Here or There, Now or Then?

Relative Reinforcement and Divided Stimulus Control in Space and Time

Stephanie Elisabeth Gomes-Ng

A thesis submitted in partial fulfilment of the requirements for the degree of
Doctor of Philosophy in Psychology, The University of Auckland, 2019.

ABSTRACT

Stimuli that signal the consequences that are likely to follow behaviour control how, when, and where organisms behave. When more than one stimulus provides such information, each may exert some control over behaviour. Recent research suggests that such divided stimulus control depends on the relative ability of each stimulus to predict future reinforcers. The present thesis extended this finding in four experiments with pigeons. Experiment 1 showed that relative reinforcer rates determine divided control between separately trained stimuli that signalled the location of future reinforcers if the stimuli were spatially separated, but not if they were combined together. Experiment 2 was an extension of Experiment 1, and showed that reinforcer rates determine divided control between spatially separated stimuli that were trained together. In Experiment 3, relative reinforcer probabilities divided control between stimuli that signalled the time of future reinforcers, but such effects were small and time-dependent, probably because elapsed time also competed for behavioural control. Experiment 4 investigated divided control between elapsed time and a visual stimulus in a concurrent-choice schedule, and showed that this division depends on elapsed time and the reliability with which visual stimuli signalled future reinforcer availability, but not on the reliability with which elapsed time signalled future reinforcer availability. Therefore, in general, stimuli that were better predictors of future reinforcers exerted stronger control over behaviour. These experiments extend the relation between relative reinforcer predictability and divided stimulus control to (1) separately trained stimuli, (2) spatially separated stimuli, (3) temporal discriminations, and (4) temporal and non-temporal stimuli. However, the effects of relative reinforcers on divided stimulus control were modulated by several factors, including the spatial configuration of the stimuli, the type of training procedure, changes in previously learned contingencies, elapsed time, and contingency discriminability. Therefore, the present thesis helps to establish the generality of the relation between relative reinforcers and divided stimulus control, but also highlights some of the potential limits of this generality. This thesis is a step towards a more comprehensive understanding of the complexities of stimulus control, and provides a platform for future research to investigate further the mechanisms underlying divided stimulus control.

ACKNOWLEDGEMENTS

First and foremost, thank you to my main supervisor, Dr. Sarah Cowie, and co-supervisor, Professor Doug Elliffe. Thank you for providing the facilities and resources to conduct this research, for the opportunities you have given me that have allowed me to grow and develop as a researcher and teacher, and for trusting me with those opportunities. I am indebted to both of you for all of the time and effort that you have invested into helping me develop and improve as a researcher over the last five years, and I credit much of my success as a graduate student and researcher to you both. Sarah, I am so grateful to have worked under your supervision. Thank you for your never-ending patience, the feedback on countless drafts of my work, and for your constant support and encouragement – you never doubted my ability to complete this PhD, and you made sure to remind me of that whenever I doubted myself. Doug, thank you for your patience in answering all of my statistics questions, and for your wise insights and contributions to every research project. You were always there to answer the questions that I thought were unanswerable. Thank you for teaching me not to split my infinitives and that data are plural – I am a better writer because of you.

Thank you to my fellow students (past and present) of the operant behaviour lab for keeping me company as I completed this thesis, and for listening patiently whenever I babbled about my research. Special mention to John (PhD), Vikki (PhD), Athena, and Peter. John, thank you for the advice you have so patiently given me over the past few years, and for always reminding me to stay calm. Vikki, thank you for your encouragement – I could always count on you to listen whenever I needed someone to talk to. Athena, thank you for your friendship, and for appreciating rice and boba as much as I do. Peter, I enjoyed our many research and non-research related chats, many of which were intellectually stimulating.

To my family, thank you for always supporting me in my academic endeavours, and for enduring the many ups and downs during this PhD. To my friends outside the operant behaviour lab, thank you for reminding me that a world outside of the lab exists. Thank you to my church family, for reminding me why I do what I do. A special thanks to the LIFE South Worship and Creative team – I have been privileged to serve alongside all of you throughout this PhD, and always felt so encouraged by the many stories each of you shared (and thank you for entertaining my many pigeon-related facts!).

Thank you to Sarah Carter, my high-school psychology teacher at UWCSEA, for your encouragement over the past 10-11 years (yes, it’s really been that long!), and for believing in me even when I was an infuriatingly quiet student.

Thank you to the University of Auckland Doctoral Scholarship, which enabled me to complete this PhD.

Finally, this thesis would not have been possible without the help of a few non-human animals. Thank you to my pigeons, without whom this thesis really would not exist. And thank you to Pepper, who kept me sane.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

School of Graduate Studies
AskAuckland Central
Alfred Nathan House
The University of Auckland
Tel: +64 9 373 7599 ext 81321
Email: postgradinfo@auckland.ac.nz

Co-Authorship Form

This form is to accompany the submission of any PhD that contains published or unpublished co-authored work. **Please include one copy of this form for each co-authored work**. Completed forms should be included in all copies of your thesis submitted for examination and library deposit (including digital deposit), following your thesis Acknowledgements. Co-authored works may be included in a thesis if the candidate has written all or the majority of the text and had their contribution confirmed by all co-authors as not less than 65%.

Please indicate the chapter/section/pages of this thesis that are extracted from a co-authored work and give the title and publication details or details of submission of the co-authored work.

Chapter 3 - Experiment 1: Relative Reinforcer Rates and Divided Control by Separately Trained Stimuli

Publication: Gomes-Ng, S., Elliffe, D., & Cowie, S. (2019, in press). Relative reinforcer rates determine pigeons' attention allocation when separately trained stimuli are presented together. Learning & Behavior.

Nature of contribution by PhD candidate	Designed the experiment, programmed all experimental conditions, analysed the data, drafted and revised the manuscript.
Extent of contribution by PhD candidate (%)	95

CO-AUTHORS

Name	Nature of Contribution
Douglas Elliffe	Provided feedback on manuscript drafts, supervised project
Sarah Cowie	Provided feedback on manuscript drafts, supervised project

Certification by Co-Authors

The undersigned hereby certify that:

- the above statement correctly reflects the nature and extent of the PhD candidate’s contribution to this work, and the nature of the contribution of each of the co-authors; and
- that the candidate wrote all or the majority of the text.

Name	Signature	Date
Douglas Elliffe		24/6/19
Sarah Cowie		24/6/19

School of Graduate Studies
AskAuckland Central
Alfred Nathan House
The University of Auckland
Tel: +64 9 373 7599 ext 81321
Email: postgradinfo@auckland.ac.nz

Co-Authorship Form

This form is to accompany the submission of any PhD that contains published or unpublished co-authored work. **Please include one copy of this form for each co-authored work**. Completed forms should be included in all copies of your thesis submitted for examination and library deposit (including digital deposit), following your thesis Acknowledgements. Co-authored works may be included in a thesis if the candidate has written all or the majority of the text and had their contribution confirmed by all co-authors as not less than 65%.

Please indicate the chapter/section/pages of this thesis that are extracted from a co-authored work and give the title and publication details or details of submission of the co-authored work.

Chapter 5 - Experiment 3: Relative Reinforcer Probabilities Divide Stimulus Control in the Multiple Peak Procedure

Publication (submitted): Gomes-Ng, S., Elliffe, D., & Cowie, S. (submitted). Timing compound stimuli: Relative reinforcer probabilities divide stimulus control in the multiple peak procedure. Manuscript submitted for publication.

Nature of contribution by PhD candidate	Designed the experiment, programmed all experimental conditions, analysed the data, drafted and revised the manuscript.
Extent of contribution by PhD candidate (%)	95

CO-AUTHORS

Name	Nature of Contribution
Douglas Elliffe	Provided advice on statistical analyses, provided feedback on manuscript drafts, supervised project
Sarah Cowie	Provided advice on data analysis, provided feedback on manuscript drafts, supervised project

Certification by Co-Authors

The undersigned hereby certify that:

- the above statement correctly reflects the nature and extent of the PhD candidate’s contribution to this work, and the nature of the contribution of each of the co-authors; and
- that the candidate wrote all or the majority of the text.

Name	Signature	Date
Douglas Elliffe		24/6/19
Sarah Cowie		24/6/19

School of Graduate Studies
AskAuckland Central
Alfred Nathan House
The University of Auckland
Tel: +64 9 373 7599 ext 81321
Email: postgradinfo@auckland.ac.nz

Co-Authorship Form

This form is to accompany the submission of any PhD that contains published or unpublished co-authored work. **Please include one copy of this form for each co-authored work**. Completed forms should be included in all copies of your thesis submitted for examination and library deposit (including digital deposit), following your thesis Acknowledgements. Co-authored works may be included in a thesis if the candidate has written all or the majority of the text and had their contribution confirmed by all co-authors as not less than 65%.

Please indicate the chapter/section/pages of this thesis that are extracted from a co-authored work and give the title and publication details or details of submission of the co-authored work.

Chapter 6 - Experiment 4: Reinforcer Predictability and Divided Control by Temporal and Non-Temporal Stimuli

Publication: Gomes-Ng, S., Elliffe, D., & Cowie, S. (2018). Environment tracking and signal following in a reinforcer-ratio reversal procedure. Behavioural Processes, 157, 208-224.

Nature of contribution by PhD candidate	Designed the experiment, programmed all experimental conditions, analysed the data, drafted and revised the manuscript.
Extent of contribution by PhD candidate (%)	90

CO-AUTHORS

Name	Nature of Contribution
Douglas Elliffe	Provided feedback on manuscript drafts, supervised project
Sarah Cowie	Helped with experimental design, provided feedback on manuscript drafts, supervised project

Certification by Co-Authors

The undersigned hereby certify that:

- the above statement correctly reflects the nature and extent of the PhD candidate’s contribution to this work, and the nature of the contribution of each of the co-authors; and
- that the candidate wrote all or the majority of the text.

Name	Signature	Date
Douglas Elliffe		24/6/19
Sarah Cowie		24/6/19

CHAPTER 1

General Introduction

> …antecedent stimuli exert a pervasive influence on operant, as well as respondent, behavior. They serve as the crucial link between current behavior and past reinforcement. They are the direct and immediate determinants of what the individual does at any given moment.
>
> Dinsmoor (1995a, p. 67)

What determines behaviour? It is well established that operant behaviour depends on antecedent stimuli and on the consequences that follow behaviour (Skinner, 1938). In the presence of a stimulus, a subject may emit a behaviour that produces an appetitive (a *reinforcer*, e.g., food delivery) or an aversive (a *punisher,* e.g., electric shock) consequence. As a result of this consequence, the probability of that behaviour re-occurring in the presence of that stimulus increases or decreases respectively[1]; the stimulus comes to 'set the occasion' for behaviour. Similarly, in respondent (classical) conditioning, a stimulus comes to signal an appetitive or aversive outcome through repeated pairings of the stimulus and outcome, and hence the stimulus eventually elicits conditioned responding (e.g., salivation, response suppression; Pavlov, 1927). Thus, in both operant and respondent conditioning, stimuli signal the availability of behaviour-contingent or -noncontingent consequences, and such information then determines how the subject will behave when it encounters the stimulus again in the future – that is, the stimulus serves as the "link between current behavio[u]r and past [consequences]" (Dinsmoor, 1995a, p. 67). This process of behaviour changing as a

---

[1] Recent research suggests that the effects of reinforcers and punishers on behaviour are more complex than this – reinforcers do not always increase the probability of a behaviour, nor do punishers always decrease the probability of a behaviour (see Cowie & Davison, 2016). However, this debate on the functions of reinforcers and punishers is not central to the present thesis, and hence the description provided here will suffice.

result of a stimulus that signals the consequences that will follow is formally termed *stimulus control* (Dinsmoor, 1995a; Morse & Skinner, 1958; Terrace, 1966).

Most empirical studies of stimulus control examine how a single stimulus, such as a coloured keylight signalling the presence or absence of reinforcer deliveries, comes to control the rate, probability, or pattern of behaviour. These studies provide a rather oversimplified view of stimulus control; environments in the natural world are multi-dimensional and hence *multiple* stimuli will typically control behaviour. For example, successful navigation to a goal requires the use of multiple landmarks, and for many humans, the additional instructions provided by a GPS. Social interactions involve a variety of cues, ranging from the words spoken to body language, facial expressions, and tone, volume, and pitch of voice. The ability to categorise objects into 'concepts' (e.g., animals, plants, vehicles, predators, prey) depends on the ability to distinguish the many characteristics that collectively define a concept from those that do not. Reading and writing are skills that depend on recognising series of letters that make up words, and series of words that make up sentences. Even the most mundane of decisions – such as what time to wake up on Saturday, whether to walk or take the bus to work, and what to eat for dinner – depend on more than one stimulus (e.g., the weather, the time of day, the bus timetable, your energy levels, your bank account balance, what ingredients you have in your pantry).

Thus, control by multiple stimuli (*divided stimulus control*) is ubiquitous in the natural world (Dall, Giraldeau, Olsson, McNamara, & Stephens, 2005; Munoz & Blumstein, 2012; Schmidt, Dall, & Van Gils, 2010). Divided stimulus control has also been demonstrated across a range of species, stimuli, and procedures in the laboratory. For example, divided stimulus control has been investigated in pigeons (e.g., Maki & Leith, 1973; Reynolds, 1961; Shahan & Podlesnik, 2006, 2007), monkeys (e.g., Cox & D'Amato, 1982; D'Amato & Fazzaro, 1966), rats (e.g., Lashley, 1938; Matell & Kurti, 2014; Swanton

& Matell, 2011), and humans (e.g., Dube & McIlvane, 1997; Galloway, 1967; Heyman, Grisanzio, & Liang, 2016), among other species. Different procedures have been used, such as go/no-go (e.g., Reynolds, 1961), identity and symbolic matching-to-sample (e.g., Davison & Elliffe, 2010; Maki & Leith, 1973; Shahan & Podlesnik, 2006), search tasks (e.g., Blaisdell, Schroeder, & Fast, 2018; Legge, Madan, Spetch, & Ludvig, 2016; Rodrigo, Gimeno, Ayguasanosa, & Chamizo, 2014), concurrent schedules (e.g., Cowie, Davison, & Elliffe, 2017; Davison & Cowie, 2019; Gomes-Ng, Elliffe, & Cowie, 2018a), and the peak procedure (e.g., Matell & Kurti, 2014), with a range of stimulus dimensions, including colours, line orientations, forms, landmarks, tone frequencies, time, space, number, and speed of movement.

In fact, many experimental paradigms used to study discrimination learning may engender divided stimulus control even if they do not explicitly intend to do so, suggesting that divided stimulus control is more common than it may first appear in basic behavioural research (Pinto, Fortes, & Machado, 2017; Pinto & Machado, 2017). All studies of concept formation or categorisation are studies of divided stimulus control; as noted above, successful concept formation requires control by multiple stimulus characteristics. In procedures that arrange changes in reinforcer location that occur with some temporal regularity (e.g., when the location of a reinforcer switches halfway during the experimental session, as in the mid-session reversal task; Rayburn-Reeves, Molet, & Zentall, 2011), behaviour may come under the control of multiple cues, including elapsed time, additional keylight-colour stimuli, and/or the location or outcome of previous responses (McMillan, Spetch, Sturdy, & Roberts, 2017; Rayburn-Reeves, Moore, Smith, Crafton, & Marden, 2018; Rayburn-Reeves, Qadri, Brooks, Keller, & Cook, 2017; Santos, Soares, Vasconcelos, & Machado, 2019). In a matching-to-sample task, in which subjects are presented with a sample stimulus and then report its identity by choosing a ‘matching’ stimulus from among several comparison stimuli, comparison choice may be controlled not only by the sample stimulus, but also by other variables such as the subject's own behaviour during the sample-stimulus presentation (e.g., Blough, 1959; Urcuioli, 1984, 1985), the sample or comparison stimuli presented in the previous trial (e.g., Edhouse & White, 1988), or events between trials or between sample and comparison presentation (e.g., Grant, 1981; Grant & Roberts, 1973; Pinto et al., 2017; Pinto & Machado, 2017). When the appearance of a discriminative stimulus changes at regular temporal intervals (e.g., alternating between two colours every 0.5 s), the frequency of stimulus change and the number of changes may both exert some discriminative control over behaviour (e.g., Fetterman & Killeen, 2011; W. A. Roberts & Boisvert, 1998; W. A. Roberts, Coughlin, & Roberts, 2000). Even in tasks in which only a single, apparently simple stimulus is presented (e.g., a red light), other features of the stimulus (e.g., its brightness) or of the wider environment may exert some control over behaviour (e.g., Emmerton, 1998; Kraemer, 1993; Lionello & Urcuioli, 1998; Reynolds, 1961).

In this way, all environments, and all stimuli within environments, are multi-dimensional. Hence, most empirical studies of discrimination learning provide the opportunity for divided stimulus control to develop, even if such division does not actually occur (Kehoe & Gormezano, 1980). The question, then, is *why* does divided stimulus control develop (or not develop), and what kinds of variables determine the division? The answer to this question would not only deepen our understanding of the basic processes that underlie learning and behaviour, but would also provide insight into why some applied behavioural interventions do or do not succeed in establishing or maintaining behaviours, or in generalising behaviours to other contexts (see e.g., Koegel & Rincover, 1976; Rincover & Koegel, 1975). Research on divided stimulus control can also help to inform the development or improvement of interventions to remediate deficits in academic, social, behavioural, and linguistic development, as such deficits have been linked to the failure of behaviour to come under the control of multiple cues (see e.g., Barthold & Egel, 2001; Birnie-Selwyn & Guerin, 1997; Chiang & Carter, 2008; Dube, McDonald, McIlvane, & Mackay, 1991; Dube & McIlvane, 1997, 1999; Dunlap, Koegel, & Burke, 1981; Lovaas, Koegel, & Schreibman, 1979; Lovaas, Schreibman, Koegel, & Rehm, 1971; Ploog, 2010; Reed, Broomfield, McHugh, McCausland, & Leader, 2009; Schreibman & Lovaas, 1973).

Therefore, a comprehensive understanding of behaviour – how it is shaped and maintained, and how it can be modified – requires a thorough investigation of how, why, and under what conditions multiple stimuli come to exert joint control over behaviour. Despite this, investigations of the processes underlying divided stimulus control are less common than investigations of the processes underlying other ubiquitous behavioural phenomena, such as concurrent choice or signal detection (Davison, 2018b). Thus, the many variables that might influence how and why stimulus control is divided between multiple stimuli are not well understood. The present thesis aims to extend such understanding. Specifically, this thesis examines the role of *reinforcer frequency* in determining the extent to which multiple stimuli come to control operant behaviour in space and time.

CHAPTER 2

Literature Review

Pavlov (1927) noted that in respondent conditioning, the associative strength of a stimulus depended on the presence of other stimuli; specifically, whereas conditioned responding to a stimulus trained alone was strong, that same stimulus supported much weaker conditioned responding when trained in compound with another stimulus. Similarly, in an operant paradigm, Reynolds' (1961) seminal study of "attention[2] in the pigeon" demonstrated strong control by one stimulus at the apparent expense of control by other stimuli. In Reynolds' experiment, key-pecks to a white triangle superimposed on a red background (the S+ stimulus) were reinforced, whereas pecks to a white circle on a green background (the S- stimulus) were never reinforced. After this *go/no-go* discrimination training, Reynolds presented the colour (red and green) and form (triangle and circle) elements individually in extinction and measured response rates to each, in order to assess the extent to which the colour and form dimensions controlled behaviour. For both pigeons in the experiment, rates of key-pecking were high to the S+ element and low to the S- element along one stimulus dimension, whereas response rates to the S+ and S- elements were low and nondifferential along the other dimension. Thus, Reynolds concluded that only one stimulus dimension controlled behaviour, colour for one pigeon and form for the other, even though both dimensions had signalled reinforcer availability during discrimination training (see also Lashley, 1938).

---

[2] The terms "attention" and "stimulus control" are sometimes used interchangeably in the behaviour-analytic literature. This is because attention is defined as the "controlling relation" (Skinner, 1953, p. 123) between a stimulus and behaviour; that is, only those stimuli that are attended to can control behaviour. Thus, the general consensus is that "any observation of stimulus control is [also] an observation of attention" (Johnson & Cumming, 1968, p. 157). Because of the connotations and assumptions that may be associated with the more cognitive term "attention", and because of the different ways that the term attention is used throughout the wider literature in psychology, the present thesis uses "stimulus control" throughout. Further discussion of the links between attention and stimulus control can be found in Dinsmoor (1985).

Reynolds' (1961) findings were some of the first to suggest that discrimination training with multiple stimuli was not necessarily sufficient to establish discriminative control by all of the stimuli (see also Born & Peterson, 1969; Born & Snow, 1970; D'Amato & Fazzaro, 1966; Johnson & Cumming, 1968; Reynolds & Limpo, 1969). In other words, the mere presence of multiple stimuli in the environment does not guarantee that all of those stimuli will control behaviour (Dinsmoor, 1985; Johnson & Cumming, 1968). This is well exemplified by a brief practical exercise: Look around you – it is highly unlikely that *all* of the stimuli that are present in your current environment are influencing your behaviour; instead, a subset of the stimuli is probably controlling your behaviour. What, then, determines which stimuli in the environment control behaviour, and the strength of control by those stimuli? Since Reynolds, many researchers of divided stimulus control have investigated the conditions under which one of several stimuli controls behaviour (*selective* stimulus control) versus those conditions under which more than one stimulus controls behaviour (*divided* stimulus control). This chapter provides an overview of such research.

2.1 Re-evaluating Reynolds (1961): Some Procedural Concerns

One of the earliest findings in the literature on operant divided stimulus control was that the method and/or measure used to assess such division could influence whether stimulus control appeared selective or divided. Reynolds (1961) took differential responding to the stimulus elements along a dimension as evidence of control by that dimension, whereas he took low and nondifferential responding to the elements along a dimension as evidence of the absence of control by that dimension. However, subsequent research called the latter into question. Farthing and Hearst (1970) arranged a go/no-go task with compound S+ and S- stimuli comprised of a colour and a line orientation, and assessed divided stimulus control in post-discrimination test sessions in which each element was presented individually (i.e., as in

Reynolds, 1961) or in which S+ and S- elements were combined into novel compounds. In tests with individual elements, only the colour dimension appeared to control behaviour, as response rates to S+ and S- colour elements differed whereas responding to the line-orientation elements did not. However, in tests with novel compounds, response rates were higher to compounds containing the S+ line orientation than to compounds containing the S- line orientation, implying some control by line orientation.

Similarly, Wilkie and Masson (1976) found that following go/no-go discrimination training with compound colour-form stimuli, stimulus control appeared to be selective – in favour of colour over form – in post-discrimination tests with individual elements. However, in subsequent training sessions in which responses to both the S+ and S- form elements were reinforced, acquisition was faster and overall response rates were higher to the previous S+ form than to the previous S- form. This difference in responding between the form elements was related to the subjects' initial experience with the go/no-go task, as rates of acquisition and overall responding to the form elements were similar for control subjects without initial go/no-go training. Based on these findings, Wilkie and Masson suggested that "conclusions about attention [i.e., divided stimulus control; see Footnote 2] based on the lack of differential pecking in the nonreinforcement test [with individual stimulus elements] may not be appropriate" (p. 207). Born, Snow, and Herbert (1969) expressed similar sentiments, suggesting that response rates in post-discrimination test sessions may not necessarily provide an accurate reflection of levels of stimulus control during discrimination training with compound stimuli (see also Born & Snow, 1970; Reynolds & Limpo, 1969).

A further, and more general, concern with using the go/no-go task to study divided stimulus control is that the strength of stimulus control is usually measured in terms of overall response rates (e.g., Born et al., 1969; Farthing & Hearst, 1970; Reynolds, 1961; Wilkie & Masson, 1976). Overall response rates mask any evidence of stimulus control in other measures, such as response latencies, inter-response times, or patterns of responding across time (see e.g., Gray, 1976; Rand, 1977). For example, overall response rates may be similar in fixed-interval (FI) and variable-interval (VI) schedules, but analyses of responding across time since each reinforcer delivery reveal substantial differences in patterns of responding between FI and VI schedules (Ferster & Skinner, 1957). In an FI schedule, response-contingent reinforcers are delivered at regular intervals, and so response rates increase gradually across time after each reinforcer delivery. This ‘scalloped’ pattern of responding does not occur in VI schedules, in which the time between reinforcers varies unpredictably; instead, VI schedules engender high and constant response rates. Hence, overall response rates alone can provide an incomplete picture of stimulus control, as stimulus control may also, or alternatively, be expressed in other measures of behaviour (see Bickel & Etzel, 1985 for further discussion).

Relatedly, overall response rates are sensitive to a range of factors unrelated to the stimuli. To illustrate, consider pacing schedules, in which reinforcers are delivered contingent on emitting few responses within a specified time frame (e.g., 5 s) or on emitting many responses within that time frame. The former will produce a low response rate and the latter a high response rate (Nevin, 1974). Thus, if two discriminative stimuli are associated with these pacing schedules, the difference in response rates will be large. Now consider VI schedules. If two stimuli are associated with two VI schedules, the difference between response rates to the stimuli will likely be smaller than if the stimuli are associated with the aforementioned pacing schedules. Does this imply that the strength of control by the stimuli differs with pacing and VI schedules? Not necessarily – the differences in response rates with pacing versus VI schedules may be related to the characteristics of the reinforcer schedules, rather than to the strength of stimulus control.

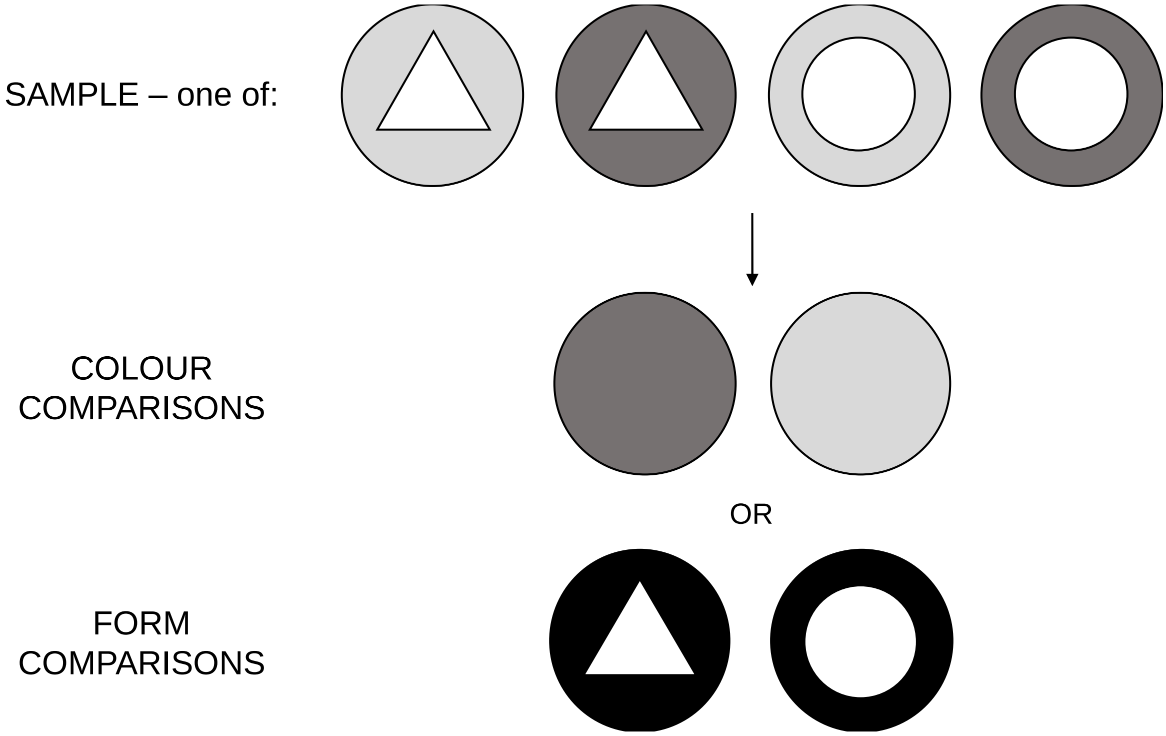

Figure 2.1. Diagram depicting the delayed matching-to-sample (DMTS) task used to study divided stimulus control. In this example, the sample stimulus is a compound comprised of a colour (represented by the light and dark grey circles) and a form (triangle or circle), and comparison stimuli are either two colours or two forms.

