

Frequency, Contingency and the Information Processing Theory of Conditioning

C. R. Gallistel

Abstract

The framework provided by Claude Shannon's (1948) theory of information leads to a far-reaching, more quantitatively oriented reconceptualization of the processes that mediate what is commonly called associative learning. The focus shifts from processes set in motion by individual events to processes sensitive to the information carried by the flow of events. The conception of what properties of the conditioned and unconditioned stimuli are important shifts from the tangible properties that excite sensory receptors to the abstract and intangible properties of number, duration, frequency and contingency, which are the carriers of the information.

Frequency is an abstraction built on abstractions—one intangible, number, divided by another intangible, time. A contingent frequency raises the pyramid of abstractions still higher. It is the rate at which an event occurs following the onset or offset of a conditioning event. Sensitivity to contingent frequency requires frequency estimation together with the estimation of contingency, which is itself a forbidding abstraction (defined below). Nonetheless, it has begun to appear that brains routinely compute conditional (contingent) frequencies. Their ability to do so may explain much that has been seen as the work of associative formation, a process that does not operate on abstractions like number, time and contingency.

In 1968, Robert Rescorla (1968) reported a simple experiment that changed in fundamental ways our conception of what has generally been called the associative process, the process that mediates Pavlovian conditioning, and, arguably, much else. The implications were so unsettling that they have not to this day been well digested by the community that studies learning, particularly by those that study it from a neurobiological perspective. What the results suggested was that the simple learning observed in Pavlovian conditioning paradigms arises from an information processing system sensitive to the contingencies between stimuli. If this implication is valid, then it changes our conception of the level of abstraction at which this basic learning process operates.

Rescorla's experiment followed from his reflections on the proper control procedures in Pavlovian conditioning experiments, published the previous year (Rescorla, 1967). Excitatory Pavlovian conditioning pairs a motivationally neutral stimulus (CS), like a tone or the illumination of a key, with a motivationally important stimulus (US),

like a mildly painful shock to the feet or the delivery of food. The tone comes on and some while later the shock is felt; or, the key is illuminated and some while later food is delivered. After a few or several such experiences, the rat freezes and defecates when the tone comes on, and the pigeon delivers food pecks on the illuminated key, suggesting in both cases that the subject anticipates the imminent appearance of the US. The learning in these paradigms appears to be a simple manifestation of the basic associative conception, which is that new conducting links are created in the nervous system by the temporal pairing of stimulus events. By these links, the CS excites the conditioned response. Therefore, Pavlovian conditioning has long been the paradigm of choice for investigating the properties of association formation.

That temporally pairing the CS and US was a *sine qua non* for association formation was long assumed to be self-evident (and it still is in some neurobiological circles--cf. Usherwood, 1993). To prove that a change in behavior was a result of association formation, experimenters typically ran an experimental condition in which two stimuli were repeatedly presented together (temporally paired) and a control condition in which they were widely separated in time (not temporally paired). If the change was seen in the experimental condition but not in the control condition, then the underlying process was said to be associative.

Temporal pairing and contingency had been tacitly assumed to be one and the same thing, until Rescorla (1967) pointed out that they could be dissociated, but that the commonly used control procedures did not do so. Control conditions in which the CS and US are never paired do not eliminate CS-US contingency, they replace one contingency with another. Rescorla (1967) pointed out that to determine whether it is temporal pairing or contingency that drives the conditioning process, one has to use the truly random control. In this control, the occurrence of the CS does not restrict in any way the time at which the US can occur, so the US must sometimes occur together with the CS, assuming that the CS is not a point event. Thus, the truly random control eliminates contingency but not temporal pairing, whereas the usual control eliminates temporal pairing but not contingency.

Rescorla (1968) realized that the most interesting version of the truly random control would be one that did not affect at all the temporal pairing in the usual experimental condition. He achieved this as follows: The experimental condition (really, now, the control) was approximately the usual condition for producing what is called the conditioned emotional reaction (more simply, conditioned fear). Hungry rats were first trained to press a lever in an experimental chamber to obtain food, until they did so readily and steadily throughout two-hour sessions. Then came five sessions during which the lever was blocked. In each of these sessions, twelve 2-minute long tones came on at more or less random intervals during each session. During these sessions, subjects also experienced occasional very short, mildly painful shocks to their feet. What Rescorla manipulated was the distribution of the shocks relative to the tones. For one group, the shocks, 12 of them per session, were completely contingent on the tone, they only occurred when it was on. The rats in a second group, a truly random group, also got 12 shocks each session while the tone was on, and they also got shocks at an equal frequency (0.5/minute) during the intervals when the tone was not on. This protocol did

not alter the number or frequency of tone-shock pairings, but it eliminated the tone-shock contingency. It also greatly increased the total number of shocks per session. To check whether that mattered, Rescorla ran a second truly random group: They got 12 shocks (the same total as the first group) but distributed at random without regard to the tone.

Before testing for the extent to which the rats in the different groups had learned to fear the tone, Rescorla first eliminated their fear of the experimental chamber, with two more sessions in which the lever was unblocked and there were no tones and no shocks. By the end of these two sessions, the rats had resumed pressing the lever for food. In several final sessions, the rats' conditioned fear of the tone was then measured by the effect of the tone on their willingness to continue pressing. If they feared the tone, they froze when it came on, and did not resume pressing until it went off. Although the rats in the first two conditions had the same tone-shock pairings, the rats in the contingent condition learned to fear the tone, whereas the rats in the truly random condition did not, nor did the rats in the other truly random condition. Thus, it is contingency and not temporal pairing that drives simple Pavlovian conditioning.

What is commonly called inhibitory conditioning also implies that simple conditioned behavior is a consequence of contingency, not temporal pairing. It is called inhibitory conditioning because the protocols for producing it are inverses of excitatory conditioning protocols. They are such that when the CS occurs, the US will not occur. The simplest protocol is the explicitly unpaired protocol: the US occurs only when the CS is absent (e.g. Rescorla, 1966). The conditioned response—what the animal learns to do or not do-- is the inverse of the excitatory response (LoLordo & Fairless, 1985; Rescorla, 1969).: The subject approaches and manipulates an inhibitory CS when it predicts a motivationally negative US. (For example, subjects approach the speaker from which the tone emanates if shocks never occur while the tone is on.) On the other hand, subjects avoid or refuse to manipulate a CS when it predicts that a positive US will not occur. (For example, if food never occurs when the key is illuminated, pigeons learn to distance themselves from the illuminated key.)

