Behavior and Philosophy, 48, 34-46 (2020). © 2020 Cambridge Center for Behavioral Studies

REPLIES TO COMMENTS ON "SELECTION BY CONSEQUENCES IN THE ONTOGENY OF BEHAVIOR: THE PROBLEM OF THE FIRST INSTANCE" [1]

Terry L. Smith
Edinboro University of Pennsylvania

Key terms: behaviorism, control, creativity, environmentalism, qualitative, quantitative, selection, supervenience

The topic of Smith (2019) was the problem of the first instance. Operant conditioning can readily address why the rate of a recurrent response increases or decreases, but how about the first instance? If the first instance is beyond its reach, there are important learned responses that conditioning cannot explain. Consider suicide. It can occur only once per subject, and therefore can never be explained as the result of reinforcement of prior instances. Is operant conditioning therefore unable to say anything significant about suicide? I argue the answer is no. Operant conditioning could explain a suicide if it were the culmination of a process of differential reinforcement of acts of increasing degree of self-harm.

This answer closely follows Skinner's characterization of operant conditioning as always selecting quantitative features. Skinner (1975) offers the example of shaping a rat to surmount a high barrier. At first, the rat is reinforced for surmounting a barrier so low that it can easily cross it. After repeated instances, its jumping and climbing responses cluster around a mean sufficient to clear this barrier. As they do, new responses occur that would be sufficient to clear a somewhat higher barrier. Raising the height slightly will now cause the mean to increase and as a result, new responses will emerge that would be sufficient to clear an even higher barrier. After repeating this process many times, the rat can clear a barrier it would not otherwise have surmounted. This is analogous to a process that occurs in evolution by natural selection. Skinner provides an example drawn from ornithologist, Albert Wolfson (1948), who discusses the origins of the Arctic tern's lengthy eastward migration. If continental drift slowly lengthened the eastward distance the terns must fly, natural selection would not only cause the average distance of eastward migration to increase, but also raise the probability that new variations of even greater distance would occur. In all three examples, selection-by-consequences not only moves the mean value of some quantitative feature, but also moves the range of variability in the same direction. It thereby raises the probability that a novel adaptive response will occur. In these instances, whether in ontogeny or phylogeny, selection can solve the problem of the first instance. This was the first thesis of the target article.

The second thesis generalized from the three examples. In each case, the feature upon which selection operates is some quantity that varies continuously along a dimension: degree of self-harm, height of barrier surmounted, distance flown. I proposed that selection can solve the

[1] I want to thank José Burgos for inviting me to let my article on the problem of the first instance be the target of commentaries. I also thank my commentators, who have been my teachers as well as my critics. Elliott Sober made helpful comments on an earlier draft of my reply to him. The experience of receiving comments, thinking about them, and writing a response has been a pleasure. Correspondence with the author should be addressed to tlsmt@msn.com or to Terry L. Smith, 1669 Columbia Road NW, Apt. 213, Washington, DC 20009.

problem of the first instance when the feature varies in this manner, but not when it varies qualitatively.

The first two commentators take issue with this second thesis. **William Baum** asserts there is no significant difference between how selection operates on quantitative versus qualitative variation and offers his own account to illustrate his point. **Elliott Sober** makes a conditional claim, arguing that if selection can solve the problem of the first instance for quantitative features, then it would be able to do so for qualitative ones as well. When two respected commentators raise the same issue, clarification, if not revision, is in order.

Let me begin with clarification. In his comments, Sober helpfully refers to a two-part article by John Beatty (2016, 2019), that discusses, in the context of evolutionary biology, whether changes in the mean of variation will shift the range of variation in the same direction. Although Beatty is concerned with natural selection and I am concerned with operant conditioning, we end up examining virtually the same set of issues in two different disciplines. Summarizing a few of his key points provides a useful context for discussing the comments.

Beatty notes that if natural selection moves the range of variation in the same direction that it moves the mean, the environment can initiate change and determine its direction. Selection would not have to wait for the first instance of an adaptive variation to occur based on some other process such as mutation or migration. This, he suggests, is the "essence" of Darwinism, which he summarizes as the claim that natural selection is "creative." Skinner makes a parallel claim about operant conditioning. He credits operant conditioning with conferring upon the environment the ability to initiate change and determine its direction. Common sense, by contrast, looks to inner causes for these effects. But Skinner's experience in the laboratory was that contingencies of reinforcement can exercise a degree of control that common sense does not anticipate, and he proposed to explain this control on the basis of selection-by-consequences. And just as Darwin argued that selection in natural settings functions in ways similar to the artificial selection practiced by breeders, Skinner argued that operant conditioning imparts to the contingencies of reinforcement of everyday life the same control over behavior that he observed in his laboratory. This is the Skinnerian counterpart of what Beatty calls the "essence" of Darwinism. In both cases, selection-by-consequences is credited with giving the environment an ability to initiate change and determine its direction. But if this is the case, selection-by-consequences can solve the problem of the first instance. The "creativity" of selection thus implies a solution to the problem of the first instance.

