




Contents lists available at ScienceDirect

Behavioural Processes

journal homepage: www.elsevier.com/locate/behavprocSelection by reinforcement: A critical reappraisal[☆]

José E. Burgos

University of Guadalajara, CEIC, 180 Fco. De Quevedo, Arcos Vallarta, Guadalajara, Jalisco 44130, Mexico

ARTICLE INFO

Keywords:

Selection by reinforcement
Pavlovian conditioning
Neural networks
Genetic algorithms
Metaphysics

ABSTRACT

This essay is a critical reappraisal of the idea of ontogenetic selection by reinforcement, according to which learning, specifically conditioning, in the individual animal is deeply analogous to phylogenetic evolution by natural selection. I focus on two general versions of this idea. The traditional Skinnerian version restricts the idea to operant conditioning and excludes Pavlovian conditioning, based on a sharp dichotomy between the two types of conditioning. The other version extends the idea to Pavlovian conditioning, based on a unified principle of reinforcement that applies to both types of conditioning, and linked to a neural-network model. I criticize both versions on the same grounds, for being: 1) unable to capture Pavlovian conditioning; 2) unnecessary to formulate said model and use it for explanation and prediction (its combination with a genetic algorithm allows for a substantive contact with the theory of evolution by selection, without the idea of selection by reinforcement), and 3) metaphysically unsound. Non-selectionist accounts of conditioning are not only possible but also more intelligible, explanatory, and heuristic.

1. Introduction

This essay is a critique of the idea of ontogenetic selection by reinforcement (*SbR*, for short), which asserts that certain forms of learning at the individual (ontogenetic) level are fundamentally equivalent to phylogenetic evolution by natural selection. More precisely, *SbR* postulates that conditioning is essentially like phylogenetic evolution in that both consist of cycles of variation, selection, and retention of the products of selection.

Using ethological techniques to record behavior, Staddon and Simmelhag (1971) proposed an early version of *SbR* in their reexamination of the phenomenon of “superstition” in hungry pigeons that were given brief access to food as reinforcement (periodic response-independent and response-dependent, as well as aperiodic response-independent). Pigeons were observed to engage in two kinds of activities in the absence of food. Interim activities tended to occur just after the food, and be more variable and non-food-related. Terminal activities tended to occur just before, seemingly in anticipation of, the food, and be more stereotypical and food-related.

The authors explained their observations by postulating a “close analogy between the Law of Effect and evolution” (p. 17), where “the Darwinian principle of selection is analogous to the process that transforms initial behavior into final behavior” (p. 18). In this analogy, reinforcement “selectively eliminates” (p. 17) responses (or behaviors or activities) that are inappropriate and leaves responses that are

“appropriate” to it (see also Staddon, 2016). Skinner (e.g., 1981) restricted his version to selection by *consequences* of emitted responding in operant conditioning. This version has been the most influential in behavior analysis.

The analogy was also used to introduce a neural-network model of conditioning (Donahoe et al., 1993). This model is a connectionist interpretation of a unified principle of reinforcement (Donahoe et al., 1982) in operant and Pavlovian conditioning (I thus abbreviate the model as UPR-NNM, for “unified principle of reinforcement neural network model”). Networks described by UPR-NNM have often been labeled “selection networks.” This version of the idea of selection by reinforcement was an attempt to extend the Skinnerian version beyond a purely behavioral selectionist account of operant conditioning.

It will thus be helpful to distinguish at the outset the two versions of *SbR* I shall discuss. I shall call the Skinnerian version “*T_{SbR}*” for “traditional (*T*) version of the idea of *SbR*,” and the version that accompanied UPR-NNM “*E_{SbR}*” for “extended (*E*) version of the idea of *SbR*.” When speaking of both versions in general, I shall call them “*SbR*” (for “the idea of ontogenetic selection by reinforcement”).

Despite initial work with UPR-NNM under *E_{SbR}*, I have grown wary of *E_{SbR}*. I explain my reasons in 3, 4, and 5, after a brief reminder of some basic distinctions in 2 that I shall use in my analysis. Like other critiques (see Tonneau and Sokolowski, 2000) of *SbR*, mine is conceptual but goes in a different direction, adding to them. It should be clear that none of these criticisms extend to the theory of evolution by

[☆] This paper is based on a presentation as a discussant in a symposium on behavioral evolution and selection by consequences, at the 42nd Annual Convention of the Association for Behavior Analysis International, Chicago, May 2016.

E-mail address: jburos@cucba.udg.mx.

<https://doi.org/10.1016/j.beproc.2018.01.019>

Received 24 July 2017; Received in revised form 19 January 2018; Accepted 30 January 2018
0376-6357/ © 2018 Elsevier B.V. All rights reserved.

natural selection. This theory does not entail *SbR*, so negating *SbR* does not negate the theory. Nor do I intend to question the notion that learning by reinforcement affects phylogenetic evolution by natural selection. Animals can be selected to reproduce *for* learning certain behaviors by reinforcement, which could be called “*phylogenetic* selection by reinforcement,” but this is not what *SbR* asserts. As I understand it, *SbR* asserts that *behaviors themselves too* are selected by reinforcement at the ontogenetic level. In *SbR*, this selection is supposed to be of the same type as that in which, say, individual organisms are selected to reproduce for having certain phenotypic traits (including certain behaviors learned by reinforcement) during phylogeny. It is this assumption what I shall examine critically.

2. Brief reminder of some distinctions

I will use the following distinctions, in no particular order of importance: Structure/function, genotype/phenotype, ontogeny/phylogeny, and selection-of/selection-for. They have been instrumental in the elaboration of the theory of evolution by natural selection beyond Darwin's (1859) initial formulation, and remain central in current evolutionary biology and its philosophy. I make them here as they are commonly made in the literature, where they are tightly related to one another and their categories are not opposite but complementary.

2.1. Structure/function

This distinction is perhaps the most general one, as it applies to all the other distinctions. Generally, a structure occupies space (has spatial parts) and persists through some time with some of its parts and features relatively unchanged (and thus give structures their identity through time; I return to this notion in 5). A function is how a structure *changes* some of its other parts and features without losing its identity through time. Some such changes can in turn change, even generate, other structures.

Certain biological structures constitute an organism's *anatomy* (the organism's spatial parts and how they relate to one another spatially). Certain biological functions constitute an organism's *physiology*, the functioning of certain anatomical structures, how some of their properties change through time and how they affect one another. The term “function” also often refers in biology to *adaptive* value of some feature, the extent to which the feature contributes to an individual organism's fitness or reproductive success.

The structure/function distinction is made at different levels of biological organization. All molecules, cells, tissues, organs, systems, and organisms are structures (although not all are anatomical, strictly speaking) that have many different functions. At all these levels, structure and function are *complementary* and tightly linked. Function is inconceivable without structure (whether the vice versa is also true might be debatable). All function is *of* some structure. Understanding anatomy is thus *integral* to understanding physiology, and vice versa. Biologists tend to specialize in one or the other, but no biologist would seriously say that one is ontologically or epistemically more fundamental than the other.

2.2. Genotype/phenotype

A genotype is an *individual* organism's genetic constitution, its DNA molecules consisting of nucleotides that determine the synthesis of certain proteins key to a cell's structure and functioning (not to be confused with the genome, the genetic constitution of a *population*). Genotypes are molecular means of transmission, from parents to offspring, of phenotypic traits that favor reproduction. Genes are molecular units of heredity transferred from parents to offspring, and replicated during the offspring's development. A gene is a specific nucleotide sequence (a locus or region of DNA) that determines the synthesis of a certain kind of protein and/or controls the functioning of

other genes. Genotypes are thus the *structural* means of *genetic* inheritance of the products of natural selection.

Although genotypes themselves are products of selection, it is common to view natural selection as acting more directly on certain phenotypic traits that, in turn, depend on certain proteins synthesized from the genotype. A phenotype is the set of an individual organism's myriad properties, traits, features, or characteristics that result from interactions between its genotype (more directly, certain phenotypic traits that depend on certain proteins) and its environment, whether or not the traits are “observable”. Phenotypic traits can be structural (e.g., anatomical) or functional (e.g., physiological). Certain kinds of organs consisting of certain kinds of cells consisting of certain kinds of molecules are examples of structural phenotypic traits. Their functioning are functional phenotypic traits. Behavior is a functional phenotypic trait.

In evolutionary biology, genotypic and phenotypic variation are both considered as necessary for evolution by natural selection. This necessity is often expressed metaphorically by saying that variation is “the raw material” for (or of) natural selection. Phenotypic variation (structural and functional) is due to complex interactions between genotypes and environments. Sources of genotypic variation include mutation and recombination.

2.3. Ontogeny/phylogeny

Ontogeny comprises all processes that occur during a *particular*, *individual* organism's *lifetime* (from its conception to its death), including its biological development (even before birth), learning and behavior, and all its physiological processes at all levels of organization (molecular, e.g., protein synthesis, neurotransmitter synthesis and release; cellular, e.g., action potentials; and anatomical, e.g., tissue regeneration, neuronal death, etc.). “Ontogenetic biology” is the name of the subfield in biology that focuses on ontogenetic processes. Ontogeny need not be restricted to a single individual organism: Processes in groups of individual organisms within a population could also qualify as ontogenetic. Biological development, as well as learning and behavior in a group of animals, also qualify as ontogenetic. A more precise characterization of most experimental studies of animal learning and behavior (whether with single subjects or groups of subjects) as being part of biology is that they belong in ontogenetic biology. This characterization does not preclude talk of the “evolution” of learning, but such talk is part of *phylogenetic* biology, which focuses on phylogeny. Research where learning is observed across several generations of individual organisms (e.g., Dunlap and Stephens, 2014; Tolman, 1924; Tryon, 1940) belongs in this category.

Phylogeny consists of all processes that transcend individual organisms (and, in a sense, particular groups of organisms) and thus occur at longer scales, over generational time across populations of organisms that constitute lineages or ancestor-descendant evolutionary lines. Generations overlap, but to a point after which later generations go on beyond previous ones. The standard thinking in biology is that evolution by natural selection with genetic inheritance is a phylogenetic process. In this thinking, and strictly speaking, talk of “evolution by natural selection” in reference to an individual organism is a biological category mistake. Strictly speaking, individual organisms do not evolve but reproduce, develop, learn, and behave in certain ways.

