

MULTIPLE TIME SCALES IS WELL NAMED

JOHN GIBBON

COLUMBIA UNIVERSITY

Staddon and Higa's article is a critique of scalar expectancy theory, and a proposed alternative, multiple time scales. The critique is generally flawed, both factually and logically. The alternative is bewildering in its flexibility, opaque in its quantitative description, and never addressed to real data.

Key words: timing, scalar expectancy theory

The article by Staddon and Higa is first a critique of, or more properly a diatribe on, scalar expectancy theory (SET), and second a proposed alternative. The paper is not satisfactory in either section, for very different reasons. The criticisms of SET are often ill taken and occasionally bizarre. On the other hand, the multiple-time-scales (MTS) alternative proposed next is a bit too aptly named. Staddon and Higa mix and match time scales to qualitative descriptions of effects in the timing literature with a bewildering variety. Power scales, log scales, summed exponential scales (MTS), and even linear scales (see below, time-left explanation) are used almost at will.

When the original version of the article was submitted to *JEAB*, the editor requested a review from me, and I submitted a signed review. The paper was subsequently revised, re-submitted, and accepted to be published with commentary. My reading of the revision is that not much has changed (with the exception of Figure 6, discussed below). I am therefore providing a commentary which is essentially portions of my original review, with discussion of minor points deleted and edited for clarity.

Poisson Variance

The authors begin with a central theme, namely that the pacemaker-accumulator idea in SET is difficult to reconcile with Weber's law. Indeed, the description implies that SET is continuously scrambling to get around this difficulty by adding parameters to the ac-

count. A central feature of a pacemaker accumulator is Poisson variance, which does not accord with the scalar property: proportional rescaling of timing distributions. Actually, from the outset the information-processing account of SET was designed with just the scalar property in mind. The idea was that a system of this kind (a) must account for the scalar property and (b) may do so in more than one way. The clock, memory, and decision process stages identified three potential sources of scalar variance, all of them with the critical property that random variation in the system be multiplicative. The three sources so identified were pacemaker rate variation (by the way, it makes no difference whether one assumes trial-to-trial variation or within-trial variation; cf. Gibbon, 1992), memory translation multipliers in storage and retrieval, and threshold variation. The authors have, I think, misunderstood the way in which the notion of a Poisson pacemaker (in both early and late versions of SET) was conceived. Gibbon (1992) showed that even small sources of multiplicative variance do indeed render Poisson variability negligible. The assumption of a Poisson pacemaker was a convenient one that is physiologically plausible, because there are many neural systems with Poisson variability. The key features of the theory did not rely on the Poisson pacemaker idea but rather on a mechanism that integrates activity over time with multiplicative variance.

Logarithmic Subjective Time Scale

Staddon and Higa then go on to posit a logarithmic perceptual subjective time scale. They argue that the temporal bisection result with indifference at the geometric mean is a natural and straightforward instantiation of equal subjective distance from the two anchor

This research was supported by NIMH Grant MH41649-12.

Address correspondence to John Gibbon, Biopsychology, Unit 50, Columbia University, 722 W. 168th St., New York, New York 10032 (E-mail: jg34@columbia.edu).

points, short and long. As Staddon and Higa state, the geometric mean finding drops out of the log time scale rather simply. As long as variance is constant and symmetric on the log scale, the midpoint between two remembered values should be at the geometric mean. In fact this was the motivation for examining the subjective time scale in the time-left experiments (see below). For bisection, given the linear scale that SET assumes, the psychophysical function is obtained not by simple differences on the subjective scale but rather by similarities on this scale, where similarity is defined, like other discriminative functions in SET, as a ratio, in this case of the probe to the referents. Staddon and Higa have not described the SET analysis of bisection accurately, in my view. Gibbon (1981) analyzed logarithmic and linear scales using a difference rule for the logarithmic scale and the ratio rule for the linear scale. Although it is obviously true that differences on a log scale reduce to ratios on a linear scale, it is not true that the form of the psychometric function is well fit by the logarithmic scale. Indeed, it is shown that symmetry of remembered distributions of the standards on the log scale results in a poor fit to the psychometric functions, albeit with an appropriate indifference point at the geometric mean.

