SPECIAL SECTION/*SECCIÓN ESPECIAL*

Revista Mexicana de Análisis de la Conducta • Mexican Journal of Behavior Analysis
2020 | Núm. 1 • junio | No. 1 • June | Vol. 46, 353 - 401
http://dx.doi.org/10.5514/rmac.v46.i1.76979

F. S. KELLER'S "SCHEDULES OF REINFORCEMENT" SUPPLEMENT TO KELLER AND SCHOENFELD (1950)

Kennon A. Lattal
West Virginia University

Iver H. Iversen
University of North Florida

Abstract

Fred Keller prepared two supplements for students to use in conjunction with Keller and Schoenfeld (1950), "Matters of History" and "Schedules of Reinforcement." The latter was found in a to-be-discarded file of Murray Sidman's reprints and other items after his death in May, 2019. After presenting evidence concerning the authorship of the supplement, the relation of the contents to the article to the teaching of behavior analysis to introductory psychology students in the course for which the Keller and Schoenfeld textbook was designed are discussed. The text is a remarkable example of the teaching of scientific principles and research methods - especially group and single-subject designs – because it is so rich with data derived from real experiments. It offered to introductory students facts and no fiction. It also is exemplary in its attention to relating scientific concepts to daily experience,

---

klattal@wvu.edu. The authors contributed equally to this article. We are indebted to John Keller for his invaluable help in determining the authorship of the "Schedules of Reinforcement" supplement (thank you, Watson), and to Charlie Catania for sharing with us an earlier draft of his reminiscence of "The course for which K &S was written" and material related to Keller's autobiography.

a critical feature of scientific material directed to such introductory-level students.

Keywords: Keller and Schoenfeld, Principles of Psychology, Columbia University, Course supplement, Schedules of reinforcement

Resumen

Fred Keller realizó dos suplementos que se usarían en conjunto con el libro de texto Keller y Schoenfeld (1950), cada uno de ellos titulado “Cuestiones de Historia” y “Programas de Reforzamiento”. Posterior a la muerte de Murray Sidman, en Mayo de 2019, el segundo de ellos fue encontrado entre sus documentos, los cuales estaban por ser desechados. A continuación se presenta evidencia de la autoría del suplemento, así como una discusión de su relación con la enseñanza del análisis de la conducta a estudiantes de psicología, inscritos en el curso para el cual el libro de Keller y Schoenfeld fue diseñado. Debido a la abundante cantidad de datos derivados de experimentos que se llevaron a cabo, el libro en cuestión es un asombroso ejemplo de la enseñanza de los principios científicos y los métodos de investigación—especialmente con diseños de grupo e intra sujetos. Le brindaba a los estudiantes hechos y no ficción. El libro es igualmente excepcional en su interés por relacionar conceptos científicos con la vida cotidiana, una característica fundamental para la presentación de material científico a estudiantes de nuevo ingreso.

Palabras clave: Keller and Schoenfeld, Principles of Psychology, Columbia University, Suplemento de curso, Programas de reforzamiento

Keller's "Schedules of Reinforcement" Supplement to Keller and Schoenfeld (1950)

After Murray Sidman's death in May of 2019, his son sent the present authors a manila folder containing mostly reprints that his father had accumulated over his long professional life. Among the items was a mimeographed document in a single-document binder. The document was labeled "Supplementary Notes," with the author listed as F. S. Keller. The notes were in two parts, separately numbered but stapled together as a single document with a single cover page. Figure 1 shows a photograph of the mimeographed cover page, dated 1956-57. The first part, titled "Matters of History" (hereafter, the *History* supplement) was 21 single-spaced typed pages. This part is reproduced in the 1995 B. F. Skinner Foundation edition of Keller and Schoenfeld's *Principles of Psychology* (1950; hereafter, K & S). The cover page in the 1995 edition of K&S lists the date of that supplement as 1958-59, apparently a later iteration of the *History* supplement that is part of the 1956-57 supplement. The second part of the latter supplement, which is the subject of this article, was not reproduced or mentioned in the 1995 Skinner Foundation edition. It is titled "Schedules of Reinforcement" (hereafter, the *Schedules* supplement) and is 11 single-spaced typed pages, followed by another five pages of cumulative records and hand-drawn graphs. Direct scans of the original mimeographed article appear on the *Mexican Journal of Behavior Analysis* website at http://rmac-mx.org/wp-content/uploads/2020/08/Original-scanned-handouts-for-KS.pdf. A typeset version of the supplement derived from a character recognition scanning of the original *Schedules* supplement appears as Appendix A at the end of this article. A narrative review of some of the *Schedules* supplement content follows an analysis of its authorship.

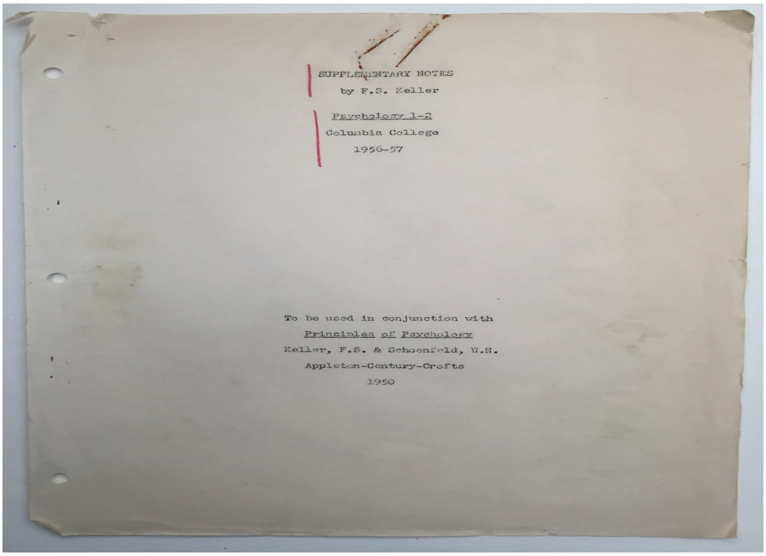

SUPPLEMENTARY NOTES
by F.S. Keller

Psychology 1-2
Columbia College
1956-57

To be used in conjunction with
Principles of Psychology
Keller, F.S. & Schoenfeld, W.N.
Appleton-Century-Crofts
1950

Figure 1. Cover page that appeared before the *History* supplement to which the *Schedules* supplement was attached. The two holes from the staple that held both supplements together is visible in the upper left corner of the page, as is the rust spot from the paper clip that at one time held the two supplements together. This same rust spot is visible on the reverse of the last page of the *Schedules* supplement.

Authorship

The question of authorship of the *Schedules* supplement arose during the course of researching its provenance. Catania (2020, this issue) observed the following:

> The earliest course materials for PSYC 1-2 were written by Keller and Schoenfeld, *as was a supplement devoted to the history of psychology*. As lab procedures were modified over the years some new material was perhaps written by teaching assistants and others. For example, in his autobiography Keller refers to supplementary readings written by Donald Bullock (Keller, 2009, p. 212). Other indirect evidence is that *K&S usually adhered to the usage that responses rather than organisms were reinforced (Catania, 1987), whereas a supplement on reinforcement schedules does not do so*. Inconsistencies in the reinforcement language also can be found in the lab handouts detailed below. Also, we might assume Keller would

> not have created a cumulative record with occasional negative slopes (Appendix 1, p. 6); on the other hand those could be attributed to a shaky hand drawing directly on an uncorrectable mimeograph master. (p. 308)

Catania acknowledged that Keller authored the 1956-57 *History* supplement in the first italicized portion, but raises the question in the second italicized portion as to whether Keller was the author of the *Schedules* supplement because of a difference in expression concerning the reinforcement of responses versus organisms (see Catania, 1969), something he also finds in the laboratory manual. It also should be clarified that his subsequent comments above about the distorted – most likely hand-drawn – cumulative record do not apply to the *Schedules* supplement, but only to the laboratory manual in which it appeared. The following evidence strongly suggests that Keller was the author of the *Schedules* supplement.

1. Figure 1 shows that Keller is listed as the sole author of the supplement. This cover page was followed by the *History* supplement, with the *Schedules* supplement following, all in one binder and bound by the single staple as already noted.
2. Comparing the original typed mimeograph pages of the *Schedules* supplement to personal letters typed by Keller to Skinner suggest that both were typed on the same manual typewriter, and furthermore, were typed by Keller himself. Keller's personal letters share at least two idiosyncratic features with the *Schedules* supplement: (1) question marks are preceded by an extra space (an idiosyncrasy his son, John, speculated might be the result of his history as a Morse code operator; J. Keller, personal communication, June 22, 2020) and (2) the lowercase "w" is often below the line of the bottom of other letters. The *History* supplement seems to have been typed on a different typewriter because it lacks the misplaced lowercase w and generally has a "cleaner" appearance. In the latter, none of the four question marks is preceded by an extra space.
3. In the second italicized portion of the above quotation, Catania (2020, this issue) suggested that K&S "usually" described responses rather than organisms being reinforced, whereas *the Schedules*

supplement does not adhere to this usage. Were this the case, it could be taken as evidence that Keller did not author the *Schedules* supplement. In the *Schedules* supplement, based on a simple word search there is reference to responses being reinforced 5 times and to organisms being reinforced 6 times. In K & S, also based on a similar word search of the text, the corresponding numbers are 18 and 13, respectively. In K & S, "organisms reinforced" includes "make his reinforcement," "get his reinforcement," "get his pellet or candy." In addition, the locution "group was reinforced" and the "group was extinguished" both appear twice. Strictly speaking, a group of subjects is not an organism, but it is not a response either. Thus, the usage in the *Schedules* supplement is consistent with that in K & S.

4. In his autobiography, Keller (2009, p. 212) stated that he wrote to Robert Yerkes, describing "our textbook and a set of supplementary reading that Donald Bullock was preparing," raising the question of whether Bullock might have written the supplement. In the first place, Keller identifies Bullock as "preparing" (not "written" as the Catania quote above indicates) the supplementary reading. The timing also is off. This quotation appears in a part of Keller's autobiography describing activities around 1948-49, several years before the date of the *Schedules* supplement. Much of the research described in the *Schedules* supplement was reported in psychological journals between 1952 and 1956. Bullock received his Ph.D. from Columbia in 1950 and apparently left soon thereafter. A 1951 publication lists his affiliation as the University of Buffalo and by 1956 he was affiliated with Smith, Kline, and French pharmaceuticals. This leads to the conclusion that Bullock was not associated with the Psychology 1-2 course at Columbia in 1956, when the supplement is dated.
5. The Fred S. Keller Papers at the University of New Hampshire lists among its holdings the following manuscript:

> Bullock, Donald H. (in collaboration with Fred S. Keller and William N. Schoenfeld). "Researches in the Science of Behavior." 1949. Approximately 100pp., typed (Box 34, Folder 3)

This manuscript was unavailable to the authors due to the 2020 COVID-19 pandemic, but its title is different from the title on the cover page of the supplements, and from the title of each of the supplements.

