

The partial-reinforcement extinction effect and the contingent-sampling hypothesis

Guy Hochman · Ido Erev

Published online: 18 April 2013
© Psychonomic Society, Inc. 2013

Abstract The *partial-reinforcement extinction effect* (PREE) implies that learning under partial reinforcements is more robust than learning under full reinforcements. While the advantages of partial reinforcements have been well-documented in laboratory studies, field research has failed to support this prediction. In the present study, we aimed to clarify this pattern. Experiment 1 showed that partial reinforcements increase the tendency to select the promoted option during extinction; however, this effect is much smaller than the negative effect of partial reinforcements on the tendency to select the promoted option during the training phase. Experiment 2 demonstrated that the overall effect of partial reinforcements varies inversely with the attractiveness of the alternative to the promoted behavior: The overall effect is negative when the alternative is relatively attractive, and positive when the alternative is relatively unattractive. These results can be captured with a contingent-sampling model assuming that people select options that provided the best payoff in similar past experiences. The best fit was obtained under the assumption that similarity is defined by the sequence of the last four outcomes.

Keywords Choice behavior · Judgment · Decision making · Reinforcement learning

G. Hochman (✉)
Duke University, 2024 W. Main Street, Erwin Mill Bldg., Bay C,
Durham, NC 27705, USA
e-mail: guy.hochman@duke.edu

G. Hochman
Interdisciplinary Center (IDC) Herzliya, Herzliya, Israel

I. Erev
Max Wertheimer Minerva Center for Cognitive Studies,
Technion—Israel Institute of Technology, Haifa, Israel

The partial-reinforcement extinction effect (PREE; Humphreys, 1939) is one of the best examples of a basic behavioral phenomenon, detected in the laboratory, with potentially important practical implications. As most introductory psychology textbooks explain, the PREE refers to the fact that learned behavior is more robust to extinction when not all responses are reinforced (partial schedules) than when 100 % of responses are reinforced in training (full schedule; see, e.g., Atkinson, Atkinson, Smith, Bem, & Nolen-Hoeksema, 1995; Baron & Kalsher, 2000). For example, Atkinson et al. stated that partial reinforcements facilitate higher performance rates, since the probability that individuals will continue responding in the absence of reinforcements is much higher under partial schedules than under full schedules.

Unfortunately, however, many empirical studies fail to support this textbook assertion. Most early demonstrations of the PREE used between-subjects laboratory designs (e.g., Grosslight & Child, 1947; Mowrer & Jones, 1945; Pavlik & Flora, 1993) to show that under partial-reinforcement schedules, individuals tend to engage in more responses during extinction than under full schedules. However, most laboratory studies using within-subjects designs (Nevin, 1988; Papini, Thomas, & McVicar, 2002; Svartdal, 2000; but see Exp. 3 of Nevin & Grace, 2005, for an exception) and field research (Latham & Dossett, 1978; Pritchard, Hollenback, & DeLeo, 1980; Yukl, Latham, & Pursell, 1976) have reported that partial reinforcements impair, rather than improve, performance. For example, Yukl et al. found that tree planters were less productive under partial- than under full-reinforcement schedules. This negative effect of partial schedules was even observed when partial reinforcements yielded higher average payoffs.

Nevin (1988, and see Nevin & Grace, 2000) proposed behavioral momentum theory in order to account for the mixed PREE results. According to this account, two effects compete under partial-reinforcement schedules. On the one hand, partial reinforcements have an overall negative effect

on the likelihood of selecting the reinforced alternative, due to a decrease in reinforcement rates. At the same time, however, partial reinforcements have a positive local effect, as they slow extinction due to a generalization decrement, which hinders detection of changes in the reinforcement schedule in some settings (see the related observations in Gershman, Blei, & Niv, 2010). Thus, momentum theory suggests that the apparent inconsistency between basic research and field studies of PREE can be explained by the assertion that classical demonstrations of PREE focused on the positive local effect of partial reinforcement in slowing extinction, whereas field studies document the overall (negative) effect of partial reinforcements.