There are other data that deal with the form of the psychophysical function as well, which again make for difficulties with the log scale. Church and Gibbon (1982) examined what form temporal generalization functions ought to take when animals discriminate whether a probe stimulus is the same or different from a reinforced standard. The scalar property is again found, but with near symmetry on the real time axis, which is not to be expected from a subjective log scale. Rather a negative skew, which is never found in real data, is predicted for the log scale.

Time-Left Experiments

The time-left experiments (Gibbon & Church, 1981) were motivated precisely by the geometric-mean finding for bisection. But Staddon and Higa do not seem to have grasped the fundamental features of these experiments, particularly Experiment 2. Subjects choosing between an elapsing alternative and another, the standard, which is fixed (beginning "right now"), should show choic-

es reflecting proximity on the log scale to food, which is much shorter in the middle of the interval for the time-left side than it is for the standard. But importantly, if a log unit is added to both sides by doubling the elapsing and standard interval, as was done in the second experiment, then a discrimination based on proximity to food on the log scale predicts *no change* in the point of indifference between the two. And this of course is ruled out by the data from Experiment 2.

Even with omission of Experiment 2, Staddon and Higa evidently find the log scale not tenable for the time-left experiments, and perform a bizarre contortion to obtain the desired result. They argue that subjects perceive time in a logarithmic fashion but then take *inverses* (i.e., antilogs) and then calculate immediacies (expectancies) in real time! That is, subjects perceive time logarithmically but have the good sense to ignore their perceptions (or the good sense to pick the antilog transform) when faced with a choice between delays to food. They use real time, as though the log scale associated with perception of time were irrelevant. SET does not need such convoluted reasoning to arrive at the appropriate result, because the point of indifference should scale linearly with real time provided that time is perceived linearly with real time.

Staddon and Higa claim a kind of "perceptual constancy mechanism" to account for the use of linear time. They argue "for a separation between the animal's capacity to assess reinforcement rates and its capacity to use a decaying memory trace as a stimulus" (p. 222). The example espoused to justify such a convoluted interpretation is that there is a smaller two-point threshold on the hand than on the back "but we do not feel that our hands are larger than our back" (p. 223). This kind of reasoning from just noticeable differences is exactly the sort that led Fechner to the log scale. Do the authors wish to challenge his inference? In any case, this example makes no sense applied to the time scale. The reason for judging that the hands are smaller than the back has nothing to do with the two-point threshold. Visual sensory data are undoubtedly used in determining size, but for temporal judgments there is no alternative receptor. What are the alternative data used when judging time durations?

The authors go on to argue for a log time scale on other grounds (p. 224). However, I confess to be completely lost on Equations 8 and 9. How is it that “internal effects” are related to time differentials by these equations? Multiplying through by t or b , they seem to be saying that a small change in t (dt) is equal to the standard deviation of t times the “internal effects” or differential on z , $dz = \sigma_t dz$. But what is z ? And where is the variance in this system? Although it is certainly true that Equations 10 and 11 follow by integration, the justification for Equations 8 and 9 remains mysterious. This is but one of several examples of what might be called mini models in this paper, none of which are developed in sufficient detail for us to evaluate them. Moreover, it is not always clear which mini model is being espoused, because power, log, and MTS are frequently given equal weight, and are introduced where convenient for one or another purpose. It is also unclear how the mini models relate to the presumably central, more developed MTS model. Sometimes, as in Figure 1, they are touted as equivalent, but of course in many other respects, such as slopes, they are clearly not.

Start/Stop/Spread Correlations in the Peak Procedure

Staddon and Higa insinuate that positive start-stop correlations (as opposed to negative start-spread correlations) are a difficulty to be overcome in SET. In fact this analysis was designed to attempt to isolate the relative contributions of memory and decision variance in this procedure. Far from being unexpected, the analysis showed that both sources of variance are present.

MTS, in contrast, appears to have only a start threshold. How would such a system handle the results of the peak procedure? If a stop threshold is to be permitted, how does the animal “know” where to place it, or even come to “expect” reinforcement at a particular time, and so detect its omission?

MTS Memory Timing Model and the Variance Problem

Figure 6 in the current article is very different from that in the original. The original Weber fraction was shown to decrease about 55% over an eightfold range in interfood in-

terval (a Poisson system would show a Weber fraction decrease of about 65% over this range). The description in the original, however, closely matched that in the current version (pp. 232–233). The text appears to describe the Weber fraction as resulting from the decreasing slope of the memory trace.