The google books website contains the following entry:

COLUMBIA UNIVERSITY. Dept. of Psychology.

Researches in the science of behavior. Pt. 1-[2] By Donald H. Bullock in collaboration with Frederick S. Keller (and) William N. Schoenfeld. 2 v. illus. ©Donald Hartmann Bullock; 10ct49, 1Feb50; AA140144, 146844.

6. Another possible author of the *Schedules* supplement is Schoenfeld. In terms of the development and implementation of the Columbia introductory psychology course, Keller and Schoenfeld were joined at the hip. They co-authored K & S, of course; wrote the 1948 *American Psychologist* article describing the first iteration of the course together (Keller & Schoenfeld, 1948); appeared with Frick (Frick et al., 1947) as co-authors of the article describing the equipment used in the laboratory portion of the course; and were listed together on the Bullock piece described above. In light of this history of co-authorship on K & S course-related material, it would seem very out of keeping with Keller's personal style of inclusion to exclude the name of someone collaborating on the supplementary notes and unfathomable that he would add his name to the exclusion of the real author.

Although some of the evidence in the above points is circumstantial, taken together it leads the present authors to conclude that Keller is the *Schedules* supplement's author, just as he is credited as being the author of the *History* supplement.

The Substance of the *Schedules* Supplement

[Page number references to the *Schedules* supplement are to the pages of the original scanned mimeograph version that appears on the *Mexican Journal of Behavior Analysis* website at http://rmac-mx.org/

wp-content/uploads/2020/08/Original-scanned-handouts-for-KS.pdf. The original page numbers are marked on the typeset version that appears in Appendix A at the end of this article by a bolded and bracketed ”p” followed by the page number (e.g.,[**p. 2**]).

The Supplement is dated 1956-1957, which falls between the publication of K&S in 1950 and the publication of *Schedules of Reinforcement* by Ferster and Skinner (1957) and Skinner's (1957) outline of an experimental analysis of behavior. Thus, from an instructional perspective the *Schedules* supplement may be seen as a stepping-stone from K&S to Ferster and Skinner. The *Schedules* supplement is impressive for its level of explanation of schedules and the behavior processes at work. It is an amazingly sophisticated account for an introductory level course. It is considerably richer in coverage of schedules of reinforcement than was K&S. Indeed, even the term “schedules of reinforcement”, used as the title of the supplement, was new compared to K&S, which mentions “reinforcement schedules” only once.

The contemporary parlance related to “schedules of reinforcement” is well established for behavior analysts in the 21st century thanks to Ferster and Skinner (1957), a groundbreaking book that introduced vocabulary, methods, and knowledge related to schedules of intermittent reinforcement of operant behavior. Before that, K&S was a similarly groundbreaking text that introduced Skinner's views and methods (Skinner, 1938) to a completely new generation of psychologists from the 1950s forward (see, e.g., Catania, 2020, this issue). K&S for decades thereafter was used as a textbook and, as has already been noted, was reprinted in 1995.

Regarding what is now termed intermittent reinforcement, K&S primarily dealt with what Skinner had called periodic reconditioning, which was contrasted with regular reinforcement where each single response is reinforced each time it occurs. Periodic reconditioning introduced some regular intermittency such as a time interval between reinforcement of responses (now called fixed-interval , FI, schedules) or a fixed count of responses between reinforcements (now called fixed-ratio, FR, schedules). K&S provided examples of periodic reinforcement and extinction after periodic reinforcement, with all six figures on periodic reconditioning in K&S showing results from the work by Skinner (1938). The main focus in K&S was on the temporal discrimination that developed on periodic schedules, where after training a pause in responding would form right after reinforcement. This resulted in scalloped cumulative curves for interval schedules and pauses followed by a high response rate for ratio schedules.

K&S interestingly realized that the periodic reconditioning schedules were somewhat limited to laboratory work. As they stated:

> Outside the laboratory, regular reinforcement [now called continuous reinforcement – each response is reinforced] is by no means the rule, but neither is strictly periodic reinforcement. It is hardly to be expected that a schedule of any fixed interval or any fixed number of responses would be scrupulously honored by an environment so crowded with different events. We may well ask, then, whether the results of *aperiodic* reinforcement are the same as those of *periodic* or *regular* reinforcement. (p. 98-99)

Aside from mentioning some studies with aperiodic reinforcement with respondents and galvanic skin responses, K&S only mentioned what now is called a variable-interval (VI) schedule, but with no data other than brief reference to the "straight-line character of cumulative-response curves" (p. 100) and the great resistance to extinction built by such a schedule. K&S did, however, observe that "in very few spheres of human activity is reinforcement either regular or strictly periodic, and, in certain cases, the effect of this aperiodicity is dramatically impressive" (p. 101). They then listed several examples like the chronic gambler, although they did not refer to any ratio-type aperiodic schedules.

Research on schedules after Skinner (1938) was slow to develop, and Skinner's own clandestine war-related work on pigeons controlling "smart" weapons in the early 1940's was not known publicly until his (Skinner, 1960) description of project ORCON. Among other methods, that work involved development of variable ratio (VR) schedules. Maybe the first formal description of VI schedules was in Skinner (1950). Although not using that exact term, he did describe a mixture of time intervals between reinforcements including a “zero second” interval where the first response after reinforcement is reinforced. The intervals ranged from zero to 2 minutes with an average of one minute and occurred in mixed order within a session. A figure in the 1950 article shows straight-line cumulative records with no pauses after reinforcement, in stark contrast to the pattern developed on periodic schedules (FI and FR), where pauses after reinforcement are the norm. By the late 1950s, the main empirical sources of schedule work were Ferster and Skinner (1957) and Skinner (1957), where the terms periodic and aperiodic reinforcement were replaced with the now familiar schedule taxonomy and illustrated with experimental data. In *Science and human behavior*, Skinner (1953) described, without experiments or data, the four now-familiar schedules (FI, FR, VI, and VR). This new terminology appears also in the *Schedules* supplement, but not in K & S.

A main theme of the *Schedules* supplement is generalizability of basic findings. The text emphasizes schedule control of behavior of different species such as rats, pigeons, cats, monkeys, and humans. The K&S textbook was an important link between basic research, primarily Skinner’s, and general education in psychology. Generalizability of basic findings was therefore a critical component in the teaching of introductory psychology students. For example, Keller emphasized that a time discrimination develops on FI schedules (as evidenced by pauses in responding after reinforcement). This happens “with all the organisms thus far studied in the laboratory, including human beings” (*Schedules* supplement, p. 3). Similarly, Keller wrote that “[h]uman subjects provide fixed-ratio curves that are often indistinguishable from those produced by monkeys, dogs, rats, and other subjects” (*Schedules* supplement, p. 7).

The *Schedules* supplement has several interesting statements regarding scientific thinking. For example, K&S reported that on FI schedules each organism emits roughly the same number of responses between reinforcements independent of the value of the FI. Thus, one rat may make about 20 responses between reinforcements on FI 2 min and also 20 responses on FI 6 min. Keller wrote: "We now believe that this suggestion was based upon too little information; and that the number of responses per interval probably increases as the intervals get longer. At least two recent studies point clearly to such a conclusion" (*Schedules* supplement, p. 2). Students were to learn from this that suggestions or principles depend on empirical information that in turn depends on research, and that one could not progress too far with "too little information." Keller articulated some issues related to experimental design as the cause of the initial erroneous conclusion. But then he emphasized that "the important thing, of course, is not the old error, but the new advance in our knowledge" (*Schedules* supplement, p. 2). In the same vein, when Keller asked how the size of the fixed interval in training would affect the response rate in extinction, he stated that "[c] ommon sense won't give you the answer to such a question. Instead, laboratory information is desirable" (*Schedules* supplement, p. 3).

The *Schedules* supplement features a curious mixture of group design and single-subject design. Dinsmoor, reflecting on the research environment at Columbia University in the late 1940s to early 1950s, observed that:

> It may come as something of a shock to those who became familiar with the experimental analysis of behavior only after [the *Journal of the Experimental Analysis of Behavior*] was founded, but almost all of the conditioning research during my stay at Columbia had been based on the traditional experimental design in which the mean performance of one group of subjects is compared with the mean performance of another group, treated differently in some way, and a statistical test is conducted to determine whether the results could have arisen by chance. The outstanding exception was Murray Sidman's dissertation. (Sidman, 1953) (1990, p. 147)

During the 1950s, research designs in behavior analysis were still influenced by the types of "group" design from which Skinner departed in his early research. In the *Schedules* supplement, Keller described experiments using groups of subjects such that each group was exposed to one level of an independent variable; for example, six groups of rats each had one FI value and groups were compared regarding the average number of responses per interval (study by Wilson, 1954;

Schedules supplement, p. 2). Yet, on the same page, Keller relates an experiment with a single-subject design. In it, each rat was exposed to different values of a light-termination schedule where a strong light was turned on. After x seconds, the first response by the rat would switch off the light for 1 minute. Each rat was exposed to 6 values of the fixed interval. Nonetheless, Keller reported the data as averages across rats with the FI value as the independent variable (Figure 2 in the *Schedules* supplement). The important emphasis by Keller was that the response rate was a function of the interval value whether responding produced food or light removal, or whether it was a group or single-subject design (Keller did not in the supplement articulate the difference in design for the emphasis on generalization of the research). Thus, this comparison reveals several layers of teaching generality to students regarding how schedules affect behavior both across consequences and methodological designs. Keller added that the light-removal schedule also could be an FR, as well as an FI schedule, as was shown in the single-subject design experiment with different FR values by Kaplan (1956), also described in the *Schedules* supplement.

