are bad questions. Reinforcement is merely a description of something that happens. In my view, that is a dereliction of our scientific duty.

There is no doubt that the *technology* of behavior generated by Skinner’s definitional system and implicit theory has been highly effective in many areas of application. We, behavior analysts, are taught that this approach, this technology, is based on a basic *science* of the underlying principles—and that is what makes it more effective than many of the alternatives. Unfortunately, relying on a closed, impenetrable system based on an implicit theory about the critical process is no way to run a science in the long run. If we want to continue to improve the efficacy of our less-than-perfect technology and have a vibrant and lasting science of behavior, we are going to have to let science

do its thing. As scientists, we must be willing to explore thinking about things in a different way. As Killeen and Jacobs (2016) put it,

Skinner's law of effect was a conditional one: Given an effective reinforcer, here is what you can do with it. There was enough in it for him and a generation—even of his students to study to yield many brilliant careers. Now those low-hanging fruits have been plucked.

I think it is time we get serious about building a better ladder.

A skeptic will surely ask whether that implicit and impenetrable theory of response strength with which I claim Skinner saddled us is still *really* hiding in the background and forcing us to think about reinforcement and operant behavior in a certain way. Looking around, it does seem to be nearly everywhere. For example, in describing operant behavior, Cooper et al. (2007) quote Skinner (1953) “in operant conditioning we ‘strengthen’ an operant in the sense of making a response more probable or, in actual fact, more frequent” (p. 65). Of course, “strengthen” has been demoted to a scare-quote quarantined *process*, but there it is. The difficulty is that what we really mean is that the response was strengthened, not made more frequent. Especially diagnostic of this linguistic *legerdemain* is that, in Skinner's writings, behavior increases in strength without actually occurring all the time (think “prepotent” or “incipient” responses), but it might occur later, finally revealing to us that such-and-such event of the past was indeed a reinforcer. I challenge the reader to make sense of *Verbal Behavior* (Skinner, 1957), *Science and Human Behavior* (Skinner, 1953), or many of Skinner's other interpretive writings without an implicit reserve-like theory of response strength. If such increases in strength without outright emission are not reflective of the use of an implicit hypothetical reserve driving later responding, what are they?

I have had many discussions with fellow behavior analysts about reinforcement and response strength, and those discussions tend to follow a similar course. It usually ends something like this “Sure, response strengthening is a vague implicit hypothetical process and is hard to defend, but I just had a lapse and erroneously started thinking about how reinforcement works—that is not really how I think about reinforcement day-to-day. I actually just stick to the descriptive definitional approach.” I do not believe it is just a lapse. I think that because we are scientists most of us cannot help but have some sort of theory about how the process of reinforcement actually works. The one Skinner set into motion, and that continues to run in the background, discourages the development of an *explicit* theory of reinforcement.

A common approach to try to avoid thinking about reinforcement as strengthening through some sort of reserve is to adopt a selectionist account. But even Skinner's (1981) *Selection by Consequences* appears to have strengthening under the hood. For example, in that paper, we learn that “Through operant conditioning, new responses could be strengthened (‘reinforced’) by events which immediately followed them” (p. 501). This time, it is “reinforced” that gets the scare quotes. Regardless, as Killeen and Jacobs (2016) note, selectionism with respect to operant behavior is as at least as poorly worked out as a reserve, and in my view, things get strange when you try to formalize it. For example, McDowell (2004) has created a computational model of selection by consequences, but the model has to include a population (or repertoire) of behavior

represented by behavioral genotypes (which must be preserved somewhere, in this case as 10 digit integers in the memory of the computer running the model), sexual reproduction of these genotypes, and some amount of random mutation of the digits comprising the genotypes. In the end, reinforced representations of the behavioral genotypes increase in frequency in the stored behavioral population and are more likely to occur later. At its core, this seems to me to be much like a reserve, with responses stored up (with variation) for later emission. It is in formalizing selection by consequences and trying to make a machine function according to this process that the need for such assumptions becomes clear—without them, selectionism is just a vague metaphor like strengthening.

