

"TRANSITIVE INFERENCE" IN MULTIPLE
CONDITIONAL DISCRIMINATIONS

J. J. HIGA AND J. E. R. STADDON

DUKE UNIVERSITY

We used multiple conditional discriminations to study the inferential abilities of pigeons. Using a five-term stimulus series, pigeons were trained to respond differentially to four overlapping pairs of concurrently presented stimuli: A+B-, B+C-, C+D-, and D+E-, where plus and minus indicate the stimulus associated with reinforcement and extinction, respectively. *Transitive inference* in such situations has been defined as a preference for Stimulus B over Stimulus D in a transfer test. We measured this and other untrained preferences (A vs. C, A vs. D, B vs. E, etc.) during nonreinforced test trials. In three experiments using a novel, rapid training procedure (termed *autorun*), we attempted to identify the necessary and sufficient conditions for transitive inference. We used two versions of autorun: response-based, in which the subject was repeatedly presented with the least well-discriminated stimulus pair; and time-based, in which the subject was repeatedly presented with the least-experienced stimulus pair. In Experiment 1, using response-based autorun, we showed that subjects learned the four stimulus pairs faster than, but at a level comparable to, a previous study on transitive inference in pigeons (Fersen, Wynne, Delius, & Staddon, 1991), but our animals failed to show transitive inference. Experiments 2 and 3 compared time- and response-based autorun. Discrimination performance was maintained, but transitive inference was observed only on the second exposure to the response-based procedure. These results show that inferential behavior in pigeons is not a reliable concomitant of good performance on a series of overlapping discriminations. The necessary and sufficient conditions for transitive inference in pigeons remain to be fully defined.

Key words: transitive inference, conditional discrimination, symbolic-distance effect, end-anchor effect, key pecking, pigeons

Ethologists are interested in animals for their own sakes; psychologists are for the most part interested in animals for what they can tell us about human beings. Despite the heroic efforts of several groups who have worked intensively with chimpanzees and other great apes, as well as parrots, the grasp of these creatures of symbolic communication, even of the most rudimentary sort, appears still to be tenuous. They speak—if they speak at all—haltingly and with the simplest sentence structure, and they use symbolic reasoning not at all. The apparently vast gap between the symbolic abilities of people and those of nonhuman animals has therefore led some cognitive psychologists to abandon the comparative method and treat humans

as animals with no forebears, or at least as so greatly evolved from those forebears that little or nothing is to be learned from the comparison.

To comparative psychologists this disconnection between humans and other animals seems fundamentally mistaken. Yet we have an obligation to show the connection, to show how the complex symbolic abilities of human beings are built upon rudimentary ancestors. One way to do this is to study the conditions under which behavior that in humans would be symbolic can be brought about, by nonverbal means, in other animals. This is the reason for continuing interest in the so-called *transitive inference* problem. This is a problem that forms the basis for Aristotelian logic, yet, with suitable nonverbal training, it can apparently be solved by young children, by chimpanzees and lesser anthropoids (Bryant & Trabasso, 1971; Gillan, 1981; McGonigle & Chalmers, 1977), and even by pigeons (Fersen, Wynne, Delius, & Staddon, 1990, 1991).

What are the necessary and sufficient conditions for the transitive inference effect in the pigeon? The present paper describes a series of experiments that take us some way towards

Research support was provided by grants to Duke University from the National Science Foundation and the National Institute of Mental Health. We thank Talitha Robinson for conducting pilot experiments, and Nancy Innis, George King, Armando Machado, and Clive Wynne for comments on an earlier version of the paper. J.E.R.S. also thanks the Alexander von Humboldt-Stiftung for support. Correspondence and reprint requests should be sent to either author at the Department of Psychology: Experimental, Duke University, Box 90086, 229 Soci/Psych Bldg., W9, Durham, North Carolina 27708-0086.

answering this question. We first describe the task and briefly summarize previous results. We then give the rationale for the experiments reported here, which differed in a couple of important ways from earlier ones. Finally, we describe the results from our experiments, which sometimes found the transitive inference effect and sometimes failed to find it.

Verbally presented, the transitive inference task is of the following form: Fred is taller than Joe (Fred > Joe); Joe is taller than Suzy (Joe > Suzy); who is taller, Fred or Suzy? (Fred ? Suzy). Responding "Fred" is taken as evidence of an inference process because the test requires that the subject go beyond the information he or she has been explicitly given—the answer to Fred ? Suzy must be deduced from the relations Fred > Joe and Joe > Suzy. Pigeons and small children are not receptive to instructions and tests like these. Some kind of nonverbal paradigm is needed. Peter Bryant devised such a task and tried it out with children. Bryant and Trabasso (1971) presented children with pairs of sticks of different lengths and colors, but each child could see only the tops of the sticks; the bottoms were hidden behind a barrier. On each trial the child was shown pairs of sticks: A and B, B and C, and so on, and asked which stick was longer. Each child was told: A is longer than B, B is longer than C, and so on, for adjacent pairs. Bryant and Trabasso found that even 4-year-olds could infer that A > C after being told A > B and B > C.

The Bryant and Trabasso (1971) task is less verbal than a paper and pencil test, but language is still used to convey information about the stimuli. Language must be excluded entirely if the task is to be applied to nonhuman animals. McGonigle and Chalmers (1977, 1992, with squirrel monkeys), Gillan (1981, with chimpanzees), and Fersen et al. (1990, 1991, with pigeons) have looked for, and found evidence of, transitive inference using nonverbal forms of the task. Although the particular stimuli and reinforcers differ, the training procedure is essentially the same: Reinforcement, rather than verbal confirmation, is used to give feedback about relationships among stimuli. More specifically, relations are expressed through conditional discriminations in which pairs of stimuli are presented successively, and responses to one stimulus are reinforced while responses to the other stimulus are either not reinforced or are followed by blackout or re-

moval of the stimulus pair. For example, the relations A > B and B > C are conveyed by training subjects on conditional discriminations of A+B- and B+C-: In the presence of A+B-, responses to A are reinforced and responses to B are not reinforced; in the presence of B+C-, responses to B are now reinforced and those to C are not (the plus indicates the stimulus that is associated with reinforcement and the minus is the one that is not associated with reinforcement)¹.

Most nonverbal forms of the transitive inference task have two additional features that are necessary so that noninferential explanations for transitive inference and related effects can be eliminated. First, the stimuli are selected so that they cannot be ordered along any physical dimension such as size, orientation, or color. Second, the series is at least five stimuli long, yielding four overlapping stimulus pairs: A+B-, B+C-, C+D-, and D+E-, so that responding to Stimulus A is always reinforced and responding to Stimulus E is never reinforced, and there are three stimuli (B, C, and D) to which responses are equally often reinforced and nonreinforced. Under these conditions, when a subject is tested with nonreinforced trials of Pair B ? D (a combination that it has never seen before), preference for Stimulus B is taken as compelling evidence for an inferential process—a process that some researchers assume produces a linear (ranked) representation of the stimuli that is used by the subject to deduce that Stimulus B is "better than" Stimulus D (e.g., McGonigle & Chal-

¹ Notice that conditional discriminations are also used in the stimulus equivalence paradigm in which classes of stimuli are established and tests of symmetry, reflexivity, and transitivity between classes determine whether an equivalence has been trained (e.g., Sidman, Kirk, & Willson-Morris, 1985; Sidman et al., 1982). Although both paradigms use conditional discriminations to establish relations among stimuli and test for transitivity, they are conceptually and procedurally different. For instance, to our knowledge, the stimulus equivalence procedure requires establishing at least three different classes of stimuli, whereas a reinterpretation and redescription of the transitive inference task in stimulus equivalence terminology implies only two, not necessarily equivalent, classes (e.g., A+B- and B+C- in one class that somehow maps onto another class with elements A, B, and C). Furthermore, tests for transitivity in the stimulus equivalence task are *between* classes, whereas the test is *within* a class for transitive inference. Given differences like these, we consider the two paradigms as addressing separate research problems for now—albeit any comprehensive explanation for how humans and nonhumans process and use relational information should explain the effects found in both areas.

mers, 1977, 1984; Trabasso & Riley, 1975). Noninferential explanations are not considered because (a) B and D cannot be ranked according to physical appearance, (b) direct information about B and D has not been given during training, and (c) the reinforcement value or associative strength of B and D has been equalized during training.

However, we argue that by considering the task as a measure of inference based on the assumed ordering of stimuli (without providing a mechanism for how the ranking develops), rather than as just another kind of choice or preference task, there is a danger of minimizing the importance of procedural details. There is much less danger of minimizing the importance of things like training schedule, how preference is to be measured (percentage correct, reaction time, response rate, etc.), and what effect the particular training history might be expected to have on the simple tendency to peck on the five different stimuli, if we regard the task as just another kind of choice procedure. We return to this theme later.

The dependent variables in most of the transitive inference studies of nonhumans are error rates (percentage correct) and (less often) reaction times. There are two notable results that accompany successful attempts to produce the transitive inference effect in animals: the *symmetrical end-anchor* effect and the *symbolic-distance* effect, which form the basis of what inferential models of transitive inference also attempt to explain. The symmetrical end-anchor effect refers to a particular feature of discrimination performance on training pairs—that performance on the two end discriminations, A+B- and D+E-, is always much better than on the two middle ones, B+C- and C+D-. The symbolic-distance effect refers to a decrease in reaction time, or an increase in the proportion of correct responses, as the “distance” between stimulus pairs increases, where distance refers to the ordinal number of stimuli interceding the stimulus pair. To illustrate, shorter reaction times and higher percentages of correct responses are typically associated with Pair BD, which is separated by one stimulus (C), than with Pair BC, which is not separated by any stimuli (e.g., Fersen et al., 1991; McGonigle & Chalmers, 1992). Symbolic-distance effects have also been reported across other stimulus combinations such as AB, AC, AD, and AE (increasingly better performance from AB through AE) or DE,

CE, BE, and AE (increasingly better performance from DE through AE), but these findings are less critical because responding to Stimulus A is always reinforced and responding to Stimulus E is never reinforced.

EXPERIMENT 1: TRAINING OVERLAPPING DISCRIMINATIONS WITH THE AUTORUN PROCEDURE

Fersen et al. (1990, 1991) recently looked for transitive-inference-like behavior in pigeons. Using a five-term series, they trained pigeons to respond differentially to Pairs A+B-, B+C-, C+D-, and D+E-. The various stimulus pairs were presented randomly, and 30% of the trials were without reinforcement. With continued training, 4 of 6 subjects reached criterion and responded with at least 80% correct choices overall. These subjects were then shifted to a testing phase in which a novel stimulus pair (B and D) was presented during the nonreinforced trials. The pigeons chose Stimulus B over D with 87.5% “accuracy.” This preference for B over D was taken as evidence for transitive inference, because B and D were associated with reinforcement and nonreinforcement equally often. Furthermore, their pigeons showed 75% correct during BC (a symbolic-distance effect, because performance on BD was better than that on BC), and they performed better on training Pairs AB and DE than on BC and CD (a symmetrical end-anchor effect).

Although this procedure sufficed to produce transitive inference behavior in pigeons, there are some problems with it. First, many sessions were required for the animals to reach criterion. Because of the ambiguous status of the three middle stimuli (B, C, and D), the animals could not learn all four pairs concurrently. Training included randomized, blocked, then randomized presentations of stimulus pairs, and a total of 125 sessions was needed. After all this, still only 4 of 6 subjects reached the 80% criterion (Fersen et al., 1991). When the number of stimuli in the series was increased from five to seven, the number of training sessions increased to around 370.