In spite of some cases showing data averaged across subjects, Keller also emphasized the individual subject's behavior as a determinant of schedule effects. He described the effect of response rate before extinction as an important variable for extinction as follows: "The way an animal behaves during reinforcement is a good indication of what he will do during extinction. A slow, steady responder during training may give a 'fixed-interval' type of curve in extinction, and a high-speed responder in training may give a 'fixed-ratio' curve in extinction, in spite of the fact both animals may have been working on a fixed-interval schedule" (*Schedules* supplement, p. 4). To this he added: "The important thing, however, is not the schedule, but the kind of behavior exhibited while the schedule is in operation" (*Schedules* supplement, p. 4). This issue of the subject's behavior at one point being a determinant of behavior at a later point was further illustrated by an early reference to work by Ferster and Skinner, about five years prior to their book on schedules of reinforcement (the reference to Ferster & Skinner, 1952, in the *Schedules* supplement). At the time, terminology was not fully developed, and what was described is now called a mixed FR FR schedule, where two different ratios alternate randomly within a session, such as FR 50 and FR 250. After considerable training, a pattern develops whereby after reinforcement there is a short pause, and if the next ratio is the large FR 250 then the animal stops after 60 of 70 responses and then responds again with another similar response run until the FR 250 is reached. As Keller wrote: "The bird acted, you might say, as if he had a crude sort of counter or inner clock by which it regulated its own pecking behavior. … … The cue for stopping a run, in each case, was suggested to be the emission of a certain number of responses, rather than by some change in the outside situation. There was, you might say, a 'feedback' from the animal's own behavior that led him to stop at a given point" (*Schedules* supplement, pp. 7-8). Keller also described experiments by Mechner, which apparently were early versions of later-published work (Mechner, 1958). Rats had to press one lever x number of times and then switch to a second lever that produced reinforcement if a sufficient number of presses were made on the first lever. Mechner studied the number of responses that exceeded the required number of responses. These experiments showed how precisely a rat's behavior on one lever can function as the only stimulus for making a switch to a second lever.

This theme of when one aspect of behavior becomes a discriminative stimulus for another aspect of behavior was further developed in Ferster and Skinner (1957). Again, one has to appreciate Keller's emphasis to his students in the *Schedules* supplement that behavior is related to the environment in an extremely complex manner such that one may not see the environment (the "outside situation" in Keller's words) acting when behavior changes (as seen in the unusual pause patterns in mixed schedules). Instead, a history of training sets up certain patterns of behavior that somehow rely on other patterns. Interestingly, the conclusion that the experiments show that "an animal's response may depend upon an inner 'clock'" (*Schedules* supplement, p. 7) actually introduces controlling variables that are not independent variables, but are embedded in training history. These "inner clocks" as causes of behavior are not entirely different from what K&S warned about when they described the problems with "inferred internal" states or drives. As K&S wrote: "… the phrase 'inferred internal state' adds nothing to our knowledge of drive, because it denotes nothing beyond that which is contained within the observations themselves. It is, once again, a convenience of expression, and we might dispense with the term altogether if it were not for the effort involved in straining for technical purity" (K&S, p. 285). This complex issue of whether internal states are a result of experimental history or are independent causal agents remains a topic of discussion in contemporary behavior analysis (e.g., Eckard & Lattal, 2020).

The *Schedules* supplement also covers a comparison of VI and VR schedules. It describes an experiment later reported by Ferster and Skinner (1957) involving VI-schedule control of pigeon's key pecking. In one case, key pecking was reinforced according to a VI 3-min schedule with intervals ranging from a few s to 6 min. On one occasion, the pigeon pecked steadily for 15 h, accumulating 30,000 responses with practically no pauses in responding during the entire period. Keller amusingly quoted the authors this way: "The reporters of this study (Ferster and Skinner) would seem to be guilty of understatement when they assert that 'the control exercised by a schedule of this sort may be very great'" (*Schedules* supplement, p. 9). We may easily assume that Keller and Skinner communicated frequently and in detail about the research they were doing at that time (Keller, 2009).

In the supplement, Keller described a VR schedule as one based on random variation in the number of responses that must be made before reinforcer delivery. Response rates become very steady, as for VI schedules, with no pauses. There is, of course, a difference from VI schedules. In the words of Keller: "the ratio schedule leads as a rule to higher response rates than does the interval schedule. This is because a response that follows any break in an interval schedule has a greater likelihood of getting rewarded; whereas a pause during a ratio schedule never increases the chances of reward for the next response" (*Schedules* supplement, pp. 9-10). Keller also related the VR schedule to gambling in humans, something Skinner (1953) had done earlier.

The *Schedules* supplement ends with some examples of conditioning extreme rates of responding (what today are sometimes called schedules of inter-response time reinforcement, among other labels). The contingency difference between VI and VR schedules, articulated above, can be used explicitly to drive response rates up when reinforcement is contingent on emitting one response within a short time after a previous response and to drive response rates down by requiring a pause in responding between successive responses. As for other examples in *Schedules,* Keller provides actual data from experiments to illustrate the reinforcement of different rates of responding.

The *Schedules* supplement does not conclude with comments or even a summary. No theory or basic principles were outlined. The overall text, however, is a remarkable example for the teaching of scientific principles, research methods - especially group and single-subject designs - because it so rich with data derived from real experiments. In short, it offered to introductory students facts and no fiction. It also is exemplary in its attention to relating scientific concepts to daily experience, a critical feature of scientific material directed to these introductory-level students.

References

Catania, A. C. (2020). The course for which K & S was written. *Mexican Journal of Behavior Analysis, 46,* 1. http://dx.doi.org/10.5514/rmac.v46.i1

Dinsmoor, J. A. (1990). Academic roots: Columbia University, 1943-1951. *Journal of the Experimental Analysis of Behavior, 54* (2), 129-149. https://doi.org/10.1901/jeab.1990.54-129

Eckard, M., & Lattal, K. A. (2020). The internal clock: A manifestation of a misguided mechanistic view of causation? *Perspectives on Behavior Science, 43*(1), 5-19. https://doi.org/10.1007/s40614-018-00189-5

Ferster, C. B., & Skinner, B. F. (1957). *Schedules of reinforcement*. Appleton-Century-Crofts.

Keller, F. S. (2009). *At my own pace: The autobiography of Fred S. Keller.* Sloan.

Keller, F. S., & Schoenfeld, W. N. (1950). *Principles of psychology: A systematic text in the science of behavior*. Appleton-Century-Crofts.

Mechner, F. (1958). Probability relations within response sequences under ratio reinforcement. *Journal of the Experimental Analysis of Behavior,* 1 (2), 109-121. https://doi.org/10.1901/jeab.1958.1-109

Sidman, M. (1953). Two temporal parameters of the maintenance of avoidance behavior. *Journal of Comparative and Physiological Psychology, 46* (4), 253-261. https://doi.org/10.1037/h0060730

Skinner, B. F. (1938). *The behavior of organisms: An experimental analysis.* Appleton-Century-Crofts.

Skinner, B. F. (1950). Are theories of learning necessary? *Psychological Review,* 57 (4), 193-216. https://doi.org/10.1037/h0054367

Skinner, B. F. (1953). *Science and human behavior*. Macmillan.

Skinner, B. F. (1957). The experimental analysis of behavior. *American Scientist,* 45 (4), 343-371. https://www.americanscientist.org/article/the-experimental-analysis-of-behavior

Skinner, B. F. (1960). Pigeons in a pelican. *American Psychologist,* 15 (1), 28-37. http://www.appstate.edu/~steelekm/classes/psy3214/Documents/Skinner1960.pdf

APPENDIX A

Converted editable copy of the original mimeographed copy of the "Schedules of Reinforcement" supplement that is reproduced in its original mimeographed format at http://rmac-mx.org/wp-content/uploads/2020/08/Original-scanned-handouts-for-KS.pdf.

[Note: The cover page is as it appeared before the "Matters of History" supplement described in the accompanying article. This supplement on schedules of reinforcement directly followed the "Matters of History" supplement in the version received by the present authors. The cover page is included here simply as an introduction to the supplement, but should not be construed as being part of the original document, which appears at the above-noted web address. The actual cover page is reproduced in Figure 1 of this article. Bolded numbers in brackets (e.g., [**p. 2**]) demarcate the beginning of the indicated page in the original document. In this appendix, the line was reset at the beginning of each new original document page. Punctuation is as in the original document. Also, in this version, as compared to the original mimeographed version at http://rmac-mx.org/wp-content/uploads/2020/08/Original-scanned-handouts-for-KS.pdf, the figures have been rearranged slightly so they are in numerical order. Specifically, Figure 11 was moved from the last page of the original document to its nominal place in order in this appendix.]

SUPPLEMENTARY NOTES

by F.S. Keller
Psychology 1-2
Columbia College
1956-57

To be used in conjunction with
Principles of Psychology
Keller, F.S. & Schoenfeld, W. N.
Appleton-Century-Crofts
1950

Schedules of Reinforcement

There are several ways of classifying the schedules of reinforcement that may be used in connection with operant and respondent behavior. In the present case, we begin with the broad, and well known, distinction between regular and intermittent schedules. Then we move to subdivisions of the intermittent class and, finally, to cases that can best be described as "rate conditioning."

Regular Reinforcement (Continuous or 100-Percent Reinforcement) This schedule may be used with both operant and respondent behavior. It is seen, in the respondent case, when the unconditioned stimulus, such as food or electric shock, always accompanies the conditioned stimulus during the training procedure. You probably think of this as the standard procedure in Pavlovian conditioning, but actually such a view is not quite correct. Pavlov and his pupils commonly tested for conditioning by presenting the conditioned stimulus alone after a certain number of pairings --- say on every tenth occasion. Even one such test would keep a reinforcement schedule from being truly "regular."

In the operant case, we speak of regular reinforcement when the response produces a reward, or gets rid of an aversive stimulus, on every occasion that it occurs. The movement of a lever may always provide food for a laboratory rat; a penny in a slot machine may always produce gum; or the adjustment of a desk lamp may always reduce glare. We may also speak of regular reinforcement when an operant is always followed by a "punisher." Thus, a lever press may always result in a shock, rather than food, for a white rat; a finger in the fire may always be burned; or the utterance of some word may always be met with disapproval.

Intermittent Reinforcement (Partial Reinforcement) In respondent conditioning, intermittent reinforcement is said to exist when an unconditioned stimulus is paired with a conditioned stimulus on certain occasions only. For example, food might be paired with the bell , in the classical example, on, say, every third trial, but omitted at all other times. Or, the food might accompany the bell after a variable number of trials without it. Thus, bell-alone might be presented on trials 2, 4, 7, 9, and 14, with bell-and-food on all the intervening trials. The trials in which pairing took place would thus occur in a more or less random order.

With respect to operant conditioning, as noted above, several subdivisions of intermittent schedules have come to be recognized. We shall consider these subdivisions in the same order as that in which they are presented in your textbook (Keller and Schoenfeld, pp. 83-102).