Beatty's two-part history recounts the debate within evolutionary biology about whether natural selection is "creative" in this sense. This thesis had influential critics within biology during the decades before and after the start of the 20th century. Figures such as Hugo de Vries and Thomas Hunt Morgan argued that the processes of variation and selection should be sharply distinguished from one another. Selection acts only as a "sieve" that reduces variation, and never creates it. A parallel development is recognizable in the history of behavior analysis. Skinner's most influential critic within behavior analysis has been John Staddon, who argues that novel behavior "requires explanation in terms of principles of variation, with only the *disappearance* of behaviors being attributable to the effects of reinforcement" (Staddon & Simmelhag, 1971, p. 23, italics in original). This is the learning-theoretic equivalent of early criticisms of the thesis that natural selection is creative.

For all the sharp disagreements between Baum and Staddon, (e.g., Baum, 2004a, 2004b and Staddon, 2004a, 2004b), Baum agrees with Staddon on the assignment of roles to variation and selection. Like Staddon, he sees selection as purely "passive." It therefore "cannot be said to produce variation." He brushes aside my suggestion that selection might solve the problem of the first instance. Selection has only one effect, to reduce variation. The source of novel variation must be found elsewhere, and he expresses confidence that the preponderance of empirical evidence supports this position. He furthermore thinks "Skinner probably … took variation in behavior for granted, as inherent."

What Skinner "took for granted" is unfortunately not easily resolved. Skinner conceived of learning as a process of selection-by-consequences and he worked out the implications over decades. He tried different ways of explaining himself. A careful reading reveals some apparent inconsistencies. He nonetheless had certain convictions that he expressed repeatedly. One such conviction was the creativity of selection. Indeed, he said "as an explanatory mode, selection is responsible only for novelty, for origins" (Catania & Harnad, 1984, p. 503). This is difficult to square with Baum's view that selection is always "passive" and "cannot be said to produce variation." Furthermore, Skinner believed the origin of change lies in the environment, exactly where it should lie if selection can be creative. These typical Skinnerian assertions support my interpretation. In Baum's favor, however, is Skinner's concession that he does not have a solution to the "problem of the first instance" (p. 609). I tried to repair that inconsistency by showing that Skinner solved this problem in his illustration of teaching a rat to clear a high barrier.

Although Baum disagrees with me about the role that selection can play, I cannot help feeling that we are talking past one another. He takes me to be entering the scientific fray over the best explanation of observed novelty. Does selection explain it or is it better explained by a principle of variation? This is an important empirical question, but it is not the question I am raising. I am asking a conceptual question—namely, is selection *capable* of explaining novelty? Perhaps as a matter of empirical fact, principles of variation initiate the majority of (if not all) changes in behavior and determine the direction of change. If so, then selection never actually explains the first instance. But that is not *my* question. I am asking whether selection offers an option for explaining the first instance, and my claim is that it does.

How is this question relevant to the debates internal to behavior analysis? Here again, Beatty's account is useful. The denial that natural selection can be creative was the dominant position in evolutionary biology until theorists began working out the implications of traits underlain by multiple genes, where gene combinations are reshuffled by sexual reproduction. Beatty (2019, pp. 716-720) discusses the result of this theoretical work under the heading, "Selection Creates the Variation It Acts Upon." The Modern Synthesis in biology began with a theoretical demonstration that selection, under these circumstances, could play the creative role Darwin suggested. Instead of waiting for novel variations to be provided by mutation or migration, selection itself could produce novel variations with a bias toward adaptiveness.

How could it do this? Consider a multigene trait such as height. Assume height is influenced at numerous loci, each of which has a combination of two possible types of allele that I will call "plus" and "zero." The plus type adds a small amount to the individual's height, the zero type has no effect. For simplicity, assume a plus/plus locus has twice the effect of a plus/zero locus. The inherited height will be the sum of the plus alleles. Assume the environment selects for greater height. Individuals of greater height will tend to have a higher probability of surviving and reproducing than individuals of lesser height. As a result, there will be a tendency for the mean to change in the direction of greater height. But there is another consequence. As a result of sexual reproduction, genes get reshuffled into new combinations from one generation to another. As the mean moves in the direction of greater height, the proportion of "plus" alleles in the population of reproducing individuals increases. Therefore, the probability increases that there will be an individual in the succeeding generation with more plus alleles than any individual in the preceding generation. The range of variation will therefore move in the same direction as the change in the mean — i.e., some individual in the succeeding generation will be taller than any individual in the previous generation. At the same time, the shortest individual in the succeeding generation will also tend to be taller than the shortest individual of the preceding generation. Selection not only moved the mean in the direction of adaptation, it also moved the range of variation in the same direction, and in doing so it solved the problem of the first instance.