Second, the sequence of stimulus pairs used during training was chosen in a relatively unconstrained fashion: The number of training trials on each pair was equal (not necessarily the number of correction trials; we address this

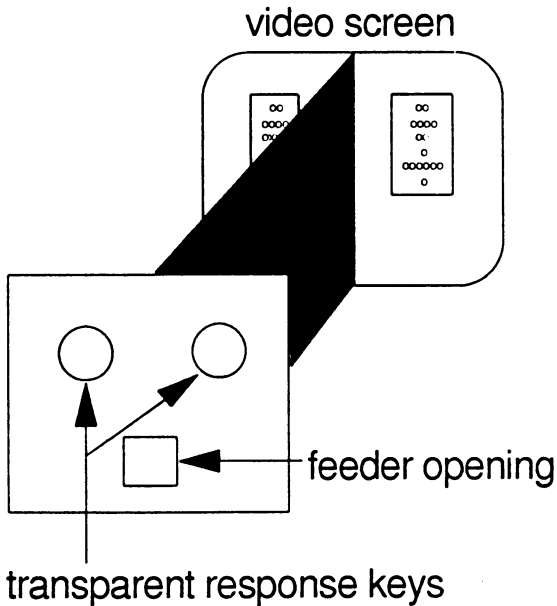


Fig. 1. Method of displaying two stimuli with a single video screen.

in a moment), but the order of presentation was decided on an ad hoc basis.

Finally, Fersen et al. (1991) used a correction procedure, which meant that although the probability of reward (hence the number of reinforced responses) was equal for Stimuli B, C, and D, the number of nonreinforced responses generally was not. Persistent incorrect responding caused the same stimulus pair to be presented repeatedly until the animal responded correctly. In addition, it is unclear how they handled nonreinforced responses during other, noncorrection, trials, for example those responses to Stimulus B during A+B- presentations, to Stimulus C during B+C-, and to Stimulus D during C+D-. Hence, the procedure made it unlikely that equally rewarded Stimuli B, C, and D would also be responded to equally often, or presented equally often or for the same total time.

These problems cannot all be solved simultaneously. In this first experiment, therefore, we addressed the first two: the need for rapid acquisition, and the need for a principled method of introducing new stimulus pairs. In this experiment we tested a procedure in which a subject's performance level is continuously monitored and an adaptive rule determines, trial by trial, which stimulus pair will be presented next. The rule was designed to ensure

that performance on all four discriminations would be maintained approximately equal throughout training. We asked three questions: Can pigeons learn a series of four overlapping conditional discriminations using this procedure? Will they learn rapidly enough to make the procedure useful for further explorations of the transitive inference problem? Will they show the transitive inference effect (i.e., a preference for B over D) in test trials?

METHOD

Subjects

Four pigeons, *Columba livia* (2 White Carneaux and 2 Silver Kings), served as subjects. All had previous experience on various reinforcement schedules, but none had been exposed to the procedure used in this experiment.

Apparatus

The experiments were conducted in a standard operant conditioning chamber (39.3 cm by 30.5 cm by 40.5 cm). There were two clear response keys (2.5 cm diameter) in the center of the front panel, 6.0 cm apart from center to center and 9.0 cm from the chamber ceiling. Access to mixed grain in a magazine was made available through an aperture (5.2 cm by 4.2 cm) located in the front panel 11.5 cm below the center of the two keys. A color video monitor (21.5 cm by 29.5 cm) presented stimuli on the right and left sides of the screen. The monitor and chamber were arranged such that a single stimulus could be viewed through each key. The monitor was placed behind the front panel at a distance of approximately 36.5 cm, and a divider was placed perpendicular to the front panel beginning at a point between the two keys and extending to the monitor. The divider visually isolated each stimulus (see Figure 1). A houselight on the ceiling, 26.5 cm from the front panel, illuminated the chamber. Both chamber and monitor were enclosed in a sound-attenuating box, with a fan to mask extraneous noise. Experimental events were controlled by a Commodore 64® microcomputer located in an adjacent room. Data were transferred to a faster and larger computer for analysis.

Procedure

Subjects were given trials of conditional discriminations that usually ended with reinforcement. On each trial, two stimuli were

presented simultaneously, one associated with reinforcement (S+) and the other not (S-), according to schedules to be described below. Selection of the pair to be presented was determined by what we call an *autorun* procedure. We next describe the stimuli and autorun procedure used here and in subsequent experiments.

Stimuli. Stimuli were symmetrical patterns, each composed of a unique symbol, color, and surface area. Each stimulus was 15 symbols high and had a unique symbol type: open circles (Stimulus A), plus signs (Stimulus B), clubs (Stimulus C), hearts (Stimulus D), or a vertical bar (Stimulus E). Rows were generated by applying the following rules once for each stimulus: (a) Place the designated symbol in the center of the top row, at Column 0; (b) move to Column 1, and if a random number is greater than p (the criterion for that stimulus), place another symbol at that column; otherwise, move to the next row, Position 0, and repeat; (c) repeat the procedure until Row 15 has been completed. These rules generate the right half of the stimulus. The left half is simply the mirror image of the right. The five bilaterally symmetric symbols we generated by applying this algorithm are shown in Figure 2.

These five stimuli were arbitrarily designated as A, B, and so forth, and were grouped into pairs, to be associated with reinforcement as follows: A+B-, B+C-, C+D-, D+E-. Each pair is referred to as a *group*,² which therefore comprised the same two stimuli in the two possible physical arrangements: reinforced stimulus on the left half of the monitor and the nonreinforced stimulus on the right half, or vice versa. This classification generates four groups and a total of eight different pairs from the five stimuli (see Table 1).

Autorun procedure. Each session consisted of several trials during which a stimulus pair appeared for some amount of time and ended with either reinforcement or nonreinforce-




Fig. 2. The five stimuli used in all experiments. The method by which these were generated is described in the text.

² The standard notation is stimulus pair, but, because our procedure ensures that the reinforced (and nonreinforced) stimulus is presented on both the right and left sides of the video screen, we needed a way to refer to the stimuli of a conditional discrimination (e.g., A+B-) and the possible configurations (A+ on left, B- on right; A+ on right, B- on left). Hence, the latter are referred to as "pairs" and the former as a "group" (of possible configurations).

Table 1

Description of stimulus groups and pairs used in Experiments 1, 2, and 3. "Position" refers to the side of the video screen on which the reinforced (+) and nonreinforced (-) stimulus appeared.

Stimulus group	Stimulus pair	Position	
		Left	Right
1	1	A+	B-
	2	B-	A+
2	3	B+	C-
	4	C-	B+
3	5	C+	D-
	6	D-	C+
4	7	D+	E-
	8	E-	D+

ment. On reinforced trials, a timer for reinforcement began with the simultaneous presentation of a stimulus pair, and a reinforcer (2-s access to mixed grain) was delivered when the time spent responding to the S+ stimulus met or exceeded a specified duration, selected pseudorandomly from a range of 3 to 10 s (variable interval, VI 6.5 s). Pecks to S- reset the timer, which began again only with a peck to S+. In this way, the interfood intervals of the VI schedule determined the minimum time for reinforcement and the minimum duration of each trial. On nonreinforced trials, the stimuli were presented for a fixed amount of time, beginning with each stimulus-pair presentation. Trials were separated by an intertrial interval (ITI) of 1.5 s, during which the monitor was blank and key pecks were not recorded.

The stimulus pairs were selected trial by trial according to an adaptive rule, which always (with a constraint to be described) favors the stimulus pair with the poorest performance history. Performance is measured as the number of total correct responses minus the total number of incorrect responses. We call this the *V value* of each stimulus pair; a peck to S+ increments the *V value* of that pair by one, and a peck to S- decrements the *V value* of that pair by one. The *V values* are continuously updated throughout the session, so that at end of each trial, the response-based autorun procedure (we used a time-based procedure in later experiments) selects the stimulus pair with the worst performance at that point in the experiment. When the lowest *V value* is the same among two or more pairs, the autorun

procedure makes a random selection among them.

The autorun procedure is also designed to prevent subjects from developing a position bias by automatically ensuring that each stimulus pair is presented in both configurations. Whereas the lowest *V-value* rule determines which stimulus pair is selected, a variable we call *momentum* controls how often stimulus pairs from one group are presented before pairs from another group are selected. We use momentum in the following way. Autorun first selects the stimulus pair with the lowest *V value*, and then determines if momentum should be applied. If the to-be-presented stimulus pair is from the *same* group as the previous pair, then no momentum is used. However, if the next pair is from a *different* group, then momentum is applied: A constant amount (the momentum value) is subtracted from the *V values* of both pairs of the to-be-presented group, and the same momentum value is added back to the *V values* of the two pairs of the previous group. *V values* are carried over from one session to the next. In this way, the momentum variable artificially (and temporarily) lowers the *V values* of *both* stimulus pairs within a group, linking the pairs and ensuring that both are presented for some amount of time—the larger the momentum value used, the longer the exposure to the pairs within a group.

To illustrate, the lowest *V-value* and momentum rules coordinate in the following way. If Pair 1 (i.e., A+B-; see Table 1) has just been presented and Pair 2 (B-A+) is selected to appear on the next trial, then no momentum is used. But if Pair 1 has just been presented and Pair 5 (C+D-) is to appear on the next trial, then a momentum value (20 in all these experiments) is added to the *V values* for Pairs 1 and 2 and is subtracted from Pairs 5 and 6. This procedure accomplishes two things. First, subtracting an amount from both pairs of the next group ensures that the subject is exposed to both pairs of that group; that is, they see the stimuli in both spatial arrangements. Second, restoring the momentum value to the last set of *V values* means that the *V values* of groups no longer "in play" truly reflect the animal's performance.

In sum, with the autorun procedure, each session begins with the stimulus pair with the lowest *V value* (for the very first session, the

Table 2

Order and description of conditions used in Experiments 1, 2, and 3. V values are based on either number of responses (rsp) or time; some conditions contained both reinforced (Sr) and nonreinforced (Sr-) trials; and the durations of test trials were either independent of (rsp ind) or dependent on (rsp dep) responding.

Experiment	V value rule	Sr schedule	Trials	Sr	Sr-	Test trials	Sessions	Test trial duration
1 Phase A	rsp	VI 6.5 s	80	80	0	none	42	none
1 Phase B	rsp	VI 6.5 s	96	72	24	rsp ind	19/20/21	20 s
1 Phase C	rsp	VI 6.5 s	96	72	24	rsp ind	5	10 s
1 Phase D	rsp	VI 6.5 s	80	80	0	none	14	none
2 Phase A	rsp	VI 6.5 s	96	72	24	rsp dep	14	10 s
2 Phase B	rsp	VI 6.5 s	96	72	24	rsp dep	5	10 s
2 Phase C	time	VI 6.5 s	96	72	24	rsp dep	14	10 s
2 Phase D	time	VI 6.5 s	96	72	24	rsp dep	5	10 s
3 Phase A	rsp	VI 6.5 s	96	72	24	rsp dep	20	10 s
3 Phase B	rsp	VI 6.5 s	96	72	24	rsp dep	5	10 s
3 Phase C	time	VI 6.5 s	96	72	24	rsp dep	14	10 s
3 Phase D	time	VI 6.5 s	96	72	24	rsp dep	5	10 s

pair is chosen randomly). As performance on that pair improves, the next pair is always the one with the lowest V value. Momentum is applied only when the lowest V rule requires a transition between stimulus groups.

We studied the effect of changing various aspects of the autorun procedure in four phases (see Table 2 for a summary of all phases and experiments).

Phase A: 100% reinforcement. Each session consisted of 80 trials, all of which ended with reinforcement. If a subject did not meet the reinforcement requirement after 9 min, the program proceeded to the next trial (this occurred on only six trials throughout the entire study). This phase was in effect for 42 sessions.