Maki and Leuin (1972; see also Maki & Leith, 1973) proposed an alternative procedure – the *delayed matching-to-sample* (DMTS) task – to study divided stimulus control (see Zentall, 2005b for a brief review). Most recent behaviour-analytic studies of divided stimulus control employ the DMTS task. In the typical DMTS task, which is depicted in Figure 2.1, subjects are presented with a compound 'sample' stimulus (e.g., a form superimposed on a coloured background), and following its offset, are presented with a choice between several 'comparison' stimuli from one of the compound-stimulus dimensions (e.g., two colours or two forms). One of the comparison stimuli matches one of the sample-stimulus elements, either identically or symbolically, and choice of the matching comparison is reinforced. The accuracy with which subjects report the identities of the elements along a dimension provides an indication of the degree of control by that dimension. Chance performance (i.e., 50% correct in a two-comparison procedure) suggests no control by that dimension, and higher accuracies indicate stronger control. Thus, unlike the go/no-go procedure, no post-discrimination testing is required, and stimulus control is not assessed using overall response rates in the DMTS task. Hence, DMTS may provide a more accurate indication of the division of control between the dimensions of a compound discriminative stimulus.

2.2 Theories of Divided Stimulus Control

Early theories of divided stimulus control were concerned with the processes underlying the acquisition of conditioned responding in associative learning (e.g., Mackintosh, 1975). These theories proposed that trial-by-trial changes in the associative strength of a conditioned stimulus depended on the sensory capacities of the subject, the characteristics of the stimuli, and the correlation between each stimulus and the trial outcome (e.g., food delivery or electric shock). For example, according to Mackintosh (1975), each stimulus has its own learning-rate parameter, ∝, which determines the amount of attention to, and hence the strength of control by (see Footnote 2), that stimulus. The initial value of ∝ depends jointly on the subject’s sensory capacities (e.g., visual versus auditory acuity) and on stimulus characteristics such as intensity (e.g., brightness or loudness). During acquisition, the value of ∝ changes depending on the outcome of each trial; ∝ increases if the stimulus is correlated with a unique change in trial outcome not signalled by the other stimuli in the environment, whereas ∝ decreases if other stimuli in the environment are better predictors of trial outcome. Thus, Mackintosh’s theory predicts that the subject’s sensory capacities and the nature of each stimulus determine the initial associability of a stimulus, and that stimuli that are better predictors of trial outcome then come to gain more associative strength during acquisition.

More recent theories of divided stimulus control have been concerned with steady-state operant behaviour (i.e., behaviour after subjects have learned the contingencies), rather than with acquisition in associative learning. Shahan and Podlesnik (2006, 2007) suggested that the Generalised Matching Law (GML; Baum, 1974), which describes the relation between reinforcer ratios and response ratios in a concurrent-choice schedule, could be applied to describe the effects of reinforcers on the division of control between stimuli. According to the GML, the effects of reinforcer ratios on choice can be described by a linear function with slope and y-intercept parameters that represent the extent to which changes in reinforcer ratios cause changes in choice (*sensitivity to reinforcement;* Lobb & Davison, 1975) and inherent bias towards one alternative, respectively. Thus, in applying the GML to divided stimulus control, Shahan and Podlesnik suggested that the strength of control by a stimulus is proportional to the reinforcer rate associated with that stimulus relative to the reinforcer rates associated with other stimuli, and that sensory properties and stimulus characteristics can bias stimulus control towards some stimuli over others.

Davison and Elliffe (2010; see also Davison, 2018a) took a similar approach to Shahan and Podlesnik (2006, 2007), except that they emphasised the importance of *contingency discriminability* (Davison & Nevin, 1999) in determining divided stimulus control. Specifically, Davison and Elliffe argued that divided stimulus control may not be perfectly proportional to arranged or obtained relative reinforcer rates because of imperfect discrimination of stimulus-reinforcer and response-reinforcer contingencies. As a consequence of such imperfect discrimination, reinforcers are occasionally ‘misallocated’ from one stimulus or response to another stimulus or response. The extent of such misallocation depends on the subject’s ability to discriminate the stimulus-reinforcer and response-reinforcer contingencies – an ability that depends on the subject’s sensory capacities, on the characteristics of the stimuli, and on the reinforcer contingencies arranged (see Davison & Nevin, 1999; see Chapter 4 of this thesis for a more detailed description of Davison and Elliffe's model).

Despite their differences, it is clear that the theories of divided stimulus control described here all assume that the same variables play a key role in determining the division of control between multiple stimuli. These variables are (1) the sensory capacities and phylogenetic history of the subject, (2) the characteristics of the stimuli, and (3) the ability of each stimulus to predict reinforcers. (1) and (2) determine the initial associability of each stimulus (Mackintosh, 1975) or overall bias towards some stimuli over others (Davison & Elliffe, 2010; Shahan & Podlesnik, 2006, 2007), and (3) modulates these initial levels of control. We now turn to empirical research that has investigated how each of these variables contributes to divided stimulus control. Because the experiments reported in subsequent chapters of this thesis examine the effects of reinforcers on divided stimulus control, a brief overview of research examining the first two variables is provided, and research that has investigated how relative reinforcer predictability divides stimulus control is described in more detail.

2.3 Sensory Capacities and Phylogenetic History

Subjects often appear to be predisposed to use some types of stimuli over others in conditional discriminations (Baron, 1965; Seligman, 1970). Such a predisposition has been attributed to the sensory capacities or phylogenetic history of the subject. Different species have different dominant sensory modalities; for example, humans are visual-dependent (Colavita, 1974; Posner, Nissen, & Klein, 1976), as are other primates (Kirk & Kay, 2004) and most birds (Hodos, 1993), whereas olfaction is the primary sensory modality of rats (Jennings & Keefer, 1969; Slotnick & Katz, 1984) and dogs (Gadbois & Reeve, 2014; Thesen, Steen, & Doving, 1993). In general, stimuli from a subject's dominant sensory modality are more likely to control behaviour than stimuli from other modalities (e.g., Colavita, 1974; Colombo & Graziano, 1994; Duncan & Slotnick, 1990; Gliner, Pick, Pick, & Hales, 1969; Hale & Green, 1976; Kraemer & Roberts, 1985; M. Meltzer & Masaki, 1973; Randich, Klein, & LoLordo, 1978). Thus, in pigeons, the visual element of a visual-auditory compound stimulus is more likely to control behaviour (Kraemer & Roberts, 1985; Randich et al., 1978).

In addition, some stimulus dimensions are more ecologically valid than others, depending on the subject’s phylogenetic history. Some findings suggest that those dimensions that are more ecologically valid may be more likely to control behaviour (Baker, 1968; Baron, 1965; Brodbeck, 1994) – for example, spatial cues tend to exert stronger control over behaviour than non-spatial cues for food-storing birds, whereas spatial and non-spatial cues appear to exert more equal control for non-food-storing birds (Brodbeck, 1994; Brodbeck & Shettleworth, 1995; Clayton & Krebs, 1994; Shettleworth & Westwood, 2002). One explanation for this difference is that food-storing birds may have evolved more complex spatial-cognitive mechanisms that enable them to store and recover a large number of food caches from various locations, and their heavy reliance on spatial over visual cues may be adaptive given that visual cues in the wild are likely to change over time and hence are probably not reliable cues for locating previously cached food (see e.g., Brodbeck & Shettleworth, 1995; Herborn, Alexander, & Arnold, 2011; Olson, 1991; Sherry & Duff, 1996; Sherry & Schacter, 1987; Shettleworth & Hampton, 1998; Shettleworth & Westwood, 2002). Therefore, some stimulus dimensions may control behaviour more readily than others due to the subject’s sensory capacities or phylogenetic history. As a result, stimulus control may be selective, or divided unequally (see also Kraemer & Roberts, 1987).

2.4 Stimulus Characteristics

2.4.1 Stimulus Intensity

The physical characteristics of a stimulus also determine the extent to which it controls behaviour. In general, stimuli that are easier to discriminate exert stronger control than those that are harder to discriminate (Davison & Nevin, 1999; Dinsmoor, 1995b; Johnson & Cumming, 1968). For example, one of the most well established findings in the literature on divided stimulus control is that more intense stimuli – such as those that are brighter, louder, or larger – exert stronger control over behaviour than less intense stimuli (*overshadowing;* Pavlov, 1927). Johnson (1970) trained pigeons in a go/no-go task with compound stimuli comprised of dim, moderate, or bright vertical or horizontal lines superimposed on a coloured background. To assess divided stimulus control, Johnson ran generalisation tests in which the line orientation or wavelength was varied systematically along its respective dimension. These generalisation tests revealed that both dimensions controlled behaviour, as response rates decreased systematically as the line orientation or wavelength became more different from the S+ element. However, the extent of change in response rates along each dimension depended on the brightness of the lines – the brighter the lines, the larger the change in response rates along the line-orientation dimension in generalisation tests, and the smaller the change in response rates along the wavelength dimension. This suggests that control by line orientation was stronger and control by wavelength was weaker when the lines were brighter. Such overshadowing by more intense stimuli has been replicated many times across species, procedures, and stimulus dimensions, in both operant and respondent conditioning paradigms (see e.g., Denton & Kruschke, 2006; Feldman, 1975; Fields, 1978, 1979; Fields, Bruno, & Keller, 1976; Hall, Mackintosh, Goodall, & Dal Martello, 1977; Leader, Loughnane, McMoreland, & Reed, 2009;

Mackintosh, 1971, 1976; Miles & Jenkins, 1973; Rodrigo, Gimeno, Ayguasanosa, & Chamizo, 2014).

2.4.2 Stimulus Disparity

Chase and Heinemann (1972) reinforced pigeons' left key-pecks in the presence of one compound stimulus comprised of a luminance (brightness) dimension and a tone dimension, and right key-pecks in the presence of a different luminance-tone stimulus. Subsequent generalisation tests indicated that the luminance dimension exerted stronger control and the tone dimension exerted weaker control over choice when the difference between the luminances used in training was larger, compared with when it was smaller. Similarly, White (1986) arranged a concurrent-choice task in which responses to one of two wavelength stimuli were reinforced in the presence of one of two line-orientation stimuli, and responses to the other wavelength were reinforced in the presence of the other line orientation. Hence, accurate performance required control by both stimulus dimensions. Across conditions, White varied the difference between the wavelengths and between the line orientations. As these stimulus disparities increased, so did the frequency of correct responses, implying stronger control by the stimulus dimensions. Thus, the physical disparity between the elements along a stimulus dimension determines the extent of control by that dimension, with larger disparities generally engendering stronger control (see also Doran & Holland, 1979; Lea, Pothos, Wills, Leaver, Ryan, & Meier, 2018; Lovejoy & Russell, 1967; Nevin, Cate, & Alsop, 1993).

2.4.3 Duration of Exposure

In the DMTS task (Figure 2.1), accuracy on both dimensions of the compound stimulus increases as the duration of the sample-stimulus presentation increases (M. F. Brown, 1991; M. F. Brown & Morrison, 1990; Cook, Riley, & Brown, 1992; Langley & Riley, 1993; Maki & Leith, 1973; Maki & Leuin, 1972; Oscar-Berman & Bonner, 1985;

Shahan & Podlesnik, 2007). Thus, control by the dimensions of a compound stimulus appears to increase as the duration of exposure to the compound increases. Additionally, Langley and Riley (1993) showed that the *difference* between accuracy on each stimulus dimension increased as the sample duration increased, suggesting that stimulus control became increasingly selective as the duration of exposure increased (see also Zentall, Sherburne, & Zhang, 1997, Experiment 1). However, other studies suggest that although overall accuracy on each dimension increases with longer sample presentation in the DMTS task, the relative difference between accuracy on each dimension remains the same (e.g., M. F. Brown & Morrison, 1990; Kraemer & Roberts, 1987; Maki & Leith, 1973). Therefore, the duration of exposure to the stimuli affects overall levels of stimulus control, but appears not to have a robust effect on relative levels of control between stimuli.

2.4.4 Stimulus Configuration

Another stimulus characteristic that may affect the division of control between multiple stimuli is their spatial or temporal configuration. Spatially, two stimuli may be combined together into a single “unified” stimulus (e.g., a coloured triangle), they may be superimposed (e.g., a white triangle on a coloured background), or they may be spatially separated by various distances (e.g., a white triangle and a coloured stimulus displayed side-by-side). Lamb and Riley (1981) found that in a DMTS task with compound stimuli comprised of a colour and line orientation, matching accuracy on both dimensions was relatively equal with unified compounds, but was higher on the colour dimension than on the line-orientation dimension when they were spatially separated. The discrepancy in accuracy between dimensions increased as the spatial separation between them increased, suggesting that stimulus control becomes more selective as the stimuli become more spatially separated (see also Cohen & Hachey, 1977; Hale & Green, 1976; Hinson, Cannon, & Tennison, 1999; Kraemer, Mazmanian, & Roberts, 1987; Lamb, 1988; Leith & Maki, 1977; Riley, 1984;

Riley & Leith, 1976; Sutherland & Andelman, 1967). However, other research has not replicated this result (e.g., Wolfgang & Richardson, 1973; Werner, Tiemann, Cnotka, & Rehkämper, 2005), and Cook et al. (1992) found that only the *proximity* between the stimulus elements, not their configuration per se (i.e., unified, superimposed, elements side-by-side), determined divided stimulus control (see also Zentall, 2005a). Hence, stimulus configuration may only influence divided stimulus control if the spatial distance between stimuli differs.

Just as stimuli may be closer together or farther apart in space, they may be closer together or farther apart in time. The effects of temporal configuration on divided stimulus control appear to be somewhat more systematic than the effects of spatial configuration. In general, the closer in time stimuli are presented, the more equal the division of control between them. This is evidenced by research demonstrating that conditional discriminations involving compound stimuli are harder for subjects to learn if the stimulus elements are presented sequentially rather than simultaneously, and are harder to learn if the elements are presented sequentially without overlap (i.e., each is removed before the presentation of the next element) than if the elements are presented sequentially with overlap (e.g., Thomas, Cook, & Terrones, 1990). Additionally, when stimuli are presented sequentially, stimulus control generally becomes more selective – typically, in favour of the most recent stimulus – as the temporal separation between them increases (e.g., B. L. Brown, Hemmes, & De Vaca, 1992; Cowie, Davison, & Elliffe, 2011, 2017; Fairhurst, Gallistel, & Gibbon, 2003; Gomes-Ng, Elliffe, & Cowie, 2018a; Matell, De Corte, Kerrigan, & DeLussey, 2016; Miranda-Dukoski, Davison, & Elliffe, 2014). Therefore, separation of stimuli in space or time appears to reduce the likelihood of divided stimulus control.

2.5 Reinforcer Predictability

In addition to sensory capacities, phylogenetic history, and the physical characteristics of the stimulus, the ability of a stimulus to predict reinforcers reliably and non-redundantly, relative to other stimuli in the environment, also determines the division of stimulus control. In fact, such *reinforcer predictability* is thought to be a crucial determinant of stimulus control (Cowie & Davison, 2016; Gallistel, 2003; Rescorla, 1968; Shahan, 2010; Shahan & Cunningham, 2015; Wagner, 1969; Wagner, Logan, Haberlandt, & Price, 1968). Indeed, the effects of learning and reinforcement history on divided stimulus control can override those of sensory capacities, phylogenetic history, and/or physical stimulus characteristics. For example, stimuli that have been previously paired with phylogenetically important outcomes (e.g., food, water, electric shock, loud noise) more readily gain strong control over behaviour compared with other stimuli (e.g., Denniston, Miller, & Matute, 1996; Miller & Matute, 1996; Oberling, Bristol, Matute, & Miller, 2000), and apparently inherent weak control by a stimulus dimension can be strengthened by differential training with that dimension (e.g., Newman & Benefield, 1968; von Fersen & Lea, 1990). Similarly, less intense stimuli can exert stronger control over behaviour than more intense stimuli if the less intense stimuli are more strongly correlated with trial outcome, or if they are pre-trained before being compounded with the more intense stimuli (e.g., Klein, Weston, McGee-Davis, & Cohen, 1984; Le Pelley, Beesley, & Griffiths, 2014). Thus, some researchers argue that differences in the readiness with which stimuli control behaviour may sometimes be better explained in terms of differences in the ability of each stimulus to predict trial outcomes, rather than in terms of biology, phylogeny, or stimulus characteristics (Segal & Harrison, 1978).

2.5.1 Stimulus Validity in Conditional Discriminations

Wagner et al. (1968) published a series of experiments that were seminal in illustrating the importance of reinforcer predictability in the division of stimulus control. In

Wagner et al.’s experiments, compound stimuli comprised of different tone elements but the same light element signalled food delivery (Experiment 1) or electric shock (Experiments 2 and 3). For one group of subjects (Group Correlated), the trial outcome was delivered in the presence of one compound and not in the presence of the other compound. Thus, the tones differentially signalled food or electric shock availability, whereas trial outcome was nondifferential with respect to the light element. For another group of subjects (Group Uncorrelated), food or electric shock was delivered equally often in the presence of either compound, and so trial outcome was signalled nondifferentially by both the tones and the light. Wagner et al. found that rates of conditioned responding to the light element were lower for Group Correlated than for Group Uncorrelated, even though the light signalled the same probability of trial outcome (.5) for both groups. They explained this difference between groups in terms of the *validity* of the light element relative to the tone elements. For Group Correlated, the tone was a better predictor of trial outcome than the light, whereas the light and tone were equally predictive of trial outcome for Group Uncorrelated. As a result, control by the light was weakened by the presence of a more valid stimulus for Group Correlated, but not for Group Uncorrelated (see also Wagner, 1969).

Subsequent studies have confirmed Wagner et al.’s (1968) general conclusion that stimuli that are more valid or relevant to a discrimination typically exert stronger control over behaviour than less valid or irrelevant stimuli. In addition, such studies have also shown that shifts in divided stimulus control can be induced by varying the relevant and irrelevant stimuli. For example, Leith and Maki (1975) found that in the DMTS task (Figure 2.1), stimulus control was divided between the compound-stimulus dimensions if both dimensions served as comparisons, whereas when only one dimension served as comparisons, only that dimension controlled behaviour. That is, divided stimulus control was evident only when multiple stimulus dimensions were relevant to the discrimination (see also e.g., Berg &

Grace, 2011; Blough, 1969; Castro & Wasserman, 2014, 2016, 2017; Chase, 1968; De Lillo, Spinozzi, Palumbo, & Giustino, 2011; Farthing, 1972; Fremouw, Herbranson, & Shimp, 1998; Heinemann, Chase, & Mandell, 1968; Lea, Wills, Leaver, Ryan, Bryant, & Millar, 2009; McMillan & Roberts, 2013; Vyazovska, Navarro, & Wasserman, 2018; Vyazovska, Teng, & Wasserman, 2014; Wilkie, 1973; B. A. Williams, 1984).

2.5.2 Effects of Prior Training on Subsequent Discrimination Learning

Findings from a variety of experiments point to the importance of previous reinforcement history in determining divided stimulus control in conditional discrimination learning. In studies of blocking (Kamin, 1969), a stimulus previously paired with a phylogenetically important outcome (e.g., food or electric shock) is compounded with a novel stimulus, and the compound is then paired with the outcome. Typically, the novel stimulus exerts much weaker control over behaviour compared with when it is trained alone (e.g., Johnson, 1970; Johnson & Cumming, 1968; Kamin, 1969; Seraganian & vom Saal, 1969; Wagner, 1969). Kamin (1969) suggested that such blocking occurs because the pre-trained stimulus comes to predict trial outcome reliably, and hence exerts strong control over behaviour. As a result, when the pre-trained stimulus is compounded with a novel stimulus, the novel stimulus does not provide any additional information about trial outcome, and so does not gain strong control over behaviour (see also e.g., Boutros, Davison, & Elliffe, 2009, 2011; Egger & Miller, 1962, 1963; McCausland, Menzer, Dempsey, & Birkimer, 1967; McLinn & Stephens, 2006 for similar discussions). In support of this, blocking is attenuated if the introduction of the novel stimulus is accompanied by a change in trial outcome (e.g., a change in rate, magnitude, or type of outcome; Dickinson, Hall, & Mackintosh, 1976; Feldman, 1971; Kamin, 1969; Mackintosh, Bygrave, & Picton, 1977; Mackintosh, Dickinson,

& Cotton, 1980; Mackintosh & Turner, 1971; Stickney & Donahoe, 1983; but see also B. A. Williams, 1994 for blocking despite a change in type of reinforcer)[3].

Studies of intradimensional and extradimensional shift learning have shown that previous experience with stimulus dimensions in one conditional discrimination task can influence learning in subsequent discriminations with those dimensions (see Dias, Robbins, & Roberts, 1996; Mackintosh, 1974 for reviews). In such studies, subjects are first trained in a conditional discrimination in which one dimension of a compound stimulus differentially signals reinforcer availability, and the other stimulus dimension is irrelevant to the discrimination. Thereafter, a new discrimination is trained using novel stimuli from the same dimensions, and the relevant and irrelevant dimensions either remain unchanged (*intradimensional* shift) or are reversed (*extradimensional* shift). Although there are exceptions (e.g., Couvillon, Tennant, & Bitterman, 1976; Hall & Channell, 1985; Shepp & Schrier, 1969), acquisition of the new discrimination is generally faster with intradimensional shifts than with extradimensional shifts (e.g., Cuell, Good, Dopson, Pearce, & Horne, 2012; Klosterhalfen, Fischer, & Bitterman, 1978; Mackintosh & Little, 1969; D. Meltzer & Robertson, 1989; A. C. Roberts, Robbins, & Everitt, 1988; Shepp & Eimas, 1964; Shepp & Schrier, 1969; Singh & Beale, 1978; Uengoer & Lachnit, 2012). Thus, stimulus dimensions that differentially signal reinforcer availability appear to control behaviour more readily in future discriminations than those dimensions that have previously not signalled reinforcer availability.

---

[3] Blocking is well accounted for in respondent conditioning by the Rescorla-Wagner model (Rescorla & Wagner, 1972), which predicts that a stimulus increases in associative strength towards an asymptote as training progresses. When that stimulus is then compounded with a novel stimulus, as in blocking preparations, the associative strength of the compound is already close to asymptote due to the previously trained element. As a result, the novel element gains little associative strength, resulting in blocking. Because the present thesis examines divided stimulus control in operant conditioning paradigms, further discussion of compound conditioning in respondent paradigms and of the Rescorla-Wagner model is beyond the scope of this thesis.

Relatedly, Ryan, Hemmes, and Brown (2011) showed that the degree of congruency between past and present reinforcer contingencies modulates divided stimulus control. Ryan et al. arranged a symbolic DMTS procedure in which human participants matched individual shape or colour stimuli to one of two nonsense syllables (“VEC” or “BUP”). One of the stimulus dimensions was designated as the target dimension, and the other as the non-target dimension. For some participants, reinforcers were only delivered following correct responses for both dimensions (Group Incongruent), whereas for other participants, some reinforcers were delivered following error responses for the target dimension (Group Partially Congruent). After this initial training, the shape and colour stimuli were compounded together, and participants matched the compounds to the same nonsense syllables. In compound training, the sample-comparison associations learned during initial training remained intact for the non-target dimension and were reversed for the target dimension. The non-target dimension exerted selective control over behaviour for Group Incongruent, whereas stimulus control was divided between the target and non-target dimensions for Group Partially Congruent. That is, when the contingencies associated with one dimension of a compound stimulus were reversed, control by the reversed dimension was stronger when there was some history of reinforcement for choosing the newly correct comparison during initial training.

2.5.3 Shifts in Divided Stimulus Control Following Revaluation of One Stimulus

In addition to prior reinforcement history, additional training following exposure to a compound stimulus can shift control between the elements of the compound. When one dimension or element of a compound stimulus exerts weak control over behaviour, such control can be strengthened by additional discrimination training with that dimension or element (Newman & Benefield, 1968; von Fersen & Lea, 1990). Similarly, control by one element of a compound stimulus may emerge following a reduction in the reinforcer rate associated with a different element of the compound. For example, Blaisdell, Schroeder, and Fast (2018) trained pigeons to locate a goal using one of two landmarks which differed in their distance to the goal, and then presented the landmarks simultaneously and found that the landmark closer to the goal exerted strong control over search behaviour. Blaisdell et al. then ran further training sessions with each individual landmark, except that no reinforcers were delivered for correct responses in trials with the closer landmark (the reinforcer rate for the farther landmark remained unchanged). Such training resulted in a shift in stimulus control from the closer landmark to the farther landmark when both landmarks were again presented simultaneously (see also e.g., Broomfield, McHugh, & Reed, 2010; Cole, Barnet, & Miller, 1995; Kelly, Leader, & Reed, 2015; Reed, Broomfield, McHugh, McCausland, & Leader, 2009; Reed, Reynolds, & Fermandel, 2012). Such emergence of control by one stimulus following ‘revaluation’ of another stimulus suggests that divided stimulus control depends on the reinforcer rate associated with one element relative to other stimuli.

2.5.4 Effects of Relative Reinforcer Predictability on Divided Stimulus Control

The studies described above demonstrate that stimuli that differentially signal reinforcer availability (i.e., that are relevant or valid to a discrimination) are more likely to control behaviour than irrelevant stimuli, and that the reinforcement history associated with a stimulus dimension can have a strong influence on the degree of control exerted by that dimension in the future. A related line of research has shown that divided stimulus control is *graded* according to the relative ability of each stimulus to predict future reinforcers. That is, the effects of relative reinforcer predictability on divided stimulus control are not ‘all or none’, in the sense that it is not simply the case that relevant stimuli or those associated with a history of reinforcement exert control and irrelevant stimuli or those not previously associated with reinforcement do not. Instead, the *degree* of stimulus validity or reinforcement history determines the degree of control exerted by a stimulus, relative to other stimuli in the environment.

2.5.4.1 Correlations Between Peck Location and Relative Reinforcer Probabilities

One way to measure the relation between relative reinforcer predictability and divided stimulus control is to analyse the distribution of responses between stimuli correlated with different probabilities of reinforcer delivery. Wasserman (1974) showed that when the elements of a compound stimulus were associated with different reinforcer probabilities, pigeons preferred to peck elements associated with higher probabilities. In his experiment, the compound stimulus was comprised of two colours, white and either red or green, and presentation of the compound was followed by probabilistic response-independent food delivery. Across conditions, Wasserman varied the probability of food deliveries signalled by the red and green colours (the probability of food deliveries signalled by the white stimulus was always .5, because it was always present). To measure divided stimulus control between the compound-stimulus elements, Wasserman analysed key-pecks to each colour. Rates of key-pecking were higher to colours associated with higher probabilities of food, suggesting that colours associated with higher probabilities of food exerted stronger control over behaviour (see also Baetu, Baker, Darredeau, & Murphy, 2005; Farthing, 1971; Hearst, 1988; Jenkins & Sainsbury, 1969, 1970; Sainsbury, 1971; Wagner et al., 1968; Wasserman & Anderson, 1974; B. A. Williams, 1984).