This example raises the question of whether Skinner discussed an analogous process in the ontogeny of behavior. When Skinner (1953) introduced the analogy with natural selection, he posited the existence of behavioral "elements," which (as I say in note #6 of the target article) "is a rare example of Skinner inferring the existence of hypothetical entities." He proposed "to identify the element rather than the response as the unit of behavior" (Skinner, 1953, p. 94). He went further than simply positing the need for such entities. He said something about their properties. He said, "the reinforcement of a response increases the probability of all responses containing the same elements" (ibid.). I take this to mean that there are multiple elements underlying any given operant response and that what contingencies of reinforcement select to be replicated is not simply the response but also the elements underlying the response. So far as I know, Skinner never went on to discuss how these elements come together to produce a response, so the analogy is incomplete. But he carries it further than his philosophy of radical behaviorism might lead one to expect. Is this a case of Skinner contradicting himself? Is he not telling an inside story? And does he not elsewhere deny the relevance of the inside story?

We need to ask: Relevant to what? He did not dismiss the relevance of hypothetical entities or processes for all scientific purposes. Skinner (1950) noted that they are useful in other sciences. His critique of inner causes has to do with the specific question of what initiates change and determines its direction. The problem with cognitive science, he said, is that it attributes the initiation and direction of change to inner states and processes (Catania & Harnad, 1984, p. 508). For him, the crucial question was the origin of behavior (p. 608). He believed the environment initiates change and determines its direction. This conviction is consistent with a need to refer to underlying entities to *explain* how the environment can do this. But if one's most important objective is to solve behavioral problems, which was Skinner's dominant aim during much of his career, the key question is what initiates change and determines its direction. For this purpose, the inside story is, in Skinner's opinion, a distraction that leads us to search in the wrong places.

That is the basis of his "environmentalism," which does not contrast with nativism (he acknowledges that much behavior—especially animal behavior—is unlearned), nor with realism (he is willing to refer to hypothetical entities when they serve one's purpose), but with a certain form of cognitivism (he insists that, insofar as behavior is learned, change is not initiated by, and its direction is not determined by, our cognitive states or processes, but by contingencies of reinforcement). That is what he believed he discovered in the animal laboratory. This is consistent with positing hypothetical entities insofar as one's aim is to explain (as opposed to control) behavior. Indeed, (as Sober notes) one clearly needs to refer to inner states and processes to explain how contingencies of reinforcement have their effect. Skinner agreed, although he expected these inner states and processes to be physiological—or perhaps, as he suggested in his discussion of the autoclitic (Skinner, 1957), they might be "behavioral," in the sense of defined based on their function.

This is probably enough exegesis of Skinner. Let us return to the problem of the first instance. Why is it of importance? Why does it matter whether the origin of change is due to the environment by way of selection-by-consequences (as proposed by Skinner) or to principles of variation (as proposed by Staddon and Baum)? Are these not just two ways of describing the same phenomena? Who cares whether the novel variant arises by conditioning or by some principle of variation? The answer is that operant conditioning is the most versatile process in the conceptual tool kit of behavior analysis. It alone can address the problem of the first instance in a wide range of non-stereotypic situations. It therefore is the best behavioral solution to the problem posed by the ordinary creativity of human beings. And this creativity is the most important challenge to behavior analysis raised by its cognitive critics. It also is the phenomenon that leads cognitive scientists to conclude that change in behavior is initiated by, and its direction determined by, cognitive states and processes. The problem of the first instance does not lie at the fringes of this dispute, but at its center.

The target article argues that operant conditioning can solve this problem for quantitative features, but not for qualitative ones. Sober's well-crafted counterexamples have caused me all kinds of grief regarding the second half of this thesis. He attempts to show that if selection can solve the problem for quantitative features, then it can solve it for qualitative ones too. If he is right, then I either overestimated the reach of selection-by-consequences (because it cannot solve the problem of the first instance for either quantitative or qualitative features) or underestimated its reach (because it can solve the problem for both). The first possibility does not worry me. I am satisfied that I have shown selection can solve the problem for quantitative features. But the second possibility is forcing me to reconsider selection's limitations.

Sober gives two counterexamples. The first assumes that height is a quantitative feature. He imagines a range of heights divided into three categories: short, medium, and tall. Suppose we start with a population of only short and medium individuals. If height is adaptive, then selection could produce the first instance of a tall individual. But Sober stipulates that tall is a qualitative feature. Hence, if selection can cause the first instance of a quantitative feature, then it can cause the first instance of a qualitative one.