Phase B: Partial reinforcement. To prepare pigeons for testing of novel combinations of stimuli without feedback, nonreinforced trials were added. As in Phase A, the autorun procedure determined the sequence of training stimulus pairs. However, the total number of trials per session was increased from 80 to 96, 24 of which were now nonreinforced trials. Nonreinforced trials lasted for a fixed amount of time (20 s), measured from the presentation of the stimulus pair; responses were recorded, V values were not calculated, but the lowest V stimulus-selection rule remained in force. One in every block of four trials (determined randomly) ended without reinforcement. Birds 158 and 160 were exposed to this phase for 19 and 20 sessions, respectively; Birds 169 and 173 received 21 sessions each.

Phase C: Testing with novel stimulus pairs. This was the same procedure as Phase B, but novel combinations of stimuli appeared during nonreinforced trials; these trials lasted for 10 s. The novel groups were AC, AD, AE, BD, BE, and CE. In total, there were six novel stimulus groups and 12 stimulus pairs (6 multiplied by 2 for each position on the screen). Each pair appeared twice within a session in randomized order. This phase lasted five sessions.

Phase D: 100% reinforcement. Repeat of Phase A for 14 sessions.

RESULTS

The data gathered across Phases A through D were analyzed into proportion of correct responses and proportion of trials. Each day, the proportion of correct responses was calculated for each stimulus group as the number of responses to S+ divided by the total number of responses to S+ and S-, for both stimulus pairs of that group. We did the same when calculating the proportion of trials in which a particular group appeared.

Phase A. Because four discriminations were being learned simultaneously, these experiments generated a great deal of data, which must be presented in highly compressed form if the pattern of acquisition is to be grasped as a whole. The method we finally settled on is shown in Figures 3 and 4. Figure 3 shows for one session (Phase A, Session 13, Bird 158) the proportion correct [p(correct); upper half

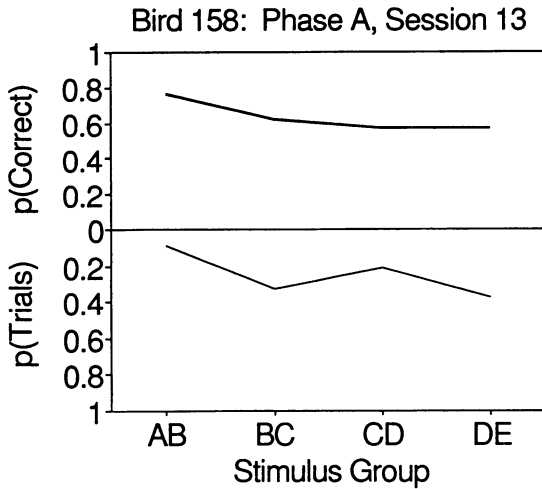


Fig. 3. The data are from Session 13 for Bird 158. The proportion correct— $p(\text{correct})$ —during training stimulus groups (AB, BC, CD, and DE) is shown as a heavy solid line; the proportion of trials— $p(\text{trials})$ —in which each stimulus group appeared is shown as a thin solid line.

of y axis] and proportion of trials [$p(\text{trials})$; lower half of y axis, plotted below, inverted, to make comparison with performance easier] associated with each stimulus group. Figure 4 shows data in the same form as Figure 3 for every session and every pigeon in Phase A. The data from Figure 3 are shown in a thin vertical box at the top left graph in Figure 4. In Figure 4, rows correspond to stages of training (Sessions 1 through 14, 15 through 28, and 29 through 42); columns correspond to individual birds. The x axis on all graphs consists of a repeated series of three-line segments, connecting four points. The four points are performances on the four training-stimulus groups in each experimental session. Thus, the first four data points indicate the proportion correct (heavy solid lines) and trials (thin solid lines) for stimulus groups AB, BC, CD, and DE on Session 1, the second set of four points indicates these two data types for the same four groups on Session 2, and so forth.

How quickly did the subjects learn the four overlapping conditional discriminations? Scanning performance across sessions in Figure 4 indicates that learning was relatively rapid. On the very first session, most birds responded well enough that at least three of the four stimulus groups appeared during that session; for example, Bird 158 saw Groups BC, CD, and DE, but not AB [its $p(\text{correct})$

and $p(\text{trials})$ for this group were both zero]. Bird 169 differs from the other subjects in that its performance on Group AB was so poor in the first session that it saw no other groups. During subsequent sessions, performance for all subjects was variable across the stimulus groups and only two or three groups appeared within a session; however, all four stimulus groups appeared within a single session by Session 8 for Birds 160 and 173, and by Session 9 for Bird 169. Bird 158 saw all stimulus groups by Session 6, but one was missed (not the same one) during Sessions 7 through 11. By Session 12 all stimulus groups appeared within the same session (for all subjects) and continued to do so, with the exception of Bird 158 on Sessions 22, 28, and 31. Its proportion of trials on these days suggests that performance on certain groups was poor and that more trials were therefore dedicated to those groups. Nevertheless, once all four stimulus groups began to be seen each day, the performance of each subject did not change dramatically across sessions. Mean proportion correct across stimulus groups was generally between .70 and .80. Only Bird 158 showed a slight improvement between Sessions 29 and 42, from about .65 to .70 proportion correct.

Phases B and C. Figure 5 presents averaged results from Phases B (partial reinforcement) and C (partial reinforcement with test stimuli). The top graph of each panel displays the mean proportion correct and trials of the last five sessions of Phase B; bottom graphs show the data from Phase C and are based on the mean of all five sessions. Performance during Phase B was comparable to that in Phase A, even though 25% of the trials did not end with reinforcement: Proportion correct was at least .70 for all stimulus groups. Only 2 birds performed below .70 on certain stimulus groups: The proportion correct for Bird 158 was .69 for BC and .67 for CD, and the proportion correct for Bird 169 was .69 for CD. The highest level of performance, for all subjects, occurred during Group DE: .85 for Bird 158, .99 for Bird 160, .90 for Bird 169, and .93 for Bird 173. As might be expected from the autorun rule, Group DE was also presented on fewer trials than the other groups: from .14 to .20 of all trials on average.

Does performance change during testing with novel stimulus combinations? Is there evidence for transitive inference, symmetrical

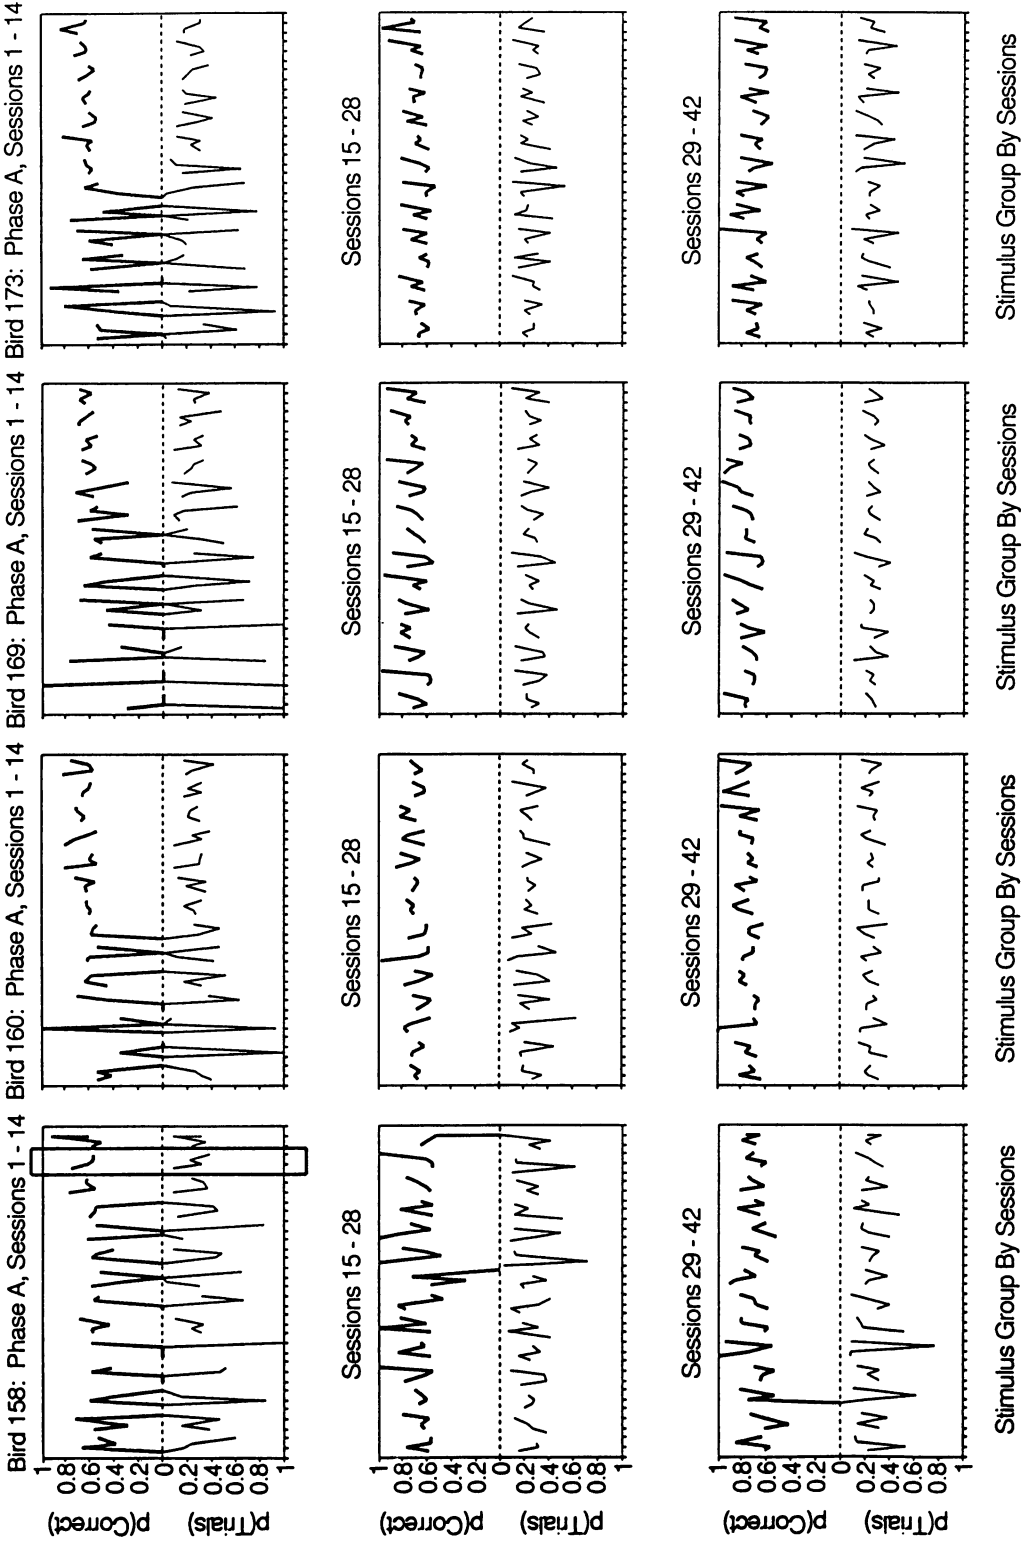


Fig. 4. Experiment 1. Mean proportion of correct responses [$p(\text{correct})$ —heavy solid lines] and trials [$p(\text{trials})$ —thin solid lines] during training stimulus groups (AB, BC, CD, and DE), by sessions, during Phase A. Proportion of trials is plotted below proportion correct for clarity. The contents of the vertical box surrounding Session 13 for Bird 158 are enlarged in Figure 3.