The new Figure 6 shows rising and falling functions. The degree of change near the origin is small for values of $\lambda \geq 1.5$. However, small on what scale? Ordinate values are omitted from the new Figure 6. Indeed, it would be useful to have more than simply ordinate values in the new Figure 6; it would be much more compelling to see actual Weber fractions from real data obtained over these ranges.

The problem with where variance arises in this system is endemic to all of the discussion. Staddon, at a recent meeting (personal communication), described the Weber fraction as obtained from a simulation of MTS with threshold variance. That is, a given threshold on the memory decay trace functions with constant variability would induce the curves shown in Figure 6. Is threshold variance then the source of variability producing the Weber fraction? The description in the text does not make this clear at all. Indeed, there is no talk of variance throughout this paper, and the quantitative mechanisms underlying variance, or even those underlying the scales, are generally opaque (MTS is never described with closed forms). And if the lack of detail on variance mechanisms is a problem for me, I suspect I am not alone.

It is almost as though animals are perfectly accurate but somehow the slope of the subjective time functions induces a Weber fraction without any variability (at least in Equations 19 and 20). If threshold variance is indeed what induces the Weber fraction here, then that description needs to be made explicit. Threshold variance is one of the sources of scalar variance in SET, but it is always explicitly so. Here it is not clear (a) whether the Weber fraction increases or decreases or (b) whether the functions in Figure 6 are a result of a simulation with true variability or are the result of some calculations on slopes, as the text suggests.

Even on a qualitative level, it seems that the MTS decay functions, which can approach arbitrarily close to zero over a rather short

range (see Gallistel's commentary), must pose processing problems for real live subjects that have internal processing noise; for example, constant threshold variance must avoid negative threshold values as the MTS subjective scale approaches zero. But truncating the threshold distribution at zero means that it is no longer constant, and is less so the closer the decay function is to zero.

Deterministic accounts are in principle error free, but a major thrust of psychophysics for many years has been to understand sources of variability and error. We need to know more about how errors are produced in this system before an evaluation can be made. On

its face, MTS appears to avoid the variance problem by simply remaining silent.

REFERENCES

- Church, R. M., & Gibbon, J. (1982). Temporal generalization. *Journal of Experimental Psychology: Animal Behavior Processes*, 8, 165–186.
- Gibbon, J. (1981). On the form and location of the psychometric bisection function for time. *Journal of Mathematical Psychology*, 24, 58–87.
- Gibbon, J. (1992). Ubiquity of scalar timing with a Poisson clock. *Journal of Mathematical Psychology*, 36, 283–293.
- Gibbon, J., & Church, R. M. (1981). Time-left: Linear vs. Logarithmic subjective time. *Journal of Experimental Psychology: Animal Behavior Processes*, 7, 87–108.

MODELING MODELING

PETER R. KILLEEN

ARIZONA STATE UNIVERSITY



Models are tools; they need to fit both the hand and the task. Presence or absence of a feature such as a pacemaker or a cascade is not in itself good. Or bad. Criteria for model evaluation involve benefit-cost ratios, with the numerator a function of the range of phenomena explained, goodness of fit, consistency with other nearby models, and intangibles such as beauty. The denominator is a function of complexity, the number of phenomena that must be ignored, and the effort necessary to incorporate the model into one's parlance. Neither part of the ratio can yet be evaluated for MTS, whose authors provide some cogent challenges to SET.

Key words: models, pacemakers, theories

If you think models are about the truth, or that there is a best timing model, then you are in trouble. There is no best model, any more than there is a best car model or swimsuit model, even though each of us may have our favorites. It all depends on what you want to do with the model. Nor are models theoreticians' guns of domination, any more than data are empiricists' bullets of assault. War

This work was supported by NSF Grant IBN 9408022 and NIMH Grant K05 MH01293.

Address correspondence to Peter Killeen, Department of Psychology, Box 1104, Arizona State University, Tempe, Arizona 85287-1104 (E-mail: Killeen@asu.edu).

games can be fun, however, especially when you do not have entangling alliances with the principles, and can just watch them swat it out. Will the grapplers or the punchers win this year? It happens to empiricists too, failing to replicate and sniffing about controls, but somehow it is more fun when it is the guys in suits, the guys who prefer *ln* to *log*, the guys who try to explain *your* data to *you*, are going at it.