2.4. Selection-of/selection-for

This distinction was introduced by Sober (1984, pp. 97–102) to distinguish between the *objects* of selection (selection-of) and the *properties* for which objects are selected (selection-for).¹ In 5.2, I will

¹ Sober identifies selection-of with the effects and selection-for with the causes of selection. The topic of causality, however, raises difficult issues I cannot properly discuss in this paper. I will thus leave the topic aside.

discuss what the term “object” could mean in this distinction metaphysically. For now, I will only say that Sober (1984) speaks mostly of “selection of objects,” and “selection for properties,” very rarely if ever of “selection of properties” or “selection for objects.” I thus assume that selection of objects is, well, only about objects, and selection-for is only about properties. Sober confirmed this interpretation: “I think I may on occasion have talked about selection of properties, but my main point has been to distinguish the selection of objects from selection for properties” (E. Sober, personal communication, December 29, 2017). I thus assume that he excludes “selection of properties” and “selection for objects”.²

Sober (1984, p. 99, Fig. 2) illustrates with a simple analogy to a “selection toy,” a closed container with levels that function as sieves. Each level has holes of equal size that varies across levels (large, medium, small). The container also encloses balls of different sizes, say, large, medium, and small, perfectly correlated with different colors, say, blue, red, and green, respectively (all and only the small balls are green). The large balls are larger than all the holes, whereas the medium balls are smaller than the larger holes but larger than the small holes. The small balls are smaller than all the holes. If the container is positioned to make the level with the smallest holes the lowest and shaken, only the small balls sift through all the levels and eventually end at the bottom.

In this toy, selection is the separation of balls according to their sizes by the sieves, with the outcome being the small balls ending at the bottom. Some balls can thus be said to “be selected” for their smallness, but this does not mean that smallness is an object of selection. *Smallness is a property, not an object* (I return to this in 5). The selected balls are also green, but they were not selected for their greenness, nor was greenness an object of selection.

Obviously, the process in the toy is just an analogy not to be taken literally. Balls are not organisms (or genes, or groups of organisms), nor is any ball selected in the biological sense that organisms are. Biologically, selection refers to phenotypic-dependent differential reproduction in a population of organisms in some environment. In the simplest case, one organism reproduces (leaves viable offspring) more (at least once) for having certain heritable (genotypic-dependent) phenotypic traits under certain environmental conditions, than at least one other organism that reproduces less (never, in the extreme case) for lacking such traits.

Darwin (1859), as is well known, thought of the term “selection” inspired by the farmers’ practice of selective breeding, where they choose some individual organisms over others to reproduce for having certain traits of interest. The difference between this “artificial selection,” as Darwin called it, and natural selection is that the latter is done by Nature and, hence, mindlessly (in contrast to the breeders, Nature does not “choose”). Of course, all of this occurs in populations, which lends itself to statistical accounts of changes in frequencies of phenotypic traits and genes across generations. Perhaps except for the statistical part (one can count balls of different kinds at different levels in the toy), nothing else in biological selection (especially the phenotype, genotype, and reproduction parts) applies to the balls in the toy.

Still, Sober’s analogy can help clarify what is involved in biological selection. The objects of selection (selection-of, what is selected) in the toy are balls. They correspond to the entities biologists have considered as objects of natural selection, such as individual organisms, genes, or groups (the dominant view in biology is that such entities are individual organisms; see Mayr (1997), pp. 2091–2092). The smallness of the toy’s

balls corresponds to a phenotypic trait (possibly but not necessarily smallness) that increases fitness qua reproductive success in biological evolution by natural selection. Greenness corresponds to a phenotypic trait that is neutral in that it does not affect fitness, but is still selected because it is highly correlated with traits that increase fitness (the panadaptationist would reject such neutral trait).

In a more natural example, suppose an ancestor of giraffes’, say, a specimen of *Canthumeryx* (estimated to have lived about 16 million years ago) was selected for having a longer neck. This specimen would correspond to a particular ball in the selection toy, whereas a certain neck length corresponds to a ball’s smallness. Greenness in the selection toy might correspond to, say, the ossicones’ shape (although the panadaptationist can concoct an adaptation story for any phenotypic trait).

As any other distinction, Sober’s is neither perfect nor universally accepted, but has been sufficiently influential among biologists and philosophers of biology to be taken seriously. I will use it throughout the paper, especially in 5 to articulate a metaphysical concern about *SbR*. Before this, I discuss other, more epistemological concerns. I do not intend them to be definitive, only as issues that I believe warrant some discussion.

3. Pavlovian conditioning

In Pavlovian conditioning, the “reinforcer” is an unconditioned stimulus (US) that normally elicits an unconditioned response (UR) without any training. The US occurs in an optimal (temporal and statistical) relation with a conditioned stimulus (CS), which does not normally evoke the UR. After several CS-US trials, the CS comes to evoke, typically, a similar response (CR). The US need not depend on any behavior, a feature many view as key to the Pavlovian-operant distinction. Another feature is that the reinforcer in Pavlovian conditioning initially and demonstrably elicits the response of interest unconditionally (in fact, this response is often defined by whatever response the US elicits), whereas in operant conditioning it is not.

It seems clear that T_{SbR} , with its exclusive emphasis on operant conditioning (where responding is emitted and reinforcement depends on some response), does not capture Pavlovian conditioning. Supporters of T_{SbR} do not even seem to care about Pavlovian conditioning. Indicative of this is the total absence of talk of Pavlovian conditioning in all formulations of T_{SbR} . As Baum (2017) put it, “...operant behavior [is] the main focus of behavioral evolution” (2017, p. 322).³

Also symptomatic is the tendency among supporters of T_{SbR} to speak more specifically of “selection by consequences,” in reference of “consequences of responses” (i.e., operant reinforcement), rather than “selection by reinforcement” (which can include Pavlovian reinforcement). Most research in Skinnerian behavior analysis, the home of T_{SbR} , remains strongly guided by a sharp dichotomy between Pavlovian and operant conditioning, with an interest only in operant conditioning. A major reason for this separation is that Skinnerians still strongly repudiate a key aspect of current Pavlovian conditioning research: The postulation of nonbehavioral (typically cognitive) entities to explain behavior.

T_{SbR} ’s exclusion of Pavlovian conditioning is problematic because the evidence strongly suggests that both types of conditioning interact

² He talks about “selection for jaws” (pp. 24, 98). If jaws are objects, my understanding that there is no selection for objects is incorrect. However, the way he discusses the distinction through the rest of his book strongly suggests that such talk is an abbreviation of “features of jaw structure” (p. 24, emphasis mine). In fact, he seems to view jaws as properties, not objects: “If jaws and chins were different characteristics” (p. 24, emphasis mine). In a personal communication about this, he replied “Yes, I meant selection for having a jaw” (E. Sober, December 29, 2017).

³ Baum (2012) has proposed to account for operant and Pavlovian conditioning in terms of the notion of induction, and cast this notion in behavioral selectionism using a behavioral interpretation of the Price equation (Baum, 2017). However, it is unclear to me whether and exactly how induction differentially explains specific Pavlovian phenomena such as interstimulus interval functions, overshadowing, blocking, second-order conditioning, and latent inhibition, among others. Equally unclear is whether and how Pavlovian conditioning fits in Baum’s (2017) use of the Price equation, which focuses on operant conditioning. He does not illustrate how the equation applies specifically to any of these phenomena.

in significant ways. This interaction has been observed in misbehavior (Breland and Breland, 1961), superstitious conditioning (Skinner, 1948), Pavlovian-instrumental transfer (e.g., Estes and Skinner, 1941; Walker, 1942), reinforcement devaluation (Colwill and Rescorla, 1990), autoshaping (Brown and Jenkins, 1968), and cue-to-consequence effects (Foree and LoLordo, 1973; Schindler and Weiss, 1982), among others (see Davis and Hurwitz (1977), for more on operant-Pavlovian interactions). It is far from clear whether and how T_{SbR} can account for these phenomena.

E_{SbR} , in contrast, was partly an attempt to extend SbR to Pavlovian conditioning. In this section, I focus on whether E_{SbR} captures Pavlovian conditioning *behaviorally*. E_{SbR} also has a strong neural import related to UPR-NM. This import raises the possibility that certain neural entities (synaptic efficacies) could be objects of SbR in E_{SbR} , but I will discuss this possibility in 4.

3.1. Environment-behavior relations as objects of selection

Key to E_{SbR} 's inclusion of Pavlovian conditioning in SbR was the assumption that, behaviorally, reinforcement did not select just responses (or behaviors, or activities) but “environment-behavior relations” (e.g., see Donahoe, 2013; Donahoe et al., 1997a, p. 194, b). A major point of contention (among others) about the validity of this assumption in operant conditioning (e.g., see Shull, 1995) arose from the behavior-analytic characterization of operant responses as being “emitted,” “free,” or “spontaneous.”

These terms only mean that experimental identification of controlling antecedent stimulation of operant responding is not necessary to obtain reliable (“ordered”) data in operant conditioning. Pavlovian conditioning, in contrast, requires (in fact, is defined by) the explicit, precise scheduling of antecedent stimuli (the CS and US) that evoke the response of interest. There are discriminated operants, which are said to be “controlled” by explicitly scheduled antecedent stimuli, but these stimuli cannot be treated as CSs in the traditional sense in Pavlovian conditioning research.

What could the controlling antecedent stimulation be in free operant responding? The answer had been in the air for a while before UPR-NM was formulated in terms of E_{SbR} , although not necessarily in these terms: *Background* or *contextual cues*. These cues are not stimuli in the traditional sense that CSs and USs are, that is to say, discrete, punctuate, or *phasic* environmental occurrences explicitly scheduled by the experimenter. Contextual cues, in contrast to standard CSs, USs, and discriminative stimuli, are more “static” (for the phasic-static distinction, see Balsam (1985), p. 9). That is to say, they remain unchanged for relatively longer periods than the CSs and discriminative stimuli (although the latter tend to last longer than CSs and, thus, approximate more closely contextual cues). More precisely, contextual cues *spatio-temporally include* all scheduled phasic stimuli (CSs, discriminative stimuli, primary reinforcers, USs) within sessions. The static features of an operant conditioning chamber, such as the walls, floor, feeder opening, and, crucially, operandum (e.g., a lever, a key, etc.), are contextual cues in that sense.

There has been extensive animal research of the role of contextual cues thus conceived (typically manipulated as a qualitative independent variable general features of the experimental environment, e.g., the color and pattern of its walls, are changed across experimental conditions). This research shows that the context exerts a strong influence on both, operant and Pavlovian conditioning (e.g., see Balsam and Tomie, 1985; Dickinson et al., 1996; Gould and Bevins, 2012).