Fixed-Interval Reinforcement (Periodic Reconditioning at Fixed Intervals). This is perhaps the best known of the intermittent schedules. It involves simply the reinforcement of a response after a set period of time during which it has not been reinforced. In one of the earliest experiments with this schedule, white rats were reinforced for pressing a lever at 5-minute intervals throughout a 1-hour session. The first response of the hour was reinforced with a pellet of food. Then followed a 5-minute period in which the responses were not reinforced. The first response after this period was reinforced. Then followed another 5-minute period, another reinforced response, and so on, for the remainder of the hour. Details of this procedure, and its effect upon rate of response and [**p. 2**] resistance of the response to extinction, have been described in your textbook (pp. 83-91). Your text account, however, requires some correction and expansion in the light of recent findings.

For example, it is suggested therein (p. 88) that the number of responses made by an organism during the periods of non-reinforcement in a fixed-interval schedule will be approximately the same with one interval as it is with another --- that the number of unreinforced responses for each reinforced response (the so-called "extinction ratio") is fairly constant for a given animal. That is, if the animal makes, on the average, 20 responses between reinforcements on a 3-minute schedule, he will also make approximately 20 responses on a 6-minute or a 9-minute schedule.

We now believe that this suggestion was based upon too little information; and that the number of responses per interval probably increases as the intervals get longer. At least two recent studies point clearly to such a conclusion. In one, by Wilson (1954), six groups of white rats were taught to bar-press, and then were [word or words missing in original] 240 fixed-interval reinforcements. Each group was trained with a different interval, and the intervals ranged in length from 10 seconds to 6 minutes. Wilson found that the number of responses per interval for the different groups climbed from 5, with the 10-second interval, to more than 20, with the 6-minute interval. In other words, the number of responses, with different fixed intervals, is by no means constant. This is readily apparent in Figure 1.

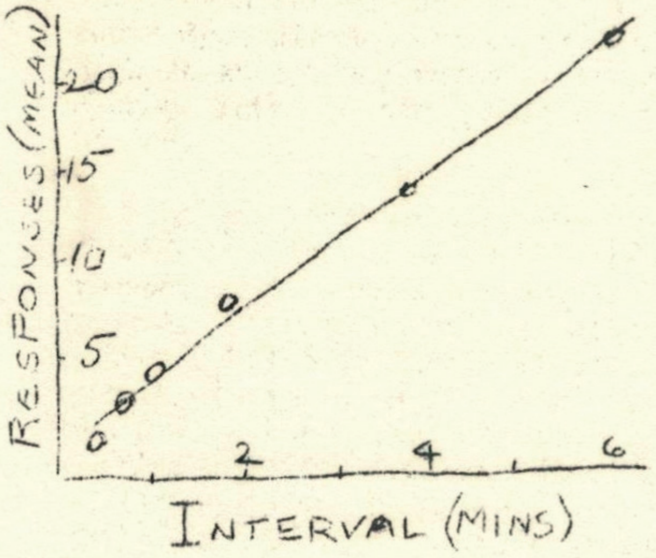

FIG. 1

In another study, by Kaplan (1952), 5 white rats were reinforced by light-removal after different intervals of non-reinforced responding. The size of the interval ranged from 12 seconds to 5 minutes, and each animal was given about 15 sessions of practice at each one of the intervals. The procedure, for a given interval, was as follows. A 25-watt light was turned on above the animal's head, and remained on for a period of, say, 30 seconds. No lever-pressing response, during this period, would turn the light off; but the first response made after the time was up would be successful, leaving the animal in darkness for about 1 minute. Then the light would come on again for the 30-second minimum and the whole cycle of events would be repeated. When each rat had been tested with all the different intervals, Kaplan found that the average number of non-reinforced responses at the 12-second interval was about 5. From then on the number increased gradually until, at the 5-minute interval, it amounted to more than 40 responses. The results are shown in Figure 2.

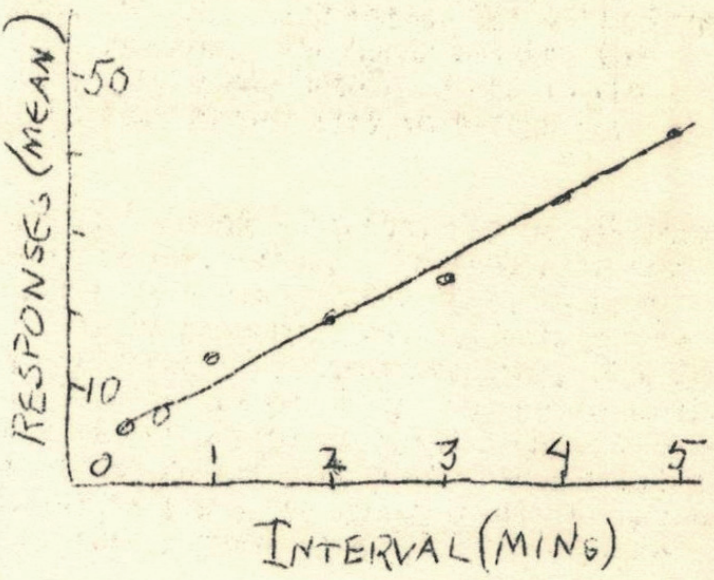

FIG. 2

Since both of these experiments point in the same direction, we should probably revise our opinion concerning the constancy in number of responses during different intervals of periodic reinforcement. Yet, if our earlier opinion was wrong, in what way were we misled , This question cannot be answered with complete assurance. However, it seems likely that the original study, in which each one of four rats was tested for bar-pressing rate at fixed intervals of 3, 5, 7, and 9 minutes, was defective in two respects. The tests at each interval were too few in number to provide reliable averages; and there was not enough sampling of rates at the short-interval end of the scale, where rate change occurs most rapidly. The important thing, of course, is not the old error, but the new advance in our knowledge.

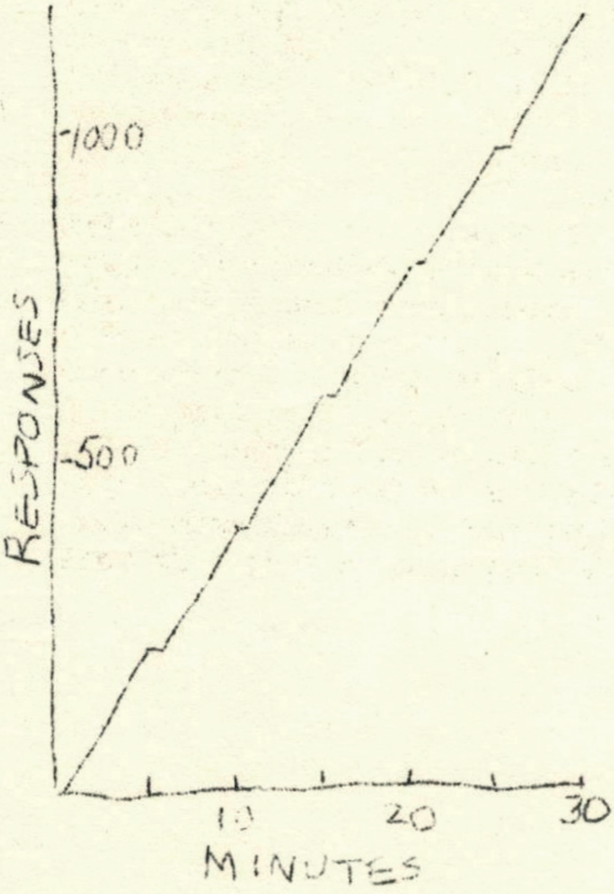

Fig. 3

When a fixed-interval schedule of reinforcement has been in force for awhile, a time discrimination usually develops (K & S, p. 86). This happens with all the organisms thus far studied in the laboratory, including human beings. The discrimination shows itself in the fact that, right after a reinforcement, the response rate drops down to zero or nearly so, whereas just before reinforcement it is usually at its maximum. Look at Figure 3. This is a cumulative [**p. 3**] curve for the pecking response of a pigeon. The bird was reinforced on a 5- minute fixed-interval schedule for pecking at a little disc on the wall of the experimental chamber. You will note that, immediately after each reinforcement, there was usually a pause of from 3O to 40 seconds. Then, rather suddenly, the bird would pick up a steady and fairly high rate of pecking that ordinarily took him all the way to his next reinforcement.

A pigeon will also show a time discrimination when the fixed interval is much longer than 5 minutes, although the rate change may then have a different appearance. Figure 4 is taken from the record of a bird that had been working for many days on a 45-minute fixed-interval schedule. This curve is for just one such interval. You can see that here, too, there is a pause right after reinforcement --- a pause that lasts for many minutes rather than seconds. Then the responses pick up in speed gradually, rather than abruptly, from zero to the highest rate.

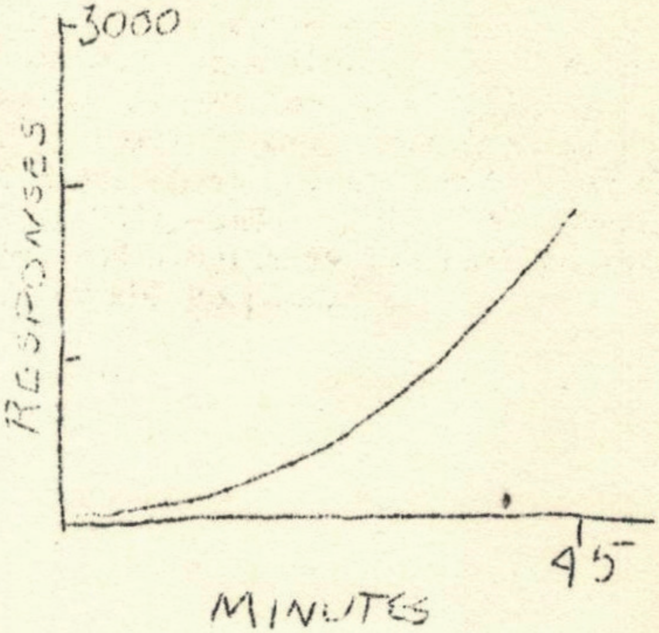

FIG 4

Extinction after Fixed-Interval Reinforcement. It has often been noted by students of behavior that an intermittently reinforced response is more difficult to extinguish than a regularly-reinforced one. This seems to be true for any type of intermittent schedule, fixed-interval or otherwise. But we still need to know a lot more about the relation between reinforcement schedule and extinction responding. For example, in the fixed-interval case, what does the size of the interval have to do with the ease of extinguishing the response ? Will a long interval establish a response more firmly than a short one, or will the reverse be true ? Common sense won't give you the answer to such a question. Instead, laboratory information is desirable.