The main goal of the present analysis was to clarify and extend Nevin and Grace's (2000) explanation of the mixed PREE findings. We relate Nevin and Grace's (2000) assertion to the suggestion that people tend to select the action that has led to the best outcomes in similar situations in the past (see Biele, Erev, & Ert, 2009; Gonzalez, Lerch, & Lebiere, 2003; as well as a related observation by Patalano & Ross, 2007), and elucidate the conditions under which partial reinforcements are likely to be effective and countereffective.

Experiment 1: evaluation of the positive and negative effects of partial reinforcement

In most previous demonstrations of the PREE (e.g., Grant, Hake, & Hornseth, 1951), the expected benefit from the reinforced choice was higher under full than under partial schedules, because the same magnitudes were administered at higher rates. In the present study, we avoided this confound by manipulating the size of the rewards to ensure equal sums of reinforcements under both partial and full schedules. The primary goal of Experiment 1 was to examine whether the overall negative effect of partial reinforcements could be observed, even when this condition did not imply a lower sum of reinforcements.

Method

Participants A group of 24 undergraduates from the Faculty of Industrial Engineering and Management at the Technion served as paid participants in the experiment. They were recruited via signs posted around campus for an experiment in decision making. The sample included 12 males and 12 females (mean age 23.7 years, $SD = 1.88$).

Apparatus and procedure For the experiment we used a clicking paradigm (Erev & Haruvy, 2013), which consisted of two unmarked buttons and an accumulated payoff counter. Each selection of one of the two keys was followed by three immediate events: a presentation of the obtained

payoff (in bold on the selected button for 1 s), a presentation of the foregone payoffs (on the unselected button for 1 s), and a continuous update of the payoff counter (the addition of the obtained payoff to the counter). The exact payoffs were a function of the reinforcement schedule, the phase, and the choice, as explained below.

Participants were instructed to repeatedly choose a button in order to maximize their total earnings. No prior information regarding the payoff distribution was delivered. The study included 200 trials, with a 1-s interval between trials. The task lasted approximately 10 min.

Design The participants were randomly assigned to one of the two reinforcement schedule conditions: full ($n = 11$) and partial ($n = 13$). Each participant faced two phases of 100 trials ("training" and "extinction") under one of the schedules and selected one of the buttons for each trial. A selection of one of the two buttons, referred to as the "demoted option," always led to a payoff of eight points. The payoff from the alternative button, referred to as the "promoted option," depended on the phase and the reinforcement schedule, as follows:

During the first phase of 100 trials (the "training phase"), the promoted option yielded a mean payoff of nine points per selection. The two schedules differed with respect to the payoff variability around this mean. The full schedule involved no variability: Each selection of the promoted option provided a payoff of nine points. In contrast, the partial schedule involved high variability: Choices of the promoted option were rewarded with a payoff of 17 points on 50 % of the trials, and with a payoff of one point on the remaining trials.

The second phase of 100 trials simulated "extinction." During this phase, the promoted option yielded one point per selection. Thus, the promoted option was more attractive than the demoted option during training (mean of nine relative to eight points), but less attractive during extinction (one relative to eight). The left side of Table 1 summarizes the experimental conditions of Experiment 1.

To motivate participants, we provided a monetary incentive. The translation from points to actual payoffs was according to an exchange rate in which 100 points = NIS 1.5 (about 33 cents). This resulted in an average total payoff of NIS 14 (about \$3.2).

Results and discussion

The left side of Fig. 1 presents the observed choice proportions for the promoted option (P -prom) in blocks of 10 trials as a function of the reinforcement schedule and phase. Table 1 presents the mean values. During training, P -prom was higher under the full ($M = .92$, $SD = .15$) than under the partial ($M = .68$, $SD = .19$) schedule. The opposite pattern was observed during extinction: P -prom was lower under the full ($M = .03$, $SD = .01$) than under the partial ($M = .06$, $SD = .03$) schedule.