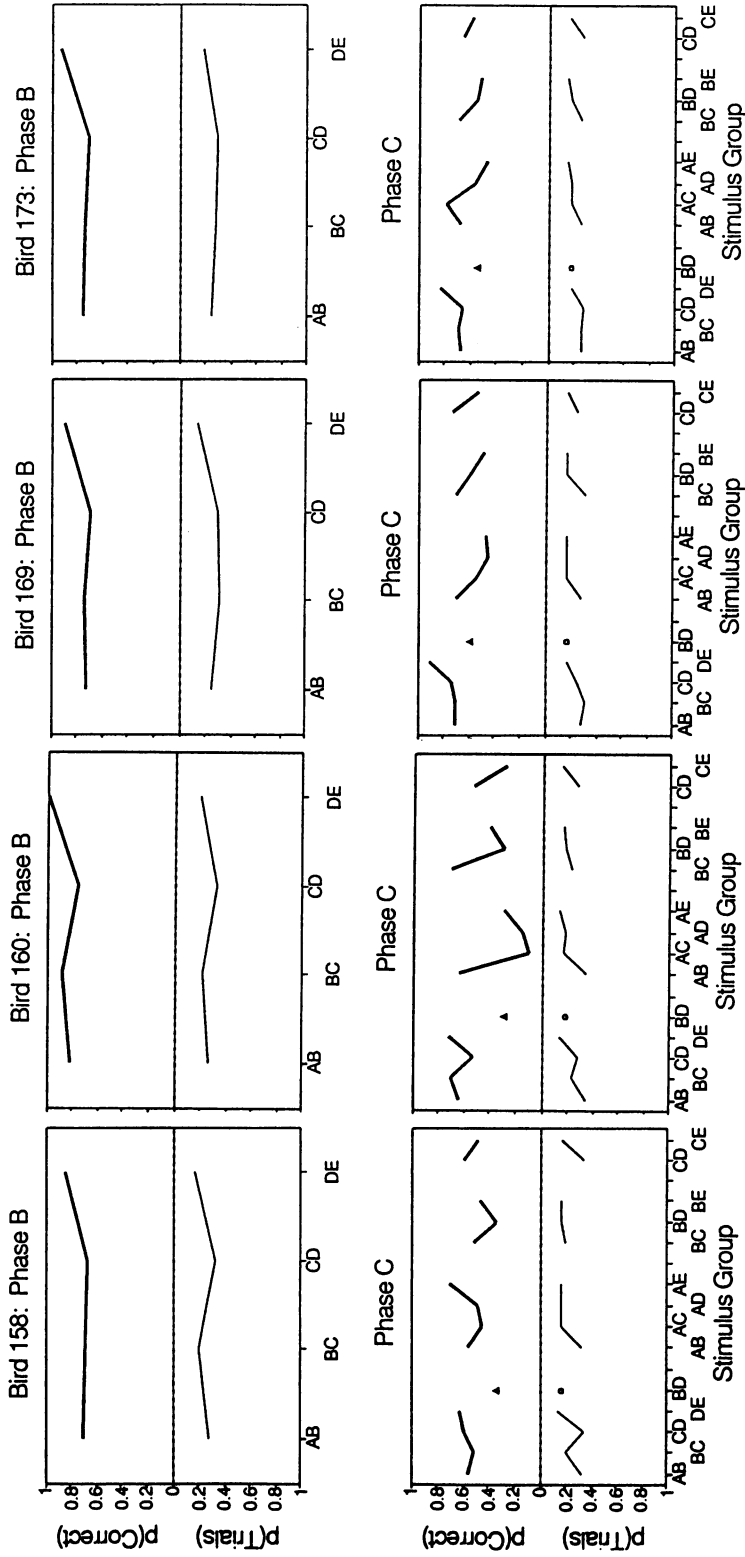


Fig. 5. Experiment 1. Top row: Mean proportion correct and trials during training stimulus groups, based on the last five sessions of Phase B. Bottom row: Mean proportion correct and trials during all stimulus groups (training and novel), based on all five sessions of Phase C.

Table 3

Number of reinforcers, rate, and probability of reinforcement for Stimuli A, B, C, D, and E, accumulated across Phases A and B of Experiment 1. Reinforcer rate (per minute) is calculated as the number of reinforcers divided by the amount of time spent responding to a stimulus. Probability is the number of reinforcers divided by the number of responses.

Subject	Stimulus					
	A	B	C	D	E	
Number	B158	1,173	1,240	1,378	1,009	0
	B160	1,101	1,202	1,314	1,111	0
	B169	1,269	1,341	1,229	1,027	0
	B173	1,223	1,231	1,471	947	0
	Group average	1,191.5	1,253.5	1,348	1,023.5	0
Rate	B158	6.407	4.328	3.913	3.152	0
	B160	7.244	4.370	4.046	3.612	0
	B169	6.331	3.756	4.100	4.021	0
	B173	7.125	3.860	4.449	2.724	0
	Group average	6.776	4.079	4.127	3.377	0
Probability	B158	.085	.053	.056	.048	0
	B160	.080	.061	.058	.053	0
	B169	.088	.062	.057	.054	0
	B173	.064	.044	.048	.035	0
	Group average	.079	.055	.055	.048	0

end-anchor, and symbolic-distance effects? The results from Phase C are shown in the bottom graphs of Figure 5, which shows proportion correct and trials on the training-stimulus groups and performance on novel stimulus combinations.

First, overall performance on the four training groups deteriorated from Phase B to Phase C for Birds 158, 160, and 173, from about .73 (Phase B) to .58 (Phase C), .86 to .65, and .78 to .71, respectively. Second, performance on novel groups was generally poor, ranging from .34 (BD) to .70 (AE) for Bird 158, .10 (AC) to .40 (BE) for Bird 160, .45 (AD) to .59 (BD) for Bird 169, and .46 (AE) to .78 (AC) for Bird 173.

Third, our subjects failed to show transitive inference during Phase C. The B ? D transfer tests reveal that none of the subjects consistently preferred B over D (compare filled triangles in Figure 5); in fact, 2 seemed to prefer D by responding at .34 (Bird 158) and .30 (Bird 160) proportion correct; the remaining subjects chose B over D at slightly over chance levels, .59 and .54 for Birds 169 and 173, respectively.

What might explain Bird 158's and Bird 160's preferences for D during BD tests? The autorun procedure was not designed to equalize reinforcement among the stimuli (instead, it equalized performance levels across groups);

it is possible, therefore, that an incidental outcome of the procedure may have been to reinforce responses to D more often than responses to B for these 2 subjects. Table 3 shows the number of reinforcers and the rate and probability of reinforcement for each stimulus for all 4 pigeons. These values were accumulated across Phases A and B, immediately prior to testing in Phase C.

The data in Table 3 do not explain why some birds (158, 160) showed a preference for D during BD presentations. These birds actually received fewer reinforcers and had a lower rate and probability of reinforcement for D. Does number, rate, or probability of reinforcement predict performance during other stimulus pairs in Phase C (Figure 5)? Table 3 suggests that reinforcement per se does not reliably predict test performance. For example, Bird 169 received many more reinforcers for C (1,229) and D (1,027) than for E (0). Yet, during tests this bird (and others) emitted more correct responses to C when it was paired with D than when C was paired with E, even though responding to E was never reinforced. Similar discrepancies can be found for rate and probability of reinforcement—neither can explain the test performance of all subjects.

Fourth, there were no symmetrical end-anchor effects (i.e., better performance on A + B— and D + E—): Birds 158 and 160 showed ap-

proximately equal performance on all stimulus groups; Birds 169 and 173 showed an end-anchor effect with a high proportion correct on Group DE. However, their performance on AB was about the same as that during BC and CD.

Finally, we looked for symbolic-distance effects, where performance level should increase as the ordinal distance between pairs of stimuli increases. Comparison of performance on BC and BD reveal an opposite effect. Instead of an increase in the proportion correct from BC to BD, the proportion correct actually *decreased* for all subjects. We also looked for symbolic-distance effects across other stimulus groups, assembled in Figure 5 in terms of the first item (AB-AC-AD-AE, BC-BD-BE, and CD-CE; see *x* axis label), and found that performance did not improve the further apart the stimuli were (ordinally and in terms of reinforcement value for these particular cases, because responses to A and E were invariably reinforced and nonreinforced, respectively). Performance on AE was not reliably better than performance on AB, for example; likewise, performance on CE was consistently lower than CD for all subjects. These types of inverse relationships were clearest for Bird 169, whose performance level decreased across AB-AC-AD-AE from .70 to .47, across BC-BD-BE from .70 to .49, and across CD-CE from .74 to .52.

Phase D. The proportion correct and trials during Phase D (return to 100% reinforcement) are presented in Figure 6; the top graphs show responding across all sessions for each subject, and the lower graphs are the mean of the last five sessions of this phase. There were no consistent changes across sessions, and by the final session performance was comparable to that obtained in Phase A: All subjects chose the S+ stimulus of each group with at least .70 accuracy, with the exception of Bird 173, whose proportion correct ranged from .66 on Group AB to .94 on Group DE. Also, as in other phases of this experiment, all subjects performed consistently best on Group DE: .84, .97, .93, and .95 for Birds 158, 160, 169, and 173, respectively.

DISCUSSION

Our results suggest that the autorun procedure is a viable method for rapidly training subjects to discriminate four overlapping stim-

ulus groups. By the end of Phase A most subjects were responding to each S+ with at least .70 accuracy. Performance on Groups BC, CD, and DE was comparable to that reported by Fersen et al. (1991). However, we did not obtain an average level of performance on Group AB as high as that reported by Fersen et al., whose subjects achieved more than .87 proportion correct—our subjects achieved only about .75. That is, they found symmetrical end-anchor effects, whereas our subjects typically showed better performance on only one end pair (DE), an asymmetrical effect. More striking is that the response-based autorun failed to produce transitive inference, and performance level decreased as the symbolic distance between pairs of stimuli increased. This was especially surprising because Fersen et al. (1990, 1991) and several other researchers (e.g., Gillan, 1981; McGonigle & Chalmers, 1977, 1992) found transitive inference and related effects following similar conditional discrimination training.

There are several differences between our study and theirs that are worth noting, however. First, they used longer intertrial intervals (4 s) and used correction trials in which incorrect responses produced 5 s of timeout in the dark and repeated presentation of that pair until a correct response was given. They also used a response-based criterion for reinforcement (eight consecutive pecks to S+) and punishment (eight not necessarily consecutive pecks to S-), whereas we used a time-based criterion (VI intervals) and pecks to S- reset the interval timer. Any or all of these differences might account for the small performance differences between the two procedures. Second, some of the subjects in the Fersen et al. (1991) study failed to reach criterion during the initial training condition and were not included in the remainder of the study; we show data from all 4 of our birds, across all experimental conditions. Third, we needed significantly fewer sessions than did Fersen et al. to reach asymptotic performance: 42 versus 125. Even so, some of our birds performed at the same level as Fersen's: Proportion correct in the .80 to .90 range was not unusual (see Figures 4, 5, and 6). Given that Fersen dropped some of his subjects from the study because they did not reach criterion, it would not be wise to conclude that our procedure produces generally worse performance, overall or on just Group




Fig. 6. Experiment 1. Top row: Mean proportion correct and trials during training stimulus groups, by sessions, during Phase D. Bottom row: Mean proportion correct and trials of the last five sessions of Phase D.

AB. Finally, our autorun procedure did not equalize how often a response to each stimulus was reinforced, and analyses of the outcome of the procedure indicate that some stimuli were associated with more reinforcers or had a higher rate or probability of being associated with reinforcement. But, reinforcement during training did not clearly predict choice on the various combinations of stimulus pairs during training, at least when absolute differences were considered. It is difficult to know if these results are unique to our procedure, because it is unclear exactly how (or whether) Fersen's procedure balanced reinforcement.

