Behavioural Processes 161 (2019) 65–72

Contents lists available at ScienceDirect

Behavioural Processes

journal homepage: www.elsevier.com/locate/behavproc

Group selection in behavioral evolution

Howard Rachlin

Psychology Department, Stony Brook University, Stony Brook, NY 11794-2500, United States

ARTICLE INFO

Keywords:
Altruism
Behavioral evolution
Behavioral variation
Biological evolution
Extended self
Group selection
Self-control
Teleological behaviorism

ABSTRACT

How may patterns of behavior change over an organism's lifetime? The answer is that they evolve (behavioral evolution) as species evolve over generations (biological evolution). In biological evolution, under certain conditions, groups of cooperative organisms would be selected over groups of non-cooperative organisms, even when cooperation imposes a cost to individuals. Analogously, in behavioral evolution, patterns of acts may be selected even when each individual act in the pattern is costly. Although there is considerable debate among biologists whether the conditions for group selection are met in biological evolution, it is argued here that they are met in behavioral evolution (as well as in cultural evolution). The article shows how selection of patterns can explain the learning of self-control and altruism.

1. Introduction

From the viewpoint of teleological behaviorism (Rachlin, 1992, 1994, 2010, 2014, 2017) the conflict between self-control and impulsiveness is not primarily a conflict among representations of intentions in some specific location in a person's brain—between internal intentions to do one thing and internal impulses to do another. Rather, the conflict is fundamentally temporal—between (on the one hand) abstract patterns of overt behavior extended in time, rewarded in the long-run, and (on the other hand) particular overt acts, rewarded in the short-run, but inconsistent with those patterns. For example, smoking a cigarette at this moment has several immediate advantages including improved concentration, alertness, a pleasurable feeling; but smoking over an extended period is highly maladaptive. Explaining self-control creates a problem for ordinary behavioristic theories of learning, in which particular acts are increased in rate by particular reinforcers; the adaptive pattern (e.g., smoking less frequently) consists of particular choices (not smoking now) each of which is unreinforced (or relatively weakly reinforced) or, in the case of addicts, actually punished. How can patterns of acts increase in rate when each of their components is unreinforced or punished?

A similar conflict occurs in biological evolution. In a given environment a group of altruistic organisms may be better adapted to its environment than is a group of selfish organisms even though, within each group, selfishness is more adaptive than altruism. How can altruistic behavior evolve when individual survival is enhanced by selfishness? A common answer to that question is that evolution occurs at the genetic level rather than at the level of individual organisms. Individuals may be altruistic to the extent that their genes overlap with those of the beneficiary of their altruism (Dawkins, 1989). For example, Rachlin and Jones (2008) asked participants to choose between $75 to be given to another person or a (usually) smaller amount of money for themselves. They could choose to accept the smaller amount (v) or forgo it for the benefit of the other person. A person might prefer to give $75 to the 10th closest person to her rather than to receive as much as $20 for herself.[1] The maximum amount forgone (v = $20 in this case) varied directly (and strongly) with the degree of relatedness (genetic overlap) between subject and receiver. However, v also varied strongly with the social distance to the receiver—even among non-relatives. We found that amount forgone (v) decreased as the ordinal social distance (between subject and receiver) increased. The relationship between social distance and amount forgone was described by a hyperbolic discount function: $v = \$75/(1 + kN)$, where v = amount forgone, N = ordinal social distance between subject and receiver, and k = a constant differing among individuals. The lower the value of k, the more altruistic the subject. When receivers were relatives (second cousins or closer), the average discount function was shallower than it was for non-relatives; however, there was a great deal of overlap. Many subjects were more generous to non-relatives to whom they felt close than they were to relatives, even close relatives, to whom they felt comparatively distant. Of course, there is a strong genetic overlap even among strangers, but that does not explain why altruism should vary so hyperbolically (in both literal and metaphorical senses of the word) among non-related individuals.[2]

E-mail address: howard.rachlin@stonybrook.edu.

[1] These rewards were hypothetical, but a further experiment with real rewards of lesser amounts found corresponding results (Locey et al., 2011).

[2] Rachlin and Jones (2008) found that the steepness of an individual's social discount function correlated with the steepness of her delay discount function. The correlation was statistically significant but not strong. That is, there were many exceptions—people generous to others but not to their future selves, and the reverse.

http://dx.doi.org/10.1016/j.beproc.2017.09.005
Received 16 June 2017; Received in revised form 5 September 2017; Accepted 5 September 2017

A second problem with the pure genetic overlap theory of altruism is that it is a genetic theory. It gives people a genetic altruistic tendency but does not by itself answer the question: How may a person learn to be altruistic? As indicated above, classic behavioral learning theories offer no answer to this question either. Those theories (e.g., Skinner, 2014/1969) say that patterns of acts are established as chains of discrete responses; feedback from each link (each act in the chain) is thought to serve as a conditioned reinforcer for the prior link and as a discriminative stimulus for the following link—until the final link, which is reinforced by an unconditioned reinforcer. But if each link of the pattern were punished (as withdrawal symptoms punish cigarette abstinence), the pattern of abstinence could not be learned or maintained in this way. How can a pattern of acts be learned when none of its components is reinforced—when some may be punished? One conceivable answer to this question would rely on hypothesized internal representations of the ultimate reward (e.g., health) to take the place of external discriminative stimuli and conditioned reinforcers, but there are both conceptual and empirical problems with the notion that a person can internally reinforce her own acts (to be discussed later). The present article presents a different mechanism—a behavioral analog to a mechanism that has been proposed by some biologists to explain the evolution of altruism—group selection.[3]

The main purpose of this article is to explore this analogy—between selection of groups of organisms in biological evolution and selection of groups (or patterns) of acts of an organism in behavioral evolution. Group selection in behavioral evolution is a mechanism by which individual organisms may learn, over their lifetimes, to pattern their behavior in temporally extended units. Such patterns constitute self-control and, in group settings, altruism (Rachlin, 2014). There has been considerable debate among biologists whether the conditions for group selection ever occur in real-world societies but most agree that, were certain conditions to obtain, biological group selection would at least be possible (Nowak, 2006; Smaldino, 2014). It is an open question whether group selection may occur in the analogous conflict between individual acts and temporally extended patterns of acts that explains the learning of self-control and altruism. That question will be addressed in the remainder of this article—and answered in the affirmative.

