Endnotes

ABRAM AMSEL
University of Texas, Austin, Texas

The reception of the book by my five distinguished colleagues was generally positive, for which I am grateful. There was, however, very little reference, positive or negative, to the specific contents of the book. The commentaries were not in any detailed sense "book reviews," nor, perhaps, were they meant to be. It was more that the commentators, each having done some work on reward-schedule effects—on alternative or related treatments of reward and nonreward (reinforcement and nonreinforcement)—chose to compare and contrast the principles that guide their research with those that guide mine. As I see it, my task in responding is not to deal with each and every point raised by my five colleagues, but to select from among them issues that have some general importance or that raise interesting questions. I would add that, had the book we are discussing been more a textbook on reward-schedule effects and less a monograph based on a particular theoretical approach to these effects, much of the material in the commentaries of my five colleagues would have been included.

Bitterman

Instrumental conditioning. I do not regard my analysis of instrumental conditioning as being "unabashedly Hullian," and not because I am not happy to acknowledge Hull's influence. I have not anywhere in the book (that I can find) specifically adopted the "S-R reinforcement principle" of instrumental learning (Bitterman, 1994, p. 297), although I don't disavow it. The point is that the "reinforcement" part of this principle is not, strictly speaking, required for my own analysis. The interesting thing in this regard, however, is that, on the basis of recent neurobiological work, a strong reinforcement principle appears to be gaining strength. I refer, for example, to recent work from Stein's laboratory (e.g., Stein, Xue, & Beluzzi, 1993) on the operant conditioning of spiking and bursting of hippocampal pyramidal neurons reinforced with a drug, cocaine, or a neurotransmitter, dopamine. This would seem to be a case not only of the operation of reinforcement, but of its operation at a very basic level. And, in a way that was probably never intended, Hebb's famous "rule," which plays such a large part in current work in the neurobiology of learning and memory, strikes me as a reinforcement mechanism at the level of synaptic transmission. The rule states, in paraphrase, that presynaptic and postsynaptic cells must fire conjointly to increase the efficiency of the firing of the latter by the former, obviously via some kind of feedback

from the post- to the presynaptic cell or from the postsynaptic cell to the synapse. Hebb's rule is also the basis for the concept of "back-propagation" in neural network theorizing, referring to an increment on the synaptic connection, the increment, again, being a product of the postsynaptic firing by the presynaptic neuron. Again, this is an approach very much like Hull's reinforcement principle (see Hintzman, 1993). In cases more removed from reinforcement, those in which responses, in Bitterman's terms, are "learned about" (remembered and anticipated), there is little or no substantive difference between Tolman's expectancy position in 1932 and Hull's fractional anticipatory goal response in 1931.

Bitterman, in referring to my stimulus-response (S-R) account of classical conditioning, points to the awkwardness of putting dashed arrows (associations) between the conditioned stimulus (CS) and the conditioned response (CR) rather than between the CS and unconditioned stimulus (UCS), implying that the connection between UCS and unconditioned response (UCR) becomes the template for the conditioned connection. I do not, as Bitterman suggests, take S-R associations to be more "objective" than S-S associations. Rather than objectivity, this is, in my case, a straightforward theoretical necessity: S-R-S (as in $S-r_f-s_f$) sequences are central to mediational mechanisms in frustration theory, because r-s feedback is at the heart of its explanations and predictions. In my classroom teaching, I always put the dashed lines both ways—between CS and CR and between CS and UCS—and take the weaker position that it may be either one or both, and that, ultimately, neurophysiological and neurochemical research will settle it.

Long-term and short-term memory. This long section of Bitterman's commentary is something he and I have discussed informally. But before I take on the central point, let me respond briefly to others. The statement that I give little weight to sensory aftereffects or short-term memories neglects our extensive recent work on "memory-based" learning—particularly the ontogenetic analysis in chapters 7 and 8, which suggests age-related growth in the duration of such sensory aftereffects or short-term memories. I surely do agree that "carry-over ... does play a broad and important role in massed trials." I point out in the book that, in this regard, there is no disagreement with Capaldi. The only disagreement at this point is with Capaldi's (and perhaps now Bitterman's) position that carry-over stimulation lasts indefinitely. This our research does not support, neither in the

adult nor (certainly) in the developing infant rat. It takes about 240 trials of single-alternation training to show that it lasts about 60 sec in the weanling rat. Is the carry-over sensory or mnemonic or both? This is a difficult distinction for me, at least operationally (see Capaldi), and I would be inclined to let neurophysiology rule on this one as well.