E_{SbR} thus hypothesized that *implicit* (experimentally nonscheduled) context-reinforcer relations reliably accompanied explicit response-reinforcer contingencies in “free” operant conditioning. These relations allowed certain contextual cues, especially tactile, proprioceptive, and visual stimulation from the operandum, to exert substantial antecedent stimulus control over emitted responding. None of this meant that operant conditioning was identical Pavlovian conditioning. Strictly,

implicit context-reinforcer contingencies in operant contingencies are not Pavlovian, because, again, contextual cues are not CSs in the traditional sense of the term, and the reinforcer still depends on a response, not a stimulus.

To call the resulting control relations “S-R” was equally inaccurate, as antecedent stimuli (S), again, have been traditionally defined as explicit (experimentally scheduled) phasic events, whereas contextual cues are implicit and more static. They were thus called broadly “environment-behavior” relations, to include “context-response” relations. Contextual cues are exteroceptive, sensory stimuli, but then again not in the traditional *operational* sense that CSs, USs, or even discriminative stimuli are. Contextual cues are stimuli *functionally* speaking, in that they are reliable antecedent environmental conditions that control responding, under operant and Pavlovian contingencies.

Terminology aside, E_{SbR} was partly inspired by a unified principle of reinforcement (Donahoe et al., 1982). According to this principle, whenever a reinforcer causes a certain stimulation and responding discrepancy (between absence in one moment and presence in the next), all other cues acquire some control over all responses present at that moment. Suppose a rat's barpress response is immediately followed by food. The occurrence of the food causes a stimulus and response discrepancy (the food, and the UR that it presumably elicits, were absent immediately before, and now are present). This discrepancy strengthens the control of all responses, including the barpress response and the UR to the food, by all other stimuli present when the food occurs. In free operant responding, only contextual cues are present, so they are primary candidates for antecedent controlling stimulation.

E_{SbR} arose from the additional assumption that the environment-behavior relations that occur at the moment of reinforcement in an individual animal are of the same type as the objects of natural selection in phylogenetic evolution (e.g., individual organisms). Environment-behavior relations are thus supposed to be “selected” by reinforcement (whether Pavlovian or operant) in the same sense that, say, individual organisms are selected for having certain phenotypic traits to reproduce during phylogenetic evolution.

3.2. Does E_{SbR} capture Pavlovian conditioning behaviorally?

It is far from obvious how this question can be answered in the affirmative. One reason has to do with the nature of the objects of behavioral selection in Pavlovian conditioning, and the nature of the behavioral variation that presumably is necessary for such selection, according to E_{SbR} . In terms of Sober's (1984) distinction, it is unclear what is selected, or what selection-of is about *behaviorally*, in Pavlovian conditioning, according to E_{SbR} .

Again, according to E_{SbR} , the behavioral objects of selection in operant and Pavlovian conditioning are environment-behavior relations, not just behaviors. But this is too broad: Specifically what are the selected environment-behavior relations in Pavlovian conditioning *as studied experimentally* in animals, with the explicit presentation of CS-US trials? The relations can only be CS-UR and US-UR, as only they occur before and during Pavlovian conditioning and can thus serve as behavioral objects of selection by Pavlovian reinforcement. CS-CR relations do not count, as they occur *after* Pavlovian conditioning, when presumably SbR has already occurred.

Donahoe and Vegas (2004) reported evidence that the critical relation in Pavlovian conditioning is not CS-US, as traditionally believed, but CS-UR, a kind of environment-behavior relation. I thus take it that CS-UR relations are the preferred candidates for behavioral objects of selection by Pavlovian reinforcement, according to E_{SbR} . However, in typical experimental demonstrations of Pavlovian conditioning in the laboratory, the CS does not vary qualitatively (in its sensory modality) or quantitatively (in its duration, intensity, light wavelength if visual, tone frequency if auditory, etc.). If variation is necessary for selection, it is unclear that the CS provides the variation that according to E_{SbR} is supposed to be necessary for the selection of CS-UR relations by

Pavlovian reinforcement.⁴

More crisply, if variation is necessary for selection in general, there can be no selection, and hence no evolution, without variation. By analogy, if according to E_{SbR} CS variation is necessary for ontogenetic selection by Pavlovian reinforcement, and if this selection is necessary for Pavlovian conditioning to occur, there should be no (or very weak) Pavlovian conditioning without CS variation. However, and this is the problem, significant Pavlovian conditioning occurs without any CS variation in the laboratory. Therefore, either there is no selection by Pavlovian reinforcement, or there is but it is unnecessary for Pavlovian conditioning, or CS variation is not necessary for such selection. All these possibilities weaken E_{SbR} .

There is variation in the UR, qualitative (e.g., Jenkins et al., 1978; Zener, 1937) as well as quantitative (e.g., UR amplitude and latency vary across US occurrences). But it is unclear whether it suffices for selection by Pavlovian reinforcement, according to E_{SbR} . If CS-UR relations are the objects of behavioral selection in Pavlovian conditioning, and if variation is necessary in *both* relata, UR variation per se would not suffice.

Quite likely, there is much CS variation in *natural* settings, but this does not solve the problem. Pavlovian conditioning in the laboratory, as artificial and simpler as it might be, shows that Pavlovian conditioning can occur without CS variation, and this suffices to conclude that CS variation is not necessary for Pavlovian conditioning. If variation in the CS is not necessary, then to postulate CS-UR *relations* as objects of selection by Pavlovian reinforcement is unwarranted. Postulating the UR as the only object of such selection should presumably suffice, but this would eliminate a key explanatory resource from E_{SbR} : Antecedent control stimulation.

3.2.1. What are CS-UR relations selected for in Pavlovian conditioning?

A related issue refers to what CS-UR relations are selected for (in the sense of Sober's notion) in Pavlovian conditioning, according to E_{SbR} . They can only be selected for their properties: CS modality, duration, and intensity; UR form, duration, and magnitude; the CS-UR temporal relation. It is far from clear which of these properties CS-UR relations are supposedly selected for, and why. It is not clear how, or even whether the matter can be resolved empirically. What kind of observation would provide unequivocal evidence that CS-UR relations are selected for all or only some of these properties, and if only some, which ones?

A similar issue afflicts selection by consequences in operant reinforcement. Behaviors in operant conditioning are defined by certain topographical (form, duration, magnitude) and functional properties (producing a reinforcer, perhaps in the presence of a discriminative stimulus). Exactly which of these properties are behaviors in operant conditioning selected for, and why? All of them? What kind of evidence would support answers to these questions empirically?

In sum, although it seems clear that T_{SbR} does not capture Pavlovian conditioning, it is unclear that E_{SbR} , which was propounded to include Pavlovian conditioning, does, at least behaviorally. I am not claiming that E_{SbR} cannot possibly capture Pavlovian conditioning. I am only claiming that it is far from clear exactly how it does, especially regarding the objects of selection (what is selected) in Pavlovian conditioning, and the supposed necessity of CS variation for selection by Pavlovian reinforcement.

4. Computational modeling

As mentioned in the introduction, UPR-NNM was formulated as a unified connectionist (formal) interpretation of Pavlovian and operant

conditioning (Donahoe et al., 1993). The interpretation was originally cast in terms of E_{SbR} , and remained central to simulation research work with UPR-NNM for years. Although I discussed E_{SbR} 's behavioral aspect separately in the previous section, E_{SbR} has always been closely related to UPR-NNM. This relation also gives E_{SbR} a strong neural import, which I discuss in this section, as formulated in UPR-NNM.

Theorizing in terms of brain structure and function has been central to E_{SbR} , another key difference with T_{SbR} (the other is E_{SbR} 's inclusion of Pavlovian conditioning). In this section, however, I raise two further concerns. One, what are the *neural* objects of selection according to UPR-NNM's interpretation of E_{SbR} ? Two, is E_{SbR} necessary to formulate and use UPR-NNM? I will dedicate most of this section to the second concern.

About the first concern, I will only say that it remains unclear to me what could the neural objects of SbR be according to UPR-NNM's interpretation of E_{SbR} , in relation to the supposed necessity of variation for SbR . One candidate is synaptic efficacies, simulated in UPR-NNM as connection weights. The assumption would be that if synaptic efficacies are the neural objects of SbR according to E_{SbR} , then variation among synaptic efficacies previous to conditioning should be necessary for SbR and, to this extent, conditioning to occur.

I do not deny the very likely possibility that there is much variation in an animal's relevant synaptic efficacies before conditioning in the laboratory. However, this possibility per se, even if it is a fact, does not imply that such variation is *necessary* for conditioning. Compelling evidence for this necessity would be significantly weaker conditioning without such variation. As far as I know, no such evidence is available. It is not even clear whether and how the key experimental manipulation (making all relevant synaptic efficacies as similar as possible) could be done.

Additionally, Burgos (2001) showed that neural networks in UPR-NNM can simulate strong conditioning (Pavlovian as well as operant, or at least a form of autoshaping) with all initial connection weights set to the same value, with no variation across them. UPR-NNM thus predicts strong conditioning in the absence of any initial variation in synaptic efficacies. In this way, the model predicts, contrary to a key assumption of SbR , that initial variation is not necessary for conditioning. Of course, the model could be wrong, but compelling evidence is needed in order to reach this conclusion. My point is that this prediction is at odds with E_{SbR} , which weakens its association with UPR-NNM.

4.1. Is E_{SbR} necessary to formulate and use UPR-NNM?

Another problem is that UPR-NNM can be formulated and used without any reference whatsoever to E_{SbR} . UPR-NNM can be formulated purely in terms of its activation rule, learning rule, and network architecture, linked to neuroscientific principles (suitably simplified for modeling purposes), and how they relate to behavior and the environment. E_{SbR} does not help formulate any of these components in any clearer, more precise way. Quite the contrary, adding E_{SbR} makes the model's formulation unnecessarily cumbersome and more speculative than it already is.

Nor is E_{SbR} necessary to use UPR-NNM for explanation and prediction purposes. All simulations of conditioning phenomena with UPR-NNM can be explained purely in terms of different dynamics of activations and weights in a network architecture trained with protocols that simulate conditioning procedures of interest. None of this calls for E_{SbR} . Quite the contrary, adding E_{SbR} makes explanations and predictions with the model unnecessarily less parsimonious. Parsimony should not be pursued at the expense of explanatory and predictive power, but thus far E_{SbR} has not increased the model's explanatory or predictive power in any way.