2. Evolution

Fig. 1 is a general model of how evolution works. It is far from a formal model; it is just a template for relating corresponding processes in different kinds of evolution to each other. An entity (a trait or habit) varies along some dimension; a constraint or contingency, external to the trait or habit, selects a specific instance or set of instances from the varying entity (The dashed line in the box labeled, SELECTION indicates the selected instance); a generative process alters the evolving entity so as to increase the relative prevalence of the selected instance, and so forth. Although the model is very simple, complexities arise because the selection process, the generative process, and the properties of variation may themselves evolve (i.e., take their place in the diagram as evolving entities) all at the same time in both biological and behavioral evolution.

Let us consider biological evolution first. The genetic composition of a species (the evolving entity) varies, resulting in corresponding variation of structure, development, and innate behavioral patterns (instincts) within the species. A selective process in the environment kills members of the species population with maladaptive traits whereas those with traits that fit better in the ecological and social niche survive and reproduce (the generative process) increasing their relative frequency in the population [Note that variation itself is a trait and may evolve to better fit the variation of the environment (Grunow and Neuringer, 2002; Johnson et al., 1994)]. The population thus changes in structure, development and behavior to fit better in its ecological niche (which itself is constantly varying).

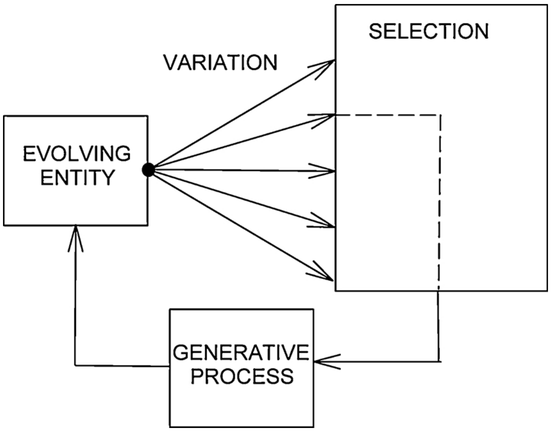

Fig. 1. A general model of an evolutionary process. An evolving entity varies (as represented by the multiple arrows). A process of selection acts as a filter, allowing (in this case) one of the variants (symbolized by the dashed line) to enter the generating process, which alters the properties of the evolving entity.

A frequently used example of a biological evolutionary process is the length of the giraffe's neck. The giraffe population is the evolving entity; like all giraffes' traits, the length of their necks varies; the height of the edible leaves on local trees is one selective process among many in the giraffes' ecological system; all else equal, giraffes with too-short necks will die earlier than those with longer necks (because it is difficult for short-necked giraffes to reach the edible leaves and to spot lions in the distance, among other advantages of height); biological reproduction (the generative process), among those longer-necked giraffes that live long enough to reproduce, increases the average neck-length of the population of giraffes as a whole until other selective forces (such as the difficulty of pumping blood from the heart to the brain) counteract the neck-lengthening mechanism. This story ignores many complications, such as the fact that the trees as well as many other aspects of the giraffes' environment are evolving at the same time. Therefore, as indicated above, Fig. 1 is a gross simplification. Nevertheless, this very wide concept of evolution may apply to many traits or habits otherwise explained in terms of non-natural forces (in the case of biological evolution) or an internal and/or spiritual force of free will (in the case of behavioral evolution).

3. Group selection in biological evolution

Imagine a (simplified) basketball league where the only factor that determined whether a player was retained on a team was points scored by that player (individual selection), and the only factor that determined whether a team was retained in the league (and the franchise not moved to another city) was games won by the team (group selection). All else equal, a team on which the individuals play unselfishly will win more games than a team on which they play selfishly. Yet, at the same time, the selfish players on each team will score more points than the unselfish players. From the individual player's point of view, the selective process (points scored) will increase playing time, salary, fame, etc., which in turn generates shots taken (as opposed to passing, getting back on defense, etc.) From the coach's point of view, however, percent of games won is the more important selective factor. Thus, individual-organism selection works, in this fictional league, to promote selfishness among the players whereas group selection works to

---

[3] Group selection is a controversial explanation for the evolution of altruism in biological evolution. Other mechanisms have been proposed; most involve kin selection and various forms of reciprocity (Nowak, 2006). But group selection allows a fairly straightforward correspondence to be drawn between groups of individuals and patterns of acts within an individual, whereas the others do not. This is not to say that evidence for group selection in behavioral evolution bears in any way on the mechanism underlying altruism in biological evolution.

promote unselfishness; the two selection processes are in conflict. Which will win out depends on the ferocity of competition within the group as compared to that between groups. If low-scoring players are replaced on teams relatively more frequently than low-winning teams are replaced in the league (as is the case in almost any real basketball league), then (other factors held constant) selfishness will come to dominate among players. However, if the selective pressure on teams in the league is more fierce than that of players on teams (for instance, if many teams were competing for fans in a relatively small demographic area), unselfishness would come to dominate (Boyd et al., 2005; Soltis et al., 1995). Another reason that players in real leagues, such as the NBA, *sometimes* come to play unselfishly, despite a very slow rate of team replacement, may be the frequent rate of coach replacement. The coach will be replaced if the team keeps losing games; he or she has an interest in the whole team's performance. The coach and the individual players are therefore to some extent in conflict.

In biological evolution, group selection tends to produce altruism among individuals. Perhaps the clearest illustration of this is a recently published experiment by Shaffer et al. (2016) with harvester ants as subjects. These ants multiply in the environment when queens leave the nests where they were born, mate, and then create new colonies (the evolving entity) where they rear their broods (the generative force).[4] In most ant species, a single queen creates a colony, but in this species multiple, unrelated queens may cooperate in colony creation, maintenance, and defense against raids by ants from other colonies. Competition among these ant colonies is very fierce; the lifespan of a colony is measured in weeks. Shaffer et al. found that in areas where colonies were densely clustered (where competition *between* colonies was conceivably fiercer than competition within them) there were higher proportions of colonies with several queens than there were in colonies that were more sparsely distributed in the environment (where competition *within* colonies was presumably fiercer than competition between them). In other words, highly dense populations of colonies tended to foster queens that cooperated with other queens. *Within* the colonies, more aggressive queens survived longer than cooperative queens but (where they were densely populated) colonies with cooperative queens survived longer than those with non-cooperative queens.