In the study by Wilson, mentioned earlier, he tried to get such information. Six groups of rats were used in his study, each trained with a different interval. The intervals, as already mentioned, ranged from 10 seconds to 6 minutes, and the members of each group received 240 reinforcements each. The same number of reinforcements was also given to a 7th group of animals ---a regularly-reinforced, or zero-interval group. At the end of training, all of the groups were put through extinction sessions of 50 minutes each on 5 successive days. The results are represented in the upper curve of Figure 5.

You can see from this curve that the bar-pressing response was hardest to extinguish in those rats that had been reinforced on the 1-minute fixed-interval schedule. The response was more readily extinguished in all the other groups, both above and below the 1-minute value. The zero-interval or regular-reinforcement group was, as usual, the least resistant to extinction.

Such results are not, at present, easy to understand. It is quite unlikely that anyone would have predicted the greatest effect from the 1-minute value in Wilson's series. Yet, this high point in the curve does not seem to be a chance affair, or limited to these particular animals or this number of reinforcements. In a second experiment, Wilson used other groups of rats and gave each group 15 reinforcements, rather than 240. The extinction results are portrayed in the lower curve of Figure 5.

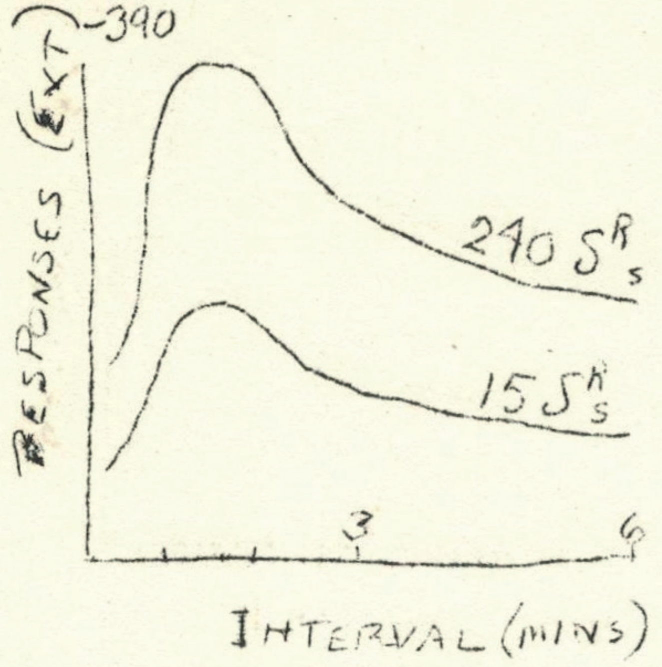

FIG. 5

“Typical” Extinction Curves. If you were to look at a number of individual curves for the extinction of an operant after some fixed-interval schedule had been in effect for a long time, you would find that some of the curves looked like the “typical” ones of your textbook (K & S, p. 90). Others, however, might [**p. 4**] look more like the extinction curves obtained after fixed-ratio training (K & S, p. 90). This might be very puzzling unless you went further and examined the rate curves for each individual before the extinction began. Then you would discover that the form of the extinction curve may depend more upon the rate of an animal’s response during training than it does upon the type of training schedule. The way an animal behaves during reinforcement is a good indication of what he will do during extinction. A slow, steady responder during training may give a "fixed-interval" type of curve in extinction, and a high-speed responder in training may give a "fixed-ratio" curve in extinction, in spite of the fact both animals may have been working on a fixed-interval schedule. This sort of thing is even more likely to occur when the fixed-interval has been a short one. In general, fixed-interval schedules produce "fixed-interval" curves, just as fixed-ratio schedules produce "fixed-ratio" curves of extinction. The important thing, however, is not the schedule, but the kind of behavior exhibited while the schedule is in operation.

Reinforcement Number and Resistance to Extinction. If you were to increase the number of fixed-interval reinforcements of a response, would you increase the difficulty of extinguishing that response ? From what you know about everyday human behavior, your answer to this question will probably be Yes, although you might add that the effect would depend upon the number of reinforcements that had already been given. And your answer would probably be the same if you recalled the experiment by Williams (K & S, pp. 72-73) in which he related the number of responses in extinction to the number of regular reinforcements given during training.

Actual research on the fixed-interval case is limited, at the moment, to one exploratory experiment. Wilson (1954) gave different numbers of reinforcements to 6 groups of white rats, all working on a 2-minute fixed-interval schedule. The actual numbers were 15, 50, 85, 240, and 500. When each group had received its appropriate number, the response was extinguished for 50 minutes a day throughout a 5-day period. The results are shown in Figure 6. The general trend is clear: the greater the number of reinforcements the greater the resistance to extinction. A straight line is suggested as the best picture of the relation between the two variables. You can see, however, that some of the data-points do not fall very close to this line, hinting that more work should be done at the small-number end of the scale. Also, points might well be added at the large-number end, to see if the effect would continue to be the same.

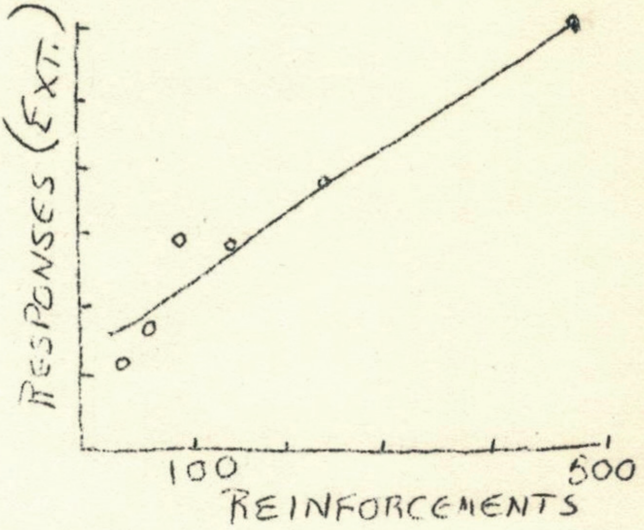

FIG. 6

Fixed-Ratio Reinforcement. With this schedule, a response is reinforced only after a certain number of responses have been emitted. The "ratio" refers to the ratio of unreinforced to reinforced responses. A ratio of 20:1 means that an organism is reinforced on every 21st response.

As noted in your text (K & S, pp. 91-98), there are several effects of long- continued fixed-ratio training. For one thing, a time discrimination develops, just as in the case of fixed-interval reinforcement. That is, the rate of response drops to a low, or zero, rate right after each reinforcement. This is presumably due to the fact that the organism on this schedule has never been reinforced for any response that occurred right after receiving a reinforcement. Saying that he 'tells time' means that he doesn't respond again until the effects of the reinforcement have largely disappeared. When responding does begin, there is usually an acceleration to a rate appreciably higher than would be found with a fixed-interval schedule. The acceleration may be rapid, as when the fixed ratio is small, or it may be slow, as when the fixed ratio is large.

[**p. 5**] The situation is much like that in which short or long fixed intervals are employed. The gradual change, during the running off of a large ratio (or interval), probably means that the organism's own behavior can furnish cues which tell him when reward is near. Besides the cue for non-response, arising from the reinforcement, he seems to have a cue for response that is based upon responding. We shall return to this a little later.

The extinction curve following fixed-ratio reinforcement is usually steeper, at the start, than the one following fixed-interval reinforcement. (As noted earlier, this initial rate may reflect the rate that was shown under reinforcement conditions.) Also, this rate usually comes to a fairly abrupt stop, as shown in your text (K & S, p. 98, Fig. 27). One shouldn't assume, however, that the responding is at an end when the first break occurs. An organism may, after a long period of no responding, start in again, and give a run that nearly equals the first one in size. And still further runs, of lessening length, may follow before extinction is complete. In analyzing this situation, you might say that when the animal is not responding there is very little reason for him to begin, since he has almost never been reinforced for starting off a run. On the other hand, when he does respond, there is more reason for him to continue, since he has often been reinforced for a response that followed closely upon one or more other responses. And, in those cases where his reward comes only after a long run of responses, still another factor may enter. The nearer he gets to the right number of responses the more he will provide of those cues from his own behavior that were present when he did receive reinforcement. Hence the increase in speed as he nears the end of his run.

In thinking about fixed-ratio reinforcement, or using the schedule in the laboratory, a number of questions may come up. For example, you might ask yourself this: Can a fixed-ratio schedule be introduced at the very beginning of training, or must it be introduced gradually by way of either regular or fixed- interval reinforcement ? Could you start right off with a 10:1, a 20:1, a 30:1, or a 40:1 ratio, without any prior strengthening of the response?

The answer to this question is easy. You may start from scratch with a fixed-ratio schedule provided that you do not make the ratio too large. A white rat, a pigeon, or a college student, will usually repeat the response for a few times at least. Now, if the second reinforcement comes along before the effect of the first is gone, the rate of responding will immediately pick up. As further reinforcements come along, the extinction effect will be barely noticeable, and a typical fixed-ratio mode of behavior will develop. An indication of the practical limits, in the case of the white rat, will soon be presented.

In actual practice, a fixed-ratio schedule is usually preceded by a little priming with regular reinforcement. (Unless you want to call regular reinforcement a fixed ratio of 0:1 !) Such priming is seen in the following study by Boren (1953), which aimed to answer the following questions: (1) Will different fixed ratios produce different rates of response; and (2) Will different fixed ratios produce different amounts of resistance to extinction ?

In answering the first of these questions, Boren used 5 groups of white rats, each on a different schedule of fixed-ratio reinforcement. At the outset of the experiment, however, he gave 20 regular reinforcements each to the members of all five groups. Then he gave his Group 1 animals 540 <u>more</u> reinforcements on the regular schedule. The animals of Group 2 were also given 540 more reinforcements, but on a 2:1 ratio --- every third response reinforced. Groups 3, 4, and 5 each [**p. 6**] received 40 reinforcements on the 2:1 schedule (to prime them for still higher ratios), after which Group 3 was put on a 5:1 schedule, Group 4 was put on a 14:1 schedule, and Group 5 on a 20:1 schedule.

Figure 7 shows how the over-all rates of response for these 5 groups was affected by the size of the ratio. An increase in size of the ratio, within the limits of this study, brought an increase in rate of response. The curve is negatively accelerated, however, and little further increase in rate would be expected if still higher ratios were employed --- at least for animals moved abruptly from a 2:1 ratio to the higher ones.