Bitterman deals with the permanence of associations as follows: "The conception of permanent association is one Hullan legacy that does not seem worth preserving" (p. 300). The question, as I would put it, is this: Does inhibition (in its several forms) subtract from excitation? (On this issue I address not only Bitterman's position but also the position on permanence of conditioning taken in the model of Rescorla and Wagner and, consequently, in the Daly model [DMOD].) The trouble is that inhibition has always meant too many things: Pavlovian nonreinforcement; suppression and avoidance based on electric shock (with respect to which it would have been better had psychologists never heard about it); and the nonreinforcement of instrumental appetitive responses, just to mention the most obvious ones. What Bitterman particularly doesn't like is Hull's distinction between conditioned inhibition (S^I_R) and spontaneously decaying reactive inhibition (I_R). The point for present purposes is that there is in the subject index of my book only one reference to I_R and S^I_R (it is on p. 41), and this one only to make the point that, in an informal paper in 1951, I proposed to add a third "inhibitory" factor (r_F), which for a long time now I have not regarded as inhibition, but as competing excitation.

My point has always been that if an appetitive Pavlovian expectancy is excitatory, then an aversive Pavlovian expectancy is *excitatory, not inhibitory*, whether the aversion is from nonreinforcement or from electric shock: it is simply a counteracting excitatory tendency. Fear is not inhibitory; it is excitatory, even though it suppresses responding, as in conditioned suppression. Approach-avoidance conflict is not the interaction of excitatory and inhibitory tendencies; both are excitatory. We also have to be able to account, even in the simplest Pavlovian case, for spontaneous recovery and, in other cases, for our old friend, reminiscence. It's possible that an inhibition concept works differently at some phylogenetic levels than others; however, I firmly believe that, in rats and people, most of the uses of the term "inhibition" are descriptive and do not in fact represent a direct, subtractive mechanism from the hypothetical excitatory state or, in the case of Woodward and Bitterman, the reduction of inhibition by reinforcement. That is why, I think, Hull had total inhibitory potential, $I_R = (S^I_R + I_R)$, subtract from S^E_R , not from S^H_R .

There is much work nowadays, in laboratories of neurobiology of learning and memory, not only on long-term potentiation (LTP), but also on long-term depression (LTD). There appear to be two kinds of LTD. One, called homosynaptic, is like reactive inhibition, in that it dissipates rapidly in time; the other is associative and lasts much longer, as long as LTP. The very interesting thing,

according to a speaker I recently heard on the subject, is that antagonists to NMDA receptor function have the same effect on LTP and on *associative* LTD; they block both of them, meaning that the NMDA receptor may be involved in the same way in both. So, even at the neurochemical level, there is evidence for counteracting, competitive excitatory tendencies rather than inhibitory reductions in or subtractions from associative strength.

Evolutionary divergence in learning. The closest we come in experimental work to evolutionary divergence is the work from Bitterman's laboratories, and this comparative approach was surely the model for our experiments on the ontogeny of the nonparadoxical and paradoxical reward-schedule effects. A special problem raised by Bitterman in our work with very young rats is the relatively early appearance of the partial reinforcement extinction effect (PREE). As I point out in the appropriate part of the book, I would not have supposed it would precede effects such as successive negative contrast (SNC). Indeed, in a concluding discussion (chapter 9), I try to rationalize the unexpected order of emergence of these effects in what is essentially a modification of frustration theory. The suggestion by Bitterman is that the early appearance of the PREE may depend on training with relatively massed trials. He might be right, but we do have experiments on the PREE in 14-day-old rats in which the intertrial interval (ITI) is 10–12 min (Letz, Burdette, Gregg, Kittrell, & Amsel, 1978). Bitterman's statement that massed-trial SNC can be accounted for in terms of primary frustration (R_F) alone is puzzling: in massed trials, primary frustration produces FE-like effects, not SNC. I don't know how to explain SNC without conditioned (anticipatory) frustration (r_F), and I don't see how counterconditioning of r_F – s_F to approach explains it either, certainly not in the rat work. In Bitterman's work with the honeybee, to which he alludes, there appears to be a way to account for SNC as well as simultaneous negative contrast in terms of primary frustration alone. As he suggests, this is the kind of difference between the ontogenetic and phyletic work that may point up unanswered questions about evolutionary divergence.

