

UC Davis

UC Davis Previously Published Works

Title

Selection by consequences, behavioral evolution, and the price equation

Permalink

<https://escholarship.org/uc/item/7d28p1cq>

Journal

Journal of the Experimental Analysis of Behavior, 107(3)

ISSN

0022-5002

Author

Baum, William M

Publication Date

2017-05-01

DOI

10.1002/jeab.256

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at

<https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

SELECTION BY CONSEQUENCES, BEHAVIORAL EVOLUTION, AND THE PRICE
EQUATION

WILLIAM M. BAUM

UNIVERSITY OF CALIFORNIA, DAVIS AND UNIVERSITY OF NEW HAMPSHIRE

Price's equation describes evolution across time in simple mathematical terms. Although it is not a theory, but a derived identity, it is useful as an analytical tool. It affords lucid descriptions of genetic evolution, cultural evolution, and behavioral evolution (often called "selection by consequences") at different levels (e.g., individual vs. group) and at different time scales (local and extended). The importance of the Price equation for behavior analysis lies in its ability to precisely restate selection by consequences, thereby restating, or even replacing, the law of effect. Beyond this, the equation may be useful whenever one regards ontogenetic behavioral change as evolutionary change, because it describes evolutionary change in abstract, general terms. As an analytical tool, the behavioral Price equation is an excellent aid in understanding how behavior changes within organisms' lifetimes. For example, it illuminates evolution of response rate, analyses of choice in concurrent schedules, negative contingencies, and dilemmas of self-control.

Key words: Price Equation, behavioral evolution, induction, concurrent schedules, induced activity, operant activity, reinforcement, negative contingencies, self-control

Often mathematics expresses clearly and concisely what many words cannot. Price's equation is an example. George Price (1970, 1972), using straightforward algebra, derived an equation that expresses evolution by natural selection in precise mathematical terms. Recognition of its significance has grown steadily since it was published in 1970 (Frank, 1995; Panchanathan, 2011). Nowadays, it is part of the standard toolkit for clarifying evolution at multiple levels, and discussions of multilevel selection often rely on it (e.g., McElreath & Boyd, 2007). A large part of the Price equation's power derives from its generality; because it describes selection in the most general terms possible, it applies to any Darwinian process of evolution by selection and even to inanimate entities (Price, 1972, 1995). Behavior analysis is a science in which evolution by selection is central, because a basic premise of behavior analysis is that behavioral change depends on selection by consequences (Skinner, 1981). Thus, Price's equation should be interesting and useful to behavior analysts. Indeed, it may be thought of as restating or even replacing the law of effect, including the "correlation-based" law of effect (Baum, 1973, 2012).

This paper shows how Price's equation may be applied to behavioral evolution (often called "shaping by consequences"). The first section gives an overview of Darwinian processes to give a framework for thinking about behavioral evolution. The second section presents three versions of the Price equation, applying to genetic evolution, cultural evolution, and behavioral evolution. The third section presents examples that illustrate the approach to behavioral evolution: response rates on variable-interval schedules, choice in concurrent schedules, self-control, and negative contingencies.

Overview of Darwinian Processes

Every Darwinian process includes three elements: Variation, Recurrence, and Selection.

Variation occurs within a pool or population. Examples are: (a) a gene pool, which is an abstraction consisting of all the genes in a population regardless of their occurrence in particular organisms; (b) a pool of genomes, in which each member of the population constitutes a "vehicle," carrying the genome around; (c) a pool of cultural practices, in which the members of the cultural group also are like vehicles, putting the practices into action; and (d) the coat colors that vary among all the mice (the "vehicles") that live in a particular meadow. Although variation is commonly thought of as operating on punctate entities

Correspond with: William M. Baum, 611 Mason #504, San Francisco, CA 94108, Email: wbaum@sbcglobal.net
doi: 10.1002/jeab.256

like genes, nothing restricts variation to punctate entities, because the required variation can be in any property of the vehicles of the population. For behavioral evolution, the “vehicles” are not organisms, but samples of behavior that vary, for example, in the rate of an operant response, as will become clear below. The Price equation (below) clarifies this flexibility.

Recurrence means “occurring again” or “coming up again.” In genetic evolution, recurrence results from reproduction, sexual or asexual, that causes traits to recur from parents to offspring. Reproduction is only one mechanism of recurrence (Baum, 2001). Recurrence is a more abstract and general term for what is required in an evolutionary process—that types within a population recur from one time to the next. When we turn to cultural evolution, for example, the mechanisms of recurrence are imitation and rule following (i.e., instruction), or a combination of the two. Such recurrence is sometimes referred to as “transmission.” Cultural practices recur from one time to another (often within a generation) and from one person to another by one person modeling behavior for another (e.g., a dance movement, a ritual, or a method of manufacture) or by one person telling another, “Behave this way” (e.g., how to dress, how to approach the king, or how to do arithmetic), which constitutes rule giving that induces rule following in the receiver (Baum, 2000). Cultural recurrence differs from genetic recurrence too because cultural recurrence is not necessarily, nor even usually, from parent to offspring. One’s cultural “parents” go beyond one’s genetic parents. One may acquire behavior from peers and from members of an older generation other than one’s genetic parents, such as teachers, coaches, employers, and ministers (Boyd & Richerson, 1985). One may acquire behavior even from a stranger viewed in an online video.

In behavioral evolution, recurrence means repetition of behavior from time to time. Evolutionary change and the Price equation only require some kind of recurrence; the mechanism of recurrence varies from one evolutionary process to another. For behavior, the traditional view ascribes recurrence to “strengthening” (Skinner, 1938; Thorndike, 1911/2012). An earlier paper suggested that the mechanism of behavioral recurrence is *induction*, the

term introduced by Segal (1972) to denote the relation between a phylogenetically important event (PIE) like food and the (adjunctive) behavior it induces (Baum, 2012). Although “induction” has had other usages, Segal’s usage is particularly useful, because it applies not only to adjunctive behavior, but also to operant behavior, the main focus of behavioral evolution. In Baum’s (2012) explication, the contingency between a PIE like food and an operant activity like lever pressing enforces a correlation that causes the activity to be induced by the PIE it produces (e.g., Reid, 1958; “reinstatement”). This insight will be useful when we come to examples of behavioral evolution below.

Selection depends on competition between different variants within the pool or population. It is not the only force by which a population’s traits may change; drift due to small population size, for example, has little to do with competition (Sober, 1984). Evolutionary change, however, depends on competition due to differences among variants in fitness or recurrence within a population.

Competition is the heart of selection. Since the pool or population typically is limited (e.g., by carrying capacity of the environment) or only grows relatively slowly (e.g., a cultural group), substitutable variants must compete through time because increases in one forces decreases in others. Their competition makes their fitness or recurrence relative rather than absolute; the fitness of any variant depends on its competitors. Short-necked giraffes might survive and reproduce well on their own, but if longer-necked giraffes recur more often by leaving more offspring, shorter-necked giraffes disappear from the population. Different alleles of a gene, different ways of worshiping or making a canoe, and different activities that produce different outcomes—pecking at one key rather than another or spending time with family versus working—all compete because of limitations on resources or time. If one increases, its competitors must decrease. Fitness is always relative.

For evolutionary change, the variants that may substitute for one another are the variants that matter; other variants may be independent of them. Darker coloration may compete with lighter coloration, but not necessarily be linked to larger body size. The practice of making internal-combustion-driven cars may

compete with the practice of making electric-motor-driven cars, but neither may have any bearing on methods of birth control. No matter what the mechanism, selection of one variant (trait, practice, or activity) over others occurs when competition results from differences in their fitness or differential recurrence through time. If larger body sizes recur more often than smaller body sizes, larger body sizes increase in the population at the expense of smaller body sizes. If members of a cultural group eat either with chopsticks or a fork, and eating with a fork recurs (spreads) in the group more often than eating with chopsticks, eating with a fork increases in frequency at the expense of eating with chopsticks. If my driving on the highway to work recurs in my drives to work more often than my driving on back roads, driving on the highway increases at the expense of driving on back roads.

Cultural practices constitute temporally extended patterns of behavior that accomplish outcomes important to the survival or cohesion of the group that practices them; they are ways of getting vital "jobs" done (Guerin, 1997). A cultural practice is an integral whole with parts that function together to bring about their desirable result. For example, manufacturing a certain kind of ceramic pot requires finding and processing clay, working the clay, decorating the pot, and firing the pot. The practice may be refined by substituting a more effective variant for one of the parts—say, finding a new type of clay or a new mode of decorating the pot. The parts make a whole that may change as the parts change.

Similarly, an operant activity is an integral whole that produces results useful to the organism (Baum, 2002). Indeed, cultural practices are operant activities—they are distinguished only by their being a property of a group. A skilled operant performance is refined the same way that a cultural practice is refined, usually by substituting a more efficient variant for one less efficient, as when an activity like swinging a baseball bat becomes more effective by the substitution of one muscular motion for another.

The combination of contingency and induction results in one operant activity having higher recurrence ("fitness") than the variants that compete with it. We will take up examples later. Right now, we only need to recognize that, even though the physiological basis of

behavioral recurrence may be unknown, the behavior selected by the contingency combined with induction may be considered fitter than its competitors and to have more "descendants" at a later time.