None of this is surprising. After all, UPR-NNM is a model. As such, it leaves much out, as any model. Input activations in the model simulate primary-sensory effects of exteroceptive stimuli, not stimuli per se as usually defined in conditioning research (e.g., electromagnetic

⁴ It might be assumed that there is variation in how animals "perceive" and "attend to" the CS. However, much hinges on how perception and attention to the CS are conceived. If conceived neurally (as they would likely be in E_{SbR}), they would be *neural*, whereas my focus here is on *behavioral* objects of selection.

radiation of a certain wavelength, magnitude, and duration). Output activations simulate primary-motor precursors of responding, not responses (or patterns of activity extended in time) per se. Therefore, the model's relevance to behavior is conjectural. Primary-sensory activity is hypothesized to be a reasonably reliable indicator of the occurrence of stimuli, and primary-motor activity to be a reasonably reliable indicator of responding (or behaving).

The model can simulate Pavlovian conditioning with constant CS and US input, and UR output activations, but this only suggests that variation in primary-sensory effects of the CS and primary-motor precursors of the UR (and CR) is not necessary for SbR (another novel prediction pending empirical test). But the model excludes everything before such effects and after such precursors. It thus says nothing about whether it could simulate Pavlovian conditioning without assuming any variation in the CS-UR relation, variation that I take E_{SbR} to postulate as necessary for Pavlovian conditioning.

It thus is not possible to determine whether UPR-NNM satisfies E_{SbR} , but this determination is not required to formulate and use the model. Whether there actually is variation in CS-UR relations, and whether it is necessary for Pavlovian conditioning, is irrelevant to UPR-NNM. Equally irrelevant is the issue of why CS-UR relations are selected for in Pavlovian conditioning. None of these issues, central to E_{SbR} , are raised in the model, nor do they need to be addressed to formulate and use it.

Overall, then, the model does not call for E_{SbR} , which does not necessarily make E_{SbR} wrong, of course, but this is not my point. My point is that using and formulating UPR-NNM does not need E_{SbR} , and that this does not necessarily make the model wrong. The model can simulate a wide range of Pavlovian conditioning phenomena without being cast in terms of E_{SbR} . Forcing E_{SbR} into the model makes its formulation and use unnecessarily intricate. The model stands on its own.

4.1.1. Going Evo-Devo without SbR

A major motivation for SbR is to make a substantive contact between conditioning theory and the theory of evolution by natural selection with genetic inheritance (the synthetic theory, for short). UPR-NNM per se does not make such contact, as this model is only about some of the neural substrates of conditioning during ontogeny, in individual networks that are handcrafted and remain fixed during training (i.e., do not develop). The handcrafted networks are supposed to simulate neural circuits in adult brains that result from evolutionary and developmental processes, but UPR-NNM says nothing about such processes. Can UPR-NNM make contact these processes without E_{SbR} ? Yes, it can. Commitment to E_{SbR} is not the only or even best way for conditioning research to make substantive contact with the synthetic theory.

Burgos (1997) developed a computational approach to the evolution of conditioning by combining UPR-NNM with a *genetic algorithm* (GA), to simulate the evolution of Pavlovian conditioning in neural networks. A detailed description of this approach is beyond the scope of this essay. I will thus only summarize it briefly, just to give a general idea of how it works. This idea will suffice to make my point that SbR is not just unnecessary for making a substantive contact between conditioning theory and the synthetic theory of evolution (by natural selection with genetic inheritance). It also makes this contact needlessly redundant and muddled.

A GA is a class of computational models of evolution by selection with genetic inheritance. Invented by Holland (1975), they are widely used to solve search and optimization problems, but can also be used to theorize about the evolution of learning and behavior. The main components of a GA are a population of virtual chromosomes or genotypes (bit strings), a genotype-to-phenotype transformation or “developmental” algorithm, and a fitness function that determines which phenotypes are selected to reproduce. Reproduction is simulated as genetic replication with recombination and (far less probably) mutation.

In the GA I coded for UPR-NNM, the developmental algorithm converts the genotypes into neural networks. This conversion follows

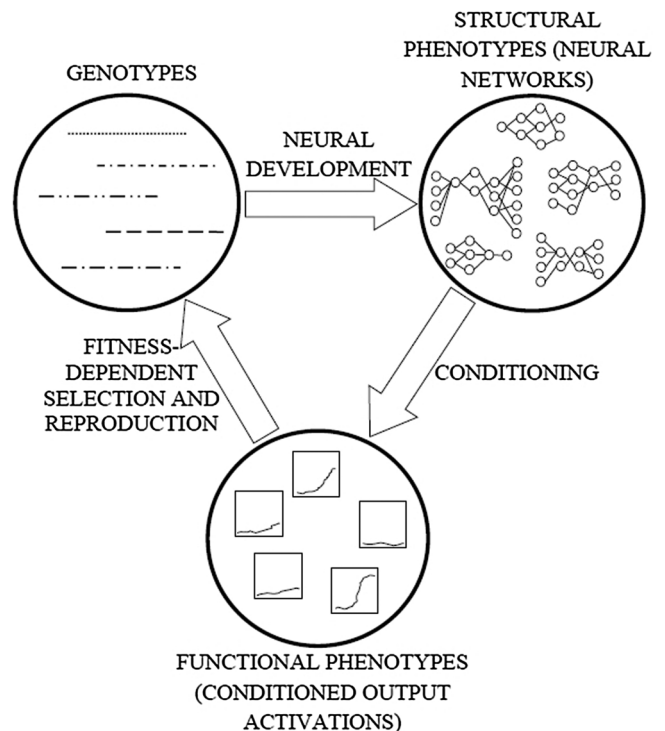


Fig. 1. An evolutionary cycle in a combination of UPR-NNM with a GA (after Burgos and Robayo, 2011, Fig. 3). First, a population of genotypes (simulated as bit strings that encode neurodevelopmental parameters) is developed into a population of structural phenotypes (neural networks), according to a neural development algorithm inspired by the phases of proliferation, migration, specification, synaptogenesis, and neuronal death, as identified in developmental neuroscience. Each network is then trained in a conditioning procedure of interest (e.g., a forward-delay procedure with a certain CS-US, interstimulus interval or ISI), and given test CSs to determine the conditioned output activations (functional phenotypes). After this, networks compete in selection tournaments based on their mean conditioned output activations during the test phase. Networks with the highest conditioned output activations win the tournaments and are selected to reproduce. Reproduction is simulated by replicating, recombining, and mutating the reproducing networks' genotypes. The result is a new population (generation) of genotypes. The cycle can be repeated for as many generations as wished.

the phases of neural development, identified in developmental neuroscience: Proliferation (production of unspecific units), migration (movement of produced units to layers), specification (when the features of the units, e.g., inhibitory or excitatory, free activation and learning parameters, etc., are specified), synaptogenesis (formation of connections, with initial weights set according to a Hebbian rule), and neuronal death (elimination of unconnected units). A genotype in this combination thus encodes in binary developmental parameters (e.g., proliferation and migration probabilities, etc.).

In contrast to other combinations of a GA with a neural-network model, genotypes in the present combination do not encode directly any particular traits of neural networks, let alone their behavior (cf. McDowell, 2004). The fitness function is just a network conditioned output activations after some training. The networks with higher conditioned output activations are more likely to be selected to reproduce.

4.1.2. A cycle for the evolution of learning without SbR

A cycle in this approach is shown in Fig. 1. First, a population of genotypes is transformed into a population of neural networks (the initial genotype population is randomly generated) through a neurodevelopmental algorithm. Each genotype is divided into segments that encode neurodevelopmental parameters in binary (e.g., proliferation, migration, and connection probabilities for the various types of units that can constitute a neural network in UPR-ANN; free activation and learning parameters for each type of unit, etc.). These parameters determine the final (“adult”) architecture of a neural network. The

genotypes do not directly encode any network structural or behavioral traits.

The result of the neurodevelopmental algorithm is a population of neural networks qua structural phenotypes. Then, each neural network is trained in a conditioning procedure, after which groups of randomly selected networks compete in selection tournaments to reproduce. The network with the highest conditioned output activation after training wins and is selected to reproduce. Reproduction consists of the replication, recombination, and (much less probably) mutation of the winners' genotypes (no inheritance of acquired characters). The result is a new population that descends from the previous one (a generation). This cycle is repeated for as many generations as desired, which results in a sequence of ancestor-descent generations, lineages, or evolutionary lines.

This cycle has all the basic components of a process of evolution by natural selection with (non-Lamarckian) genetic inheritance, as propounded in the synthetic theory: Genotypic and structural and functional phenotypic variation in a population, phenotype-dependent (in this case, CR-dependent) fitness function, and selection with reproduction (genetic recombination and mutation) that generates a new genotype population. The cycle also includes a developmental component that interacts with the evolutionary process.

The way the approach simulates the ontogeny/phylogeny, genotype/phenotype, and structure/function distinctions summarized in 2 is clear from Fig. 1. The approach captures the selection-of/selection-for distinction as follows. The most direct objects of selection are individual neural networks (structural phenotypes). They thus correspond to the balls in Sober's selection toy. The networks' genotypes are also objects of selection, but less directly, as reproductive success (winning selection tournaments) does not directly depend on them. The approach thus allows for at least two objects of selection, direct (neural networks qua structural phenotypes) and indirect (genotypes).

Networks are most directly selected for higher conditioned output activations (functional phenotypic traits). These activations are thus analogous to the smallness of the small balls in Sober's toy (with an important difference I will discuss in 5.3). Such higher activations are mediated by a structural phenotype, so networks are also selected for having certain structural traits (e.g., a certain number of input, hidden, and output units; a certain connectivity, etc.). Are any network structural traits analogous to the greenness of the selected balls in Sober's toy? No answer to this question is currently available, as the necessary quantitative analyses remain to be done. However, the approach provides all the necessary means to answer the question, albeit some preliminary conjectural work will likely be required (e.g., a criterion to decide whether a structural network trait, like the greenness of the balls in Sober's toy, piggybacks on those which allow networks to have higher conditioned output activations).