Table 1 Experiment 1: Observed mean proportions of selecting the promoted option (P -prom), as a function of the phase and the reinforcement schedule. The two rightmost columns present the predictions of the contingent-sampling (CS) and fictitious-play (FP) models

| Schedule | Phase | Demoted | Promoted | Observed | CS ($m = 4$) | FP ($w = .05, \sigma = 4$) |
|----------|------------|---------|----------|-----------|----------------|------------------------------|
| Full | Training | 8 | 9 | .93 (.06) | .98 | .93 |
| | Extinction | | 1 | .03 (.11) | .02 | .02 |
| | Overall | | | .48 (.08) | .50 | .47 |
| Partial | Training | 8 | 17 or 1 | .67 (.15) | .58 | .72 |
| | Extinction | | 1 | .06 (.14) | .03 | .03 |
| | Overall | | | .34 (.11) | .31 | .37 |
| | | | | MSD | .002 | .001 |
| | | | | r | .992* | .997* |

Correlations and mean squared deviations (MSDs) are averaged across participants. Standard deviations appear in parentheses. The r statistic represents the Pearson correlation coefficient between the actual and predicted P -prom values in each block of trials. $p < .001$.


A 2×2 repeated measures analysis of variance (ANOVA) was conducted to test the effects of the phase and reinforcement schedule on P -prom. This analysis revealed significant main effects for both the phase [$F(1, 22) = 604.65, p < .0001, \eta_p^2 = .965$] and the schedule [$F(1, 22) = 10.07, p < .0005, \eta_p^2 = .314$], as well as a significant interaction between the two factors [$F(1, 22) = 22.69, p < .0001, \eta_p^2 = .508$]. These results suggest that the negative effect of partial reinforcements in the training phase and the positive replication of the PREE in extinction were significant.

Over the 200 trials, P -prom was significantly higher under the full ($M = .48, SD = .08$) than under the partial ($M = .36, SD = .11$) schedule, $t(22) = 2.74, p < .01, d = 1.17$. These results support Nevin and Grace's (2000) resolution of the mixed PREE pattern. The observed effect of partial reinforcements was positive at the transition stage, but negative across all trials.

The contingent-sampling hypothesis Previous studies of decisions from experience have highlighted the value of models that assume reliance on small samples of experiences in similar situations. Models of this type capture the conditions that facilitate and impair learning (Camilleri & Newell, 2011; Erev & Barron, 2005; Yechiam & Busemeyer, 2005), and

have won recent choice prediction competitions (Erev, Ert, & Roth, 2010). Our attempt to capture the present results starts with a one-parameter member of this class of models. Specifically, we considered a contingent-sampling model in which similarity was defined on the basis of the relative advantage of the promoted option in the m most recent trials (m , a nonnegative integer, is the model's free parameter). For example, the sequence G–G–L implies that the promoted option yielded a relative loss (L = lower payoff, relative to the demoted option) in the last trial, but a relative gain (G = higher payoff) in the previous two trials, and with $m = 3$, all trials that immediately follow this sequence are “similar.” The model assumes that the decision is made on the basis of comparing the average payoffs from the two options in all past similar trials (and of random choice, before gaining relevant experience).

To clarify this logic, consider the contingent-sampling model with the parameter $m = 2$, and assume that the observed outcomes in the first nine trials of the partial condition provided the sequence L–L–L–G–L–L–G–L–L. That is, the payoff from the promoted option was 17 (a relative gain) in Trials 4 and 7, and one (a relative loss) in the other seven trials. At Trial 10, the agent faces a choice after the sequence L–L. She will therefore recall all of her experiences after identical sequences (Trials 3, 4, and 7), compute the average experience

**Fig. 1** Observed and predicted proportion of selecting the promoted option (P -prom) in blocks of 10 trials as a function of the reinforcement schedule conditions in Experiment 1. Extinction has started at the 11th block

from the promoted option in this set to be $(1 + 17 + 17)/3 = 11.67$, and since the average payoff from the demoted option is only nine, will select the promoted option.