Students of basic learning have always recognized that inhibitory conditioning was just as fundamental as excitatory conditioning (Hull, 1929; Pavlov, 1928), but they have not faced the full implications of the fact that it results from systematically not pairing the CS and the US. Between them, inhibitory conditioning experiments and the truly random control experiments demonstrate that temporal pairing is neither necessary nor sufficient to produce associative learning (for reviews of the evidence, see Gallistel, 1990; Gallistel & Gibbon, 2000). What does appear to be necessary and perhaps sufficient is contingency.

Why has evidence that contingency not temporal pairing drives simple conditioning been so unsettling? Rescorla (1972, p. 10) put his finger on the problem, when he wrote, "We provide the animal with individual events, not correlations or information, and an adequate theory must detail how these events individually affect the animal. That is to say that we need a theory based on individual events." In other words, within the framework in which we have been wont to think about learning, the process had to operate at the level of events. It could not operate at the level of correlations and

information. We do in fact provide the animal with correlations and information in conditioning experiments, because when we construct a conditioning protocol we correlate the events in such a way that the CS provides information about the temporal locus of the US. What Rescorla presumably meant was not that we do not provide the animals with information, but rather that we take it for granted that the learning processes does not operate at that level of abstraction; we assume that it sensitive only to the individual (physically definable?) events, not to the information they convey.

In sum, the evidence that conditioning is driven by contingency unsettles us because it challenges what we have taken for granted about the level at which basic learning processes operate. Contingency, like number, arises at a more abstract level than the level of individual events. It can only be apprehended by a process that operates at the requisite level of abstraction--the level of information processing. A US is contingent on a CS to the extent that the CS provides information about the timing of the US.

The information an event provides to a subject is measured by the reduction in the subject's uncertainty (Shannon, 1948). This way of defining information leads directly to a rigorous quantification of information. It also accords with our every day intuition that the more information we have, the less uncertain we are. However, a disturbing feature of this definition is its subjectivity. The information conveyed by an event cannot be specified without reference to the subject's prior knowledge. What that means in the present case is that if the CS does not reduce the subject's uncertainty about the timing of the US, then it does not provide information.

The conceptual upheaval in the study of learning in the late 1960s was the result of a series of experiments demonstrating that whether a conditioned response developed or not depended on the information about the US provided by the CS (L. J. Kamin, 1969; Rescorla, 1968; Wagner, Logan, Haberlandt, & Price, 1968). These experiments were unsettling because they suggested that the conditioning process manifests the subjectivity inherent in Shannon's definition of information. Learning theorists have been struggling with these implications ever since (Barnet, Grahame, & Miller, 1993; Mackintosh, 1975; Pearce & Hall, 1980; Rescorla & Wagner, 1972; Wagner, 1981). What they have not done, however, is adopt the conceptual framework proposed by Shannon, despite its seeming relevance, its mathematical rigor, and its success in physics and engineering. The reluctance to adopt this framework appears to flow from two sources: dislike of the subjectivity inherent in it and a reluctance to believe that simple learning processes operate at so abstract a level. Here, I analyze some central issues in the study of conditioning from an information theoretic perspective.

Understanding Conditioning from an Information Theoretic Perspective

The subject's prior estimate of US frequency is the baseline from which the information provided by a CS may be measured. The information provided by the CS depends on the difference between its uncertainty about the timing of the next US before the CS comes on, and its uncertainty after it has come on. Its uncertainty before the CS comes on

depends on the information about US frequency provided by other stimuli experienced up to and including the experience of the CS onset (or offset).

Knowledge of the expected US frequency pure and simple (expected number of USs per unit of time) constitutes one end of an informational continuum about the timing of the US. The other end of the continuum is knowledge of the exact moment at which the US will occur. These two states of knowledge are represented graphically in Figure 1, which plots the cumulative probability that the US will have occurred as a function of the time that will have elapsed. The curve for the random rate process represents the maximum state of uncertainty (given any knowledge at all). At the other of the continuum is perfect certainty, which gives the step function (dashed line) in Figure 1.

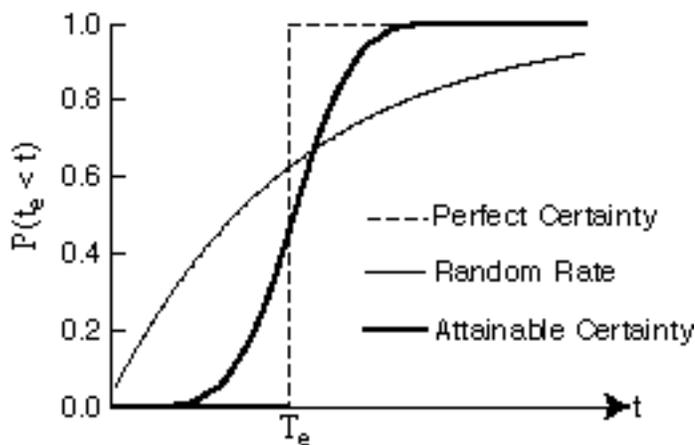


Figure 1. These curves represent the limits of certainty regarding the time of occurrence of a future event (t_e) when the expected interval to the next occurrence is T_e . The random rate curve represents the state of affairs when T_e is all that the subject knows. The perfect certainty curve represents the case when the subject knows that the next occurrence is a Gaussian random variable with expectation T_e and zero standard deviation (zero temporal uncertainty). Perfect certainty is not attainable. Attainable certainty is limited by the minimum psychologically attainable value for the standard deviation of the distribution of expected intervals. The attainable certainty curve is the integral of a Gaussian probability density function with expectation T_e and a standard deviation equal to the coefficient of variation (in essence, the Weber fraction) times the expectation.