This example fails for a trivial reason. I defined qualitative features as having gaps between them that no individual could instantiate. For a population to move in a quantitative manner from medium to tall, one or more individuals would have to occupy the very gap that no individual could occupy if medium and tall are qualitative features. Hence, if height is a quantitative feature, then medium and tall are not qualitative. Conversely, if medium and tall are qualitative features, then height is not quantitative. In this second case, height would just be a measure used to determine which qualitative feature an individual has. But no individual could occupy the gap between medium and tall.

His second counterexample is the one that worries me. He imagines a mutation, B, at a locus that previously had only the AA genotype. Now it also has an AB genotype. If AB is more adaptive than AA, then selection in conjunction with sexual reproduction will tend to increase the proportion of AB individuals in the population. This in turn raises the probability that a pair of AB individuals will mate and produce the first instance of a BB individual. Assuming the phenotypes associated with these three genotypes are qualitative, his example shows that natural selection can cause a novel qualitative feature to evolve. Furthermore, if BB is more adaptive than AB, then the novel phenotype associated with BB would also mark an adaptive improvement. Does this not show that selection can solve the problem of the first instance for a qualitative feature?

To answer this question, we need a clearer understanding of the problem of the first instance. When I speak of this problem, I am referring to a pattern of reliably producing novel features that are adaptive. Merely causing the first instance of one appropriate feature is insufficient for a process to solve this problem. Instead, it must exercise control over production of the first instance. Controlling a given outcome implies causing it, but causing a given outcome does not imply controlling it. If I trip on a rock while revising Ned's entry in my phone directory, this could cause me to call Ned. But tripping over rocks does not control whom I call or when I call them. Calling Ned was an accidental effect of stumbling. Everything had to be "just so" for stumbling to cause the call. The phone had to be in my hand, it had to be open to Ned's entry in the directory, the rectangle marked "call" had to be on the screen, and my stumble had to cause me to contact the rectangle. That is an example of cause, but not of control. Control implies an ability to cause a certain type of outcome repeatedly in a variety of settings. We saw that selection is capable of conferring upon the environment control over height. Variability in height will move in whatever direction the environment selects for, and it will do so in a sustained way, not just once but repeatedly.

Skinner's deepest convictions were about environmental *control*, not simply about environmental *cause*. Sober's example shows that natural selection can *cause* the first instance of a novel qualitative feature, but selection causes the first instance of the adaptive BB genotype only because the relation between the AB genotype and the BB genotype is "just so." It is an example of causing an adaptive novel variant to occur, but not of controlling its occurrence. The first instance of the BB genotype would have been just as likely to occur if it were not adaptive. If BB happens to be adaptive, then selection can cause an adaptive variant. But the fact that it is adaptive is not part of the reason why selection causes it. It is an incidental aspect of the circumstances, similar to the fact that my directory happened to be open to Ned's entry when I stumbled over the rock. Control requires more than incidental success. Sober has thus shown that I need to make it clear that both claims of the target article are about control (in the sense of initiating change and determining its direction). The first claim should be understood to assert that selection makes it possible for the environment to control quantitative features, and the second to assert that it cannot control qualitative features.

Although Sober's example does not refute my revised second thesis, it raises the question of whether a valid counterexample exists. Perhaps I can provide one of my own. Suppose, then, that a given raptor's talons are either sharp or dull, with a gap between the two that no talon can instantiate. Therefore, sharp and dull would be qualitative features. Suppose also that its talons are either curved inward or curved outward, again with a gap between these two extremes that no individual can instantiate. These also would be qualitative features. Now suppose that these two pairs of contrasting features are determined at two different loci. The first locus has two alleles, S (sharp) and D (dull); the second locus has two other alleles, I (inward) and O (outward). The effect of S or I by itself is to make it easier to grasp prey, which would be adaptive. In combination they are even more adaptive. Would not the environment under conditions of sexual reproduction be capable of reliably producing, by means of natural selection, the first instance of the combination of these features? And would not that combination itself (sharp, inwardly curved talons) be a qualitative feature? This appears to be a valid counterexample to my revised second thesis. Combining qualitative features would form novel qualitative features that could be under environmental control. In this way, the environment could put together a complex qualitative feature that is more adaptive than any of its components in isolation. The evolution of such combinations of features could be under environmental control in the same way that the evolution of height could be. Evidently, I have underestimated the capacity of selection-by-consequences. It is capable of conferring upon the environment control over the evolution of qualitative traits.

This is relevant to behavior analysis, however, only if learning is a form of selection-by-consequences. I have been taking this assumption for granted, but this is exactly the assumption **François Tonneau** calls into question. As he insightfully remarks in his first footnote, "The question that remains at this stage is whether behavior analysts prefer models that are selectionist but not behavioral to models that are behavioral but not selectionist." The context of this quotation makes clear what he means by "at this stage." He means the stage at which it has been shown on empirical grounds that "behavioral" selectionist models are unworkable. The main point of his commentary is to argue that this is the stage at which behavior analysis currently finds itself, which is to say, that empirical results have shown that if contingencies of reinforcement select something, it cannot lie exclusively at the level of operant responses.