The model's predictions were derived using a computer simulation in which virtual agents, programmed to behave in accordance with the model's assumptions, participated in a virtual replication of the experiment. Two thousand simulations were run with different m values (from 1 to 5) using the SAS program. During the simulation, we recorded the P -prom statistics as in the experiment.

The results revealed that the main experimental patterns (a negative overall effect of partial reinforcement and a small positive effect after the transition) were reproduced with all m values. The best fit (minimal mean squared distance between the observed and predicted P -prom rates) was found with the parameter $m = 4$. The predictions with this parameter are presented in Table 1 and Fig. 1.

Notice that with $m = 4$, the model implies reliance on small samples: Since 16 sequences are possible, the decisions during the 100 training trials will typically be based on six or fewer "similar" experiences. This fact has no effect in the full schedule (in which a sample of one is sufficient to maximize), but it leads to deviations from maximization in the partial schedule (since some samples include more "1" than "17" trials). During extinction, however, after the fourth trial all decisions are made after the sequence L–L–L–L. In the full schedule, participants never experience this sequence during training. Thus, all of their experiences after this sequence lead them to prefer the demoted option (and the first experience occurs in the fifth extinction trial). In contrast, the typical participants in the partial schedule experience this sequence six times during training, and these experiences can lead them to prefer the promoted option in the early trials of the extinction phase.

Inertia, generalization, noise, and bounded memory The present abstraction of the contingent-sampling idea is a simplified variant of the abstraction that won Erev, Ert, and Roth's (2010) choice prediction competition (Chen, Liu, Chen, & Lee, 2011). The assumptions that were excluded were inertia (a tendency to repeat the last choice), generalization (some sensitivity to the mean payoff), a noisy response rule, and bounded memory. A model including these assumptions revealed a slight fit improvement in the present settings but did not change the main, aggregate predictions. Thus, the present analysis does not rule out these assumptions, it only shows that they are not necessary in order to capture the aggregate effect of partial reinforcements documented here.

Fictitious play and reinforcement learning In order to clarify the relationship between the contingent-sampling hypothesis and popular learning models, we also considered a two-parameter smooth fictitious-play (SFP) model (Fudenberg & Levine, 1999). SFP assumes that the

propensity to select option j at trial $t + 1$ after observing the payoff $v(j, t)$ at trial t is

$$Q(j, t + 1) = (1 - w)Q(j, t) + (w)v(j, t),$$

where w is a free weighting parameter and $Q(j, 1) = 0$. The probability of selecting j over k at trial t is

$$P(j, t) = 1 / \left(1 + e^{\sigma[Q(k, t) - Q(j, t)]} \right),$$

where σ is a free response strength parameter. SFP is an example of a two-parameter reinforcement learning model (Erev & Haruvy, 2013). Table 1 and Fig. 1 show that this model fits the aggregate choice rate slightly better than the one-parameter contingent-sampling model.

Experiment 2: the relative importance of the two effects

The contingent-sampling model implies that the overall effect of partial reinforcements depends on the difference between the two alternatives during training. When the advantage of the promoted option is sufficiently large, partial reinforcement can increase the overall choice rate of this option. The fictitious play model predicts a weaker effect: it implies that a large advantage of the promoted option will only eliminate the negative effect of partial reinforcements. Experiment 2 was designed to clarify and compare these predictions by examining the effect of the payoff from the demoted option. Two payoff environments were compared: The payoff from the demoted option was five in Environment 5, and two in Environment 2. In all other respects, Experiment 2 was identical to Experiment 1.