More recent studies support Wasserman’s (1974) finding that pigeons prefer to peck stimuli that are better predictors of reinforcers. Castro and Wasserman (2014, 2016, 2017) arranged DMTS tasks with compound sample stimuli comprised of multiple elements, which predicted the correct comparison with varying probabilities ranging from .5 to 1. These probabilities determined the distribution of pigeons’ key-pecks during sample-stimulus presentations; more pecks were allocated to elements that predicted the correct comparison with higher probabilities than to elements that predicted the correct comparison with lower probabilities. Likewise, in one of the experiments reported in this thesis (Experiment 1, Chapter 3), pigeons preferred to peck stimulus elements that had previously been associated with higher reinforcer rates (see also M. F. Brown, Cook, Lamb, & Riley, 1984; Dittrich, Rose, Buschmann, Bourdonnais, & Güntürkün, 2010; Gomes-Ng, Elliffe, & Cowie, in press; see Awh, Armstrong, & Moore, 2006; Perez, Endemann, Pessôa, & Tomanari, 2015 for similar findings in humans).

However, responses made directly to a stimulus provide only a tentative measure of the strength of control by that stimulus. Indeed, M. F. Brown et al. (1984) found that the dimensions exerting stronger discriminative control could not always be predicted from sample-stimulus peck locations. Specifically, although Brown et al. found a positive correlation between peck location and the elements of a compound stimulus that controlled comparison choice in DMTS, shifts in matching accuracy on each dimension were not always accompanied by shifts in peck location during sample-stimulus presentations. Similarly, in Experiment 2 (Chapter 4) in this thesis, there was no correlation between the location of pigeons' key-pecks to the elements of compound stimuli and the elements that exerted discriminative control over comparison choice in DMTS (see also Furrow & LoLordo, 1975; Perez et al., 2015). Therefore, subjects may not always allocate more responses to the features of the discriminative stimuli that exert stronger control over behaviour, suggesting that other measures should also be used to assess divided stimulus control.

2.5.4.2 Relative Reinforcers and Divided Stimulus Control in DMTS

2.5.4.2.1 Early Evidence for the Effects of Relative Reinforcer Predictability

Studies that have explicitly manipulated the relative ability of each stimulus dimension to predict future reinforcers in DMTS tasks provide stronger evidence that relative reinforcers determine divided control between stimuli that signal the location of reinforcers.

Galloway (1967) presented children with a compound sample stimulus containing a colour and a form element, and then presented those two elements as comparison stimuli and measured preference for one element over the other. When choice of either element was reinforced with the same magnitude (number of marbles), subjects were biased to choose the form element, suggesting that the form dimension exerted stronger control over choice. When the reinforcer magnitude associated with the form dimension was then reduced and the reinforcer magnitude associated with the colour dimension was increased, choice shifted away from the form elements and towards the colour elements, implying that the change in relative reinforcer magnitudes resulted in a decrease in control by the form dimension and an increase in control by the colour dimension (see also Galloway & Petre, 1968).

In a similar experiment, Dube and McIlvane (1997) examined divided control between stimuli that had been previously associated with a high or low rate of reinforcers in an identity DMTS task with individuals with intellectual or developmental disabilities. Like Galloway (1967), Dube and McIlvane found that comparison choice depended on the relative reinforcer rate associated with each stimulus. When two high-rate comparisons were presented, participants were more likely to choose the correct comparison than when two low-rate comparisons were presented. When one high- and one low-rate stimulus were presented as comparisons, participants preferred the high-rate comparison. These findings appear consistent with the conclusion that stimuli associated with higher reinforcer rates exert stronger control over behaviour.

However, some aspects of Galloway's (1967; see also Galloway & Petre, 1968) procedure and of Dube and McIlvane's (1997) results suggest that the apparent shifts in divided stimulus control may not have been related to divided stimulus control at all. In Galloway's experiment, the comparison stimuli were always the two elements that had comprised the previously presented sample stimulus, and hence there was no correct or incorrect comparison. As a result, discriminative control by the sample-stimulus elements was not actually required, and so comparison choice may have been alternatively controlled by the comparison stimuli themselves (see e.g., Nevin, Davison, & Shahan, 2005; Sidman, 1969). That is, shifts in comparison choice may not have been caused by shifts in the division of control between the sample-stimulus elements; instead, participants may have simply preferred one comparison over the other independent of the sample stimulus. Similarly, comparison choice may not have been controlled by the identity of the sample-stimulus elements in Dube and McIlvane's experiment. Dube and McIlvane's participants always preferred the high-rate comparison when presented with a choice between comparison stimuli associated with different reinforcer rates, suggesting that they were simply more *motivated* to choose high-rate comparisons regardless of the sample-stimulus elements (c.f. Lamb, 1991; Wasserman & Miller, 1997). These issues with Galloway's and Dube and McIlvane's experiments highlight the importance of arranging experimental procedures that ensure that the sample stimulus exerts *discriminative* control over behaviour.

2.5.4.2.2 Shahan and Podlesnik (2006): Relative Reinforcement, Divided Stimulus Control, and the GML

To date, the strongest evidence for a relation between relative reinforcer predictability and divided stimulus control comes from a series of studies conducted by Shahan and Podlesnik (2006, 2007), with further replications by Davison and Elliffe (2010) and Davison (2018a). Shahan and Podlesnik (2006)[4] arranged a DMTS task in which the compound

---

[4] In the same year, Della Libera and Chelazzi (2006) published a seminal paper suggesting that reinforcers determine attention allocation in the cognitive-science literature. Their investigation was the catalyst for hundreds of subsequent investigations of the effects of reinforcers on attention allocation, and such 'value-driven attentional capture' has been replicated many times (see e.g., Anderson, 2013, 2016; Anderson, Laurent, & Yantis, 2011, 2013; Chelazzi, Perlato, Santandrea, & Della Libera, 2013; Della Libera & Chelazzi, 2009; Krebs, Boehler, & Woldorff, 2010; Sali, Anderson, & Yantis, 2014; Shomstein & Johnson, 2013). This research shares similarities with the behaviour-analytic literature, in that both demonstrate that stimuli associated with higher relative reinforcer rates are more likely to control behaviour. Della Libera and Chelazzi's paper is mentioned here to make the point that converging evidence from different areas of psychology suggests that reinforcer contingencies play an important role in divided stimulus control (or divided attention).

sample stimuli consisted of a vertical or horizontal line superimposed on a blue or green background, and comparison stimuli were either two lines or two colours. Across conditions, they varied the probability of reinforcer deliveries for correct responses according to the colour dimension (the probability of reinforcers for correct responses according to the line-orientation dimension was the complement). As the probability of reinforcer deliveries for correct colour choices increased, the percentage of correct colour choices also increased, while the percentage of correct line-orientation choices decreased. That is, as the relative reinforcer rate associated with one dimension increased, so did control by that dimension, while control by the other dimension decreased. Furthermore, the relation between relative reinforcer rates and relative matching accuracies was well described by the GML (Baum, 1974), suggesting that the extent of control by a stimulus dimension was proportional to the relative reinforcer rate associated with that dimension.

In order to determine whether their findings were related to shifts in divided stimulus control, or whether their subjects were simply more motivated to choose the correct comparison when the probability of reinforcer deliveries associated with a dimension was higher, Shahan and Podlesnik (2007) conducted a second experiment in which they varied the duration of the sample-stimulus presentation. They reasoned that if relative reinforcer rates affected motivational processes during the comparison phase, then varying the sample-stimulus duration should have little effect on matching accuracy on each dimension. However, if relative reinforcer rates determined divided stimulus control, then increasing the sample-stimulus duration should increase control by each dimension, and to a greater extent for the dimension associated with the higher reinforcer rate. Shahan and Podlesnik's findings were consistent with the latter prediction, thus supporting the conclusion that relative reinforcer rates determine the division of control between stimuli.

Davison and Elliffe (2010) provided further evidence that relative reinforcer rates determine divided stimulus control in a systematic replication of Shahan and Podlesnik's (2006, 2007) experiments. Davison and Elliffe arranged a symbolic DMTS procedure, in which compound sample stimuli were fast or slow colour-alternating stimuli that were presented for either a short or long duration. Pigeons reported either the speed of alternation (Fast or Slow) or the duration (Long or Short) of the sample stimulus by choosing a left comparison key for Fast or Long stimuli and a right key for Slow or Short stimuli. Thus, unlike Shahan and Podlesnik's experiments, the comparison stimuli were identical for both dimensions in Davison and Elliffe's experiment. Hence, any effects of relative reinforcers could not be attributed to motivational processes during the comparison-choice phase. Davison and Elliffe replicated Shahan and Podlesnik's result: As the probability of reinforcers for correct responses according to one dimension increased, so did accuracy on that dimension, while accuracy on the other dimension decreased. In a systematic replication of Davison and Elliffe's experiment, Davison (2018a) demonstrated an important invariance – changing the frequency of the sample stimulus did not systematically affect matching accuracy on each dimension if relative reinforcer rates remained unchanged. Thus, shifts in divided stimulus control appear to be driven by changes in relative reinforcer rates (see Heyman, Grisanzio, & Liang, 2016 for a similar result in humans).

Podlesnik, Thrailkill, and Shahan (2012) extended these findings by assessing the persistence of divided stimulus control following the introduction of a disruptor (pre-session feeding or the removal of all reinforcers during the session). In sessions of disruption, matching accuracy for both stimulus dimensions was higher in a context previously associated with a higher overall reinforcer rate than in a context associated with a lower overall reinforcer rate, suggesting that control by the stimulus dimensions persisted for longer in the context associated with a higher overall reinforcer rate. Additionally, within each context, accuracy was higher on the dimension associated with the higher relative reinforcer rate. Thus, Podlesnik et al.'s results provide further evidence that divided stimulus control depends on relative reinforcer contingencies.

The effects of relative reinforcers on divided stimulus control in the DMTS task have also been studied by manipulating the *reliability* with which the stimuli signal the correct location of future reinforcers, rather than by manipulating the probability of reinforcer deliveries. Rubi and Stephens (2016b) presented blue jays with a compound stimulus comprised of a colour and a pattern, and each dimension signalled, with varying reliability, which of two keys would deliver a reinforcer in each trial. The more reliable dimension – which signalled the correct location of the reinforcer more often – exerted stronger control over choice than the less reliable dimension (see also e.g., Lea et al., 2009; Nicholls, Ryan, Bryant, & Lea, 2011; Rubi & Stephens, 2016a). Similarly, in studies of landmark use, in which subjects locate a goal using several landmarks which signal the goal location, more reliable landmarks (e.g., those that are closer to the goal or that are more stable) generally exert stronger control over search behaviour (e.g., Chen, McNamara, Kelly, & Wolbers, 2017; Cheng, Shettleworth, Huttenlocher, & Rieser, 2007; Du, McMillan, Madan, Spetch, & Mou, 2017; Lechelt & Spetch, 1997; Legge, Madan, Spetch, & Ludvig, 2016; Spetch, 1995; Spetch & Wilkie, 1994).

2.5.4.3 Relative Reinforcers and Divided Stimulus Control in Temporal Discriminations

Thus far, this section has focused on research investigating the division of control between stimuli that provide information about the location of future reinforcers. There is limited evidence that relative reinforcer rates also determine the division of control between stimuli that provide information about the *time* of future reinforcers (Davison, 2018b). Matell and Kurti (2014) arranged a multiple peak procedure (Catania, 1970; S. Roberts, 1981) in which tone and light stimuli signalled a short or long interval to response-contingent reinforcer delivery, and each stimulus was also associated with a different probability of reinforcer delivery. The tone and light stimuli were occasionally presented alone or simultaneously in longer, unreinforced ‘peak’ trials. In peak trials with one stimulus, response rates increased gradually to the usual time of reinforcer delivery signalled by the stimulus, and then decreased gradually thereafter. When both stimuli were presented simultaneously, patterns of responding across time were more similar to response-rate patterns in peak trials in which only the stimulus associated with the higher reinforcer probability was presented. Thus, the stimulus previously associated with the higher reinforcer probability appeared to exert stronger control over interval-timing behaviour (see also Experiment 3, Chapter 5, of this thesis).

In a similar experiment to Matell and Kurti (2014), Delamater and Nicolas (2015) showed that when the tone and light stimuli were associated with different types of reinforcers, pre-session satiation with one reinforcer shifted response-rate functions in compound peak trials away from the interval signalled by the stimulus associated with that reinforcer and towards the interval signalled by the other stimulus. That is, when the reinforcer associated with one stimulus was devalued (Skinner, 1932) by pre-session satiation, that stimulus appeared to exert weaker control over behaviour. Thus, the relative reinforcer value associated with the tone and light stimuli determined patterns of responding in the presence of both stimuli.

Research investigating divided stimulus control between temporal and non-temporal stimuli also suggests that relative reinforcers determine such division. For example, as described previously, Davison and Elliffe (2010) found that the division of control between the duration of sample-stimulus presentation (Long or Short, a temporal stimulus) and the speed of keylight-colour alternation (Fast or Slow, a non-temporal stimulus) depended on relative reinforcer rates. Likewise, Aum, Brown, and Hemmes (2007) found that temporal discrimination in the peak procedure was reduced by the insertion of an S+ stimulus previously associated with reinforcers – suggesting a shift in stimulus control from elapsed time to the S+ stimulus – whereas the insertion of an S- stimulus did not disrupt temporal discrimination (see also McMillan & Roberts, 2013). Thus, when elapsed time since a marker event (e.g., trial start) correlates with a change in reinforcer availability, divided control between elapsed time and other informative non-temporal stimuli depends on the relative ability of each stimulus to predict future reinforcers.

Furthermore, Rayburn-Reeves, Qadri, Brooks, Keller, and Cook (2017) found that the division of control between elapsed time and non-temporal stimuli changed across time, in accordance with time-based changes in the relative reliability of each stimulus. Rayburn-Reeves et al. arranged a mid-session reversal (MSR) task, in which the location of reinforcers changed abruptly after a fixed number of trials since session start and additional keylight-colour stimuli signalled the location of the next reinforcer. Although the change in reinforcer location was not contingent on elapsed time, it was correlated (albeit imperfectly) with time since session start (see Rayburn-Reeves & Cook, 2016 for a review), and hence choice was jointly controlled by elapsed time and by the keylight-colour stimuli. Such divided stimulus control changed across time: At early and late times in each session, elapsed time exerted stronger control and keylight colour exerted weaker control, whereas the reverse was true towards the midpoint of the session. Rayburn-Reeves et al. suggested that this time-based change in divided stimulus control occurred because elapsed time signalled the next-reinforcer location with more certainty at early and late times (because at those times, the next reinforcer would definitely occur at either one or the other location), whereas elapsed signalled the next-reinforcer location with less certainty towards the midpoint of the session because of the change in reinforcer location (see also Cowie, Davison, Blumhardt, & Elliffe, 2016c; Cowie, Davison, & Elliffe, 2014; McMillan, Sturdy, Pisklak, & Spetch, 2016; Rayburn-Reeves, Molet, & Zentall, 2011; Smith, Beckmann, & Zentall, 2017).

However, in contrast to Rayburn-Reeves et al. (2017), both Cowie et al. (2017) and Experiment 4 (Chapter 6) of this thesis found no systematic change in the division of control between elapsed time and informative keylight-colour stimuli across time, despite systematic changes in the reliability with which elapsed time signalled the location of future reinforcers (see Chapter 6 and Gomes-Ng et al., 2018a for further discussion). Instead, keylight-colour stimuli exerted strong control over behaviour regardless of the reliability of elapsed time. Thus, at present, the effects of relative reinforcer predictability on divided control between elapsed time and informative non-temporal stimuli appear to be somewhat mixed.

2.5.5 Relative Reinforcement and Contingency Discriminability

Relative reinforcer predictability can only control behaviour to the extent that it can be accurately discriminated by the subject (Davison & Nevin, 1999). Thus, a stimulus that differentially signals reinforcer availability will only exert strong control over behaviour if subjects can discriminate such differential reinforcers, and likewise a stimulus that nondifferentially signals reinforcer availability will only exert weak control if subjects can discriminate such nondifferential reinforcers (Davison, 2018a; Davison & Elliffe, 2010). Reliable stimuli will exert less behavioural control or unreliable stimuli will exert greater control to the extent that differential and nondifferential reinforcers, respectively, are imperfectly discriminated (see Chapters 4 and 6 of this thesis for further discussion).

Thus, although the general consensus is that reliable stimuli that provide non-redundant information exert stronger control over behaviour – as the findings described in this chapter illustrate – there are cases where relevant stimuli fail to gain strong control over behaviour, and where redundant or irrelevant stimuli additionally or alternatively exert strong control (e.g., Aust & Huber, 2002; Born & Snow, 1970; Davison & Cowie, 2019; Emmerton, 1998; Gomes-Ng et al., in press; House & Zeaman, 1963; Nicholls et al., 2011; Pinto, Fortes, & Machado, 2017; Pinto & Machado, 2017; Reinhold & Perkins, 1955; Thomas, Burr, & Eck, 1970; D. I. Williams, 1967; see also Experiment 2, Chapter 4, of this thesis). Indeed, changes in the reliability of elapsed time may have had little to no effect on divided stimulus control in Cowie et al.'s (2017) and Gomes-Ng et al.'s (2018a) experiments because subjects may not have discriminated such changes (see Chapter 6 for further discussion). Such cases demonstrate the importance of considering contingency discriminability from the subject's perspective; the experimenter-arranged relevant and irrelevant stimuli are not necessarily the same as the subject-discriminated relevant and irrelevant stimuli (Davison & Nevin, 1999; Ray, 1972).

2.6 Chapter Summary

The introductory paragraphs of this chapter emphasised that divided stimulus control is a matter of "this, that, or these?" and of "how much?". That is, which stimulus or stimuli control behaviour, and how strong is the control? This chapter provided an overview of research that has examined the variables that determine which stimuli control behaviour, and the strength of control of those stimuli. These variables include the sensory capacities and phylogenetic history of the subject; characteristics of the stimuli such as intensity, physical disparity, and spatial and temporal configuration; and the ability of each stimulus to predict future reinforcer availability. To summarise, some stimuli are more likely to control behaviour because they come from a subject's dominant sensory modality, they are more ecologically valid to the subject, they are easier to discriminate due to physical characteristics such as intensity or disparity, or they provide more reliable information about the consequences that are likely to follow a behaviour.

Focusing on the latter finding, the research described in the preceding section overwhelmingly suggests that divided stimulus control depends on the relative reinforcer rate associated with, or the reliability of, each stimulus. Is there any evidence of the contrary? Presently, little research has examined the generality of the relation between relative reinforcers and divided stimulus control, and hence the conditions under which relative reinforcers do and do not determine divided stimulus control are poorly understood (see Davison, 2018b for a brief discussion). Nevertheless, as described earlier, irrelevant, redundant, or less reliable stimuli sometimes exert strong control over behaviour, perhaps because such stimuli are, or appear to be, additionally correlated with reinforcers. Also, to pre-empt the findings of the four experiments reported in this thesis (Chapters 3 to 6), each experiment showed that relative reinforcers determine divided stimulus control in a novel procedure, but also suggested some potential limits to the conditions under which relative reinforcers determined divided stimulus control.

2.7 The Present Thesis

The aim of the present thesis was to investigate further how relative reinforcers determine the division of control between the dimensions of compound stimuli. Presently, only a few studies provide strong evidence for a systematic relation between relative reinforcer rates and divided stimulus control (Davison, 2018a; Davison & Elliffe, 2010; Heyman et al., 2016; Podlesnik et al., 2012; Shahan & Podlesnik, 2006, 2007; see also Rubi & Stephens, 2016a, 2016b). These studies arrange the same general procedure, in which a compound stimulus is presented as a sample in a DMTS task (see Figure 2.1), and the probability of reinforcer deliveries for correct responses according to each sample-stimulus dimension is varied across conditions. Therefore, beyond this procedure, the generality of the relation between relative reinforcers and divided stimulus control is not well established

(Davison, 2018b). As Davison (2018b) puts it, “We know … not nearly enough about how reinforcer frequency divides stimulus control” (p. 136). Each experiment in this thesis thus investigated the effects of reinforcer contingencies on divided stimulus control in a novel procedure.

2.7.1 Here or There? Relative Reinforcers and Divided Stimulus Control in Space

The first experiment[5] (Chapter 3) addressed a gap in the basic behaviour-analytic literature regarding the effects of previously trained stimulus-response-reinforcer relations on divided stimulus control. At present, it is unclear whether previously trained stimulus-response-reinforcer relations divide stimulus control in the same way as stimulus-response-reinforcer relations trained in compound do. Experiment 1 therefore asked whether relative reinforcer rates determine the division of control between the dimensions of a compound stimulus comprised of elements previously associated with high or low reinforcer rates. In addition, the configuration of the compound stimulus varied across conditions, in order to investigate the combined effects of spatial configuration and previously trained stimulus-response-reinforcer relations on divided stimulus control.

The results of Experiment 1 suggested that the spatial configuration and/or the type of training procedure may affect whether or not a relation between relative reinforcer rates and divided stimulus control is observed. Experiment 2 (Chapter 4) investigated how these two procedural factors may have contributed to the results of Experiment 1. Experiment 2 was a replication of Davison and Elliffe (2010) with spatially separated stimulus dimensions. Thus, in addition to investigating the potential explanations for the results of Experiment 1, Experiment 2 also investigated whether relative reinforcer rates determine divided stimulus

---

[5] Portions of this chapter are included in the publication Gomes-Ng, S., Elliffe, D., & Cowie, S. (in press). Relative reinforcer rates determine pigeons’ attention allocation when separately trained stimuli are presented together. *Learning & Behavior.*

control with a different spatial configuration from the typical unified or superimposed configuration, and whether a quantitative model of divided stimulus control proposed by Davison and Elliffe could describe any effects of relative reinforcer rates on divided stimulus control between spatially separated stimuli.

2.7.2 Now or Then? Relative Reinforcers and Divided Stimulus Control in Time

Experiments 1 and 2 examined the effects of reinforcer contingencies on divided stimulus control in space. Experiment 3 (Chapter 5) asked whether relative reinforcer probabilities determine divided stimulus control between the dimensions of compound stimuli when those dimensions provided information about the time of future reinforcers. Experiment 3 was a systematic replication of Davison and Elliffe (2010), but each dimension signalled a 2-s or 8-s interval to response-contingent reinforcer delivery, rather than to choose a left or right key. Across conditions, the relative reinforcer probability associated with each stimulus dimension varied. This experiment is the first to extend previous research examining the relation between relative reinforcer rates and divided stimulus control to temporal discriminations, and also provides insight into the mechanisms that may underlie temporal discrimination when multiple stimuli signal discrepant times to reinforcer delivery.

Finally, the fourth experiment[6] (Chapter 6) investigated further the nature of divided stimulus control in time. Experiment 4 took a slightly different approach than Experiments 1 to 3. Experiment 4 arranged a concurrent schedule in which the reinforcer ratio changed across time since a reinforcer delivery, and hence elapsed time since a reinforcer delivery signalled the likely location of the next reinforcer. In Experiment 4, one alternative was always more likely to deliver reinforcers than the other alternative, and across conditions the

---

[6] This chapter is an edited version of the publication Gomes-Ng, S., Elliffe, D., & Cowie, S. (2018a). Environment tracking and signal following in a reinforcer-ratio reversal procedure. *Behavioural Processes, 157,* 208-224.

probability of reinforcer deliveries on this locally richer alternative varied. This manipulation effectively changed the *reliability* with which elapsed time signalled the likely location of the next reinforcer – the higher the probability of reinforcer deliveries on the locally richer alternative, the more reliably elapsed time signalled the next-reinforcer location. In some conditions, an additional keylight-colour stimulus signalled the likely or definite location of the next reinforcer, and the reliability of the keylight-colour stimuli depended on the information provided by the stimuli. Therefore, Experiment 4 asked how relative reliabilities divide stimulus control between temporal (elapsed time) and non-temporal (keylight colour) stimuli. This extends previous research that has asked the same question (e.g., Rayburn-Reeves et al., 2017; see also Cowie et al., 2017), but unlike those studies, we explicitly manipulated the reliability of elapsed time, and changes in the reliability of elapsed time were not confounded with changes in elapsed time, in Experiment 4.

Taken together, the experiments in this thesis provide new insights into the conditions under which relative reinforcer predictability does and does not determine the division of control between multiple stimuli. These insights, as well as their applied relevance and the future research directions that they generate, are discussed in Chapter 7.

CHAPTER 3

Experiment 1: Relative Reinforcer Rates and Divided Control by Separately Trained Stimuli[7]

3.1 Introduction

The division of control between the dimensions of a compound stimulus depends on the relative reinforcer rate associated with each dimension (Davison, 2018a; Davison & Elliffe, 2010; Podlesnik, Thrailkill, & Shahan, 2010; Shahan & Podlesnik, 2006, 2007). The generality of this relation between relative reinforcer rates and divided stimulus control is not well established, because the studies demonstrating such a relation employ the same general procedure – a delayed matching-to-sample (DMTS) task in which the sample stimuli are compounds comprised of two elements from different stimulus dimensions, and control by each dimension is assessed by asking subjects (typically, pigeons) to report the identity of one of the sample-stimulus elements during the comparison phase. One situation in which the effects of relative reinforcer rates on divided stimulus control have not yet been thoroughly examined is when stimuli are first trained separately and then encountered together. That is, when stimuli are trained individually, will the relative reinforcer rate associated with each stimulus during training determine divided stimulus control when the stimuli are later presented simultaneously?

Investigating how stimulus control is divided between separately trained stimuli, and how relative reinforcer rates might affect such division, is particularly important from an

---

[7] This experiment was conducted under Approval 001396 granted by the University of Auckland Animal Ethics Committee.
Portions of Experiment 1B in this chapter are included in the publication Gomes-Ng, S., Elliffe, D., & Cowie, S. (in press). Relative reinforcer rates determine pigeons' attention allocation when separately trained stimuli are presented together. *Learning & Behavior.* Permission to include Experiment 1B in this thesis was obtained (see Appendix B3 for license).

applied perspective. In applied settings, selective stimulus control (termed *stimulus overselectivity*) poses a challenge to the success of behavioural interventions, especially in individuals diagnosed with intellectual or developmental disabilities. In addition to reducing the success of behavioural interventions, stimulus overselectivity can hamper all aspects of an individual's social, cognitive, and linguistic development, as many of these skills rely on control by multiple relevant cues (see Ploog, 2010 for a review). Some research suggests that stimulus overselectivity can be remediated by separately training stimulus-reinforcer associations – for example, reducing the reinforcer rate associated with the overselected stimulus in separate training trials has been shown to reduce control by the overselected stimulus and increase control by underselected stimuli (e.g., Broomfield, McHugh, & Reed, 2008, 2010; Kelly, Leader, & Reed, 2015; Reed, Broomfield, McHugh, McCausland, & Leader, 2009; Reed, Reynolds, & Fermandel, 2012). If relative reinforcer rates determine the division of control between separately trained stimuli, this would provide additional support for such an intervention, and would also suggest that pre-training with individual stimuli before presenting them together might help to reduce the likelihood of overselectivity occurring (see also e.g., Dube & McIlvane, 1997; Newman & Benefield, 1968; von Fersen & Lea, 1990).

More broadly, investigating the variables that determine divided control between separately trained stimuli is important because organisms in the natural world may sometimes encounter informative cues separately, and at other times those same cues may be encountered together. This is especially likely given that natural environments are ever-changing. Blaisdell, Schroeder, and Fast (2018) provide an example of such a situation: Sometimes, all of the landmarks that a human or non-human animal uses to navigate to a location will be within their field of vision, whereas at other times, only a subset of those landmarks will be visible (e.g., due to the individual's position, the terrain, or the weather).

Another example, relevant to social interactions in humans, is that individuals may sometimes interact with another individual one-on-one, and at other times they will interact with multiple individuals in a group. Examining the role of relative reinforcer rates in determining divided control between separately trained stimuli will help to elucidate the processes that underlie behaviour when separate stimuli (e.g., landmarks, people) are sometimes encountered in isolation, and at other times are encountered together.