Price's Equation

Genetic Evolution

Price (1970) derived the equation by proving a theorem. He made three assumptions—variation, "transmission" (Price's word for recurrence), and selection. First, variation: a population of N individuals at two time periods or "generations," Time 1 and Time 2. A trait or property, x , of these individuals varies among them, and the variable x_i indicates the level of x in individual i . At Time 1, the mean of x across the N individuals is denoted $\bar{x}$. Second, transmission: offspring resemble their parents, which means that offspring of a parent with trait-level x_i also have trait-level x_i . Third, selection: Each individual i may contribute surviving offspring to the population at Time 2. The fitness or relative recurrence of individual i depends on how many of its offspring survive to Time 2. An individual that contributes no surviving offspring has a fitness of zero. Others have fitness proportional to the number of offspring they contribute. The fitness of individual i is symbolized by w_i , which varies across individuals. Price's great insight was to see that the change in the mean of x in the population from Time 1 to Time 2, symbolized as $\Delta\bar{x}$, depends simply on whether individuals high in the trait x are also high in fitness w and individuals low in the trait x are also low in fitness w —that is, whether x_i and w_i vary together—their covariance.

Figure 1 illustrates Price's reasoning with a simple example. The first row of the table shows a population of 10 individuals, with i ranging from 1 to 10. Suppose all these individuals possess either one of two variants of a certain gene—its alleles, A or B. Seven of the individuals possess the A allele, and 3 possess the B allele, as we see in the second row of the table. If an individual possesses the A allele, we set $x_i = 0$, and if an individual possesses the B allele, we set $x_i = 1$, as shown in the third row of the table in Figure 1. The bottom row gives various arbitrarily selected numbers of offspring contributed to the population at Time 2, w_i , for the 10 individuals. For




Fig. 1. A simple example illustrating Price’s reasoning. In the top row of the table, i indexes individuals in a population of 10 individuals. Each individual possesses either the A allele or the B allele of a gene, as shown in the second row of the table. If an individual possesses the A allele, x_i is set to zero, and if an individual possesses the B allele, x_i is set to 1.0, as shown in the third row of the table. The 10 individuals vary in fitness, which is the number of offspring or copies of the allele they possess at Time 2, w_i , shown in the bottom row of the table. The graph below shows fitness w_i plotted against x_i (X symbols). The squares represent the average fitnesses of the two alleles. The slope of the line shows that covariance is greater than zero and selection occurs.

example, Individual 1 contributed one offspring, Individual 2 contributed no offspring, Individual 3 contributed two offspring, and so on. Since the number of offspring contributed by the A-allele individuals total to 6, their mean fitness equals $6/7$. Since the number of offspring contributed by the B-allele individuals total to 4, their mean fitness equals $4/3$, greater than for the A allele. In the graph below the table, these fitnesses w_i are plotted against x_i , and the line connects the two mean fitnesses. The key point of the graph is that the line has a positive slope; individuals’ fitness w_i covaries with individuals’ trait-level x_i .

Figure 2 shows two additional examples of covariance between fitness and a trait x . The top graph shows an extension of the reasoning behind the graph in Figure 1 to diploid organisms. The x-axis represents the number of copies of the B allele that different genotypes possess: Individuals homozygous in A have no copies of B; heterozygous individuals have one copy of B; and individuals homozygous in B have two copies of B. If the more copies of allele B an individual possesses the higher the fitness, the covariance between w_i and x_i would




Fig. 2. Two examples of covariance between fitness w_i and x_i . Top: an example with two alleles of a gene, as in Figure 1, but with diploid organisms. Individuals homozygous in the A allele possess no copies of the B allele. Heterozygous individuals possess one copy, and individuals homozygous in the B allele possess two copies. If fitness tends to increase with number of copies of B, the slope of the line connecting the average fitnesses (squares) shows that covariance is positive and selection occurs. Bottom: an example in which x is a continuous variable and x_i may take on any level within its range. The positive slope of the line fitted to the fitnesses or recurrences w_i at Time 2 shows that covariance is greater than zero and selection occurs.

look like that shown. The bottom graph shows the situation in which x is a continuous variable like body size, height, or coat color. Then x_i can take on any number of levels, resulting in a scattering of points like that shown.

This covariance is the heart of the Price equation:

$$\Delta \bar{x} = cov(w_i, x_i) \tag{1}$$

which states that the population change in x equals the covariance between w and x across individuals.

The full Price equation contains another term:

$$\Delta \bar{x} = \text{cov}(w_i, x_i) + \sum_i w_i \Delta x_i / N \quad (2)$$

where Δx_i is the change in x_i from Time 1 to Time 2 due to imperfect fidelity of recurrence; if recurrence were just the passing on of x_i from one time to the other, Δx_i would be zero. Since the second term on the right is a mean, it is usually rewritten as an expected value $E_w(\Delta x_i)$, leading to the simplest form of the Price equation (see Price, 1970, 1972, and McElreath and Boyd, 2007, chapter 6, for more detail):

$$\Delta \bar{x} = \text{cov}(w_i, x_i) + E_w(\Delta x_i) \quad (3)$$

Equation 3 says that the population's mean change in x equals the dependence of fitness on x_i plus an expected value that gives the mean change in x apart from the covariance. The covariance represents selection; if relative fitness varies positively with x_i , $\Delta \bar{x}$ is positive, and the trait increases in the population. If the covariance is negative, the trait decreases, and if the covariance is zero, the trait remains stable. The second term on the right, which expresses the tendency of x to change as it recurs from Time 1 to Time 2, due to factors usually rare or negligible (e.g., mutation or meiotic drive), is usually considered close to zero for genetic evolution (El Mouden, André, Morin, & Nettle, 2014; Henrich, 2004).

Put in more general terms, Equation 3 says that, if recurrence is differential with respect to various levels of a trait, and fidelity of recurrence is high, then the trait recurs increasingly (positive covariance) or recurs decreasingly (negative covariance) in the population from one time to the next. Suppose that the population of mice that live in a certain meadow vary in coat color and that darker coat colors are less visible to predators. This lower visibility means that coat color varies with fitness. If we think of x as darkness of coat color, the covariance between darkness and fitness is positive, and Equation 3 says that mean coat color in that population grows darker from one generation to the next. Since the mice in the meadow constitute a population not structured into groups, however, this example

applies Equation 3 without any reference to group selection.

Three points about Equation 3 should be apparent. First, the equation is true, because it is a derived identity (see Appendix A). It describes or defines rather than predicts, but it is a valuable analytical tool for thinking about selection, for example, group selection, when a population is structured into competing groups. Its theoretical aspect lies in the assumptions made when applying the equation. The basic assumptions appear to be minimal—variation in a population of individuals, recurrence, and covariance—but some may be hidden, as for example, the assumption that the environment remains stable from one time to the next. The concept of fitness is crucial to the equation. Fitness is relative and, because it is relative, fitness is dimensionless (Appendix A). As a result, the two sides of Equations 1, 2, and 3 have the same units. Second, it neatly includes all three ingredients of the Darwinian process: Variation appears in x_i as it varies from individual to individual or vehicle to vehicle of the population; recurrence appears in the term w_i ; and selection appears in the covariance between x_i and w_i —in differential recurrence across variation in x_i . Third, Price's equation is extremely general. It doesn't specify the mechanism of recurrence, the nature of the individuals that vary in x_i , whether x_i is a discrete or continuous variable, or the factors that favor higher x_i over lower x_i (differential recurrence or fitness w_i).

Like the law of effect, like many scientific laws, the Price equation is a tautology. For example, Newton's law $F = ma$ is a tautology. Rachlin (1971) argued that the choice matching law might be seen as a tautology. Indeed, natural selection may be viewed as a tautology (Sober, 1984). Just as an activity that increases when in contingency with a reinforcer is said to be reinforced, and the reinforcer is known by its ability to increase activities when in contingency with them, so a variant in a population that increases has higher fitness, and fitness is known by its ability to increase variants within a population. Tautology is not a defect. Tautological laws may be highly useful generalizations, conceptual frameworks, and explanatory frameworks (Sober, 1984). Among other possibilities, a tautology may allow us to grasp and apply selection by consequences.

Group Selection

Suppose that a group under consideration is one group among many, because the population is structured into groups. A good example is a bee colony; many colonies might inhabit a geographical area, and they compete for resources and succeed or fail as wholes. Another example is a human social group, such as a baseball team, which competes with other teams and wins or loses as a whole. Human groups compete both directly, as in warfare, and indirectly, as in exploiting limited resources poorly or well. Since Equation 3 does not depend on the sort of individuals, Equation 3 would also apply to selection among bee colonies and other social groups that function as integrated wholes. For the sake of clarity, we index groups with the subscript g , instead of i , and apply Equation 3 to get an expression for $\Delta\bar{x}$, now the change in average trait x across all the groups in the population. The equation says that if the covariance between x_g and fitness w_g is positive, the prevalence of x in groups with a higher mean of x increases, and such groups survive, grow in size, and multiply relative to groups with lower mean x . For example, bee colonies that better resist invasion by mites will tend to survive and grow better than those with less resistance, and baseball teams that coordinate well together will tend to win and gain resources against teams that are less coordinated.