And it is a good thing for them to do, too. Do *you* want to attempt to tell a Gibbon that his integral is improper? Or a Staddon that he might have one too many layers in his

leaky cascade? Besides, if John ****on shows that John ****on's models are wrong or irrelevant, well, then, so much less for us to feel guilty about not mastering; and if we wait a few years, with any luck J*** H*** will upset the survivor from this round.

So, what do we mortals hope to learn from this paper, other than that, apparently, heavy experimental equipment cannot be operated after doing too many integrals? Perhaps that conceptual analysis is always relevant, no matter how mathematical a theory. Or how verbal. Pictures, either explicit (e.g., Staddon & Higa's Figures 3 and 4) or implicit (e.g., Skinner's reflex reserve, pictured in Killeen, 1988), are the foundation of models. They are the skeletons that get muscled by mathematics and padded by words. Until recently models were valued to the extent that they had such a visualizable (*anschaulich*) foundation (Miller, 1984). One is never wrong to refer back to pictures for guidance.

Models are go-betweens: They go between the data and our sense of understanding. This is as true of Skinner's models (e.g., the three-term contingency, reinforcement, the mand) as it is of the paraphernalia of timing. To the extent that models fail to make contact with the data or with our sense of understanding, they fail. In the first case their death is quick; in the second, the models languish until Burked by boredom.

If we sit out too many rounds, we might wind up as mere paying spectators, couch potatoes, Monday morning quarterbacks. This is safer, retrospectively smarter, but ultimately a less reinforcing posture. After all, it is *we* who are paying *them* with our attention. This is a game worth getting into.

And play we must, because all understanding involves models—reference to systems that exist in a different domain than the thing studied. Loose models make vague reference to ambiguous or ad hoc causes. Tighter models are more careful about definitions and avoid gratuitous entities. Models of phenomena are not causes of phenomena; they are descriptions of hypothetical structures or functions that aid explanation, prediction, and control. Their promulgators may believe that such structures exist, or that they have a particular instantiation, or that they were sighted last year in the medulla oblongata; interesting, certainly, but such beliefs are not

what is most important about the models. A model may be functionally brilliant without being structurally enlightening. Such was the case for Newton's mechanics. To get to the heart of a matter does not require that the matter have ventricles.

Washing the Baby

To use models effectively, and to understand those of others, scientists must undergo some training. In many cases this means studying some math. In *The Behavior of Organisms*, Skinner (1938) formulated rich verbal descriptions of acquisition, discrimination, and extinction that presaged Hull's more mathematical treatments, and were directly translatable into them. He inveighed against the "paper dolls" of models only after his logarithmic extinction curves were criticized (by Ellson, 1939). It was perhaps easier for Skinner to throw out his model baby than to learn how to bathe it. It is an advantage of a model and its packaging, however—not a defect—that its emissions and omissions can be discriminated. Most other parents would have held their breath and cleaned.

Staddon and Higa detect an odor arising from Gibbon and Church's baby, and don't like the way it is dealt with. "Too many layers of diapers" they seem to say. But it ain't their kid, and because they are raising one of their own, it is best just to sluice out the nest.

Watching the Baby

So, how about their neonate, the pacemaker-free model of timing? This issue records the village tour, as we peer into the pram, listen to the neighbors' comments, and make a few of our own. But the new infant is scarcely old enough to have soiled itself, let alone to have done much with its limbs; we will have to judge it from its looks, not its parents' promises. Analysis of its sibling (Staddon & Higa, 1997) would be worthwhile; but this one and its commentaries are already demanding enough, and we just do not have enough attention left in our wallets, what with the other game we hoped to catch. In fact, it is easier to listen to the local gossip about the kid than to study its features, and that is why some of you skipped ahead to here just now.

Making Paces

So then, what about some of its parents' claims for the multiple-time-scale (MTS)

model: Is it possible to have a pacemaker-free model of timing? What does a pacemaker add that makes it worth having, and with what have Staddon and Higa replaced it?

A pacemaker is a model generator. Say you want to boil an egg for 4 minutes, and have to decide when it is time to quench it. You need to subtract the time of day you started it from the current time of day, and compare the difference to a model—a model of 4 minutes. A prosthesis helps. You can buy an egg timer that is a sand clock; invert it and it takes 4 minutes for the sand to empty into the bottom vessel. Inverting it resets the origin to zero, and its completion represents a model of 4 minutes.