Some research suggests that relative reinforcer rates might not determine divided control by separately trained stimuli. Relative reinforcer rates can only determine divided stimulus control if differences in reinforcer rates between stimuli are discriminated by the subject (see e.g., Davison & Nevin, 1999; Lobb & Davison, 1975). Such discrimination is probably harder when stimuli are trained separately, compared with when they are trained together. In support of this, sensitivity to relative reinforcer rates is lower in multiple schedules, in which stimuli are presented separately, than in concurrent schedules, in which stimuli are presented simultaneously (Davison & Ferguson, 1978; Lander & Irwin, 1968; Lobb & Davison, 1977; McLean & White, 1983; McSweeney, Farmer, Dougan, & Whipple, 1986; Reynolds, 1963). Similarly, compared with concurrent schedules, sensitivity to reinforcement is lower in concurrent-chains schedules, in which the contingencies associated with each discriminative stimulus are experienced individually in separate ‘terminal links’ (Davison, 1976; Fantino & Davison, 1983). Thus, it appears that subjects are better able to discriminate relative reinforcer rates when stimuli are encountered simultaneously rather than separately. This suggests that sensitivity to the reinforcer rate associated with one stimulus *relative to* other stimuli may be lower when stimuli are trained separately than when they are trained together in compound. If so, then relative reinforcer rates will have smaller effects, or perhaps even no effect at all, on divided stimulus control when stimuli are trained separately, compared with when they are trained in compound.

Nevertheless, some research tentatively suggests that relative reinforcer rates determine divided control between separately trained stimuli, in the same way as they do when the stimuli are trained in compound. Stimuli trained separately, but in the same context, can compete with each other for control over behaviour (e.g., Escobar, Matute, & Miller, 2001; Matute & Pineño, 1998), suggesting that control by separately trained stimuli may depend on each stimulus' ability to predict uniquely an outcome in a given context. Kalafut and Church (2017) found that the probability of reinforcer deliveries previously associated with tone and light stimuli predicted overall response rates in the presence of both stimuli. Similarly, as described in Chapter 2, Matell and Kurti (2014) showed that in the presence of a compound light-tone stimulus, rats' patterns of responding across time were modulated by previously trained light-reinforcer and tone-reinforcer relations; the stimulus associated with the higher probability appeared to exert stronger control over responding (see also Delamater & Nicolas, 2015). In a DMTS task, Dube and McIlvane (1997) found that comparison choice depended on the reinforcer rate previously associated with each stimulus; participants preferred comparisons previously associated with higher reinforcer rates.

However, some aspects of Kalafut and Church's (2017), Matell and Kurti's (2014) and Dube and McIlvane's (1997) studies make it difficult to determine the nature of the relation between previously trained stimulus-response-reinforcer relations and divided stimulus control. Both Kalafut and Church and Matell and Kurti arranged single-alternative schedules, and so it was impossible to determine precisely which stimulus controlled behaviour more strongly because the same response was associated with both stimuli (see Bushnell & Weiss, 1977). Additionally, Matell and Kurti found that stimulus modality and the interval to reinforcer delivery also affected patterns of responding, and so the extent to which relative reinforcer rates versus these other variables contributed to response patterns is unclear. Although these were not issues for Dube and McIlvane's study, some of Dube and

McIlvane's findings are better explained in terms of the motivation to respond rather than shifts in divided stimulus control, as their participants preferred the comparison associated with the higher reinforcer rate regardless of whether or not it correctly matched the sample stimulus.

Therefore, in order to investigate thoroughly whether relative reinforcer rates determine divided control between separately trained stimuli, the stimuli should be associated with different responses (Bushnell & Weiss, 1977), and any potential effects of other variables such as stimulus modality or interval to reinforcer delivery on stimulus control, or of relative reinforcer rates on processes during the comparison phase (e.g., motivation), should be eliminated as best as is possible. The present experiment addressed these issues. This experiment asked whether relative reinforcer rates determine divided stimulus control when separately trained visual stimuli, each associated with a left- or right-key response, are presented simultaneously. Specifically, pigeons learned to match two colours (red and green) and two frequencies at which a white keylight alternated on and off (fast and slow) with either a left or right comparison key in a symbolic DMTS procedure. The relative reinforcer rate associated with the colour and flash-frequency dimensions varied across conditions during training. After training, test trials were introduced, in which a colour and a flash-frequency stimulus were presented simultaneously. In test trials, the stimulus elements were either compounded together (Experiment 1A) or were presented on separate stimulus displays (Experiment 1B). Behaviour in such test trials provided an indication of whether the previously learned stimulus-response-reinforcer relations during training influenced the division of control between the colour and flash-frequency dimensions.

3.2 Method

3.2.1 Subjects

Six experimentally naïve pigeons, numbered 71 to 76, served as subjects. The pigeons were maintained at 85% ± 10 g of free-feeding body weight by supplementary feeding of mixed grain at around 10:00 a.m. daily. Water and grit were freely available. The pigeons were housed in a colony room with a time-shifted environment (lights on from 12:00 a.m. until 4:00 p.m.).

3.2.2 Apparatus

The pigeons were housed individually in home cages (375 mm high x 375 mm deep x 370 mm wide), which also served as the experimental chambers. Each cage contained two wooden perches, spaced 135 mm apart and parallel to the back wall. The back perch was 110 mm away from the back wall. An operant panel was mounted on the back wall, 115 mm above the cage floor. The operant panel contained five circular response keys, each 20 mm in diameter, arranged in three rows. The top and bottom rows each contained two keys spaced 113 mm apart from centre to centre. The middle row contained one response key, which was centred horizontally between the keys in the top and bottom rows. The perpendicular distance from the bottom of the keys in one row and the top of the keys in the row below was 25 mm. Only the top and bottom keys were used; these could be illuminated red, green, or white. Responses exceeding about 0.1 N to illuminated keys were recorded. Centred on the back wall and 53 mm below the operant panel was a magazine aperture measuring 40 mm by 40 mm. A hopper, filled with wheat, was located behind the magazine aperture. During a reinforcer delivery, the hopper was raised for 2 s, the magazine aperture was illuminated, and the keys were darkened. In an adjacent room, a computer running MED-PC® IV software ran the experiment and recorded all experimental events.

3.2.3 Procedure

3.2.3.1 Pretraining

The pigeons were slowly deprived of food until they reached 85% of their free-feeding body weight, and then trained to eat from the food hopper. Once the pigeons were reliably eating during 2-s hopper presentations, they were autoshaped (P. L. Brown & Jenkins, 1968) to peck the top and bottom response keys when the keys were illuminated red, green, or white. Once the pigeons were pecking all keylight colours and response keys, the schedule of reinforcer deliveries was changed from reinforcing every response to a variable-interval (VI) 1-s schedule, and the schedule was gradually thinned to VI 30 s. The experiment proper began once all pigeons were reliably pecking illuminated keys on a VI 30-s schedule.

3.2.3.2 Experimental Sessions

Experimental sessions were run daily, beginning at 1 a.m. The pigeons were run successively and in numerical order. Each session lasted until 136 trials or 65 min had elapsed, whichever occurred first. All pigeons completed all 136 trials in every session. The same procedure was used in both Experiments 1A and 1B.

The pigeons were trained to associate four sample stimuli with either a left or right comparison key using a symbolic 0-s DMTS task. The sample stimuli were a red keylight, a green keylight, and a white keylight that alternated on and off every 0.1 s or every 0.5 s (hereafter, for brevity, these stimuli are termed ‘red’, ‘green’, ‘fast’, and ‘slow’ respectively). There were thus two stimulus dimensions – colour and flash-frequency.

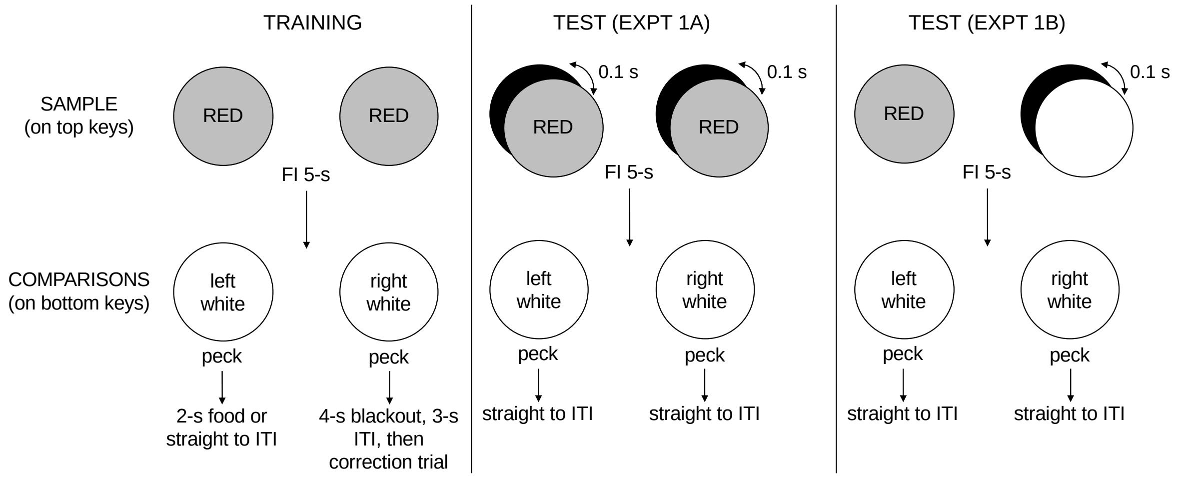

Figure 3.1. Diagram showing examples of a red training trial (left panel) and red-fast test trials in Experiments 1A (middle panel) and 1B (right panel). In this example, the red stimulus matches the left comparison. In red-fast test trials in Experiment 1A, the stimulus elements were compounded together so that the red sample keys alternated on and off every 0.1 s. In red-fast test trials in Experiment 1B, the red and fast stimuli were presented on separate sample keys and their locations were counterbalanced across trials. See text for further details.

The left panel of Figure 3.1 depicts an example of a training trial. Each trial began with the illumination of the sample stimulus, which was displayed on the top-left and top-right keys (hereafter, collectively termed the *sample keys*). The sample stimulus was chosen probabilistically ($p = .25$), with the constraints that there were an equal number of trials per sample in each session and the same stimulus could not occur more than twice consecutively. The sample stimulus was presented for at least 5 s, after which a response to either sample key resulted in the offset of the sample stimulus and the onset of the comparison keys. This fixed-interval (FI) 5-s requirement ensured that the pigeons observed the sample in each trial. The comparison keys were the bottom-left and bottom-right keys, which were illuminated white. Depending on the previously presented sample, one of the comparisons was defined as correct, and the other as incorrect. If the program arranged a reinforcer, a response to the correct comparison was immediately reinforced. If no reinforcer was arranged, a correct response turned off the comparison keys and started the inter-trial interval (ITI), which was a

3-s blackout period. An incorrect response resulted in a 4-s blackout followed by the ITI, after which the same trial repeated except that only the correct comparison was illuminated during the comparison phase (i.e., a correction procedure). Correction trials did not count towards the 136 total trials per session.

The sample-comparison associations were counterbalanced across pigeons. For Pigeons 71 and 72, the red and fast stimuli matched the left key; for Pigeons 73 and 74, the green and fast stimuli matched the left key; for Pigeon 75, the green and slow stimuli matched the left key; and for Pigeon 76, the red and slow stimuli matched the left key. The other stimuli matched the right key.

The relative probability of a reinforcer delivery associated with each stimulus dimension varied across conditions. The probabilities of reinforcer deliveries for each dimension were always complementary, in order to keep the overall reinforcer rate the same in all conditions (arranging such complementary probabilities also replicates previous work examining the relation between relative reinforcer rates and divided stimulus control). Table 3.1 shows the sequence of conditions, which was the same for all pigeons. In Conditions 1 and 4, the probability of a reinforcer delivery was .2 for colour stimuli and .8 for flash-frequency stimuli. Conditions 2, 5, and 7 arranged the reverse probabilities. In Conditions 3 and 6, the probability of reinforcer deliveries was .5 for both stimulus dimensions.

Training sessions in each condition ran for at least 31 sessions and until accuracy for each sample stimulus was above 80% for five consecutive sessions. All pigeons generally met this accuracy criterion within 31 sessions. The exact number of training sessions for each pigeon in each condition is shown in Appendix Table A3.1. Due to experimenter error, training sessions ended after 31 sessions despite accuracy for the green stimulus being 76%

in Session 29 and 79% in Session 31 for Pigeon 73 in Condition 1. However, this had no apparent effect on that pigeon's data.

Table 3.1.

Sequence of conditions in Experiments 1A and 1B.

Experiment	Condition	*p*(food) for colour	Stimulus type in test trials
1A	1	.2	Compound
1A	2	.8	Compound
1A	3	.5	Compound
1B	4	.2	Separated
1B	5	.8	Separated
1B	6	.5	Separated
1B	7 (replication)	.8	Separated

Note: Compound means the colour and flash-frequency dimensions were compounded together. Separated means the stimulus dimensions were presented separately on different sample keys.

After training, four test sessions were conducted. Test sessions were identical to training sessions, except that 16 unreinforced training trials (taken equally from the four sample stimuli) were replaced with unreinforced test trials. Test trials were distributed randomly throughout the session, and never occurred consecutively. In test trials, a colour and a flash-frequency element were presented simultaneously. There were thus four types of test trials: red-fast, red-slow, green-fast, and green-slow. Four trials of each type occurred in each session. For two of the test-trial types, both elements matched the same comparison key (*compatible* trials), whereas for the other two trial types, the elements matched opposite keys (*incompatible* trials).

In Experiment 1A (Conditions 1 to 3), the stimulus elements were compounded together into a single stimulus in test trials. That is, the sample stimulus was a red or green keylight that alternated on and off every 0.1 s or 0.5 s. This compound sample stimulus was presented on both sample keys. The middle panel of Figure 3.1 depicts an example of a red-fast test trial in Experiment 1A.

In Experiment 1B (Conditions 4 to 7), the stimulus elements were presented separately on different sample keys in test trials. Thus, a red or green stimulus was presented on one key, and a fast or slow stimulus on the other key. The location of the colour and flash-frequency elements was counterbalanced so that the colour was presented on the left sample key in half of test trials, and on the right sample key in the other half. The right panel of Figure 3.1 depicts an example of a red-fast test trial in Experiment 1B.

3.3 Results

3.3.1 Obtained Reinforcer Rates

In all conditions and for all pigeons, the obtained proportion of reinforced trials for each sample stimulus was always within .02 of its arranged value (see Table 3.1 for arranged reinforcer probabilities). Thus, obtained and arranged reinforcer rates were equal.

3.3.2 Matching Accuracy in Training and Compatible Test Trials

We calculated the proportion of correct responses following each sample stimulus (*matching accuracy*) in training trials across the last five training sessions, and in training trials and compatible test trials across all four test sessions. Correction trials were excluded from analyses. Accuracy in training trials was similar during the last five training sessions and during test sessions; binomial sign tests on accuracy in training trials during training versus test sessions were not significant for all conditions (all $p > .3$). Hence, introducing test trials did not systematically affect training-trial accuracy. Therefore, only training-trial data from test sessions is presented here, and hereafter 'training trials' always refers to training trials during test sessions.

Table 3.2 shows matching accuracy in training trials for each condition. The data shown in Table 3.2 are averaged across pigeons because accuracy was similarly high for all pigeons. In all conditions, matching accuracy in training trials was close to 100% for all sample stimuli. Relative reinforcer rates had no systematic effect on accuracy in training trials; two-tailed nonparametric trend tests (Elliffe & Elliffe, 2019; Kendall, 1955) on matching accuracies averaged across stimuli were not significant (ΣS = -4 and 2 for Experiments 1A and 1B respectively, both, $p > .4$, $N = 6$, $k = 3$).

In compatible test trials, both stimulus elements matched the same comparison. The rightmost two columns of Table 3.2 show the proportion of correct comparison responses in compatible test trials. Nonparametric Friedman analyses of variance by ranks comparing accuracy in each type of compatible test trial (i.e., red-compatible, green-compatible) with accuracy in training trials with the stimuli that comprised the compatible compound were all not significant ($N = 6$, $df = 4$, $p > .05$), although the comparison between green-compatible trials and training trials approached significance in Condition 5 ($\chi^2 = 8.63$, $p \approx .07$). Therefore, overall, matching accuracy in compatible test trials was similar to accuracy in training trials. Like training trials, relative reinforcer rates had little effect on accuracy in compatible test trials (two-tailed nonparametric trend tests: $\Sigma S = 1$ for both red- and green-compatible trials in Experiment 1A; ΣS = -1 and 2 for red- and green-compatible trials in Experiment 1B, all $p > .8$, $N = 6$, $k = 3$).

Table 3.2.

Proportion of correct responses (matching accuracy), averaged across pigeons, for each sample stimulus in training trials and compatible test trials. Standard deviations are shown in parentheses.

Condition	Red	Green	Fast	Slow	Red-Compatible	Green-Compatible
1	.97 (.04)	.97 (.04)	.97 (.03)	.99 (.01)	.91 (.17)	.96 (.05)
2	.97 (.05)	.97 (.05)	.90 (.12)	.98 (.02)	.97 (.05)	.96 (.08)
3	.99 (.01)	.98 (.03)	.97 (.05)	.99 (.01)	.99 (.03)	.96 (.07)
4	.99 (.03)	.97 (.04)	.97 (.02)	.99 (.01)	.94 (.07)	.97 (.06)
5	.99 (.02)	.99 (.03)	.94 (.10)	.98 (.02)	.96 (.06)	.95 (.05)
6	.99 (.01)	.99 (.01)	.98 (.02)	.99 (.01)	.98 (.05)	.97 (.05)
7	.99 (.01)	.99 (.01)	.98 (.02)	.98 (.02)	.90 (.16)	.98 (.05)

Note: Red-compatible and green-compatible refer to compatible test trials in which the colour element was red or green, respectively.

Compatible test trials served as control trials, as they provided an indication of whether the pigeons treated the simultaneous presentation of a colour and flash frequency as novel, and whether the pigeons learned over time that test trials were unreinforced. The high matching accuracies in compatible test trials in all conditions (Table 3.2) suggest that neither of these occurred. If the pigeons treated the compound stimulus as novel in test trials, then accuracy should have been closer to chance (i.e., proportion correct = .5), and if the pigeons learned that test trials were always unreinforced, accuracy in compatible test trials should have been systematically lower than in training trials.

Now that we have established that matching accuracy was high in training and compatible test trials, we can evaluate divided stimulus control between the colour and flash-frequency dimensions in test trials. The remainder of this results section presents several analyses of test-trial data aimed at determining the division of control between the stimulus dimensions in test trials.

3.3.3 Experiment 1A

To assess divided stimulus control between the colour and flash-frequency dimensions in test trials, we calculated a bias-free measure of preference (Davison & Tustin, 1978; see also Davison & Nevin, 1999) for the comparison key matching the colour or flash-frequency element (log $d_{matching}$):

$$\log d_{matching} = 0.5 \log\left(\frac{B_{colour\,|\,S1}}{B_{flash-freq\,|\,S1}} \cdot \frac{B_{colour\,|\,S2}}{B_{flash-freq\,|\,S2}}\right). \qquad \text{(Equation 3.1)}$$

In Equation 3.1, $S1$ and $S2$ represent two different trial types (e.g., red- and green-incompatible), and $B_{X=colour\ or\ flash-freq\,|\,Y=S1\ or\ S2}$ represents the number of responses made to the comparison key matching element X in trial type Y. To correct for $B_{X|\,Y}$ counts that equalled zero (i.e., all responses were made to one key), we added 0.5 to all response counts (Hautus, 1995; see also Davison & Nevin, 1999). Positive values of Equation 3.1 indicate stronger preference for the comparison key matching the colour element and hence greater control by that element. Negative values indicate stronger preference for the comparison key matching the flash-frequency element. Thus, if relative reinforcer rates determine divided stimulus control between separately trained stimuli, values of log $d_{matching}$ should increase as the reinforcer probability associated with colour stimuli increases. It is important to note that log $d_{matching}$ only provides a meaningful indication of divided stimulus control in incompatible test trials, because the colour and flash-frequency elements matched different comparison keys in only those test trials. Therefore, only log $d_{matching}$ values for incompatible test trials were calculated.

The filled data points in Figure 3.2 show log $d_{matching}$ (Equation 3.1) values, plotted as a function of the probability of reinforcer deliveries associated with colour stimuli, in incompatible test trials in Experiment 1A. log $d_{matching}$ increased as the reinforcer probability associated with colour stimuli increased for two pigeons (Pigeons 71 and 76), whereas it varied unsystematically for the remaining four pigeons. One-tailed nonparametric trend tests on log $d_{matching}$ were not significant for Experiment 1A ($\Sigma S = 6$, $p > .1$, $N = 6$, $k = 3$). Thus, relative reinforcer rates did not systematically affect the degree of discriminative control exerted by each element of a compound stimulus when those elements were trained separately.

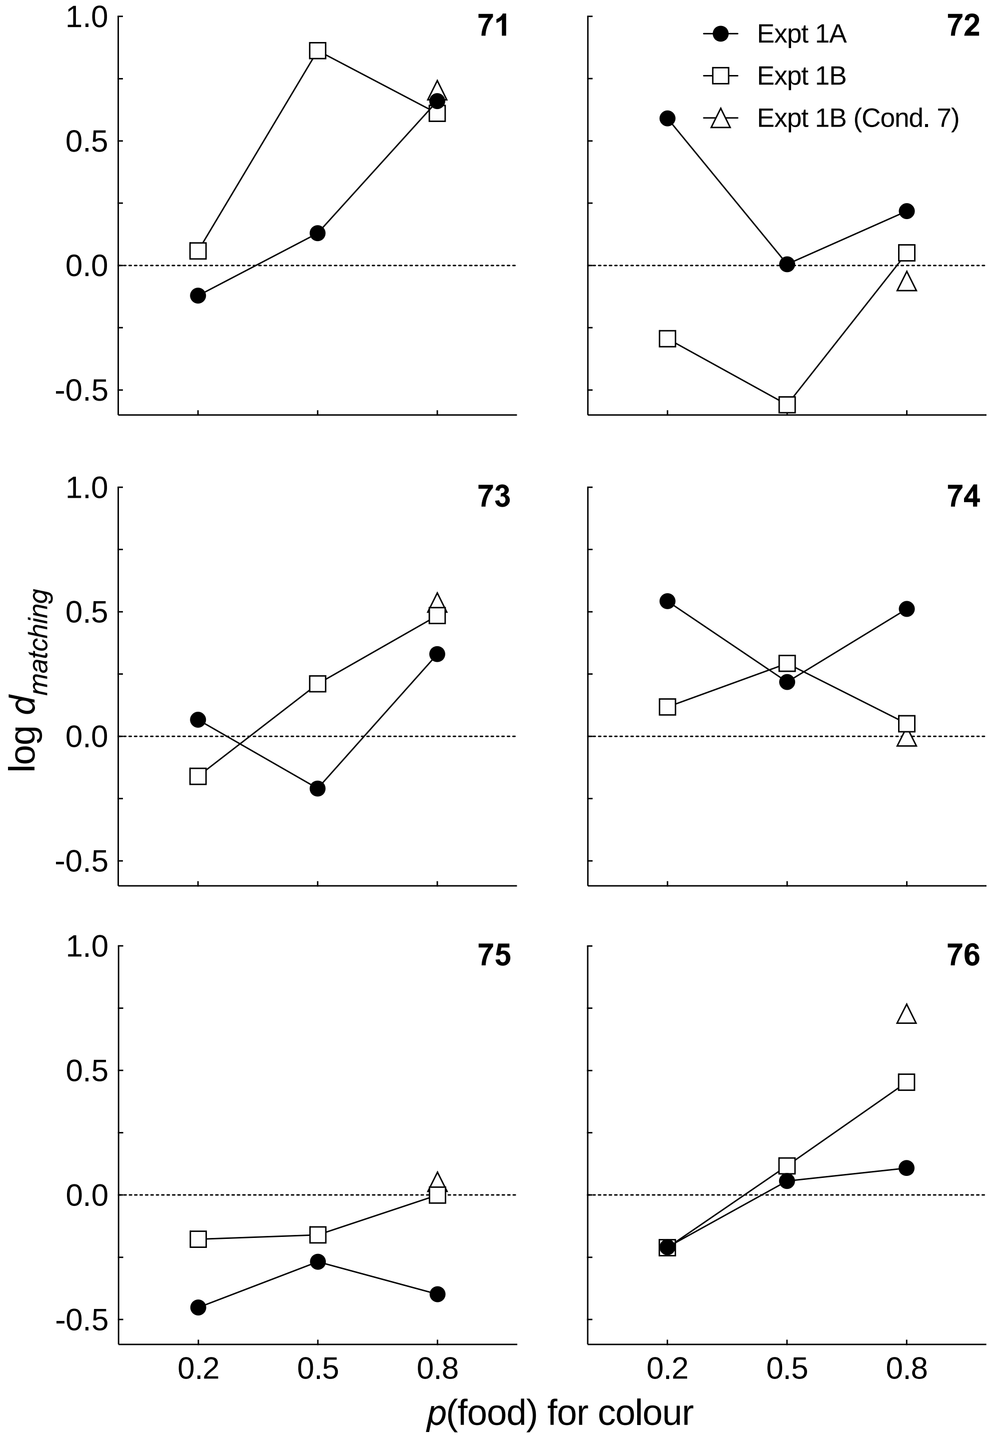

Figure 3.2. Preference for the comparison key matching the colour or flash-frequency element (log $d_{matching}$) in incompatible test trials in Experiment 1A (filled symbols) and Experiment 1B (unfilled symbols), plotted as a function of the probability of reinforcer delivery associated with colour stimuli. Data from Condition 7 (triangles), which was a replication of Condition 5, are shown separately from data from Conditions 4 to 6 (squares).

3.3.4 Experiment 1B

The unfilled data points in Figure 3.2 show log $d_{matching}$ (Equation 3.1) values for incompatible test trials in Experiment 1B. Unlike Experiment 1A, log $d_{matching}$ generally increased as the probability of reinforcer deliveries associated with colour stimuli increased for all pigeons in Experiment 1B ($\Sigma S = 10$, $p < .03$, $N = 6$, $k = 3$). Therefore, relative reinforcer rates appeared to determine divided stimulus control when the colour and flash-frequency elements were presented separately (Experiment 1B), but not when they were compounded together (Experiment 1A).

Although there was an increasing trend in log $d_{matching}$ values as the probability of reinforcer deliveries associated with colour stimuli increased in Experiment 1B (Figure 3.2), changes in log $d_{matching}$ as relative reinforcer rates changed were somewhat unsystematic across pigeons. A clear monotonic increase in log $d_{matching}$ values as the reinforcer probability associated with colour stimuli increased was only apparent for Pigeons 73 and 76. For the other pigeons, log $d_{matching}$ values increased non-monotonically (Pigeons 71 and 72) or hardly changed (Pigeons 74 and 75) as the relative reinforcer rate associated with colour stimuli increased. Therefore, relative reinforcer rates appeared to have only a small effect on divided stimulus control for most pigeons in Experiment 1B.

One explanation for why relative reinforcer rates appeared to have a small effect on log $d_{matching}$ in Experiment 1B (Figure 3.2) is that the relation between relative reinforcers and divided stimulus control may have been expressed in behavioural measures other than comparison choice. One such measure is choice during sample-stimulus presentations; because the colour and flash-frequency elements were presented on separate sample keys (Figure 3.1), analyses of sample-key choice provided an additional measure of divided control between the stimulus dimensions. Thus, we calculated preference for the colour and flash-frequency elements during sample-stimulus presentations (log $d_{samples}$) using Equation 3.1, except that responses made to the colour and flash-frequency elements were used.