Of course, response strengthening is not always just an implicit theory shielded from the limelight by a do not-look-under-the-hood descriptive approach. Indeed, two of our most revered formal theories are explicitly based upon response strength. Herrnstein's (1970) matching law is a relative law of effect, and it describes how the relative strength of a response varies with relative reinforcement. But as Rachlin (1971) noted, the matching law in its most general form is really just a restatement of the descriptive definition of reinforcement and, as a result, is arguably tautological and unfalsifiable. Behavioral momentum theory (e.g., Nevin & Grace, 2000) is undoubtedly a theory of response strength relating strength or "mass" to a history of reinforcement. One decided benefit of behavioral momentum theory is that it attempts to formalize how a history of reinforcement impacts the future persistence of behavior, and it does so with what can be characterized as a state variable (i.e., the mass term). Unlike a reserve, this state does not drain out; it is static, and it governs the strength (i.e., persistence) of behavior as revealed by disruption caused by any number of challenges (i.e., extinction, satiation). Unfortunately, this use of a static state has led to difficulties for the equations of the theory in describing an increasing number of phenomena involving changes in reinforcement conditions across time (for examples and elaboration see Craig, Cunningham, & Shahan, 2015; Craig, Nevin, & Odum, 2014; Craig & Shahan, 2016; Nevin et al., *in press*; Shahan & Craig, *in press*). Regardless, it is important to keep in mind that the equations used to describe behavioral data by the matching law and behavioral momentum theory do not themselves require the use of the metaphor of response strength. The metaphor is only our interpretation of what the relations between environment and behavior described by the equations mean. To the extent that they are successful in accounting for behavior, the equations might be interpreted in other ways (see Shahan, 2010, for discussion), and the way in which we interpret them in no way changes the relevant relations in the data, just how we interpret them. The question is this: Could thinking about the data and relevant processes in a different way lead to better theories, novel questions, and an improved behavioral technology? Like Killeen and Jacobs (2016), I think the answer is yes.

### **What Is the Problem with Strengthening Anyway? What is an Alternative?**

For me, the problem with a response-strength (or even a selectionist) theory is not that it relies on hypotheticals to relate past experience to future behavior by having something somewhere accumulated or stored up—such preservation is a logical and physical

necessity for solving the problem of bringing the past into the present. Instead, the first problem is that behavior analysts have belittled other so-called “organism-based” or “cognitive” (or, more derogatorily, “mentalist”) approaches for relying on such mechanisms, while we have similarly relied on them implicitly and without admitting that we are doing so (see Staddon, 1993, for discussion of why the distinction between environment-based and organism-based accounts is misguided). It seems to me that we are being dishonest with ourselves, and as a result, we are prevented from seriously acknowledging and evaluating our assumptions and exploring other approaches that might be helpful. A second problem is that a simple response-strengthening theory does not appear to work particularly well as an account of the acquisition and performance of behavior (as noted above, even Skinner encountered this early on with the operant reserve).

As noted by Timberlake (1988),

Skinner’s use of the term “strengthening,” although explicitly excluded from referring to an internal connection between stimulus and response, is obviously indebted to connectionist ideas in classical learning theory and serves much the same function. (p. 311)

In other words, Skinner’s implicit strengthening theory (and indeed the simple law of effect) is at its core really just a temporal-contiguity-based connectionist or associative learning theory. Any response that closely precedes a reinforcer is strengthened. This is clear in Skinner’s (1948) superstition paper and throughout his subsequent work. Although behavior analysts are taught not to say that an association is formed somewhere, the structure of the account is really the same. It is as if reinforcers have some special stuff (let us call it *reinforcementstuff*) that responses or stimuli gain as result of their temporal proximity (but only forward temporal proximity for some reason that is a mystery within this approach). The more reinforcementstuff a response or stimulus acquires, the greater the reserve or the stronger the bond. In the Pavlovian literature, this reinforcementstuff is called associative value. In the operant world, especially subsequent to the matching law, we tend to refer to it as just reinforcing value—the ability of the reinforcer to increase or maintain behavior.

From this perspective, reinforcement drives the entire show. *Assuming* the relevant motivational conditions are present, reinforcers are responsible both for learning to engage in a particular behavior or learning the relation between two stimuli (i.e., acquisition) and for the future likelihood of behaving with respect to what is learned (i.e., performance). Contiguity is the means by which reinforcement accomplishes both of these tasks, and poorer contiguity results in less buildup of the strength of the learned relation (i.e., the bond). Both acquisition and performance are based on the accumulation of individual experiences of contiguous pairings, and occurrences of either responses or stimuli in the absence of the reinforcer (i.e., extinction) lead to weakening of the bond and less frequent performance because the stuff maintaining the bond is not being replenished.

The problems with such an approach are manifold, and presenting them all is far beyond the scope of the present paper. But, many others (including Killeen & Jacobs, 2016) have described serious problems with using the concept of strengthening-by-contiguity as a means to understand the acquisition and regulation of behavior (for

some relatively modern examples, see Davison & Baum, 2006; Baum, 1973, 2012; Gallistel, 2005; Gallistel, Mark, King, & Latham, 2001; Gallistel & Gibbon, 2002; Staddon & Simmelhag, 1971; Staddon, 1973). I too have discussed some of the issues (Shahan, 2010), so here I will just briefly review a few of my favorites before moving on.

First, many phenomena studied extensively in the area of Pavlovian conditioning (e.g., blocking, overshadowing, relative validity, temporal context effects) suggest that learning and/or performance seem to depend on how well stimuli *predict* unconditioned stimuli (for reviews, see Rescorla, 1988; Wasserman & Miller, 1997; Williams, 1983). Yet, as Williams (1983) summarizes, many of these same phenomena have been demonstrated with operant learning and performance. These and related findings have suggested that learning and/or performance may not be merely the result of the concatenation of individual stimulus-reinforcer or response-reinforcer contiguities over time, as suggested by the implicit strengthening-based theory employed by behavior analysts. Rather, predictiveness itself seems to play an important role.