To experimentally test these conclusions, the authors collected queens from predominantly single-queen, and predominantly multi-queen populations of nests. Then, as they say (pp. 2–3):

> *In laboratory observation nests we created six-member foundress associations of two types: 1) pure groups in which all queens came from the [multi-queen] population, and 2) mixed groups with one queen from the [single-queen] population… We observed the queens for 60 days as they excavated nests and reared their first workers. We made daily records of queen survival, brood presence, and any aggressive behavior…. [I]n groups of six containing one or more aggressive queens, initiators of aggression survived significantly longer than their non-aggressive nestmates. In 15 of 27 nests with aggression, the most aggressive queen was the last survivor. Tolerant queens formed the bulk of the victims of aggression, and only 2 of 35 aggressive queens were clearly killed by a cofoundress. However, individuals in groups with only tolerant queens survived significantly longer than individuals in groups with aggression. At the group level, associations with no aggressive members survived significantly longer than those having one or more aggressive queens, with group survival defined as having at least one living queen at the end of the experiment.*

That is, selection at the individual level occurred within groups while selection at the group level proceeded; although there were exceptions, on the whole it was better for an individual to be a member of a group with all-cooperative queens than one with any aggressive ones. Cooperation among queens was altruistic: Any individual queen had a better chance of survival by being aggressive than by cooperating with other queens, but the group had a better chance of survival if that queen cooperated. This is evidence that biological evolution may occur at more than one level.

4. The behavioral evolution of self-control

Consider a squirrel saving nuts during the fall so as to have food available during the coming winter. Because the squirrel prefers nut-storing to other available activities, the squirrel's nut storing is not self-controlled. Mother Nature, in the form of biological evolution by natural selection, arranged it in the past; squirrels that failed to store nuts in the fall tended to die during the coming winter, whereas squirrels that stored nuts (and remembered where they were) tended to live and reproduce. Over generations, squirrels evolved so that they tended to value nut storing for its own sake (rather than for the sake of having food to tide themselves over the winter months). But Mother Nature has not had enough time to work the same sort of magic on humans living in a modern society; it is a rare morning when we wake up and just feel like putting our money in the bank. What she has arranged in our case is a capacity to learn that putting money in the bank is generally a good thing to do, even though we might not feel like doing it at the present moment.

Evolutionary theorists (for instance, Dawkins, 1989; Sober and Wilson, 1998; Wilson and Wilson, 2008) typically focus on biological and cultural levels of evolution—innate mechanisms and cultural rules. They often ignore changes of behavior within an organism's lifetime in response to environmental contingencies (behavioral evolution). Where such changes are not ignored, evolutionary theorists tend to attribute them to inherited developmental and cognitive processes or simply to rational thought. But learning within an organism's lifetime (behavioral evolution) is as much an evolutionary process as biological and cultural evolution (Baum, 2005; Skinner, 1981; Staddon and Simmelhag, 1971). [See discussion of cultural evolution in Section 6]

If a pigeon is food-deprived and gets a pellet of food each time it pecks a lit key, its pecking rate increases. The food (the reinforcer) allows the act it follows to expand in the available time, whereas other acts, not reinforced, die out (are "extinguished"). Once the keypeck appears, it may be molded or "shaped." Just as the giraffe's long neck evolved by the enhanced survival of long-necked giraffes (as they were able to reach higher and higher into trees for food and spot lions over long distances on the African plains, etc.), so the pigeon's key-peck evolves. If the number of pecks needed to obtain the food is gradually increased, the pattern of pecking (the evolving entity) changes. Just as giraffes could conceivably have become extinct if, in the distant past, the height of the edible leaves on the trees had suddenly jumped from ten to twenty feet, so a pigeon's key-pecking would extinguish if the number of pecks needed to obtain the food (the "ratio" requirement) suddenly increased by 100. To shape pecking in numbers near the limit of the pigeon's ability, the ratio must be raised in small stages. The variation of the pigeon's pecking at each stage must include patterns that will be reinforced at the next stage. Selection by reinforcement, like natural selection, works only on behavior that is already present because of variation at a previous stage. That is, selection by reinforcement works on top of natural selection; natural selection provides the mechanism for selection by reinforcement to work. An innate adaptation to temporally extended contingencies comes to us by biological evolution, but the amazing degree of self-control we humans sometimes show is the result of behavioral evolution working on that innate form.

Let us consider a more complex evolving entity. A temporally extended response pattern (think of it as a dance performance) varies in

---

[4] In some instances of biological group selection, migration, along with reproduction, may increase the size of successful groups and thus serve as an additional generative force. However, migration may also destabilize group boundaries, reducing between-group variation, hence lessening the effect of group selection.

small ways along several dimensions. These patterns, as performed by a professional dancer, are highly sensitive to the response of the audience or the feedback from a teacher, or the degree to which they fit some ideal pattern (perhaps unexpressed). The pattern is not selected by the feedback itself. That is, a burst of applause (for instance) does not simply preserve a particular series of steps. Rather, the *relationship over time* of steps to applause (the reinforcement contingency) determines the preserved pattern. The generative process in behavioral evolution is ultimately determined by biological evolution. In an economic model, this process would be maximization of utility under the constraints imposed by the reinforcement contingency. Utility maximization is an answer to the question: “Why is she doing this?” Another answer to this question is given by melioration (Herrnstein and Vaughan, 1980): She is moving from a condition of low reinforcement rate to a condition of high reinforcement rate. According to Baum (2012) her behavior is induced by signals for phylogenetically important events. Which of these proposed generative processes is most useful is an important question, but it is not at issue here. The issue is whether temporally extended patterns of behavior (such as dances) may evolve as such or whether they must be constructed in a chain-like fashion from more basic units as described in the Introduction. In dealing with this question it is important to remember that even a pigeon’s peck on a lit key is a highly complex, temporally extended pattern of behavior. One may ask, “Why should the pigeon bring its head back in a direction away from the key, and away from the food, before it pecks the key? The answer is that pecking is part of an innate and extremely valuable pattern for a pigeon, and bringing the head back is a necessary part of pecking. If you are teaching a pigeon to peck a key, you never need to reward it for bringing its head back.

Like other species, humans are born with a set of innate patterns that may be modified. Post-Darwinian psychologists noted that habits evolve over the lifetimes of individuals in the same way that the structure of a species evolves over generations (Thorndike, 1911/2000; Skinner, 1981). Taking the biological and the behavioral processes of evolution together, if we observe any consistent pattern in the behavior of a human or non-human animal, that pattern must have evolved by natural selection over the history of the species or must have evolved by reinforcement over the history of the individual. (The mechanism of biological evolution has shaped the mechanism of behavioral evolution.) Much of the time, these two evolutionary processes work together—as in the normal development of speech. But, in many cases, especially in our modern society, they work at cross-purposes. Such cases constitute a very fundamental problem of self-control. Behavior that was generally adaptive in the past often conflicts with current contingencies. In the history of the human race, eating as much as you could whenever food became available was a valuable habit; it might be a long time before food became available again. Now, with food continually available, eating whenever you have the chance is bad in many ways (Logue, 1988).