When all that a subject knows is the number of USs so far observed and the duration of the interval of observation, its uncertainty about the time of occurrence of the next US is maximal because the probability that the US will occur in the next instant, is—so far as the subject knows(!)—constant, no matter how long it has been since the last US. This is the maximum possible state of uncertainty about the timing of the next US. It may or may not correspond to reality. If a random rate process generates the USs, then greater certainty than this is not attainable. The cumulative probability of the next US

under random-rate conditions is given by the random rate curve in Figure 1. The abscissa is time (t), with the origin (time 0) at the present moment. The longer the subject waits (the greater t is), the more certain it is that the next event will have occurred, that is, that t_e , the time of occurrence, is less than t . The probability that the US will have occurred approaches 1 exponentially, at a rate specified by the rate parameter, λ . Random rate processes are specified entirely by this one parameter (in contrast, for example, to Gaussian renewal processes, which require two parameters for their specification, a mean and standard deviation). The reciprocal of λ is T_e , the expected time to the next event.

A defining feature of a random rate process is that the cumulative probability function does not change as the present moment advances into the future. Put another way, the expected interval to the next event, T_e , is that same at every moment, it does not decrease as the interval since the last event lengthens. This property of a random rate process is counter-intuitive. Even when we know we are observing an approximation to a random rate process such as the rate of jackpots from a slot machine being played steadily, it is hard to resist the conviction that the longer it has been since the last jackpot, the shorter the time until the next one.

A corollary of this defining property of a random rate process is that the interval to the first observed event gives an unbiased (if often inaccurate) estimate of the expected (average) interval between events, regardless of the moment at which one begins to observe the process. A further corollary is that the subject begins to gain information about a random rate process as soon as it begins to observe it, even before the process generates an observable event. The observer/subject gains information even when nothing is happening, simply from the act of observing!

This startling manifestation of the subjectivity of information arises for the following reason: The subject's uncertainty about the rate it will observe is maximal at the moment it begins to observe, because—absent any prior knowledge about what it is going to observe—any observable rate is as likely as any other. Thus, at the outset of observation, the chances of observing a rate of about 1/s are the same as the chances of observing a rate of about 1/day—so far as the subject knows! (What the chances may in be in point of fact is neither here nor there. When we focus on information, we are concerned not with what is, but with what the subject knows about what is.) After 2 hours of observation without an event, these possibilities are no longer equally likely. The a posteriori probability that the rate is anywhere near 1/s given that one has observed an interval of 2 hours without an event is smaller by many orders of magnitude than the prior probability of such a rate (the probability at the outset of observation). By contrast, this same observation--this same stretch of nothingness—has little effect on the likelihood that the rate the subject is observing is on the order of 1/day, because an interevent interval of two hours or longer is highly likely when the rate being observed is on the order of 1/day (but wildly unlikely when the rate being observed is on the order of 1/s). Thus, the subject gains knowledge of the probable range of base rates even before it observes the first event. The length of an initial interval of observation during which nothing occurs reduces the uncertainty about what value the frequency being observed will turn out to have. Given an interval without any events, the subject can construct a

curve representing the lower limit of its current uncertainty about how long it will be to the next event by assuming that the base rate is no higher than 1 divided by the interval of observation.

In sum, an estimate of the base frequency, which estimate can be partly formed even in the absence of any observed events, establishes the maximum uncertainty about the timing of the next event. Further events can only reduce that uncertainty. The limiting case at the other end of the certainty continuum is when the subject knows exactly the expected interval, T_e , from a temporal marker event—CS onset, for example—to the US. If the origin in Figure 1 is at CS onset, then in this case, the cumulative probability that the US has occurred is zero for all $t < T_e$ and 1 for all $t \geq T_e$. In this case, of course, the expected interval to the next event is not constant as the present moment advances into the future; it decreases as t increases, that is, as the moment T_e gets closer.

Perfect certainty about the timing of a future event is unattainable because expected intervals cannot be known with infinite precision. This is a physical truth—or at least a truth of physics—but it takes a form in psychology that helps explain our intuitive sense that temporal pairing is critical to learning. One of the important findings to emerge from the study of human and animal timing is that Weber's law applies to remembered temporal intervals (Gibbon, 1977; Killeen & Weiss, 1987): The uncertainty regarding the true value of a remembered interval is proportionate to the remembered interval. Expected intervals are always remembered intervals, because expectations depend on information provided by past experience and carried forward in memory. Thus, the shorter the expected interval between a CS and a US, the more the CS reduces our uncertainty about the timing of the US. This information-theoretic consideration may explain why close forward temporal pairing has always seemed to be somehow important in Pavlovian conditioning—despite the extensive experimental evidence showing that it does not drive the conditioning process. If the CS does not precede the US, then it cannot reduce the subject's uncertainty about the timing of the US. That is why forward pairing is important (Barnet, Arnold, & Miller, 1991; Barnet, Cole, & Miller, 1997; Matzel, Held, & Miller, 1988; Miller & Barnet, 1993). Secondly, the residual uncertainty about the timing of the US is proportional to the CS-US interval (Weber's law), so the shorter that interval, the smaller the residual uncertainty. More generally, these considerations suggest an explanation for our intuitive sense that close forward temporal pairing creates a strong connection between any two stimuli. This intuition probably accounts for some of the enduring appeal of the associative theory of learning.

Weber's law establishes the limit of attainable certainty, because the Weber fraction establishes the limit to the accuracy with which we can know an unvarying interval (without resort to artificial methods and aids). Because of this irreducible uncertainty, our knowledge of a fixed temporal interval takes the form of a Gaussian probability density function, with a standard deviation proportionate to its mean. The coefficient of variation (in essence, the Weber fraction) is the ratio of the width of the distribution (as measured by its standard deviation) to its mean. The “attainable certainty” curve in Figure 1 is the integral of this distribution.

These information-theoretic considerations also allow us to understand in a rigorous but intuitively accessible way the experimental results that inspired the conceptual upheaval in learning theory at the end of the 1960's. Rescorla's (1968) experiment, which has already been discussed, was one of them. A second set of results published at the same time came from Kamin's (1967; 1969; 1969) experiments establishing the phenomena of blocking and overshadowing.

In a blocking experiment, one CS is paired with the US and when the conditioned response to this CS is well established, a second CS is introduced, which is always presented together with the first. Although it is thereby repeatedly paired with the US, the subject does not learn to respond to it (when it is presented alone on test trials). The conditioning to the first CS is said to block conditioning to the second CS.

As Kamin realized, the essence of the blocking phenomenon is that the second CS does not tell the subject anything it does not already know. More technically, it does not convey any information--in Shannon's (1948) rigorous sense. Recall that the information conveyed by an event—in this case a CS onset—depends on the subject's prior knowledge, because information is defined in terms of its effect on the recipient's uncertainty. The second CS does not reduce the subject's uncertainty about the time of US occurrence, because its onset and offset coincide with those of the first CS.