Figure 3.3 shows log $d_{samples}$ in compatible and incompatible test trials in Experiment 1B. Positive values indicate stronger preference for the colour element, and negative values indicate stronger preference for the flash-frequency element. In both compatible and incompatible test trials, log $d_{samples}$ (Figure 3.3) increased as the probability of reinforcer deliveries associated with colour stimuli increased (both $\Sigma S = 18$, $p < .001$, $N = 6$, $k = 3$). Thus, as the relative reinforcer probability associated with colour stimuli increased, so did preference for the sample key displaying the colour element. This effect of relative reinforcer rates on sample choice was larger in incompatible test trials than in compatible test trials, as changes in log $d_{samples}$ were larger in incompatible trials.

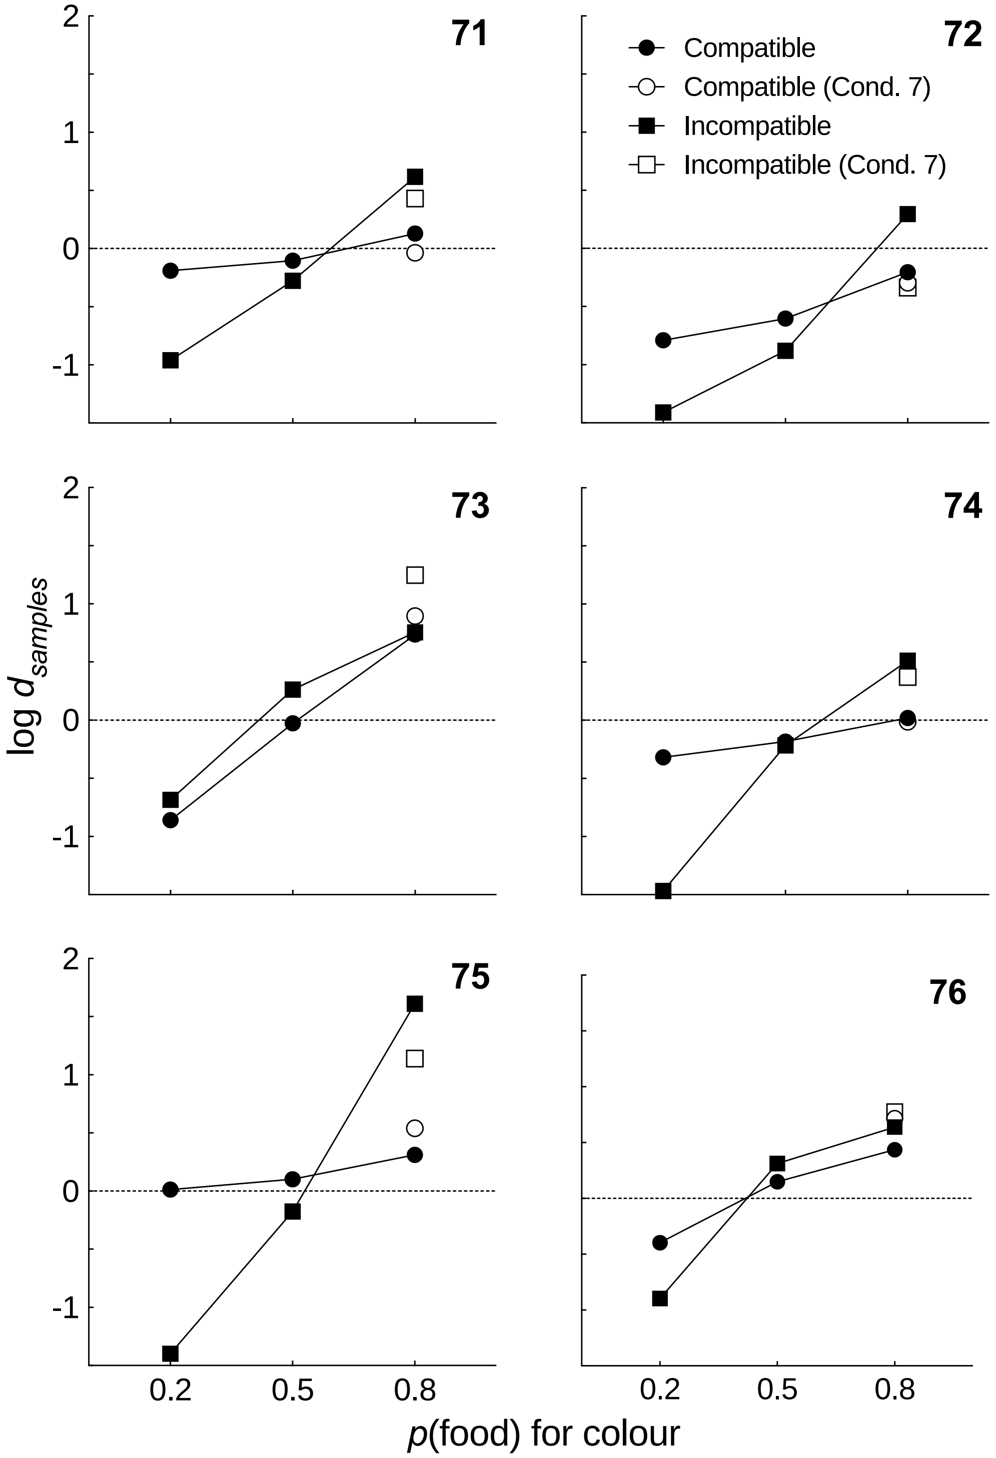

Figure 3.3. Preference for the colour or flash-frequency element during sample-stimulus presentations (log $d_{samples}$) as a function of the probability of reinforcer deliveries associated with colour stimuli in compatible (circles) and incompatible (squares) test trials in Experiment 1B. Filled symbols show data from Conditions 4 to 6, and unfilled symbols show data from Condition 7, which was a replication of Condition 5.

However, if the pigeons strongly preferred to peck the element associated with the higher reinforcer rate during sample-stimulus presentation (Figure 3.3), why did they not also prefer the comparison *matching* that element (Figure 3.2)? One possible explanation for this discrepancy is that the pigeons may have preferred the comparison that was *directly below* the element associated with the higher reinforcer rate. That is, they may have chosen the comparison that was physically closest to the element associated with the higher reinforcer rate, regardless of whether or not it matched that element. To determine whether this was the case, we reanalysed comparison-choice data using responses made to the comparison key below the colour and flash-frequency elements, rather than responses that matched those elements.

Figure 3.4 shows the results of this reanalysis (log $d_{comparison\text{-}side}$). Positive values of log $d_{comparison\text{-}side}$ indicate preference for the comparison key below the colour element, and negative values indicate preference for the comparison below the flash-frequency element. Clearly, log $d_{comparison\text{-}side}$ remained at or close to zero in compatible test trials (one-tailed nonparametric trend test: $\Sigma S = 7$, $p > .1$, $N = 6$, $k = 3$). This is unsurprising, because the pigeons preferred the comparison that matched the sample-stimulus elements in compatible test trials regardless of relative reinforcer rates (Table 3.2). That comparison would have been below the colour element in half of test trials and below the flash-frequency element in the other half, resulting in log $d_{comparison\text{-}side}$ values close to zero. In contrast, log $d_{comparison\text{-}side}$ values increased as the probability of reinforcer deliveries associated with colour stimuli increased ($\Sigma S = 18$, $p < .001$, $N = 6$, $k = 3$). Thus, in incompatible test trials, all pigeons preferred to peck the element associated with the higher reinforcer rate during sample-stimulus presentation, and then pecked the comparison below that element regardless of whether that comparison matched that element. This was further confirmed by analyses of the proportion of responses to the comparison key below the last-pecked sample key; the pigeons strongly preferred the comparison key that was below the last-pecked sample (Table 3.3). Data from training trials, and from Experiment 1A, are also shown in Table 3.3 for completeness.

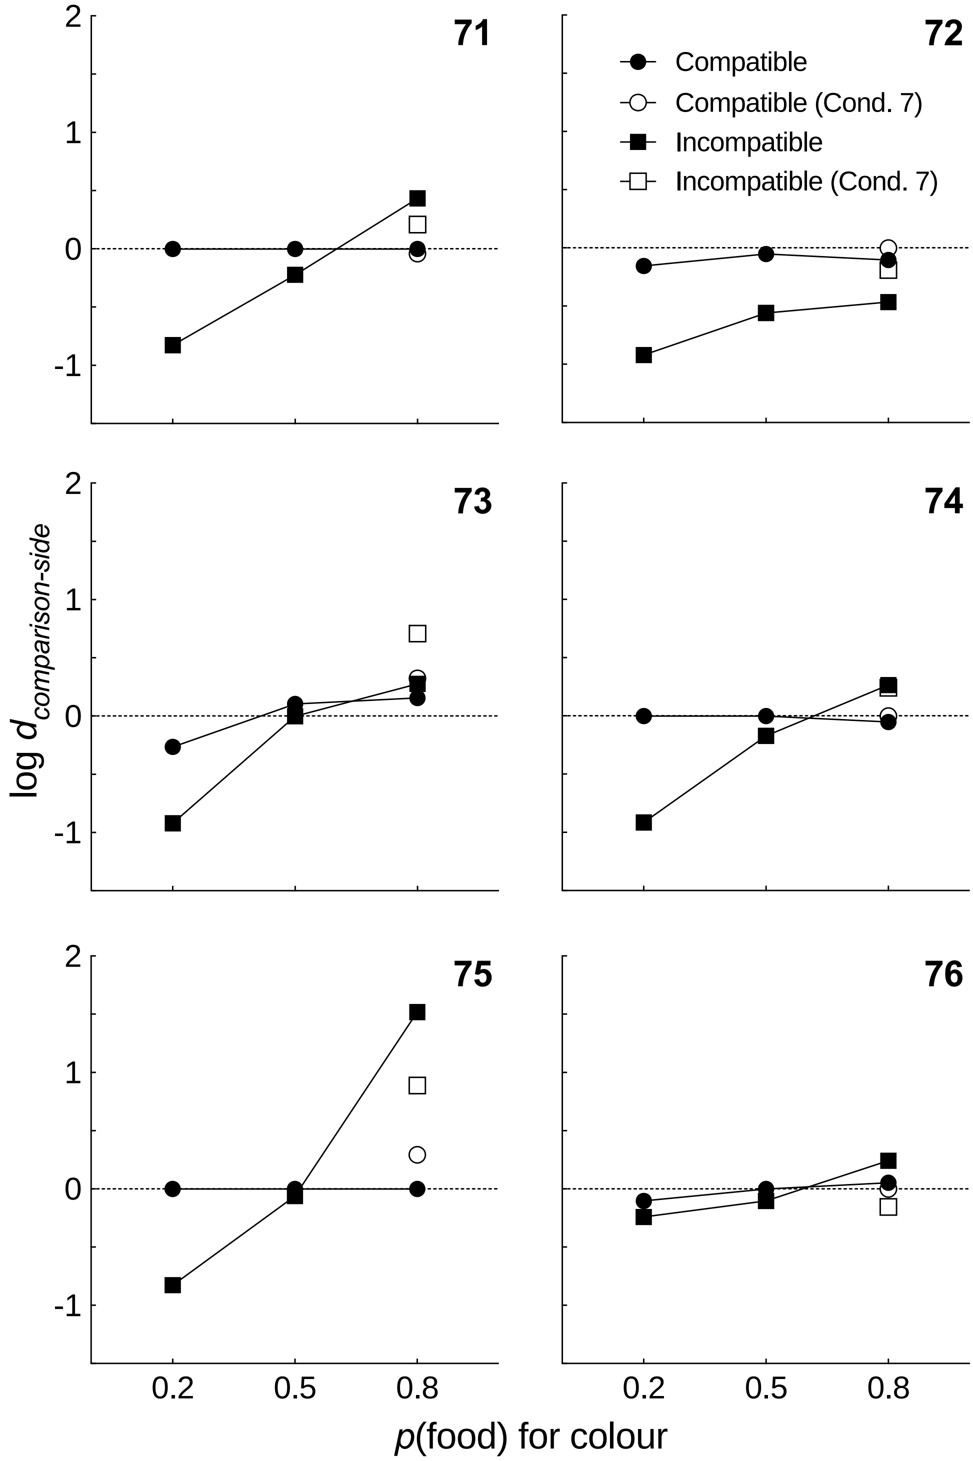

Figure 3.4. Preference for the comparison key on the same side as the colour or flash-frequency element (log $d_{comparison\text{-}side}$) as a function of the probability of reinforcer deliveries associated with colour stimuli in compatible (circles) and incompatible (squares) test trials in Experiment 1B. Filled symbols show data from Conditions 4 to 6, and unfilled symbols show data from Condition 7, which was a replication of Condition 5.

Table 3.3.

Mean (averaged across pigeons) proportion of comparison-key responses below the last-pecked sample key in training trials, compatible test trials, and incompatible test trials. Standard deviations are shown in parentheses.

	Training Trials				Compatible Test Trials		Incompatible Test Trials	
Condition	Red	Green	Fast	Slow	Containing Red	Containing Green	Containing Red	Containing Green
1	.82 (.38)	.85 (.35)	.83 (.40)	.99 (.01)	.83 (.41)	.97 (.04)	.84 (.32)	.86 (.24)
2	.72 (.41)	.82 (.40)	.82 (.37)	.99 (.02)	.76 (.39)	.90 (.17)	.71 (.42)	.72 (.38)
3	.77 (.40)	.84 (.37)	.79 (.40)	.99 (.01)	.81 (.40)	.86 (.24)	.66 (.37)	.73 (.42)
4	.90 (.20)	.68 (.47)	.87 (.23)	.98 (.03)	.95 (.07)	.80 (.22)	.82 (.21)	.82 (.21)
5	.90 (.21)	.52 (.52)	.78 (.34)	.99 (.02)	.92 (.15)	.73 (.32)	.84 (.21)	.72 (.32)
6	.87 (.29)	.58 (.47)	.81 (.28)	.99 (.03)	.94 (.15)	.77 (.29)	.79 (.22)	.75 (.25)
7	.85 (.36)	.64 (.42)	.77 (.36)	.97 (.03)	.94 (.13)	.75 (.28)	.79 (.27)	.72 (.31)

Note: “Containing Red” or “Containing Green” refer to test trials in which the colour element was red or green, respectively.

Therefore, in compatible test trials, the pigeons generally preferred to peck the element associated with the higher reinforcer rate (Figure 3.3) and they also preferred the correct comparison key (Table 3.2 and Figure 3.4), suggesting that element location and element identity jointly controlled behaviour. Control by the location of the correct comparison, as signalled by element identity, was much stronger than control by the location of the sample-stimulus elements. It is also possible that the identities and locations of the sample-stimulus elements jointly determined behaviour in incompatible test trials. To determine whether this was the case, we calculated log $d_{matching}$, log $d_{samples}$, and log $d_{comparison\text{-}side}$ values separately for incompatible test trials in which the elements were on the same side as their matching comparisons and trials in which the elements were on the other side. The extent of control by each variable was determined by linear regression, as the slope of the line provides an indication of how much log d_x values changed as the probability of reinforcer deliveries associated with colour stimuli increased. Table 3.4 shows these slopes. The positive slopes indicate that control by the colour element increased and control by the flash-frequency element decreased as the probability of reinforcer deliveries associated with the colour dimension increased. The negative slopes, which are mostly for log $d_{matching}$ values when the elements were on the opposite side from their matching comparisons, indicate that the pigeons preferred the comparison below the element associated with the higher reinforcer rate (see Figure 3.4).

For each measure of stimulus control (log $d_{matching}$, log $d_{samples}$, and log $d_{comparison\text{-}side}$), the absolute value of the slope (Table 3.4) was larger when the elements were on the same side as their matching comparisons compared with when they were on the other side for all pigeons except Pigeon 74. Thus, changes in relative reinforcer rates had larger effects on divided stimulus control when the elements were on the same side as their matching comparisons. However, the differences between the absolute values of the slopes were not statistically significant for all three measures of stimulus control (binomial sign tests: all $p > .1$). Nevertheless, the fact that such a difference was observed for five out of six pigeons suggests that in general, the location of the correct comparison also exerted some control over behaviour in incompatible test trials, although such control was weaker than control by the location of the sample-stimulus elements.

Table 3.4.

Slope of linear regression on log $d_{matching}$, log $d_{samples}$, and log $d_{comparison\text{-}side}$ as a function of the probability of reinforcer deliveries associated with colour stimuli when the colour and flash-frequency elements were on the same side as their matching comparisons, and when they were on the other side.

	log $d_{matching}$		log $d_{samples}$		log $d_{comparison\text{-}side}$	
Pigeon	Same	Other	Same	Other	Same	Other
71	2.80	-1.26	3.51	1.62	2.80	1.26
72	1.82	-0.17	2.66	2.06	1.82	0.17
73	3.01	-0.69	4.32	2.06	3.01	0.69
74	1.44	-1.86	2.83	2.96	1.44	1.87
75	3.84	-2.57	5.62	3.78	3.84	2.57
76	1.83	1.04	3.61	3.59	1.83	-1.04

3.3.5 Summary of Main Results

When separately trained colour and flash-frequency elements were compounded together (Experiment 1A), relative reinforcer rates did not determine divided stimulus control, as values of log $d_{matching}$ changed unsystematically with changes in relative reinforcer rates for all but two pigeons (Figure 3.2). In contrast, when the colour and flash-frequency elements were presented simultaneously on separate sample keys (Experiment 1B), relative reinforcer rates did determine divided stimulus control (Figures 3.2 to 3.4). However, this relation between relative reinforcer rates and divided stimulus control in Experiment 1B did not result in strong preference for the comparison key matching the element associated with the higher reinforcer rate in incompatible test trials (Figure 3.2). Although some control by the location of the correct comparison was evident in incompatible test trials (Table 3.4), such control was weak, whereas control by the location of the colour and flash-frequency elements was strong. That is, in Experiment 1B, the pigeons preferred to peck the element associated with the higher reinforcer rate during sample-stimulus presentation (Figure 3.3), and then pecked the comparison key on the same side as that element, regardless of whether that comparison matched that element (Figure 3.4 and Table 3.3).

3.4 Discussion

The present results extend previous research examining the relation between relative reinforcer rates and the division of control between the dimensions of compound stimuli in DMTS tasks (Davison, 2018a; Davison & Elliffe, 2010; Podlesnik et al., 2012; Shahan & Podlesnik, 2006, 2007) to a novel situation. Unlike previous research, stimuli (colours and flash-frequencies) were first trained individually, and then compounded together (Experiment 1A) or presented simultaneously on separate sample-stimulus keys (Experiment 1B) in the present experiments. The present experiments also extend research examining the effects of previously trained stimulus-response-reinforcer relations on divided stimulus control (Dube & McIlvane, 1997; Kalafut & Church, 2017; Matell & Kurti, 2014). Compared with such previous research, the present experiment provides stronger evidence that relative reinforcer rates determine divided control between separately trained stimuli. Here, the stimuli were associated with different responses, the comparison keys were the same for both dimensions, both stimulus dimensions were visual, and reinforcers were always delivered immediately after a correct response. These procedural factors made it easier, compared with Dube and McIlvane (1997), Kalafut and Church (2017), and Matell and Kurti (2014), to determine the extent of control by each stimulus dimension, and also reduced any potential effects of stimulus modality, interval to reinforcer delivery, or motivational processes during comparison choice.

In the present experiments, the effects of relative reinforcer rates on divided stimulus control varied depending on the configuration of the stimulus elements: Relative reinforcer rates did not determine divided stimulus control when the colour and flash-frequency elements were compounded together (Experiment 1A), but they did when the elements were presented separately (Experiment 1B; Figures 3.2 to 3.4). Thus, the results of Experiment 1B replicate the general conclusion of previous research that relative reinforcer rates divide stimulus control. However, the *behavioural expression* of that relation in Experiment 1B differed from previous research. In Experiment 1B, the effects of relative reinforcer rates on divided stimulus control were most evident in analyses of sample choice, and the pigeons preferred the comparison that was physically closest to the element associated with the higher reinforcer rate regardless of whether or not it matched that element (Figures 3.2 to 3.4). In contrast, pigeons preferred the comparison that matched the element associated with the higher reinforcer rate in previous research. The remainder of this discussion considers explanations for the differences between the results of Experiments 1A and 1B, and between our results and those of past research.

The high matching accuracies in compatible test trials (Table 3.2) suggest that the previously learned stimulus-response-reinforcer relations during training generalised to test trials in both Experiments 1A and 1B. Therefore, the results of Experiment 1A and 1B appear not to be related to a generalisation decrement from training to test trials, discriminability of the compound-stimulus elements, nor to our pigeons processing the compound stimulus as a separate entity (i.e., a novel stimulus) from its constituent elements. Although these factors may have contributed to the present results, any such contribution was probably small. Instead, taken together, the results of Experiment 1A, Experiment 1B, and those of past research suggest that the training procedure and stimulus configuration interact to determine the nature of the relation between relative reinforcer rates and divided stimulus control. Specifically, the differences between the results of Experiments 1A and 1B, and between Experiment 1B and previous research, appear to be related to differences in the configuration of the stimulus elements (compounded together vs. spatially separate). Additionally, the training procedure arranged in the present experiments differs from previous research, suggesting that the training procedure may also modulate the behavioural expression of the relation between relative reinforcers and divided stimulus control.

The results of Experiment 1B provide a possible explanation for the effects of the type of training procedure and stimulus configuration on the relation between relative reinforcer rates and divided stimulus control. In Experiment 1B, there were two variables competing for control over comparison choice in test trials: The location of the sample-stimulus elements, and the location of the comparison(s) matching those elements. This was evident in compatible test trials, in which preference for the sample key displaying the element associated with the higher reinforcer rate was attenuated by preference for the sample key on the same side as the correct comparison (Tables 3.2 and 3.3; Figures 3.3 and 3.4). Similarly, in incompatible test trials, the extent to which relative reinforcer rates determined divided stimulus control depended on whether the stimulus elements were on the same or opposite side as their matching comparisons. When the locations of the elements were *congruent* with the locations of their matching comparisons, the effects of relative reinforcer rates on divided stimulus control were *enhanced*, whereas when those locations were *incongruent*, the effects of relative reinforcer rates on divided stimulus control were *attenuated* (Table 3.4).

This competition between the location of the stimulus elements and the location of the matching comparisons in Experiment 1B likely occurred as a result of both the training procedure arranged and the stimulus configuration. During training, stimuli were always presented separately, and hence the location of the correct comparison key was predictable. As a result, the pigeons learned to orient towards the correct comparison during sample-stimulus presentations, and this orienting behaviour then determined which comparison key the pigeons chose (Table 3.3). Thus, comparison choice depended on the location of the last-pecked sample key. Although such orienting behaviour was not required or explicitly reinforced, it was probably maintained by the reinforcers obtained for choosing the correct comparison (i.e., by adventitious reinforcement; Skinner, 1948), especially given the close temporal proximity between the sample and comparison presentations (see e.g., Blough, 1959; Chatlosh & Wasserman, 1987; Chudasama & Muir, 1997; Hunter, 1913). When a colour and a flash-frequency element were then presented simultaneously on separate sample keys in incompatible test trials (see Figure 3.1), the pigeons oriented towards the element associated with the higher reinforcer rate (Figure 3.3), and then chose the comparison below that element (Figure 3.4).

Why did the pigeons orient towards the element associated with the higher reinforcer rate, and not towards the comparison key matching that element in incompatible test trials in Experiment 1B? One explanation is that as a consequence of the stimulus configuration, there was an additional source of information – the location of the sample-stimulus elements – signalling which comparison key to choose. In incompatible test trials, the colour and flash-frequency elements matched different comparisons, and so there was no clear “correct” choice in these trials (although the comparison matching the element associated with the higher reinforcer rate may appear to be more correct, this was apparently not the case for the pigeons; see Figure 3.2). Hence, the sample-stimulus elements provided conflicting information about which comparison key to choose. In contrast, the information provided by the element location was probably less ambiguous, because it signalled to choose either the left or right comparison, depending on the location of the element associated with the higher reinforcer rate. Additionally, when multiple stimuli are correlated with different reinforcer rates, pigeons tend to prefer to peck the stimuli that are generally better predictors of reinforcers (Castro & Wasserman, 2014, 2017; Dittrich, Rose, Buschmann, Bourdonnais, & Güntürkün, 2010; Jenkins & Sainsbury, 1969, 1970; Sainsbury, 1971; Wasserman, 1974). This may explain why the pigeons used the location of the sample-stimulus elements to guide choice in incompatible test trials (Figures 3.3 and 3.4; Tables 3.3 and 3.4). Thus, when stimuli provide conflicting information about the location of future reinforcers, subjects may use any additional sources of information available, even those not explicitly arranged, to guide behaviour.

In contrast to Experiment 1B, the sample-stimulus elements were compounded together and the compound stimulus was displayed on both sample keys in Experiment 1A. Hence, the element associated with the higher reinforcer rate was presented on both sample keys, and there was no additional information provided by element location to guide choice systematically in incompatible test trials. It is possible that the pigeons also preferred to peck the element associated with the higher reinforcer rate in Experiment 1A. If so, then sample-key choice would have been unsystematic, and as a result, comparison choice would have been equally unsystematic because comparison choice depended on the location of the last-pecked sample key (Table 3.3). Thus, the difference in the results of Experiments 1A and 1B may be an artefact of the present procedure; it may be related to the fact that the sample stimuli were presented on two discrete response keys that were positioned directly above the comparison keys.

The strong control over comparison choice by orienting behaviour in both of the present experiments, and the control by element location in Experiment 1B, can also explain differences between the present results and those of past research. In previous research (e.g., Davison & Elliffe, 2010; Dube & McIlvane, 1997; Matell & Kurti, 2014; Shahan & Podlesnik, 2006, 2007), the training procedure did not facilitate the development of orienting behaviour because the location of the correct comparison was always unpredictable. Additionally, unlike Experiment 1B, there was no competition between the location of the sample-stimulus elements and the location of the correct comparison in previous research. Thus, comparison choice was controlled by the identity of the sample-stimulus elements in previous research, whereas it was controlled by both the sample-stimulus identity and the location of the elements in Experiment 1B.

Therefore, the differences between the findings of Experiment 1A, Experiment 1B, and past experiments appear to be related to differences in the variables controlling comparison choice, which arose as a consequence of the training procedure (training the stimuli separately versus together) and of the stimulus configuration (presenting the stimulus elements on two keys positioned directly above the comparison keys versus compounding the elements together into a single stimulus). This line of reasoning implies that if the location of the correct comparison had been unpredictable in the present experiment (e.g., if we had arranged a symbolic matching procedure between colour samples and form comparisons, and counterbalanced the location of the forms), then comparison choice would have been controlled by the sample-stimulus elements themselves, rather than by the location of the last-pecked sample key. Alternatively, if the sample-stimulus elements had been presented on a single display centred between the comparison keys, this might have eliminated both the pigeons’ preference to peck the element associated with the higher reinforcer rate and the control by the location of the last-pecked sample key, because the elements would be about equidistant from the comparison keys.

The present experiments also share similarities with a study by Yokoyama, Dailey, and Chase (2006), who investigated divided control by compound colour-form stimuli following training with the individual stimulus elements in DMTS with pigeons. Unlike the preceding discussion, Yokoyama et al. concluded that stimulus control in test trials was *selective*. They based this conclusion on (1) the finding that matching accuracy in compatible test trials was comparable to accuracy in training and (2) the finding that pigeons sometimes chose the comparison matching colour, and other times the comparison matching form, in incompatible test trials. However, their conclusion may not be entirely correct. The first of the aforementioned findings was based on the prediction that if stimulus control were divided in test trials, then the probability of choosing the correct comparison should be higher in compatible test trials, due to the presence of two stimuli signalling the same comparison, than in training trials. However, this ‘summation’ effect (see also Chase & Heinemann, 1972; Cook & Wixted, 1997; Zuckerman, 1973) is limited by a ceiling, as matching accuracy cannot exceed 100%. Therefore, the failure to obtain such a summation effect cannot be taken as evidence of selective stimulus control, especially when matching accuracy is high, as it was in Yokoyama et al.’s and the present experiments (see Table 3.2). Second, because the comparison keys were discrete responses (as in the present experiments), pigeons could only choose one comparison in incompatible test trials, regardless of whether stimulus control was selective or divided.