There is an apparent problem in extending the analogy with natural selection to self-control. Where is the reward when I pass by the bakery without entering, even though I am hungry and the smell wafting onto the street is extremely enticing and I have money in my pocket? What in general reinforces self-control? It is clear enough what reinforces lack of self-control. For an alcoholic, drinking is highly rewarding in itself. Drinking, in fact, may reinforce virtually any act the alcoholic is capable of doing. But refusing a drink is far from rewarding in itself and is rarely followed by anything pleasant; in fact, it may entail a great deal of pain. An alcoholic’s single act of drink refusal is not reinforced by better health, better job performance, better social relationships, either immediately nor later. Nor, is it possible to reinforce ones’ own behavior, say by a feeling of pride. Feelings of pride occur, but (as will be argued below) they cannot, *as such*, act as reinforcers. To repeat, if an alcoholic forgoes a single drink, she does not wake up two weeks later healthier or happier, nor is a single drink refusal intrinsically rewarding—just the opposite. The value of drink refusal—in the form of health, social acceptability, job performance, and so forth—depends on an already established habit of refusing drinks; to attain its high value, drink refusal must occur consistently over a period of time. To attain their high value, bakery avoidance and its equivalents must also occur consistently over a period of time. It seems as if self-control would be impossible because the pattern you need to establish (perhaps a very complex pattern) has to exist before its value is attained. Given this problem, how can a complex pattern of behavior evolve by reinforcement?

Let us consider a similar question as it appears in evolution by natural selection: How can very complex structures, such as the human eye, evolve by natural selection? Each part of the eye depends on other parts. It seems as if, until each part develops, the other parts would be useless. The retina cannot function as it does without a lens to focus an image on it; the lens cannot function without a pupil to regulate the light (and what is the function of the pupil if there is no retina?) Finally, an optic nerve has to develop simultaneously with the eye to encode and transmit the image to the brain. And once encoded in the brain, then what? How does a pattern of neural states get converted into a pattern of behavioral discrimination? It is hard to conceive how the human eye could have developed in stages from primitive structures. For creationists, such development is impossible to conceive. *God had to have done it*, they think. But evolutionary biologists, studying primitive eyes in other organisms, have been able to tell a convincing story of how the human eye developed. [A light-sensitive patch of skin (the future retina) would be protected if it was in a depression, and would be more protected the deeper the depression. The more protected, the more sensitive it could be. As the depression grew deeper, into a hole with a light-sensitive bottom, the smaller and smaller entranceway, like a pinhole camera, would begin to focus light onto the retina. Further protection would come from a clear membrane over the hole. That would develop into a lens, and so forth.] Regardless of how complex and seemingly interdependent the parts of a behavioral pattern may be, they can be traced in stages to simpler patterns in development.[5]

To argue that the immense complexities observed in human self-controlled behavior could not have evolved through reinforcement is a kind of creationist argument. The creator, this time, is not God but something within the person—the soul, the self (as an internal entity), or the nervous system (see Baum, 2005, for an argument for the evolution of complex behavior over a person’s lifetime).

One version of the creationist solution to the problem of self-control has been posited by some (I believe misguided) behavioral psychologists themselves. Failing to find extrinsic reinforcers for each component of self-controlled behavior, they posit events inside the person, usually inside the brain, representing reinforcement. A person may be conceived to influence her own choices by exerting an internal ‘will-power’ or relating them to some positive internal concept such as a feeling of pride of accomplishment. In an alcoholic, for instance, such concrete internal representations would immediately reinforce individual drink refusals and counteract the innate immediate value of the drink. Another way to put it is that the reinforcers of self-controlled behavior become “internalized.” It is as if, each time he refuses a drink, the alcoholic internally pats himself on the back for a job well done. An internal vision of the addict’s future enhanced health, for instance, wells up inside him and supposedly reinforces each drink (or cigarette or dessert) refusal. But the concept of self-reinforcement, internal or external, has numerous empirical and theoretical problems. According to a well-supported concept of how reinforcement works (Premack, 1965), a reinforcer is a highly valued activity (such as eating) contingent on the performance of a less-valued activity (such as running). An internal or even an external activity that may be initiated or withheld at any time is by stipulation not contingent on any other activity.

---

[5] Nilsson and Pegler (1994) show that, with current models of biological evolution, these processes could have occurred over surprisingly short time periods.

But, pragmatically, self-reinforcement does seem to sometimes work. In an experiment in my laboratory (Castro and Rachlin, 1980), for example, dieters who took money from a dish next to a scale, proportional in amount to weight lost, lost more weight on average than did dieters who did not take money. However, a third group of dieters, who *put money into the dish* proportional in amount to weight lost, lost even more weight than did those who took money out. It seems that self-punishment was a better self-reinforcer than was self-reinforcement. How can this be? The answer is that self-reinforcement works, where it does appear to work, by emphasizing—making more salient—the response itself. And, putting money into a dish (a loss) is a more salient act than taking an equivalent amount out (a gain).

Another putative problem with the concept of behavioral evolution of self-control (aside from the notion that human behavior is too complex to have evolved) lies in the abstract nature of the pattern. The alcoholic who cuts down from a quart of whisky a day to a glass of wine each night with dinner is not doing anything in particular during the time she is now not drinking; her drinking is just occurring at a much slower rate than before. Each day, the time spent not drinking may be filled with different activities. Every one of these activities cannot individually be more valuable than having a drink. What sustains the lowered drinking rate? The answer is that over time the lowered drinking rate itself is more valuable in the life of the reformed alcoholic than is her old high drinking rate. But, like a probability, the reduced drinking rate is an abstract entity that has no existence at any one moment. You cannot point to it as you would point to an eye. Some alcoholics do learn to abstain or drink moderately. How do the highly abstract patterns in their behavior hold together, and how do they (sometimes) resist being broken up by temptation? The answer is that, like organisms, acts may be selected in groups.