In applying Equation 3 to group selection, however, the second term on the right, which would be written $E_w(\Delta x_g)$, is no longer negligible, because the change in x_g apart from group selection includes changes in individual members due to selection at the level of individual members. Recognizing that x_g is a mean across members in the group that may change from Time 1 to Time 2, we treat Δx_g the same way as we treated Δx (see Appendix A) and arrive at:

$$\Delta\bar{x} = \text{cov}(w_g, x_g) + E_g[\text{cov}(w_{ig}, x_{ig}) + E_w(\Delta x_{ig})] \quad (4)$$

where the subscript ig indicates member i in group g (see the derivation in Appendix A and Henrich, 2004, and McElreath & Boyd, 2007, chapter 6, for more detail). The second term on the right of Equation 4 is really just the average across groups (E_g) of Equation 3

for each group. It accounts for all change within the population at the level of the individuals making up all the groups.

The two terms on right-hand side of Equation 4 represent two components of change in x in the population: the first for selection between groups and the second for selection between members within groups, the expected value term. For genetic evolution, we may ignore the expected value of Δx_{ig} in Equation 4 as we did the expected value of Δx_i in Equation 3, because the change would be due to rare or negligible factors such as meiotic drive and mutation (El Mouden et al., 2014; Henrich, 2004). Equation 4 then simplifies to the sum of a covariance for between-group selection and expected value (or average) of covariance for within-group selection:

$$\Delta\bar{x} = \text{cov}(w_g, x_g) + E_g[\text{cov}(w_{ig}, x_{ig})] \quad (5)$$

An alternative form of the Price equation that uses the algebraic relation between covariance and the regression coefficient β may further clarify group selection (see Henrich, 2004, and McElreath & Boyd, 2007, chapter 6, for more detail):

$$\Delta\bar{x} = \beta(w_g, x_g) \text{var}(x_g) + E[\beta(w_{ig}, x_{ig}) \text{var}(x_{ig})] \quad (6)$$

Equation 6 makes clear that the dependence (β) between w and x and the variance in x must both be greater than zero for selection to occur. Each term on the right implies a role for each of the three ingredients of evolution: variation ($\text{var}(x)$), recurrence (w), and selection (β).

When applied to genetic evolution, Equation 6 clarifies why genes for altruistic and cooperative behavior are unlikely to be selected by genetic evolution on its own. Some traits may have both between-group and within-group components that act in concert—both have positive β —such as body size (helpful in intragroup and intergroup combat) or tendency to stay close to the rest of the group (avoiding predators). Genes for altruism or cooperation, however, though helpful for the group, reduce the fitness of individuals that behave so—at least in the short term (Baum, 2016). Thus, for altruism or cooperation, β is positive for between-group

selection, but negative for within-group selection. Since variance across groups $var(x_g)$ in particular is lowered by migration between groups, the groups have to be practically completely isolated genetically for group selection to overcome negative individual selection. For a bee colony, variance within the colony $var(x_{ig})$ is small, because the workers are close relatives, and no migration occurs between colonies, so altruistic behavior dominates. For human groups, however, negative β is likely to be high for any individual incurring an immediate cost for helping or cooperating with others, and within-group variance $var(x_{ig})$ is likely to be substantial. Even a small amount of migration between groups will ensure that between-group selection is smaller than the within-group selection, making $\Delta\bar{x}$ for that altruistic behavior negative. In contrast to selection on genes, however, when we come to cultural evolution and behavioral evolution, a means for altruistic behavior to be readily selected will become clear. Another way of using the Price equation to describe individual and group selection, known as “contextual analysis,” is briefly explained in Appendix B.

Cultural Evolution

Equation 3 may be applied to cultural evolution, too, with suitable changes in definition. A cultural practice that might be considered all-or-none (e.g., allowing marriage between first cousins or not, eating pork or not, and primogeniture or not) would take the place of an allele that is present or not (as in Fig. 1). As with genetic evolution, however, quantitative variation in a practice (e.g., a ritual performed daily or less often, for longer or shorter duration, with more or less of some ingredient, etc.) would do just as well for the variable x (as in Fig. 2, bottom).

Recurrence of a cultural practice from Time 1 to Time 2 doesn't go just from parent to offspring, as in genetic evolution. Teachers, coaches, ministers, and peers also both model practices and instruct practices that have proven successful in gaining the proximate tokens of ultimate fitness (reinforcers and avoidance of punishers). Particularly if these influential people are prestigious and display the trappings of success (many reinforcers and few punishers), this modeling and instructing spur the recurrence of the practice at Time

2, denoted, as before, w_i . As before, if covariance between w_i and x_i is positive, the practice spreads in the population. The second term on the right of Equation 3 might be considered non-negligible if one wished to incorporate errors on the part of recipients, but usually may be considered negligible because transmission errors would likely be insignificant compared with selection (see El Mouden et al., 2014; supplemental Document S1; for a detailed derivation of Eq. 4 for culture and further discussion).

Cultural Group Selection

Equation 3 would be about the spread of a practice through a cultural group (e.g., spread of planting hybrid corn among American farmers; Rogers, 1995). If we think of cultural groups as parts of a larger population, the groups may compete with one another for resources and members. Equations 4, 5 and 6 apply then, with the same changes in meaning of the variables. Equation 5 contains the same two components of intergroup selection and within-group selection (El Mouden et al., 2014; Henrich, 2004). If a cultural practice benefits the group relative to other groups, the more members practice it, or the greater the recurrence (w_g) in the group at Time 2, the more positive is the covariance in the first term in Equation 5. Whether the covariance in the second term is positive or negative depends on whether the practice is reinforced or is costly for the individual who engages in it.

To understand how an altruistic or cooperative practice can spread in a population, Equation 6 shows what happens if the within-group covariance for members is negative. As with genetic group selection, the variances $var(x_g)$ and $var(x_{ig})$ are crucial. In cultural group selection, biases in imitation and instruction (“transmission biases”) that promote conformity lower $var(x_{ig})$ to near zero, even in the face of considerable migration between groups. For example, if immigrants into a group tend to adopt the practice that is most frequent or above average in the group or tend to adopt the practice of prestigious or successful group members, their conformity will keep $var(x_{ig})$ low (Henrich, 2004). Low $var(x_{ig})$ allows the cost to the individual represented by the negative covariance (β) of the second term to be

offset by the positive covariance (β) for the group, which allows altruistic and cooperative practices to spread in the population. The group in which members sacrifice for the group will gain more resources than other groups, will increase in numbers at the expense of other groups, will be more likely to survive, and will be more likely to multiply (Henrich, 2004; Richerson et al., 2016).

Equations 5 and 6 represent “multi-level” selection, but just at two levels—groups and individuals (McElreath & Boyd, 2007; chapter 6). A practice may have many parts, and some of these parts may change without the practice taking on a new identity. For example, a manufacturing process or a ritual may drop or add elements; the American Pledge of Allegiance after many years changed to include the words “under God,” but it is still the Pledge of Allegiance, and process improvement constantly upgrades manufacture of automobiles, but the practice remains the manufacture of automobiles (Baum, 2000, 2002). When practices change in their parts, one might consider a third level of selection: within-practice selection. In Equation 4, Δx_{ig} may be treated as change due to competition among alternative parts within the practice and expanded as before to result in a three-term Price equation (McElreath & Boyd, 2007; chapter 6).

Behavioral Evolution

Price (1995; written circa 1971) opens a general discussion of selection, “Selection has been studied mainly in genetics, but of course there is much more to selection than just genetical selection. In psychology, for example, trial-and-error learning is simply learning by selection” (p. 389). If we substitute for “trial-and-error learning” instead “shaping by consequences,” we see that Price was thinking of operant behavior. Thus far, we have seen how the Price equation can clarify a variety of selection processes. As far as I know, no one has applied the equation to behavioral evolution. The entities that may vary would no longer be individual organisms, because we are dealing with the behavior of a single organism. An organism’s behavioral patterns vary, recur in time, and are selected when they recur differentially—i.e., when they vary in fitness.

To think of behavioral evolution with the Price equation, we require a population of vehicles or individuals that vary in a trait x . If one were thinking in molecular terms, one could think of a population of responses that differed in a trait like force, for example. Glenn, Ellis, and Greenspoon (1992) explained that an operant like the lever press, as it actually occurs, constitutes a population of responses. One could also think of a population of interresponse times that differed in duration.

The law of effect. All of Equations 1 – 6 can be applied to behavioral evolution. Depending on one’s view of Thorndike’s (1911/2012) writings, Equation 1 either precisely restates or replaces Thorndike’s original law of effect. Thorndike studied cats’ behavior in escaping from an enclosure with a mechanism inside, such as a loop of string or a lever, attached to a catch on the door to the box. A bowl of food usually sat outside. A trial began with placement of the cat in the box, and when the cat’s actions caused the string to be pulled or the lever to be pushed, the door to the box opened, ending the trial. Thorndike noted that the context of confinement and food induced a variety of escape- and food-related activities. Instead of recording the activities, however, Thorndike measured the time until the cat escaped and considered its decrease with trials to indicate formation of an association between the situation of confinement and the action that resulted in escape. In doing so, he focused on successful actions and largely ignored the other, unsuccessful actions that occurred before escape. Thorndike thought of the law of effect as a law of association, like the laws of frequency and contiguity. It stated that satisfiers that followed a stimulus–response sequence strengthened the association between the stimulus and the response. He acknowledged that, when successful actions occurred sooner, fewer alternative actions occurred and wrote that successful actions were “stamped in” while unsuccessful actions were “stamped out.” Thorndike included comments in his writings that indicate he considered the change due to selection (Donahoe, 1999), but he never explained the analogy between behavioral evolution and genetic evolution at length.

