

Conditioning from An Information Processing Perspective

C. R. Gallistel

Abstract

The framework provided by Claude Shannon's (1948) theory of information leads to a quantitatively oriented reconceptualization of the processes that mediate conditioning. 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 to the intangible properties of number, duration, frequency and contingency. In this view, a stimulus becomes a CS if its onset substantially reduces the subject's uncertainty about the time of occurrence of the next US. One way to represent the subject's knowledge about the time of occurrence of the CS is by the cumulative probability function. This function has two limiting forms: 1) The state of maximal uncertainty (minimal knowledge) is represented by the inverse exponential function associated with the random rate condition in which the US is equally likely at any moment. 2) The state of maximal certainty is represented by the cumulative normal function whose expectation is equal to the CS-US latency minus the time elapsed since CS onset and whose standard deviation is the Weber fraction times the CS-US latency.

The late 1960s saw the publication of three experimental papers that demanded far reaching revision in our conception of conditioning. Rescorla (1968) showed that if one removed the contingency between a to-be-conditioned stimulus (CS) and the unconditioned stimulus (US) without altering the temporal pairing, the subject never developed a conditioned response to the CS (Figure 1). This result implied that it was contingency not the temporal pairing that generated conditioned responding.

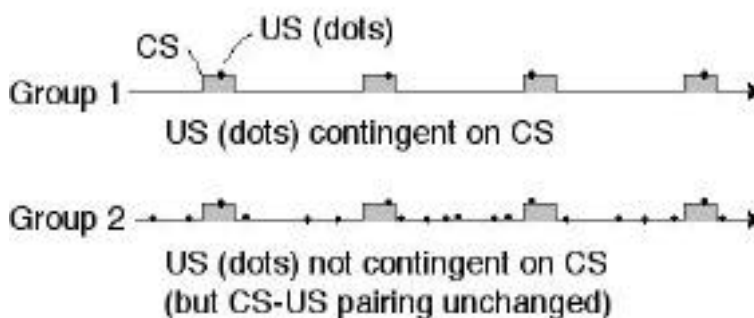


Figure 1. Rescorla's (1968) truly random control experiment. The temporal pairing of USs with CSs is identical in the two groups but in the second group, there is no CS-US contingency. Subjects in the second group do not develop a conditioned response to the CS.

Kamin (1969) showed that pairing a CS and a US did not produce a conditioned response to the CS if it was always presented together with another, already conditioned CS (Figure 2). The already conditioned CS blocked conditioning to the newly introduced CS. This implied that when a subject "expected" the US, that is when it was not "surprised" by it, then pairing it with a CS did not produce a conditioned response. I have followed Kamin in putting scare quotes around 'surprise.' He did so presumably to warn that it was unclear what the scientific meaning of the word might be. And, as a behaviorist, he was chary of the notion of expectation, without which there can be no surprise.

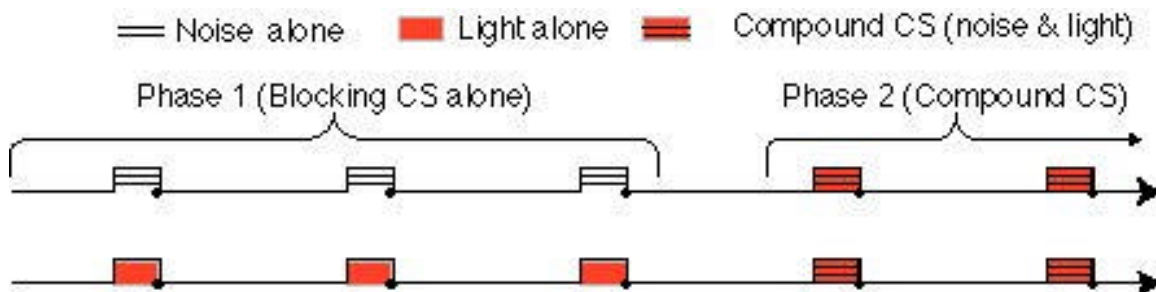


Figure 2. Kamin's blocking protocol. Subjects are conditioned with one CS in Phase 1. In Phase 2, a second CS is given coincident with the already conditioned CS. Subjects do not develop a conditioned response to the second CS.

Kamin (1969) further showed that when two CSs were presented together from the outset of conditioning, the conditioned response that developed to one of them was much stronger than the conditioned response that developed to the other. Kamin called this 'overshadowing.' A striking example of it had been published some years before by Reynolds (1961, see Figure 3). He interpreted it in terms of selective attention, an interpretation that has continued to be popular (Mackintosh, 1975; Pearce & Hall, 1980). These experiments showed that when two stimuli are redundant predictors the conditioning process eliminates one of them from consideration. And, of course, they also show that temporal pairing is not sufficient to produce conditioning.