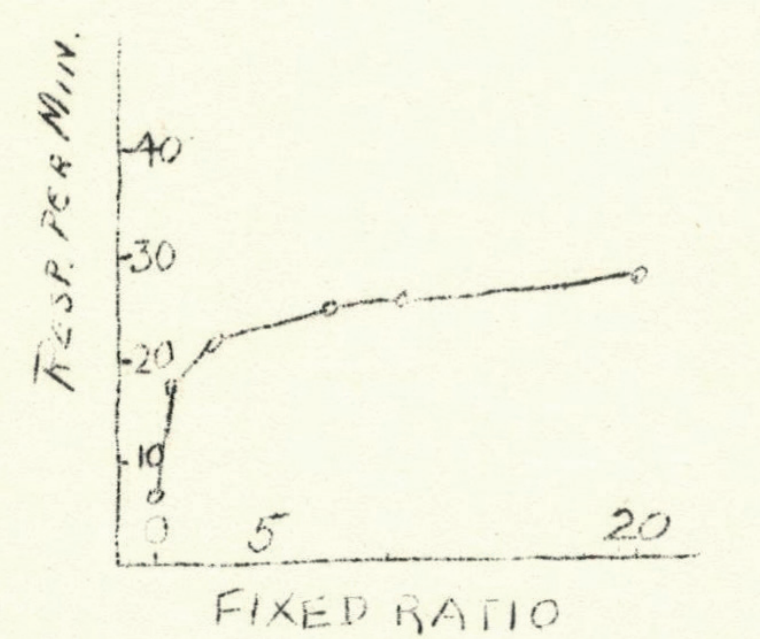

FIG. 7

After their fixed-ratio training, Boren's rats were given 5 daily sessions of extinction, to answer his second question. Figure 8 shows how the average number of responses made in 5 hours of extinction is affected by the different ratio schedules of training. Here is another curve with a (mild) negative acceleration. You can see that it is in sharp contrast with the one obtained by Wilson with different fixed intervals. Interval and ratio schedules clearly produce different effects, on both rate and resistance to extinction.

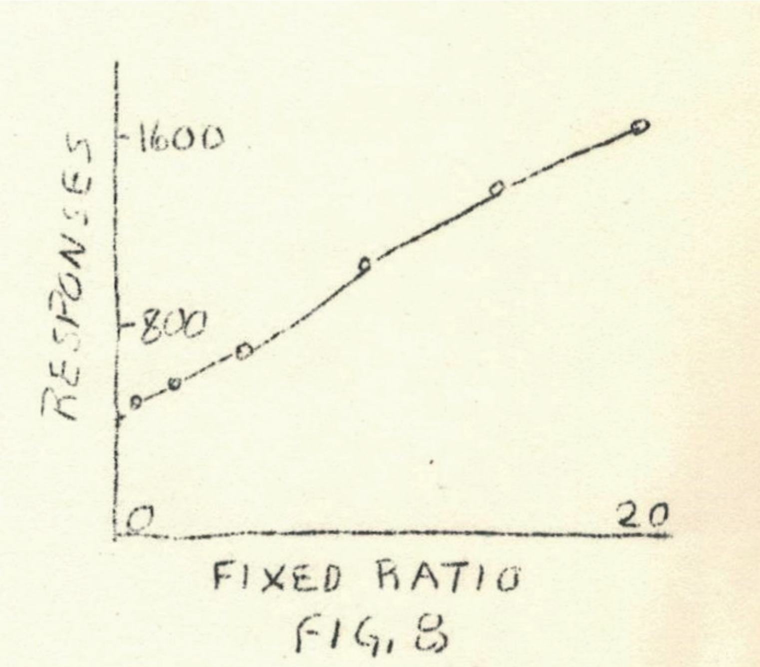

FIG. 8

Fixed-ratio schedules may also be used when the reinforcement consists in the removal of an aversive stimulus, or negative reinforcer, such as an intense light, an intense sound, or an electric shock. Experiments on human subjects are rare in this area, as you might expect when you consider the kind of stimulation that would have to be endured. At the animal level, however, some work has already been done. Kaplan (1956), for example, has recently completed a study with white rats, in which he showed that the rate of response tends to increase as the size of the ratio increases. He used light-removal as a reward for the rat's response of pushing down on a pedal in the center of a small chamber. He began his experiment with 3 hours of regular-reinforcement schedule [sic], during which each response in the presence of a 25-watt light was reinforced by light-removal for 1 minute. Then, he put each of 3 subjects through a series of gradually increasing fixed ratios, from 1:1 to 15:1. In order to stabilize the response at the four ratios plotted in Figure 9, there were at least 5 experimental sessions of 30 minutes each at that point. This average rate is also what Kaplan called the "terminal" rate. That is, it is the rate for the last responses of each ratio run, just before the animal turned off the light. At this point the rate was usually at its peak. (The open circles in these plots are based on second tests at the same ratio, run off after the scale had already been covered from 1:1 to 15:1.)

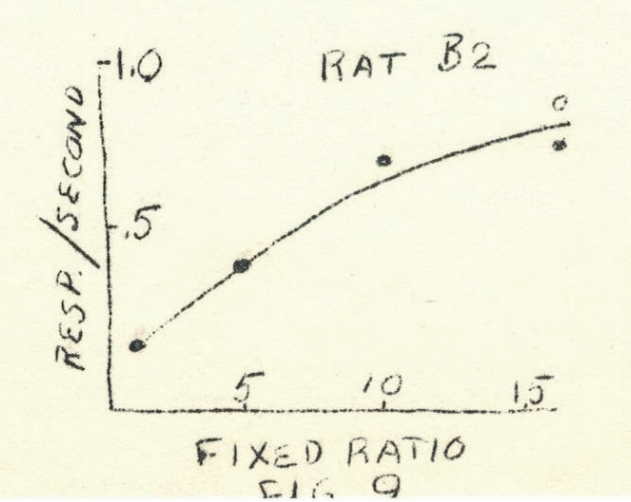

In this study, as in Boren's, there is apparently an increase in response rate that goes with an increase in fixed ratio. But, in both experiments, there is another factor to be considered in connection with the way in which one determines a rate. The factor is this: the higher the ratio the longer it takes an animal to get started just after receiving a reinforcement. (This is probably also true of starting timed after different fixed intervals.) The picture of the way ratio influences rate will be different, depending on whether we leave out or include those starting-times or "breaks" right after reinforcement. When Boren subtracted these times in estimating his rates, he found that the increase in rate with increase in ratio was much less dramatic; and when Kaplan treated his data in the same way that Boren did, he found an even greater leveling effect. (Kaplan's "terminal" rates, of course, left out these starting-times.)

Other Fixed-Interval and Fixed-Ratio Studies. To indicate the variety of experiments in which such schedules of reinforcement may be employed, a few [**p. 7**] additional studies may be mentioned here. In one of these, Sidman and Stebbins (1954) investigated the effect of liquid-satiation upon bar pressing, using fixed-ratio schedules with 4 rats, 2 cats, and a monkey. The fixed ratio for the rats and the monkey was 20:1; for the cats it was 25:1. The satiation effect was studied, in one part of their experiment, by the simple procedure of continuing the ratio schedules on a given day until each subject had received all the liquid (water, milk, or sugar-water) he could drink. Typical results are shown in Figure 10. You can see that the principal change during the entire period of satiation was an increase in the length of break in response following reinforcement. The rate while responding was practically unaltered in each subject throughout the period of observation.

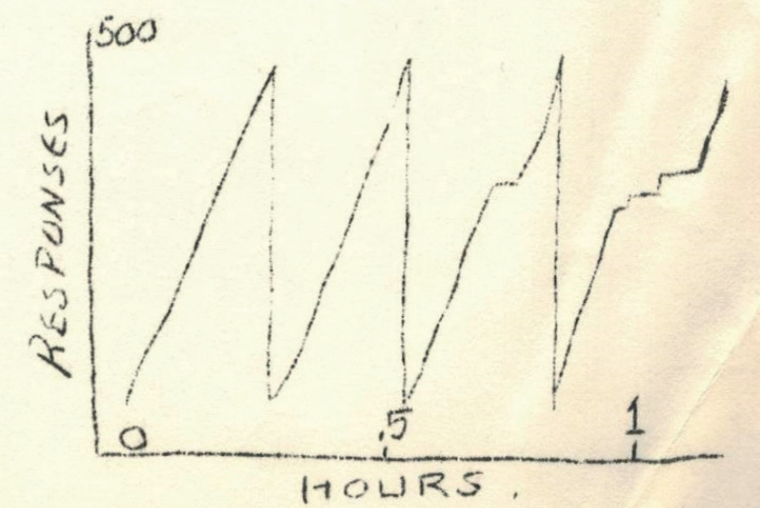

FIG. 10

Human subjects provide fixed-ratio curves that are often indistinguishable from those produced by monkeys, dogs, rats, and other subjects. A dramatic example, although not necessarily typical, is shown in Figure 11. These records were obtained from a 32-year-old man, diagnosed as schizophrenic, who had been for 16 years a patient in a Massachusetts mental hospital. The record at the left shows the rate at which he pulled the plunger of a vending machine for a candy reward. The one at the right is for a button-pressing response that produced a nude-art picture on the wall of his experimental chamber. Both responses were rewarded on a 19:1 fixed-ratio schedule and this was the 35th 1-hour session in each of the experimental situations for this subject. You will observe, in both records, the characteristic breaks that follow reinforcement, although these breaks are very uneven in length. We are told that, during the long pauses, the patient usually engaged in the kind of activity that marked him as "psychotic" --- dressing and undressing, 'picking' at himself, and so on. For this patient, an increase in plunger-pulling and button-pressing meant a decrease in the amount of his abnormal conduct.

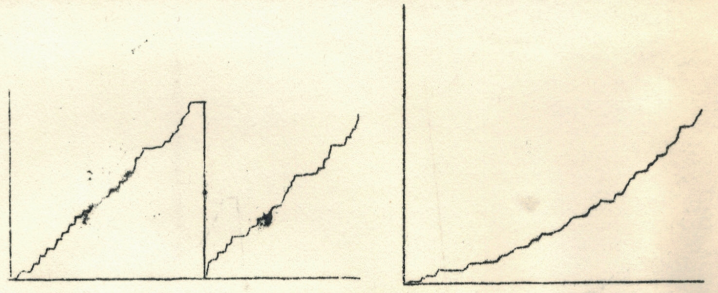

FIGURE 11

Two further Instances of fixed-ratio responding may be mentioned here because they throw light upon an ability of animals that has not often been reported. That is, they show how an animal's response may depend upon an inner 'clock,' and they suggest the sort of clock that this must be.

The first example is from an experiment by Ferster and Skinner (1952), in which two different fixed-ratio schedules were used for a single animal in each experimental session. The subject, a pigeon, was rewarded with food for pecking at a disc, on either a 50:1 or a 250:1 fixed ratio. On one occasion, a reinforcement might follow a run of 50 responses; on another, it might follow a run of 250; and the arrangement of ratios was such that the bird had no way of telling which was coming next.