If you want to grill a fish for 10 minutes, you need a different model. Imagine a kitchen with a row of sand clocks, sitting under the row of spices, ranging from a garlicky 15 seconds up to a turkey 4 hours. Hinged, perhaps, to invert upon a touch. Quaint, but inefficient. There is a better prosthesis available: one of those wind-up timers that we use to alert us when to change the pigeons. It contains all of the sand clocks in a single portable device. Buy one of those, and you have bought yourself a pacemaker. It can model any interval up to an hour, and ring a bell as a discriminative stimulus for you to leave the meeting. Those are handy devices in colloquia, even if you do not have any pigeons running.

What most clocks do, and what all accurate clocks do, is make and measure paces. Water clocks and candles do not have pacemakers; they move a measured quantity of matter at a relatively constant rate, and signal when it is all gone. They are inaccurate. Mechanical and electrical clocks meter the number of cycles of an oscillator. They measure them by mapping them to a spatial interval, such as the circuit of the clock face; or they count them by incrementing an integer, as do digital watches. The pacemaker is an oscillator, as is anything that rotates or revolves or swings—anything, that is, that can be modeled as an oscillator.

To use the pacemaker, you need to count or measure its output. The finer the scale and the more accurate the period, the more model intervals you can generate. The earth is a pacemaker. It rotates around its axis. Its course can be subdivided with a sundial to

time the day. Its rotations can be counted on a calendar to time the year. Calendar reforms have made this a fine art, primarily by knowing when to reset the year counter to maintain a correspondence between the months and the seasons of the year, the latter generated by an independent pacemaker—earth's revolutions around the sun.

Counting them. Our location in time is recorded as years (from an arbitrary origin) plus months plus days plus hours, and so forth. Whenever a second elapses, our pacemaker increments a seconds counter and resets to zero; when that reaches 60 it increments a minutes counter and resets to zero; when that reaches 12 it increments a p.m. counter and resets to zero. And so on. This is a counting cascade, with the output from faster registers being passed to slower registers.

The Good News

Pacemaker-counter systems are a great invention because the more oscillations you count, the less important is accuracy in the pacemaker. The subdivision of the period of the sun by a sundial is inaccurate. Better to multiply a fast oscillator than divide a slow one. In a pacemaker-counter system, even the worst pacemaker—one whose n th tick gives absolutely no information about when its next tick will come—can be part of an accurate clock, if the pacemaker is fast enough. The Poisson emitter is the mathematical model of the worst possible pacemaker. It has a constant probability of ticking at any instant in time, and this probability is independent of when the last tick occurred. If you count enough of these ticks, however, you can be as accurate as desired. In particular, your relative inaccuracy—the standard deviation divided by the mean—decreases as $1/\sqrt{n}$.

If I ask you to say when 10 seconds have elapsed without use of a watch, you are likely to count “one-thous-and-one, one-thous-and-two, . . .” You are counting quarter seconds, which you produce by rhythmic motions of mouth and tongue. This is more accurate than if you just guessed. You are timing by counting: You are using a pacemaker-counter system. Skinner forgive, for you have just gainfully employed one instantiation of the insidious hypothetical construct, *pacemaker*. If you refuse to play my game again and use

your watch, well, by this point you should know what I would say.

The Bad News

The Poisson system—the worst possible constant-rate pacemaker and the best possible counter—is mathematically a relatively simple model. That is why it is used. It is a subset of recurrent processes; you can replace the Poisson emitter with a more accurate oscillator, and the relative accuracy will improve further. But it cannot get worse. That is also the bad news, because when animals time intervals, their accuracy does not increase as the inverse square root of the interval to be timed. In fact their relative accuracy is relatively constant. This is Weber's law. It is also a keystone of scalar expectancy theory (SET). So, whereas the Poisson system is a good way to model time, it is a bad way to model animals that time. To do that, the clock has to be detuned.

Detuning the Clock

Where can error be added? The pacemaker is already almost as bad as it can get. One way to make it worse is to add error to the average rate of the pacemaker, a source of variance invoked by SET. Another is to add error to the counter. SET does that by having the storage of the count add an error proportional to the magnitude of the count. The behavioral theory of timing (BeT; Killeen & Fetterman, 1988) does it by having the pacemaker run faster for shorter intervals and slower for longer intervals, as though there were a ceiling on how high the counter could go. (The mechanism for this adjustment is assumed to be the higher levels of arousal typically associated with the higher rates of reinforcement on short intervals.) Machado's (1997) learning to time (LeT), a continuous version of BeT, also lets the speed of the pacemaker vary with the interval to be timed. The connectionist model of timing (Church & Broadbent, 1991) adds error to a register of counters that is proportional to the magnitude of the numbers they count.