Figure 3. Overshadowing. Reynolds (1961) trained to peck the red key with a white triangle for food reward. (Pecking the other key did not yield reward.) When he tested each of the four stimulus components in isolation, he found that one bird was conditioned to the white triangle, while the other was conditioned to the red background.

Wagner and his collaborators (Wagner, Logan, Haberlandt, & Price, 1968) showed that when one CS is more reliably informative of US delivery, it is the only one to which the conditioned response develops (Figure 3). It does not develop to the other less reliably informative CSs, even though they are frequently paired with the US. Once again, this showed that it is informational considerations determine conditioning not processes set in motion by temporal pairing.

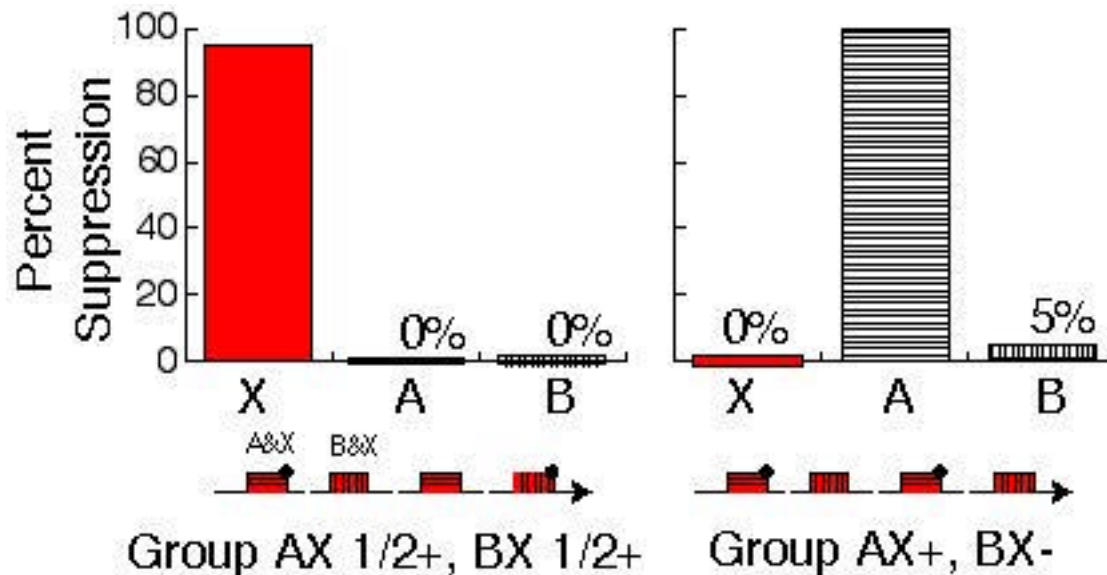


Figure 4. For both groups, there are an infinity of rates of US occurrence that could be ascribed to the different CSs to give rates indistinguishable from the experimentally programmed rates, under the assumption that the rates predicted by the different CSs are additive. However, in both groups, subjects settle on the only solution that ascribes everything to one CS. That is, a conditioned response develops only to the relatively more valid predictor.

When the history of this period is written, it will be a challenge to explain why psychologists were surprised to discover that temporally pairing a CS and US was not sufficient to produce conditioned behavior because it had long been known that this was not necessary. When a CS predicts the omission of an otherwise to be expected US, a conditioned response develops just as readily as it does when a CS predicts a not otherwise to be expected US (Figure 4). This is called inhibitory conditioning, because a CS that predicts the omission of the US suppresses the conditioned response to a CS that predicts that CS. The effects of inhibitory conditioning are not, however, limited to antagonizing the effects of excitatory conditionings. Inhibitory conditioning produces conditioned responses of its own (Hearst & Jenkins, 1974; Wasserman, Franklin, & Hearst, 1974). In inhibitory conditioning, the CS and US are never paired, so this phenomenon demonstrates that CS-US pairing is not necessary.

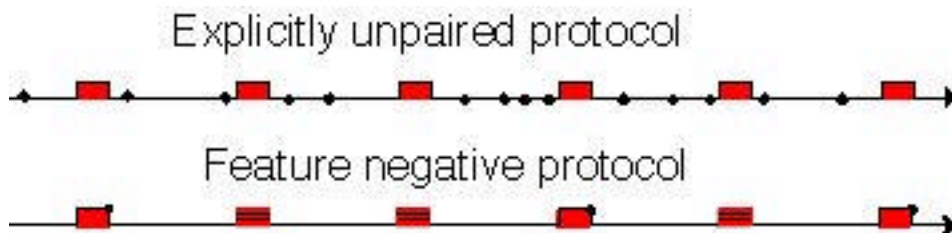


Figure 5. Two protocols for producing conditioned inhibition. In both, the US is never paired with the CS. Nonetheless a conditioned response develops.