After long training with this combination of schedules, the pigeon developed a special, and quite revealing, mode of response. Throughout each experimental session, his rate of pecking had a step-wise appearance. The steps were made by short pauses after runs of 50, or somewhat more than 50, responses. When the ratio was 250:1, the bird would not respond steadily throughout the entire run, but would halt after 60 or 70 pecks, as if it were "time for food." Then would come another short run, another stop, and so on, until the 250 responses had been made and the food appeared. The bird acted, you might say, as if he had a crude sort of counter or inner clock by which it regulated its pecking behavior. (This was especially obvious during extinction, when the responses came for some time in runs of approximately 50 each.) The cue for stopping a run, in each case, was apparently provided by the emission of a [**p. 8**] certain number of responses, rather than by some change in the outside situation. There was, you might say, a "feedback" from the animal's own behavior that led him to stop at a given point.

This is more clearly shown in the performance of seven rats in an experiment by Berryman and Mechner (1956). During training, each animal had two levers before him, either of which might provide water reinforcement when pressed at certain times. The arrangements for reinforcement were as follows: when the rat was first placed in the experimental chamber, 12 pressings on one of the levers would bring the reward, but the animal had no way of knowing which lever was the correct one. Suppose he picked the right-hand lever, and received a reinforcement after 12 pressings. The <u>next</u> reinforcement might come from 12 more responses on the same lever, or it might come from 12 responses on the other one. Suppose he picked the wrong lever. How many responses would he waste on this one before shifting to the correct one ? How far would he have to go before he 'saw' that he was wrong ?

In this experiment, there was, at the start, a lot of confusion in the subject's behavior. Sometimes the rats shifted frequently from bar to bar, and sometimes they responded to the wrong bar for a long time. In a few hours of training, however, they adjusted nicely to the situation. Beginning on one bar, a rat might respond steadily until reinforcement came; or, if that were the wrong bar, until he had responded more than 12 times. Beyond the 12th response on the wrong bar, he would soon begin to slow down and, by the 17th or 18th response, he was ready to shift to the correct one (where he would get his reinforcement after 12 responses). On the first day of training, the average amount of overshoot (responses beyond 12) for all the animals amounted to 28 responses. On the 10th day of training, this had dropped to 6 responses, with no individual performance deviating from this average by more than 1 or 2 responses. Again it is quite clear that the emission of a certain number of responses can become the cue for a change in behavior.

Variable-Interval Reinforcement. Both fixed-interval and fixed-ratio reinforcement schedules may be called *periodic*. This is because the fixed- ratio schedule, although dealing with *number* of responses, will also provide reinforcements at fairly regular intervals. The schedules now to be treated are *aperiodic*, since they provide reinforcements after uneven lengths of time, but in one case we speak of variable-interval reinforcement, and in the other of variable-ratio. Let's begin with the former.

In variable-interval reinforcement, the time between reinforcements is varied about some average value. This value may never be known in the case of most everyday human behavior, but it is carefully specified and arranged for in experimental studies. Thus, a 5-minute variable-interval schedule might involve intervals from 0 to 10 minutes, at 1-minute steps of difference. (A 0-minute interval refers to the case in which one reinforcement is obtained immediately after another, with no delay imposed between them by the experimenter.) These intervals would then be "randomized" in their order of "succession," by means of some clearly stated procedure that might be followed by anyone who desired to repeat the experiment. Several such procedures have been used, but a detailed account of them need not be given here.

The general effect of a variable-interval schedule (see K & S, pp. 98- 101), after prolonged training, is to produce a steady rate of response that shows itself as a straight line in the usual cumulative response curve. The [**p. 9**] rate will be high or low depending upon the size of the average interval and the range of the intervals employed. There will be no breaks or 'scallops' in the rate curve like those that are found when rewards are periodic. The organism does not pause after reinforcement, because he is just as likely to be reinforced for responding then as he is at some later time.

The way in which a variable-interval schedule will maintain a steady rate of response in a number of different organisms is illustrated in Figures 12 to 15, inclusive, which are shown below. Figure 12 contains rate curves for a monkey in an experiment by Conrad and Sidman (1956). This animal was reinforced with small sips of sugar-water, on a schedule that permitted one reinforcement every 37 seconds, on the average, with a range of intervals from 6 to 69 seconds. The lower curves, for a 60-percent sucrose solution, show [stricken s in original] the effect of satiation during the last part of the experimental session. (The right-hand curve in each pair is merely a continuation of the left-hand curve, which drops to the baseline after reaching its maximal height on the recording paper.)

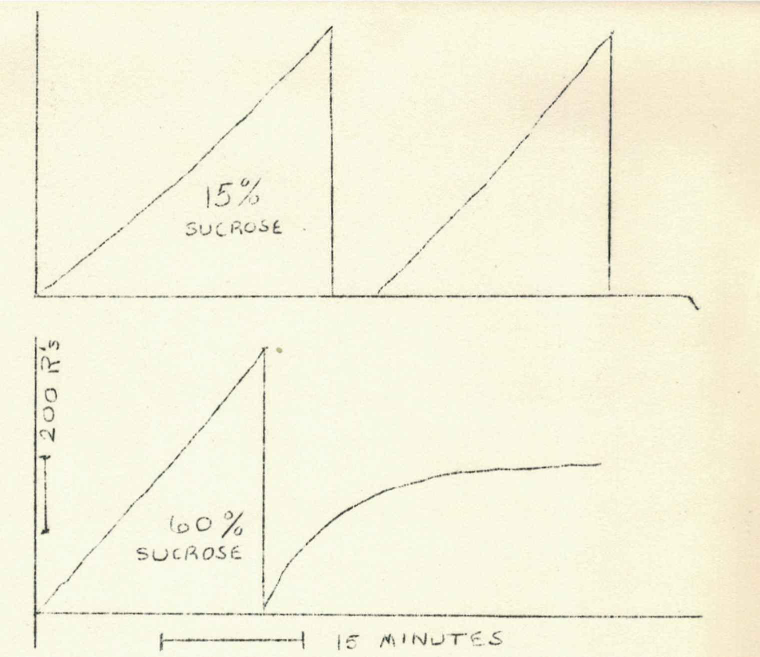

FIG. 12

Figure 13 is the record of a pigeon, working for food on a 3-minute variable-interval schedule. In this case, the intervals ranged from a few seconds to 6 minutes. The record covers a period of more than 2 hours, in which the bird made about 7,000 pecking responses. On other occasions, this same bird responded steadily for 15 hours, totaling 30,000 responses. Except for one pause of about 1 minute, the animal never stopped between responses for more than 15 seconds during the entire stretch of time. The reporters of this study (Ferster and Skinner) would seem to be guilty of understatement when they assert that "the control exercised by a schedule of this sort may be very great."

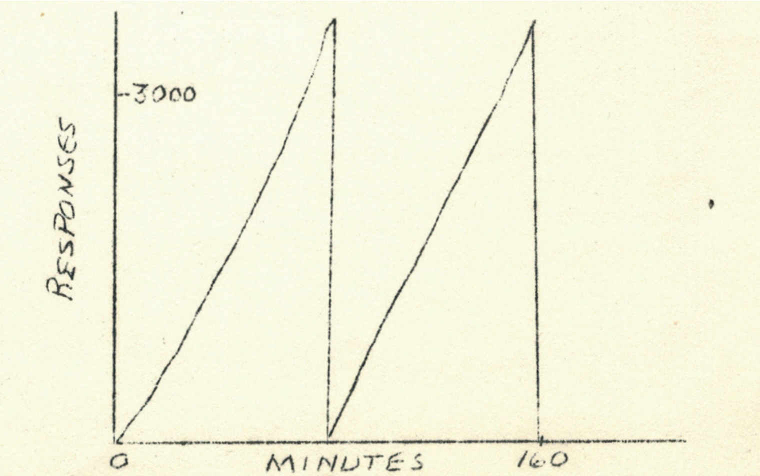

FIG. 13

With respect to Figure 14, the two upper records are for normal human subjects, attendants in a mental hospital, reinforced with nickels on a 1-minute variable-interval schedule. This amount of difference between individuals is not uncommon, either at the animal or human level. Another sample of individual differences is seen in the middle records of Figure 14. These are rate curves for ”normal” beagle dogs, working for food on the same schedule as that employed with the hospital attendants. The lower curves in this Figure are for chronic psychotic patients, diagnosed as cases of paranoid schizophrenia, reinforced on the same 1-minute Schedule but with candy or cigarettes. Of these records and several others it was noted that “the one-minute variable-interval schedules yielded performances which resembled the performances of ’normal’ attendants, dogs, rats, and pigeons on the same schedule. The rates of response were of intermediate value and any ‘breaks’ that occurred were not correlated with the time of reinforcement.” There were more of those breaks, however, in the patients’ curves than in those for other subjects, and day-to-day variability in rate was far greater.

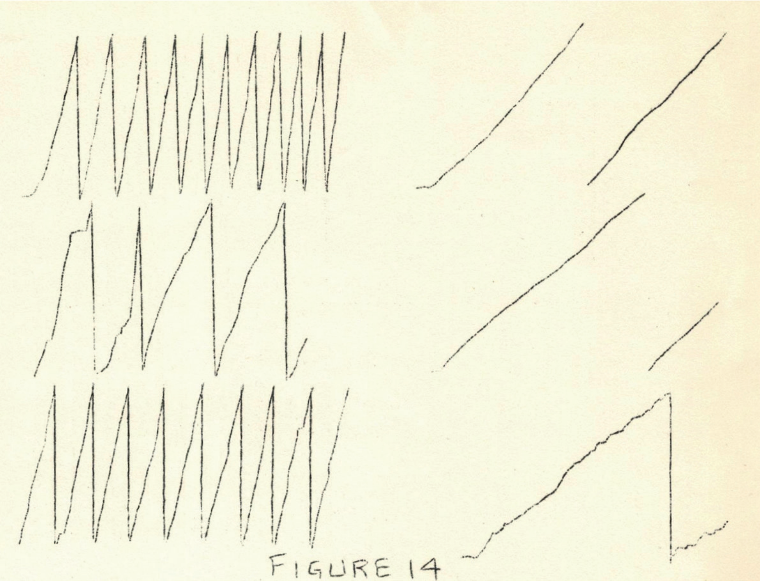

FIGURE 14

Variable-Ratio Reinforcement. This schedule is based upon random variation in the number of responses that an organism must make before being rewarded. Very high response rates are built up with such a program, and they are built up very quickly when the average ratio is small. The fully- established response curve looks a lot like the one obtained with a variable-interval schedule, in that it shows no breaks or pauses following a reward. But the ratio schedule leads as a rule to higher response rates than does the interval schedule. This is because a response that follows any break in an [**p. 10**] interval schedule has a greater likelihood of getting rewarded, whereas a pause during a ratio schedule never increases the chances of reward for the next response. Or, it might be said this way: With an interval schedule, reinforcement of response-after-pause will strengthen 'pausing' behavior as well as the response itself. And this effect may operate to slow down the organism's responding to a rate that is well below that of a fixed-ratio program.