The present experiments offer an alternative explanation for Yokoyama et al.’s (2006) results. In test trials, stimulus control may have been *divided* between the colour and form stimuli, resulting in a proportion of comparison-key responses matching each stimulus in incompatible test trials. Indeed, this is consistent with the present findings, and also with previous research on divided stimulus control (e.g., Davison & Elliffe, 2010). Furthermore, such divided control may have been related to relative reinforcer rates. Yokoyama et al. arranged equal reinforcer rates on the colour and form dimensions, and so the obtained reinforcer rates depended on matching accuracy on each dimension. In their first experiment, matching accuracies on both dimensions during training were similar, and hence the obtained reinforcer rates on each dimension were probably equal. In contrast, in their second experiment, matching accuracies were higher for colour stimuli than for form stimuli, and hence the obtained reinforcer rate on the colour dimension was probably higher. In the former experiment, comparison choice in incompatible test trials was about indifferent, whereas in the latter experiment, comparison choice in incompatible test trials was biased towards the comparison matching colour. That is, stimulus control appeared to be divided equally when the relative reinforcer rates were probably equal, and divided unequally when the relative reinforcer rates were probably unequal.

3.5 Chapter Summary and Conclusions

The present experiments asked whether relative reinforcer rates determine the division of control when separately trained stimuli are encountered together. The present results suggest that relative reinforcer rates determine divided stimulus control between separately trained stimuli if those stimuli remain separated when they are encountered together (Experiment 1B), whereas this is not the case if the stimuli are combined into a compound stimulus (Experiment 1A; Figures 3.2 to 3.4). Thus, Experiment 1B is the first to demonstrate clearly a systematic relation between relative reinforcer rates and the division of control between separately trained stimuli. These findings extend previous research demonstrating the same relation in DMTS tasks with compound stimuli (Davison, 2018a; Davison & Elliffe, 2010; Podlesnik et al., 2012; Shahan & Podlesnik, 2006, 2007) and also extend research examining the division of control between separately trained stimuli (Dube & McIlvane, 1997; Kalafut & Church, 2017; Matell & Kurti, 2014).

However, unlike previous studies, the stimulus dimensions did not exert strong discriminative control over behaviour in Experiment 1B. Instead, the relation between relative reinforcer rates and divided stimulus control was most evident in analyses of sample-key, rather than comparison-key, choice. This was because the identities of the sample-stimulus elements exerted much weaker control over behaviour than their locations. Thus, the present findings suggest that the type of training procedure and stimulus control interact to determine whether or not a systematic relation between relative reinforcer rates and divided stimulus control is observed, as well as the behavioural expression of that relation. These findings also suggest that subjects will use additional sources of information, such as the location of the stimulus elements in Experiment 1B, to guide choice when the location of future reinforcers is ambiguous.

The preceding discussion on the variables controlling comparison choice highlights some of the complexities of measuring divided stimulus control in the DMTS procedure. The present results show that comparison choice may not always provide a complete picture of divided stimulus control. If subjects learn to orient consistently towards one location (as in the present experiments) during training, such orienting behaviour may control comparison choice, and potentially mask control by other variables such as sample-stimulus identity or relative reinforcer rates (see e.g., Urcuioli, 1984, 1985; Urcuioli & Honig, 1980; Weaver, Dorrance, & Zentall, 1999). If this occurs, the effects of such variables may appear to be smaller than they actually are in analyses of comparison choice, and hence any conclusions about divided stimulus control based solely on comparison choice may be inaccurate or incomplete. This argument is similar to those put forth by early researchers of divided stimulus control, who suggested that conclusions about selective and divided stimulus control may differ depending on the type of test or measure used (e.g., Born, Snow, & Herbert, 1969; Farthing & Hearst, 1970; Wilkie & Masson, 1976; see Chapter 2). In such cases, additional analyses, such as analyses of sample choice, may help to clarify how stimulus control is divided and the variables controlling the division.

The present findings also suggest that manipulating relative reinforcer rates associated with individual stimuli may be a viable method to shift stimulus control away from irrelevant stimuli and towards relevant stimuli, or to reduce stimulus overselectivity and encourage divided stimulus control in conditional-discrimination tasks in applied interventions (see also Dube & McIlvane, 1997; Kelly et al., 2015; Reed et al., 2009, 2012). However, the present findings also suggest that there may be limits to the conditions under which such an intervention will succeed. When stimuli are compounded together, separately training stimulus-response-reinforcer relations may not determine the subsequent division of control between stimuli (Experiment 1A), whereas when stimuli are presented simultaneously but separately, the division of control may depend on prior training (Experiment 1B).

The present experiments bring us one step closer to establishing the conditions under which relative reinforcer rates determine divided stimulus control, and those conditions under which they do not. Future research investigating the effects of training procedures and stimulus configuration on divided stimulus control will help to elucidate how these procedural variables modulate divided stimulus control, and hence to establish further the generality of the relation between relative reinforcer rates and divided stimulus control. Thus, the experiment in the following chapter investigates further the effects of relative reinforcer rates on divided stimulus control between spatially separated stimuli.

Appendix A3

Number of training sessions for each pigeon in each condition

Table A3.1.

Number of training sessions for each pigeon in each condition. Training sessions ran for at least 31 sessions, and until accuracy for each sample stimulus was above 80% for five consecutive sessions.

	Condition						
Pigeon	1	2	3	4	5	6	7
71	31	31	31	31	31	31	31
72	38	31	31	51[†]	31	31	31
73	31*	36	35	31	44	31	31
74	31	31	31	31	31	31	31
75	31	31	31	31	31	31	31
76	31	31	31	34	31	31	31

*For Pigeon 73 in Condition 1, training sessions ended after 31 sessions despite accuracy for one sample stimulus being slightly below 80%.

[†] For Pigeon 72 in Condition 4, training sessions ran for longer than typical due to an equipment issue.

Appendix B3

License granting permission to include Experiment 1B in this thesis chapter

SPRINGER NATURE LICENSE
TERMS AND CONDITIONS

May 02, 2019

This Agreement between Ms. Stephanie Gomes-Ng ("You") and Springer Nature ("Springer Nature") consists of your license details and the terms and conditions provided by Springer Nature and Copyright Clearance Center.

License Number	4581060959823
License date	May 02, 2019
Licensed Content Publisher	Springer Nature
Licensed Content Publication	Learning & Behavior
Licensed Content Title	Relative reinforcer rates determine pigeons' attention allocation when separately trained stimuli are presented together
Licensed Content Author	Stephanie Gomes-Ng, Douglas Elliffe, Sarah Cowie
Licensed Content Date	Jan 1, 2019
Type of Use	Thesis/Dissertation
Requestor type	academic/university or research institute
Format	print and electronic
Portion	full article/chapter
Will you be translating?	no
Circulation/distribution	<501
Author of this Springer Nature content	yes
Title	Here or there, now or then? Relative reinforcement and divided stimulus control in space and time
Institution name	The University of Auckland
Expected presentation date	Jul 2019
Requestor Location	Ms. Stephanie Gomes-Ng New Zealand Attn: Ms. Stephanie Gomes-Ng
Total	0.00 USD

Terms and Conditions

Springer Nature Terms and Conditions for RightsLink Permissions

Springer Nature Customer Service Centre GmbH (the Licensor) hereby grants you a non-exclusive, world-wide licence to reproduce the material and for the purpose and requirements specified in the attached copy of your order form, and for no other use, subject to the conditions below:

1. The Licensor warrants that it has, to the best of its knowledge, the rights to license reuse of this material. However, you should ensure that the material you are requesting is original to the Licensor and does not carry the copyright of another entity (as credited in the published version).

CHAPTER 4

Experiment 2: Do Relative Reinforcer Rates Determine Divided Control Between Spatially Separated Stimuli?[8]

4.1 Introduction

Experiment 1B of the present thesis (Chapter 3) demonstrated that relative reinforcer rates determine the division of control between separately trained and spatially separated stimuli. However, unlike previous studies demonstrating a systematic relation between relative reinforcers and divided stimulus control (e.g., Davison & Elliffe, 2010; Shahan & Podlesnik, 2006), the sample-stimulus elements did not exert strong discriminative control over comparison choice in Experiment 1B; instead, comparison choice depended on the location of the sample-stimulus elements. Thus, whereas the frequency of correct responses according to a stimulus dimension increased as the reinforcer rate associated with that dimension increased in previous research, pigeons in Experiment 1B preferred to peck the sample-stimulus element associated with the higher reinforcer rate and then the comparison below that element regardless of whether it matched that element (see Figures 3.2 to 3.4). The present experiment was conducted as a follow-up of Experiment 1B, in order to investigate explanations for the results of Experiment 1B.

As discussed in Chapter 3, the differences between the results of previous studies and of Experiment 1B may be related to two key procedural differences: The configuration of the stimulus elements, and/or the type of training procedure arranged. To reiterate briefly, the stimulus elements (red, green, fast, and slow) were trained separately in Experiment 1B, and the pigeons learned to orient towards the correct comparison during training. As a result,

---

[8] This experiment was conducted under Approval 001396 granted by the University of Auckland Animal Ethics Committee.

comparison choice was controlled by such orienting behaviour (see Table 3.3). When the colour and flash-frequency elements were then presented simultaneously on separate response keys in incompatible test trials, the pigeons oriented towards and pecked the element associated with the higher reinforcer rate, and then simply pecked the comparison below that element. In previous studies of divided stimulus control (e.g., Davison & Elliffe, 2010), such orienting behaviour and control by element location was unlikely to develop because the dimension to match was unpredictable and the sample-stimulus dimensions were compounded into a single stimulus. Therefore, the orienting behaviour during sample-stimulus presentations and the strong control by element location may have masked any effects of relative reinforcer rates on comparison choice in Experiment 1B.

Additionally, or alternatively, it is possible that when the elements of a compound stimulus are spatially separated and are differentially correlated with reinforcers, as in Experiment 1B, sample choice, rather than comparison choice, better reflects the relation between relative reinforcers and divided stimulus control. Indeed, previous research suggests that the location of pigeons' key-pecks to complex discriminative stimuli (e.g., photographs of humans) provides an indication of which aspects of the stimuli exert the strongest control over behaviour (e.g., Allan, 1993; M. F. Brown, Cook, Lamb, & Riley, 1984; Castro & Wasserman, 2014, 2016; Dittrich, Rose, Buschmann, Bourdonnais, & Güntürkün, 2010), and that pigeons allocate relatively more pecks to those features of a stimulus that are more predictive of reinforcers (e.g., Castro & Wasserman, 2014, 2016, 2017; Jenkins & Sainsbury, 1969, 1970; Sainsbury, 1971; Wasserman, 1974; Wasserman & Anderson, 1974).

One way to determine the extent to which the configuration of the stimulus elements and the type of training procedure contributed to the results of Experiment 1B is to replicate systematically a typical divided-stimulus-control experiment (e.g., Davison & Elliffe, 2010) with spatially separated stimuli. If the spatial separation of the stimulus elements was responsible for the results of Experiment 1B, then such a replication should produce results similar to those of Experiment 1B. However, if the training procedure was responsible for the results of Experiment 1B – because it facilitated the development of orienting behaviour towards one location – then such a replication should produce results similar to previous studies (e.g., Davison & Elliffe, 2010; Shahan & Podlesnik, 2006, 2007), because the location of the correct comparison would be unpredictable.

Replicating a typical study of divided stimulus control with spatially separated stimuli also provides the opportunity to test a quantitative model of divided stimulus control proposed by Davison and Elliffe (2010). Recall that Davison and Elliffe's procedure involved compound stimuli consisting of a Fast-Slow dimension and a Long-Short dimension, and choice of a left comparison key was correct following Fast or Long stimuli and choice of a right comparison was correct following Slow or Short stimuli (see Chapter 2 for a more detailed description). Davison and Elliffe noted the similarities between this procedure and a "reinforcement-for-errors" procedure, in which a proportion of reinforcers is delivered following incorrect responses in a conditional-discrimination task (e.g., Davison & McCarthy, 1980; Nevin, Jenkins, Whittaker, & Yarensky, 1982; Nevin, Olson, Mandell, & Yarensky, 1975). Because the dimension to match was unsignalled in Davison and Elliffe's procedure (the comparison keys were always white), some reinforcers following correct responses according to one stimulus dimension were also reinforcers following error responses according to the other dimension. To illustrate, consider trials in which the sample stimulus was Fast and Long, or Slow and Short. In these trials, both dimensions matched the same key, and so a correct reinforced response according to one dimension was also a correct reinforced response according to the other dimension. Now consider trials in which the sample stimulus was Fast and Short or Slow and Long. In these trials, the dimensions matched different keys, and so a correct response according to one dimension was an incorrect response according to the other dimension. Therefore, any reinforcers obtained following correct responses according to one dimension were inconsistent with the other dimension. As a result, increasing the frequency of reinforcers following correct responses according to one dimension (*differential* reinforcers) also increased the frequency of reinforcers following error responses according to the other dimension (*nondifferential* reinforcers).

Differential reinforcers on a dimension increase the degree of discriminative control by that dimension, whereas nondifferential reinforcers decrease the degree of control by that dimension (Davison, 2018a; Davison & Elliffe, 2010; Davison & McCarthy, 1980; Nevin et al., 1975, 1982). Thus, Davison and Elliffe (2010) suggested that changes in relative reinforcer rates produce shifts in divided stimulus control in the DMTS task because of changes in the rates of differential and nondifferential reinforcers. That is, the dimension associated with the higher reinforcer rate is associated with a higher rate of differential reinforcers *and* a lower rate of nondifferential reinforcers, resulting in strong discriminative control by that dimension. In contrast, the dimension associated with the lower reinforcer rate is associated with a lower rate of differential reinforcers and a higher rate of nondifferential reinforcers, resulting in weak discriminative control by that dimension (see also Davison, 2018a).

To account for the effects of differential and nondifferential reinforcers on divided stimulus control, Davison and Elliffe (2010) applied Davison and Nevin's (1999) model of conditional-discrimination performance to their data. According to the Davison-Nevin model, behaviour is a function not of arranged contingencies, but of the contingencies *as they are discriminated by the subject.* The extent of mismatch between the experimenter-arranged and discriminated contingencies depends on the extent to which the discriminative stimuli are 'confusable' (i.e., *indiscriminable*), and also on the extent to which the response-reinforcer contingencies are indiscriminable. This confusability results in *reinforcer generalisation* between stimuli and between responses. That is, reinforcers obtained following correct responses for one stimulus may generalise to other stimuli and to other responses. The extent of such reinforcer generalisation increases as stimuli and responses become more indiscriminable. For example, as the physical disparity between S_1 and S_2 decreases, the degree of reinforcer generalisation between them increases (e.g., Davison & Jenkins, 1985; Davison & Nevin, 1999; Godfrey & Davison, 1998).

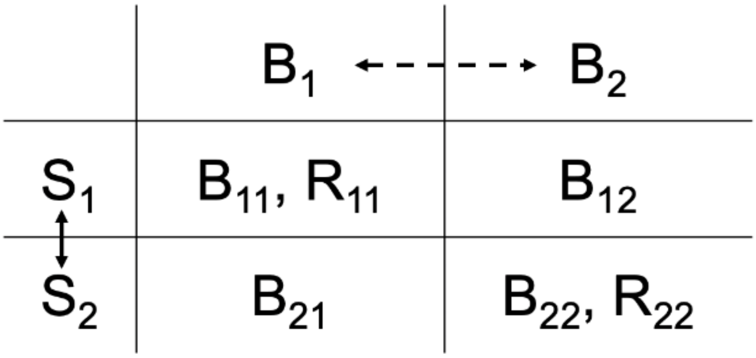

Figure 4.1. A 2x2 signal-detection matrix. S_1 and S_2 represent two discriminative stimuli, and B_1 and B_2 represent two responses. The dashed horizontal arrow represents reinforcer generalisation between responses, and the solid vertical arrow represents reinforcer generalisation between stimuli.

Reinforcer generalisation can be summarised using a 2x2 signal-detection matrix (Figure 4.1). In Figure 4.1, S_1 and S_2 represent two stimuli, and B_1 and B_2 represent two response alternatives. B_1 responses following S_1 are reinforced, whereas B_2 responses following S_2 are reinforced. R_{11} and R_{22} represent these reinforcers. The solid vertical arrow between S_1 and S_2 represents generalisation of reinforcers between stimuli (i.e., R_{11} reinforcers to S_2, and R_{22} reinforcers to S_1), and the dashed horizontal arrow between B_1 and B_2 represents generalisation of reinforcers between responses (i.e., R_{11} reinforcers to B_2

responses, and R_{22} reinforcers to B_1 responses). Davison and Nevin (1999) quantified such reinforcer generalisation between stimuli and between responses with the parameters d_{sb} and d_{br}, respectively, which range from 1 (no discrimination, complete generalisation) to ∞ (perfect discrimination, no generalisation).

Thus, in the Davison-Nevin model (Figure 4.1), the effects of differential and nondifferential reinforcers on behaviour are modelled by generalising a proportion of reinforcers from one stimulus to other stimuli (i.e., S_1 to S_2 and vice versa), and from correct responses to incorrect responses (i.e., B_1 to B_2 and vice versa). Therefore, in applying Davison and Nevin's (1999) model to their data, Davison and Elliffe (2010) suggested that in addition to differential and nondifferential reinforcers, reinforcer generalisation between the stimuli along a dimension and between responses also contributes to divided stimulus control (see also Davison, 2018a). In support of this, the Davison-Nevin model described Davison and Elliffe's data well. However, at present, the model's applicability to divided stimulus control is not well established, because only Davison and Elliffe have fitted the Davison-Nevin model to divided-stimulus-control data.

Therefore, the present experiment asked whether relative reinforcer rates determine divided stimulus control between spatially separated stimuli, and whether the Davison-Nevin model could describe any such effects, in a systematic replication of Davison and Elliffe's (2010) procedure. We arranged a symbolic DMTS procedure in which sample stimuli were compounds consisting of two elements from different dimensions – a colour (red or green) and a speed of keylight on/off alternation (fast or slow). The two sample-stimulus elements were presented on spatially separate response keys, as in Experiment 1B. The comparison stimuli were the left and right keys, and one element from each dimension matched the left comparison, while the other elements matched the right comparison. Across conditions, the probability of reinforcers following correct responses according to each stimulus dimension varied. This procedure allowed us to determine (1) whether the results of Experiment 1B were primarily related to stimulus configuration or to the type of training procedure, and (2) whether Davison and Nevin's (1999) model can describe the effects of relative reinforcers on divided control between spatially separated stimuli.

4.2 Method

4.2.1 Subjects and Apparatus

The subjects and apparatus were the same as in Experiment 1 (Chapter 3).

4.2.2 Procedure

Sessions were run daily, beginning at 1 a.m. The pigeons were run successively and in numerical order. Each session lasted until 144 trials or 65 min had elapsed, whichever occurred first.

A symbolic 0-s DMTS procedure was arranged, in which each sample stimulus was a compound consisting of two elements from different dimensions, and the pigeons matched one of those elements to a left or right comparison key in each trial. The compound-stimulus dimensions were colour (red or green) and the speed at which a white keylight alternated on and off (every 0.1 or 0.5 s, hereafter termed 'fast' and 'slow'). There were thus four sample stimuli: red-fast, red-slow, green-fast, and green-slow. One colour and one flash-frequency matched the left comparison, and the other colour and flash-frequency matched the right comparison. The element-comparison associations were counterbalanced across pigeons. For Pigeons 71 and 75, the red and slow elements matched the left key; for Pigeons 74 and 76, the red and fast elements matched the left key; for Pigeon 72, the green and fast elements matched the left key; and for Pigeon 73, the green and slow elements matched the left key. The other stimuli matched the right key. Thus, for three pigeons (Pigeons 71, 73, and 76), the element-comparison associations for the colour dimension were unchanged from Experiment

1 but were reversed for the flash-frequency dimension. For the other three pigeons (Pigeons 72, 74, and 75), the element-comparison associations for the flash-frequency dimension remained unchanged from Experiment 1 but were reversed for the colour dimension.

At the beginning of each trial, one of the four compound sample stimuli was selected probabilistically ($p = .25$), with the constraints that each sample was presented an equal number of times in each session and the same stimulus could not occur more than four times consecutively. The sample stimulus dimensions were presented separately on the top-left and top-right keys (collectively termed the ‘sample’ keys). The location of the colour and flash-frequency elements was counterbalanced; for each sample stimulus, the colour was presented on the left sample key and the flash-frequency on the right sample key in half of the trials in each session, and vice versa in the other half.

The first response to either sample key that occurred at least 5 s after the onset of the sample stimulus resulted in the offset of the sample and the onset of the comparison keys, which were the bottom-left and bottom-right keys. The comparison keys were always lit white. One of the comparison keys was defined as correct and the other as incorrect, depending on the dimension to match (colour or flash-frequency), which was selected with $p = .5$, with the constraints that each dimension was selected an equal number of times per sample-stimulus configuration (i.e., colour on left or right sample key). If the colour dimension was selected, then the correct comparison (left or right) was the key that matched the colour of the sample stimulus. If the flash-frequency dimension was selected, then the correct comparison was the key that matched the flash-frequency. A correct response was immediately followed by either 2-s access to food or a 2-s blackout, depending on whether the program had arranged a reinforcer. Incorrect responses were followed by a 4-s blackout. Thereafter, a 3-s blackout (the inter-trial interval) occurred. Figure 4.2 provides an overview of trial-type selection in the present experiment.

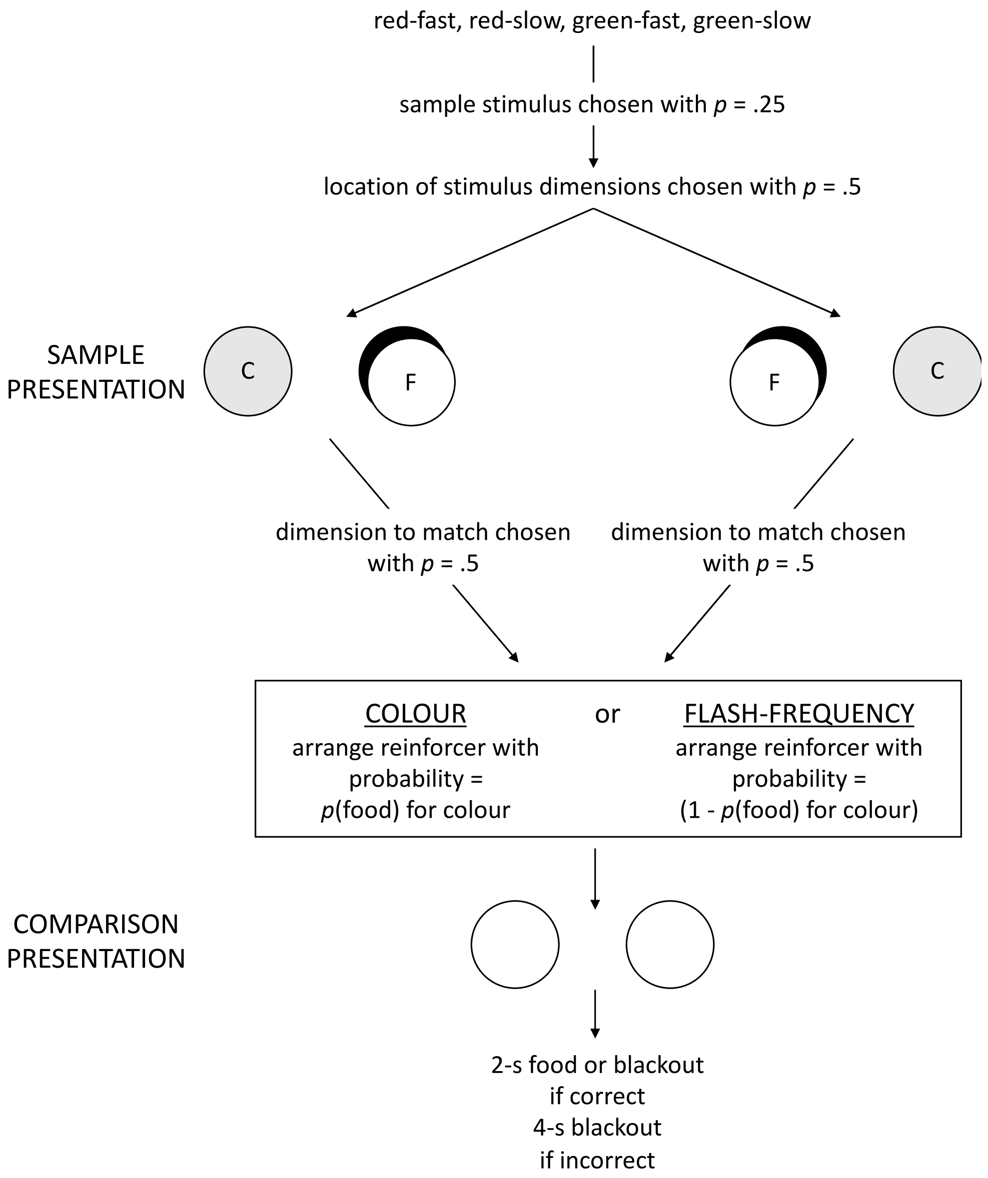

Figure 4.2. Overview of trial-type selection. ‘C’ and ‘F’ represent ‘colour’ and ‘flash-frequency’, respectively. Correct and incorrect comparison-key responses depended on the dimension to match. See text for further details.

4.2.2.1 Pretraining

The goal of pretraining was to reduce the percentage of reinforced trials from 100% to 50% before the experiment proper. The pigeons were trained with the procedure described above and with all correct responses reinforced for 22 sessions. Thereafter, the number of arranged reinforced trials was reduced to 120 out of 144 for three sessions, and then further reduced by 16 trials every third session until the number of arranged reinforced trials reached 72. In each session, reinforced and unreinforced trials were intermixed randomly, and the number of arranged reinforced and unreinforced trials was split evenly between the compound stimuli. After five sessions with 50% of trials reinforced, the experiment proper began.

4.2.2.2 Experimental Sessions

Experimental sessions were identical to pretraining sessions, except that 50% of trials were reinforced and the probabilities of reinforcer deliveries following correct responses according to the colour and flash-frequency dimensions were varied across conditions. Reinforcers were arranged independently; after the dimension to match was selected in each trial, the program decided whether or not to arrange a reinforcer according to the reinforcer probability associated with the selected dimension (see Figure 4.2). In Conditions 1 and 4, the probability of reinforcer deliveries associated with both dimensions was .5. In Condition 2, the probability of reinforcer deliveries was .8 for the colour dimension and .2 for the flash-frequency dimension. Condition 3 arranged the reverse probabilities of Condition 2. The order of Conditions 2 and 3 was counterbalanced across pigeons; Pigeons 71 to 73 experienced Condition 2 first, and Pigeons 74 to 76 experienced Condition 3 first. Each condition lasted for 30 sessions. Responding was stable, determined by visual inspection, by the 15th session of each condition. Hence, the last 15 sessions were used for all analyses.

4.3 Results

4.3.1 Proportions of Obtained Reinforcers

Table 4.1 shows the proportion of obtained reinforcers on each dimension, calculated across all trials and separately for each compound stimulus. Because the element-comparison associations were counterbalanced across pigeons, the compound stimuli are labelled according to whether the elements matched the same comparison (‘Compatible’ trials) or different comparisons (‘Incompatible’ trials) in Table 4.1.