The right panels of Figs. 2A and B and Table 2 presents the predictions of the contingent-sampling and fictitious-play models to Experiment 2 with the parameters that best fit the results of Experiment 1. The predictions were derived from computer simulations in which virtual agents (that behaved in accordance with the models' assumptions for the parameters estimated in Exp. 1) faced 100 training and then 100 extinction trials. The contingent-sampling model predicts that the overall effect of partial reinforcements on the choice rates of the promoted option will be negative in Environment 5, but positive in Environment 2. In contrast, the fictitious-play model predicts similar learning patterns in the two schedules.

Method

Participants A group of 49 undergraduates from the Faculty of Industrial Engineering and Management at the Technion served as paid participants in the experiment. They were recruited through signs, as in the previous setting. The sample included 25 males and 24 females (mean age 24.56 years, $SD = 2.7$).

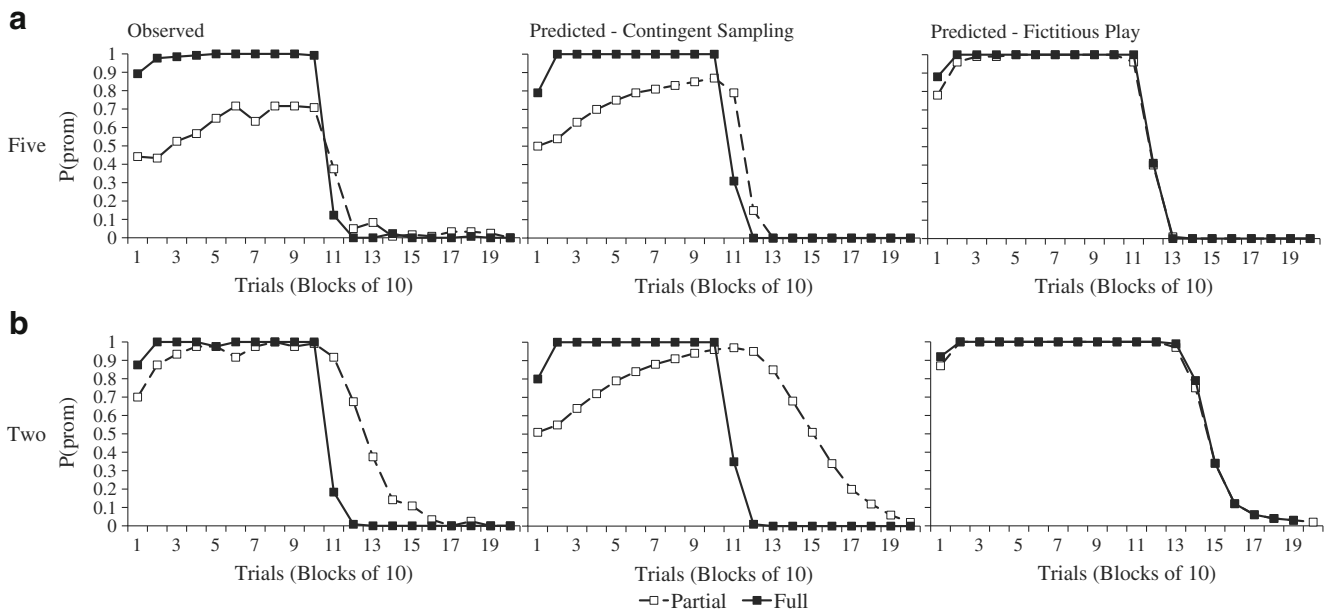


Fig. 2 (a) Observed and predicted proportion of selecting the promoted option (*P*-prom) in blocks of 10 trials as a function of the reinforcement schedule conditions in Environment Five. (b) Observed and

predicted proportion of selecting the promoted option (*P*-prom) in blocks of 10 trials as a function of the reinforcement schedule conditions in Environment Two. Extinction has started at the 11th block

Apparatus and procedure The apparatus was the same computerized money machine as in Experiment 1.