Capaldi

Memory versus expectancy. Capaldi refers to memory-expectancy theory and characterizes frustration theory as being one theory of this kind and a recent revision of his own sequential theory as being another (Capaldi, 1994). In this paper, Capaldi treats memory of a reward (or a nonreward, presumably) as a CS, the CR being the expectation of a future reward (or nonreward), and suggests that "stimuli that elicit expectancies, as well as the expectancies themselves, may change progressively over a series of learning trials" (p. 303). This is, in my experience, a new theoretical strategy in sequential theory, and it is close to some of our own recent thinking, the difference being only in the relative duration of the carried-over memory. Our developmental research on patterned single-alternation (PA) learning shows that

these memories can last (the ITIs can grow) from about 8 sec in 11-day-old rats to about 60 sec in weanlings. Working exclusively with adult rats, Capaldi has put virtually no limit on the duration of the effective ITI, and it is a long-term disagreement between his laboratory and ours that we have not found that effective carried-over stimulation in adults (memories of R and N) has an indefinite duration. Capaldi points out that he has "attached as much importance to the memories of reward events as Amsel has to expectancies of reward events." This statement appears to refer to Capaldi's earlier position in which the entire emphasis was on the carry-over stimulation from R and N, which he came to regard as memories. In frustration theory, expectancies (r_R , r_F and their feedback stimuli s_R and s_F) always bore an explanatory emphasis that was absent in Capaldi's earlier theorizing, but that seem now to be a part of his memory-expectancy theory. Capaldi indicates that these classically conditioned expectancies are "the workhorses ... of frustration theory"; apparently, they are now a part of sequential theory as well.

In an experiment published some years ago (Lobaugh, Greene, Grant, Nick, & Amsel, 1989), we proposed an idea to account for the learning of PA at increasing ITIs as infant rats grow older. It is similar to Capaldi's new memory-expectancy approach. The proposal was that the carried-over cues from R and N (the memory system), which can support PA learning in very young rats only at relatively short ITIs (at longer ITIs, in older—e.g., weanling—animals), the memory system engages the conditioned anticipation or expectancy system to construct a more complex discriminative stimulus, permitting learning at longer ITIs.

It is the carry-over stimulation arising out of *primary* rewards and nonrewards that I take to be not long-lasting, even though, as long as they do last and are not interfered with, they can function as "memories," like other short- or intermediate-term memories. (The same can be said of the feedback stimulation from anticipatory rewards and frustrations, e.g., r_F-s_F .) I agree that both intratrial and intertrial mechanisms are required if we are to understand instrumental learning: the difference is only in the effective duration of the ITI in the latter. It is by now a part of long-term history that they have been longer lasting in Capaldi's laboratory than in mine. In Capaldi's (1994) new memory-expectancy approach, however, the importance of this discrepancy appears to be minimized. Let me emphasize again that, as I understand it, Capaldi's recent suggestion is that memories of prior reward and nonreward (S^R and S^N in his notation) "may serve as CSs for the expectancies of reward (E_R) and of nonreward (E_N)," and that both memories and expectancies "may come to elicit the instrumental reaction" (p. 304). Here, then, our theoretical constructs have become more similar, even though the phenomena they predict and explain are mostly very different. A major remaining difference is that Capaldi has become more "cognitive" in his approach and refers to S-S associations, while my position is still S-R.