The overall proportion of obtained reinforcers associated with the colour and flash-frequency dimensions deviated from the arranged reinforcer probabilities of .2, .5, or .8 (Table 4.1). Nevertheless, the proportions of obtained reinforcers on the colour and flash-frequency dimensions were more equal when the arranged probabilities were equal (i.e., *p*(food) for colour = .5), and were unequal when the arranged probabilities were unequal (i.e., *p*(food) for colour = .2 or .8). In the latter conditions, a greater proportion of reinforcers was obtained on the dimension with an arranged reinforcer probability of .8 (i.e., flash-frequency when *p*(food) for colour = .2, and colour when *p*(food) for colour = .8). Therefore, despite differences between the absolute proportions of arranged versus obtained reinforcers, the *relative* differences between the proportion of obtained reinforcers associated with the colour and flash-frequency dimensions were consistent with the arranged relative differences.

Closer inspection of the proportions of obtained reinforcers for each individual compound stimulus (labelled Red-Compatible, Green-Compatible, Red-Incompatible, and Green-Incompatible in Table 4.1) helps to clarify the reason for the differences between overall arranged and obtained proportions of reinforcers. When the colour and flash-frequency elements matched the same comparison key (compatible stimuli), the proportions of obtained reinforcers were close to the arranged reinforcer probabilities for all pigeons and in all conditions. However, this was not the case when the stimulus elements matched different comparisons (incompatible stimuli). With such stimuli, the proportion of obtained reinforcers was much greater for one dimension in conditions in which the arranged probabilities were equal. In conditions in which the arranged probabilities were unequal, relative differences between the proportion of obtained reinforcers on the colour and flash-frequency dimensions were generally less extreme than the arranged differences. These deviations in the arranged and obtained reinforcer contingencies are unsurprising, given that the dimension to match was always unsignalled and that reinforcers were arranged independently. Thus, if one dimension exerted stronger control than the other, the proportion of obtained reinforcers would be greater on that dimension.

Table 4.1.

Proportion of obtained reinforcers for each stimulus dimension, across all trial types (overall), and separated based on the compound stimulus.

	Overall		Red-Compatible		Green-Compatible		Red-Incompatible		Green-Incompatible	
Pigeon	Colour	Flash-Freq	Colour	Flash-Freq	Colour	Flash-Freq	Colour	Flash-Freq	Colour	Flash-Freq
p(food) for colour = .5										
71	.49	.33	.49	.50	.49	.50	.49	.12	.47	.21
72	.37	.46	.49	.48	.50	.50	.25	.46	.24	.41
73	.49	.29	.50	.50	.48	.49	.50	.11	.49	.14
74	.36	.43	.44	.45	.50	.50	.07	.50	.42	.27
75	.35	.48	.50	.49	.50	.50	.26	.46	.14	.48
76	.49	.31	.50	.50	.49	.50	.49	.17	.49	.09
p(food) for colour = .2										
71	.16	.55	.17	.83	.17	.83	.17	.18	.15	.36
72	.14	.72	.17	.81	.17	.78	.13	.57	.10	.72
73	.17	.45	.17	.83	.17	.79	.17	.01	.17	.16
74	.11	.74	.15	.75	.17	.81	.04	.78	.09	.63
75	.12	.80	.17	.82	.17	.82	.09	.76	.04	.81
76	.16	.48	.17	.76	.17	.75	.16	.21	.16	.18
p(food) for colour = .8										
71	.79	.11	.83	.17	.80	.16	.83	.00	.71	.10
72	.72	.13	.78	.17	.80	.16	.70	.10	.60	.12
73	.80	.10	.82	.16	.79	.16	.83	.01	.77	.04
74	.59	.15	.76	.16	.81	.17	.14	.16	.64	.11
75	.58	.16	.83	.17	.83	.17	.36	.16	.29	.16
76	.80	.10	.81	.17	.79	.17	.77	.04	.81	.04
p(food) for colour = .5 (replication)										
71	.48	.36	.50	.50	.50	.50	.50	.21	.50	.23
72	.39	.46	.50	.50	.50	.49	.31	.41	.26	.42
73	.50	.28	.49	.50	.50	.50	.50	.02	.50	.09
74	.38	.43	.43	.44	.50	.50	.18	.49	.42	.27
75	.35	.47	.50	.50	.50	.50	.19	.45	.19	.41
76	.48	.31	.48	.47	.50	.48	.43	.20	.50	.09

Note: ‘X-Compatible’ and ‘X-Incompatible’ refer to trials in which the colour element was X, and the elements matched the same or different comparisons, respectively.

4.3.2 Discriminative Control by Colour and Flash-Frequency

Responses to each comparison key in the last 15 sessions of each condition were separated according to the colour of the stimulus (i.e., red or green) or according to its flash-frequency (i.e., fast or slow). We then used the former to calculate a measure of discrimination, log d_{colour}, for the colour dimension, and the latter to calculate log $d_{flash\text{-}freq}$ for the flash-frequency dimension, using the following equation (Davison & Nevin, 1999; Davison & Tustin, 1978):

$$\log d_x = 0.5 \log\left(\frac{B_{corr\,|\,S1}}{B_{incorr\,|\,S1}} \cdot \frac{B_{corr\,|\,S2}}{B_{incorr\,|\,S2}}\right), \qquad \text{(Equation 4.1)}^{9}$$

where *x* is a placeholder for the colour or flash-frequency dimension. $B_{corr\,|\,Sx}$ and $B_{incorr\,|\,Sx}$ represent the number of correct and incorrect responses, respectively, in trial type Sx. For log d_{colour}, *S*1 and *S*2 are trials in which the colour element was red and green, respectively, and for log $d_{flash\text{-}freq}$, *S*1 and *S*2 are trials in which the flash-frequency element was fast and slow, respectively. Thus, for example, for log d_{colour}, $B_{corr\,|\,S1}$ consisted of responses to the comparison key that matched red following both red-fast and red-slow stimuli and $B_{incorr\,|\,S1}$ consisted of responses to the other comparison. Likewise, for log $d_{flash\text{-}freq}$, $B_{corr\,|\,S1}$ consisted of responses to the comparison that matched fast following red-fast and green-fast stimuli, and $B_{incorr\,|\,S1}$ consisted of responses to the other comparison. A log d_x value of zero indicates chance performance (i.e., no discrimination) on a dimension, and hence no control by that dimension. Higher values of log d_x indicate better discrimination

---

[9] This log *d* measure differs from that used in Experiment 1 (see Equation 3.1) because data from compatible trials (which served as control trials) and incompatible trials (which provided an indication of divided stimulus control) were separated for analyses in Experiment 1. No such separation of compatible and incompatible trials was required in the present experiment because all trials involved compound stimuli. Thus, Equation 4.1 is a *dimensional* measure of stimulus control (see e.g., Davison & Elliffe, 2010).

along, and hence stronger control by, a dimension. We also calculated a related measure, log b_x, which is a measure of response bias in favour of one key (e.g., left) over the other:

$$\log b_x = 0.5 \log \left(\frac{B_{corr\,|\,S1}}{B_{incorr\,|\,S1}} \cdot \frac{B_{incorr\,|\,S2}}{B_{corr\,|\,S2}} \right), \qquad \text{(Equation 4.2)}$$

where the variables are as in Equation 4.1. For log b_{colour}, more positive values indicate a stronger bias towards the comparison that matched red in all trials (e.g., towards left for Pigeon 71), and more negative values indicate a stronger bias towards the comparison that matched green. For log $b_{flash\text{-}freq}$, more positive values indicate a stronger bias towards the comparison that matched fast in all trials and more negative values indicate a stronger bias towards the comparison that matched slow.

Figure 4.3 shows log d_{colour} and log $d_{flash\text{-}freq}$ values (Equation 4.1) as a function of the probability of reinforcer deliveries associated with the colour dimension for all pigeons. Clearly, one dimension exerted much stronger control over behaviour than the other dimension for all pigeons – for Pigeons 71, 73, and 76, the colour dimension exerted stronger control, and for Pigeons 72, 74, and 75, the flash-frequency dimension exerted stronger control. This strong control by one dimension appears to be related to carryover effects from the pigeons’ previous experience in Experiment 1. In the present experiment, the element-comparison associations for one stimulus dimension remained unchanged from Experiment 1, whereas the element-comparison associations for the other dimension were reversed. For all pigeons, the unchanged dimension exerted stronger control than the reversed dimension in the present experiment (Figure 4.3). Furthermore, the strong control by one dimension is consistent with the obtained proportions of reinforcers on each dimension (Table 4.1). The dimension exerting stronger control was also the dimension associated with a higher proportion of obtained reinforcers in conditions in which the arranged reinforcer probabilities were equal. When the arranged reinforcer probabilities were unequal and favoured the dimension exerting stronger control, the proportions of obtained reinforcers were close to their arranged values, whereas when the arranged reinforcer probabilities favoured the other dimension, the strong control by the unchanged dimension attenuated the proportion of obtained reinforcers on the other dimension. Therefore, control by the unchanged dimension may have carried over from Experiment 1, and then been maintained by the generally higher proportion of obtained reinforcers on that dimension due to the dynamical relation between behaviour and reinforcers.

Nevertheless, relative reinforcer probabilities systematically affected log d_x in the present experiment. In general, log d_{colour} increased and log $d_{flash\text{-}freq}$ decreased as the probability of reinforcer deliveries associated with the colour dimension increased (Figure 4.3). One-tailed nonparametric trend tests (Elliffe & Elliffe, 2019; Kendall, 1955) on log d_{colour} and log $d_{flash\text{-}freq}$ confirmed these trends ($\Sigma S = 12$ for colour, -12 for flash-frequency, both $p < .01$, $N = 6$, $k = 3$). The only exception was Pigeon 73, for whom log d_x values changed little across conditions. Response bias (log b_x, Equation 4.2; data not shown here) did not change systematically with changes in relative reinforcer probabilities (two-tailed nonparametric trend test: $\Sigma S = -6$, $p > .2$, $N = 12$, $k = 3$). Thus, relative reinforcer rates systematically affected discrimination along a dimension (i.e., the degree of control by a dimension), but not response bias. The results of Conditions 1 and 4, which were replication conditions, were similar, hence, the effect of relative reinforcers on divided stimulus control was replicable.

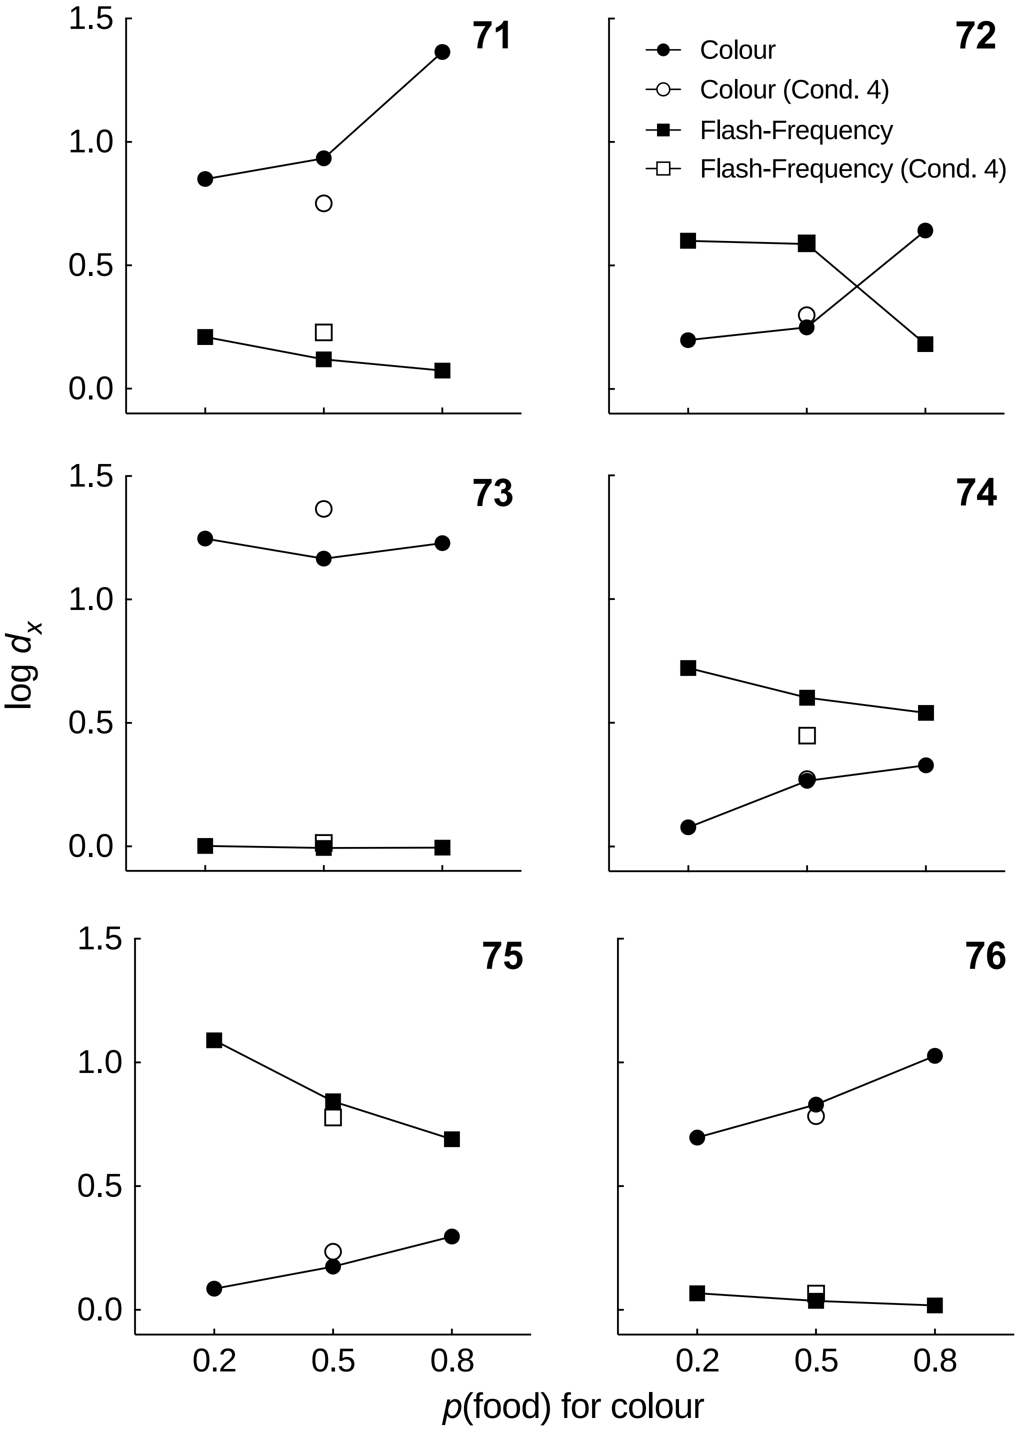

Figure 4.3. Point estimates of log d_x for the colour (circles) and flash-frequency (squares) discriminations as a function of the probability of reinforcers associated with the colour dimension. Filled data points show data from Conditions 1 to 3, and unfilled data points show data from Condition 4, which was a replication of Condition 1.

4.3.3 Preference for Colour and Flash-Frequency during Sample Presentations

To determine whether the pigeons preferred to peck the element associated with the higher reinforcer rate as they did in Experiment 1B, responses to the colour and flash-frequency elements during sample-stimulus presentations were aggregated across the last 15 sessions of each condition, and were entered into Equation 4.1 to calculate preference for each dimension. $B_{corr\,|\,Sx}$ and $B_{incorr\,|\,Sx}$ consisted of responses made to the colour and flash-frequency elements in trial type Sx for the colour dimension, and vice versa for the flash-frequency dimension. Figure 4.4 shows log d_x values for sample-key analyses. For Pigeon 76, preference for the colour dimension generally increased and preference for the flash-frequency dimension generally decreased as the probability of reinforcers associated with colour stimuli increased. A similar pattern was evident for Pigeons 71 and 73, although changes in preference across conditions were much smaller for these two pigeons. For the remaining three pigeons, preference changed little (Pigeons 74 and 75), or the opposite pattern was observed (Pigeon 72). Thus, unlike comparison choice (Figure 4.3), relative reinforcer rates had little systematic effect on sample choice across pigeons (Figure 4.4). One-tailed nonparametric trend tests on sample choice were not significant (ΣS = 6 for colour, -6 for flash-frequency, both $p > .2$, $N = 6$, $k = 3$). As for comparison choice, sample choice in Condition 4 was similar to sample choice in Condition 1.

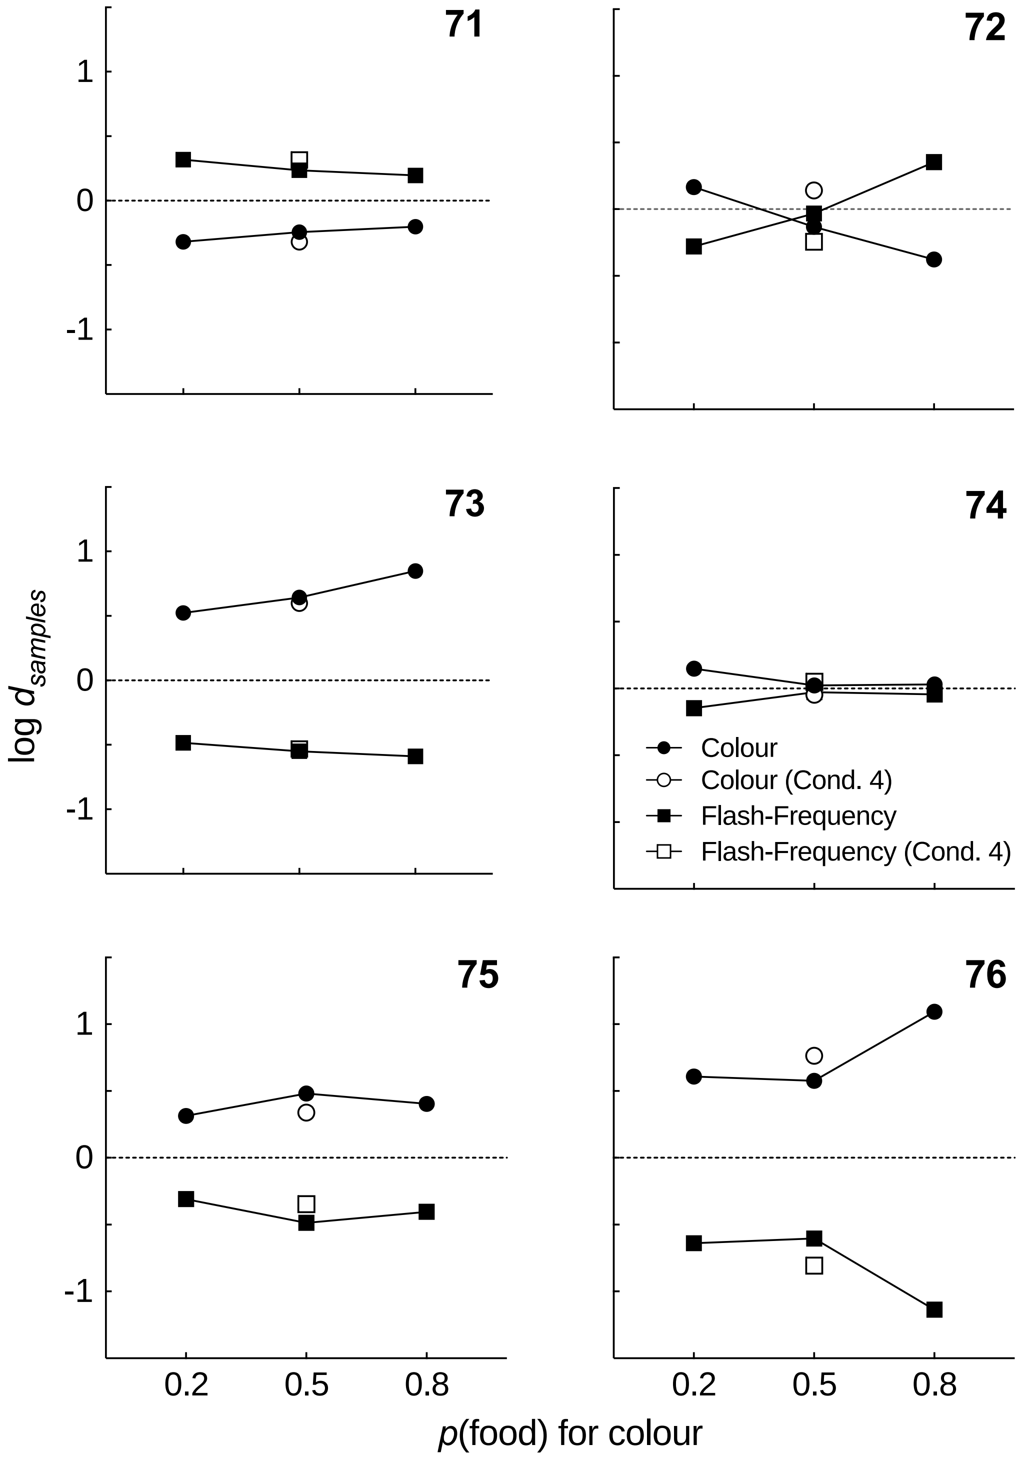

Figure 4.4. Point estimates of log d_x for sample-key choice as a function of the probability of reinforcers associated with the colour dimension. Filled data points show data from Conditions 1 to 3, and unfilled data points show data from Condition 4.

4.3.4 The Davison-Nevin Model

Finally, we fit Davison and Nevin's (1999) model to log d_{colour} and log $d_{flash\text{-}freq}$ values (Figure 4.3), in order to obtain estimates of d_{sb} and d_{br}, which quantify the degree of generalisation between stimuli and response-reinforcer contingencies, respectively (see Figure 4.1). The model included two d_{sb} parameters, one for each stimulus dimension, and one d_{br} parameter (because the comparison stimuli were always the same; see Davison & Elliffe, 2010). To fit the Davison-Nevin model to the data, we calculated, for each stimulus dimension, the effective reinforcer counts for each cell of the matrix shown in Figure 4.1. For the colour dimension, R_{11} consisted of reinforcers obtained for correct B_1 responses following red stimuli, and R_{22} consisted of reinforcers obtained for correct B_2 responses following green stimuli. R_{12} and R_{21} are not shown in the matrix in Figure 4.1, but they represent reinforcers for errors. Thus, R_{12} consisted of reinforcers obtained for a B_2 response following red stimuli, and R_{21} consisted of reinforcers obtained for a B_1 response following green stimuli. The logic is similar for the flash-frequency dimension. Using these effective reinforcer counts, we fit the following equation to log d_{colour} and log $d_{flash\text{-}freq}$ values, where x is a placeholder for the stimulus dimension (see Davison & Elliffe, 2010):

$$\log d_x = 0.5 \log \frac{R_{11} + \frac{R_{12}}{d_{br}} + \frac{R_{21}}{d_{sbx}} + \frac{R_{22}}{d_{sbx}d_{br}}}{R_{12} + \frac{R_{11}}{d_{br}} + \frac{R_{21}}{d_{sbx}d_{br}} + \frac{R_{22}}{d_{sbx}}} \times \frac{R_{22} + \frac{R_{21}}{d_{br}} + \frac{R_{12}}{d_{sbx}} + \frac{R_{11}}{d_{sbx}d_{br}}}{R_{21} + \frac{R_{22}}{d_{br}} + \frac{R_{12}}{d_{sbx}d_{br}} + \frac{R_{11}}{d_{sbx}}}. \qquad \text{(Equation 4.3)}$$

The fits were done simultaneously for the colour and flash-frequency dimensions using Microsoft® Excel Solver.

Table 4.2 shows log d_{sb} and log d_{br} values, as well as the proportion of variance accounted for (VAC), from fits of Equation 4.3 to individual-pigeon data. VACs were all at least .90, and at least .98 for 5 of the 6 pigeons, indicating that the fits were excellent. Values of log d_{sb} for the colour dimension ranged from 0.47 to 4.62, and log d_{sb} for the flash-frequency dimension ranged from 0.00 to 5.32. Values of log d_{br} ranged from 0.89 to 1.99. For all pigeons, log d_{sb} was much higher for one dimension (colour for Pigeons 71, 73, and 76; flash-frequency for Pigeons 72, 74, and 75), consistent with the higher log d_x values on that dimension (Figure 4.3). Therefore, the degree of reinforcer generalisation appeared to be larger for one dimension than for the other dimension for all pigeons.

Table 4.2.

Best-fitting values for log d_{sb} (colour), log d_{sb} (flash-frequency), and log d_{br}, and the proportion of data variance accounted for (VAC), for individual pigeons.

	Pigeon					
	71	72	73	74	75	76
log d_{sb} (colour)	4.62	1.31	1.80	0.66	0.47	1.51
log d_{sb} (flash-freq)	0.35	4.03	0.00	2.91	5.32	0.11
log d_{br}	1.73	0.89	1.99	0.92	1.46	1.46
VAC	.90	.98	.99	.99	.99	.99

4.4 Discussion

The present experiment investigated whether relative reinforcer rates determine the division of control between two dimensions of a compound stimulus when those dimensions are presented on two separate response keys, in order to understand better the results of Experiment 1B (Chapter 3). The present results differ from those of Experiment 1B. Here, discrimination (i.e., log d_x values) on a stimulus dimension generally increased as the probability of reinforcer deliveries for correct responses on that dimension increased, while discrimination on the other dimension decreased (Figure 4.3). Additionally, choice between the colour and flash-frequency elements during sample-stimulus presentations varied unsystematically (Figure 4.4). In contrast, in Experiment 1B, relative reinforcer rates had less systematic effect on the degree of discriminative control by a stimulus dimension (see Figure

3.2), whereas they strongly determined choice between the sample-stimulus elements (see Figure 3.3).

Before continuing, it is important to address the apparent carryover effects from the pigeons' previous experience in Experiment 1. For each pigeon in the present experiment, one stimulus dimension exerted strong control – in some cases, near-selective control – over behaviour (Figure 4.3). The dimension exerting stronger control was always the dimension for which the element-comparison associations remained unchanged from Experiment 1, and hence the dimension exerting little to no control was the dimension for which the element-comparisons associations were reversed. Therefore, the unequal division of control between the colour and flash-frequency dimensions in the present experiment was related to the pigeons' experience with the same dimensions in Experiment 1. Indeed, previous research has shown that when previously trained colour and line-orientation stimuli are compounded together and the contingency associated with one dimension is reversed, that dimension exerts little to no control over behaviour, whereas the dimension for which the contingency remains unchanged exerts near-selective control (Ray, 1969; see also Huguenin, 1987; Huguenin & Touchette, 1980; Leith & Maki, 1977; Ryan, Hemmes, & Brown, 2011).

Although the carryover effects from Experiment 1 somewhat confound the present results, the pigeons behaved differently in the present experiment compared with Experiment 1B, despite the identical stimulus configuration. Therefore, although the pigeons' previous experience in Experiment 1 was probably responsible for the near-exclusive control by the unchanged stimulus dimension (Figure 4.3) and for the deviation in obtained reinforcers from arranged reinforcers (Table 4.1), behaviour was also controlled by the contingencies arranged in the present experiment. Specifically, the change in training procedure – from separately training each stimulus element in Experiment 1B to training the elements in compound in the present experiment – eliminated the strong preference for one sample-stimulus element over the other and the control by element location (see Figures 3.2 to 3.4), and resulted in the development of discriminative control by the sample-stimulus elements (Figures 4.3 and 4.4).