The blocking phenomenon makes it clear once again that what matters in conditioning is the information conveyed by a CS, not whether it is temporally paired with the US or not. And this is just another way of saying that conditioning is driven by contingency, because a US is contingent on a CS to the extent that the CS conveys information about the timing of the US. Although Kamin often wrote that the CS did not become conditioned to the US unless the subject was "surprised" by the US, he made no use of the conceptual framework created by Shannon; hence he was not able to give a rigorous, let alone quantitative, formulation of what it meant to say that a US was or was not surprising.

Overshadowing protocols are like blocking protocols except for the omission of the initial phase in which one of the two CSs is alone paired with the US. The two CSs are presented together from the outset of the experiment, but the subject learns to respond to one or the other but not both (Kamin, 1969; Mackintosh, 1971; Reynolds, 1961). Again, because they have exactly the same temporal relation to the US, both CS convey exactly the same information. This means that when the information conveyed by one of them alone is taken into account, the other conveys no information. What the conditioning process cares about is the information, not the carriers of that information, so it fastens on one carrier or the other.

Finally, there was a more complex experiment from the laboratory of Allan Wagner (Wagner et al., 1968), which showed that when the same information was conveyed both by a single CS and by some combination of CSs, the conditioning process picked out the one CS that by itself conveyed all the information rather than responding to two other CSs that between them conveyed the same information. This result is usually summarized by saying that subjects in conditioning experiments pick out the more valid

CS, but what exactly is meant by valid is unclear. A more rigorous formulation is to say that subjects respond to the minimum set of cues that conveys all the available information about US timing. If one cue by itself reduces the uncertainty about the timing of the US to a point where the other cues cannot between them effect any further reduction, then that is the only cue the subject responds to.

What the above amounts to is turning Rescorla's perplexity on its head: what matters in conditioning are not the individual events but rather the information that they convey. We need is a theory that deals with the information in the stream of events.

The information theoretic perspective also resolves a perplexity that has confronted thoughtful students of learning for so long without being resolved that they have generally decided to simply ignore it. The perplexity concerns how we are to understand the process of extinction if we take it as given that "we need a theory based on individual events" rather than on information. When we stop pairing the US with the CS, the subject eventually stops responding to the CS. This is the phenomenon of extinction. Long before it stops responding to the CS, the subject gives unmistakable evidence that it knows that the CS is no longer predicting the US. If the CS predicts a US that the subject likes, for example, food, then the subject gives clear signs of what we empathetically call frustration (Amsel, 1962). The perplexity is that the subject in extinction reacts not to events but rather to the failure of events (USs) to occur.

But cannot we conceive of this failure as itself an event? This has indeed been suggested as a remedy for our perplexity; it has been suggested that the failure of a US to occur can itself be conceived of as a US—a no-US (Dickinson, 1989). On this view, extinction results when we stop pairing a US with the CS with and start pairing a no-US. The first objection to this remedy is that it does violence to what has historically been considered the sine qua non of an event, namely, that it generate stimulus fluxes capable of exciting sensory receptors. There can be no sensory receptors sensitive to no-USs, because a no-US has no physical properties. In a materialist neurobiology, only physical stimuli can excite sensory receptors. Part of the motivation for seeking a theory that operates at the level of individual events rather than at the level of information is to have a neurobiologically transparent theory. In the theory as usually understood, the temporal pairing of events that excite different sensory inputs produces changes in synaptic conductances so that signals generated by CSs can excite conditioned responses by way of these new neuronal connections. Introducing events that have no physical properties into such a theory is awkward.

There isa deeper problem with this suggestion. Whatever other properties an event may or may not have, it surely has a locus in time. One can say when it happened. The problem with no-US events is that they may not have a specifiable time of occurrence. Consider an experiment in which we reinforce a CS at a random rate, which is what Rescorla (1968) did in the experiment with which we began this analysis. So long as the CS was on, the US it occurred at a random. When the CS was not on, there was another random rate of US occurrence (the background rate). The CS provided information about the timing of US occurrences to the extent that these two rates differed.

Under Rescorla's conditions, the CS did not specify where the US would occur within the CS. It was equally likely at any moment. Suppose we put a subject thus conditioned into extinction; that is, when the CS is on, USs no longer occur. If we think that extinction can be explained by the occurrence of no-USs, we must ask when the no-USs occur? It is impossible to say because only moments when the CS is on are relevant and the US was equally likely at all such moments. But if we cannot in principle say when the no-USs occur how can we talk about the processes set in motion by these no-USs? When are they set in motion? Must we conclude—falsely—that extinction of the conditioned response to the CS will not be observed under these conditions?

The phenomenon of extinction ought all by itself to convince us that the information theoretic perspective is a profitable perspective from which to understand conditioning, because the perplexity it poses vanishes when we consider it from this perspective. Recall, that intervals when nothing happens are themselves carriers of information about rate. Rieke, et al. (1997) apply information theory to the analysis of neural coding in simple preparations and show rather convincingly that the information is carried not by the spikes but rather by the intervals between the spikes, which is why a short neural signal can convey many more bits than there are spikes in the signal. When we put Rescorla's (1968) subjects into the extinction phase of the experiment, they experience a lengthening cumulative interval during which the CS has been on but there has been no US. As already explained, this lengthening interval when nothing happens conveys information about the upper limit on the current rate of US occurrence. There comes a time when that limit implied by this ever-lengthening interval lies well below the rate at which the US used to occur. Thus, the certainty that the rate of US occurrence in the presence of the CS is not what it used to be gets greater and greater the longer the extinction phase of the experiment continues. There is nothing mysterious about extinction when we explain it at the level of information, the mystery only arises when we try to explain it by processes set in motion by individual events. We need to focus instead on processes set in motion by the duration of the intervals in which nothing happens, processes sensitive to the flow of information.

Computing Contingent Frequency

What is required for a process to be sensitive to the information conveyed by the stream of events? Among other things, the process must time the intervals between events and count the numbers of events, because it is these abstract properties of the event stream that convey the information. Second, it must be capable of the computational operations required to extract from the intervals and counts the information in them. What might these operations look like?