Design For the experiment, we used a 2 × 2 between-subjects design (see the left side of Table 2). The participants were randomly assigned to one of two environments

and one of two reinforcement schedules. The environments differed with respect to the value of the demoted option: five points in Environment 5, and two in Environment 2. As in Experiment 1, we used partial and full reinforcements, and each participant faced 100 training trials first, and then 100 extinction trials. The assignment to conditions was random,

Table 2 Experiment 2: Observed mean proportions of selecting the promoted option (*P*-prom), as a function of the environment, the reinforcement schedule, and the training phase, as well as the predictions of the contingent-sampling (CS) and fictitious-play (FP) models

| Schedule | Phase | Demoted | Promoted | Observed | CS ($M = 4$) | FP ($w = .05, \sigma = 4$) |
|---------------|------------|---------|----------|-----------|----------------|------------------------------|
| Environment 5 | | | | | | |
| Full | Training | | 9 | .99 (.02) | .98 | .99 |
| | Extinction | 5 | 1 | .01 (.01) | .02 | .13 |
| | Overall | | | .50 (.01) | .50 | .56 |
| Partial | Training | | 17 or 1 | .61 (.16) | .73 | .97 |
| | Extinction | 5 | 1 | .06 (.14) | .08 | .13 |
| | Overall | | | .34 (.11) | .40 | .55 |
| Environment 2 | | | | | | |
| Full | Training | 2 | 9 | .99 (.02) | .98 | .99 |
| | Extinction | | 1 | .03 (.01) | .03 | .43 |
| | Overall | | | .51 (.01) | .51 | .71 |
| Partial | Training | 2 | 17 or 1 | .94 (.07) | .78 | .99 |
| | Extinction | | 1 | .23 (.13) | .47 | .43 |
| | Overall | | | .58 (.09) | .62 | .71 |
| | | | | MSD | .009 | .038 |
| | | | | <i>r</i> | .97* | .94* |

Correlations and mean squared deviations (MSDs) are averaged across participants. Standard deviations appear in parentheses. * $p < .001$.

with 13 participants under the full schedule in Environment 5, and 12 participants in each of the remaining conditions.

Results and discussion

The experimental results are presented in the left panels of Fig. 2 and in Table 2. The main results are consistent with the predictions of the contingent-sampling model: Partial reinforcements reduced the overall choice rate of the promoted option in Environment 5 (from .50 under the full schedule to .34 under the partial schedule), but increased it in Environment 2 (from .51 to .58). A 2×2 ANOVA was conducted to test the effects of the environment and the schedule on P -prom. This analysis revealed significant main effects of both the environment [$F(1, 48) = 59.106, p < .0001, \eta_p^2 = .568$] and the schedule [$F(1, 48) = 6.974, p < .05, \eta_p^2 = .134$], as well as a significant interaction between the two factors [$F(1, 48) = 54.123, p < .0001, \eta_p^2 = .546$].

Analysis of the learning curves revealed the robustness of the pattern documented in Experiment 1. Partial reinforcements slowed learning during training and slowed extinction. In addition, the curves show that the effect of the environment could be observed in both phases: A weak demoted option (Environment 2) reduces the negative effect of partial reinforcements during training and enhances their positive effect during extinction. Table 2 clarifies the advantage of the contingent-sampling model.

The effect of multiple transitions The practical implications of these results could be questioned on the grounds that our experiments included only one transition from training to extinction (at Trial 101), whereas most applied studies have examined multiple transitions. To address this critique, we ran a replication of Experiments 1 and 2 in a multitransition setting. In this study, transitions occurred with a 5 % probability after each trial. The results replicated those in Experiments 1 and 2, and thus are not discussed further. Partial reinforcement reduced the overall choice rate of the promoted option when the demoted option was relatively attractive (Environments 8 and 5), but increased it when the demoted option was unattractive (Environment 2). In addition, we found high agreement between these results and the predictions of the contingent-sampling model.