Contingency Not Temporal Pairing

In summary, by 1970, it was known that temporal pairing is neither necessary nor sufficient to produce conditioned responding. What is necessary and sufficient is CS-US contingency, the CS must provide information about the US. The importance of the experimental findings just reviewed was immediately recognized; they have been frequently and widely replicated, and they are recounted at length in every textbook on conditioning and learning written in the last quarter century. Nonetheless, the conviction that temporal pairing is the key requirement in associative learning persists undiminished (Krasne, Hndbk chap., (Gluck & Thompson, 1987; Hawkins & Kandel, 1984; Miller & Escobar, 2002; Usherwood, 1993)). Why?

Rescorla (1972, p. 10) put his finger on the problem that has blocked our progress in coming to terms with the implications of these experiments, 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."

He could not, on reflection, have meant to say that we do not provide the animal with correlations and information, because when we construct our experimental protocols we manipulate the correlations between the various CSs and the US so as to vary the information that the CSs provide about the US. What he presumably meant was that we take it for granted that the conditioning process does not operate at this level of abstraction, because we take it for granted that it involves the creation or modification of associative connections. Events—and more particularly the temporal and spatial convergence of events—are what increment associations, not information.

For most psychologists and neurobiologists, information does not seem substantial enough to do the job of making connections. Information must be extracted from the flow of events by computational processes that are driven by the intervals between events and the numbers of events—rather than by their more tangible properties, the properties that excite sensory receptors and are therefore the natural focus of empiricist theories of mind and brain.

Information Drives Conditioning

From an information-processing perspective, however, it is the information in the protocol that drives the conditioning process. The assumption that conditioning involves the forging of connections is a conceptual roadblock. It prevents our coming to terms with the implications of these experiments and the many other experimental results which imply that conditioned behavior is a manifestation of the brain's information processing activity.

About the time he published his paper on blocking and overshadowing, Leon Kamin moved from Macmaster to Princeton. Fifty miles to the north of him there, Claude Shannon worked in the Bell Laboratories, where, twenty years earlier, he published a famous paper that could have enabled Kamin to remove the scare quotes from around surprise (Shannon, 1948). In it, he showed how to define information in such a way as to make it quantifiable. This made it a respectable scientific concept. It also laid a foundation for the modern communications industry, for computer science, and for the modern understanding of thermodynamics. It should, I argue, also become a foundation for our understanding of the brain (cf Rieke, Warland, de Ruyter van Steveninck, & Bialek, 1997) and the process of conditioning (Gallistel & Gibbon, 2002).

In Shannon's definition, a signal conveys information to the extent that it reduces the receiver's uncertainty about the state of the world. This leads directly to a rigorous quantification of information. It also accords with two everyday intuitions: 1) The more information we have, the less uncertain we are. 2) Something that tells us what we already know does not convey information to us.

This definition introduces what, for some, is a discomfiting subjectivity into the foundations of the science. The information conveyed by a signal cannot be measured in the absence of knowledge about the receiver's (that is, the subject's) state of uncertainty before and after the signal. To apply Shannon's insight to our understanding of conditioning, we have to consider what it is about the world that the subject represents (knows), what the limits are on the precision with which the subject knows it, and how those limits are captured in the subject's representation, that is, how the subject represents its uncertainty about the state of the world.

Some Simple Principles

In a first pass, the application of Shannon's definition to conditioning might adopt the following principles: 1) Only informative CSs elicit conditioned responses. 2) A CS informs a subject to the extent that it reduces the subject's uncertainty about the timing of the next US. 3) The brain's information processing is efficient; it maximizes bandwidth—the amount of information carried by a given signal. 4) The information provided by a CS is not a property of the CS itself but rather of its temporal and numerical distribution relative to the US. Thus, the brain needs timers and counters in order to extract from an experimental protocol the information that the CS carries about the timing of the US. This is the level of abstraction at which the conditioning process operates. 5)

The uncertainty in the representation of temporal and numerical magnitudes is proportional to their magnitude (Weber's law). 6) Poisson (random rate) processes provide the other limit, the limit on uncertainty. In a random rate process, the US is equally likely at every moment.

Some Simply Constructed Functions

The cumulative probability function is a felicitous way for a subject to represent what it knows about the timing of the next US (Figure 5). It emerges naturally from the cumulative distribution of US experiences, which is a simple tabulation of intervals and counts. It is the fraction of experiences as a function of the durations of the intervals from one experience to the next.

The simplest case to consider is when the US functions as its own CS, that is, when the timing of the previous US provides all the information that the subject has about when to expect the next US. In a fixed time schedule, the next US is delivered at a fixed interval after the previous one. In such a schedule, each US is fully informative about when the next US will occur. The subject's uncertainty is limited not by uncertainty in the world but by the precision with which it can represent the world—in this case, by the precision with which it can represent the US-US interval. In a random rate schedule, the next occurrence of the US is equally likely at any moment after the previous occurrence, the US. In such a schedule, one US is completely uninformative about the timing of the next US.