Second, in terms of learning, there is evidence that reinforcement may not, in fact, be driving the entire show. For example, the phenomenon of sensory preconditioning is well established empirically (for reviews, see Thompson, 1972; Wasserman & Miller, 1997). In phase 1 of a sensory preconditioning procedure, one neutral stimulus is made to be predictive of a different neutral stimulus ( $S1 \rightarrow S2$ ) in the absence of a US. In phase 2,  $S2$  is arranged so as to be predictive of an unconditioned stimulus ( $S2 \rightarrow US$ ). Finally, in phase 3,  $S1$  is presented alone and generates a conditioned response. Thus, although reinforcement (i.e., the US) was never present in phase 1, organisms nevertheless learn the predictive relation between  $S1$  and  $S2$ . It appears that such a predictive relation itself is sufficient for learning, but this learning is not apparent in performance until  $S2$  is made motivationally important in phase 2.

Now, lest the reader think that sensory-preconditioning in Pavlovian conditioning might have little to say about learning the relation between responses and consequent stimuli characteristic of operant conditioning, operant response-preconditioning has also been demonstrated (St. Claire-Smith & MacLaren, 1983). In phase 1 of this experiment, rats without pretraining were placed in operant chambers with a lever and with the food magazines covered. For the response-preconditioning group, each of the first 20 presses of the lever produced only a 1-s neutral, diffuse stimulus presentation (white noise or house light) and the rats were then removed from the chamber. The next day, in phase 2, the lever was removed and the neutral stimulus was presented 30 times and followed by a food presentation each time. The subsequent day, in phase 3, rats were placed in the chamber with the lever available. Importantly, presses to the lever had no effect (i.e., neither the stimulus nor the food was presented). The results showed that the rats in the response-preconditioning group pressed the lever more than groups of control rats that had received either (a) the same treatment in phase 1 but pairings of a different stimulus with food in phase 2 or (b) the same number of presentations of the neutral stimulus (i.e.,  $S1$ ) but uncorrelated with responses in phase 1. Thus, it appears that rats in the response-preconditioning group had learned the relation between lever presses and neutral stimulus presentations in phase 1 in the absence of reinforcement. It was only after  $S2$  was made motivationally relevant in phase 2 that it became apparent that the relation between the response and the stimulus that was neutral at the time had been learned.

Being trained as a behavior analyst allows me to anticipate how the reader might interpret these response-preconditioning results. Step 1 is to find the reinforcer in phase 1 that could be responsible for the apparent response strengthening. Well, that is easy. All that is required is to assert that the neutral stimulus produced by the response was not actually neutral—the stimulus in fact had sensory-reinforcing effects. In this account, those inferred sensory reinforcing effects are responsible for the learning that occurred.<sup>1</sup> There is no doubt that this is consistent with the standard behavior analytic approach of spotting reinforcement as something that happens. Although it is certainly possible to make this assertion, the critical point is that doing so renders the approach unfalsifiable. The assumption at the very core of this approach is that there *must* be a reinforcer somewhere. If a reinforcer is not readily apparent, then we must make one up. Consider, for instance, our use of “automatic reinforcement,” or even the behavior analytic interpretation of the latent learning effect (see Jensen, 2006)—a can of worms I will not open here. No matter how plausible it may seem to assert that there must be a reinforcer somewhere, and to invent one if we must, this means that our implicit theory of operant learning as driven by response strengthening is nothing more than an unfalsifiable assumption that the account is correct. The entire framework rests on the inference of reinforcement and response strengthening, and if nothing else, this assumption must be acknowledged. But why should we make *this particular* assumption that organisms learn to engage in behavior via reinforcement and infer *this particular* hypothetical process (i.e., strengthening by reinforcement)? Because that is what Skinner taught us to do. For a scientist, is that a good enough reason? I would say no. Relying on this assumption might be considered parsimonious, but parsimony is not an end unto itself. Parsimony is desirable when accompanied by precision, specificity, and falsifiability. It is fool’s gold when, as with pseudoscientific explanations and our implicit theory of strengthening, it derives from dogmatic authority and directs attention from a search for explicit mechanisms.

Third, there is considerable evidence from the study of timing that animals learn the temporal relationships they encounter in the environment, and that learning of these relationships can be demonstrated prior to the development of measured responding (see Balsam, Drew, & Gallistel, 2010, for a review). Although I cannot possibly meaningfully review the relevant evidence here, based on that evidence, Balsam et al. (2010) conclude:

The intervals between events are no longer simply the aspect of experience that conditions the formation of associations: rather the durations of those intervals and the proportions between them are the content or substance of the learning itself. (p.5)

---

<sup>1</sup> It is worth noting that such an approach would still also have to account for the difference between the response-preconditioning group and the control group that received the same number of response-stimulus pairings in phase 1 (and thus, the same number of assumed sensory reinforcers) but a different stimulus paired with food in phase 2. That is, for the response-preconditioning group, the approach somehow would be required to provide a principled explanation of the transfer of the effects of the stimulus-food pairings in phase 2 to the previously established response-stimulus relation (presumably acquired via sensory reinforcement in phase 1) without the stimulus having ever served as a consequence for the response (in phases 2 and 3) after it had been paired with food in phase 2.