First, let us consider the problem of computing contingency in an environment in which there is more than one CS, more than one stimulus whose presence may affect the expected time to the next US. This is often the case. It was the case, for example, even in Rescorla's (1968) simple experiment. We have so far discussed this experiment as if there was only one CS, the tone. In fact, there was a second CS, namely, the experimental chamber itself. Subjects never experienced foot shocks in their home cages, where they spent most of every day; they only experienced them in the experimental chamber, where

they spent only two hours each day. Thus, entry into the experimental chamber changed the rate at which shocks could be expected to occur. One might expect them, therefore, to have learned to fear the chamber in Rescorla's experiment, and, in fact, they did. On the first session following the sessions in which they got tones and shocks, they were so afraid of the chamber that they did little lever pressing. That is why Rescorla gave two sessions with the lever unblocked and no tones or shocks before the sessions in which he tested the effect of the tone; he had to eliminate their fear of the chamber.

Whenever the tone came on, the subjects for which the shocks were contingent on the presence of a tone were also in the experimental chamber, so the shocks they experienced were paired with both the tone and the chamber. By what process is it possible to determine that the shocks were contingent on the tone, not the chamber? Before addressing this question, let us return to an experiment mentioned briefly earlier because it suggests the subtlety and power of the process that computes contingency. Wagner, et al. (1968) ran two groups of rats in an experiment very much like Rescorla's but one that made the problem of multiple possible predictors more apparent and more complex. In their experiment, there were, in addition to the experimental chamber, three intermittent CSs, which, like the tone in Rescorla's experiment, came on and stayed on for 2 minutes at various times during the sessions when the rats got foot shocks. What Wagner, et al. (1968) manipulated was the correlations among these 3 CSs and the shock. Two of the CSs (designated A & B) were mutually exclusive: each occurred on half of the occasions (trials) on which a CS was presented, but never together with the other. A third CS (designated the X CS) always occurred together with the first two; it came on every time that either of them did, and never otherwise. For one group of rats, the shocks only occurred when the A and X stimuli were on, which means that they occurred on half the occasions on which the X stimulus was present. For a second group, shocks also occurred on half the occasions on which the X stimulus was present, but half the time it was the A stimulus that was also present and half the time the B stimulus. The first group learned to fear only the A stimulus, while the second group learned to fear only the X stimulus. What does it take to extract the "true" contingencies in the face of this kind of complexity?

Not as much as one might suppose. On the assumption that CSs have independent effects on the observed rates of US occurrence, there is a simple analytic solution to the contingency-computing problem confronted by the rats in Rescorla's (1968) experiment and Wagner, et al.'s (1968) experiment and, more generally, by the subjects in any basic Pavlovian conditioning experiment (Gallistel, 1990; Gallistel & Gibbon, 2000). It is given by the matrix equation,

$$\vec{r}_t = \mathbf{T}^0 \vec{r}_t,$$

where \vec{r}_t is the vector (the ordered list) of the rates of US occurrence predicted by each CS acting in isolation, \vec{r}_t is the raw rate vector,

$$\vec{Q}_r = \begin{pmatrix} \frac{N_1}{T_1} \\ \frac{N_2}{T_2} \\ \vdots \\ \frac{N_m}{T_m} \end{pmatrix}$$

and \mathbf{T}^{-1} is the inverse of the temporal coefficient matrix:

$$\mathbf{T} = \begin{vmatrix} 1 & \frac{T_{1,2}}{T_1} & \dots & \frac{T_{1,m}}{T_1} \\ \frac{T_{2,1}}{T_2} & 1 & \dots & \frac{T_{2,m}}{T_2} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{T_{m,1}}{T_m} & \frac{T_{m,2}}{T_m} & \dots & 1 \end{vmatrix}.$$

What comes out of this computation is the rate of US occurrence to be ascribed to each CS acting alone. The US is contingent on any CS to which a non-zero rate is attributed, whether negative or positive. When any such CS comes on, it changes the baseline expectation regarding the interval to the next US. If the rate ascribed to the CS is positive, the expected interval to the next US gets shorter; that is, the cumulative probability curve in Figure 1 approaches 1 more rapidly. If the rate ascribed to the CS is negative (as it will be in inhibitory conditioning paradigms), then the expected interval to the next US gets longer; that is, the cumulative probability curve in Figure 1 approaches 1 more slowly.

The quantities that go into the computation are the cumulative numbers of US occurrences ($N_1, N_2 \dots N_m$) in the presence of each of the CSs ($CS_1, CS_2, \dots CS_m$), the cumulative amounts of time that each CS has been observed to be present ($T_1, T_2 \dots T_m$) and the cumulative durations of the pairwise combinations of CSs (the $T_{i,j}$ s). These quantities must be obtained by cumulative timing of the intervals when individual CSs and pairwise combinations of them are present and by running counts of how often the US has occurred when each CS was present. These times and counts plus the machinery for doing the requisite computation are all that is needed to solve the contingency problem.

It will often happen that there is more than one solution to the contingency problem, because the animal's experience is inherently ambiguous, that is, it could have arisen from more than one state of the world. For example, in the Wagner, et al. (1968) condition where the shock occurred on half the AX trials and half the BX trials, the rats concluded that the US was contingent only on X. There are, however, an infinite number of other conclusions consistent both with the data they were given during the experiment and the assumption that the CSs were acting independently. It could have been that A and B were doing the predicting and not X. In that case, the rates of US occurrence ascribed to A and B would each be twice the rate ascribed to X in the solution that the rats

avored. Or, A could have accounted for 1 out of three USs, B for 1 out of 3, and X for 1 out of 3. And so on. There is no way of telling which was in fact the case.

The matrix computation automatically sounds an alarm whenever there is more than one solution, because then the determinant of the temporal coefficient matrix is 0, so the computation cannot be performed. The computation can only be performed when enough alternative predictors (CSs) have been omitted from consideration so as to make the determinant no longer 0. Or, if we prefer to think of the machinery as considering the simplest solutions first and adding to the set of predictors only when more predictors are required, then we can imagine the system doing the computation with successively more inclusive matrices (a successively broader canvass of the potential predictors) until it gets to matrices that have 0 determinant, at which point it has accounted for all the contingency that can be accounted for under the assumption of independent effects. At that point, the system will often have discovered more than one solution. However, there is an information-theoretic consideration that dictates a preferred solution in such cases. Interestingly, this consideration appears to dictate the solution the rats in fact favor in these inherently ambiguous situations.