We saw earlier that Skinner (1953) realized that there must be units at a level below the operant and there must be a process by which they interact as a result of reinforcement. But he went on to say, "We lack adequate tools to deal with the continuity of behavior or with the interaction among operants attributable to common atomic units" (p. 95). He therefore proposed an interim solution, according to which the operant "represents a valid level of analysis" because it has properties that "may be given a functional unity." He foresaw, however, that "methods must eventually be developed which will not emphasize units at this level." These methods, however, "are not necessary to our understanding of the principal dynamic properties of behavior." Tonneau, in effect, says that behavior analysis has reached the point at which remaining exclusively at the operant level of analysis fails to advance "our understanding of the principal dynamic properties of behavior." His way of putting this is to say that a model that is selectionist and at the same time "behavioral" is empirically unsound. I have no reason to reject this conclusion, so let me provisionally accept it as true.

Skinner believed that "methods must eventually be developed" that would address "the interaction among operants attributable to common atomic units." Is behavior analysis now able to develop such methods? **Jack McDowell** and **Steven Riley** argue that it is. McDowell's (2010) Evolutionary Theory of Behavior Dynamics (ETBD) would qualify as a theory that is (in Tonneau's phrase) "selectionist but not behavioral." According to ETBD, each response is underlain by multiple gene-like entities. Reinforcement of a response replicates these entities and, through a process analogous to sexual reproduction, reshuffles them into novel combinations that can later occur. As a result of selection by contingencies of reinforcement, the probability that these subsequent novel combinations will be adaptive is increased. In this respect, ETBD is (as McDowell and Riley note) analogous to the biological theories of the Modern Synthesis.

Tonneau has used the word "behavioral" to mean absence of an inside story. ETBD is not "behavioral" in this sense. It tells an inside story. It is, however, "behavioral" in a different sense—namely, in the sense of positing underlying entities that are defined by the consequences they have the function of causing, and not by the content they represent.[2] Behavioral units defined in this way would be analogous to genes. A gene is a stretch of DNA that has a function in the causation of traits. This function might entail the production of a protein or the regulation of another gene, but a stretch of DNA with no function is not a gene. A gene does not "represent" its function or have that function as its "content." A gene for (say) sickle-shaped blood cells does not represent

[2] I am using the term "function" here in a biological sense. This is a subtly different sense of "function" than employed by behavior analysts when they say, "operant response" and "reinforcing stimulus" are given functional definitions. To see how different these uses of "function" are, consider that a reinforcing stimulus, on the so-called functional definition classically given by behavior analysts, always by definition increases the rate of the response it reinforces. But an organ that has a certain biological function can fail on a given occasion to carry out that function. Having a function in the behavior analytic sense always implies success but having a function in the biological sense does not.

sickle-shaped blood cells, nor it does have the content of 'sickle-shaped blood cell.' Instead, it has the function of being (in Tonneau's useful phrase) a "causal antecedent" of sickle-shaped blood cells. Furthermore, the biochemical process by which genes are reshuffled to create new combinations is not an inference or computation. ETBD thus provides a non-cognitive alternative to information processing explanations.[3]

The concept of a level of scientific explanation is central to what is at stake here. Skinner embraced two quite different levels of scientific explanation, both of which have important roles. The first is exemplified by a cumulative record, which expresses the control an independent variable (a contingency of reinforcement) exercises over a dependent variable (rate of response). Skinner called this a "functional analysis," using the word "functional" in a quasi-mathematical, non-biological sense, to mean that the value of the dependent variable is determined by (is a function of) the value of the independent variable. Ferster & Skinner (1957) compiled many such functional analyses. These functional analyses, however, leave a wide causal gap between independent and dependent variables. The animal must interact with the contingencies of reinforcement for an extended time for the independent variable to determine the value of the dependent variable. Explanatory completeness requires that this causal gap be filled. Ferster and Skinner never proposed to fill it completely, but they did attempt to narrow it. They did so with an account of selection-by-consequences that remained at the ontological level of the operant. They thus proposed to offer an account at this second *explanatory* level that nonetheless remained at the same *ontological* level as the functional analysis they were trying to explain. This is the explanatory strategy that Tonneau says is no longer empirically defensible. Even if their attempt to narrow the causal gap had succeeded, however, smaller causal gaps would have remained. Skinner always contended that physiology would pick up where behavior analysis leaves off and complete the explanation.