What other factors might be responsible for the absence of transitive inference and related effects in our experiment? One possibility is that our birds failed to attain the level of performance of those of Fersen et al. (1991). But, in fact the differences are rather small, and even our "best" bird (160) failed to show the transitive inference effect (see Figure 5); thus, level of discrimination performance does not seem to be the crucial factor. Other possibilities should also be considered. First, our failure to find a transitive inference effect might be somehow related to the lack of symmetrical end-anchor effects enforced by the response-based autorun procedure: Because stimulus groups were selected according to the "rights minus wrongs" V -value rule, performance varied little among stimulus groups. In Fersen et al.'s study, both performance and amount of exposure were free to vary (because of their correction procedure), and performance on the end groups was much better than performance on the two middle groups (i.e., presence of symmetrical end-anchor effects). Second, the lack of transitive inference effects may also be related to the absence of symbolic-distance effects. In fact, our subjects generally performed worse as the distance between stimulus pairs increased. It is especially difficult to understand why performance on AE, BE, and CE was generally poor, given that responding to E was never reinforced. Perhaps our subjects needed more training trials to learn about E. However, this is unlikely because performance on DE, immediately prior to testing (Phase B), was high for all subjects (see top of Figure 5). Third, we used many more test trials and novel stimulus combinations than Fersen et al. did. Perhaps our subjects detected differences between training (reinforced) and test (non-

reinforced) trials, because in addition to the fact that novel stimulus combinations were used, nonreinforced (response-independent) trials were significantly longer (20 s) than the typical VI interval value (VI 6.5 s).

We addressed these problems in the next experiment (a) by giving additional training to see if performance would improve further, (b) by arranging all nonreinforced trials to be of the same duration during training and testing, and (c) by equating exposure to all stimulus groups and letting performance level vary across groups.

EXPERIMENT 2: EFFECTS OF RESPONSE- VERSUS TIME-BASED V VALUES

In this experiment, we studied the effect of holding the amount of exposure to a stimulus group constant while allowing performance to vary. "Amount of exposure" might be defined as "trials" or "time." In this experiment we held constant the amount of time that subjects saw a stimulus group.

METHOD

Subjects and Apparatus

These were the same as those used in Experiment 1.

Procedure

Subjects were exposed to four different experimental phases, A through D. Phases A and B involved response-based assessment of V values and momentum, in which performance was held approximately constant and exposure was free to vary; Phases C and D based V values and momentum on time, so that performance was free to vary. See Table 2 for a summary.

Phases A and B (response-based V values). In Phase A, animals were trained to respond correctly to $A+B-$, $B+C-$, $C+D-$, and $D+E-$. Each session (for a total of 14) comprised 96 trials, 24 of which were nonreinforced; these trials lasted for 10 s and began with a response to $S-$ or $S+$. V values and momentum were response-based, as in Experiment 1. Following this phase, subjects were exposed to Phase B in which tests with "novel" combinations of stimuli were given. This phase was identical to Phase A, with the exception that novel stimulus groups appeared during

nonreinforced trials. This phase lasted for five sessions.

Phases C and D (time-based V values). These phases were similar to Phases A and B described above, with training during Phase C and testing during Phase D. However, the V values and momentum were calculated and based on the time spent responding. At the end of each trial, the time between the first response and reinforcement was added to the V value, and the lowest V-value rule was applied as in Experiment 1, favoring the stimulus pair with the least amount of exposure. Momentum was used as before, with a value of 20 s subtracted from a new stimulus group and added to the old stimulus group. The long-term effect of this procedure was to equalize the time spent responding to all eight stimulus pairs.

Responses were recorded but had no effect on the V value. The subjects received 14 sessions of training (Phase C) followed by five sessions of testing (Phase D).

RESULTS

Response-Based Conditions

Figure 7 presents the data for each bird during training (Phase A) and testing (Phase B). The data are calculated as in Experiment 1, and each graph shows mean proportion correct and trials across the various stimulus groups.

During training, most subjects allocated at least .70 of their responses to the correct stimulus on training groups with a few exceptions: .68 to A in the AB comparison and .69 to C in CD for Bird 158, and .67 to C in CD for Bird 173. All subjects performed well on the End Group DE, correctly choosing D with .80 to .96 accuracy. When shifted to testing conditions, the proportion correct on training stimulus pairs decreased, but responding to DE remained high, ranging from .90 to .96 proportion correct.

There was little evidence for transitive inference or for symmetrical end-anchor and symbolic-distance effects. Only Bird 169 showed signs of a symmetrical end-anchor effect, with higher proportions correct on End Groups A+B- and D+E-. Given symmetrical end-anchor effects, one might expect to observe transitive inference in this bird because Fersen et al. (1991) found both, but the bird selected Stimulus B over D during BD with only .47 accuracy. The remaining birds failed

to show both symmetrical end-anchor and transitive inference effects; they selected B over D with .51 (Bird 158), .57 (Bird 160), and .64 (Bird 173) accuracy. Finally, there were no symbolic-distance effects across Groups BC-BD or any other combination of groups; if there was any trend at all, it was in the opposite direction: Performance on AB was generally better than AC or AE, BC was better than BD or BE, and so forth.

Table 4 presents the number of reinforcers and the rate and probability of reinforcement for each stimulus during the two training phases of this experiment (Phases A and C); the measures were calculated as in Experiment 1. Although some aspects of performance seemed to depend on the birds' reinforcement history, some important features cannot be explained. For example, during testing all birds performed well on DE: Responding to Stimulus E was never reinforced (i.e., had zero rate and probability of being reinforced), thus when presented with DE, all subjects consistently chose D. But this analysis suggests that subjects should also do well on all stimulus pairs containing Stimulus E. Instead, performance on AE, BE, and CE was consistently lower than that during DE. Moreover, performance on CE should have been higher than that during CD, which might be predicted by the conditions of reinforcement associated with the stimuli: Responses to C and D were equally often reinforced and nonreinforced and the stimuli should have been responded to about the same number of times (i.e., poor performance on CD); whereas because responding to E was never reinforced, subjects should have responded exclusively to C during CE (good performance). Figure 7 shows that the opposite was found.

Time-Based Conditions

The results from training (Phase C) and testing (Phase D) are presented in Figure 8. How does changing from a response- to a time-based rule affect performance? Comparison of the results from Phases A (response-based autorun) and C (time-based autorun) indicates little change in the proportion correct during training. Although performance deteriorated for some birds during certain stimulus groups, the proportion correct was generally above .70. The most obvious changes were seen in the data from Birds 158 and 160, whose level of

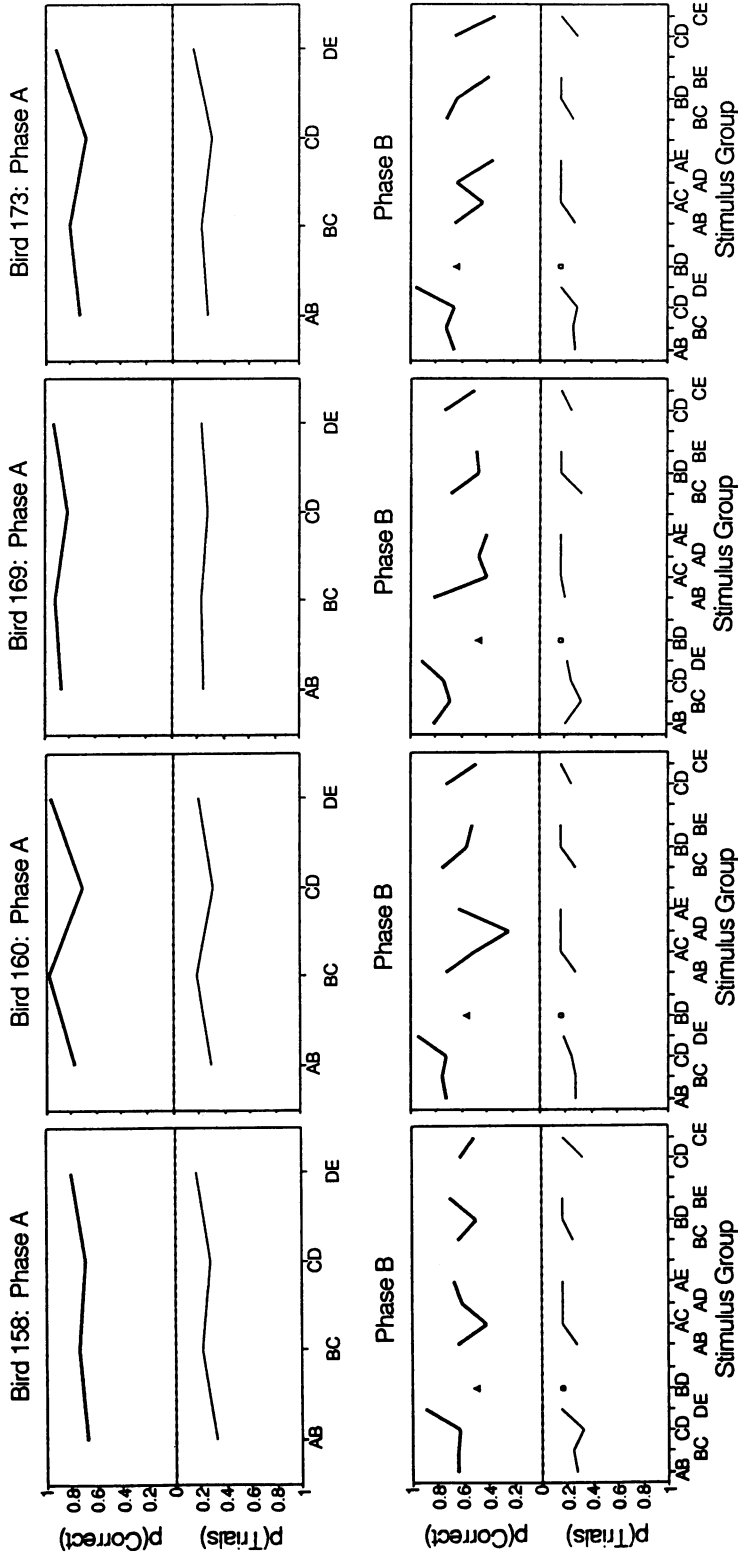


Fig. 2. Experiment 2. Top row: Mean proportion correct and trials during training stimulus groups, based on the last five sessions of Phase A. Bottom row: Mean proportion correct and trials of all stimulus groups (training and novel), based on all five sessions of Phase B.

Table 4

Number of reinforcers, rate, and probability of reinforcement for Stimuli A, B, C, D, and E during the training phases of Experiment 2 (Phases A and C). Reinforcer rate (per minute) is the total number of reinforcers obtained while responding to a stimulus divided by the total amount of time spent responding to that stimulus. Probability is the number of reinforcers divided by the number of responses.

Phase	Subject	Stimulus				
		A	B	C	D	E
A (training, response-based)						
Number	B158	387	391	335	231	0
	B160	307	225	279	197	0
	B169	256	272	256	224	0
	B173	269	268	295	176	0
	Group average	304.75	289.00	291.25	207.00	0
Rate	B158	8.204	5.756	5.084	4.587	0
	B160	6.575	3.532	5.067	3.210	0
	B169	5.718	4.640	4.646	4.336	0
	B173	6.360	4.204	5.051	2.935	0
	Group average	6.714	4.533	4.962	3.767	0
Probability	B158	.077	.051	.049	.048	0
	B160	.061	.040	.049	.038	0
	B169	.104	.089	.087	.091	0
	B173	.044	.034	.035	.025	0
	Group average	.072	.054	.055	.051	0
C (training, time-based)						
Number	B158	284	231	174	319	0
	B160	239	244	185	340	0
	B169	233	225	243	307	0
	B173	229	248	211	320	0
	Group average	246.25	237.00	203.25	321.50	0
Rate	B158	6.406	3.930	2.954	3.587	0
	B160	6.198	4.379	4.327	4.343	0
	B169	5.480	4.308	4.033	4.367	0
	B173	6.575	4.418	4.339	4.281	0
	Group average	6.165	4.259	3.913	4.144	0
Probability	B158	.037	.030	.022	.027	0
	B160	.058	.042	.041	.045	0
	B169	.116	.106	.088	.102	0
	B173	.047	.038	.035	.038	0
	Group average	.065	.054	.046	.053	0

responding on CD decreased, from Phase A to Phase C, from .69 to .44 (Bird 158) and from .70 to .63 (Bird 160). The proportion correct on DE remained high across conditions for all birds, from .89 to 1.00 proportion correct.