In preliminary simulations with this approach (Burgos, 1997), networks were selected for higher CR output activations under Pavlovian contingencies (in a forward-delay procedure) with different inter-stimulus intervals (ISIs) for different groups of networks, which resulted in different lineages. The results showed a negatively accelerated increase in the mean population fitness (CR output activations) across generations of networks, for all lineages. These results are consistent with Baldwin' (1896) conjecture that learning could be a factor in evolution without assuming Lamarckian inheritance. Different lineages resulted in phenotypically different network populations. In particular, shorter ISIs resulted in smaller networks, and longer ISIs in larger networks. There also was a progressive reduction in variation (genotypic and phenotypic), which is consistent with the observation that natural selection reduces variation within a lineage.

The approach is not restricted to Pavlovian conditioning strictly conceived. The approach also applies to a key aspect of operant conditioning, the acquisition of emitted responding, not elicited by the reinforcer. Burgos and Robayo (2011) showed that the approach could also simulate the evolution of conditioned emitted responding by

selection for higher emitted-responding output activations (primary-motor precursors of responding not caused by reinforcer input activations) under Pavlovian contingencies. The study thus simulated the evolution of a sort of autoshaping (without response directedness; cf. Burgos, 2007). The results, as expected, showed a substantial increase in the mean fitness (emitted-responding output activations under Pavlovian contingencies) from the first to the last generation.

My key point with all this is that a substantive contact between conditioning theory and the synthetic theory of evolution is possible without any form of *SbR*. A primary motivation for *SbR*, to make such a contact, thus vanishes. *SbR* is unnecessary for the purpose of making that contact. I see no need to assume that conditioning at the ontogenetic level, whether Pavlovian or operant, is deeply like evolution by natural selection.

5. A metaphysical concern

My concerns about *SbR* thus far have been epistemological. They point at *SbR*'s difficulties to account for Pavlovian conditioning, and be justified in view of a computational approach to the phylogeny and ontogeny of Pavlovian and operant conditioning. In this section, I supplement these concerns with a metaphysical one. It arises from the assumption in *SbR* that reinforcement selects behaviors (responses, actions, activities, response patterns, etc.), or environment-behavior relations. Specifically, are behaviors the kinds of entities of which it makes metaphysical sense to say that are selected? In terms of Sober's (1984) notion of selection-of, are behaviors *objects* of selection? I will argue that they are not, at least under a standard metaphysical view of objects, even if animals can be selected for behaving in certain ways (a different assertion I do not question). I begin with a reflection on the metaphysics of natural selection.

5.1. Natural selection as a process

Natural (biological) selection minimally involves heritable (genetically-determined) phenotypic-dependent differential reproduction in a population. Some organisms reproduce (leave viable offspring) more than others do for having certain heritable phenotypic traits under certain environmental conditions. A key aspect is implicit in the more reproductively successful organisms: Natural selection is a *process*. This diachronic character of natural selection is mentioned quite often in evolutionary biology, but seldom discussed. Darwin (1859) used several times the expression "the process of natural selection." A search for this expression in Google Books returns over 75,000 results, and "natural selection is a process" returns over 7000 results. Sober (1984) talks of "selection process" multiple times. A recent book gives this definition: "*Natural selection* is a process of sorting by reproductive success that occurs in populations of replicating units, whether those units are molecules, cells, organisms, or larger units" (Stearns, 2017 p. 196).

This process character of natural selection has never been put in metaphysical focus, but I will do it here for my analysis. What does it mean that natural selection is a process? We must first ask what a process is, to which Rescher (1996) answers:

A process is a coordinated group of changes in the complex of reality, an organized family of occurrences that are systematically linked to one another either causally or functionally. It is emphatically not necessarily a change in or of an individual thing, but can simply relate to some aspect of the general "condition of things." A process consists in an integrated series of connected developments unfolding in conjoint coordination ... Processes are correlated with *occurrences or events*: Processes always involve various events, and events exist only in and through processes. Processes develop over time (p. 38, emphasis added).

This answer includes all the key metaphysical aspects of processes, of which *change* may be the most central: Processes consist

fundamentally of change in that they develop over time and crucially involve a series of related events or occurrences (these will also be central to my analysis in 5.4). Processes may (although need not) involve one or more changes in an individual thing (more on things in 5.2).

One way to capture the process character of natural selection (there might be other ways) is to view natural selection as involving *episodes of changes* between two critical times: Before reproduction and during reproduction. The durations and separations of these times are irrelevant to the interpretation. Following Mayr (1997), I focus on individual organisms as objects of natural selection: “... most evolutionists agree that the individual organism is the principal object of selection” (p. 2091). I will thus assume that individual organisms can be correctly said to be selected. The interpretation (suitably modified) also works with other objects (as conceived in 5.2) of natural selection (e.g., genes, groups of organisms).

Formally, to say that an individual organism O was selected, in a way that captures the process character of natural selection, means *minimally* that O does not reproduce in time t_1 (before reproduction) and then reproduces in t_2 (during reproduction), both times before stating that O was selected (or that there was selection of O). Reproduction in t_2 is the culmination of the episode and a criterion (among others) for that statement. The change that gives an episode of natural selection a process character is from O 's not reproducing in t_1 to O 's reproducing in t_2 (the repetition of O in t_1 and t_2 will be crucial in 5.2). Fig. 2 depicts a diagram of this interpretation, where the arrows represent time as passing from left to right.

An episode of natural selection is more than a single reproduction incident. It *also* is genetically-heritable phenotypic-dependent differential reproduction in a population. I thus am not equating natural selection to reproduction of a single organism. I am only saying that the change from not reproducing to reproducing is partly *constitutive* of an episode of natural selection *as a process*. To say that O was selected means *additionally* that O (usually among others) reproduced in t_2 for having certain phenotypic traits under certain environmental conditions in at least t_1 (likely also in t_2). Other organisms reproduced less for having such traits to a lesser degree, under comparable conditions. For simplicity, I leave this additional meaning of “ O was selected” implicit to focus on said change.

This interpretation captures the basic intuition behind the breeders' artificial selection that inspired Darwin. Artificial selection is minimally a two-stage process. First, the breeder examines a number of candidates for reproduction. In this initial stage, there has been no selection yet,

but the stage is constitutive of the process of selection. Second, the breeder chooses at least one of the candidates over the others, and makes the selected candidate(s) mate to reproduce. Only then can these organisms be truthfully said to have been (artificially) selected to reproduce.

The same intuition is behind Sober's (1984) selection toy. In the first stage, the toy is upside down and all the balls are placed at the bottom. Then, in the second stage, the toy is turned and shaken, the small balls sieve through the levels, and end up at the other bottom. Only then, the small balls can be truthfully said to have been selected. In both cases, as in natural selection, there is a process consisting at least of an initial and a final stage.

Similarly, in an episode of natural selection, interactions between certain organisms' phenotypic traits and their environments occur during a first stage, which allow organisms to survive. At this point, there has not been selection yet, but this stage is constitutive of the process (the process *also* includes other organisms that do not survive for having such traits to a lesser degree). Then, sometime after these interactions, the *same* organisms mate and reproduce, while others do not. Only after this, the organisms that reproduced can be truthfully said to have been (naturally) selected.

5.2. Objects as continuants

In Fig. 2, what is selected, O , is an individual organism. In Sober's (1984) distinction, there is selection of O . If selection-of is only about objects, O is an object. However, I need to discuss in more depth the metaphysics of objects, because the term is used in several ways. Philosophers often use “object” to refer to properties (features, traits, characteristics), but if Sober did this, there would be little if any difference between selection-of and selection-for. I thus doubt that he uses “object” as synonymous with “property.”

What does he mean by “object,” then? One possibility is that he means the traditional metaphysical sense of “thing.” The balls in his selection toy are objects in this sense. There are many other examples: A particular, specific chair, table, ball, boulder, lightbulb, clock, car, pencil, tree, dog, cat, person, glass, ship (e.g., Titanic), bridge (e.g., the Golden Gate), city (e.g., New York City), building (e.g., the Empire State, the building as well as the State of New York), spoon, and computer, among myriad others.

Examples of things also abound in science. In astronomy and astrophysics, a specific celestial body (e.g., planet, satellite, asteroid, comet, and star) is a thing, as is a particular subatomic particle in quantum mechanics, and atom and molecule in chemistry. Jupiter is a thing, as is each one of its moons, our own Moon, our planet, and our Sun. In biology, a particular protein, cell, organ, animal, and plant is an object-qua-thing. O in Fig. 2 is an object-qua-thing as well, as is a particular neuron, neural circuit (natural or artificial), and brain in neuroscience. In behavior analysis, a particular experimental subject (e.g., Rat 221, Pigeon 143) is an object-qua-thing, as is any particular lever, key, food pellet, grain, and water drop (a stimulus or response is not an object-qua-thing, as I argue in 5.3).

In general, all objects-qua-things are specific, particular, concrete, *spatiotemporally restricted* entities. Grammatically, many things (e.g., people, pets, places) are denoted by proper names. Many spatiotemporally restricted entities are not objects-qua-things. An important class of such entities was anticipated in Rescher's (1996) quotation above: *Occurrences*, which will be central to my analysis and I will discuss in 5.3.

The terms “object” and “thing” are too equivocal. Following Simons (1987), I will use the more technical term “continuant” for “object-qua-thing.” Another term that philosophers use is “individual,” but some readers might misread it as “person” or “animal.” Philosophers have also used the term “substance,” but it brings back bad memories (e.g., Cartesian thinking substances). The term “continuant” does not suffer from these problems. I thus interpret Sober's (1984) notion of selection-

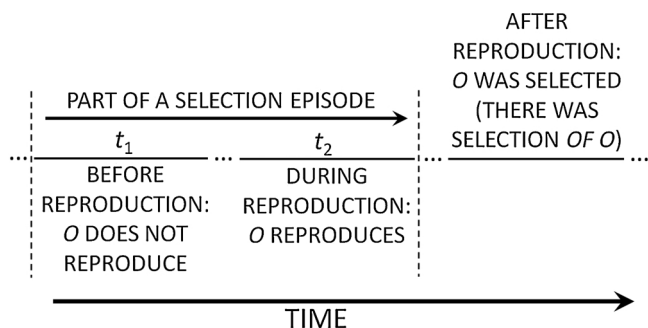


Fig. 2. An interpretation of part of an episode of natural selection as a process. Time (arrows) passes from left to right. An individual organism O does not reproduce in t_1 (before reproduction) and reproduces sometime afterwards in t_2 (during reproduction, which ends the episode). This change gives natural selection a process character. Sometime afterwards, O is said to have been selected (i.e., it is asserted that there was selection of O). The durations and separations (ellipses) of t_1 and t_2 are irrelevant. To say that O was selected additionally means that O reproduced for having certain phenotypic traits under certain environmental conditions during t_1 , and other individual organisms did not reproduce (or reproduced less) for having such traits to a lesser degree, in a comparable environment. To simplify, I leave this additional meaning implicit and focus on said change.

of to be only about continuants. All the above examples of objects-qua-things are continuants.