Instead of focusing on an isolated response that gets more strongly connected to the stimulus situation, Equation 1 assumes a

population of various responses or actions, in keeping with evolutionary population thinking (Mayr, 1970). If we take Thorndike's trials to constitute individuals in Price's reasoning (Fig. 1) and take a block of trials as the population at Time 1, then the different levels of x_i represent variation in responses or actions. Some of these variants operate the mechanism in Thorndike's box (e.g., clawing at the string) and have higher "fitness" or recurrence w_i than others that fail to do so, because these successful actions correlate with escaping from the box. For successful actions, let x_i equal 1.0, analogously to allele B. For unsuccessful actions, let x_i equal zero, analogously to allele A. From one block of trials to another (i.e., a population at Time 2), the successful actions (e.g., "string-pulling") increase in the population, and the unsuccessful actions decrease, and covariance between x_i and w_i is positive (i.e., $\Delta\bar{x}$ in Eq. 1 is positive).

Equation 1 similarly restates or replaces Skinner's (1938, 1953) version of the law of effect, which Skinner (1981) explicitly compared to genetic and cultural evolution. Like Thorndike, he focused only on successful actions—say, those that succeeded in depressing a lever—which he defined as a class. In a free-operant situation, the operant responses that actually occur constitute a population of responses, as explained by Glenn et al. (1992). The population consists of responses both successful and unsuccessful at pressing the lever, as suggested by Catania's (1973) phrase "functional operant," which includes all responses induced by the contingency with food. If we take Time 1 as a time interval of observation—say, an hour-long session—then the various responses or actions may be equated with levels of x_i , and we can call any variants that get the lever depressed "successful" with x_i equal to 1.0, whereas those that fail to depress the lever we can call "unsuccessful," with x_i equal to zero, analogously to alleles B and A in Figure 1. If the mean fitness or recurrence of the successful actions ($\bar{w}_1$) is greater than the mean fitness or recurrence of the unsuccessful actions ($\bar{w}_0$), comparing Time 1 to another session at Time 2, then the covariance between x_i and w_i is positive and the operant population evolved. Catania's "functional operant" is better described by Equation 2, because the expected value term allows the possibility of new variants—new levels of x_i —

arising at Time 2. These new variants, which allow shaping and give rise to novelty, would be like mutations.

Response rate. Elsewhere I have argued that a multiscale molar view of behavior based on time allocation among temporally extended activities is more powerful and more general than a molecular view (Baum, 2012, 2013; Baum & Rachlin, 1969). How can the multiscale view square with "selection by consequences" and Price's equation?

The answer requires thinking in temporal terms. The varying entities or vehicles would be, for example, time samples within a larger time interval, and these time samples would constitute the population—seconds within a minute, minutes within an hour, hours within a day, weeks within a year, and so on (Baum, 1973, 2012). The variable x_i that differs across time samples may be occurrence or not of an activity—if the time samples are short—or rate or time taken up by the activity—if the time samples are longer—or any other quantitative aspect of an activity, such as force or specifics of topography.

The same considerations apply whether one adopts the multiscale view that behavior consists of extended activities (Baum, 2013) or the molecular view that behavior consists of discrete responses (Skinner, 1938, 1969). Not only activities, but discrete responses also take up time; that is why response rate cannot increase indefinitely, but always must reach an upper limit (Baum, 1981, 2015; Baum & Davison, 2014b; Herrnstein, 1970). Whether discrete responses or extended activities, their total takes up the observation time interval and thus limits the size of the population to the observation time interval. The time samples constitute the varying individuals or vehicles, and the observation time interval acts analogously to carrying capacity to force competition among substitutable variants (i.e., the time samples). They compete because they take up time—whether discrete responses or different activities and parts within activities. If one variant of x increases in frequency, others must decrease. Thus, whichever view one takes, the varying individuals or vehicles would be the same: smaller time samples within the larger observation interval. Figure 3 illustrates this approach graphically.

Figure 3 applies the multiscale approach to hypothetical data. The top graph shows an




Fig. 3. Hypothetical plot of an activity at two different times, T1 and T2. Top: Each horizontal bar represents an episode of the activity with pauses in between. Time moves from left to right. The activity occurs more frequently at Time 2 than at Time 1. Bottom: The two observation time intervals are divided into equal time samples (vertical lines), which constitute the individuals in this example of behavioral evolution. The two circles on the left show time samples empty of the activity. The two circles in the middle show time samples that were completely filled with the activity. The two circles on the right show time samples in between the two extremes, partially taken up by the activity.

activity occurring at two different times, Time 1 and Time 2. The horizontal axis represents time going from left to right. Each horizontal bar represents an episode or bout of the activity. The activity is occurring at a higher rate at Time 2 than at Time 1; the bouts are longer and the pauses between bouts are shorter (e.g., Shull, Grimes, & Bennett, 2004). The bottom graph in Figure 3 shows the two time intervals divided into small, equal time samples. The time samples vary in the proportion of the sample taken up by the activity. The two circled time samples on the left are empty, containing none of the activity—the proportion of time taken is zero. The two circled time samples in the middle are completely filled by the activity—the proportion of time taken is 1.0. The two circled time samples on the right contain proportions in between the two extremes of 0.0 and 1.0.

When the analysis illustrated in Figure 3 is applied to actual data, the result is a frequency distribution like that shown in Figure 4 (Baum, 2012). The data were gathered by John Van Syckel when he was a graduate student at University of New Hampshire. Several sessions of a rat's stable performance on variable-interval (VI) 30 s of food delivery, in which lever presses produced foods after

varying amounts of time, were broken into successive 2-min time samples, as in the lower panel of Figure 3, and the number of presses in each time sample counted. The 2-min samples were sorted according to the number of lever presses in the sample, each sample representing a certain press rate (presses per min). The solid line shows the frequency distribution of press rates. Most were in the range from 3 to 12 per min, but occasionally rates above 15 per min occurred. The modal response rate was about 5 per min, but the mean (dotted line) is higher, because the distribution is skewed to the right.

If the proportions of activity illustrated in Figure 3 and shown in Figure 4 recur exactly from one time to another, mutation might be irrelevant when considering response rate. Then Equation 1 would suffice. If new time-sample response rates occurred at Time 2, however, then the analogy to mutation might apply, and one would describe the change with Equation 3, to include the expected value term expressing change in x (Δx_i) due to distortion in recurrence. Equations 1 and 3 may be applied in at least two ways.

Directional selection and stabilizing selection. If one changed the situation in Figure 4,




Fig. 4. Variation in response rate within a rat's stable performance on a VI 30 schedule. Stable sessions were divided into 2-min intervals as in Figure 3 and sorted according to number of presses in the interval. The jagged line shows the frequency distribution of press rates. The dotted vertical line indicates the mean.

say, by changing the VI schedule average duration, one might expect the whole distribution to shift. The top panel of Figure 5 illustrates such a shift, from Time 1 with a leaner schedule to Time 2 with a richer schedule. The change in schedule causes increased covariance between higher rates and recurrence. Higher rates recur until the distribution shifts to a new, higher mean at Time 2. A whole-distribution shift like this is termed "directional selection" (Mayr, 1970).

Equation 1 or 3 applied to directional selection means, for example, that if time samples (the vehicles or individuals indexed by i) including more of an activity in Time Interval 1 correspond to relatively more such time samples in Time Interval 2 (i.e., positive covariance with w_i in Eq. 1), then the activity ($\bar{x}$) increases from Time Interval 1 (Population 1) to Time Interval 2 (Population 2)— $\Delta\bar{x}$ is positive. In the laboratory example, Figure 4, if the rat's rate of lever pressing (x_i) varies, and the covariance between x_i and w_i (recurrence) is positive, the mean rate across an hour ($\bar{x}$) increases from one hour to the next (Fig. 5, top). If a person's life changes by having children, spending more time with family per day (x_i) in one month may vary positively with relative recurrence per day (w_i) in the next month, and the monthly amount of time spent with family ($\bar{x}$) increases from month to month (i.e., $\Delta\bar{x}$ is positive).




Fig. 5. Directional selection and stabilizing selection. Top: directional selection results in a shift of the population's frequency distribution of x from Time 1 to Time 2. Bottom: stabilizing selection keeps a population stable by selecting against deviations from the mean in both directions. The black curve represents the population's frequency distribution of x . The vertical dotted line represents the mean of x in the population. The gray curve represents variation in fitness, decreasing in both directions away from the mean. The arrows indicate the direction of selection. The plus sign indicates positive covariance between fitness and x . The minus sign indicates negative covariance between fitness and x .

Selection does not cease, however, when a population stabilizes around a particular mean. Selection then acts to keep the population stable. The lower panel of Figure 5 illustrates this second application of Equation 1, known as "stabilizing selection" (Mayr, 1970). The black curve represents the frequency distribution of x around the mean (dotted line). The gray curve shows how fitness varies with x . Fitness decreases with deviations from the mean in both directions. The gray arrows indicate how selection drives the population back toward the mean, and the circled plus and minus show the sign of the covariance between x and fitness above the mean.