Tests with novel stimulus groups showed several interesting features. First, the overall proportion correct on the training groups decreased from training to testing phases. Birds 158 and 160 had similar patterns, with higher performance on Groups BC and DE and much lower performance on AB and CD. The other birds maintained their accuracy on End Group

DE, but this was not the case for other stimulus groups.

Second, only Bird 173 showed a transitive inference effect. This bird selected B over D, during presentations of BD, with .83 accuracy. Yet, this subject still lacked symmetrical end-anchor effects. In contrast, Bird 169 had symmetrical end-anchor effects but not transitive inference. The other birds (158 and 160) showed neither effect. Could the reinforcement history account for finding transitive inference in the performance of Bird 173? Table 4 does not support this: Bird 173 had approximately the same rate and probability of reinforcement

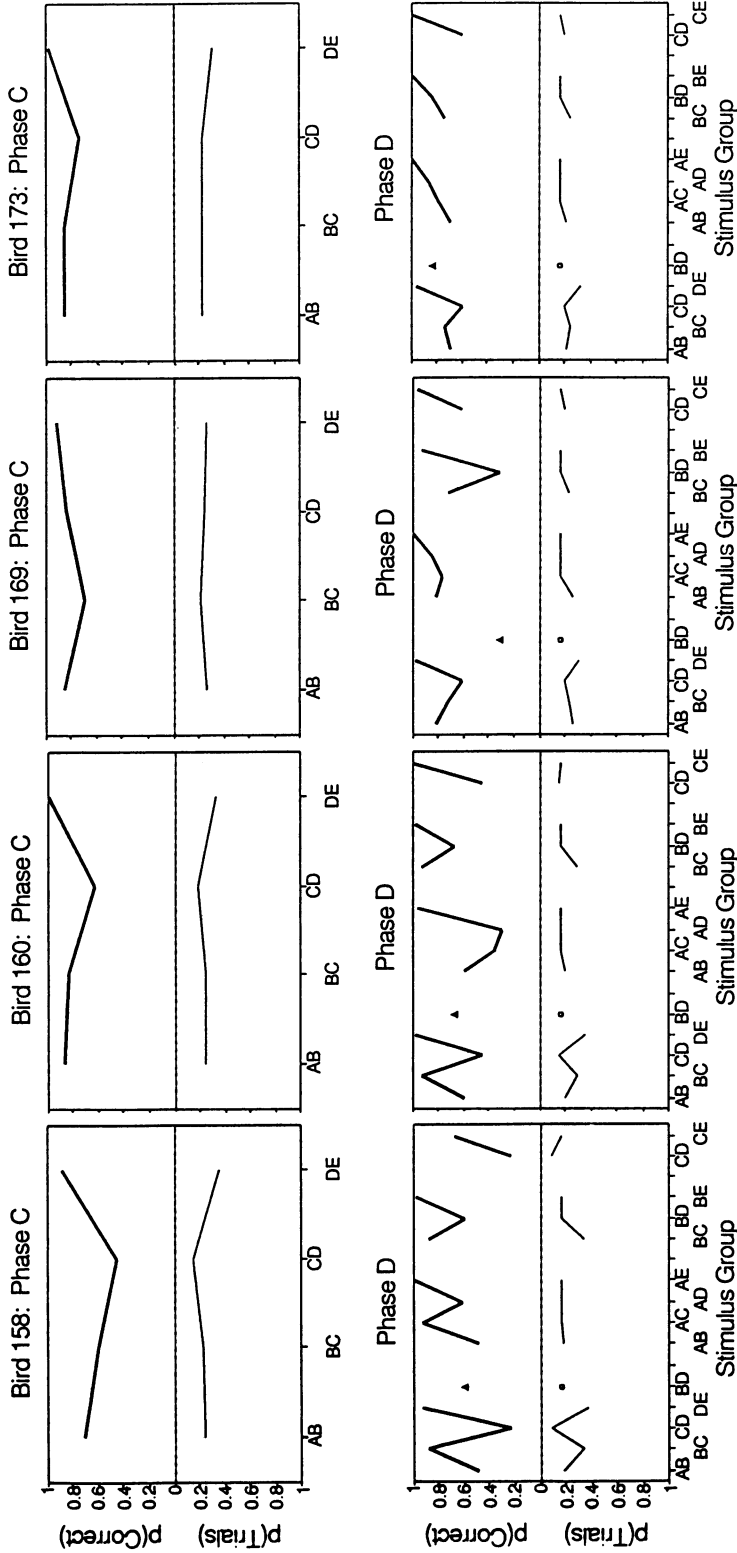


Fig. 8. Experiment 2. Top row: Mean proportion correct and trials during training stimulus groups, based on the last five sessions of Phase C. Bottom row: Mean proportion correct and trials of all stimulus groups (training and novel), based on all five sessions of Phase D.

for B and D during training, and had a *higher* number of reinforcers for D (248 for Stimulus B vs. 320 for Stimulus D).

Third, there was some evidence for symbolic-distance effects for some of the birds and across certain stimulus groups. For example, all birds showed a large increase in the proportions correct from CD to CE, and most birds showed general increases across Groups AB-AC-AD-AE. This was clearer in Birds 169 and 173, and less so in others. However, performance on Groups BC-BD (comparable groups because responses to B, C, and D were about equally reinforced) did not reveal symbolic-distance effects. As in Experiment 1, most birds (158, 160, 169) performed poorer on BD than on BC. Only Bird 173 did better on BD, which is also evidence of transitive inference. This bird also had clear symbolic-distance effects across Groups BC, BD, and BE and Groups AB-AC-AD-AE and CD-CE.

Fourth, there is an inverse correspondence between performance and the proportion of trials across training stimulus groups only under the time-based rule: Stimulus groups associated with better performance appeared on fewer trials (compare the training series in Figures 7 and 8).

Finally, not only was accuracy highest on Group DE, but it was also relatively higher on all other stimulus combinations with Stimulus E. For instance, Bird 160 performed on Groups AE, BE, CE, and DE with .96, .97, .99, and 1.00 accuracy, respectively.

DISCUSSION

Additional training under a response-based autorun rule did not significantly improve performance on Groups AB, BC, CD, and DE (compare top panels of Figure 5 from Experiment 1 and Figure 7 from this experiment); nor did this procedure reliably produce transitive inference or symmetrical end-anchor and symbolic-distance effects. But, replications of the response-based conditions did produce several consistent effects. First, the subjects learned the training stimulus groups with about .70 accuracy, and the proportion correct was always highest on End Group DE—most subjects selected D over E with an accuracy exceeding .90. Second, unlike Fersen et al. (1991), we consistently failed to find transitive inference and symmetrical end-anchor effects. Instead, there was a slight decline in performance

on training groups during testing, and performance on AB was about the same as that on BC and CD, although performance remained consistently higher on Group DE. Finally, there was relatively little evidence for a symbolic-distance effect: In most cases there was a decrease in the proportion correct as the distance between stimulus pairs increased.

Changing the autorun rule from response-based to time-based did not seem to improve performance on training stimulus groups (compare top panels of Figures 7 and 8). Reliable transitive inference and symmetrical end-anchor effects were also lacking. But, unlike the response-based condition there was some, though weak, evidence of symbolic-distance effects for certain combinations of groups. Especially noticeable were the results for Bird 173 (see bottom right panel of Figure 8).

Although we did not reliably observe transitive inference or other related effects under either the response-based or the time-based autorun rule, these conditions did produce different reinforcement histories and patterns of responding. First, under the response-based autorun rule, Stimulus B was associated with a larger number of reinforcers than Stimulus D, although this effect was not statistically significant for the group of subjects. The time-based autorun rule, however, consistently produced the opposite effect: There were more reinforcers for Stimulus D than B (because performance on unambiguous discrimination D+E- was always better than on ambiguous discrimination B+C-). Second, the time-based autorun rule caused a decline in performance on certain stimulus groups, but did not affect performance on Group DE, which was always good. Third, the two autorun rules also produced different correspondences between proportion correct and trials. Under a response-based rule, better performance was associated with fewer trials; with a time-based rule, better performance was associated with more trials. Fourth, time-based autorun produced two suggestive results. One subject (169) had symmetrical end-anchor effects and no transitive inference. This is unusual, because prior studies that find transitive inference also find symmetrical end-anchor effects. In contrast, another subject (173) was the only one with transitive inference and symbolic-distance effects, but it did not show a symmetrical end-anchor effect. These results suggest that per-

haps the presence of symmetrical end-anchor effects is not necessary for transitive inference and that symbolic distance might be. We do not know to what extent this dissociation among effects is significant, or whether the results are unique to this experiment or might have been found by Fersen et al. (1991), because they did not test their pigeons with all possible combinations of their five-term series. Furthermore, the study of Fersen et al., as well as most other studies, do not provide detailed analyses of performance for subjects without transitive inference.

Why, unlike Fersen et al. (1991), did we fail to find consistent transitive inference and symmetrical end-anchor and symbolic-distance effects, even under the time-based procedure? It seems unlikely that reinforcement history alone explains the various effects we observed. For instance, although the reinforcement history of a particular bird suggested that it should not have shown transitive inference, during testing it did show a preference for Stimulus B over D. It also seems unlikely that lack of experience is a reason: By the end of this experiment our subjects had received 119 to 120 sessions of training and testing. Moreover, towards the end of Phase D (under a time-based autorun rule), we observed some symbolic-distance effects and also found that subjects performed consistently better on Groups ?E (where ? is either A, B, C, or D). Despite signs of symbolic-distance effects, the pigeons still failed to show consistent transitive inference effects in tests. We did find small, differential effects of the time- and response-based procedures on the performance profile, the presence of symbolic-distance effects, and some suggestion that transitive inference performance was emerging towards the end of the experiment. To be sure that the apparent differential effects were real and not just the result of accumulated experience, and to allow further development of transitive inference performance if, as Fersen's results seem to suggest, extended experience is essential, we repeated both the response- and time-based procedures in Experiment 3.

EXPERIMENT 3: REPLICATION OF RESPONSE- AND TIME-BASED MOMENTUM EFFECTS

Experiment 2 did not control the order in which the animals received the time- and re-

sponse-based procedures; all 4 birds got the response-based procedure first, followed by the time-based procedure. We therefore repeated all four conditions in this experiment, to determine whether performance changes were reliably correlated with changes in procedure or merely reflected increasing experience.

METHOD

Subjects and Apparatus

These were the same as in the previous experiments.

Procedure

The conditions (Phases A through D) and order of exposure were identical to those used in Experiment 2. See Table 2 for a summary.

RESULTS

Response-Based Phases

Figure 9 shows the results from the training (Phase A) and testing (Phase B) conditions. Overall, performance on the training stimuli during Phase A was comparable to that obtained during other experiments (see Figures 5 and 7) and did not improve throughout the study; specifically, the proportion correct was at least .75 for all subjects and training groups, with the single exception of Bird 158 during Groups BC (.69) and CD (.67).

When shifted to testing (Phase B), pigeons' performance on the training pairs declined slightly; although the proportion correct remained high on Group DE, performance on AB was no better than on other stimulus groups, revealing an asymmetrical end-anchor effect (the same as that found in Experiments 1 and 2). Unlike the previous experiments, we did find clear evidence for transitive inference in 3 of 4 subjects: Birds 158, 169, and 173 selected Stimulus B over D during BD presentations with .93, .80, and .85 accuracy, respectively; Bird 160 achieved only .56 accuracy on BD presentations. There was also some evidence of symbolic-distance effects, clearer for Birds 158 and 173 and less so for Birds 160 and 169.