We know that subjects distinguish between these two cases. When the timing of the next US is predictable, the conditioned response anticipates it (LaBarbera & Church, 1974). The question from an information processing perspective is, What representation of its experience and what processing of that representation would be required in order for the subject to distinguish between the two cases?

Figure 5A gives the cumulative record from the case in which the US occurs at random (heavy line) and the case in which it occurs at a fixed interval. The expectations of the generating processes were 60 seconds in both cases. The slopes (l and m) of the cumulative records estimate those expectations. The slopes of the lines drawn from the origins of these cumulative records to points somewhere on them differ somewhat depending on where in the record those termination points are, because these slopes are estimates of the expectations not the expectations themselves. On average, the longer the cumulative records, that is, the more events they record, the less these slopes differ from one another and from the expectations of which they are estimates. Put another way, the longer the subject's experience with the event-generating process is, the more accurate its estimate of its expectation will tend to be.

Figure 5B shows the construction of the cumulative distribution functions—by sorting the intervals between events, stacking (cumulating) the sorted intervals, and dividing the stack by the number of intervals in it (normalizing).

The inverse exponential curve in Figure 5B is the limit to which the cumulative distribution function for the random rate process converges as the number of intervals in the stack becomes arbitrarily large. The curve is entirely specified by the expectation of the random rate process, that is, by the quantity estimated by l.

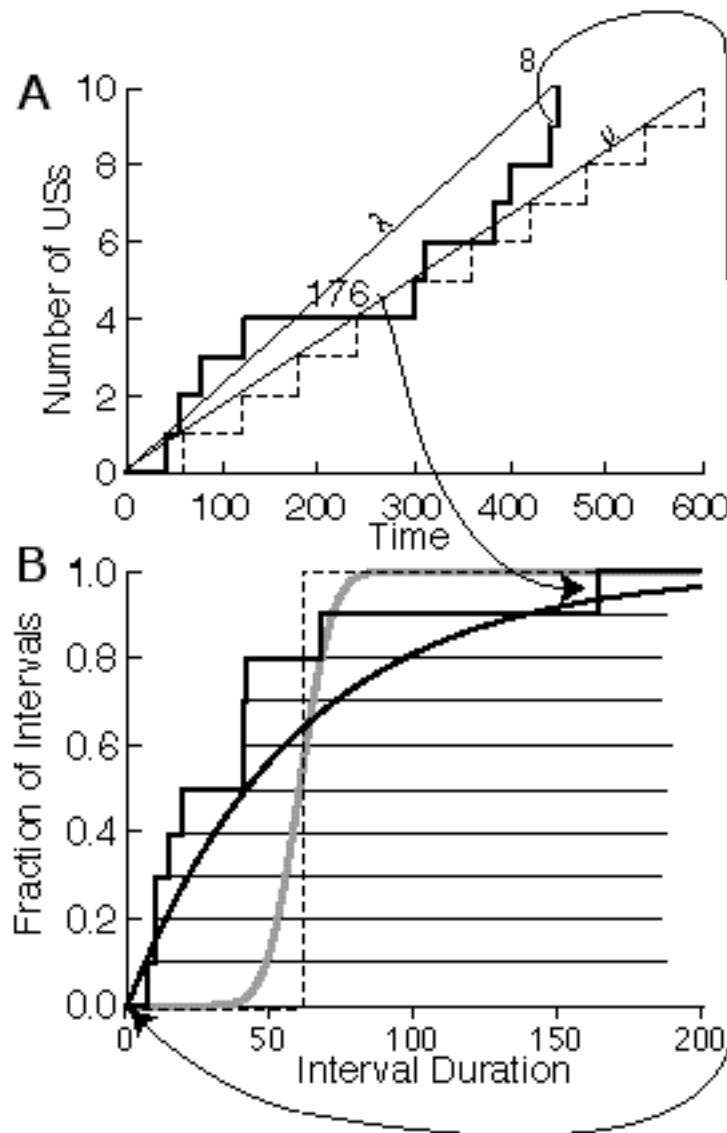


Figure 5. A. Two cumulative records. The dashed line is the cumulative record of the first 10 events generated by a fixed time schedule. The solid line is from a random rate schedule with the same expectation. **B.** Empirical and theoretical cumulative distribution functions. The dashed line is the empirical distribution for the fixed time process, the solid black line with steps is for the random rate process. Each event produces a unit step in the cumulative record. These unit step functions differ only in the length of the interval before the step. Sorting the empirical functions on the basis of the length of these intervals and then stacking the steps produces the cumulative distribution function. Arrows from A to B show fate of two of the intervals. The black curve is the inverse exponential function, which the empirical function approaches in the limit as the number of events goes to infinity. The gray curve is the limit on the representation of a step function imposed by the scalar variability (Weber-law characteristic) in the brain's representation of magnitudes.