(negative) and below the mean (positive). These positive and negative covariances stabilize the population.

When a behavioral performance like that in Figure 4 is stable, selection within the stable performance acts as in the lower panel of Figure 5 to keep the population of response rates stable. Selection must favor higher rate for rates below the modal rate and must favor lower rate for rates exceeding the modal rate. Covariance between rate and fitness (recurrence) must be positive below the modal rate and negative above the modal rate, as in Figure 5.

The multiscale view applies to behavior at any time scale, whether seconds, days, months, or years. The Price equation allows one to think about selection acting at various time scales. As with cultural group selection, Equations 4 and 5 apply to operant activities at a level or time scale of selection that may include selection at a lower level or shorter time scale. The analog to group selection is selection of whole activities extended in time, like playing tennis well or poorly and relating to one's spouse well or poorly. The analog to within-group selection is the selection of parts referred to above, because every activity is composed of parts that are themselves activities on a smaller time scale (Baum, 2002, 2012, 2013; Baum & Davison, 2004). Thus, the second, expected value, term on the right in Equation 5 for behavioral evolution represents selection among parts of an activity. A tennis player's serve may improve when the ball is thrown up in a new way, and that variant may be selected. Active listening may improve one's relationship with one's spouse and be selected to replace a passive stance.

Theory of response rate. Let us now see how the Price equation may mesh with some concepts about response rate—feedback functions (i.e., contingency) and induction (Baum, 2012). Even after an organism has been exposed to a single VI schedule for enough days to stabilize its performance—that is, response rate varies only unsystematically from day to day—we may assume that within the stable performance, dynamics of selection still occur. Otherwise, were the schedule to change, performance might not change. Looked at in this light, the feedback function for a VI schedule itself suggests dynamics within stable performance (Baum, 1973, 2012).

The top panel in Figure 6 shows the same data as in Figure 4 with the VI feedback function added. The jagged solid line shows the frequency distribution of press rates. The filled square indicates the average press rate. Each press rate is associated with a certain food rate, calculated by adding up the number of food pellets delivered for the press rate and dividing by the total time at that press rate (number of intervals times 2 min). The diamonds show these food rates plotted as a function of the press rates. The broken curve shows a feedback function fitted to the diamonds:

$$r = \frac{60}{t + \frac{a}{B}} \quad (7)$$

where r is food rate in pellets/minute (FPM), t is the VI in minutes (equal to 0.52), B is the press rate in presses/minute, and the parameter a equals 0.95 (Baum, 1992). If we assume that differential reinforcer rate (food pellet rate here) selects among response rates (press rates here), then the slope of the feedback function around any particular press rate suggests the degree of selection in that region of the curve. The slope is positive throughout, but becomes progressively flatter as press rate increases. Selection for higher response rate is highest for low response rates and becomes progressively less for higher and higher response rates. Intuitively, the curve must be showing that covariance between response rate and fitness is greater below the mean than above the mean, but how to make the connection to fitness?

A feedback function by itself cannot predict stable performance, fitness, or change. It may be part of a theory, but something more is needed, because the fitness of a response rate is not measured by the reinforcers it garners, as comparison of the frequency distribution in Figure 5 with the feedback function in Figure 6 shows. For VI performance, some factor must deter high response rates, and that factor must be either an effect of response rate itself or an effect of reinforcers on the response rate that wins them. In an earlier paper, I suggested that the limiting factor might be the growing cost of responding as response rate increases (Baum, 1981). In the present, another possibility presents itself: the inducing effect of reinforcers on response rate (Baum, 2015).




Fig. 6. Selection dynamics within a rat's stable performance on a VI 30 schedule. Stable sessions were divided into 2-min intervals and sorted according to number of presses in the interval. Top: the jagged line shows the frequency distribution of press rates (right vertical axis), the diamonds show corresponding food rates, the broken curve shows the feedback function, and the filled square shows the mean across sessions. Bottom: squares show food pellets per min (FPM) transformed by a power function modeling induction to presses per min (PPM). The filled circle shows the preserved mean performance (square in top graph; assuming overall press rate is stable). The solid curve shows fitness (right vertical axis) calculated by treating the induced pressing as a second population. The broken line indicates fitness equal to 1.0—selection favors increasing rate only for the rates with relative fitness greater than 1.0.

To apply Price's equation to selection within this stable VI performance, we need to make two assumptions. The first is that the variable x_i in Equation 1 is press rate indexed by the 2-min intervals (i). The second is that recurrence of any press rate is directly related to its associated food rate. The relation between recurrence and food rate is induction (Baum, 2012, 2015; Baum & Davison, 2014b) and may be given by a power function:

$$B(I) = cr^s \quad (8)$$

where $B(I)$ is induction, which might be thought of as induced response rate, and r is pellet rate. $B(I)$ corresponds to Time 2 in Price's equation. The concave-downward shape of the feedback function tells us that covariance between relative recurrence (fitness) and response rate is high for lower response rates and decreases for higher response rates. Selection for increased response rate is high for lower response rates and low for higher response rates.

The lower graph in Figure 6 shows the results of applying Price's equation to the performance in the upper graph. The squares show the feedback function (Eq. 7; the dashed curve in the upper graph) transformed according to Equation 8. This transformed food rate, $B(I)$, is shown by the squares. The parameter c (equal to 92.4) and the exponent s (equal to 3.56) were derived by requiring the average induced response rate (filled circle) to equal the average response rate in the upper graph (filled square), on the grounds that average response rate is stable. The large exponent causes induction to be low for low food rates but to increase rapidly as food rate increases (Baum, 2015). Fitness (w) of each press rate was calculated as the ratio of $B(I)$ to press rate and dividing by the mean; if the ratio $\left[\frac{B(I)}{B}\right]$ were called v , then fitness was calculated as $[v/\bar{v}]$. The solid curve shows fitness across press rates. It is less than 1.0 (dashed line) for low response rates, increases above 1.0 in the region from about 2.5 to 10 presses/min, and then drops below 1.0 again for rates above 10 presses/min. Selection favors response rates in the midrange and goes against both lower and higher response rates.

The analysis in Figure 6 suggests an explanation of performance on VI schedules—why response rate is never extremely high as in ratio schedules, in which a reinforcer is produced by a number of responses, and never extremely low. If we compare the fitness curve in Figure 6 with the fitness curve in the lower panel of Figure 5, we see that, as in the illustration in Figure 5, covariance between response rate and fitness is positive for response rates below the mean and negative for response rates above the mean. The dynamics are such as to select moderate rates and select against both extremely high and extremely low rates, but the overall average

rate remains the same from time to time. Thus, VI schedules will always maintain moderate response rates.

In contrast, ratio schedules, for which the feedback function is a straight line, always select higher response rate for all response rates (Baum, 1993, 2015). On ratio schedules, response rate will always tend to be the maximum possible, which is correct to an approximation (Baum, 1981; Fig. 10). Why rate is ever less than the maximum—as with ratio “strain”—requires more assumptions (e.g., Baum, 1981, 2015).

This explanation of VI and variable ratio (VR) performance may be compared to ones based on the molecular view that behavior consists of discrete responses. These molecular-based explanations rely on differential reinforcement of interresponse times (IRTs; e.g., Morse, 1966). Since the probability of a reinforcer goes up with time in a VI schedule, the longer between responses, the more likely the next response will produce a reinforcer—that is, long IRTs are differentially reinforced. Such selection would work to keep response rate from becoming extremely high, but it fails to explain why VI response rate is not extremely low—low enough that every response would produce a reinforcer. The loss in reinforcer rate at low response rates might seem like the reason that response rate stays in the moderate range, but such an appeal to molar relations illustrates the inadequacy of the molecular account. The present explanation based on selection of response rates accounts for the moderate VI rates by showing selection against both extremely high and extremely low response rates.

The treatment of response rate offered here is at least incomplete, because it fails to incorporate competition between the operant activity (pressing) and other activities induced by the reinforcer (PIE; e.g., food). As a result, it fails to make the correct prediction of stable response rates across VI schedules. These stable rates are often fitted with a hyperbola that incorporates a form of competition by viewing response rate as dependent on relative reinforcer rates (Baum & Davison, 2014b; Herrnstein, 1970). An earlier paper showed, however, that response rates across VI schedules are well described by competition between two power functions (Baum, 2015). A more correct model will recognize that the

response rates don't just compete with one another, but also with the other activities induced.

Choice in concurrent schedules. Equation 2 may be applied to performance on concurrent VI schedules. Much research has focused on concurrent VI VI schedules in procedures that entail maintaining a pair of schedules for many sessions, until performance appears to be “stable”—that is, varying from session to session but with no apparent trend. In reporting such an experiment, the whole session is usually treated as a unit, and the measure of choice or preference between the two VI alternatives is computed as either $\frac{B_1}{B_1 + B_2}$ or $\log \left[\frac{B_1}{B_2} \right]$, where B_1 and B_2 are time or behavior spent in responding at the two VI schedules. For Equation 3, such data analysis would indicate that the session constitutes the individual and that the trait x is choice, measured one way or the other. A block of earlier sessions would constitute a population at Time 1, and a later block of sessions would constitute Time 2. Selection favoring some levels of x over others would result in choice changing across sessions, until choice becomes stable and stabilizing selection takes over (Fig. 5; see Baum, 2010, for a tutorial in dynamics of choice). If analysis remains at the level or time scale of whole sessions, Equation 3 suffices, but if analysis proceeds beyond whole sessions to consider smaller time scales within sessions, then Equation 5, the multiscale equation, applies.