General discussion

The present analysis starts with a pessimistic observation and concludes with an optimistic one. The initial pessimistic observation involves the inconsistency between basic and field research: The partial-reinforcement extinction effect (PREE), one of the best-known examples of a behavioral regularity discovered in basic experimental research,

appears to be inconsistent with the field research. The optimistic conclusion is the observation that the knowledge accumulated in basic decision research can help resolve this apparent inconsistency. Specifically, the notion of contingent sampling (Biele et al., 2009; Gonzalez et al., 2003), which implies reliance on small sets of experiences (e.g., Fiedler & Kareev, 2006; Hertwig, Barron, Weber, & Erev, 2004), can be used to predict the magnitude of the positive and negative effects of partial reinforcements and to clarify the conditions under which the overall effect of partial reinforcements is likely to be positive.

As was suggested by Nevin and Grace (2000), our results demonstrate that the effect of the reinforcement schedule is highly sensitive to the evaluation criteria. In many cases, partial reinforcement has a small positive effect on selecting the promoted alternative during extinction, and a larger negative effect on this behavior during training. This sensitivity can explain the apparent inconsistency between laboratory and field studies of partial reinforcements: Most laboratory studies have focused on the positive effect of partial reinforcements during extinction, whereas most field studies have focused on the fact that the overall effect of partial reinforcement tends to be negative.

In addition, our results show that when the attractiveness of the demoted option is sufficiently low, the positive effect of partial reinforcements is larger than their negative effect. This observation implies that the negative effect of partial reinforcements found in field studies might have resulted from the fact that the attractiveness of the demoted option in these experiments (slacking off) was higher than that of the promoted option (exerting high physical effort; e.g., beaver trapping in Latham & Dossett, 1978, and tree planting in Yukl et al., 1976).

In order to clarify the practical implications of the present results, we conclude with a concrete example. Assume that in a particular factory, workers must carry out a quality control test after the production of each unit. To promote this quality control test, the supervisor can reinforce this behavior. She can reinforce all of the tests that she observes (full reinforcement) or only some of the tests that she observes (partial reinforcement). Since she cannot see all of the workstations at all times (but only monitor them sometimes), the workers are alternating between a training phase (when the supervisor monitors the workstation) and an extinction phase. Our results suggest that the optimal reinforcement schedule depends on the relative attractiveness of the demoted option (i.e., not performing the test). When the test is short and relatively unobtrusive (and skipping it does not save much time and/or effort), partial reinforcements can be effective. However, partial reinforcements are likely to be counterproductive if the test is demanding (cf. Brewer & Ridgway, 1998).