Herein lies part of the answer for what an alternative to response-strengthening might look like. Important events such as USs or reinforcers are really just providing us with a means to measure the learning of such predictive relationships via performance, rather than driving the learning themselves. Importantly, our measures of behavior are typically based on anticipatory responding. As Balsam et al. note, we tend to observe anticipatory behavior when a stimulus is relatively temporally proximal to a US (i.e., it precedes it), but if a stimulus is temporally distant from the next US (e.g., as in backward conditioning) and we employ other measures, we tend to see other behavior appropriate for a long wait to the US (e.g., withdrawal, other activities). Thus, it appears that organisms have learned the predictive relations, but “Because we have historically used anticipatory behavior as our index of learning we have been misled into equating learning and anticipation. They are not the same.” (p.5).

The alternative approach suggested by Gallistel and colleagues (e.g., Balsam & Gallistel, 2009; Balsam et al., 2010; Gallistel & Balsam, 2014; Gallistel & Matzel, 2013) is that organisms learn the structure (e.g., what, when, where) of events in their environment, and that the patterns of behavior we observe are the result of what that structure predicts about important events. But, predictive relationships between non-motivationally relevant events might similarly be learned (as in sensory preconditioning). The goal from this perspective then becomes to determine how to formalize those predictive relations in the environment, so we might be better equipped to discover how detection of such relations and the subsequent effects on performance might be accomplished in the absence of strengthening.

Although a behavior analyst might consider such an approach as alien and from “outside” of the field, similar arguments are being made much closer to home. Based on experiments examining how the predictive relations arranged by reinforcers and by stimuli predictive of those reinforcers impact the allocation of responding within concurrent schedules of reinforcement (i.e., choice), Davison and Baum (2006) suggest:

The most general principle, rather than a strengthening and weakening by consequences, may be that whatever events predict phylogenetically important (i.e., fitness-enhancing or fitness-reducing) events, such as food and pain, will guide behavior into activities that produce fitness-enhancing events and into activities that prevent the fitness-reducing events. (p.281).

Here, as with Gallistel and colleagues above, predictiveness itself, rather than strengthening, is driving the show. Baum (2012) has further suggested that what he calls as “induction” might be a better way to understand what we commonly refer to reinforcement, and that such induction is actually responsible for the ultimate effects of predictiveness on guiding behavioral allocation across time.

With respect to choice and concurrent schedules, Baum (2010) provides a nice review of how shifting reinforcement distributions can produce rapid shifts in the allocation of behavior that appear to be too fast for a reinforcement-based strengthening/weakening approach. Gallistel et al. (2001) provide similar evidence. Indeed, Gallistel et al. (2007) provide additional evidence that the pattern of allocation of behavior that we know as the matching law might be innate and guided by detection of the predictive relations about relative payoffs (i.e., “reinforcement”) across available options. This suggestion is consistent with an earlier suggestion that the matching law

might reflect such induced, innate, or unconditioned behavior (Heyman, 1982). Given, as noted above, that the matching law is really just a restatement of the law of effect and the definition of reinforcement, such findings and the alternative view to which they point represent fundamental challenges to the necessity of the law of effect and the response strengthening it implies (see Gallistel, 2005, for discussion).

The above considerations are contributing to the ongoing development of a more general alternative view, more accurately a family of alternative views, and the approach of Killeen and Jacobs (2016) can be seen as one member of this family. Their paper is calling our attention back to the organism and back to motivation and the states or dispositions it creates for understanding what events are important for that organism at any given time. Such events might be phylogenetically important, or they might acquire their importance because they are useful for predicting where to find phylogenetically important events (PIEs; Baum, 2010), and ultimately allowing the organism to engage in consummation (what we have traditionally called reinforcement). Organisms are moved (or induced) to approach or work for both PIEs and PIE-predictive events. The sorts of events or stimuli organisms are approaching or working for at any given time depends on the particular motivational mode or behavioral system active at that time (e.g., Timberlake, 1993)—which is itself a function of the organism's history and the relevant signals in the current environment. Organisms learn what signals predict which PIEs, and they learn what actions allow them to approach (or avoid) PIEs or signals predictive of how to get to or avoid PIEs. As noted above, reinforcement strengthening (by PIEs) has always been the means by which we have explained all of this; it has been the process driving learning about signals and actions and the performance of what is learned. But, as the findings and approach described in the paragraphs above suggest, “reinforcers” might just be performance modulators, what we have attributed to reinforcement strengthening might be better characterized as reflecting innate, or induced response allocations to perceived predictive relations between important events in the environment.