The cumulative distribution function for the fixed time process would be a step function if the intervals between events were always exactly the same and if the subject could represent them with infinite precision—neither of which is true. With the aid of modern technology, however, it is easy for us to create an event generating process (scheduling mechanism) in which the variation in the intervals between USs is much smaller than the variation in the subject's representation of those intervals, which is roughly $\pm 15\%$ in pigeons, rats and mice (Church, Meck, & Gibbon, 1994; Gibbon & Church, 1992; King, McDonald, & Gallistel, in press). That the variation in the subject's representation of an unvarying interval is proportional to the duration of that interval is a central assumption of Gibbon's (1977) scalar expectancy theory. It is justified by extensive experimental findings in the animal timing literature (see Gallistel & Gibbon, 2000, for review). The cumulative Gaussian curve in Figure 5B (gray sigmoid) is the representation of a fixed interval schedule that results from scalar variability in the representing process, with a Weber fraction (coefficient of variation) of 0.15. It is the cumulative Gaussian distribution with a standard deviation equal to 15% of its mean.

Acquisition of a Timed CR

The decision whether the timing of one US provides information about the timing of the next US can be based on a comparison of the cumulative distribution function to its exponential approximation, because the odds against the no-information assumption are an a priori knowable function of the maximum deviation between the function and that approximation together with the number of events represented in the distribution function (Gallistel & Gibbon, 2002). This gives an acquisition theory for timed responding.

Once it has been decided that the time of occurrence of the next US can be estimated from the time of occurrence of the last US, the cumulative distribution function can be used to set decision criteria for when to anticipate the next US. If a US predicts the time of occurrence of the next US then it constitutes a temporal landmark, by reference to which the animal can locate itself in time. Once, it has perceived the landmark it can estimate its progress through time toward the next occurrence of the US, and it can use the cumulative distribution function to set decision criteria (response thresholds) that enable it to anticipate the next US with whatever degree of certainty it wants. In Figure 6, it has set its start criterion so as to be 90% certain that it has begun to respond before the next US occurs. It has advanced past the point in time corresponding to this criterion, and so it has begun to respond. Its location in time corresponds approximately to the location of the steepest portion of the distribution function, so its expectation of the US is maximal at this location in time.

The past is (alas!) a fallible guide to the future, and so the subject's expectation may be disappointed. It must therefore have a stop criterion as well as a start criterion; otherwise it would continue to respond indefinitely if and when the US failed to materialize. In Figure 6, it has set its stop criterion at a level such that it

stops responding when there is less than 1 chance in 20 that the anticipated US is going to happen within the limits of its uncertainty regarding the anticipated time. That is, it stops only when it becomes highly likely that its expectation has been violated. This account of the timing of a conditioned response is Gibbon's (1977) Scalar Expectancy Theory for the peak procedure.

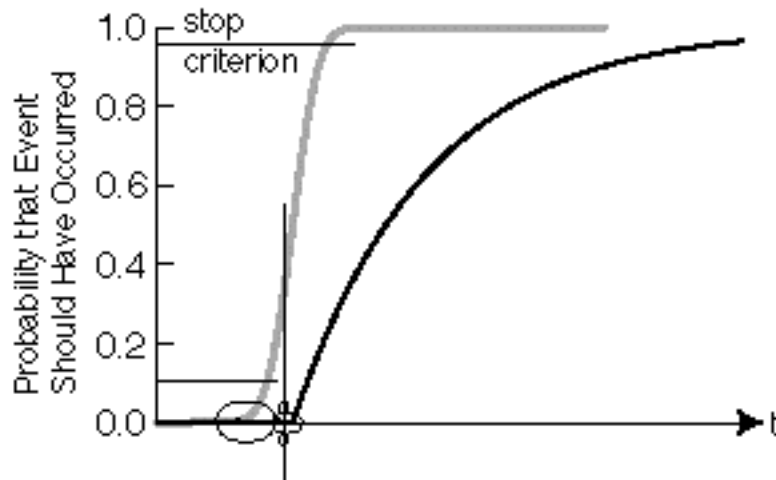


Figure 6. Cumulative distribution functions as temporal frameworks. When the next US is located in time (gray Gaussian), then the subject (the mouse on the time axis) can use its timer to measure its advance toward and past that time and it can set decision criteria that govern the timing of its anticipatory behavior (conditioned behavior). When the next US has not been temporally localized, the mouse remains always at the origin of the function (black inverse exponential) as the mouse advances through time; it never catches up with it.