Couching such models in mathematical terms can be useful, because they can then make more precise predictions that are susceptible to test. If they fail those tests, they must be fixed or discarded. For instance, one obscure mathematical model of pacemaker-

counter systems shows that under many different kinds of component error, the optimal period of the pacemaker should be independent of the interval to be timed; it predicts the value of that optimal interval in terms of the variances of the components; and it shows that, under the stated assumptions, the origin of Weber error must reside in the counter (Killeen & Weiss, 1987).

Why maintain a pacemaker-counter model if it must be fixed? Well, it is still a simple model, and a simple model that is accurate is much better than no model at all. All that is required is the replacement of a simplifying assumption (error-free counters) with a more realistic one (error in counting). Many feel that it is best to start with the simplest models and adjust them only when the data require it. This tactic keeps models understandable until some of that simplicity starts to exact a cost. Furthermore, the pacemaker-counter system has some face validity, at least in the realm in which intervals are long enough to be counted. Finally, it has some ecological validity: All of the best clocks do it that way.

How MTS Works

How do Staddon and Higa get worse than Poisson timing? They have a cascade of different interval timers, each of different (increasing) natural periods (much like the egg timers, arranged in order). The response output is the positive difference between the current input and the "memory" (V_i ; their Equation 15). The effect of a single input decays proportionally with its magnitude (Equation 16), which entails an exponential decay of its strength. One such unit feeds into (cascades into) the next, with the next having a slower rate constant (a_i).

The cascade acts as a low-pass filter. If temporal stimuli are coming at high frequencies, the earliest-fastest units in the cascade get saturated, and that blocks the signals from getting through to the slower units downstream. The first-fastest do all the work: It is variance in their level of saturation that accounts for most of the variance in the animal's response. If signals come more slowly, the fast units are not saturated when they recur; their response output to the next level of the cascade is therefore greater (Equation 15), and those slower levels start to pick up some of the burden. This is a good thing, because the accu-

racy of any one level decreases exponentially with its level of saturation. One can see this by noting that the change in time corresponding to a small change (error) in trace strength V grows as $dt/dV = e^{at}/a$. This is much worse than Poisson timing; much too much worse.

So, how do Staddon and Higa move in the other direction, and improve model performance to get better than exponential error? The secret is in the cascade. By passing the ball to the slower units at longer intervals, the system can operate in a closer to optimal range. Furthermore, by passing the ball most forcibly when the units are most accurate (Equation 15 shows that the signal strength decreases as a unit saturates) the bad performance near saturation is not propagated downstream. Their Figure 6 shows that by selecting the rate at which these downstream units are slowed, one can generate a range of Weber functions.

The MTS Counter

The MTS model is a counting cascade, with the output from faster registers being passed to slower registers. A start-time marker must discharge all the units so they can start timing an interval afresh, like a bar that restores all the egg timers to upright. (The authors profitably speculate on empirical consequences of weak start signals.) It must simultaneously initiate a start pulse; because all registers have been cleared, this will get through equally (weighted by the b_i parameters) to all units. For the system to come to equilibrium, a series of input pulses must then be sent that condition (fill) the various reservoirs, layers, or units to their various asymptotic levels.

The state of each unit is summed by Equation 17. Whereas counters are digital all-or-none devices, Equation 17 sums continuous variables for a continuous aggregate memory strength. It is more accurate to call the Staddon-Higa counter an “accumulator.” Same function, slightly different capabilities: If you count things small enough, then counters are not discriminable from accumulators. In the process of accumulation, MTS does not retain information about the individual registers. Those are weighted (by the b parameters) and summed. This is just as well, because animals usually make binary temporal decisions and primarily need to discrimi-

nate “not yet” from “now.” Any more information could be counterproductive.

This is a dynamic system, so it is not surprising that it is path dependent: The effect of a single time interval depends on those which came before, and on the disequilibrium state of the system. To be understood, it really has to be simulated. This is not hard to do. Equations 15 through 18 are simple when you work them through. Just pick a discrete time interval small enough (say, 0.1 s) and input timing signals at longer intervals (say, 10 s) and play with the parameters. Try it. That is the only way you will understand it. If it doesn't work, you can blame the authors. You might even get a publication out of it.