The consideration is a version of Occam's razor, the principle of explanatory parsimony: that solution is preferred that minimizes the number of predictor variables (CSs to which an information-transmitting capacity is attributed). It is preferred because it is the most powerful solution, where power is measured by the average amount of information conveyed per CS.

The Rapidity of Conditioning and Extinction

Traditional models of classical conditioning have nothing to say about how rapidly conditioning and extinction ought to proceed. In associative models, the rate of conditioning is governed by physically meaningless free parameters, so there are no considerations that suggest even approximately what the values of these parameters should be. The information-theoretic perspective, however, has a natural, meaningful metric, which dictates approximately how rapidly conditioning and extinction should occur. The metric is the amount by which the subject's a priori uncertainty about the current contingency between the CS and the US has been reduced by the flow of events up to a given point in the experiment.

This is the metric that experimentalists use when they test a null hypothesis, the hypothesis that there is no contingency (no reliable relation) between the independent and the dependent variables in their experiment. The assumption is that in absence of observation (experiment), a zero contingency is as likely as any other. The question is whether the observations made in the course of the experiment make the assumption of zero-contingency (no effect) relatively unlikely. If the answer is, Yes, very unlikely, then we proceed to act on the assumption that the independent variable conveys information about the dependent variable. If the answer is, No, not very unlikely, then we proceed on the assumption that the independent variable conveys little information about the dependent variable.

The information gained at any point in a conditioning experiment about the contingency between a CS and a US depends on the durations of the observations in the presence and the absence of the CS and the numbers of USs observed in each condition. Elementary probabilistic reasoning with these quantities gives us a suitable measure, at least in the case where there are only two CSs, one of which is the steady background CS or context (hereafter B), and the other is the usual kind of intermittent CS (a tone or light). (The generalization to more complex cases is not difficult.) Call the intermittent CS A. We want to know how unlikely it is that there is no contingency between A and the US—the odds against the possibility that the rate of US occurrence predicted by A acting alone is 0.

Let T be the total interval of observation, T_a , the interval when A was present, and $T_b = T - T_a$ the interval when only the background was present. Likewise, let N be the total number of USs observed, N_a , the number when A was present, and $N_b = N - N_a$, the number when only the background was present. On the null hypothesis that A had no effect on the rate of US occurrence, the probability that an arbitrarily chosen one of the N USs (ignoring their order) happened to occur when A was present is T_a/T . The probability P_f of observing N_a or fewer events is given by the cumulative binomial probability function, as is the probability P_m of observing N_a or more events. When the number of events in T_a is approximately the expected number, the ratio P_f/P_m is approximately unity and the log of this ratio approximates 0. When the number of USs while A is present becomes improbably high relative to the number observed when A is not present, the ratio between the two cumulative probabilities becomes very large and its log approaches infinity. The growth of this ratio (or its logarithm) measures the extent to which the null hypothesis has become unlikely.

This measure of the information about the CS-US contingency that has so far accumulated works equally well in the case of a negative contingency. In that case, the number of USs observed while A is present becomes improbably low as the experiment progresses, so the ratio of the two cumulative probabilities becomes very small and its log approaches minus infinity. Thus, the absolute value of this logit (log of the odds) measures the growing certainty that there is a contingency; the farther this quantity is from 0, the more certain it is that the CS affects the rate of US occurrence.

Moreover, we know a priori what constitute reasonable values for this measure, because it has an objective probabilistic meaning. Speaking very conservatively, if the conditioned response appears at the point where the subject should be reasonably certain that there is a CS-US contingency, then the response should appear when this measure (the absolute value of the log of the odds) is less than 6, because when the logit equals 6, the odds are already 1,000,000:1 against the null hypothesis.

The empirical problem for the information-theoretic approach is that in the conditioning paradigms from which we have the best data, this measure is unreasonably high at the point in the experiment where the conditioned response is first seen. The odds against the null hypothesis at that point are on order of 10^{65} :1 (Gallistel & Gibbon, 2002). Thus, it cannot reasonably be argued that the conditioned response appears when the

possibility that the US is not contingent on the CS has become very unlikely. The problem is that it often does not appear until this possibility has become absurdly unlikely. The same is true for extinction. Reasoning very similar to that given above tells us how long it should take for us to observe the cessation of condition responding when we stop delivering USs in the presence of the CS (Gallistel, et al., 2001). In fact, however, it often takes much longer than this (Gallistel & Gibbon, 2002, Feb).

There are, however, paradigms in which subjects demonstrate a much greater sensitivity to the accumulation of information about current contingencies. In the matching paradigm, rewards (USs) become available at random moments at each of two foraging sites. When a reward becomes available at a site, it remains available until the subject comes to that site and harvests it. (Typically, rodent subjects harvest available rewards by pressing a lever, while bird subjects harvest rewards by pecking at keys.) The intervals until the next reward become available are determined by independent random rate (Poisson) processes. This means, as already explained, that the likelihood that a reward will be made available in the next instant is flat. Once made available, a reward remains available until it is harvested; therefore, the longer the subject has gone without visiting a site, the more certain it is that there is a reward waiting for it. The cumulative probability that a reward has become available during the subject's absence is described by the random rate curve in Figure 1, which approaches 1 at a rate determined by the availability rate (usually called the programmed reward rate) for that site. This characteristic means that subjects can increase the total amount of reward harvested per unit of foraging time by moving back and forth between the sites, even when the availability rates at the two sites differ greatly. The longer the subject has stayed at the richer site, the more certain it is that there is a reward waiting to be harvested at the poorer site, and the cost of a quick visit to the poorer site, measured in average time lost in harvesting the next reward at the richer site, is very small.

In a matching paradigm, the experimenter sets the rate parameters of the two sites, usually to unequal values. It is called a matching paradigm because the general finding is that the ratio of the amounts of time the subject spends at the two locations approximately equals the ratio of the rates at which it obtains rewards from them (Davison & McCarthy, 1988; Herrnstein, 1961). Importantly, the reward rates in this equation are defined by reference to total foraging time (usually session time, time spent in the chamber), not the time the subject spends at a given location. In economists' terms, the rates are incomes (rewards obtained per unit of time), not returns (rewards obtained per unit of time invested in foraging at a site). The relative amounts of time the subject spends at the two sites is determined by the ratio of the expected (average) stay durations at the two sites. These stay durations appear to be themselves determined by random rate processes within the subject, because they exhibit the flat hazard function that is the signature of a random rate process (Gibbon, 1995; Heyman, 1979, 1982). That is, the likelihood of the subjects leaving a site is independent of how long it has been there.