There are, as we have been told repeatedly in recent years, many ways to classify memory: implicit versus explicit, dispositional versus representational, and so on, and in these classifications no great distinction can be made between memory and expectancy, when both have the status and the function of classically conditioned responses. However, when we are dealing with reward-schedule effects, the distinction is clear and it makes sense: the carried-over stimulus traces from the previous trial outcome in learning based on reward schedules can be regarded as a form of working or short-term or intermediate-term memory (explicit and representational in any case), whereas the classical or Pavlovian conditioning of expectancies is more implicit or non-representational. Mishkin and Petri's (1984) memory versus habit distinction makes the point perfectly for me, as I point out in the précis. In their case, the habit is one of expectancy.

Daly and Daly

The experimental work of Helen Daly, even including her recent work on the toxic effects of Lake Ontario salmon, has contributed more to the validation of frustration theory than any other body of research, with the possible exception of the work in my own laboratory. If anyone was in a position to develop a mathematical model of the theory it was she, in collaboration with John Daly, her mathematician husband, and they have done this admirably, using as a basis for DMOD the mathematical form and assumptions of the Rescorla-Wagner model of Pavlovian conditioning.

Rather than being a commentary on the book, I regard the Dalys' (Daly & Daly, 1994) contribution as an addendum to it. It could easily have been a separate chapter in which I reported their work, except that it was already so well known that I thought it best just to refer to it. And I couldn't possibly have come up with the 60 phenomena from frustration theory that they say they have modeled.

DMOD, being a combination of the Rescorla-Wagner model of Pavlovian conditioning and frustration theory, addresses instrumental learning and the paradoxical reward-schedule effects, for which Rescorla-Wagner, itself, was not intended to account. Helen Daly's recent application of DMOD to our extensive research on the ontogeny of the paradoxical reward-schedule effects was, as she points out, accomplished during a 4-month visit to my laboratory, during which she examined meticulously all the relevant data that led her to the conclusion that the order of first appearance of those effects in ontogeny can be attributed to a single factor—the growth of anticipatory frustration with increasing age.

The premise of DMOD that is alien to frustration theory is the one I reacted to in Bitterman's commentary: negative V values in this case. This is, however, a problem for me only as long as V continues to be defined as associative strength rather than as excitatory potential, in Hull's sense. As a matter of evidence as well as strategy, I regard decrements in associative strength as un-

likely, the reasons being those I addressed in my lengthier reaction to Bitterman's commentary on this issue. And, on this issue and others, mathematical models, like computer models and nerve-network models—theories of all kinds, even frustration theory—can do a great job of organizing diverse findings in a given area of investigation, and, in the Dalys' case, at least, they may lead to interesting and testable predictions; but nobody ever said they necessarily represent reality. Beyond that, the experimental work of Helen Daly and the models of Helen and John are for me, as the old saying goes, "pure gold."

Hulse

Representational versus dispositional learning. Hulse (1994) refers in several places to his interest in cognitive, representational learning and memory, and he contrasts it with my interest in dispositional learning and memory, implying that his interest (and not mine) addresses the "signaling properties of reinforcement and nonreinforcement" (p. 318). I think this overstates the case. Frustration theory depends as much on signaling as on emotional and motivational properties. Obviously, there is no way an S-R theory can get along without Ss, large and small, external and internal, real and hypothetical, and, in Hulse's terms, all of these have "signaling" properties. They enter into most of the predictions in frustration theory, along with energizing, generalized-drive contributions that define the primary frustration effect, and, for example, the start and run segments of the Goodrich effect (see my commentary on Killeen). Hulse prefers to put his signaling properties of R and N into "a cognitive framework," and he suggests that calling his position cognitive might give me justification for "ignoring or dismissing ... [his] thoughts altogether" (p. 318). Clearly, a word does not make a theory, or even a hypothesis. There is no great difference theoretically or semantically between a signal and a stimulus—nothing that cannot in principle (if necessary) be incorporated into frustration theory. It may have given readers a misleading impression to present as Table 1 a set of distinctions taken from Squire (1987), all of which relate to the difference between explicit and implicit learning and memory. Hulse asks why I ignore those on the explicit (left-hand) side of the table. The simple answer is that I do not ignore the left-hand side of the table. I presented it to establish the fact that my particular approach is more a theory of the right-hand side than of the left. The book is indeed mainly about implicit, dispositional rather than explicit, representational learning. My defense is that, simply by highlighting the distinction, I am cleared of any accusation of neglecting it.