The term “continuant” also conveys a key feature that objects qua things are widely assumed to have. This feature thickens the plot slightly, but it is quite intuitive: Objects as things have an *identity* that makes them what they are and remains unchanged (“continues,” hence the term “continuant”) throughout their existence. This notion of identity is fundamental to all ontological considerations. As Quine (1957-58) aptly said, “No entity without identity” (p. 20): No existential claim can be meaningfully made without a criterion to decide unequivocally whether the claim refers to one or several entities (he uses the term “entity” in its broadest sense of a member of *any* ontological category; I use it in this way).

Exactly what determines a continuant’s identity is complicated, but I do not need to discuss this here. I need only say that a continuant’s identity is partially determined by the continuant’s *essential parts* (not to be confused with *essential properties*), without which the continuant would cease to exist or, at least, become a different continuant. Exactly what the essential parts of a thing are can also be complicated, especially in such complex things as organisms, but intuitively, it seems reasonable to assume that my left thumb, for instance, is not an essential part of me. I could lose it and still be me, or so I hope. Perhaps my brain (continuants themselves), or a certain part of it (e.g., the prefrontal cortex) is an essential part of me. Of course, there is much recycling of parts (molecules, cells, tissues) over time, so the term “identity” does not mean “perfectly identical throughout time.” Rather, it denotes more flexibly an *approximately* constant conservation of a certain *physical configuration*.

Whatever a continuant’s essential parts are, the standard account in metaphysics is that they all are *only spatial*, in the sense of having all three spatial dimensions. Continuants exist in spacetime, but all of their parts are only spatial. Continuants thus exist completely (have all their essential parts) at every moment of their existence (if a continuant loses an essential part, it ceases to exist or be the continuant it is). Continuants are thus said to persist through time by *enduring*, for which they are also called “endurants.”

5.2.1. Continuant self-sameness through time

All this can be expressed as the ontological principle of (approximate) *continuant self-sameness*: An object as a thing as a continuant is (approximately, sufficiently) the same entity at every moment of its existence. I am the same entity now as the entity I was an hour ago or yesterday, my cat Bolt is the same entity today as he was yesterday, the day before yesterday, and so on. The importance of this principle cannot be overstated. For example, it underlies all practices of punishment for bad deeds and rewards for good deeds. To punish a person today for having committed a crime last year assumes it is the same person today and last year. If two numerically distinct people were assumed here, it would be an injustice to punish the person today for a crime that a different person committed last year. Again, the sameness here is approximate, not strict: A continuant is the same entity (or *something close enough*) at every moment of its existence.

Similarly, in an FR-20 schedule, the rat that receives a food pellet for emitting the 20th barpress response is supposed to be the same rat that emitted the first 19 barpresses, and so on. The logic of within-subjects designs relies on this principle: *The* subject, in singular, is said to be its *own* control. The assumption is that the same animal receives various treatments through time. In this sense, then, these designs also implement processes. The same logic is found in the process character of natural selection.

In my interpretation in 5.1 (Fig. 2), *O* denotes an individual organism as a continuant. *O* is thus (approximately) the same entity in t_1 (before reproduction) and t_2 (during reproduction). This condition is necessary for saying truthfully in t_3 that *O* was selected (other conditions must also be met, as I have clarified). If numerically different entities were involved in t_1 and t_2 , this statement would be false, as it

refers to only one entity (*O*). It is unclear what alternative selection statement made in t_3 could refer to more than one entity and was true of this scenario while capturing the process character of natural selection. If two numerically different entities O_1 and O_2 are assumed to be involved in t_1 and t_2 , respectively, the only possible alternative statement is that O_2 was selected in t_2 , but this statement does not capture the key change that gives natural selection a process character.

5.3. Particular behaviors as occurrents

All continuants are spatiotemporally restricted entities, but not vice versa. Many entities are spatiotemporally restricted but not continuants, for which ontologists have another ontological category. This category is anticipated in Rescher’s (1996) quotation above: Occurrences. This category arises naturally from the precision that many continuants have many properties only sometimes, *occasionally* or *temporarily*. Behavioral properties are primary examples; other examples are physiological properties (e.g., a certain membrane potential or synaptic efficacy). No behavioral property is continuously possessed throughout any animal’s entire life. A particular behavior *occurs* when a particular animal has a behavioral property temporarily, in a particular time and place. Hence, all particular behaviors too are temporary. During one time, an animal, say a rat, does something (e.g., pressing a bar); during another time, the *same* animal does something else (e.g., eating a food pellet).

It makes no substantive difference whether a particular behavior is a single response, a sequence or pattern of responses, a stimulus-response or environment-behavior relation, or an activity extended in time (a particular baseball game, a particular piano recital). A particular response is as temporary as a particular action potential (even if the latter is much shorter). Both are as temporary as last night’s piano recital, Tuesday’s baseball game, my writing this essay for the past few hours, and this morning’s walk in the park and my waving my neighbor, all of which are longer and more complex than a particular barpress response or action potential. The key notion here is that all particular behaviors and their relations to their environments, no matter how short or long, simple or complex, are temporary possessions of behavioral properties. The molar-molecular distinction is irrelevant to this temporary character of particular behaviors.

In a word, all particular behaviors, no matter how short or long, simple or complex, are *events*, *occurrences*, or *happenings*. It is this that makes behavior *dynamical* in nature. I shall use the more technical and general term “occurrent” (e.g., Simons, 1987) to avoid the confusion that events are short, simple occurrences. I will also use the term to include related entities such as states and processes.

In general, occurrents involve *change*. They happen when an object ceases to have a certain property during some time and has a different property in some future time (e.g., when a chameleon turns from red to green, a patch of a neuron’s membrane changes polarity from -45 to 65 mv, a rat stops barpressing to eat, etc.). It is key to understand that occurrents, as traditionally conceived in metaphysics, belong in a different ontological category from properties (whether temporarily or permanently possessed) and continuants.

In natural selection, organisms can be selected for having properties permanently or temporarily. The small green balls in Sober’s (1984) toy are selected for being small, and they possess this property permanently. In contrast, consider a particular chameleon that turns green when it crawls into a green background (a dense forest tree). This occurrent camouflages the chameleon, allowing it to be missed by a predator (or a prey) in a particular time and place. In contrast to the small green balls in Sober’s toy, which are permanently green (and small and spherical), the chameleon is only temporarily green, but still it can be selected for having been temporarily green. This temporary possession of greenness might have allowed the chameleon to live another day by escaping a predator, or eating a pray, in which case, it also temporarily extended its tongue, another occurrent (a behavioral one).

This example illustrates that the notion of selection-for applies to temporarily possessed as much as it does to permanently possessed properties. Behavioral properties are not the exception to this. The chameleon's temporary extension of its tongue is a particular behavior (an occurrent, a temporary possession of the property of tongue-extensionhood), for which the chameleon can be selected (i.e., survive and mate afterwards). A cheetah's running during a particular time in a particular place in chase of a prey is a particular behavior (an occurrent) where the cheetah is an object that temporarily possesses the property of running-hood. If the cheetah gets the prey, this occurrent allowed the cheetah to obtain nourishment and thus live another day, increasing its chances to mate and reproduce. If the cheetah mates afterwards, it can be said to have been selected for *having* run fast.

5.3.1. Behaviors-qua-occurrents have temporal parts

The main outcome of the preceding section is that particular behaviors are occurrents. According to a standard metaphysical view, a distinctive feature of all occurrents, particular behaviors included, is that they (in contrast to continuants) have *temporal parts*, stages, or phases. The notion of a temporal part might sound strange, as the term "part" is commonly used to refer to a continuant's spatial parts (e.g., a table's legs, an organism's cells, a rat's, chameleon and pianist's limbs, etc.; see 5.2). 5.2.

As an example, consider a particular behavior like a waving episode in a particular time and place by a continuant called "Sticko," as depicted in Fig. 3. Like all other occurrents, this behavior is divided into indefinitely many temporal segments, but I will illustrate with six, labeled as t_1 through t_6 . The horizontal arrow represents the passage of time, from left to right. Sticko's name is repeated throughout the occurrent to indicate that Sticko (like *O* in Fig. 2) is a continuant and hence (approximately) the same entity throughout the occurrent (see 5.2.1).

Each segment is a snapshot as a temporal part of the occurrent. During t_1 , Sticko's right arm is at rest. During t_2 , Sticko raises its right arm a bit, and so on. This episode consists of shorter occurrents (e.g., Sticko's right arm position change from t_1 to t_2 is also an occurrent, but a shorter one). All these segments (and infinitely many more) are temporal parts of the occurrent. It is the occurrent which has temporal parts, not Sticko. Sticko is a continuant and hence satisfies the principle of continuant self-sameness discussed in 5.2.1. During each temporal part, Sticko possesses a certain behavioral property temporarily. The




Fig. 3. A waving episode, an occurrent during which a continuant, Sticko, possesses certain behavioral properties (right-arm positions) temporarily in particular times and a particular place. Time passes from left to right (horizontal arrow), where t_1, \dots, t_6 are some temporal segments or parts of the occurrent. If any of these segments, say, t_4 , is the present (highlighted by the vertical arrow), the occurrent exists incompletely during that segment, in that some of its parts (t_1, t_2 , and t_3) are in the past, others (t_5 and t_6) in the future. The occurrent is complete only when it ends (all of its parts are in the past), and cannot occur again (it is unrepeatable). Only *similar* occurrents can occur (after this particular occurrent), but they would still be *numerically* different occurrents. Sticko's name is repeated across the entire episode to indicate that Sticko is (approximately) the same entity throughout the occurrent (i.e., satisfies the principle of continuant self-sameness discussed in 5.2.1).

behavioral properties that Sticko possesses temporarily are the different right-arm positions (at rest at a certain angle at t_1 , slightly higher at a different angle at t_2 , even higher at another angle at t_3 , etc.).

Having temporal parts implies that the occurrent (like all other occurrents) is *incomplete* or, as is more commonly put in metaphysics, *exists incompletely* during any of its temporal parts. Assume for the sake of argument that Sticko's right arm position during t_4 (a temporal part of the occurrent) is the present (highlighted by the vertical arrow). During t_4 , the occurrent is incomplete because some of its temporal parts (t_1, t_2 , and t_3) are in the past, others in the future (t_5 and t_6). This feature obtains for any time during this (or any other) occurrent. However, Sticko, as a continuant, exists completely during any temporal part of the occurrent.