For my part, I agree with Gallistel and colleagues (as described above) that such an approach suggests that it is of critical importance to understand what it is that organisms are perceiving with respect to such predictive relations in the environment. Thus, a potentially useful first step it to try to provide formal descriptions of the structure of those predictive relations. By doing so, perhaps we can be better equipped to determine what sort of mechanisms would allow learning about those relations that are ultimately responsible for the guidance and allocation of behavior.

Let me give a couple of examples to flesh out this approach. In earlier papers (Shahan, 2010, 2013), I argued that the available evidence suggests that the typical interpretation of conditioned reinforcement in terms of acquired response strengthening might be misguided. Like many others before, I came to the conclusion that the stimuli we have called conditioned reinforcers might be better understood as signals or signposts that are useful as a means-to-an-end with respect to PIEs. Although this was admittedly a rather vague conceptual approach to the problem of conditioned reinforcement, it did inspire my colleagues and I to empirically demonstrate that a conditioned reinforcer could be established through backward second-order conditioning (Thrailkill & Shahan, 2014)—a finding that is rather perplexing from a traditional strengthening interpretation. In addition, the notion that conditioned reinforcers might function as temporal *signals* also suggested a more formal approach. Specifically, Gallistel and colleagues (e.g., Balsam & Gallistel, 2009; Balsam et al., 2010;

Gallistel & Balsam, 2014; Gallistel & Matzel, 2013; Ward et al., 2012; Ward, Gallistel, & Balsam, 2013) have demonstrated how the quantitative methods of information theory might be used to formalize the predictive signaling relations common in Pavlovian conditioning arrangements. I am well aware that even the use of the word “information” can elicit disgust and anger from behavior analysts, but this is unfounded. Yes, the concept of information, used in a particular way, has played a prominent role in the history of cognitive psychology, but information theory in its most specific sense is nothing more than a quantitative method for formalizing patterns and quantifying *signals*—it is a branch of applied mathematics. There is nothing magical or mentalistic about it (see Jensen, Ward, & Balsam, 2013, for a discussion).<sup>2</sup> Regardless, using information theory, Gallistel and colleagues have made considerable progress in providing a formal quantitative characterization of what it means for one stimulus to be temporally predictive of (to be a *signal* for) another. Based on their successes and the fact that operant-conditioned reinforcement is generally understood to be the result of Pavlovian conditioning, Shahan and Cunningham (2015) applied these same methods to formalize what it means for a conditioned reinforcer to be a signal. In doing so, we resolved a problem with previous applications of information theory to conditioned reinforcement and showed that temporal informativeness might provide a reasonable account of so-called conditioned reinforcement effects. In short, the approach suggests that conditioned reinforcers in our usual experimental paradigms might attract and maintain behavior because they signal a reduction in uncertainty about the expected time to a PIE. In addition, we showed that perhaps the most venerable theory of conditioned reinforcement in the response-strengthening operant tradition (i.e., delay reduction theory; see Fantino, 1977) provides a rather close approximation to relative temporal informativeness. This approach also suggests that the same methods might be extended more generally to formalize reductions in uncertainty about predictive relations between stimuli and PIEs in other non-temporal dimensions (e.g., spatial, relational)—dimensions that our common theories tend to ignore.

The potential promise of the approach described above is that by formalizing predictive relations between stimuli and other stimuli (some of which are PIEs), we might be better equipped to derive what it is we mean when we say “stimulus control” (the “signs” of Killeen & Jacobs, 2016). Our field’s conceptualizations of conditioned reinforcement and discriminative stimuli have usually been closely related. But the notion of the discriminative stimulus is complex and tends to carry with it both an occasioning function and an instigative or motivational function (see Falk, 1994; Nevin, Tota, Torquato, & Shull, 1990; Weiss, 1978; for discussions). Although we all know to say that discriminative stimuli “set the occasion” for a response, it is much harder to say what that really means. Often, as when teaching the concept to new students, we might “inappropriately” resort to saying that a discriminative stimulus serves as a signal for when the response will be reinforced. Maybe signaling this relation is exactly what a discriminative stimulus does. It serves as a conditional; it provides information about when a response will lead to PIEs and how often. As a result, organisms will approach or work for such signals and follow them or perform

<sup>2</sup> To help get over these potential uneasy feelings, I also strongly recommend Gleick (2011). *The Information: A History, A Theory, A Flood*. Pantheon Books: New York, as a non-technical review of the history of information theory and its importance and impact on modern science and society.

actions that they signal are required. Perhaps it is possible to provide a formal characterization of such a signaling effect for discriminative stimuli in a manner similar to that applied to conditioned reinforcement. Time will tell, but this approach certainly seems worthy of consideration. Having said that, it is important to note that the additional instigative/motivational effects of discriminative stimuli are another issue—that is, why and how such stimuli get the organism moving and give its movement a sort of momentum. It seems to me that Killeen and Jacobs (2016) make a fairly convincing case for how such motivational effects might be understood; in short, they are dependent on the affective effects of the stimuli resulting from what they have signaled about PIEs in the past. The key to a more fully developed approach of this sort will be in how it formally integrates these signaling and motivational/affective effects. To do so may also require formalizing how a history of experiences with PIEs in the presence of predictive stimuli is carried forward in time (for one possible approach, see Devenport, Hill, Wilson, & Ogden, 1997; Shahan & Craig, *in press*) and is related to both changes in motivation/affect and behavioral allocation over time. There is much work to be done here. Please join the effort.