For the past half-century, mainstream philosophy of psychology has focused on cognitive strategies to narrow the causal gap between independent and dependent variables. The dominant position has been that only a cognitive story can do so. In other words, success at this second explanatory level is held to require us to posit entities at a different ontological level. The most important reason for this opinion is the assumption that only a system of formal rules operating upon inner representations has the capacity to solve the problem of the first instance. Such an account is analogous to describing the operation of a computer. One can explain a computer's operation without saying anything about the physics of its operation. This is because computational states "supervene" upon physical states, meaning that two different computers performing the same computations could be going through different sequences of physical states, whereas any two computers going through exactly the same sequence of physical states would be performing the same computations. Supervenience thus offers an ontological level that is determined by physical reality (two computers going through the same sequence of physical states will perform the same computations) but is not itself a physical description (a description of the computations being performed is not a physical description of the computer's operations). A computer thus provides a model of how to explain behavior without relying upon neurophysiological concepts to do so.

Support for such a level of ontology is provided by its ability to solve the problem of the first instance. Philosophical allies of cognitive science have argued that to solve this problem, an information processing explanation is "the only game in town." They sometimes contend that

[3] The concept of information being used in such explanations is one that requires the representation of a content. This is not the same concept of information that is used in physics, which is simply the inverse of entropy. Every time we pour milk into a cup of coffee, information in this second sense is being processed.

explanations given at this level of ontology are even better than physiological ones, because they capture broad regularities that a purely physiological explanation would miss. By analogy, we have a better understanding of what a computer is doing if we understand its program than if we know only the sequence of physical states it goes through to carry out the program.[4]

McDowell and Riley propose an alternative to the ontology of cognitivism. ETBD's explanations are nonetheless at the same explanatory level as cognitive explanations and they can solve the problem of the first instance. Indeed, ETBD's capacity to solve this problem is a plausible explanation of its apparent empirical superiority over its behavioral competition: It can produce the first instance of an adaptive response based on selection, and therefore does not have to wait for some other process to provide it. Yet, it stands well within the tradition started by Skinner, in that selection-by-consequences plays the central role in accounting for functional analyses. This means that information processing explanations are no longer "the only game in town."

I am a fan of ETBD. I do, however, have a few reservations about some of the comments McDowell and Riley make about this innovative theory. My first reservation is about the following claim by them: "*That* selection occurs is not in doubt, for reinforced behaviors in fact are observed to survive and become more plentiful in an organism's repertoire." This surprises me. Is it really that easy to demonstrate that behavior is the result of selection? Everyone can of course agree that reinforcement increases the rate of responses that are reinforced. That is true by definition. But that is not sufficient to demonstrate selection. Nor is selection demonstrated by the fact that schedules of reinforcement control rate of responding, as specified in a functional analysis.

Ferster & Skinner (1957) documented the control exercised by contingencies of reinforcement over rate of response. They also, however, proposed a "theoretical analysis" (p. 3) to account for this control. Selection-by-consequences was at the core of this account, supplemented by stimulus induction, extinction, mediating behavior, respondent conditioning, etc. All their processes were at the ontological level of the operant, and this made it possible for them to avoid an inside story, meaning their explanation was consistent with radical behaviorism. If Tonneau is right, their proposal and all others like it have failed on empirical grounds. One implication of this failure would be that Skinner was unable to show that selection explains the control that contingencies of reinforcement exercise over rate of response. McDowell and Riley propose an alternative selectionist account, thereby "improving on Skinner." They argue that ETBD offers an empirically sound explanation of the relation between independent and dependent variables, and they refer to a series of peer reviewed publications in support of this claim. The occurrence of selection-by-consequences, however, would be the *conclusion* of this complex argument, not its *premise*.

I also am surprised to find them claiming there is no tension between ETBD and cognitive psychology. I would have expected them to claim exactly the opposite, that ETBD increases the tension with the latter. Information processing theories can offer an alternative account of how the independent variable controls the dependent variable. Broadly speaking, such an account portrays the organism as constructing inner representations of its environment, performing formal operations on these representations, and then modifying its rate of response in a manner the outcome indicates would satisfy its desires. If that is what explains the relationship between independent and dependent variables in a functional analysis, then selection-by-consequences has not occurred. Information processing is the rival of selection. To the extent it explains behavior, selection does not.

[4] There are many philosophical sources for such a claim, but Dennett (1981, 1984, 1987) is a well-known proponent.

Contrast ETBD with behavioral research on “derived stimulus relations,” a topic of current interest among behavior analysts (see the June 2018 issue of *Perspectives on Behavior Science* devoted to this topic). An example is provided by the following hypothetical procedure to teach reading. First, reinforce the utterance of the word “cat” in response to a picture of a cat. Now reinforce pointing to the picture of a cat in response to the written word “cat.” Subjects have acquired a derived stimulus relation if they then respond to the written word “cat” by uttering the word “cat.” This relation between stimulus and response is “derived” in the sense that the response of uttering “cat” to the written word “cat” has never been reinforced.