Some experiments on choice have taken a multiscale approach (e.g., Aparicio & Baum, 2006, 2009; Baum & Davison, 2004, 2014a; Davison & Baum, 2000). Taking log-ratio of behavior as x , these experiments suggest interpretation in the light of Equation 5. Typically, each experimental session is composed of seven components, each consisting of a different pair of VI schedules, one for responding at the left key or lever, and one for responding at the right key or lever. Blackouts separate the components, no signals indicate which pair of schedules is in effect, and each component lasts for a certain number of food deliveries (e.g., 10 or 12).

For example, Davison and Baum exposed pigeons to such conditions for 50 sessions per condition (Baum & Davison, 2004; Davison & Baum, 2000). By the end of 15 sessions, performance appeared stable at the time scale of

whole sessions, as indicated, for example, by calculating log ratio of key pecks at the left and right keys across whole sessions. The remaining 35 sessions were taken as stable at that time scale and pooled to measure performance at smaller time scales. To do this, all approximately 245 components from the 35 sessions could be taken to define populations of log behavior ratios within components. If the data from six pigeons are pooled, the population is about 1470. Although performance appeared stable at the time scale of whole sessions, the procedure assured that evolution would occur within components, because no signals indicated the different food ratios and only the obtained food deliveries themselves could guide responding.

Analysis begins at the time scale of the component, because choice, measured from the beginning of a component up to the first food delivery, started out near indifference and changed from delivery to delivery as the component continued. The top panel in Figure 7 shows a result from an experiment with pigeons and these frequently changing components (Baum & Davison, 2004). It shows the effect on interfood peck ratio (from the beginning of the component to the first food, from the first to the second food, etc.) of repeated food deliveries from the same alternative, called "continuations." Peck ratios were calculated as pecks on the just-productive key (P) divided by pecks at the not-just-productive key (N) during the interfood interval. Choice increased up to about the sixth continuation and then leveled off at a preference close to 4:1. Although the graph only shows averages and not distributions of individuals, the increase indicates directional selection of choice at this time scale. To think of the increase in relation to Equation 5, we think of x_g as log peck ratio in one component, with g indexing components. For six pigeons, 35 sessions, and seven components, the population size n starts at 1470 for the interval from component beginning to first food and necessarily decreases across interfood intervals, because sometimes alternative N produces food (a discontinuation). Though not shown, for each sequential interfood interval, a frequency distribution of x_g exists, with g ranging from 1 to n (decreasing from 1470 to roughly 1200 after one reinforcer and ultimately to about




Fig. 7. Behavioral evolution at different time scales within pigeons' key pecking on concurrent VI VI schedules. Sessions consisted of seven unsignaled components, each with a different reinforcer ratio. Both graphs show evolution within components with continuing food deliveries from the same alternative. Top: evolution of interfood peck ratio at the just-productive key (P) to the not-just-productive key (N). Each diamond shows the ratio of all pecks at P to all pecks at N from the beginning of the component or the previous food delivery to the next food delivery. Bottom: evolution of peck ratio in the first changeover cycle within interfood intervals. The diamonds and triangles represent the length of the first visit (#1), which was nearly always to the just-productive key (P1), and the length of the second visit (#2), which was to the not-just-productive key (N2). The lengths in pecks/visit are indicated on the secondary (right) vertical axis. The filled squares represent the ratio of the two visits (P1/N2). Selection at the smaller time scale (bottom) acted in concert with selection at the larger time scale (top).

500 after six reinforcers). Each point in the graph is the average of a population of log peck ratios following a certain number of continuing food deliveries, and the sequential interfood intervals correspond to Time 1, Time 2, Time 3, and so on. The first term on the right of Equation 5, $cov(w_g, x_g)$, represents the covariance between log peck ratios (x_g) across components (the individuals, indexed by g) with fitness (w_g) for one interfood interval (Time 1) to the next interfood interval

(Time 2). As long as peck ratio increases from food delivery to food delivery, this covariance is positive.

These interfood peck ratios, however, tell only what happens at their time scale. No allocation of behavior is specified completely by the ratio of summed behavior. At least two more parameters must be specified: the length of visits to or bouts of the alternative activities.

All activities are episodic. Working may occur for some hours each day, and time with family for a certain time too, varying from day to day. Similarly, in an experiment with pigeons' choice, pecks accumulate in episodes of pecking at the alternatives. Since interfood log peck ratio was calculated by summing all pecks at each alternative, it says nothing about any regularity at the smaller time scale of within-interfood-interval activities. The second term on the right in Equation 5 represents this smaller time scale.

One way to think about log peck ratio at this smaller time scale is to compare pairs of sequential visits. The first visit was nearly always to the just-productive alternative (P) and the second to the not-just-productive alternative (N). We may call this pair of visits one changeover cycle, and because the visits must strictly alternate between alternatives, we may call the changeover cycles the individuals within an interfood interval, in analogy to the members within a group. If we have measured 10 sequential visits in each interfood interval, we have five changeover cycles. For each changeover cycle, we may calculate the log ratio of pecks per visit in the first visit to pecks per visit in the second visit. This log ratio would be x_{ig} in Equation 5, with g indexing the roughly 1200-500 interfood intervals at each ordinal position in the component and i indexing the (five) changeover cycles within the interfood interval. The analysis is analogous to a population of 1200-500 groups with five members each. For illustration, we focus on the log peck-per-visit ratio between the first visit and the second visit—the first changeover cycle. The bottom graph in Figure 7 shows the first postfood visit (P1; broken line and diamonds), the second visit (N2; broken line and triangles), and their ratio (P1/N2; solid line and squares). Each diamond or triangle represents an average of visits. The first visit increased from about 5-6 pecks to about 10-11 pecks (vertical axis to the right). Even the first visit in a component, before any

food was delivered, was longer than the second visit in the component (zero on the x-axis), either because of a carry-over from the previous component or simply because it was induced by the light coming on at the beginning of the component. The second visit remained about the same throughout the component, at about 1-2 pecks. The brevity and invariance of the second visits indicates that responding during the interfood interval took on a fix-and-sample pattern (Aparicio & Baum, 2006; Baum, 2002; Baum, Schwendiman, & Bell, 1999). Since only the first visits increased, we may conclude that only those visits evolved. The line for the visit ratio confirms this reasoning by following the increase in first visits closely; the ratio increases from about 2.3 to about 6.0 (left vertical axis). The squares represent the averages of x_{ig} across components (for i just equal to 1).

A similar curve for log ratio of Visits 3 and 4 (the second changeover cycle) could be calculated. It too increased, but not as much as the log ratio of Visits 1 and 2, because Visit 3 was always shorter than Visit 1. In the light of Equation 5, we may conclude that the covariance $\text{cov}(w_g, x_g)$ at the time scale of interfood intervals (group level) and the average covariance $E_g[\text{cov}(w_{ig}, x_{ig})]$ at the time scale of changeover cycles were both positive. Selection on changeover cycles (members) acted in concert with selection on interfood log ratios (groups).

Given this clear qualitative result, we might be encouraged to calculate the actual fitnesses and covariances as choice changed. To do this, one may calculate from frequency distributions of the populations of choice measures or visits. Appendix C shows a worked example, using the visit lengths (pecks/visit) of the first visit following food—the diamonds in the lower graph of Figure 7. The negative curvature seen there indicates that the changes in visit length were initially positive and decreased across food deliveries. The changes in visit length (Δx , corresponding to $\Delta \bar{x}$ in Eq. 1) and the covariances between visit length and fitness across food deliveries appear in Figure A2. The changes and covariances are equal, as they must be in accord with Equation 1.

The conclusion that selection acted in concert at the different time scales shown in Figure 7 raises a question: In what situation

would selection at different time scales oppose one another?

Self-control and impulsivity. Equation 6 for behavioral evolution illuminates problems in self-control, including altruism and cooperation. The conflict between impulsivity and self-control translates into a positive first term on the right (self-control) and a negative second term (impulsivity). Avoiding a bad habit or cultivating a good habit (e.g., refusing a drink or a piece of cake or visiting the dentist) on any specific occasion entails the cost (punisher) of forgoing immediate enjoyment or of enduring immediate discomfort, and these are the parts of a whole extended pattern (x_g ; sobriety, dieting, or good health). The β of the second term is negative, because the individual parts or occasions x_{ig} of activity x_g incur this immediate punishment. Thus, the extended pattern (self-control, cooperation, or altruism) has positive β , but the positive β must suffice to offset the negative β of the second term (for x_{ig}) if the good habit is to survive. This offsetting may occur if $\text{var}(x_{ig})$ is reduced to zero by following a rule that enforces good behavior (refusing the drink or cake; doing the right thing) on particular occasions. Seen this way, altruism and cooperation are examples of good habits that incur a cost in the short term (helping out) but are reinforced in the long term (Baum, 2016; Rachlin, 1995, 2002).