The notion of a temporal part qua temporal segment and its implication that any occurrent exists incompletely during any of its temporal parts should not be counterintuitive. A particular baseball game is temporally segmented into the first inning, second inning, third inning, and so on. Each inning is a temporal part of game. The first inning is an initial segment (a temporal part of the beginning of the game), the second inning another segment that is in the future of the first inning, and so on (each inning is another, shorter occurrent, consisting in turn of other, even shorter occurrents). During the fourth inning, the game exists incompletely, as some of its parts (the first three innings) lie in the past, others (the rest of the innings) in the future. A soccer game is temporally segmented into the first half and second half (other segments include various "stoppage" times and a halftime), and so on. During the second time, the game exists incompletely, as some of its temporal parts are in the past, and so on.

An occurrent is complete (exists completely) only when it is over, and hence, *all* of its temporal parts are in the past. Sticko's waving episode is over after t_6 , a chameleon's turning green occurrent is over when the chameleon changes color, a baseball game is over after the last inning, a rat's barpressing response is over as soon as the bar micro-switch closes, and so on. Once an occurrent ends, it is *unrepeatable*. The reason is clear: The time of an occurrent is constitutive of the occurrent. An occurrent that has ended lies entirely in the past. Hence, to assert that an occurrent is repeatable (i.e., that the *same* occurrent can occur again) implies that the past could be in the present, which violates basic intuitions about time.⁵ Sticko can thus be truthfully said to *have waved* only *after* the occurrent is over, not before.

Similar occurrents often take place, but they would still be *numerically* different occurrents qua particular happenings. The first time a chameleon turns green is a different occurrent from the second time the *same* chameleon turns green (after having turned brown), and this occurrent is different from the third time a chameleon turns green, and so on. There are three different occurrents of the same type, but they are unrepeatable. They are similar, but this does not make them numerically one and the same occurrent (even if there is only one occurrent type). Likewise, a rat's first barpress is a different occurrent from the *same* rat's second barpress (I explain the emphasis in 5.2), and the second barpress is different from the third barpress, and so on. They are similar (which is traditionally interpreted in metaphysics as being of the same type) but still count as numerically three distinct responses, and they are unrepeatable.

5.3.2. Is there selection of occurrents?

In 5.2.1, Fig. 2, to say in t_3 that *O* was selected means partly that *O* is the same entity before (in t_1) and during reproduction (in t_2). That is to say, *O* satisfies the principle of continuant self-sameness, which is necessary (but not sufficient) for saying truthfully in t_3 that *O* was

⁵ A related issue arises from assuming that a reinforcer reinforces the response that produced it. A present reinforcer cannot affect a past response, which is why Skinner (1935) postulated response *classes* as what reinforcers reinforce. However, I have never understood what this could mean metaphysically, whether classes are conceived as properties or as sets.




Fig. 4. A particular chameleon, C_1 , is temporarily green in time t_1 . This temporary possession of greenness by C_1 is an occurrent (E_1) that allowed C_1 to survive, but happens before C_1 mates and reproduces in t_2 . Between t_1 and t_2 (thick arrow), C_1 is depicted as having turned from green to brown (another occurrent, not labeled), to emphasize the fact that E_1 is over by, and hence lies entirely in the past of, t_2 , the time when C_1 reproduces. To further emphasize this, C_1 changes to a different color in t_2 . C_1 is a continuant and thus the same entity in t_1 and t_2 . C_1 (like O in Fig. 2) can thus be truthfully said to have been selected. The same cannot be said about E_1 . Unlike C_1 (and O in Fig. 2), E_1 is not present during reproduction.

selected. Sticko in Fig. 3 too satisfies this principle. Do occurrents satisfy this principle? In this section I argue that they do not. I will illustrate with a more complex version of Fig. 2.

Fig. 4 shows this version, where the key continuant (corresponding to O in Fig. 2), an individual chameleon, is labeled as C_1 . For simplicity, I have omitted the part of Fig. 2 when the organism is said to have been selected after reproduction. In t_1 , before reproduction, C_1 turns green, an occurrent labeled as E_1 (E for “event,” O for “occurrent”) could be confused with O for “organism”). Assume that E_1 allowed C_1 to survive another day (e.g., by catching a prey, not depicted).

Sometime after E_1 , C_1 turns brown, another occurrent (not labeled), and stays brown for some time (thick arrow). This occurrent indicates the fact that E_1 is over and thus lies entirely in the past of this other occurrent. In t_2 , C_1 reproduces (for simplicity, only the mating is depicted), and has a darker color, to emphasize that fact.

Like O in Fig. 2, C_1 in the present example is a continuant: All of its parts are only spatial and thus exists completely in every time throughout its existence, the reproduction episode included. C_1 is thus the same entity in t_1 and t_2 , and E_1 and E_2 . Like O in Fig. 2, then, C_1 satisfies the principle of continuant self-sameness, which meets a condition (among others) for truthfully stating after the episode that C_1 was selected (that there was selection of C_1), in a way that captures the process character of selection.

But is there selection of E_1 , the occurrent for which C_1 was selected? Can E_1 be truthfully said after the episode to have been selected? Clearly not, because a truth condition for this statement is that the selected entity must be present before and during reproduction. However, E_1 is not present during reproduction (E_1 is over when C_1 reproduces; C_1 ’s reproduction occurs after E_1). Therefore, E_1 cannot be truthfully said (after the episode) to have been selected, at least not in the same metaphysical sense that C_1 and O (Fig. 2) can be said to have been selected. If there is another sense, I do not know what it could be. The same outcome obtains if E_1 extended beyond t_1 to t_2 , which is possible: C_1 could remain green during the entire episode. However, E_1 can still not be truthfully said to have been selected because it exists incompletely during reproduction (part of it lies in t_1). What is present in t_1 is not E_1 but a temporal part of E_1 , which is different than what is present in t_2 , another temporal part of E_1 . E_1 either is not present at all in t_1 , or only partially present.

A key implication for my present purposes is that particular behaviors (whether responses, or activities) qua occurrents cannot be meaningfully said to be selected, or be objects of selection. There is no selection of particular behaviors. The assumption that reinforcement selects behaviors, or that behaviors are the objects of *SbR*, or that there is selection of behavior, is metaphysically unsound, at least in the traditional metaphysical interpretation of objects-qua-continuant and particular behaviors-qua-occurrents I have adopted here. A different metaphysical interpretation might change this implication, but I cannot

envision what such alternative interpretation could be. The interpretation I have adopted is widely used across philosophy, so there better be very good reasons to replace it. Avoiding that implication per se does not strike me as a good reason.

5.4. Other possibilities

In sum, the standard metaphysical interpretation of objects I have used here restricts Sober’s (1984) notion of selection-of to continuants. Therefore, there is no selection of occurrents. If reinforcement does not select particular behaviors, what does it select? No obvious answer presents itself, which puts *SbR* at risk of being left without a key component of the idea of selection: A plausible object of selection. The only alternatives left are to view particular behaviors as continuants, or assume that reinforcement selects behavioral properties rather than particular behaviors. Both possibilities face very difficult issues. Viewing particular behaviors as continuants amounts to viewing them as things or substances, a reification that has no metaphysical justification whatsoever. Again, postulating this (or any other) view just to save *SbR* does not seem like a good reason.

As for the second possibility, much hinges on how behavioral properties are metaphysically conceived. Traditionally, they have been conceived as *universals* or *types*, meaning that they can (many in fact do) have multiple examples or tokens. That is to say, they are *abstractnesses*. The problem is that it remains most unclear what it could mean metaphysically for abstract entities to be selected. Suppose some animals (e.g., ants), like the small balls in Sober’s selection toy, are selected for being small. In this case, ants are the objects of selection, and smallness is the property for which they were selected. But what does it mean that this property too is an object of selection? Aside from putting ants and smallness in the same ontological category (a recipe for confusion), it remains unclear what does it mean that smallness is selected? Many other animals are selected for being large (e.g., elephants), which is to say, for *not* being small. By the same logic, this means that smallness is *not* selected. Hence, smallness is *and* is not selected, which is contradictory. Besides, all of it blurs the distinction between selection-of and selection-for too much, making it much harder to formulate and use. If a property can be both selection-of and selected-for, what would be the point of making the distinction in the first place? The distinction is easier to formulate and use if the object-property distinction is made sharply (as indeed it is in traditional metaphysics).

6. Concluding remarks

Because of all of the above reasons, I do not pursue *SbR* anymore and will not until they are compellingly addressed. I have two more reasons. One, *SbR* has been propounded as an explanation of complex behavior (e.g., Donahoe and Palmer, 1994). Although the theory of

evolution by natural selection is the best explanation of biological complexity, I am not so sure that this implies that explanations of complex behavior at the ontogenetic level requires *SbR*.

Complex behavior can be explained as the result of learning by interaction between certain complex environmental conditions and a complex animal with certain complex structures that have evolved by natural selection during phylogeny. An adult brain, in particular, especially human, is a very complex continuant, the structure and functioning of which allows its bearer to learn to execute many different behaviors under many different environmental conditions. A brain's development, structure, and functioning is the result of millions of years of evolution by natural selection of only continuants (whether genes, individual organisms, or groups), neither properties nor occurrences. When such a complex organ functions under complex environmental conditions (as part of an animal), it mediates learning of complex behavior. This explanation does not call for any deep analogy between evolution and learning.

Two, *SbR* is also founded on the assumption that selection is "a causal mode found only in living things" that "replaces explanations based on the causal modes of classical mechanics" (Skinner, 1981, p. 501). No behavioral selectionist thus far has questioned this formulation, so I assume it remains foundational of *SbR*. My problem with it is that I do not understand it. It raises too many unanswered questions that need answers if *SbR* is to be intelligible. What is a "causal mode"? What is the causal mode of "classical mechanics" and how does it differ from selection by consequences? Why does selection by consequences "replace" the causal mode of classical mechanics? Exactly what is the problem with the classical-mechanics causal mode vis-à-vis living things? Are explanations of living things in terms of this causal mode impossible? Why? Or they are possible but worse than explanations in terms of behavioral selectionism? Why? I have found no answers to any of these questions in discussions of *SbR*, which makes me even more doubtful about it.

Compliance with Ethical Standards

This study was not funded by any grant.