When, by contrast, the time of occurrence of a US does not predict the time of occurrence of the next US, then it does not serve as a landmark by reference to which the subject may assess its approach to the next US. The subject has only an egocentric temporal framework, so its relation to the cumulative US distribution function does not change as it advances through time (Figure 6, inverse exponential curve). This is a counterintuitive property of random rate processes: the expected time to the next event is independent of how long it has been since the last event. That is why, the time between the onset of observation and the first occurrence of an event gives an unbiased (but, of course, noisy) estimate of the expectation of the process (the interval between events). All of these are corollaries of the fact that in a random rate process, the time at which one event occurs provides no information about the time at which the next event will occur.

Detecting Changes in the Rate of US Occurrence

To say that the occurrence of an event in a random rate process provides no information about when the next event will occur is not to say that it provides no information about anything. It provides information about the value of λ . The objective uncertainty about the true value of λ is inversely proportional to the square root of the number of USs so far experienced. The more events, the less the uncertainty; hence the more information the subject has acquired from the

stream of events. But this holds true only so long as the precision with which the subject can represent an estimate of I 's value is not the limiting factor. The decrease in the inherent or unavoidable uncertainty or a rate estimate as the square root of the number of events increases means that the first few events the subject observes convey substantial information about the value of I , but the amount conveyed by further events diminishes rapidly.

Of course, the rate may change. This violation of expectation cannot be detected by the just described process for detecting the failure of the US to occur at an expected time, because the defining feature of a random time schedule is that there is no expected time; the US is equally likely at every moment. This poses an interesting conceptual challenge to associative theories of extinction and conditioned inhibition, because they are driven by events rather than by the information conveyed by a stream of events. In an associative theory, whether trial based (Rescorla & Wagner, 1972) or real time (see also Brandon, Vogel, & Wagner, 2002, this volume; Wagner, 1981), the mechanism of extinction is activated by the failure of an expected event to occur. This "event"—the failure—must itself have a time of occurrence. But when there is no expected time of occurrence for an expected event, how can its failure to occur have a time of occurrence?

In one information processing approach, changes in the rate of US occurrence are hypothesized to be detected by a simple ongoing computation (Gallistel & Gibbon, 2002, p. ? ff). Despite its simplicity, it has been shown to be an ideal detector of such changes (Gallistel, Mark, King, & Latham, 2001); it uses all the information available in the sequence of interevent intervals. What it does is look for seeming changes in the slope of the cumulative record and calculate how unlikely they would be to have arisen through chance fluctuations (Figure 7). Comparing the behavior predicted by this ideal detector to the behavior observed when rats adjust to unexpected but frequent changes in the relative rates of reward in a Herrnstein matching paradigm show that the rat is an ideal detector of changes in random rates (Gallistel et al., 2001). It adjusts to them about as fast as is in principle possible.

Explaining Cue Competition

The explanations of basic conditioning so far outlined all assume that the brain has built into it mechanisms that embody in their structure the statistics of random rate processes and mechanisms that detect systematic deviations from randomness. This assumption constitutes a much stronger commitment to rationalism and nativism than empiricist learning theorists have traditionally been comfortable with. However, to the biologically inclined, at least, it seems no more exceptional than saying that the eye embodies in its structure the laws of optics. Why should not the structure of learning mechanisms reflect the structure of the problems they solve just as much as does the structure of sensory or metabolic organs (see Gallistel, 1999, for an elaboration of this argument).

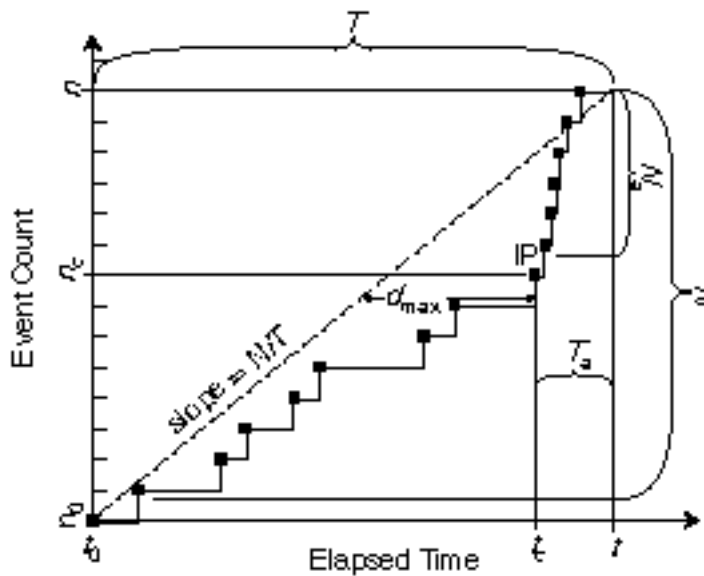


Figure 7. Changing the rate of a random rate process creates an inflection point (IP) in the cumulative record, which will be at the maximum deviation (d_{\max}) from the straight line approximation to the record (dashed line with slope N/T). The creditworthiness of the hypothesis that a putative inflection point corresponds to a genuine change in rate is a simple calculation of binomial likelihood ratios. The calculation compares the proportion N_e/N to the proportion T_e/T to test whether they differ by more than is to be expected from chance fluctuations.