How does this relate to cognitive psychology? Jan De Houwer (2018) cogently argues that research on derived stimulus relations *complements*, rather than competes with, cognitive science. What makes derived stimulus relations of special interest to cognitive science is the fact that they offer well-defined instances of novelty, the preferred target of cognitive explanation. It is a different kind of novelty, however, than that provided by the problem of the first instance. The learner already knew how to say “cat,” so the response is not novel. What is novel is uttering “cat” when stimulated by the written word, “cat.” De Houwer’s point is that cognitive explanations of this phenomenon are at a different “explanatory level” from the explanation offered by a functional analysis (pp. 231-232). The latter describes the control exercised by an independent variable (e.g., the learning procedure just described) over a dependent variable (e.g., the disposition to utter “cat” in response to the written word “cat”). A cognitive explanation, by contrast, purports to give an account of how the independent variable controls the dependent variable (e.g., by means of a system of formal operations on inner representations).

Now turn to ETBD. It is not a functional analysis, but instead it occupies the same explanatory level as a cognitive explanation. But it is not a cognitive explanation, because it makes no appeal to inner representations or information processing. ETBD therefore offers a behavioral explanation that can compete with a cognitive explanation of how an independent variable controls a dependent variable.

Radical behaviorists will object that ETBD tells an inside story. If the central claim of scientific behaviorism, however, is that mental states and processes are unnecessary for the explanation of learned behavior, then ETBD is within the scope of behaviorism. ETBD exemplifies a Neo-Skinnerian, non-radical version of behaviorism. It is not the only example of such a philosophy. Rachlin (1991), Hayes et al. (2001), Baum (2005), and Staddon (2014), have offered their own versions. What sets ETBD apart, however, is the precision with which it joins issue with the leading cognitive critique of behaviorism. The extent of its empirical success in countering this critique will depend upon its ability to expand the phenomena to which it applies, but it alone among Neo-Skinnerian, non-radical behaviorisms, addresses the problem of the first instance at the same explanatory level as information processing accounts.[5] Or so it seems to me.

ETBD has implications not only for psychology, but also for philosophy. Georges Rey (1997) is the most prominent mainstream philosopher over the past quarter-century to have taken behavior analysis seriously, devoting an entire chapter of his book to the topic. His treatment follows the broad outlines of Chomsky (1959), arguing that the significance of the behavior

[5] Howard Rachlin embraces final causes. The relational frame theory of Steven Hayes is meant to address derived stimulus relations, but it is not clear whether relational frame theory is at the same explanatory level as ETBD or if it is simply a special type of functional analysis. Hayes is interested in prediction and control but not explanation (he calls this solution-oriented set of interests, “pragmatism”), but this lack of interest in explanation makes the question of the explanatory level of his account difficult (if not impossible) to answer. Baum is the closest of the four to being a traditional radical behaviorist, but unlike Skinner, he does not think selection explains individual responses, and like Staddon, he sees the need for principles of variation. Staddon emphasizes the need for imaginative hypotheses and criticizes Skinner for limiting himself to inductive methods.

analytic program of research is that it is the most advanced attempt to explain intelligent behavior without appeal to cognitive states or processes. He argues that it fails to do so and takes this to demonstrate how utterly impossible the task is. Rey's (2006) article on "Behaviorism" for *The Encyclopedia of Philosophy* is a condensed version of this account. He writes, "Radical behaviorism is essentially the bold hypothesis that all intelligent human and animal behavior can be explained by the Law of Effect" (p. 521). He argues that this hypothesis faces a series of problems, the final one being "improvisation," which is a response "appropriate to achieve reinforcement but not previously produced" (p. 522). In other words, selection-by-consequences cannot solve the problem of the first instance. This is the proposition that I have been at pains to refute in the target article and in this reply. I consider ETBD to be an ally in this effort.

Many of the comments by **David Stahlman** and **Charles Catania** overlap with those made by previous commentators. What is distinctive is their extensive disagreement with my interpretation of Skinner. On my reading, Skinner is not a doctrinaire Skinnerian and his philosophical position is more complex and nuanced than the one usually attributed to him. I feel I have made a reasonable case for this point of view, and I am content to let it stand or fall on its own merits. One comment, however, calls for reply. Stahlman and Catania warn against confusing the Skinner of the 1930s with the mature Skinner whose work underwent a shift from being modeled on physics to being modeled on biology. I actually am in full agreement. That is why the earliest publication by him that I cite in the target article was written in 1953, when he introduced the analogy between natural selection and operant conditioning.

It is useful, in this regard, to recall the philosophical use Skinner (1953) first made of the analogy. He criticized the notion that a learned response occurs because it has a forward-looking purpose or final cause. To the contrary, he said, it occurs because similar responses had similar consequences in the past, and these past consequences selected that kind of response to recur. He suggested this is the same explanation that natural selection gives for why the spider spins its web—not because of the fly this web will catch in the future, but because of the flies that similar webs have caught in the past (pp. 87-90).