5. Evidence for group selection of patterns of acts in behavioral evolution

As previously indicated, there is much current debate within biology about the extent to which group selection has prevailed within the biological evolution of currently existing species. But if biological group selection were rare or nonexistent in nature, it would be because the particular conditions for it are rare or nonexistent, not because it is impossible. In *behavioral* evolution it seems clear that acts naturally (that is, as a consequence of biological evolution) come and go in groups or patterns rather than as individual muscular movements. Group selection (i.e., multilevel selection) may thus be more common in behavioral than in biological evolution.[6]

Among non-humans there is evidence that reinforcement selects not individual responses but patterns of individual responses. For example, if a rat is rewarded for running faster or slower down an alley or for pressing a lever faster or slower, it changes its rate not by adjusting its speed of running or pressing the lever but by periodically pausing more or less frequently between bursts of running or lever pressing, which, within bursts, remain constant in rate (Gilbert, 1958; Shull et al., 2001). Infants vary their rate of sucking not by altering pauses between individual sucks but by altering pauses between *bursts* of sucks (Wolff, 1968). That is, we have a natural tendency to group acts into patterns. Rats vary their rate of licking at a drinking spout and pigeons vary rates of pecking in a pause-and-burst way (Teitelbaum, 1977). Ethologists have discovered and studied more elaborate "*fixed" action patterns* in consumption and mating. Such patterns are not actually fixed; they are more or less modifiable over the organism's lifetime by environmental contingencies that select patterns or sequences of acts, not individual movements (Hinde, 1966). Over the lifetime of the organism, some patterns are reinforced and may be modified by reinforcement.

Neuringer and his colleagues have shown how patterns of behavior may be directly shaped by environmental contingencies. For example, Grunow and Neuringer (2002) rewarded rats for different levels of variability in sequences of presses on four levers. Rats rewarded for high variability emitted highly variable patterns; rats rewarded for low variability emitted more tightly clustered patterns. The rats were then rewarded in addition for a particular, ordinarily rare, sequence. Those rats initially rewarded for high variability (therefore occasionally emitting the rare sequence) differentially increased the rate of that sequence. The initially rare pattern emerged as a unit from a population of other patterns by a process of differential reinforcement like an island emerging from a receding sea. The environment created by Grunow and Neuringer selected the rewarded sequence, just as the natural environment selects the fittest organisms. Over the course of the experiment, individual presses on each of the four levers were equally reinforced. It is highly unlikely that any sequence could have been learned as a chain of individual presses. This significant experiment and others like it (e.g., Neuringer, 1993) show that patterns of responses may be reinforced as whole units.

An experiment by Locey and Rachlin (2013) with food-deprived pigeons shows how more complex patterns may be shaped from simpler ones. A trial consisted of eight pecks distributed across two keys (L and R) in various patterns. The simplest pattern was at least one switch between the two keys (LR or RL) anywhere during the sequence; the next simplest was at least one instance of LLRR or RRLL anywhere during the sequence; the next was LLLRRR or RRRLLL; the most complex was LLLLRRRR or RRRRLLLL. Note that each more complex pattern contains the simpler ones within it. Initially, all patterns were reinforced, but amount of reinforcement varied directly with complexity of pattern: 1, 2, 4, and 8 s of food access respectively as complexity increased; we represent this reinforcement pattern as: 1-2-4-8. In subsequent conditions, simpler patterns were progressively unreinforced (0-2-4-8; 0-0-4-8; 0-0-0-8) until finally only the most complex pattern (exactly four pecks on one key followed by exactly four pecks on the other) was reinforced (0-0-0-8). Three of the 4 pigeons tested maintained responding under this contingency; responding of the 4th pigeon extinguished. That is, 3 of the 4 pigeons learned and maintained the most complex pattern when it was the only one reinforced. The pigeons typically began the eight-peck sequence by pecking on their *dispreferred* key and then switched to their preferred key during the sequence. They emitted the LLLLRRRR or RRRRLLLL pattern (and received the reward of 8 s of food access) on average on about one-fourth of the 40 trials per daily session for the 40 sessions of this condition. To test whether the most complex pattern evolved from simpler ones, a second group of pigeons was exposed immediately after training to extinction of all patterns except the most complex one: 0-0-0-8. Three of the pigeons failed to maintain responding and the 4th maintained responding at a very low level (often receiving no food during a session) for the duration of the experiment. These results are evidence that response patterns can be shaped (that is, can evolve) as such. The results may serve as a model of how self-controlled behavior can arise through reinforcement.

The question is, can longer and much more complex patterns—those constituting human self-control—evolve from simpler patterns over a person's lifetime, just as the complex human eye evolved from simpler light-sensitive organs? It is not possible to prove that every complex behavioral pattern we exhibit has evolved over our lives from simpler

---

[6] In biological evolution, groups of traits may work together synergistically without group selection (as in heart-lung coordination, for example). In such cases the individual elements of the synergistic adaptation may each be non-adaptive by itself. But if they occurred simultaneously by natural variation their combination could be adaptive. In behavioral evolution of self-control, each element of the adaptive (high valued) pattern is itself of lower value than its alternatives. But these elements are part of a temporal pattern; they cannot change simultaneously. In the analogy with biological evolution, the traits that need to work together to form an adaptive whole would each be not just non-adaptive but maladaptive if they occurred alone. Moreover, each instance (say of cigarette refusal) makes subsequent instances less and less likely. In behavioral evolution, these instances could not occur simultaneously. It is not clear how a synergism between *successive* choices, each one *maladaptive* by itself, could take place without group selection.

patterns [and was not, as in a cognitive explanation, created by an internal logic mechanism overcoming an equally internal "visceral" force (Loewenstein, 1996)]. But when behavior is organized into coherent patterns, self-control may increase. For example, in choosing between an immediate money reward and a non-immediate but larger earning rate, with no 1:1 relationship to any particular choice, people forced to make four choices at a time (which were then played out in sequence), chose the larger overall rate significantly more frequently than did people choosing on a case-by-case basis (Rachlin, 1995). That is, we do know that reinforcement *may* act directly on patterns of acts (without necessitating that chains be built up from individual acts, each followed by its own reinforcer).

In other words, patterns are primary and are, in a sense, causes of their components. I call the view that patterns of (overt) acts may be causes of their components, *teleological behaviorism* (Rachlin, 2017). Teleological causes are answers to the question, "Why…?" For examples, Q: Why are they pitching (or catching or running the bases)? A: Because they're playing baseball. Q: Why are you buying that scarf? A: It's a present for my mother. Q: Why aren't you drinking this drink? A: Because I'm on the wagon. Rachlin (2017) defends the scientific validity of teleological causes in psychology. An important aspect of that defense is that teleological causes are seen not as arising from inside the organism (as efficient causes do), but as patterns in the organism's overt behavior. Taking a teleological and behavioral view of self-control problems goes contrary to our normal way of thinking about self-control. The very term *self-control* seems to imply a force arising from inside—a force opposed to control by the environment. This article argues, however, that a change in such thinking—self-control as adaptation to abstract aspects of the environment, impulsiveness as adaptation to particular aspects of the environment—may be worth the effort involved. This form of behavior analysis fits psychology neatly into evolutionary biology by extending Darwinian evolution to behavioral adaptation within a person's lifetime.