Can some aspect of each bird's reinforcement history predict whether or not it will show transitive inference? Table 5 presents the number of reinforcers and the rate and probability of reinforcement for all birds and stimuli, calculated as in Table 4. The values in Table 5 suggest that those birds which showed transitive inference (158, 169, and 173) during




Fig. 9. Experiment 3. Top row: Mean proportion correct and trials during training stimulus groups, based on the last five sessions of Phase A. Bottom row: Mean proportion correct and trials of all stimulus groups (training and novel), based on all five sessions of Phase B.

Table 5

Number of reinforcers, rate, and probability of reinforcement for Stimuli A, B, C, D, and E during the training phases of Experiment 3 (Phases A and C). Reinforcer rate (per minute) is the total number of reinforcers obtained while responding to a stimulus divided by the total amount of time spent responding to that stimulus. Probability is the number of reinforcers divided by the number of responses.

Phase	Subject	Stimulus				
		A	B	C	D	E
A (training, response-based)						
Number	B158	324	377	521	218	0
	B160	393	307	416	324	0
	B169	355	422	376	287	0
	B173	359	423	364	294	0
	Group average	357.75	382.25	419.25	280.75	0
Rate	B158	6.479	4.093	4.836	2.339	0
	B160	6.195	4.184	5.733	4.014	0
	B169	5.590	4.721	4.389	4.380	0
	B173	6.671	4.260	4.597	4.105	0
	Group average	6.234	4.314	4.889	3.710	0
Probability	B158	.042	.034	.040	.021	0
	B160	.052	.036	.047	.040	0
	B169	.119	.114	.099	.092	0
	B173	.047	.043	.036	.036	0
	Group average	.065	.057	.055	.047	0
C (training, time-based)						
Number	B158	247	230	179	352	0
	B160	224	264	213	307	0
	B169	231	250	242	285	0
	B173	227	217	256	308	0
	Group average	232.25	240.25	222.50	313.00	0
Rate	B158	6.406	3.719	3.307	3.685	0
	B160	6.237	4.825	5.519	4.713	0
	B169	5.607	4.458	4.413	4.499	0
	B173	6.219	4.322	4.648	4.896	0
	Group average	6.117	4.331	4.472	4.448	0
Probability	B158	.034	.030	.025	.029	0
	B160	.053	.039	.044	.041	0
	B169	.129	.098	.105	.107	0
	B173	.050	.045	.038	.040	0
	Group average	.067	.053	.053	.054	0

BD pairs received more reinforcers for Stimulus B during training; Bird 160, which did not show transitive inference, had slightly more reinforcers for Stimulus D. Measures of reinforcer rate and probability were not predictive.

Furthermore, as noted by one insightful reviewer, because the autorun procedure did not control for which stimulus group preceded a BD test group, recency effects from one stimulus pair to the next might predict the presence or absence of transitive inference in this condition. That is, preference for Stimulus B over D during BD presentations might somehow depend on the previous stimulus pair and re-

inforcement conditions. For example, if the previous trial contained B+C-, reinforced pecks to B might explain selection of B during a subsequent test trial with BD; similarly, if C+D- appeared on the previous trial, then nonreinforced pecks to D might explain avoidance of D during BD test trials. The overall effect would look like transitive inference.

We looked for recency effects by comparing the mean proportion correct for BD trials and the immediately preceding trial. For each subject, all BD trials [trial($n + 1$)] and the immediately preceding trial [trial(n)] were extracted from the data set [BD trials preceded by other test groups (e.g., AD or BE) were

not included in this analysis]. If the reinforcement conditions of the preceding trial affect BD performance, then performance on BD should be better when preceded by B+C- or C+D- trials than by A+B- or D+E- trials—the former favors transitive-inference-like behavior, whereas the latter does not. Table 6 shows the frequency and mean proportion correct for BD preceded by AB or DE trials (grouped together) or by BC or CD trials (grouped together).

Of the 3 birds that showed transitive-inference-like performance (158, 169, and 173), the data from only 2 (169 and 173) were in the right direction—better performance on BD following BC/CD trials. For instance, Bird 173's proportion correct on BD was .90 when preceded by BC/CD and only .71 when preceded by AB/DE. The remaining subject that displayed transitive inference (158) and the other subject that did not produce transitive inference both had lower proportions correct during BD trials preceded by BC/CD. Because of the relatively small number of available test trials (15 out of 20 at most) and because we do not have similar information from comparable studies (i.e., Fersen et al., 1991), it is probably wise to assume, for now, that the analysis of recency effects is inconclusive.

Time-Based Phases

The results from training (Phase C) and testing (Phase D) are shown in Figure 10. The training results were similar to those obtained during Experiment 2 (see Figure 8, Phase C) and slightly worse than the response-based conditions used here and in Experiment 2. Performance deteriorated between training and testing, especially during Group CD for Birds 158, 160, and 169, and during Group BC for Bird 173.

Transitive inference was found in only 2 birds: Bird 160 selected B over D with .77 accuracy, and Bird 169 performed with .83 accuracy. These birds did not have symmetrical end-anchor effects, although performance on DE was consistently higher than that on other training groups. There was also some evidence of symbolic-distance effects. Performance for these birds tended to improve with distance among the elements of a group, especially across CD-CE, and the proportion correct was generally higher on AE than AB. But, changes in performance on BC-BD were

Table 6

The number and mean proportion of correct trials with Test Group BD [trial($n + 1$)] and the immediately preceding trial [trial(n)], from Experiment 3 (Phase B). The data from Groups AB and DE and BC and CD are combined (see text for details). BD trials that were preceded by other test groups (e.g., AD or BE) are not included.

Subject	Trials	Trial		Trial		
		Trial(n) AB/DE	Trial($n + 1$) BD	Trial(n) BC/CD	Trial($n + 1$) BD	
B158	4	1.00	.94	15	.80	.89
B160	7	.73	.54	7	.94	.47
B169	9	.97	.70	10	.95	.88
B173	6	.98	.71	13	.89	.90

in the opposite direction. Table 5 indicates that all birds received more reinforcers during Stimulus D than B; thus, reinforcement number did not predict the occurrence of a transitive inference effect in the 2 birds that showed it. As in Experiment 2, 1 bird (173) showed clear symmetrical end-anchor effects but no transitive inference. In fact, this bird seemed to prefer D over B by responding with about .22 correct on BD tests.

GENERAL DISCUSSION

Our primary aims at the start of this study were to establish a rapid training procedure for multiple conditional discriminations and to determine the necessary and sufficient conditions for the transitive-inference-like behavior that develops after training on such tasks. In three experiments, we trained pigeons on four overlapping conditional discriminations by using a novel training procedure, the autorun procedure. Two rules comprise autorun, both operating on a trial-by-trial basis: (a) The lowest V-value rule determined which stimulus pair was to be presented, and (b) momentum prevented subjects from developing a position bias. In Experiment 1, the V values were response-based and kept track of performance (the number of correct minus the number of incorrect responses), and the lowest V-value rule favored the stimulus pair with the poorest performance history. In Experiment 2, we varied the autorun procedure and compared the effects of (a) holding performance on training stimulus pairs constant and letting the amount of exposure to each stimulus group vary and the reverse (response-based autorun), and (b) holding the amount of exposure to each stim-

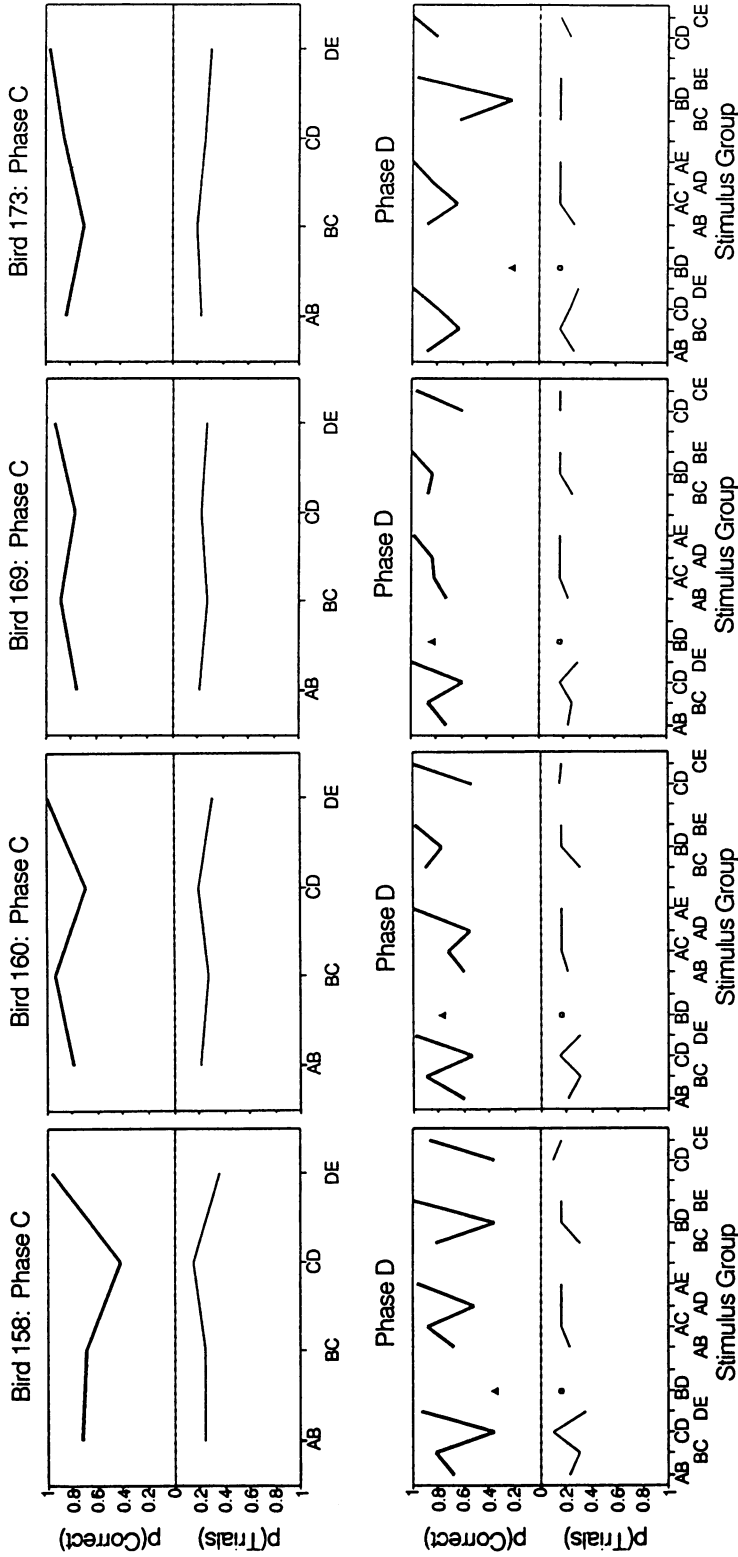


Fig. 10. Experiment 3. Top row: Mean proportion correct and trials during training stimulus groups, based on the last five sessions of Phase C. Bottom row: Mean proportion correct and trials of all stimulus groups (training and novel), based on all five sessions of Phase D.

ulus pair constant and letting performance vary (time-base autorun)—here, the lowest V-value rule favored the stimulus pair with the least amount of exposure. Finally, in Experiment 3, we repeated the conditions used in Experiment 2 to control for order effects and to determine whether further training might allow the development of consistent transitive inference performance.