A notion of the high rates that can be generated with the ratio schedule is conveyed by the curve of Figure 15. This curve is for a pigeon on a variable-ratio schedule in which the average number of responses required for reinforcement was 110, but with a possible range from 0 ('the very next' response) to 500. This bird worked at a rate of about 12,000 responses per hour --- which is 200 responses per minute, or a little more than 3 responses per second. Is it any wonder that the makers of the world's gambling devices always build in a variable-ratio "pay-off" ?

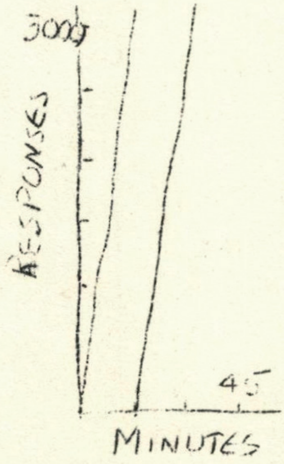

FIG. 15

Extinction after Variable-Ratio and Variable-Interval Reinforcement. There is nothing very new to be added here. After performing under a variable-interval or variable-ratio schedule for a long period of time, an organism's extinction curve is essentially a continuation of the same. That is, the curve begins at the same rate that existed under reinforcement. This may go on for quite a while before the subject shows any sign of let-up in response rate. Look at Figure 16, for example. This is an extinction curve for the same pigeon that made the variable-ratio record shown in Figure 15. The curve begins with a run of about 7,500 responses at almost the identical rate that had been maintained under reinforcement conditions. Extinction was probably not yet complete in 3 hours, during which at least 15,000 responses in all had been emitted. Other birds, on similar training schedules, have been known to respond as many as 35,000 times before coming to a stop. If you did not know the reinforcement history of such behavior, you might easily be led to call it pathological.

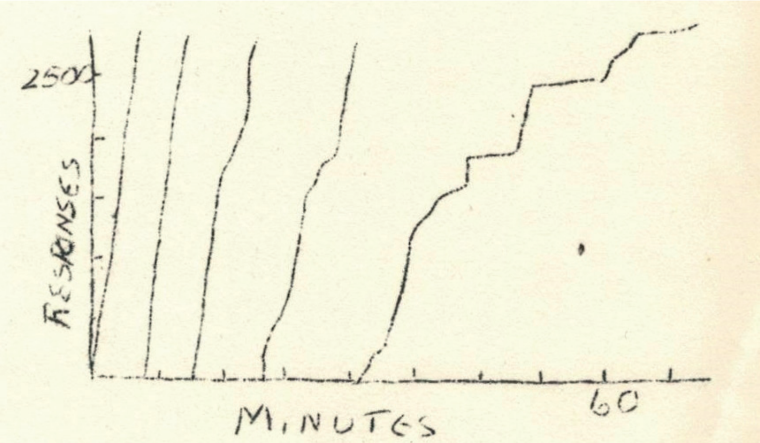

FIG 16

Conditioning Extreme Rates. It is possible to build up even higher rates of response than the ones just mentioned. This may be done in more than one way, but the simplest method, perhaps, is one in which responses are reinforced only when they are coming out at a rate that is above-average for the schedule then in effect. For example, suppose you had, as your experimental subject, a white rat, pressing a lever at a rate of approximately 30 responses per minute on a 1-minute variable-interval schedule of reward. To increase this rate, you would first estimate the average time elapsing between any pair of successive responses being emitted by your rat. The value, in this case, would be slightly less than 2 seconds, since several seconds of eating time must be subtracted from each minute. Some pairs of responses will, of course, be separated by more than 2 seconds, and some by less. Your next step would be to reinforce responses on approximately the same variable-interval schedule as before, but with this difference: you would now demand that the response to-be-reinforced came within 2 seconds of the response that preceded it. If this did not happen, you would withhold reinforcement until a pair of responses did occur within that period. (This is not an easy thing to do in practice , unless reinforcements are delivered automatically when the time requirements are met.) This selective reinforcement of responses that are close on the heels of other responses will soon decrease the average between-response period, hence increase the number of responses per minute. The new average may then be [**p. 11**] used to set up a new requirement and the rate may be raised again, after which even further advances may be made. The limiting rates, for white rats, under such circumstances, are at present unknown. For human beings, the highest individual rates of tapping a key may range between 8 and 13 per second, although the 'preferred' rate is much slower (1.5 to 5 taps per second.)

"Low-rate" responding, in both human and animal subjects, may be established with somewhat less effort than high-rate responding. The procedure is simple. You reinforce responses only when they are farther apart in time than is 'normal' under some schedule of reinforcement. The schedule may be either regular or intermittent, and it need not be of long standing. You may even begin low-rate training after reinforcing but one response, provided that you have a rough idea of the organism's probable rate of responding when on a regular reinforcement schedule, and that you do not require too much spacing of responses at the outset. Your spacing will then be great enough to exceed the regular-reinforcement distance, and not so great as to permit complete extinction between reinforcements.

In an experiment by Wilson and Keller (1953), three white rats were given 30 regular reinforcements each for lever-pressing on the first day of experimentation. For the next 5 days, reinforcement was provided for every response that followed another by 10 seconds or more. Thereafter, the time requirement was systematically increased, in 5-second steps, until each subject had been given 6 or 7 hours of training with minimal intervals of 15, 20, 25, and 30 seconds. The effect of this upon the average bar-pressing rate for the three animals is seen in Figure 17. There you can see that the rate of responding dropped from a little more than 6 responses per minute, with a 10-second requirement, to approximately 3 responses per minute, with a 30-second requirement. Although a lower rate than this might have been reached by demanding still more delay between responses, a limiting value would ultimately be found. The number of reinforcements obtained by these animals dropped from a daily average of about 80 to one of about 14 as the delay intervals were lengthened from 10 to 30 seconds. A further lengthening would mean fewer reinforcements per hour, and eventually the number would be less than enough to keep up the strength of the response.

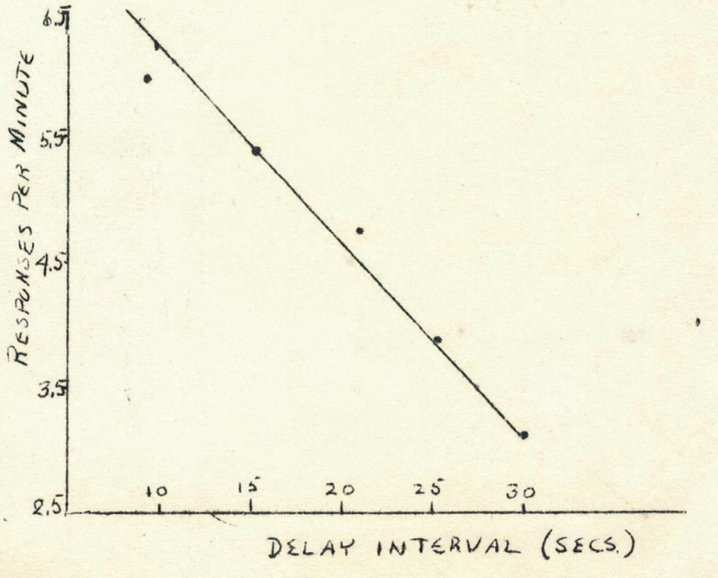

FIG. 17

References and Notes on References to Material in the *Schedules* supplement

(The original "Schedules of reinforcement" supplement did not have a reference list. These notes list the names of all the experimenters cited in the supplement, along with the article we believe to be the cited reference.)

Berryman, R., & Mechner, F. (1956). A reference to a published article by these two authors for this year was not found. The reference perhaps is to research that either led to or later appeared in Mechner's 1957 doctoral dissertation at Columbia University, published as Mechner, F. (1958). Probability relations within response sequences under ratio reinforcement. *Journal of the Experimental Analysis of Behavior, 1 (2),* 109-121. https://doi.org/10.1901/jeab.1958.1-109 Berryman's suggestions and encouragement are acknowledged in a footnote to the latter work.

Boren, J. J. (1953). Response rate and resistance to extinction as a function of the fixed ratio. Unpublished doctoral dissertation, Columbia University. This research later was published as Boren, J. J. (1961). Response rate and resistance to extinction as a function of the fixed ratio. *Journal of Experimental Psychology, 61*(4), 304-308. https://doi.org/10.1037/h0040208

Conrad, D. G., & Sidman, M. (1956). Sucrose concentration as reinforcement for lever pressing by monkeys. *Psychological Reports,* 2, 381-384. https://doi.org/10.2466/pr0.1956.2.3.381

Ferster, C. B., & Skinner, B. F. (1952). A reference to a published article by these two authors for this year was not found. The reference may be to a research grant report of some kind or perhaps to an early partial draft of material that would be published by Ferster and Skinnner in *Schedules of reinforcement* (1957). Appleton-Century-Crofts.

Kaplan, M. (1952). The effects of noxious stimulus intensity and duration during intermittent reinforcement of escape behavior. *Journal of Comparative and Physiological Psychology,* 45(6), 538-549. https://doi.org/10.1037/h0055989

Kaplan, M. (1956). The maintenance of escape behavior under fixed-ratio reinforcement. *Journal of Comparative and Physiological Psychology,* 49(2), 153-157. https://doi.org/10.1037/h0048735

Keller, F. S., & Schoenfeld, W. N. (1950). *Principles of psychology*. Appleton Century Crofts.

Sidman, M. , & Stebbins (1954). Satiation effects under fixed-ratio schedules of reinforcement. *Journal of Comparative and Physiological Psychology,* 47(2), 114-116. https://doi.org/10.1037/h0054127

Wilson, M. P. (1954). Periodic reinforcement interval and number of periodic reinforcements as parameters of response strength. *Journal of Comparative and Physiological Psychology,* 47(1), 51-56. https://doi.org/10.1037/h0057224

Wilson, M. P., & Keller F. S. (1953). On the selective reinforcement of spaced responses. *Journal of Comparative and Physiological Psychology,* 46(3), 190-193. https://doi.org/10.1037/h0057705