Ethical Approval

No animals were used or participated in this essay.

Acknowledgements

I thank William Baum for his kind invitation to submit a manuscript for review in this issue. I also thank him, two anonymous reviewers, John Donahoe, David Palmer, and Elliott Sober for insightful commentaries to a previous draft.

References

- Baldwin, J.M., 1896. A new factor in evolution. *Am. Nat.* 441–451 XXX.
- Balsam, P.D., 1985. The functions of context in learning and performance. In: Balsam, P.D., Tomie, A. (Eds.), *Context and Learning*. Erlbaum, Hillsdale, NJ, pp. 1–21.
- Balsam, P.D., Tomie, A. (Eds.), 1985. *Context and Learning*. Erlbaum, Hillsdale, NJ.
- Baum, W.M., 2012. Rethinking reinforcement: allocation, induction, and contingency. *J. Exp. Anal. Behav.* 97, 101–124. [10.1901/jeab.2012.97-101](https://doi.org/10.1901/jeab.2012.97-101).
- Baum, W.M., 2017. Selection by consequences, behavioral evolution, and the price equation. *J. Exp. Anal. Behav.* 107, 321–342. [10.1002/jeab.256](https://doi.org/10.1002/jeab.256).
- Breland, K., Breland, M., 1961. The misbehavior of organisms. *Am. Psychol.* 16, 681–684. [10.1037/h0040090](https://doi.org/10.1037/h0040090).
- Brown, P.L., Jenkins, H.M., 1968. Auto-shaping of the pigeon's key-peck. *J. Exp. Anal. Behav.* 11, 1–8. [10.1901/jeab.1968.11-1](https://doi.org/10.1901/jeab.1968.11-1).
- Burgos, J.E., 1997. Evolving artificial neural networks in Pavlovian environments. In: Donahoe, J.W., Dorsel-Packard, V. (Eds.), *Neural-Network Models of Cognition: Biobehavioral Foundations*. Elsevier, Amsterdam, pp. 58–79. [10.1016/S0166-4115\(97\)80090-8](https://doi.org/10.1016/S0166-4115(97)80090-8).
- Burgos, J.E., 2001. A neural-network interpretation of selection in learning and behavior. *Behav. Brain Sci.* 24, 531–532.
- Burgos, J.E., 2007. Auto-shaping and automaintenance: a neural-network approach. *J. Exp. Anal. Behav.* 88, 115–130. [10.1901/jeab.2007.75-04](https://doi.org/10.1901/jeab.2007.75-04).
- Burgos, J.E., Robayo, B., 2011. Condicionamiento y evolución: Una aproximación computacional. In: Gutiérrez, G., Papini, M.R. (Eds.), *Darwin y las Ciencias del Comportamiento*. Universidad Nacional de Colombia/Colegio Colombiano de Psicólogos, Bogotá, pp. 337–361.
- Colwill, R.M., Rescorla, R.A., 1990. Effect of reinforcer devaluation on discriminative control of instrumental behavior. *J. Exp. Psychol. Anim. Behav. Processes* 16, 40–47. [10.1037/0097-7403.16.1.40](https://doi.org/10.1037/0097-7403.16.1.40).
- Davis, H., Hurwitz, H.M.B. (Eds.), 1977. *Operant-Pavlovian Interactions*. Erlbaum, Hillsdale, NJ.
- Darwin, C., 1859. *On the origin of species by natural selection, or the preservation of favoured races in the struggle for life*. Murray, London. <https://archive.org/details/onoriginofspec00darw>.
- Dickinson, A., Watt, A., Varga, Z.I., 1996. Context conditioning and free-operant acquisition under delayed reinforcement. *Quart. J. Exp. Psychol.* 49B, 97–110.
- Donahoe, J.W., 2013. Reinforcement learning, psychological perspectives. In: Pashler, H. (Ed.), *Encyclopedia of the Mind*. Sage, Los Angeles, pp. 643–646.
- Donahoe, J.W., Burgos, J.E., Palmer, D.C., 1993. A selectionist approach to reinforcement. *J. Exp. Anal. Behav.* 60, 17–40. [10.1901/jeab.1993.60-17](https://doi.org/10.1901/jeab.1993.60-17).
- Donahoe, J.W., Crowley, M.A., Millard, W.J., Stickney, K.A., 1982. A unified principle of reinforcement. In: Commons, M.L., Herrnstein, R.J., Rachlin, H. (Eds.), *Quantitative Analyses of Behavior: vol. 2: Matching and Maximizing Accounts*. Balinger, Cambridge, MA, pp. 493–521.
- Donahoe, J.W., Palmer, D.C., 1994. *Learning and Complex Behavior*. Allyn & Bacon, Boston.
- Donahoe, J.W., Palmer, D.C., Burgos, J.E., 1997a. The S-R issue: its status in behavior analysis and in Donahoe and Palmer's learning and complex behavior. *J. Exp. Anal. Behav.* 67, 193–211. [10.1901/jeab.1997.67-193](https://doi.org/10.1901/jeab.1997.67-193).
- Donahoe, J.W., Palmer, D.C., Burgos, J.E., 1997b. The unit of selection: what do reinforcers reinforce? *J. Exp. Anal. Behav.* 67, 259–273. [10.1901/jeab.1997.67-193](https://doi.org/10.1901/jeab.1997.67-193).
- Donahoe, J.W., Vegas, R., 2004. Pavlovian conditioning: the CS-UR relation. *J. Exp. Psychol. Anim. Behav. Processes* 30, 17–33.
- Dunlap, A.S., Stephens, D.W., 2014. Experimental evolution of prepared learning. *Proc. Natl. Acad. Sci. U. S. A.* 111, 11750–11755. [10.1073/pnas.1404176111](https://doi.org/10.1073/pnas.1404176111).
- Estes, W.K., Skinner, B.F., 1941. Some quantitative properties of anxiety. *J. Exp. Psychol.* 29, 390–400. [10.1037/h0062283](https://doi.org/10.1037/h0062283).
- Foree, D.D., LoLordo, V.M., 1973. Attention in the pigeon: differential effects of food-getting versus shock-avoidance procedures. *J. Comp. Physiol. Psychol.* 85, 551–558.
- Gould, T.J., Bevins, R.A., 2012. Context conditioning. In: Seel, N.M. (Ed.), *Encyclopedia of the Sciences of Learning*. Springer, New York, pp. 794–797.
- Holland, J.H., 1975. *Adaptation in Natural and Artificial Systems: An Introductory Analysis With Applications to Biology, Control, and Artificial Intelligence*. The MIT Press, Cambridge, MA.
- Jenkins, H.M., Barrera, F.J., Ireland, C., Woodside, B., 1978. Signal-centered action patterns of dogs in appetitive classical conditioning. *Learn. Motiv.* 9, 272–296. <https://www.sciencedirect.com/science/article/pii/0023969078900103>.
- Mayr, E., 1997. The objects of selection. *Proc. Natl. Acad. Sci. U. S. A.* 94, 2091–2094.
- McDowell, J.J., 2004. A computational model of selection by consequences. *J. Exp. Anal. Behav.* 81, 297–317. [10.1901/jeab.2004.81-297](https://doi.org/10.1901/jeab.2004.81-297).
- Quine, W.V.O., 1957. Speaking of objects. 1957–58. *Proc. Addr. Am. Phil. Assoc.* 31, 5–22.
- Rescher, N., 1996. *Process Metaphysics: An Introduction to Process Philosophy*. State University of New York Press, Albany, NY.
- Schindler, C.W., Weiss, S.J., 1982. The influence of positive and negative reinforcement on selective attention in the rat. *Learn. Motiv.* 13, 304–323. [10.1016/0023-9690\(82\)90012-1](https://doi.org/10.1016/0023-9690(82)90012-1).
- Shull, R.L., 1995. Interpreting cognitive phenomena: review of Donahoe and Palmer's learning and complex behavior. *J. Exp. Anal. Behav.* 63, 347–358. [10.1901/jeab.1995.63-347](https://doi.org/10.1901/jeab.1995.63-347).
- Simons, P., 1987. *Parts: A Study in Ontology*. Oxford University Press, Oxford.
- Skinner, B.F., 1935. The generic nature of the concepts of stimulus and response. *J. Gen. Psychol.* 12, 40–65.
- Skinner, B.F., 1948. 'Superstition' in the pigeon. *J. Exp. Psychol.* 38, 168–172. [10.1037/0096-3445.121.3.273](https://doi.org/10.1037/0096-3445.121.3.273).
- Skinner, B.F., 1981. Selection by consequences. *Science* 213, 501–504. [10.1126/science.7244649](https://doi.org/10.1126/science.7244649).
- Sober, E., 1984. *The Nature of Selection: Evolutionary Theory in Philosophical Focus*. The University of Chicago Press, Chicago.
- Staddon, J.E.R., 2016. *Adaptive Behavior and Learning*, second edition. University Press, Cambridge, Cambridge.
- Staddon, J.E.R., Simmelhag, V.L., 1971. The "superstition" experiment: a reexamination of its implications for the principles of adaptive behavior. *Psychol. Rev.* 78, 3–43. [10.1037/h0030305](https://doi.org/10.1037/h0030305).
- Stearns, S.C., 2017. Natural selection, adaptation, and fitness: overview. In: Losos, J.B. (Ed.), *The Princeton Guide to Evolution*. Princeton University Press, Princeton, NJ, pp. 195–201.
- Tolman, E.C., 1924. The inheritance of maze-learning ability in rats. *J. Comp. Psychol.* 4, 1–18. [10.1037/h0071979](https://doi.org/10.1037/h0071979).
- Tonneau, F., Sokolowski, M.B.C., 2000. In: Tonneau, F., Thompson, N.S. (Eds.), *Perspectives in Ethology Volume 13: Evolution, Culture, and Behavior*. Springer, New York, pp. 155–180.
- Tryon, R.C., 1940. Genetic differences in maze-learning ability in rats. 39th Yrbk. *Natl. Soc. Stud. Educ.* 1, 111–119.
- Walker, B.Y.K.C., 1942. The effect of a discriminative stimulus transferred to a previously unassociated response. *J. Exp. Psychol.* 31 (4), 312–321. [10.1037/h0062929](https://doi.org/10.1037/h0062929).
- Zener, K., 1937. The significance of behavior accompanying conditioned salivary secretion for theories of the conditioned response. *Am. J. Psychol.* 50, 384–403. <https://www.jstor.org/stable/1416644>.