Our study differs from those by Fersen et al. (1990, 1991) in that we used a principled method for incorporating training stimulus pairs and more carefully controlled for how often stimulus pairs were rewarded or the amount of exposure to each stimulus. Our results with the autorun procedure confirm some features of the Fersen et al. (1990, 1991) studies, but not others. We were able to get acceptable discrimination performance on all four of the overlapping stimulus groups, A+B-, B+C-, C+D-, D+E-. In fact, we obtained results that are comparable to those of Fersen et al. The comparison is made in Figure 11, which shows the results from Fersen et al. (1991) along with our results from Experiments 1 and 3. The average performance of our pigeons was as good or better than those of Fersen on Groups BC, CD and DE, and was worse only on Group AB, despite the fact that we included all birds in our average and Fersen eliminated the 2 worst birds from his experiment. Furthermore, autorun worked as we had hoped, producing relatively rapid discrimination learning under the difficult conditions of these experiments. However, we failed to find reliable symmetrical end-anchor effects (although performance on end discrimination D+E- improved to a high level across conditions in all birds), symbolic-distance effects, or transitive inference effects.

One of these results is not surprising. The failure to find symmetrical end-anchor effects is almost forced by the response-based autorun procedure, which tends to equate discrimination performance across the four stimulus pairs. (Response-based autorun nevertheless failed to prevent the eventual emergence of maximal performance on the D+E- discrimination.) The obvious inference from the point of view of previous studies is, therefore, that the lack of transitive inference and symbolic-distance effects in our study is linked in some way to the absence of the AB end-anchor effect, so that if performance on AB is not significantly

better than performance on BC and CD, these effects do not appear. After all, almost all studies that report transitive inference also report symmetrical end-anchor and symbolic-distance effects (e.g., Bryant & Trabasso, 1971; Fersen et al., 1991; Gillan, 1981; McGonigle & Chalmers, 1977, 1992). Unfortunately, our data do not strongly support this view. The birds that did occasionally show symmetrical end-anchor effects (i.e., performance on A+B- and D+E- much better than performance on B+C- and C+D-), particularly under the response-based autorun, still failed to show strong transitive inference (e.g., Bird 169, Figures 7 and 8). And some of the birds that showed occasional transitive inference and symbolic-distance effects did not show symmetrical end-anchor effects (e.g., Bird 173, Figure 8; Bird 158, Figure 9; Birds 160 and 169, Figure 10). A tentative conclusion is that transitive inference does not depend on the presence of symmetrical end-anchor effects. The role of symbolic-distance effects is less clear and more difficult to evaluate.

The presence of a transitive inference effect was also not reliably correlated with high performance on the D+E- discrimination: Bird 173 (Figure 10) showed very strong performance on D+E- but preferred D over B in the BD test. Moreover, most birds most of the time showed excellent performance on D+E-, so this correlation does not readily differentiate conditions that do and do not produce the transitive inference effect.

Transitive inference also does not seem to depend on the level of reinforcement associated with each stimulus. In Experiments 2 and 3, the response- and time-based autorun rule produced differential effects on the number of reinforcers for Stimuli B and D—larger for Stimulus B during response-based autorun and larger for Stimulus D under time-based autorun. Hence, one might expect to find a preference for B over D (i.e., transitive inference) when the autorun rule was response-based and a preference for D over B during a time-based autorun rule. Although the only time we found strong, reliable transitive inference was during the response-based condition (Experiment 3, Phase B), some birds showed transitive inference during the time-based condition (Experiment 3, Phase D). Furthermore, it is not clear why earlier exposure to a response-based autorun rule (Experiment 1 and Experiment 2,

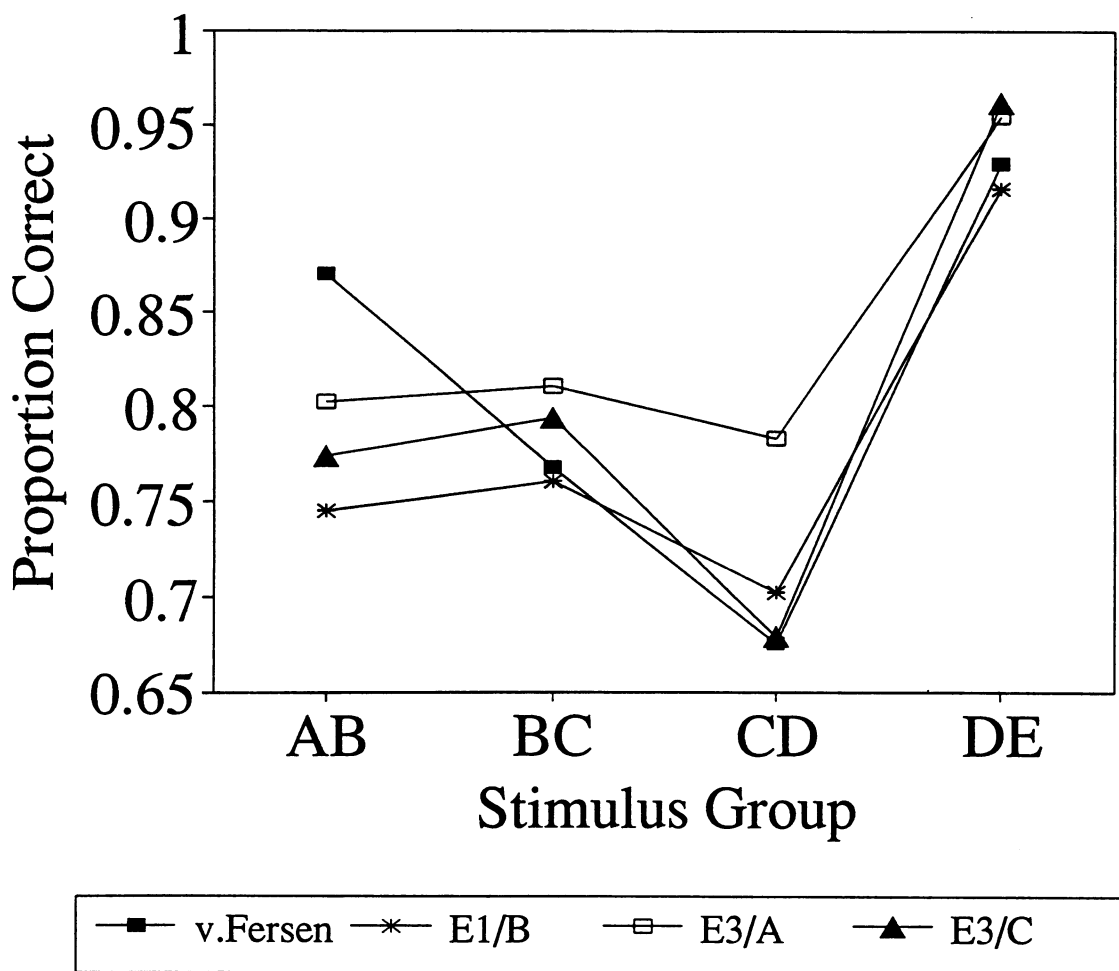


Fig. 11. Filled squares show the proportion of correct responses replotted from Fersen et al. (1991). Remaining series are from the present experiments: asterisks (Experiment 1, Phase B, response-based condition); open squares (Experiment 3, Phase A, response-based condition); and filled triangles (Experiment 3, Phase C, time-based condition).

Phase B) did not generate transitive inference. Finally, analysis of whether transitive inference during BD trials (Experiment 2, Phase B) depended on the stimulus pair in the immediately preceding trial yielded inconclusive results.

The major difference between our data and those of Fersen et al. (1991) is the worse performance of our birds on A+B-. But this difference does not seem to be the clue to transitive inference performance. In Experiment 3 our pigeons often showed good performance on the transitive inference test at the same time they were showing only average A+B- performance (e.g., Birds 158, 169, and 173 in Figure 9 and Birds 160 and 169 in Figure 10).

Thus, these experiments have not succeeded in identifying the necessary and sufficient conditions for the transitive inference effect in a series of overlapping conditional discriminations. What they do show is (a) that the autorun procedure is an effective way to get rapid learning of a series of difficult conditional discriminations; (b) that the response-based autorun eliminates symmetrical end-anchor effects during training, but the time-based procedure does not; (c) that transitive inference does not depend on the presence of symmetrical end-anchor effects; (d) that reinforcement per se does not reliably predict when transitive inference or other effects will occur; and (e) that good discrimination performance during

training with the four stimulus pairs is not sufficient to produce transitive inference effects during testing.

A question remains. Our birds' steady-state performance on training stimulus groups was comparable to that found by Fersen et al. (1991) (see Figure 11), yet we did not always find transitive inference and they did. How can the results from these two studies be reconciled? Recent theoretical work by Davis, Staddon, Machado, and Palmer (in press) on models of recurrent choice suggests a way. In short, they propose a shift away from models of steady-state properties of behavior towards models that emphasize and capture the dynamic processes underlying choice behavior. They argue (theoretically and through specific examples) that dynamic models will help us to understand historical effects and individual differences in behavior, as well as what they call the "transfer properties of a particular training history." The latter highlights a problem that is often disregarded and difficult to study: Two organisms with the same steady-state level of performance at a particular point in time might, in fact, have different histories of reinforcement, and subsequent probes (like extinction or test trials) will reveal the nature of these differences. Analogously, although the subjects in our study and that of Fersen et al. reached similar levels of steady-state behavior (percentage correct on training trials), they obviously attained these levels through very different training procedures and reinforcement histories. Fersen et al.'s birds showed transitive-inference-like behavior during test trials preceded by a training procedure that did not control for the number of unreinforced responses as well as the amount of exposure to each training pair. Test trials following our training procedure, which controlled for these variables, did not consistently produce transitive inference.

We still do not understand why or how transitive inference depends on performance level or the amount of exposure to training stimulus pairs. But, the perspective offered by Davis et al. (in press) suggests that we should study the historical effects of training and carefully ex-

amine variables such as the presentation order of stimulus pairs and reinforcement contingencies, as well as the role of nonreinforced responses, all of which can be studied by using variations of the autorun method. We think that the next step towards understanding the process of transitive inference will involve dynamic procedures like autorun and the testing of particular behavioral models of how these variables produce seemingly complex, inference-like behavior.

REFERENCES

- Bryant, P. E., & Trabasso, T. (1971). Transitive inferences and memory in young children. *Nature*, **232**, 456-458.
- Davis, D. G. S., Staddon, J. E. R., Machado, A., & Palmer, R. G. (in press). The process of recurrent choice. *Psychological Review*.
- Fersen, L., Wynne, C. D. L., Delius, J. D., & Staddon, J. E. R. (1990). Deductive reasoning in pigeons. *Naturwissenschaften*, **77**, 548-549.
- Fersen, L., Wynne, C. D. L., Delius, J. D., & Staddon, J. E. R. (1991). Transitive inference formation in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, **17**, 334-341.
- Gillan, D. J. (1981). Reasoning in the chimpanzee: II. Transitive inference. *Journal of Experimental Psychology: Animal Behavior Processes*, **7**, 150-164.
- McGonigle, B. O., & Chalmers, M. (1977). Are monkeys logical? *Nature*, **267**, 694-696.
- McGonigle, B., & Chalmers, M. (1984). The selective impact of question form and input mode on the symbolic distance effect in children. *Journal of Experimental Child Psychology*, **37**, 525-554.
- McGonigle, B., & Chalmers, M. (1992). Monkeys are rational! *Quarterly Journal of Experimental Psychology*, **45B**, 189-228.
- Sidman, M., Kirk, B., & Willson-Morris, M. (1985). Six-member stimulus classes generated by conditional-discrimination procedures. *Journal of the Experimental Analysis of Behavior*, **43**, 21-42.
- Sidman, M., Rauzin, R., Lazar, R., Cunningham, S., Tailby, W., & Carrigan, P. (1982). A search for symmetry in the conditional discrimination of rhesus monkeys, baboons, and children. *Journal of the Experimental Analysis of Behavior*, **37**, 23-44.
- Trabasso, T., & Riley, C. A. (1975). On the construction and use of representations involving linear order. In R. L. Solso (Ed.), *Information processing and cognition: The Loyola symposium* (pp. 381-410). Hillsdale, NJ: Erlbaum.

Received February 19, 1992
Final acceptance September 8, 1992