The MTS Pacemaker

MTS does not have a pacemaker. It has a series of continuous modelettes of temporal intervals (each of the layers) whose accuracy decreases exponentially with the interval to be timed. It is as though there were as many pacemakers as layers, each with a different resonance. This aspect is similar to the spectral approach of Grossberg and Schmajuk (1989) and the delay lines of Moore and Desmond (1992). The layers hum as an increasing function of the proximity of the stimulus to their resonant frequency, and the accumulator sums the chorus. They achieve their ability to predict Weber functions, as well as functions not so Weberian, from the tuning of this aeolian harp. Octaves (Figure 6, $\lambda = 2$) do quite well. Deviations from a perfectly horizontal line are also interesting, as they may come closer to real data than a perfect Weber's law.

Do We Keep It or Go for the Hose?

The parents tell us that this kid is going to grow up to be president, after having served in the space corps and composing some tone poems of local reknown. Furthermore, the kid will be inexpensive to keep, not needing a pacemaker, which the village can now sell to a used clock shop up the coast. Best of all, it is unsullied, so far. So, what do we think?

Nice thing is, a village needs more than one model. And this one is quite different; diversity is important, which we knew even before our bureaucrats told us so, and remains true despite them. But more than models, a village needs teams: an MTS team

and an SET team and an LeT team. The MTS coaches have scored some points on the SET coaches, but we have yet to see how well they can defend their own goal. It is worth nurturing this little model for a while; see if it lives up to its parents' expectations. It has the potential, but more players are desperately needed for a meaningful competition.

Modeling Modeling

To think about models, we need models of the process; restrictions on form and criteria for success and handicapping for assumptions. We have yet to evolve consensus on this most central preoccupation of scientists. I think of models as structures that exist in a domain different than the subject, and are constructed according to design principles that are dictated by theories. Parsimony (simplicity), power (the range of phenomena that can be covered), and accuracy (not too many mispredictions or omissions) are all important. But so also is excitement, the sense that models are comprehensible and testable, and that each of us can use them to score goals of our own. Science has many game-like aspects; if we model it as a game, with referees, teams, rules, matches, prizes, and penalties, we can come to some important conclusions: One match is not a season; a loss is less shameful than refusing to play; clear rules and agreed-upon goals are important; unbiased referees are essential. Perhaps most important, it is not a zero-sum game. Models

that can take the field are to everyone's advantage; contests such as this will not only improve all teams, but they will draw participants from other domains to refresh the experimental analysis of behavior.

Now then, anyone for an expansion team?



REFERENCES

- Church, R. M., & Broadbent, H. A. (1991). A connectionist model of timing. In M. L. Commons, S. Grossberg, & J. E. R. Staddon (Eds.), *Neural network models of conditioning and action* (pp. 225–240). Hillsdale, NJ: Erlbaum.
- Ellson, D. G. (1939). The concept of reflex reserve. *Psychological Review*, *46*, 566–575.
- Grossberg, S., & Schmajuk, N. A. (1989). Neural dynamics of adaptive timing and temporal discrimination during associative learning. *Neural Networks*, *2*, 79–102.
- Killeen, P. R. (1988). The reflex reserve. *Journal of the Experimental Analysis of Behavior*, *50*, 319–331.
- Killeen, P. R., & Fetterman, J. G. (1988). A behavioral theory of timing. *Psychological Review*, *95*, 274–295.
- Killeen, P. R., & Weiss, N. (1987). Optimal timing and the Weber function. *Psychological Review*, *94*, 455–468.
- Machado, A. (1997). Learning the temporal dynamics of behavior. *Psychological Review*, *104*, 241–265.
- Miller, A. I. (1984). *Imagery in scientific thought*. Boston: Birkhäuser.
- Moore, J. W., & Desmond, J. E. (1992). A cerebellar neural network implementation of a temporally adaptive conditioned response. In I. Gormezano & E. A. Wasserman (Eds.), *Learning and memory: The behavioral and biological substrates* (pp. 347–368). Hillsdale, NJ: Erlbaum.
- Skinner, B. F. (1938). *The behavior of organisms*. New York: Appleton-Century-Crofts.
- Staddon, J., & Higa, J. (1997). Multiple time scales in simple habituation. *Psychological Review*, *103*, 720–733.