In information-theoretic terms, the relative amounts of time invested in different foraging sites when the subject is free to move back and forth between them at will is equal to the relative strengths of the contingencies between those sites and food availability. The strength of the contingency between a site and the availability of food is

measured by the rate parameter for the cumulative probability curve in Figure 1; the faster this curve approaches 1, the stronger the contingency between that site and the availability of food. We can use this paradigm to ask how rapidly subjects can detect and adjust to changes in the strengths of contingency. When we suddenly decrease the rate of reward at one site and increase it at the other, we produce a step change in the relative strengths of the contingencies. We know that the subject will eventually adjust the expected durations of its stays at the two sites so that the ratio of those expectations approximately equals the ratio of the new rates of reward. How rapidly it does so will depend in part on how sensitive it is to the accumulation of information indicating that the rates of reward have changed.

In more traditional terms, the change in rates of reward will lead to a strengthening of the conditioned response to the site that has become more richly rewarding and a weakening of the conditioned response to the site that has become poorer. These strengthenings and weakenings of the conditioned responses to these sites would appear to be closely analogous to if not identical with the strengthening of the conditioned response that occurs during simple acquisition and the weakening that occurs during simple extinction. Thus, this paradigm gives us another approach to the question of the rapidity of conditioning and extinction.

It turns out that when changes in contingency in the matching paradigm are rare, subjects adjust slowly to them (Gallistel, Mark, King, & Latham, 2001, October; Mazur, 1992, 1995; Mazur, 1996), but when they are frequent, subjects adjust about as fast as is in principle possible (Gallistel et al., 2001). This means that rats can approximate ideal detectors of changes in contingency; they have the maximum possible sensitivity to the flow of information carried by a sequence of events.

These findings suggest that the surprisingly sluggish appearance of conditioned behavior in some paradigms (e.g., pigeon autoshaping and rabbit eye-blink paradigms) and the equally surprising slow rates of extinction) do not reflect failures to detect the contingency between CS and US and changes in that contingency, but rather some strategy having to do with waiting to see how stable the presently observed contingency will prove to be.

Contingent Frequencies are not Conditional Probabilities

Before concluding, I want to call attention to the important distinction between frequency and probability, or, equivalently, between discrete and continuous frequency. Throughout this essay, frequency and rate mean number of events per unit of time. Because time itself is a continuous quantity (a real-valued variable), this quantity is also continuous (real valued). Frequency is often also used to mean the number of instances per number of observations, in which case it is a discrete (rationally valued) variable. More often than not, when psychologists talk about frequency, they seem to have discrete frequency in mind, although which sense of frequency is intended is often unclear. Also, the driving variable in simple conditioning is usually taken to be the conditional probability of the US, which is a discrete frequency.

It is important to keep the distinction between discrete and continuous frequencies in mind, because the two quantities behave differently. Discrete frequency is equivalent to probability, the conversion involving no more than converting from the rational to the decimal form of the number (for example, from 6/10 to 0.6). Continuous frequency converts to probability only by integrating a probability density function over a finite interval, and the value obtained depends on the interval of integration. When discrete probabilities are independent, they combine multiplicatively to determine joint probabilities. They combine additively only if they refer to mutually exclusive observations, and then their sum cannot be greater than 1. When continuous frequencies are independent, they combine additively, and their sum can be any real number.

One reason for recalling this distinction is that the driving variable in simple associative conditioning is often taken to be the conditional probability of the US. Indeed, Rescorla (1967, 1968) defined contingency in terms of conditional probabilities, and many have followed his lead. The contingent frequencies that are the foundation of the analysis given here are not conditional probabilities. Gallistel and Gibbon (2000, p. 333) review the difficulties that arise when one tries to apply an analysis based on conditional probabilities to the many conditioning protocols in which it is difficult or impossible to say what constitutes a trial. The problem is that if trials cannot be defined, then the observed discrete frequencies cannot be defined, because, by the usual definition, the probability of the US is the number of US occurrences in a given number of trials. In an analysis based on contingent frequencies rather than conditional probabilities, these difficulties do not arise.

Conclusions

Information theory provides a powerful, rigorous and quantitative framework within which to analyze the process of conditioning. Contingent frequency is a fundamental notion in such an analysis. A contingent frequency is the rate of US occurrence to be expected in the presence of a given CS (conceived of as operating alone). Working within this framework turns on its head the traditional assumption that conditioning has to be understood in terms of processes set in motion by individual events, more particularly, processes set in motion by the temporal pairing of events. On this analysis, the events are relevant only as carriers of information. Moreover, the information resides in the abstract quantities defined by the events, the intervals between the events, and the numbers of events.

On the usual analysis, the conditioning process does not operate at this level of abstraction; the timing of intervals and the counting of events play no role in traditional models. Although the intervals between CSs and USs and the number of CS-US pairings are assumed to affect the strengths of the resulting associations, the learning mechanism is not imagined to do computations on symbols representing durations and numbers.

If, however, the conditioning process is in fact driven by the information conveyed by the stream of events, then the process must be doing just that. Thus, the validity of the information-theoretic approach to the understanding of conditioning bears strongly on the more general theory of mind and brain that will prevail within

psychology. Should we conceive of the brain as a plastic net whose connectivity is molded by its experience so as to adapt its input-output function to the exigencies of that experience, without symbolically representing abstract properties of that experience? Or should we conceive of it as an information processing organ doing computations on symbols that represent abstract properties of its experience, like number and duration?