References

- Atkinson, R. L., Atkinson, R. C., Smith, E. E., Bem, D. J., & Nolen-Hoeksema, S. (1995). *Introduction to psychology* (10th ed.). Orlando: Harcourt College.
- Baron, R. A., & Kalsher, M. J. (2000). *Psychology* (5th ed.). Boston: Allyn & Bacon.
- Biele, G., Erev, I., & Ert, E. (2009). Learning, risk attitudes and hot stoves in restless bandit problems. *Journal of Mathematical Psychology*, *53*, 155–167.
- Brewer, N., & Ridgway, T. (1998). Effects of supervisory monitoring on productivity and quality of performance. *Journal of Experimental Psychology: Applied*, *4*, 211–227.
- Camilleri, A. R., & Newell, B. R. (2011). When and why rare events are underweight: A direct comparison of the sampling, partial feedback, full feedback and description choice paradigm. *Psychonomic Bulletin & Review*, *18*, 377–384.
- Chen, W., Liu, S., Chen, C., & Lee, Y. (2011). Bounded memory, inertia, sampling and weighting model for market entry games. *Games*, *2*, 187–199.
- Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*, *112*, 912–931.
- Erev, I., Ert, E., & Roth, A. E. (2010). A choice prediction competition for market entry games: An introduction. *Games*, *1*, 117–136.
- Erev, I., & Haruvy, E. (2013). Learning and the economics of small decisions. In J. H. Kagel & A. E. Roth (Eds.), *The handbook of experimental economics* (2nd ed.). Princeton: Princeton University Press (in press).
- Fiedler, K., & Kareev, Y. (2006). Does decisions quality (always) increase with the size of information samples? Some vicissitudes in applying the law of large numbers. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*, 883–903.
- Fudenberg, D., & Levine, D. K. (1999). Conditional universal consistency. *Games and Economic Behavior*, *29*, 104–130.
- Gershman, S. J., Blei, D. M., & Niv, Y. (2010). Context, learning, and extinction. *Psychological Review*, *117*, 197–209.
- Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, *27*, 591–635.
- Grant, D. A., Hake, H. W., & Hornsath, J. P. (1951). Acquisition and extinction of a verbal conditioned response with differing percentages of reinforcements. *Journal of Experimental Psychology: General*, *42*, 1–5.
- Grosslight, J. H., & Child, I. L. (1947). Persistence as a function of previous experience of failure followed by success. *The American Journal of Psychology*, *60*, 378–387.
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*, 534–539. doi:10.1111/j.0956-7976.2004.00715.x
- Humphreys, L. G. (1939). The effect of random alternation of reinforcement on the acquisition and extinction of conditioned eyelid reactions. *Journal of Experimental Psychology: General*, *25*, 141–158.
- Latham, G. P., & Dossett, D. L. (1978). Designing incentive plans for unionized employees: A comparison of continuous and variable-ratio reinforcement schedules. *Personnel Psychology*, *37*, 47–61.
- Mowrer, O. H., & Jones, H. (1945). Habit strength as a function of the pattern of reinforcement. *Journal of Experimental Psychology: General*, *35*, 293–311.
- Nevin, J. A. (1988). Behavioral momentum and the partial reinforcement effect. *Psychological Bulletin*, *103*, 44–56.
- Nevin, J. A., & Grace, R. C. (2000). Behavioral momentum and the law of effect. *The Behavioral and Brain Sciences*, *23*, 73–130.
- Nevin, J. A., & Grace, R. C. (2005). Resistance to extinction in the steady state and in transition. *Journal of Experimental Psychology: Animal Behavior Processes*, *31*, 199–212.
- Papini, M. R., Thomas, B. L., & McVicar, D. G. (2002). Between-subject PREE and within-subject reversed PREE in spaced-trial extinction with pigeons. *Learning and Motivation*, *33*, 485–509.
- Patalano, A. L., & Ross, B. H. (2007). The role of category coherence in experience-based prediction. *Psychonomic Bulletin & Review*, *14*, 629–634.
- Pavlik, W. B., & Flora, S. P. (1993). Human responding on multiple variable interval schedules and extinction. *Learning and Motivation*, *24*, 88–99.
- Pritchard, R. D., Hollenback, J., & DeLeo, P. J. (1980). The effects of continuous and partial schedules of reinforcement on effort, performance and satisfaction. *Organizational Behavior and Human Performance*, *16*, 205–230.
- Svartdal, F. (2000). Persistence during extinction: Conventional and reversed PREE under multiple schedules. *Learning and Motivation*, *31*, 21–40.
- Yechiam, E., & Busemeyer, J. R. (2005). Comparison of basic assumptions embedded in learning models from experience-based decision making. *Psychonomic Bulletin & Review*, *12*, 387–402. doi:10.3758/BF03193783
- Yukl, G. A., Latham, G. P., & Pursell, E. D. (1976). The effectiveness of performance incentives under continuous and variable ratio schedules of reinforcement. *Personnel Psychology*, *29*, 221–231.