Selection is thus capable of explaining why learned behavior has an outcome that serves the interest of the organism. This explanatory potential is independent of Skinner's evolving thought. The process of selection has certain logical features of its own. What matters here is not the right interpretation of Skinner but the right interpretation of selection. Since the 1970s, the dominant opinion among academic philosophers has been that significant aspects of normal human behavior lie beyond behaviorism's reach. If the philosophical competition between behaviorism and cognitivism were a game of chess, most contemporary philosophers would say that cognitivism is in a position to force checkmate in a few simple moves. The leading reason for this opinion is the assumption that selection-by-consequences cannot solve the problem of the first instance. This also is the leading reason why cognitivists believe the origins of novel behavior could not possibly lie in the environment. This underestimates the explanatory potential of selection. Does selection therefore account for the novelty observed in human behavior? Is the origin of such behavior under the control of the environment? These are empirical questions that lie beyond the reach of philosophy—or at least of this philosopher. Demonstrating that selection can solve the problem of the first instance, however, shows that the answer to both questions might conceivably be yes.

REFERENCES

Baum, W. M. (2004a). The accidental behaviorist: A review of *The new behaviorism* by John Staddon. *Journal of the Experimental Analysis of Behavior*, *82*, 73-78.

Baum, W. M. (2004b). Responses to Staddon, Shimp, Malone, and Donahoe. *Journal of the Experimental Analysis of Behavior*, *82*, 117-120.

Baum, W. M. (2005). *Understanding behaviorism*, 2nd ed. Malden, MA: Blackwell.

Beatty, J. (2016). The creativity of natural selection? Part I: Darwin, Darwinism, and the mutationists. *Journal of the History of Biology*, *49*, 659-684.

Beatty, J. (2019). The creativity of natural selection? Part II: The synthesis and since. *Journal of the History of Biology*, *52*, 705-731.

Catania, A. C. & Harnad, S., eds. (1984). Canonical papers of B. F. Skinner. *Behavioral and Brain Sciences*, *7*, 473-724.

Chomsky, N. (1959). Review of the book *Verbal behavior*, by B. F. Skinner. *Language*, *35*, 26-58.

De Houwer, J. (2018). A functional-cognitive framework for cooperation between functional and cognitive researchers in the context of stimulus relations research. *Perspectives on Behavior Science*, *41*, 229-240.

Dennett, D. (1981). *Brainstorms*. Cambridge, MA: MIT Press.

Dennett, D. (1984). *Elbow room*. Cambridge, MA: MIT Press.

Dennett, D. (1987). *The intentional stance*. Cambridge, MA: MIT Press.

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

Hayes, S. C., Barnes-Holmes, D., & Roche, B. (2001). *Relational frame theory: A post-Skinnerian account of human language and cognition.* New York: Plenum Press.

McDowell, J. J (2010). Behavioral and neural Darwinism: Selectionist function and mechanism in adaptive behavioral dynamics. *Behavioural Processes*, *84*, 358-365.

Rachlin, H. (1991). *Introduction to modern behaviorism*, 3rd ed. New York: W. H. Freeman.

Rey, G. (1997). *Contemporary philosophy of mind: A contentiously classical approach.* Oxford, UK: Blackwell.

Rey, G. (2006). Behaviorism. In D. M. Borchert (Ed.), *The encyclopedia of philosophy*, 2nd ed. (Vol. 1, pp. 520-526). Farmington Hills, MI: Thomson Gale.

Skinner, B. F. (1950). Are theories of learning necessary? *Psychological Review*, *57*, 193-216.

Skinner, B. F. (1953). *Science and human behavior.* New York: Macmillan.

Skinner, B. F. (1957). *Verbal behavior*. Englewood Cliffs, NJ: Prentice-Hall.

Skinner, B. F. (1975). The shaping of phylogenic behavior. *Journal of the Experimental Analysis of Behavior*, *24*, 117-120.

Smith, T. L. (2019). Selection by consequences in the ontogeny of behavior: The problem of the first instance. *Behavior and Philosophy*, *47*, 1-14.

Staddon, J. E. R. (2004a). The old behaviorism: A response to William Baum's review of *The new behaviorism*. *Journal of the Experimental Analysis of Behavior, 82*, 79-83.

Staddon, J. E. R. (2004b). Responses to commentators. *Journal of the Experimental Analysis of Behavior*, *82*, 121-124.

Staddon, J. (2014). *The new behaviorism*, 2nd ed. New York: Psychology Press.

Staddon, J. & Simmelhag, V. (1971). The superstition experiment: A reexamination of its implications for the principles of adaptive behavior. *Psychological Review*, *57*, 3-43.

Wolfson, A. (1948). Bird migration and the concept of continental drift. *Science*, *108*, 23-30.