Sequential complexity and cognition. However, despite leaning to the right-hand side of the table, and despite the absence in my book of results such as Professor Hulse's and his colleagues' on sequential complexity, such results are not in principle impossible for frustration theory to deal with. This is also true of a large part of Capaldi's important work, although, as I have indi-

cated, some aspects of the conceptualization of Capaldi's work and ours seem already to be closer than they once were. I wish now to make my own position on the sequential, cognitive abilities of the rat clear, if it isn't already. I am *not* interested in the rat's cognitive abilities; in fact, I am not, as most investigators are not, interested in the rat, except as a model system of non-cognitive, dispositional function, the kind of function that exists in humans—nonverbally, implicitly rather than explicitly. If one wants to study cognitive function, the rat may not be the worst of all possible subjects, but it is surely not even close to being the best. This is why I find the rat a reasonable subject for the study of the development of dispositions; of reward-schedule effects that do not depend primarily on particular sequences; of transfer of persistence effects over days and months with interpolated phases of continuous reinforcement. This is an enterprise clearly quite different from Hulse's and, to the extent that Capaldi has become more cognitive in his approach, from Capaldi's, both of whom have done beautiful and important work in domains quite different from mine.

Thinking ahead coolly (rats). For example, Hulse describes an experiment with increasing and decreasing sequences of amounts of reward (0, 1, 3, 7, and 14 pellets), and suggests that this kind of experiment, in which rats can apparently anticipate (at least, react to) these sequences, can be characterized, in human terms, as "thinking ahead." This, in a sense, is the hypothesis confirmed in these experiments. On the other hand, rats given food in random sequences "opt" (Hulse's word) to run quickly on all trials (as they do in a variable-magnitude-of-reward experiment or under partial reinforcement [PRF] conditions). Hulse is not unfriendly to motivational-emotional concepts: he observes that after a number of nonreinforced runs there are signs of frustration—biting the experimenter, urination, and so forth.

"Thinking ahead" is not something that characterizes frustration theory and the experiments to which it gives rise. In fact, the theory that I have developed and some experiments that we pursued early on involved PRF training at long ITIs—sometimes very long ones (24 or 72 h)—followed by long "vacations," and then by continuous reinforcement (CRF) training and extinction. The idea in these experiments was that in the rat "thinking ahead" would be unlikely.

In the developmental work we have used the single-alternation experiment extensively and have shown that this kind of "memory-based learning" can be accomplished at ITIs from about 8 sec to 60 sec as the animal grows from 11 to about 20 days of age. We have not attempted, either with infant or with adult rats, to do in a runway the kind of experiment Hulse's group has conducted in a T-maze. Would a rat respond "appropriately" to a sequence of reward sizes in a runway? And if it did, would emotional-motivational (frustrative) factors be required, or even involved, in the explanation? The only approach to an answer to this question is the Fountain and Hulse (1981) experiment, and the answer appears to

be, yes, probably. Could frustration theory accommodate the idea of a rat “hungry and under pressure, who nevertheless comes to survey and control the discrimination situation coolly” (Hulse, 1994, p. 321)?—only in the sense that the guiding stimulation from the primary frustration drive stimulus, and particularly the feedback stimulation from anticipatory reward (r_R-s_R) and frustration (r_F-s_F), operate in the absence of the energizing motivational properties with which they normally interact, not the normal situation in the rat “hungry and under pressure,” but perhaps in the rat tranquilized. I have referred on several occasions, as well as in the book, to Terrace’s experiments on “errorless” discrimination, involving fading and shaping or the administration of a major tranquilizer. Terrace (see Amsel, 1992a, p. 134) showed that administering a tranquilizer greatly increased “error” responses, increased responding to S—after perfect discrimination had been learned in the normal way. However, when a discrimination was managed by fading and shaping methods, so that few S—responses (and nonrewards) occurred, the drug, chlorpromazine, produced few “error” responses to S—. My guess is that factors such as these would have to be involved for the rat “to survey and control the discrimination situation coolly.”