Everything that has been developed above for the univariate case—the case where the US serves as its own predictor—applies with minimal modification to the multivariate case, where the US is predicted by other stimuli (conventionally called CSs). However, when we allow other predictors, we confront the essential problem in multivariate time series analysis—determining what predicts what. The classic experiments on the truly random control (or background conditioning), blocking, overshadowing, and relative validity (reviewed in the introduction) show that subjects do a sophisticated job of solving this problem. The assumption that the information processing mechanisms that mediate conditioned behavior are structured by the properties of random rate processes leads once again to a simple model for the mechanism that solves what is commonly called the cue competition problem, because random rates are additive.

Rate Estimation Theory (Gallistel, 1990, Chap. 13) assumes that for the purpose of determining which CSs predict which US rates, the brain treats all rates as random rates. It does so, because under that assumption—which renders the effects of CSs additive—there is a simple universal solution to the problem of estimating the rates of US occurrence to be associated with each CS, that is, with

each possible predictor of those rates. The essence of the solution is given by the matrix algebra formula for the vector (list), \vec{r}_t , of true rate estimates

$$\vec{r}_t = \mathbf{T}^{-1} \vec{r}$$

where

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

and

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

This formula assigns to each CS an estimate of the rate of US occurrence that would be observed if that CS were the sole determinant of I, the rate of US occurrence, except in cases where there are redundant predictors, as there are in the overshadowing and relative validity protocols. In those cases, the additive solution to the rate estimation is not unique, that is, there are an infinity of solutions consistent with the assumption of additivity. When that happens, the determinant of the matrix is zero. In such cases, the bandwidth maximization principle comes into play. The rate estimating mechanism eliminates possible predictors so as to find the smallest set of predictors that can fully account for the observed rates of US occurrence. In cases, where there is more than one such set (e.g., in overshadowing), it either "flips a coin", as was apparently done by the process that generated the results shown in Figure 3, or it relies on other criteria to decide among predictors that are, from an information processing standpoint, equally valid potential predictors. (An Excel spreadsheet implementing this calculation and the calculation that detects a change in the rate of US occurrence ascribed to the influence of a CS may be downloaded from the SQAB website.)

Analyticity, Simplicity, and Intuitiveness

Two attractions of the information processing approach are that its predictions are analytically determined and they do not depend on free parameters. By traditional scientific criteria, this means that these models are simpler than associative models. Associative models abound in free parameters (see, for example, Brandon et al., 2002), and their predictions can generally be determined only by numerical methods, that is, by computer simulation, the details of which are rarely available for scrutiny and are hard to comprehend when they are available. A further attraction is that the predictions of the information processing models follow from information processing principles in an intuitively obvious way. It does not require neither mathematics nor faith in inscrutable computer simulations to see why the theory predicts what it predicts:

Background conditioning (truly random control) The CS does not affect expected time to the next US. Therefore, the information conveyed by the CS is zero. This follows directly from the application of Shannon's definition of information to conditioning and the principle that only informative CSs elicited conditioned responses.

Blocking. Again, the addition of the second CS, in compound with the already conditioned CS, does not add any information. The US occurs at the time already predicted by the first CS. Because only informative CSs elicit conditioned responses, there is no response to the second CS, no matter how often it is paired with the US.

Overshadowing. The CSs are redundant; either CS alone can reduce the subject's uncertainty about the time of US occurrence by as much as the two together. Bandwidth maximization (aka parsimony or Occam's razor) applies: the amount of information conveyed per conveying signal is maximized by letting one of them do all the work.

Relative validity. Again, the CSs are redundant. One could partition the information conveyed among two or more of the CSs, but bandwidth is maximized by letting one do all the work.

Conditioned inhibition. When the CS comes on, the time at which cumulative probability function rises moves away from the subject's current location in time. This, too, is informative. When there is no information about when the next US will occur, all moments are equally likely. When an inhibitory CS comes on, more proximal temporal moments become much less likely than later moments.

The importance of temporal pairing. Temporal pairing is neither necessary nor sufficient for conditioning, but that does not mean that it is unimportant. It is important because the degree to which a temporal landmark like CS onset reduces a subject's uncertainty about the timing of the next US is proportional to its proximity to the US. The closer it is, the more it reduces the subject's uncertainty. This follows from the fact that the degree of imprecision in a subject's representation of an interval is proportional to the duration of the interval—Gibbon's (1977) principle of scalar variability in the representation of temporal intervals.