Negative contingencies. In a negative contingency, one activity prevents a reinforcer or punisher (i.e., a good or bad phylogenetically important event) while all other activities allow the event to occur. In negative reinforcement, an activity prevents injury or some other fitness-reducing event, and that activity is selected. For example, Herrnstein and Hine-line (1966) exposed rats to a procedure in which pressing a lever canceled upcoming electric shocks, thereby reducing the shock rate to a lower level. Across sessions, lever pressing increased, as illustrated from T1 to T2 in Figure 3. The Price equation tells us that recurrence of lever pressing increased across sessions because the negative contingency between lever pressing and electric shock selected pressing in competition with other activities. The covariance between rate of lever pressing and recurrence of lever pressing was positive because of the selection imposed by the negative reinforcement. In problems of

self-control, a bad habit like smoking incorporates both a short-term positive contingency with small reinforcers (nicotine) and a long-term negative contingency with large reinforcers, whereas the competing good habit incorporates both a short-term negative contingency with small reinforcers and a long-term positive contingency with large reinforcers. Drunkenness provides short-term enjoyment and long-term catastrophic losses, whereas sobriety entails short-term loss of enjoyment and long-term large reinforcers. Which contingency wins determines the alcoholic's quality of life.

In negative punishment, an activity prevents a fitness-enhancing event (a good phylogenetically important event). In a laboratory example, Sanabria, Sitomer, and Killeen (2006) exposed pigeons to a procedure in which pecking at a lit key was induced by following a brief lighting of the key with food ("autoshaping"). The researchers arranged that pecking the key would cancel the food delivery, thereby negatively punishing key pecking. The negative contingency between pecking and food resulted in large reductions in pecking across sessions, indicating that pecking was selected against. The Price equation tells us that the negative punishment resulted in lower recurrence of key pecking in competition with other activities across sessions; the covariance between rate of pecking and pecking recurrence was negative—the opposite of Figure 3.

Conclusion

Although Price's equation is not a theory to be proven or disproven, it is a useful analytical tool and may be a useful basis for theory. It can elucidate selection at different levels in genetic evolution and at different time scales in behavioral evolution, including cultural evolution and shaping of behavior in individual organisms within a lifetime. The examples in the present paper aim to shed light on dynamics within concurrent VI VI performance and allocation on single schedules. The analysis of concurrent performance shows evolution of the fix-and-sample pattern that often emerges (Aparicio & Baum, 2006; Baum, 2002; Baum & Davison, 2004). The analysis of response rate suggests a path to a more complete understanding of changing response rates than the

older molecular view afforded. Other applications of Price's equation may elucidate dynamics of behavioral change in other contexts, particularly if experiments are designed in advance to allow the equation to be applied easily.

References

- Aparicio, C. F., & Baum, W. M. (2006). Fix and sample with rats in the dynamics of choice. *Journal of the Experimental Analysis of Behavior*, 86, 43-63.
- Aparicio, C. F., & Baum, W. M. (2009). Dynamics of choice: Relative rate and amount affect local preference at three different time scales. *Journal of the Experimental Analysis of Behavior*, 91, 293-317.
- Baum, W. M. (1973). The correlation-based law of effect. *Journal of the Experimental Analysis of Behavior*, 20, 137-153.
- Baum, W. M. (1981). Optimization and the matching law as accounts of instrumental behavior. *Journal of the Experimental Analysis of Behavior*, 36, 387-403.
- Baum, W. M. (1992). In search of the feedback function for variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, 57, 365-375.
- Baum, W. M. (1993). Performances on ratio and interval schedules of reinforcement: Data and theory. *Journal of the Experimental Analysis of Behavior*, 59, 245-264.
- Baum, W. M. (2000). Being concrete about culture and cultural evolution. In N. Thompson & F. Tonneau (Eds.), *Perspectives in ethology* (Vol. 13, pp. 181-212). New York: Kluwer Academic/Plenum.
- Baum, W. M. (2001). Two stumbling blocks to a general account of selection: Replication and information. *Behavioral and Brain Sciences*, 24, 528.
- Baum, W. M. (2002). From molecular to molar: A paradigm shift in behavior analysis. *Journal of the Experimental Analysis of Behavior*, 78, 95-116.
- Baum, W. M. (2010). Dynamics of choice: A tutorial. *Journal of the Experimental Analysis of Behavior*, 94, 161-174.
- Baum, W. M. (2012). Rethinking reinforcement: Allocation, induction, and contingency. *Journal of the Experimental Analysis of Behavior*, 97, 101-124.
- Baum, W. M. (2013). What counts as behavior: The molar multiscale view. *The Behavior Analyst*, 36, 283-293.
- Baum, W. M. (2015). The role of induction in operant schedule performance. *Behavioural Processes*, 114, 26-33.
- Baum, W. M. (2016). Driven by consequences: The multiscale molar view of choice. *Managerial and Decision Economics*, 37, 239-248.
- Baum, W. M., & Davison, M. (2004). Choice in a variable environment: Visit patterns in the dynamics of choice. *Journal of the Experimental Analysis of Behavior*, 81, 85-127.
- Baum, W. M., & Davison, M. (2014a). Choice with frequently-changing food rates and food ratios. *Journal of the Experimental Analysis of Behavior*, 101, 246-274.
- Baum, W. M., & Davison, M. (2014b). Background activities, induction, and behavioral allocation in operant performance. *Journal of the Experimental Analysis of Behavior*, 102, 213-230.
- Baum, W. M., & Rachlin, H. C. (1969). Choice as time allocation. *Journal of the Experimental Analysis of Behavior*, 12, 861-874.
- Baum, W. M., Schwendiman, J. W., & Bell, K. E. (1999). Choice, contingency discrimination, and foraging theory. *Journal of the Experimental Analysis of Behavior*, 71, 355-373.
- Boyd, R., & Richerson, P. J. (1985). *Culture and the Evolutionary Process*. Chicago: University of Chicago Press.
- Catania, A. C. (1973). The concept of the operant in the analysis of behavior. *Behaviorism*, 1, 103-115.
- Davison, M., & Baum, W. M. (2000). Choice in a variable environment: Every reinforcer counts. *Journal of the Experimental Analysis of Behavior*, 74, 1-24.
- Donahoe, J. W. (1999). Edward L. Thorndike: The selectionist connectionist. *Journal of the Experimental Analysis of Behavior*, 72, 451-454.
- El Mouden, C., André, J. -B., Morin, O., & Nettle, D. (2014). Cultural transmission and the evolution of human behaviour: a general approach based on the Price equation. *Journal of Evolutionary Biology*, 27, 231-241.
- Frank, S. A. (1995). George Price's contributions to evolutionary genetics. *Journal of Theoretical Biology*, 175, 373-388.
- Glenn, S. S., Ellis, J., & Greenspoon, J. (1992). On the revolutionary nature of the operant as a unit of behavioral selection. *American Psychologist*, 47, 1329-1336.
- Goodnight, C.J., Schwartz, J. M., & Stevens, L. (1992). Contextual analysis of models of group selection, soft selection, hard selection, and the evolution of altruism. *The American Naturalist*, 140, 743-761.
- Guerin, B. (1997). How things get done: Socially, non-socially; with words, without words. In L. J. Hayes & P. M. Ghezzi (Eds.), *Investigations in behavioral epistemology* (pp. 219-235). Reno, NV: Context Press.
- Henrich, J. (2004). Cultural group selection, coevolutionary processes and large-scale cooperation. *Journal of Economic Behavior and Organization*, 53, 3-35.
- Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior*, 13, 243-266.
- Herrnstein, R. J., & Hiline, P. N. (1966). Negative reinforcement as shock-frequency reduction. *Journal of the Experimental Analysis of Behavior*, 9, 421-430.
- Mayr, E. (1970). *Populations, species, and evolution*. Cambridge, MA: Harvard University Press.
- McElreath, R., & Boyd, R. (2007). *Mathematical models of social evolution: A guide for the perplexed*. Chicago: University of Chicago Press.
- Morse, W. H. (1966). Intermittent reinforcement. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application* (pp. 52-108). New York: Appleton-Century-Crofts.
- Panchanathan, K. (2011). George Price, the Price equation, and cultural group selection. *Evolution and Human Behavior*, 32, 368-371.
- Price, G. R. (1970). Selection and covariance. *Nature*, 227, 520-521.
- Price, G. R. (1972). Extension of covariance selection mathematics. *Annals of Human Genetics*, 35, 485-490.
- Price, G. R. (1995). The nature of selection. *Journal of Theoretical Biology*, 175, 389-396.
- Rachlin, H. (1971). On the tautology of the matching law. *Journal of the Experimental Analysis of Behavior*, 15, 249-251.