Killeen

Frustration as an acceptable scientific term. Killeen (1994) makes an interesting historical point about, as he puts it, “establishing the scientific credentials of frustration.” He gives as an example B. F. Skinner’s position that “frustration . . . seemed more of a nuisance variable than a key motivational force; he believed that with repeated conditioning and extinction, such responses would ‘adapt out’” (p. 323). I have lived with this attitude of Skinnerians for many years, not quite understanding it and trying, without total success, to ignore it. It gives me some satisfaction that Killeen, a close intellectual relative of Skinner, has perhaps found something more in the concept of frustration than his mentor, and spent some time on a sabbatical in Austin, during which he contemplated the ballistics of the Goodrich effect.

Starting in the mid-seventies, to Skinner’s obvious dismay, several Skinnerians began to see cognitive potential in the pigeon, and to use terms and concepts that were not in the “dictionary” of the experimental analysis of behavior, terms that in older times would have been regarded as even “spookier” than frustration. Their pigeons turned out to have not only cognitions, and conscious and unconscious representations, but even self-awareness and metaknowledge (see Amsel, 1989, pp. 70–73).

At the same time, according to Killeen’s narrative, the “uphill battle” to establish the “scientific credentials” of frustration was still going on. In this, Killeen is absolutely correct—and this has always puzzled me. It wasn’t only the older style Skinnerians whom we faced on that battlefield, though they were very much in the front lines. Fear, anxiety, conflict—with these terms

there was relatively little difficulty; even anger was acceptable, but not frustration. In their commentary, the Dalys tell a relevant story, which I have borrowed a time or two, about a Skinnerian colleague who learned what frustration meant when a grant of his had a high priority score and was not funded. Helen Daly herself, surely one of my closest intellectual and personal colleagues, now avoids the term *frustration*, and, like Jim McGaugh and his group at Irvine, refers instead to “aversive non-reward,” presumably because they regard “aversive non-reward” to be a relatively more acceptable part of the language than “frustration”; the story is that study sections react better to it. (That’s like saying study sections react negatively to the word *chaos*.)

All of this reminds me of a story told to me by a graduate student after a colloquium visit to Indiana many years ago. She was driving me to the airport and screwed up her courage to tell me the following: She had been running an experiment for a professor—rats in a runway on a PRF schedule. The experiment, as I recall, was conducted to test a prediction from statistical learning theory. She noticed that at an intermediate stage of the experiment, particularly on nonrewarded trials, the rats became agitated and urinated and defecated in the goal box. She asked her advisor how statistical learning theory could account for things like that. The reply was, “If you’re that impressed with it, take a picture of it and send it to Amsel.” Clearly this is not a story about Skinnerians, but it does touch on Peter Killeen’s commentary, because it represents an attitude that prevailed at a time when both mathematical-theoretical and Skinnerian-inductive influences were very strong in psychology (where are they now?), both beginning to overpower the influences of Hull and Tolman. Aside from the Dalys in their DMOD, neither math modelers nor, as we have seen, the Skinnerians could find room for terms connoting motivation or emotion, and particularly, for some reason, frustration, no matter how well defined.