Time scale invariance. Gallistel and Gibbon (2000) review a variety of findings that suggest that the conditioning process is, within broad limits, time scale invariant. Changing the intervals in a conditioning protocol has no effect on the results, provided they are all changed by the same factor (the scaling factor). Information processing models are naturally time scale invariant because the information in the flow of events is carried by the proportions among the intervals not by their absolute durations. Scaling the flow up or down changes the amount of time it takes to deliver the information, but it does not change the amount of information delivered by a given number of events.

Associative models, by contrast, are not time scale invariant. If they are trial based models, then they must assume a trial duration. Rescaling the protocol will then move events (e.g., CS and US onsets) that were within the same trial into different trials, or vice versa, with strong consequences for the predictions of the model. Real-time associative models make complex assumptions about the dynamics of stimulus traces (see, for example, Brandon et al., 2002). Rescaling the protocol changes where events fall relative to these trace dynamics, again with strong consequences for the predictions of the model.

Summary

Shannon's definition of the information conveyed by a signal in terms of the reduction in the subject's uncertainty about the state of the world applies naturally to conditioning paradigms, where CSs (and/or the subject's responses) reduce its uncertainty about when the next US (next reinforcer) will occur. Simple principles rooted in conventional information theoretic considerations, such as bandwidth maximization, give intuitively obvious predictions for phenomena that are generally agreed to be central to our understanding of the conditioning process—the effects of background conditioning, inhibitory conditioning, blocking, overshadowing, and relative validity. They also explain aspects of conditioned behavior that associative theories rarely attempt to explain, for example, the rarely emphasized fact that the latency of the conditioned response is proportional to the CS-US interval (Gallistel & Gibbon, 2000, for review). Finally, they explain something that it appears deeply difficult for associative models to explain, the time scale invariance of the conditioning process (Gallistel & Gibbon, 2000, for review).

References

Brandon, S., Vogel, E., & Wagner, A. (2002). Stimulus representation in SOP:I. Theoretical rationalization and some implications. Behavioural Processes (SQAB '02 special issue)(?).

Church, R. M., Meck, W. H., & Gibbon, J. (1994). Application of scalar timing theory to individual trials. Journal of Experimental Psychology: Animal Behavior Processes, 20(2), 135-155.

Gallistel, C. R. (1990). The organization of learning. Cambridge, MA: Bradford Books/MIT Press.

Gallistel, C. R. (1999). The Replacement of General-Purpose Learning Models with Adaptively Specialized Learning Modules. In M. S. Gazzaniga (Ed.), The cognitive neurosciences. 2nd ed. (2 ed.). Cambridge, MA: MIT Press.

Gallistel, C. R., & Gibbon, J. (2000). Time, rate and conditioning. Psychological Review, 107, 289-344.

Gallistel, C. R., & Gibbon, J. (2002). The symbolic foundations of conditioned behavior. Mahwah, NJ: Lawrence Erlbaum Associates.

Gallistel, C. R., Mark, T. A., King, A. P., & Latham, P. E. (2001). 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, 27, 354-372.

Gibbon, J. (1977). Scalar expectancy theory and Weber's Law in animal timing. Psychological Review, 84, 279-335.

Gibbon, J., & Church, R. M. (1992). Comparison of variance and covariance patterns in parallel and serial theories of timing. Journal of the Experimental Analysis of Behavior, 57, 393-406.

Gluck, M. A., & Thompson, R. F. (1987). Modeling the neural substrates of associative learning and memory: a computational approach. Psychol. Rev., 94(2), 176-191.

Hawkins, R. D., & Kandel, E. R. (1984). Is there a cell-biological alphabet for simple forms of learning? Psych. Rev., 91, 375-391.

Hearst, E., & Jenkins, H. M. (1974). Signtracking: The stimulus-reinforcer relation and directed action. Austin, TX: Psychonomic Society.

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.

King, A. S., McDonald, R., & Gallistel, C. R. (in press). Screening for mice that remember incorrectly. International Journal of Comparative Psychology.

LaBarbera, J. D., & Church, R. M. (1974). Magnitude of fear as a function of the expected time to an aversive event. Animal Learning and Behavior, 2, 199-202.

Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. Psychol. Rev., 82, 276-298.

Miller, R. R., & Escobar, M. (2002). Laws and models of basic conditioning. In C. R. Gallistel (Ed.), Stevens handbook of experimental psychology. 3rd ed. Vol. 3. Learning and motivation (Vol. 3, pp. 47-102). New York: Wiley.

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. (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. (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.

Wasserman, E. A., Franklin, S. R., & Hearst, E. (1974). Pavlovian appetitive contingencies and approach versus withdrawal to conditioned stimuli in pigeons. Journal of Comparative and Physiological Psychology, 86, 616-627.