6. Cultural evolution

In behavioral evolution, the spreading-out arrows of Fig. 1 represent patterns of individual behavior extended in time (think of a violinist playing at a recital at different times). In cultural evolution, the arrows of Fig. 1 represent patterns of social behavior, extended in social space, where everyone has a part to play (think of an orchestra playing a symphony at different times). Just as I have been arguing that individual temporal patterns are selected as wholes, so proponents of cultural evolution argue that social patterns are selected as wholes. According to Smaldino (2014), institutions such as schools, training centers, etc. allow patterns of cultural traits to be transmitted "wholesale," which in turn allows for selection to operate on the group-level traits that emerge from "cohesive collections" of individual-level traits. In fact, the justification for these arguments are identical; whereas the conditions for group (i.e., multilevel) selection may rarely be present in biological evolution, they are frequently present in behavioral and cultural evolution. Primary among these conditions is that the natural variation of whole patterns is greater than that of their components. Just as the patterns of behavioral evolution may be selected at various levels of abstraction (hammering a nail, attaching a board to a beam, building a floor, building a house, providing shelter for ones' family, etc.) so social patterns may be selected at more abstract levels (roles of individuals in a squad, squads in a platoon, platoons in a company, companies in a regiment, etc.).

In behavioral evolution, the box of Fig. 1 labeled "generative process" may stand for an innate maximizing tendency or induced behavior brought about by biological evolution. In cultural evolution, the generative process is held to be the teaching that occurs early in life through parents and peers and later in schools and training centers, or by imitation. Just as behavioral evolution rests on biological evolution, so cultural evolution rests on behavioral evolution. The learning that occurs in schools and training centers establishes patterns (often very complex patterns) of individual behavior. In those places, such behavior is reinforced, but it fades away ("extinguishes") unless it is supported by reinforcement outside those institutions. Imitation is an inherited trait, but behavioral patterns acquired solely through imitation extinguish unless reinforced.

7. Learning to be altruistic

As pointed out in the Introduction, altruistic behavior presents a problem for reinforcement-based theories of learning. By definition, an altruistic act is not reinforced; an act cannot be altruistic if the actor is rewarded for doing it. Yet altruistic acts are common in everyday life. Biological evolution can only go so far; it cannot explain how altruistic behavior may be learned over an organism's lifetime. We have discussed how behavioral group selection may explain self-control. We now extend the same argument to altruism. Imagine an environment—say, a jungle—with tribes of people all competing against each other for limited natural resources. Just as with basketball teams (or the ants in Shaffer et al.'s (2016) experiment) tribes within which individuals act altruistically (i.e., possess genes that generate altruistic behavior) will tend to out-compete tribes within which individuals act selfishly. At the same time, within each tribe, selfish individuals (where selfishness is that of the gene—including sacrifice to preserve the lives of blood relatives) will out-compete unselfish individuals. Individuals in successful groups may have predispositions to behave selflessly, cooperatively, altruistically. These predispositions exist along with our selfish tendencies—just as our predispositions to appear attractive, succeed at work, and be healthy exist along with our tendencies to overeat, overdrink, oversleep, and so on.

Because altruistic tendencies involve many genes, and those genes vary, altruistic tendencies vary too. Anyone with eyes to see will note that some people are more altruistic, generous, and cooperative than others. The question I shall now consider is: If you are born with a tendency to be selfish, are you fated to be selfish forever, or can you learn to be altruistic? Are we fated to be altruistic or selfish to whatever degree we were born with—regardless of our experiences? Change in behavior over a person's lifetime (learning) proceeds by means of reinforcement and punishment through the environment. Since this is so, and since altruistic acts are by definition unreinforced by external contingencies, how can altruism possibly be learned? As with self-control, the way to approach this question is to recognize that, parallel to the idea of group selection of *organisms* (responsible for inheritance of altruistic tendencies), there is another kind of group selection—selection by reinforcement of (patterns of) responses over the lifetimes of individuals.

Let us reconsider the question: Does selection by reinforcement act exclusively on individual actions, or can it act on groups of actions? Is there such a thing as group selection of behavior over the lifetime of an individual organism? Again, the answer depends on the ferocity of competition among individual actions and that among groups of actions. In this respect, altruism is no more mysterious than self-control. Although it is not possible to sacrifice your life in the present for your future benefit (except sacrificing yourself to save your children may benefit your genes), it is possible to sacrifice your life for the sake of maintaining a consistent pattern of behavior in which you have heavily invested. Individual acts of altruism that seem inexplicable when considered on a case-by-case basis, such as that of a New York City man who risked his own life by jumping onto the subway tracks to save a stranger's life, are no more mysterious (in principle) than a person stopping for a red light when there are no cars in the cross-street and no police in sight. In both cases, a valuable pattern is being maintained. Making temporal or social decisions on a case-by-case basis is generally a bad policy because, with such a policy, we tend to overestimate the value of a single immediate reward relative to the value of patterns of behavior spread out in time or social space.[7]

The goal of behavioral analysis is to identify environmental reinforcers of acts. But, as just noted, an individual altruistic act has no apparent reinforcer; if it did, it would not be altruistic. Altruism thus seems to defy behavioral analysis. Altruistic acts have been defined, in economic terms, as "… costly acts that confer economic benefits on other individuals" (Fehr and Fischbacher, 2003). This apt definition does not say that the cost to the actor and the benefit to others must be equally valued. And, it does not say whether the "other individuals" are relatives of, friends of, or complete strangers to the actor. If you put a dollar in a vending machine and someone else gets (and eats) the candy bar, your act would be altruistic according to the definition above. Very few of us would pay a dollar to give a perfect stranger a candy bar, but we might very well pay a penny. Or, if the benefit to the stranger were very high (say, he was starving) we might pay a dollar—or even more than a dollar—to give him a candy bar. Or, if the beneficiary were not a stranger but our own child (and we were not at the moment worrying about cavities or obesity), many of us would pay the dollar. Such acts, fitting within Fehr and Fischbacher's definition of altruism, are extremely common in everyday life.

To illustrate how common altruistic behavior is, consider the multiperson prisoner's dilemma game that I have played with audiences over the last 15 years at public lectures. I call it "the lecture game." At the start of the lecture game, blank index cards are handed out to 10 random members of the audience, and the others are asked (as I ask the reader) to imagine that they had received a card. Each of the 10 players is then asked to choose option *X* or option *Y* subject to the following rules (displayed on a slide):

1. If you choose *Y* you will receive $100 times *Z*.
2. If you choose *X* you will receive $100 times *Z* plus a bonus of $300.
3. *Z* equals the number of (the 10) players who choose *Y*.