The partial reinforcement acquisition effect (PRAE). This behavioral effect, measured in a runway, was established at Iowa in experiments by Goodrich, Haggard, and Wagner. It is a highly replicable effect in rats, and our developmental timetable shows that it is observable for the first time in preweanling rats (18–20 days of age). The Killeen and Amsel mathematical treatment of the PRAE (almost entirely Killeen’s work) is a kind of “icing-on-the-cake” thing. Examining the existing relevant data in terms of acceleration rather than speed, the paper Killeen refers to simply strengthens the interpretation of a finding that is at the heart of the explanatory power of frustration theory, partly because it is the kind of result that appears to be outside of the explanatory field of competing theories of reward-schedule effects (e.g. Capaldi, Hulse) perhaps because they do not include a generalized-drive component. The successful theory must account for how at asymptote, within the space of 1.5 to 2.0 sec, a rat, running in a three- or a five-foot runway under PRF relative to CRF conditions, runs and accelerates faster in the first and second thirds of the

runway and slower in the last third of the runway. The frustration theory explanation, first offered by Wagner (and as the book explains in detail), is in terms of the arousal and suppression features of anticipatory frustration that operate under PRF but not CRF conditions in the first two thirds and last third of the runway, respectively. And a consequence of this combination of arousal and suppression is the PREE, the greater resistance to extinction (persistence) in rats in the PRF relative to the CRF condition. It occurred to me at the time that, in selecting the PRAE for this kind of "microscopic" examination, Peter Killeen was pointing to its special importance as one among the many reward-schedule effects that are organized by frustration theory (see Amsel, 1992a, Appendix, and 1992b).

REFERENCES

- AMSEL, A. (1989). *Behaviorism, neobehaviorism, and cognitivism in learning theory: Historical and contemporary perspectives*. Hillsdale, NJ: Erlbaum.
- AMSEL, A. (1992a). *Frustration theory: An analysis of dispositional learning and memory*. Cambridge: Cambridge University Press.
- AMSEL, A. (1992b). Frustration theory—many years later. *Psychological Bulletin*, **112**, 396-399.
- BITTERMAN, M. E. (1994). Amsel's analysis of reward-schedule effects. *Psychonomic Bulletin & Review*, **1**, 297-302.
- CAPALDI, E. J. (1994). The relation between memory and expectancy as revealed by percentage and sequence of reward investigations. *Psychonomic Bulletin & Review*, **1**, 303-310.
- DALY, H. B., & DALY, J. T. (1994). Persistence and the importance of nonreward: Some applications of frustration theory and DMOD. *Psychonomic Bulletin & Review*, **1**, 311-317.
- FOUNTAIN, S. B., & HULSE, S. H. (1981). Extrapolation of serial stimulus patterns by rats. *Animal Learning & Behavior*, **9**, 381-384.
- HINTZMAN, D. L. (1993). Twenty-five years of learning and memory: Was the cognitive revolution a mistake? In D. E. Meyer & S. Kornblum, *Attention and performance XIV* (pp. 359-391). Cambridge, MA: MIT Press.
- HULSE, S. H. (1994). The other side of the coin: Cognitive properties of nonreward. *Psychonomic Bulletin & Review*, **1**, 318-322.
- KILLEEN, P. R. (1994). Frustration: Theory and practice. *Psychonomic Bulletin & Review*, **1**, 323-326.
- LETZ, R., BURDETTE, D. R., GREGG, G., KITTRELL, E. M. W., & AMSEL, A. (1978). Evidence for a transitional period for the development of persistence in infant rats. *Journal of Comparative & Physiological Psychology*, **92**, 856-866.
- LOBAUGH, N. J., GREENE, P. L., GRANT, M., NICK, T., & AMSEL, A. (1989). Patterned (single) alternation in infant rats after combined or separate lesions on hippocampus and amygdala. *Behavioral Neuroscience*, **103**, 1159-1167.
- MISHKIN, M., & PETRI, H. L. (1984). Memories and habits: Some implications for the analysis of learning and retention. In L. R. Squire & N. Butters (Eds.), *Neuropsychology of memory* (pp. 287-296). New York: Guilford.
- SQUIRE, L. R. (1987). *Memory and brain*. New York: Oxford University Press.
- STEIN, L., XUE, B. G., & BELLUZZI, J. D. (1993). A cellular analogue of operant conditioning. *Journal of the Experimental Analysis of Behavior*, **60**, 41-53.

(Manuscript received May 19, 1994;
accepted for publication May 26, 1994.)