The last few paragraphs have been about learning the predictive relations between stimuli and other stimuli (i.e., PIEs). But without the notion of reinforcement as response strengthening, what are we to do about learning the relation between responses and consequent stimuli? We all recognize the importance of contingencies because as Killeen and Jacobs (2016) note, “if an additional contingency is imposed, the organism will make  $R_1$ —an instrumental response”. Organisms do respond for PIEs, and they do respond for PIE predictive stimuli (i.e., conditioned reinforcers). Outside of simple induced changes in allocation, it is clear that the specifics of behavior can be modified or shaped by the imposition of more specific contingencies. The effects of such contingencies are the bread-and-butter of behavior analysis. So how does one characterize the structure of what is being learned via such contingencies in the absence of response strengthening? One possibility my colleagues and I have explored is to again use the quantitative tools of information theory to characterize the structure of operant contingencies (Gallistel, Craig, & Shahan, 2014). Although the approach is technically complex and incomplete (and I will not torture the reader with it here), this application of information theory suggests that it is possible to formalize the relations arranged by operant contingencies in such a way as to incorporate the temporal relations between reinforcers and the responses that preceded them (yes, contiguity plays a role, but by contributing to relative predictiveness). Thus, the approach provides clues about how to solve problems arising in previous attempts (e.g., Gibbon, Berryman, & Thompson, 1974) to formalize operant contingencies “because of the uncertainty as to how temporal parameters are to be incorporated into the contingency framework” (p. 70, Williams, 1983). The goal of trying to work out the details of such an approach is as above with the relations between stimuli. That is, how do we formally characterize the structure of the predictive relations arranged between PIEs (i.e., “reinforcers”) and responses so that we are better equipped to determine what sort of mechanisms might permit detection of such relations? What sort of mechanisms might those be? That is a good question. The problem is that it is not the sort of question we have been trying to answer. Instead, we have *assumed* a likely faulty reinforcement-strengthening mechanism implicitly and/or pretended that there is no question to be answered here.

One may reasonably ask what adopting such an approach would mean for our technology of behavior. Wouldn't this shift toward the organism and away from reinforcement-based response strengthening pull our attention away from the factors that we have made such good use of in our technology? I do not think so. The reason is that from this approach, contingencies and the important events that we call discriminative stimuli, conditioned "reinforcers," and primary "reinforcers" are still critically important for producing behavior change. The efficacy of well-*structured* arrangements of such events in changing behavior cannot be denied, regardless of how we conceptualize them. But, as Lit and Mace (2015) noted about the traditional approach in applied behavior analysis:

...one problem with such a common sense, pragmatic approach is that researchers can become overly focused on recognizing expected results and may overlook other un-anticipated effects. Similarly, the consequence of designing experiments solely to demonstrate a functional relation between intervention techniques and intended outcomes is that the analysis of behavioral processes that control the discriminated operant of concern is lost and potentially crucial side effects of the operation in question may be ignored. (p. 273)

The approach considered above is just a different way of thinking about and, as noted by Killeen and Jacobs (2016), augmenting our understanding of what we are doing. Perhaps by augmenting the way we look at our interventions, we can improve our technology by increasing our ability to see unintended effects and/or adapt our interventions to take greater advantage of more subtle predictive relationships and organismic/motivational/affective factors.

Finally, I suspect that many behavior analysts will be unmoved by my arguments. But, even if an alternative approach like that described here is not deemed acceptable, I hope that at least some readers will pause long enough to consider that maybe they have been reliant on an implicit, potentially unfalsifiable assumption based on an inferred hypothetical process. Maybe some will feel at least a little discomforted when resorting to a purely definitional and/or implicit response strengthening approach. If so, perhaps a subset of those readers will also feel disposed to keep moving toward something a little more satisfying.

**Acknowledgements** Thanks to the behavior analysis seminar group at Utah State University for many conversations on this topic, especially Andy Craig, Paul Cunningham, Greg Madden, and Jay Hinnenkamp. Thanks also to Stéphanie Cousin for her comments on a previous version of the paper.

#### **Compliance with Ethical Standards**

**Conflict of Interest** The author declares that he has no conflicts of interest.

## **References**

Balsam, P. D., & Gallistel, C. R. (2009). Temporal maps and informativeness in associative learning. *Trends in Neurosciences*, *32*, 73–80.