The audience is told, regretfully, that the money is purely hypothetical; then several properties of the game are pointed out. First, for any particular player, it is always better to choose *X*. By choosing *X*, a player subtracts 1 from *Z* and thereby loses $100 but more than makes up for that loss by the $300 bonus. The net gain for choosing *X* is therefore $200, *regardless of what anyone else chooses.* The point is then emphasized further by saying that any lawyer would advise choosing *X*.

It is then noted that if everyone obeyed their lawyers and chose *X*, *Z* would equal zero and each person would earn just $300, whereas if everyone disobeyed their lawyers and chose *Y*, *Z* would equal 10 and each person would earn $1000; hence the dilemma. It is then pointed out that there is no right or wrong answer and that all choices will be forever anonymous. The 10 audience members with cards are then asked to mark them with *X* or *Y* as they would if the money were real, and the cards are collected.

Over the years I have played this game dozens of times: with college students, with economists (American-capitalist and Italian-socialist), with philosophers, with professors of game theory, with computer scientists, and with psychologists (American, Japanese, and Polish). The median response is 5 *Y*s to 5 *X*s. If there is any bias it is usually in the direction of more *Y*s than *X*s. Standard decision-theory terminology calls choosing *X* "defecting" and choosing *Y* "cooperating," but these loaded terms are not used with the audience.

Although the money earned in the lecture game is hypothetical, laboratory experiments with real money (some with large stakes) have found significant numbers of cooperators in "one-shot," multiperson games such as this one (Camerer, 2003; Fehr et al., 2014; Locey et al., 2011). Because a *Y*-choice always earns $200 less than an *X*-choice, choosing *Y* is a "costly act." Because a *Y*-choice increases *Z* by 1, and each of the 9 other players earns $100 more than he or she would have earned otherwise, choosing *Y* "confers economic benefits on other individuals." Because the choices are completely anonymous, it cannot be claimed that a player's reputation would be enhanced by choosing *Y*. Thus, according to Fehr and Fischbacher's (2003) definition, *Y*-choices in the lecture game are altruistic. Altruism towards one's family members may be explained in terms of common genes (Hamilton, 1964), but it is unlikely that lecture-game players were closely related. Although there is indeed a bias in altruistic behavior toward relatives over non-relatives (Jones and Rachlin, 2008), numerous instances of altruism in everyday life are directed toward friends, acquaintances, and even complete strangers.

8. The extended self

The philosopher Parfit (1984) listed a sample of situations from everyday life modeled by multiperson prisoner's dilemma games such as the lecture game:

> *Commuters: Each goes faster if he drives, but if all drive each goes slower than if all take busses;*
>
> *Soldiers: Each will be safer if he turns and runs, but if all do more will be killed than if none do;*
>
> *Fishermen: When the sea is overfished, it can be better for each if he tries to catch more, worse for each if all do;*
>
> *Peasants: When the land is overcrowded, it can be better for each if he or she has more children, worse for each if all do;*
>
> *There are countless other cases. It can be better for each if he adds to pollution, uses more energy, jumps queues, and breaks agreements; but if all do these things, that can be worse for each than if none do. . . In most of these cases the following is true. If each rather than none does what will be better for himself, or his family, or those he loves, this will be worse for everyone. (pp. 61–62)*

Some of the situations Parfit cites (fishermen, peasants) may be described in terms of Hardin's (1968) "tragedy of the commons"—overuse by individuals of a common resource; all are instances of conflict between individual and social good.

Where does one person end and another begin? Skinner (1969) claimed that the skin is not important as a boundary. I agree with this claim, but for reasons different from Skinner's. Skinner meant that events *within* the skin are subject to behavioral investigation (Zuriff, 1979). Teleological behaviorism is consistent with Skinner's earlier position that the proper sphere of behavior analysis is the organism as a whole (Skinner, 1938). But teleological behaviorism (Rachlin, 2017) sees the organism as extending *beyond* the skin.

Individual actions take time and overlap with each other, giving us the illusion of a continuous self. But, aside from that overlap, there is no "further thing"—either spiritual or physical—to a person's self. You may have as few interests in common with yourself 20 years from now as you do currently with a distant cousin. A person's continuous character is, according to Parfit, no more (or less) real than the character we ascribe to groups of people such as families, clubs, firms, or nations. He says (p. 211): "Most of us are reductionists about nations… Nations exist. Though nations exist, a nation is not an entity that exists separately, apart from its citizens and its territory."

What, then, does tie our actions together, if not a central self? The answer is that abstract, broadly based social situations signaling consistent reinforcement contingencies ("meta-discriminative stimuli") can tie a person's actions together over extended periods. If, as I maintain, the self were actually a set of patterns in a person's behavior controlled by meta-discriminative stimuli, those patterns, hence the self, could extend over significant intervals in a person's life.

If your future self is in principle no closer to your present self than is another person, it follows that there is no *essential* difference between

---

[7] This is also true because the value of many actions in a complex world are causally opaque. Thus the importance of ritual and tradition in cultural evolution (Richerson and Boyd, 2005).

your relations to your future self and your relations to other people. Since your concern or lack of concern for other people involves *moral* issues, Parfit says, so does your concern or lack of concern for your future self. Therefore, issues of social cooperation and altruism on the one hand and self-control on the other are treated in the same way. A motorcyclist's refusal to wear a helmet would be, for Parfit, a moral issue, not just because other people may have to pay a price for her brain injury, but also because her future self may have to pay a price.

For a teleological behaviorist, a person's self consists of the temporal extension and overlap of that person's various overt behavioral patterns (Rachlin, 1992, 1994, 2010, 2014). In human as well as non-human existence, many behavioral patterns are coordinated with those of others). Their overlap—their common interest—extends our selves nearer or further into our society. Thus, altruism, like self-control, may be learned over a person's lifetime by group selection in behavioral evolution.

Acknowledgements

This article is revised from Chapter 7 of: Rachlin, H. (2014) *The escape of the mind*. New York: Oxford University Press, with permission of the publisher. Preparation was supported by Grant 1630036 from The National Science Foundation.

References

Baum, W.M., 2005. Understanding Behaviorism: Behavior, Culture, and Evolution, 2nd ed. Blackwell, Oxford.

Baum, W.M., 2012. Rethinking reinforcement: allocation, induction, and contingency. J. Exp. Anal. Behav. 97 (1), 101–124.