References

- Amsel, A. (1962). Frustrative nonreward in partial reinforcement and discrimination learning. Psychological Review, *69*, 306-328.
- Barnet, R. C., Arnold, H. M., & Miller, R. R. (1991). Simultaneous conditioning demonstrated in second-order conditioning: Evidence for similar associative structure in forward and simultaneous conditioning. Learning and Motivation, *22*(253-268).
- Barnet, R. C., Cole, R. P., & Miller, R. R. (1997). Temporal integration in second-order conditioning and sensory preconditioning. Animal Learning and Behavior, *25*(2), 221-233.
- Barnet, R. C., Grahame, N. J., & Miller, R. R. (1993). Local context ^{and} the comparator hypothesis. Animal Learning and Behavior, *21*, 1-13.
- Davison, M., & McCarthy, D. (1988). The matching law: A research review. Hillsdale, NJ: Erlbaum.
- Dickinson, A. (1989). Expectancy theory in animal conditioning. In S. B. Klein & R. R. Mowrer (Eds.), Contemporary learning theories: Pavlovian conditioning and the status of traditional learning theory (pp. 279-308). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Gallistel, C. R. (1990). The organization of learning. Cambridge, MA: Bradford Books/MIT Press.
- Gallistel, C. R., & Gibbon, J. (2000). Time, rate and conditioning. Psychological Review, *107*, 289-344.
- Gallistel, C. R., & Gibbon, J. (2002, Feb). The symbolic foundations of conditioned behavior. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Gallistel, C. R., Mark, T. A., King, A. P., & Latham, P. E. (2001, October). The Rat Approximates an Ideal Detector of Changes in Rates of Reward: Implications for the Law of Effect. Journal of Experimental Psychology: Animal Behavior Processes(October)??-??
- Gibbon, J. (1977). Scalar expectancy theory and Weber's Law in animal timing. Psychological Review, *84*, 279-335.

- Gibbon, J. (1995). Dynamics of time matching: Arousal makes better seem worse. Psychonomic Bulletin and Review, *2*(2), 208-215.
- Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. Journal of the Experimental Analysis of Behavior, *4*, 267-272.
- Heyman, G. M. (1979). A Markov model description of changeover probabilities on concurrent variable-interval schedules. Journal of the Experimental Analysis of Behavior, *31*, 41-51.
- Heyman, G. M. (1982). Is time allocation unconditioned behavior? In M. Commons & R. Herrnstein & H. Rachlin (Eds.), Quantitative Analyses of Behavior, Vol. 2: Matching and Maximizing Accounts (Vol. 2, pp. 459-490). Cambridge, Mass: Ballinger Press.
- Hull, C. L. (1929). A functional interpretation of the conditioned reflex. Psychol. Rev., *36*, 498-511.
- Kamin, L. J. (1967). "Attention-like" processes in classical conditioning. In M. R. Jones (Ed.), Miami symposium on the prediction of behavior: aversive stimulation (pp. 9-33). Miami: Univ. Miami Press.
- Kamin, L. J. (1969). Predictability, surprise, attention, and conditioning. In B. A. Campbell & R. M. Church (Eds.), Punishment and aversive behavior (pp. 276-296). New York: Appleton-Century-Crofts.
- Kamin, L. J. (1969). Selective association and conditioning. In N. J. Mackintosh & W. K. Honig (Eds.), Fundamental issues in associative learning (pp. 42-64). Halifax: Dalhousie University Press.
- Killeen, P. R., & Weiss, N. A. (1987). Optimal timing and the Weber function. Psychological Review, *94*, 455-468.
- LoLordo, V. M., & Fairless, J. L. (1985). Pavlovian conditioned inhibition: The literature since 1969. In R. R. Miller & N. E. Spear (Eds.), Information processing in animals. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Mackintosh, N. J. (1971). An analysis of overshadowing and blocking. Quarterly Journal of Experimental Psychology, *23*, 118-125.
- Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. Psychol. Rev., *82*, 276-298.
- Matzel, L. D., Held, F. P., & Miller, R. R. (1988). Information and expression of simultaneous and backward associations: Implications for contiguity theory. Learning and Motivation, *19*, 317-344.

- Mazur, J. E. (1992). Choice behavior in transition: Development of preference with ratio and interval schedules. JEP:ABP, 18, 364-378.
- Mazur, J. E. (1995). Development of preference and spontaneous recovery in choice behavior with concurrent variable-interval schedules. Animal Learning and Behavior, 23(1), 93-103.
- Mazur, J. E. (1996). Past experience, recency, and spontaneous recovery in choice behavior. Animal Learning & Behavior, 24(1), 1-10.
- Miller, R. R., & Barnet, R. C. (1993). The role of time in elementary associations. Current Directions in Psychological Science, 2, 106-111.
- Pavlov, I. V. (1928). Lectures on conditioned reflexes: The higher nervous activity of animals (H. Gantt, Trans.). London: Lawrence & Wishart.
- Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variation in the effectiveness of conditioned but not of unconditioned stimuli. Psychol. Rev., 87, 532-552.
- Rescorla, R. A. (1966). Predictability and the number of pairings in Pavlovian fear conditioning. Psychonomic Science, 4, 383-384.
- Rescorla, R. A. (1967). Pavlovian conditioning and its proper control procedures. Psychol. Rev., 74, 71-80.
- Rescorla, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. Journal of Comparative and Physiological Psychology, 66(1), 1-5.
- Rescorla, R. A. (1969). Pavlovian conditioned inhibition. Psychological Bulletin, 72, 77-94.
- Rescorla, R. A. (1972). Informational variables in Pavlovian conditioning. In G. H. Bower (Ed.), The psychology of learning and motivation (Vol. 6, pp. 1-46). New York: Academic.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), Classical conditioning II (pp. 64-99). New York: Appleton-Century-Crofts.
- Reynolds, G. S. (1961). Attention in the pigeon. Journal of the Experimental Analysis of Behavior, 4, 203-208.
- Rieke, F., Warland, D., de Ruyter van Steveninck, R., & Bialek, W. (1997). Spikes: Exploring the neural code. Cambridge, MA: MIT Press.

- Shannon, C. E. (1948). A mathematical theory of communication. Bell Systems Technical Journal, 27, 379-423, 623-656.
- Usherwood, P. N. R. (1993). Memories are made of this. Trends in Neurosciences, 16(11), 427-429.
- Wagner, A. R. (1981). SOP: A model of automatic memory processing in animal behavior. In N. E. Spear & R. R. Miller (Eds.), Information processing in animals: memory mechanisms (pp. 5-47). Hillsdale, NJ: Lawrence Erlbaum.
- Wagner, A. R., Logan, F. A., Haberlandt, K., & Price, T. (1968). Stimulus selection in animal discrimination learning. Journal of Experimental Psychology, 76(2), 171-180.