- Balsam, P. D., Drew, M. R., & Gallistel, C. R. (2010). Time and associative learning. *Comparative Cognition and Behavior Reviews*, 5, 1–22.
- Baum, W. M. (1973). The correlation-based law of effect. *Journal of the Experimental Analysis of Behavior*, 20, 137–153.
- Baum, W. M. (2010). Dynamics of choice: a tutorial. *Journal of the Experimental Analysis of Behavior*, 94, 161–174.
- Baum, W. M. (2012). Rethinking reinforcement: allocation, induction, and contingency. *Journal of the Experimental Analysis of Behavior*, 97, 101–124.
- Catania, A. C. (2005). The operant reserve: a computer simulation in (accelerated) real time. *Behavioural Processes*, 69, 257–278.
- Catania, A. C. (2013). Learning. In *Cornwall-on-Hudson* (5th ed.). NY: Sloan.
- Cooper, J. O., Heron, T. E., & Heward, W. L. (2007). *Applied behavior analysis* (2nd ed.). Upper Saddle River: Pearson.
- Craig, A. R., & Shahan, T. A. (2016). Behavioral momentum theory fails to account for the effects of reinforcement rate on resurgence. *Journal of the Experimental Analysis of Behavior*, 105, 375–392.
- Craig, A. R., Nevin, J. A., & Odum, A. L. (2014). Behavioral momentum and resistance to change. In F. K. McSweeney & E. S. Murphy (Eds.), *The Wiley Blackwell handbook of operant and classical conditioning* (pp. 249–274). Oxford: Wiley-Blackwell.
- Craig, A. R., Cunningham, P. J., & Shahan, T. A. (2015). Behavioral momentum and the accumulation of mass: effects of duration of exposure to stimulus-reinforcer relations on relative resistance to extinction. *Journal of the Experimental Analysis of Behavior*, 103, 437–449.
- Davison, M., & Baum, W. M. (2006). Do conditional reinforcers count? *Journal of the Experimental Analysis of Behavior*, 86, 269–283.
- Devenport, L. D., Hill, T., Wilson, M., & Ogden, E. (1997). Tracking and averaging in variable environments: a transition rule. *Journal of Experimental Psychology: Animal Behavior Processes*, 23, 450–460.
- Falk, J. L. (1994). The discriminative stimulus and its reputation: role in the instigation of drug abuse. *Experimental and Clinical Psychopharmacology*, 2, 43–52.
- Fantino, E. (1977). Conditioned reinforcement: choice and information. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 313–339). Englewood Cliffs: Prentice-Hall.
- Gallistel, C. R. (2005). Deconstructing the law of effect. *Games and Economic Behavior*, 52, 410–423.
- Gallistel, C. R., & Gibbon, J. (2002). *The symbolic foundations of conditioned behavior*. Mahwah: Lawrence Erlbaum Associates.
- Gallistel, C. R., & Matzel, L. D. (2013). The neuroscience of learning: beyond the Hebbian synapse. *Annual Review of Psychology*, 64, 169–200.
- Gallistel, C. R., & Balsam, P. D. (2014). Time to rethink the neural mechanisms of learning and memory. *Neurobiology of Learning and Memory*, 108, 136–144.
- Gallistel, C. R., Mark, T. A., King, A. P., & Latham, P. E. (2001). The rat approximates an ideal detector of changes in rates of reward: implications for the law of effect. *Journal of Experimental Psychology: Animal Behavioral Processes*, 27, 354–372.
- Gallistel, C. R., King, A. P., Gottlieb, D., Balci, F., Papachristos, E. B., Szalecki, M., & Carbone, K. S. (2007). Is matching innate? *Journal of the Experimental Analysis of Behavior*, 87, 167–199.
- Gallistel, C. R., Craig, A. R., & Shahan, T. A. (2014). Temporal contingency. *Behavioural Processes*, 101C, 89–96.
- Gibbon, J., Berryman, R., & Thompson, R. L. (1974). Contingency spaces and measures in classical and instrumental conditioning. *Journal of the Experimental Analysis of Behavior*, 21, 585–605.
- Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior*, 13, 243–266.
- Heyman, G. M. (1982). Is time allocation unconditioned behavior? In M. Commons, R. Herrnstein, & H. Rachlin (Eds.), *Quantitative analyses of behavior, Matching and maximizing accounts* (Vol. 2, pp. 459–490). Cambridge: Ballinger Press.
- Jensen, R. (2006). Behaviorism, latent learning, and cognitive maps: needed revisions in introductory psychology textbooks. *The Behavior Analyst*, 29, 187–209.
- Jensen, G., Ward, R. D., & Balsam, P. D. (2013). Information: theory, brain, and behavior. *J Exp Anal Behav*, 100, 408–431.
- Killeen, P. R. (1988). The reflex reserve. *Journal of the Experimental Analysis of Behavior*, 50, 319–331.
- Killeen, P. R., & Jacobs, K. W. (2016). Coal is not black, snow is not white, food is not a reinforcer: the roles of affordances and dispositions in the analysis of behavior. *The Behavior Analyst Advance online prepublication doi*. doi:10.1007/s40614-016-0080-7.