Boyd, H., Gintis, H., Bowles, S., Richerson, P.J., 2005. The evolution of altruistic punishment. In: Gintis, H., Bowles, S., Boyd, R., Fehr, E. (Eds.), Moral Sentiments and Material Interests: The Foundations of Cooperation in Economic Life. MIT Press, Cambridge, MA, pp. 215–228.

Camerer, C.F., 2003. Behavioral Game Theory: Experiments in Strategic Interaction. Princeton University Press, Princeton, NJ.

Castro, L., Rachlin, H., 1980. Self-reward, self-monitoring, and self-punishment as feedback in weight control. Behav. Ther. 11 (1), 38–48.

Dawkins, R., 1989. The Selfish Gene, 2nd ed. Oxford University Press, New York.

Fehr, E., Fischbacher, U., 2003. The nature of human altruism. Nature 425, 785–791.

Fehr, E., Tougareva, E., Fischbacher, U., 2014. Do high stakes and competition undermine fair behaviour? Evidence from Russia. J. Econ. Behav. Organ. 108, 354–363.

Gilbert, T.F., 1958. The fundamental dimensions of the operant. Psychol. Rev. 65, 272–285.

Grunow, A., Neuringer, A., 2002. Learning to vary and varying to learn. Psychon. Bull. Rev. 9, 250–258.

Hamilton, W.D., 1964. The genetical evolution of social behaviour. I. J. Theor. Biol. 7, 1–16.

Hardin, G., 1968. The tragedy of the commons. Science 162, 1243–1248.

Herrnstein, R.J., Vaughan, W., 1980. Melioration and behavioral allocation. In: Staddon, J.E.R. (Ed.), Limits to Action: The Allocation of Individual Behavior. Academic Press, New York, pp. 143–176.

Hinde, R.A., 1966. Animal Behaviour: A Synthesis of Ethology and Comparative Psychology. McGraw-Hill, New York, pp. 331–360.

Johnson, R.A., Rissing, S.W., Killeen, P.R., 1994. Differential learning and memory by co-occurring ant species. Insectes Soc. 41 (2), 165–177.

Jones, B.A., Rachlin, H., 2008. Altruism among relatives and non-relatives. Behav. Process. 79, 120–123 (PubMed Central #56111).

Locey, M.L., Rachlin, H., 2013. Shaping behavioral patterns. J. Exp. Anal. Behav. 99, 245–259.

Locey, M.L., Jones, B.A., Rachlin, H., 2011. Real and hypothetical rewards in self-control and social discounting. Judgm. Decis. Mak. 6, 552–564 (NIHMS 370506).

Loewenstein, G., 1996. Out of control: visceral influences on behavior. Organ. Behav. Hum. Decis. Process. 65, 272–292.

Logue, A.W., 1988. Research on self-control: an integrating framework. Behav. Brain Sci. 11, 665–679.

Neuringer, A., 1993. Reinforced variation and selection. Anim. Learn. Behav. 21, 83–91.

Nilsson, D., Pegler, S., 1994. A pessimistic estimate of the time required for an eye to evolve. Proc. R. Soc. 256, 53–58. http://dx.doi.org/10.1098/rspb.1994.0048/rspb.1994.0048.

Nowak, M.A., 2006. Five rules for the evolution of cooperation. Science 314 (5805), 1560–1563.

Parfit, D., 1984. Reasons and Persons. Oxford University Press, Oxford.

Rachlin, H., Jones, B.A., 2008. Altruism among relatives and non-relatives. Behav. Process. 79 (2), 120–123.

Rachlin, H., 1992. Teleological behaviorism. Am. Psychol. 47, 1371–1382.

Rachlin, H., 1994. Behavior and Mind: The Roots of Modern Psychology. Oxford University Press, New York.

Rachlin, H., 2010. How *should* we behave: a review of reasons and persons by derek parfit. J. Exp. Anal. Behav. 94, 95–111.

Rachlin, H., 2014. The Escape of the Mind. Oxford University Press, New York.

Rachlin, H., 2017. In defense of teleological behaviorism. J. Theor. Philos. Psychol. 37 (2), 65–76.

Richerson, P.J., Boyd, R., 2005. Not by Genes Alone. University of Chicago Press, Chicago.

Shaffer, Z., Sasaki, T., Haney, B., Janssen, M., Pratt, S.C., Fewell, J.H., 2016. The foundress's dilemma: group selection for cooperation among queens of the harvester ant, *Pogonomyrmex californicus*. Sci. Rep. 6, 29828. http://dx.doi.org/10.1038/srep29828.

Shull, R.L., Gaynor, S.T., Grimes, J.A., 2001. Response rate viewed as engagement bouts: effects of relative reinforcement and schedule type. J. Exp. Anal. Behav. 75 (3), 247–274.

Skinner, B.F., 1938. The Behavior of Organisms. Appleton-Century, New York.

Skinner, B.F., 1969. Behaviorism at Fifty. Appleton-Century-Crofts, New York.

Skinner, B.F., 1981. Selection by consequences. Science 213, 501–504.

Skinner, B.F., 2014/1969. Contingencies of Reinforcement: A Theoretical Analysis (Vol. 3). BF Skinner Foundation.

Smaldino, P.E., 2014. The cultural evolution of emergent group-level traits. Behav. Brain Sci. 37 (3), 243–254.

Sober, E., Wilson, D.S., 1998. Unto Others: The Evolution and Psychology of Unselfish Behavior. Harvard University Press, Cambridge, MA.

Soltis, J., Boyd, R., Richerson, P.J., 1995. Can group-functional behaviors evolve by cultural group selection? An empirical test. Cult. Anthropol. 36, 473–494.

Staddon, J.E.R., Simmelhag, V., 1971. The superstition experiment: a reexamination of its implications for the principles of adaptive behavior. Psychol. Rev. 78, 3–43.

Teitelbaum, P., 1977. Levels of integration of the operant. In: Honig, W.K., Staddon, J.E.R. (Eds.), Handbook of Operant Behavior. Prentice-Hall, Englewood Cliffs, NJ, pp. 7–27.

Thorndike, E.L., 1911/2000. Animal Intelligence. New Brunswick, NJ: Transaction Publishers.

Wilson, D.S., Wilson, E.O., 2008. Evolution for the good of the group. Am. Sci. 96, 380–389.

Wolff, P.H., 1968. The serial organization of sucking in the young infant. Pediatrics 42, 943–956.

Zuriff, G.E., 1979. Ten inner causes. Behaviorism 7, 1–8.