- Rachlin, H. (1995). Self-control: Beyond commitment. *Behavioral and Brain Sciences*, 18, 109-159.
- Rachlin, H. (2002). Altruism and selfishness. *Behavioral and Brain Sciences*, 25, 239-296.
- Reid, R. L. (1958). The role of the reinforcer as a stimulus. *British Journal of Psychology*, 49, 202-209.
- Richerson, P. J., Baldini, R., Bell, A., Demps, K., Frost, K., Hillis, V., ... Zefferman, M. R. (2016). Cultural group selection plays an essential role in explaining human cooperation: A sketch of the evidence. *Behavioral and Brain Sciences*, 39, e30.
- Rogers, E. M. (1995). *Diffusion of innovation*, 4th ed. New York: Simon and Schuster.
- Sanabria, F., Sitomer, M. T., & Killeen, P. R. (2006). Negative automaintenance omission training is effective. *Journal of the Experimental Analysis of Behavior*, 86, 1-10.
- Segal, E. F. (1972). Induction and the provenance of operants. In R. M. Gilbert & J. R. Millenson (Eds.), *Reinforcement: behavioral analyses* (pp. 1-34). New York: Academic.
- Shull, R. L., Grimes, J. A., & Bennett, J. A. (2004). Bouts of responding: The relation between bout rate and the rate of variable-interval reinforcement. *Journal of the Experimental Analysis of Behavior*, 81, 65-83.
- Skinner, B. F. (1938). *Behavior of organisms*. New York: Appleton-Century-Crofts.
- Skinner, B. F. (1953). *Science and human behavior*. New York: Macmillan.
- Skinner, B. F. (1969). Operant behavior. In *Contingencies of reinforcement: A theoretical analysis* (pp. 105-132). New York: Appleton-Century-Crofts.
- Skinner, B. F. (1981). Selection by consequences. *Science*, 213, 501-504.
- Sober, E. (1984). *The nature of selection: Evolutionary theory in philosophical focus*. Cambridge, MA: MIT Press.
- Thorndike, E. L. (1911). *Animal intelligence*. New York: Macmillan. Facsimile reprint by Forgotten Books, 2012.

Received: June 27, 2016

Final Acceptance: April 14, 2017

Appendix A

Derivation of the Basic Price Equation

Assume a population or group of N members at two time periods or “generations,” Time 1 and Time 2. A trait or property of these members x varies among them. The trait could be the presence ($x = 1$) or absence ($x = 0$) of a particular allele ($x = 0, 0.5$, or 1.0 for diploidy), or a quantitative trait like body size or coloration. At Time 1, the mean of x across members is $\frac{1}{N} \sum_{i=1}^N x_i$, denoted $\bar{x}$. Each member i contributes y_i surviving offspring or copies to the population at Time 2— i ’s absolute recurrence. Since the sum of y is not necessarily equal to N , because the population might increase or decrease from Time 1 to Time 2, absolute recurrence is not a good measure of fitness. In Figure 1, we could equate fitness with y_i because the population size ($N = 10$) didn’t change. Fitness or relative recurrence of a member i , w_i , equals $y_i/\bar{y}$. Since $\bar{y}$ equals $\frac{1}{N} \sum_i y_i$, w_i equals $\frac{Ny_i}{\sum_i y_i}$. The sum of w_i equals N , and the mean $\bar{w}$ equals 1.0 . That is, w_i is a relative measure and dimensionless. The mean of x at Time 2 $\bar{x}'$ equals:

$$\bar{x}' = \sum_i w_i x_i / N + \sum_i w_i \Delta x_i / N \quad (\text{A1})$$

where the first term on the right is the portion of the new $\bar{x}$ at Time 2 due to selection, and Δx_i is the change in x_i for member i from Time 1 to Time 2 due to imperfect fidelity of recurrence. Since each $\frac{w_i}{N}$ is a weight (i.e., they sum to 1.0), the second term on the right can be rewritten as the expected value of Δx_i , $E_w(\Delta x_i)$. Using the definition of covariance ($\sum_i w_i x_i / N - \bar{w} \bar{x}$), we may substitute for the first term on the right $\text{cov}(w_i, x_i) + \bar{w} \bar{x}$. Recalling that $\bar{w}$ equals 1.0 , and taking the difference $\Delta \bar{x} = \bar{x}' - \bar{x}$ leads to the simplest form of the Price equation (see Price, 1970, and McElreath & Boyd, 2007, chapter 6, for more detail):

$$\Delta \bar{x} = \text{cov}(w_i, x_i) + E_w(\Delta x_i) \quad (\text{A2})$$

Derivation of the Group Price Equation

We may first just represent group selection in which the groups constitute individuals:

$$\Delta \bar{x} = \text{cov}(w_g, x_g) + E_w(\Delta x_g) \quad (\text{A3})$$

where g indexes the groups in the population, instead of the individual members.

The first term on the right in Equation A3, the covariance, represents group selection alone. We take the second term $E_w(\Delta x_g)$, the mean change in a group apart from group selection, as a mean of change in the individual members of each group. For one group, we treat it the same way as we treated $\Delta \bar{x}$. So, for one group

$$\Delta x_g = \text{cov}(w_{ig}, x_{ig}) + E_w(\Delta x_{ig}) \quad (\text{A4})$$

where i indexes members within the group.

We substitute Equation A4 into Equation A3 and arrive at:

$$\Delta \bar{x} = \text{cov}(w_g, x_g) + E_g[\text{cov}(w_{ig}, x_{ig}) + E_w(\Delta x_{ig})] \quad (\text{A5})$$

which includes the change in notation E_g to indicate that the second term on the right is an expected value across groups. It is Equation 4 in the paper.

Appendix B

An alternative approach to individual and group selection, called “contextual analysis,” de-emphasizes covariance and instead partitions variance into within-group and between-group components, not entirely unlike Equation 6 (Goodnight, Schwartz, & Stevens, 1992). The “between-group” component, however, is conceived of as affecting variance among individual members as a result of their membership in the group. Contextual analysis would assume that the fitness of a bee colony is just the sum of the fitnesses of the individual bees and would not assign a fitness to the colony in competition with other colonies. In other words, contextual analysis treats groups as if they were collections of individuals and makes no distinction between collections and groups as integral wholes. The key difference between contextual analysis and Price’s approach is that contextual analysis assigns all selection, fitness, and change to individual organisms’ phenotypes—that is, group change is just the sum of the changes in the members as a result of selection acting on them. The




Fig. A1. Computation of fitness of first visits (pecks/visit) going from one interfood interval (Time 1) to the next (Time 2) or first visit in a component (Time 1) to the first interfood interval (Time 2; top left graph). Diamonds show the frequency distribution of pecks/visit for Time 1 (left vertical axis). Squares show the frequency distribution for Time 2. Each pair allows a calculation of fitness in accord with Price's equation, from the beginning of the component to the first food, from the first food to the second, and so on, up to the eighth food delivery. Triangles show mean fitnesses of the different visit lengths (right vertical axis). The broken line in each graph is a regression line fitted to mean fitness (triangles) as a function of pecks/visit. Its equation appears in the graph.

two approaches may be appropriate to different situations. Contextual analysis appears to work well for collections of flour beetles and of leafy plants (Goodnight et al., 1992), but

Price's approach might make more sense when a population is structured into true groups that are integral wholes. Since we are here concerned with behavior, which is

organized into activities that function as integral wholes, Price’s approach seems more appropriate (Baum, 2002).

Appendix C

Figure A1 shows frequency distributions of the first postfood visit (averages in Fig. 7, lower panel). Each graph contains two frequency distributions (Time 1 and Time 2): for a number of same-alternative food deliveries (“continuations”) and the next higher number of food deliveries. The frequency distribution for the smaller number of food deliveries is shown with diamonds, and the next one is shown with squares. For example, the upper left graph shows the frequency distribution of the first visit in the component (0 Food; diamonds) and the frequency distribution of the first visit following the first food (1 Food; squares), the second graph on the left shows the frequency distribution of the first visit following first food (1 Food; diamonds) and the first visit following the second food (2 Food; squares), and so on. Only visits in the range from 1 peck to 24 pecks were included; these constituted the vast majority. The fitness of a visit length going from one interfood interval to the next (from Time 1 to Time 2) equals the ratio of the relative frequencies in the two frequency distributions (Appendix A). The fitnesses of the different visit lengths are represented on the right vertical axes and appear in each graph as triangles.

The broken line in each graph is a regression line fitted to the fitnesses (triangles). Its equation is given in the graph. The line gives only a rough indication of the relation between fitness and visit length, because it takes no account of the differences in frequency, but only shows how average fitness of the different visit lengths varied with visit




Fig. A2. Change in the mean length of the first visit following food (Δx) and covariance between visit length and fitness plotted against the second ordinal food delivery of the graphs in Figure A1. The first point shows the change and covariance from no food to the first food delivery, the second point shows the change and covariance from the first food to the second food delivery, and so on. Change in visit length (Δx representing $\Delta \bar{x}$ of Eq. 1; square) and covariance (X symbol) are equal, in accord with Price’s equation.

length. The slopes roughly accord, however, with the trend one would expect from the negative curvature of the visit lengths shown in the bottom graph of Figure 7.

Figure A2 shows the changes in mean visit length (Δx) as squares. They decrease as expected from Figure 7. The abscissae indicate the second of the pair of interfood intervals paired in Figure A1. For example, the 1 on the x-axis shows the mean change from 0-1 Food in the upper left graph of Figure A1, the 2 on the x-axis shows the mean change from 1-2 Food in the second graph on the left, and so on. The X symbols show the computed covariance between fitness and visit length for each pair of interfood intervals (Time 1 and Time 2). As expected from Equation 1, the mean change and the covariance for each pair are equal.