- Lit, K., & Mace, F. C. (2015). Where would ABA be without EAB? An example of translational research on recurrence of operant behavior and treatment relapse. *Mexican Journal of Behavior Analysis*, *41*, 269–288.
- McDowell, J. J. (2004). A computational model of selection by consequences. *Journal of the Experimental Analysis of Behavior*, *81*, 297–317.
- Nevin, J. A., & Grace, R. C. (2000). Behavioral momentum and the law of effect. *Behavioral and Brain Sciences*, *23*, 73–90.
- Nevin, J. A., Tota, M. E., Torquato, R. D., & Shull, R. L. (1990). Alternative reinforcement increases resistance to change: Pavlovian or operant contingencies? *Journal of the Experimental Analysis of Behavior*, *53*, 359–379.
- Nevin, J.A., Craig, A.R., Cunningham, P.J., Podlesnik, C.A., Shahan, T.A., & Sweeney, M.M. (in press). Quantitative models of persistence and relapse from the perspective of behavioral momentum theory: fits and misfits. *Behavioural Processes*.
- Oxford English Dictionary Online. (December 2016). *Oxford University Press* <http://www.oed.com/view/Entry/161609?redirectedFrom=reinforcement> (accessed February 21, 2017).
- Rachlin, H. (1971). On the tautology of the matching law. *Journal of the Experimental Analysis of Behavior*, *15*, 249–251.
- Rescorla, R. A. (1988). Pavlovian conditioning: it's not what you think it is. *American Psychologist*, *43*, 151–160.
- Shahan, T. A. (2010). Conditioned reinforcement and response strength. *Journal of the Experimental Analysis of Behavior*, *93*, 269–289.
- Shahan, T. A. (2013). Attention and conditioned reinforcement. In G. J. Madden (Ed.), *APA handbook of behavior analysis, Vol. 1: Methods and principles* (pp. 387–410). Washington, DC: American Psychological Association.
- Shahan, T. A., & Craig, A. R. (in press). *Resurgence as choice*. doi:10.1016/j.beproc.2016.10.006. *Behavioural Processes*
- Shahan, T. A., & Cunningham, P. (2015). Conditioned reinforcement and information theory reconsidered. *Journal of the Experimental Analysis of Behavior*, *103*, 405–418.
- Skinner, B. F. (1938). *The behavior of organisms*. New York: Appleton-Century-Crofts.
- Skinner, B. F. (1948). Superstition in the pigeon. *Journal of Experimental Psychology*, *38*, 168–172.
- Skinner, B. F. (1953). *Science and human behavior*. New York: Macmillan.
- Skinner, B. F. (1957). *Verbal behavior*. Englewood Cliffs: Prentice-Hall.
- Skinner, B. F. (1981). Selection by consequences. *Science*, *213*, 501–504.
- St. Claire-Smith, R., & MacLaren, D. (1983). Response preconditioning effects. *Journal of Experimental Psychology: Animal Behavior Processes*, *9*, 41–48.
- Staddon, J. E. R. (1973). On the notion of cause, with applications to behaviorism. *Behaviorism*, *1*, 25–63.
- Staddon, J. E. R. (1993). The conventional wisdom of behavior analysis. *Journal of the Experimental Analysis of Behavior*, *60*, 439–447.
- Staddon, J. E. R., & Simmelhag, V. L. (1971). The “superstition” experiment: a reexamination of its implications for the principles of adaptive behavior. *Psychological Review*, *78*, 3–43.
- Thompson, R. F. (1972). Sensory preconditioning. In R. F. Thompson & J. F. Voss (Eds.), *Topics in learning and performance* (pp. 105–129). New York: Academic.
- Thrailkill, E. A., & Shahan, T. A. (2014). Temporal integration and instrumental conditioned reinforcement. *Learning & Behavior*, *42*, 201–208.
- Timberlake, W. (1988). The behavior of organisms: purposive behavior as a type of reflex. *Journal of the Experimental Analysis of Behavior*, *50*, 305–317.
- Timberlake, W. (1993). Behavior systems and reinforcement. An integrative approach. *Journal of the Experimental Analysis of Behavior*, *60*, 105–128.
- Ward, R. D., Gallistel, C. R., Jensen, G., Richards, V. L., Fairhurst, S., & Balsam, P. D. (2012). Conditioned stimulus informativeness governs conditioned stimulus unconditioned stimulus associability. *Journal of Experimental Psychology: Animal Behavior Processes*, *38*, 217–232.
- Ward, R. D., Gallistel, C. R., & Balsam, P. D. (2013). It's the information! *Behavioural Processes*, *95*, 3–7.
- Wasserman, E. A., & Miller, R. R. (1997). What's elementary about associative learning? *Annual Review of Psychology*, *48*, 573–607.
- Weiss, S. J. (1978). Discriminated response and incentive processes in operant conditioning: a two-factor model of stimulus control. *Journal of the Experimental Analysis of Behavior*, *30*, 361–381.
- Williams, B. A. (1983). Revising the principle of reinforcement. *Behaviorism*, *11*, 63–88.