Thus, despite carryover effects related to the element-comparison associations, the present results suggest that the type of training procedure was responsible for the results of Experiment 1B. Because the training procedure facilitated the development of orienting behaviour in Experiment 1B and the pigeons oriented towards the element associated with the higher reinforcer rate in test trials, behaviour was jointly controlled by the identity and location of the sample-stimulus elements in test trials (Figures 3.3 and 3.4, and Tables 3.3 and 3.4). In contrast, in the present experiment, the training procedure did not support the development of orienting behaviour towards a specific location (whether that be towards the element associated with the higher reinforcer rate, or towards one side of the operant panel), because the location of the correct comparison was unpredictable. Thus, in the present experiment, comparison choice was controlled by the identities of the sample-stimulus elements, rather than by their locations.

In addition to revealing the likely cause of the results of Experiment 1B, the present experiment also provides some insight into the conditions under which relative reinforcer rates determine divided stimulus control. Despite the strongly selective stimulus control, the present findings replicate the general relation between relative reinforcers and divided stimulus control obtained in past research (e.g., Davison & Elliffe, 2010; Shahan & Podlesnik, 2006, 2007). These findings suggest that the relation between relative reinforcers and divided stimulus control may hold across substantially different stimulus configurations, and when stimulus control is highly selective overall.

However, relative reinforcer rates appeared to have much smaller effects on divided stimulus control in the present experiment than in previous research (compare Figure 4.3 with e.g., Figure 2 in Davison & Elliffe, 2010). This difference may be related to the near-selective stimulus control in the present experiment. In order for relative reinforcer rates to shift stimulus control, subjects must discriminate the rate of differential reinforcers on each stimulus dimension (Davison, 2018a; Davison & Elliffe, 2010; Davison & Nevin, 1999). As stimulus control becomes more selective, subjects choose the comparison matching one dimension more frequently and the comparison matching the other dimension less frequently. As a result, subjects are less likely to experience the arranged contingencies on the latter dimension, and so changes in relative reinforcer rates will likely have smaller effects on divided stimulus control.

Additionally, as stimulus control becomes more selective, the frequency of differential reinforcers with respect to the controlling dimension increases, while the frequency of nondifferential reinforcers associated with the other dimension also increases. For example, consider an extreme case, in which only the colour dimension controls behaviour. In this case, subjects choose the comparison that matches the colour of the compound stimulus in all trials, and hence all obtained reinforcers will follow correct colour responses. Those same reinforcers are nondifferential with respect to the flash-frequency dimension, because half of them follow correct responses (in trials in which the elements match the same comparison) and the other half follow error responses (in trials in which the elements match different comparisons). Such a high frequency of nondifferential reinforcers on one dimension would likely attenuate any effects of relative reinforcer rates on divided stimulus control, and would probably also maintain the strong selective control by the other dimension. The results of Experiment 1B lend support to this explanation; in that experiment, no reinforcers for errors occurred because each stimulus element was trained individually, and relative reinforcer rates appeared to have larger effects on divided stimulus control (as expressed in sample choice; see Figure 3.3).

Indeed, the strong control by the unchanged dimension and weak control by the changed dimension in the present experiment may have been maintained by the frequencies of *apparent* (i.e., discriminated) differential and nondifferential reinforcers on each dimension. For the unchanged dimension, any reinforcers obtained were consistent with the previously learned discrimination, and hence obtained reinforcers may have appeared differential according to that dimension. In contrast, for the changed dimension, those same reinforcers were essentially error reinforcers with respect to the previously learned discrimination, and hence may have appeared nondifferential. The selective control by the unchanged dimension may been maintained by such high frequencies of apparent differential and nondifferential reinforcers associated with the unchanged and changed dimensions, respectively (see Davison & McCarthy, 1980; Davison & Nevin, 1999; Nevin et al., 1975). In support of this, Ryan et al. (2011) found that when the contingencies associated with one dimension of a compound stimulus were reversed, control by the changed dimension was stronger when some reinforcers had been obtained for choosing the newly correct response during initial training, compared with when the newly correct response was never reinforced during initial training. That is, the extent of control by the changed dimension was stronger when there was a history of reinforcement for choosing the newly correct response.

In further support of an explanation of our results based on reinforcers for errors, the present findings are similar to those from Davison and Elliffe's (2010) conditions in which only correct responses on one dimension were reinforced (i.e., differential reinforcers were only arranged on one dimension, and hence those reinforcers were completely nondifferential with respect to the other dimension). Also, like Davison and Elliffe, Davison and Nevin's (1999) model described the present data well (Table 4.2). Values of log d_{sb} were much higher for the dimension exerting stronger control over behaviour (i.e., the unchanged dimension) than for the other dimension, reflecting the near-exclusive control by the unchanged stimulus dimension (Table 4.2 and Figure 4.3). Thus, there was less reinforcer generalisation (i.e., a higher frequency of differential reinforcers) for the unchanged dimension than for the reversed dimension. Values of log d_{br} reflected generally good discrimination between the left and right response keys for all pigeons. Therefore, consistent with Davison and Elliffe (see also Davison, 2018a), the present findings suggest that divided control between the dimensions of compound sample stimuli in DMTS depends on both the frequency of differential and nondifferential reinforcers obtained on each dimension, and on reinforcer generalisation between the stimuli along a dimension and between responses.

Taken together, the present results and those of Experiment 1 are relevant to the applied behaviour-analytic literature on stimulus overselectivity (i.e., selective stimulus control). One intervention that may reduce stimulus overselectivity and encourage divided stimulus control is to *revalue* the overselected stimulus by reducing the reinforcer rate associated with that stimulus in separate training trials (see e.g., Broomfield, McHugh, & Reed, 2008, 2010; Kelly, Leader, & Reed, 2015; Reed, Broomfield, McHugh, McCausland, & Leader, 2009; Reed, Reynolds, & Fermandel, 2012). Experiment 1B provided further support that separately training stimulus-response-reinforcer relations can shift divided stimulus control (see Figures 3.2 to 3.4). However, in natural environments, it may not always be possible to separate stimuli that strongly control behaviour from those that do not, as not all compound stimuli are comprised of separable elements. The present results suggest that under such conditions, changing the reinforcer rate associated with the overselected stimulus may have small effects on overselectivity.

The present results also highlight how differential and nondifferential reinforcers may maintain and weaken, respectively, control by different stimulus dimensions. In the presence of multiple stimuli, differential reinforcers arranged for one conditional discrimination may appear nondifferential with respect to other discriminations, even if they are not explicitly arranged to be nondifferential. As a result, those stimuli that appear to be differentially correlated with reinforcers will exert strong control over behaviour, whereas other stimuli will exert little to no control. Due to the dynamical relation between behaviour and contingent consequences, obtained reinforcers may then further serve to maintain the strong control by some stimuli and weak control by other stimuli. Thus, the relative frequencies of apparent differential and nondifferential reinforcers may help to explain why some stimuli exert strong control over behaviour while other stimuli fail to gain control. This may be especially true when the contingencies associated with some stimuli change and the contingencies associated with other stimuli remain the same, as obtained reinforcers may continue to be differential with respect to the unchanged contingencies, but may appear nondifferential with respect to the changed contingencies.

4.5 Chapter Summary and Conclusions

The present experiment asked whether relative reinforcer rates determine the division of control between the dimensions of compound stimuli when those dimensions are spatially separated. This question is similar to that asked in Experiment 1 of the present thesis (Chapter 3), except that the stimulus elements from each dimension were trained separately and then presented together in Experiment 1. The present experiment also shares similarities with previous research on divided stimulus control (Davison, 2018a; Davison & Elliffe, 2010; Podlesnik, Thrailkill, & Shahan, 2012; Shahan & Podlesnik, 2006, 2007), except that the dimensions in question were compounded together into a single stimulus in previous research.

What do the present results suggest about the variables that determine divided stimulus control? In conjunction with the findings of Experiment 1 and of past research, the present results help to elucidate some of the stimulus characteristics and training conditions under which relative reinforcers do and do not determine divided stimulus control. Experiment 1 suggested that when stimulus-response-reinforcer relations are trained separately, the configuration of the stimuli determines whether or not a systematic relation between relative reinforcers and divided stimulus control is observed; in that experiment, relative reinforcers divided stimulus control when the stimuli were spatially separated, but not when they were compounded together. In contrast, past research and the present experiment suggest that when stimuli are trained together, relative reinforcer rates divide stimulus control both when stimuli are compounded together and when they are spatially separated. Additionally, the type of training procedure may affect the behavioural expression of stimulus control. In Experiment 1, the training procedure facilitated the development of orienting behaviour during sample-stimulus presentations, and hence the relation between relative reinforcer rates and divided stimulus control was evident in sample, not comparison, choice. In contrast, in the present experiment, the sample-stimulus elements exerted discriminative control over comparison choice.

The present results also suggest that previous experience contributes to the division of control between stimulus dimensions; when the contingency associated with one stimulus dimension is reversed but the contingency associated with the other dimension is left unchanged, control by the unchanged dimension exceeds control by the reversed dimension (see also Huguenin, 1987; Huguenin & Touchette, 1980; Leith & Maki, 1977; Ray, 1969; Ryan et al., 2011). The present findings further suggest that such selective stimulus control can attenuate the effects of relative reinforcers on divided stimulus control. This attenuation may be related to the frequency of differential and nondifferential reinforcers on each stimulus dimension, as Davison and Nevin’s (1999) model described our findings of near-selective stimulus control well. Thus, divided stimulus control depends on differential reinforcers, nondifferential reinforcers, and reinforcer generalisation (see also Davison, 2018a; Davison and Elliffe, 2010).

Overall, Experiments 1 and 2 of the present thesis (Chapters 3 and 4) highlight some important considerations – namely, the stimulus configuration, the type of training procedure, additional sources of information (e.g., the location of the sample-stimulus elements), the effects of changes in contingencies, and the frequencies of differential and nondifferential reinforcers – for both basic and applied researchers of divided stimulus control. These factors may determine whether stimulus control is selective or divided, the behavioural expression of such stimulus control, and the success of procedures that aim to shift control from one stimulus to another.

CHAPTER 5

Experiment 3: Relative Reinforcer Probabilities Divide Stimulus Control in the Multiple Peak Procedure[10]

5.1 Introduction

Several studies, including Experiments 1 and 2 of this thesis (Chapters 3 and 4), suggest that relative reinforcer probabilities determine the division of control between the dimensions of compound stimuli when those dimensions signal the *location* (left or right) of future reinforcers in a delayed matching-to-sample (DMTS) task (e.g., Davison, 2018a; Davison & Elliffe, 2010; Podlesnik, Thrailkill, & Shahan, 2012; Shahan & Podlesnik, 2006, 2007). Another situation in which relative reinforcer probabilities might determine divided stimulus control is when stimuli signal the *time* of future reinforcers. Presently, little research has examined divided stimulus control in time, and hence the variables that determine divided control in time are not well understood (Davison, 2018b). Thus, the present experiment asked whether relative reinforcer probabilities[11] determine the division of control between compound-stimulus dimensions when those dimensions signal the time of the next reinforcer.

The peak procedure provides a convenient method to study control by stimuli that signal the time of the next reinforcer. In the peak procedure (Catania, 1970; S. Roberts, 1981), each trial begins with the onset of a stimulus signalling the beginning of a to-be-timed interval, after which a response produces a reinforcer delivery and ends the trial. In these

---

[10] This experiment was conducted under Approval 001967 granted by the University of Auckland Animal Ethics Committee.
This chapter is a lightly edited version of the paper 'Timing compound stimuli: Relative reinforcer probabilities divide stimulus control in the multiple peak procedure', submitted for publication to *Journal of Experimental Psychology: Animal Learning and Cognition.*

[11] The word "probability" is used in this chapter, rather than "rate" as in Chapters 3 and 4, because probabilities and rates are not equivalent when intervals to reinforcer delivery differ.

‘fixed interval’ (FI) trials, response rates are temporally controlled; they increase gradually during the to-be-timed interval (Ferster & Skinner, 1957). To assess temporal discrimination more thoroughly, ‘peak’ trials, which are longer in duration and terminate without reinforcer delivery, are interspersed among FI trials. Typically, response rates in peak trials increase gradually across time to reach a maximum at around the usual time of reinforcer delivery, and then decrease gradually thereafter. The time at which response rates reach a maximum (*peak time*) and the width of the response-rate function (*peak spread*) in peak trials provide measures of timing accuracy and precision, respectively. The closer the peak time to the arranged reinforcer time, the more accurate the temporal discrimination, and the smaller the peak spread (i.e., the sharper the peak), the more precise the temporal discrimination.

Stimulus control over timing may be assessed in the peak procedure by varying the characteristics of the time-marker stimulus that signals the beginning of the to-be-timed interval (see e.g., Fox & Kyonka, 2016). Thus, to examine divided stimulus control in time, a *compound* stimulus may be used as a time marker in the peak procedure. Few studies have investigated temporal discrimination with compound time-marker stimuli. Nevertheless, a series of recent experiments by Matell and his colleagues provides some insight into divided stimulus control in the peak procedure (De Corte & Matell, 2016b; Kurti, Swanton, & Matell, 2013; Matell, De Corte, Kerrigan, & DeLussey, 2016; Matell & Kurti, 2014; Swanton, Gooch, & Matell, 2009; Swanton & Matell, 2011; see also Delamater & Nicolas, 2015). Matell and colleagues arranged a multiple peak procedure in which tone and light stimuli signalled different FI schedules, and both stimuli were occasionally presented simultaneously in compound peak trials. Response rates in compound peak trials followed a similar pattern to response rates in single-stimulus (i.e., just tone or light) peak trials, except that the peak time occurred between the FI durations associated with the tone and light stimuli. That is, their subjects (rats) appeared to “average” the FI durations signalled by the tone and light stimuli together in compound peak trials. Such *temporal averaging* implies that stimulus control was divided between the tone and light.

Temporal averaging appears to depend on the probability of reinforcer deliveries associated with each stimulus. Matell and Kurti (2014) varied the probability of reinforcer deliveries associated with the tone and light stimuli across groups of rats, and found that response-rate functions in compound peak trials generally appeared more similar to response-rate functions in peak trials in which only the stimulus associated with the higher reinforcer probability was presented. That is, during the temporal-averaging process, each element appeared to be weighted by its relative reinforcer probability. In a similar study, Delamater and Nicolas (2015) showed that when the tone and light stimuli were associated with different types of reinforcers, pre-session satiation with one reinforcer shifted peak response functions away from the interval signalled by the stimulus associated with that reinforcer and towards the interval signalled by the other stimulus. Taken together, Matell and Kurti’s and Delamater and Nicolas’ findings suggest that when multiple stimuli signal different times to reinforcer delivery, stimuli associated with a higher reinforcer probability or with a more ‘valuable’ reinforcer exert stronger control over behaviour. Therefore, these findings parallel those of studies demonstrating a systematic relation between relative reinforcer probabilities and divided control between stimuli that signal the location of reinforcers (e.g., Davison & Elliffe, 2010; Shahan & Podlesnik, 2006, 2007).

To date, only Matell and Kurti (2014) have investigated how relative reinforcer probabilities affect divided stimulus control in the temporal domain. Hence, it is unclear whether their results are specific to the procedure that they arranged – in which the stimulus-interval-reinforcer relations were trained separately and then stimuli were presented simultaneously – or whether temporal averaging is a more general timing strategy that subjects use when faced with a compound stimulus comprised of elements that signal discrepant times to reinforcer delivery. Indeed, temporal averaging is not the only timing strategy consistent with divided stimulus control; an alternative strategy is to time *both* durations signalled by the stimulus elements. The results of Experiments 1 and 2 of this thesis (Chapters 3 and 4) suggested that the behavioural expression of divided stimulus control may differ depending on whether stimuli are trained separately or together in DMTS (see also Du, McMillan, Madan, Spetch, & Mou, 2017). The same may be true in the peak procedure; that is, temporal averaging may not occur if the stimuli are trained in compound, rather than separately.

Furthermore, Matell and Kurti (2014) investigated the effects of relative reinforcer probabilities on temporal averaging across groups, rather than within subjects. Therefore, it is unclear how response-rate patterns in compound peak trials would have varied if relative reinforcer probabilities had been varied within, rather than between, subjects. Within-subjects replications enhance the reliability and generality of the effect(s) under investigation (Perone, 2019). Hence, stronger evidence for a relation between relative reinforcer probabilities and divided stimulus control in temporal discriminations would be obtained from an experiment in which all subjects experience all of the conditions.

The aim of the present experiment was twofold. First, we investigated whether relative reinforcer probabilities divide control between stimuli that signal the time of the next reinforcer in the same way as they divide control between stimuli that signal the location of the next reinforcer. Second, we aimed to provide further insight into timing strategies in the presence of compound stimuli whose elements signal different times to reinforcer delivery. To that end, the present experiment arranged a multiple peak procedure that was an analogue of the DMTS procedure arranged in Davison and Elliffe (2010). In each trial, a compound stimulus comprising elements that signalled either an FI 2-s or FI 8-s schedule was presented. The compound stimulus was a white keylight that alternated with either red or green at a fast or slow frequency (hence, the two stimulus dimensions were colour, Red-Green, and flash-frequency, Fast-Slow). One element from each dimension signalled an FI 2-s schedule, and the other elements signalled an FI 8-s schedule. The probability of reinforcer deliveries associated with each dimension varied across conditions. Interspersed among these trials were peak trials, which allowed us to assess how variations in relative reinforcer probabilities affected temporal discrimination of the intervals signalled by the compound-stimulus elements.

5.2 Method

5.2.1 Subjects and Apparatus

The subjects and apparatus were the same as in Experiments 1 and 2 (Chapters 3 and 4), except that only the centre key was used. The centre key could be illuminated red, green, or white.

5.2.2 Procedure

Sessions were run daily, beginning at 1 a.m. The pigeons were run successively and in numerical order. Each session lasted until 160 trials or 65 min had elapsed, whichever occurred first.

A multiple peak procedure was arranged, in which colour-alternating stimuli signalled an FI 2-s or 8-s schedule. Each stimulus alternated between red and white or green and white every 0.1 s or 0.5 s. Thus, each stimulus was a compound consisting of a colour (red or green) and a flash-frequency (fast or slow), and there were four stimuli in total: red-fast, red-slow, green-fast, and green-slow. Within each stimulus dimension, one of the stimuli signalled the FI 2-s schedule, and the other stimulus signalled the FI 8-s schedule. Table 5.1 shows the stimulus-interval associations for each pigeon.

Table 5.1.

Stimulus-interval associations and the order of conditions for individual pigeons.

Pigeon	FI 2 s	FI 8 s	Order of conditions (probability of food for colour)
71	Red, Slow	Green, Fast	.2, .8, .5, .9, .1
72	Green, Slow	Red, Fast	.2, .8, .5, .9, .1
73	Green, Slow	Red, Fast	.2, .8, .5, .9, .1
74	Green, Fast	Red, Slow	.8, .2, .5, .1, .9
75	Red, Fast	Green, Slow	.8, .2, .5, .1, .9
76	Red, Fast	Green, Slow	.8, .2, .5, .1, .9

Note: The probability of food for the flash-frequency dimension was always the complement of the probability of food for the colour dimension.

To ensure that trials of each type were distributed close to equally across experimental sessions, each session was divided into two halves consisting of 80 trials each. Before each trial, a stimulus was selected randomly with $p = .25$, with the constraints that each stimulus was selected an equal number of times and a stimulus could not be selected more than four times consecutively. Thereafter, one of the stimulus dimensions (colour or flash-frequency) was selected randomly ($p = .5$), with the constraint that each dimension was selected an equal number of times per compound stimulus. Thus, in each half of each session, there were 20 trials per compound stimulus, and of these 20 trials, 10 were trials in which the colour dimension signalled the FI schedule and the other 10 were trials in which the flash-frequency dimension signalled the FI schedule. For example, consider Pigeon 71, for whom the red and slow stimuli signalled the FI 2-s schedule and the green and fast stimuli signalled the FI 8-s schedule (see Table 5.1). In each session half for this pigeon, there were 20 trials in which the red-slow stimulus signalled the FI 2-s schedule and 20 trials in which the green-fast stimulus signalled the FI 8-s schedule (because both dimensions signalled the same FI for these two trial types). For each of the red-fast and green-slow stimuli, there were 10 trials in which the stimulus signalled the FI 2-s schedule and 10 trials in which it signalled the FI 8-s schedule. The program selected one of these trial types probabilistically at the beginning of each trial.

The compound stimulus was presented on the centre response key, and remained on for the duration of the trial. After the onset of the compound stimulus, the first peck to the centre key that occurred after the FI duration had elapsed resulted in the offset of the stimulus and either a 2-s food delivery or a 2-s blackout depending on whether the program had arranged a reinforcer. All trials were followed by a variable-time 5-s inter-trial interval, during which the centre key remained unlit.

5.2.2.1 Pretraining

The goal of pretraining was to reduce the percentage of reinforced trials to 50%. The pigeons were first trained on the above procedure with 75% of trials reinforced, split evenly between the four compound stimuli. After 10 sessions, the percentage of reinforced trials was reduced to 60% for five sessions, and then to 50% for five sessions. Thereafter, the experiment proper began.

5.2.2.2 Experimental Sessions

Experimental sessions were similar to pretraining sessions, except that 50% of all trials were reinforced and four trials per stimulus (two in each session half) were replaced with unreinforced peak trials. Peak trials lasted for 24 s, terminated independently of responding, and were interspersed randomly among FI trials.

Across conditions, the probability of reinforcer delivery at the end of the interval signalled by the colour dimension varied (the probability of reinforcer delivery at the end of the interval signalled by the flash-frequency dimension was the complement of this probability). The probabilities of reinforcer deliveries associated with the colour dimension were .1, .2, .5, .8, or .9 (hereafter, for brevity, *p*[food] for colour). Thus, for example, when *p*(food) for colour was .2, this meant that for each of the four compound stimuli in each half of the session, two of the 10 trials in which the colour dimension signalled the FI schedule ended in reinforcer delivery and the remaining eight trials ended in a blackout, whereas the reverse was true (i.e., eight trials reinforced) for the 10 trials in which the flash-frequency dimension signalled the FI schedule. The order of conditions was counterbalanced across pigeons; the rightmost column of Table 5.1 shows the order of conditions for each pigeon in terms of the reinforcer probability associated with the colour dimension. Each condition lasted for 30 sessions. Response patterns were stable, determined by visual inspection, by the 15th session of each condition. Hence, data from the last 12 sessions of each condition were used for all analyses.

5.2.3 Data Analysis

To assess control by the colour and flash-frequency stimulus dimensions, we analysed patterns of responding across time in peak trials, peak times and spreads, and the times at which significant changes in the rate of responding (from low to high or vice versa) occurred. Because the stimulus-interval associations were counterbalanced across pigeons (Table 5.1), data were separated according to whether the compound-stimulus elements signalled the same or different FI schedules for all analyses. Hereafter, trials in which both compound-stimulus elements signalled the FI 2-s schedule or the FI 8-s schedule are termed ‘Both FI 2’ and ‘Both FI 8’ trials respectively. Trials in which the colour element signalled the FI 2-s schedule and the flash-frequency element signalled the FI 8-s schedule or vice versa are termed ‘Colour FI 2’ and ‘Colour FI 8’ trials, respectively.

5.2.3.1 Peak Response Functions

Responses in peak trials were aggregated according to the compound stimulus and separated into 0.5-s time bins, and were used to calculate responses per min in each bin by dividing the number of responses made in a bin by the number of times that bin was reached and then multiplying by 120. Visual inspection of the resulting peak response functions suggested that they were either unimodal or bimodal. Unimodal peak response functions are well described by a single ramped Gaussian function (Cheng & Westwood, 1993; Subramaniam & Kyonka, 2019), and bimodal functions are well described by the sum of two Gaussian functions and a ramp (Subramaniam & Kyonka, 2019; Whitaker, Lowe, & Wearden, 2008). Thus, to obtain measures of peak times and spreads, we fit a single ramped Gaussian (Equation 5.1) and the sum of two Gaussian functions and a ramp (Equation 5.2) to peak-trial data. In Equations 5.1 and 5.2, *a* is a scaling constant representing the peak height, *M* is the time at which response rates reach a maximum (i.e., the peak time), *SD* is the width of the response-rate function (i.e., the peak spread), and *s* and *b* are the slope and y-intercept, respectively, of the ramp.

$$f(t)_{single} = a \cdot e^{\left(-\frac{(t-M)^2}{2SD^2}\right)} + (s \cdot t + b) \qquad \text{(Equation 5.1)}$$

$$f(t)_{summed} = a_1 \cdot e^{\left(-\frac{(t-M_1)^2}{2{SD_1}^2}\right)} + a_2 \cdot e^{\left(-\frac{(t-M_2)^2}{2{SD_2}^2}\right)} + (s \cdot t + b) \qquad \text{(Equation 5.2)}$$

Equations 5.1 and 5.2 were fit to obtained peak-trial data using nonlinear regression by maximum likelihood in Microsoft® Excel Solver. The number of responses made in a time bin must be an integer greater than or equal to 0, hence, we assumed that the number of responses in each time bin was distributed according to a Poisson distribution for the curve-fitting procedure (see Gomes-Ng, Elliffe, & Cowie, 2018b).

Because Equation 5.1 is the reduced version of Equation 5.2, the proportion of data variance accounted for (VAC) for fits of Equation 5.2 was always equal to or greater than the VAC for fits of Equation 5.1. To determine whether the increase in VAC as a result of three additional parameters (a_2, M_2, and SD_2) was justified, the models were compared using Akaike's Information Criterion (Burnham & Anderson, 2002), corrected for small sample sizes (AICc). The VACs for some of the fitted functions were quite low, but visual inspection of the fits suggested that this was because there was considerable variability in response rates from bin to bin, and that the functions still provided an adequate indication of peak times and spreads. Appendix Figures A5.1 to A5.6 show obtained peak-trial data and the fits of Equations 5.1 and 5.2 for Pigeons 71 to 76, respectively, and the parameter estimates, VAC, and AICc for each of the model fits are provided in Appendix Tables B5.1 to B5.12. The AICc favoured the single model in 21 out of 30 cases (70%) for Both FI 2 trials, in 20 out of 30 cases (66.6%) for Both FI 8 trials, and in 13 out of 30 cases (43.3%) for both Colour FI 2 and Colour FI 8 trials.

There were two difficulties with assessing temporal discrimination using the peak times and spreads from the single or summed Gaussian model (Equations 5.1 and 5.2, respectively). First, longer intervals are generally timed less precisely (scalar property of timing; Gibbon, 1977), and so peak spreads are usually larger for longer intervals. Therefore, increases in peak spreads may occur either because the same interval is timed less precisely, or because a longer interval is timed. This makes it difficult to assess the effects of relative reinforcer probabilities on timing precision. To mitigate this issue, we calculated *coefficients of variation* (CVs) by dividing peak spreads by peak times. This normalised peak spreads relative to peak times, and hence provided a measure of timing precision relative to the interval being timed. Decreases and increases in the CV imply increases and decreases, respectively, in relative timing precision.

Second, because the AICc sometimes favoured the single model and other times the summed model, the number of peaks differed between trials, conditions, and pigeons. This made it difficult to assess changes in timing accuracy and precision of the 2-s and 8-s intervals, because peak times and spreads were only available for one interval when the AICc favoured the single model. This was especially true for Colour FI 2 and Colour FI 8 trials.

Therefore, for Colour FI 2 and Colour FI 8 trials, we chose to use the peak times and spreads from the summed Gaussian model (Equation 5.2) for all pigeons. For Both FI 2 and Both FI 8 trials, we used the parameters from the model favoured by the AICc. For the nine (Both FI 2) and 10 (Both FI 8) cases in which the AICc favoured the summed model, visual inspection of the fitted summed Gaussian functions (see Appendix Figures A5.1 to A5.6) indicated that the first peak time and spread was closer to the arranged reinforcer time (i.e., 2 or 8 s), and had a higher amplitude. Thus, the first peak probably reflected temporal discrimination of the time to reinforcer delivery, whereas the second peak time and spread described responding at later times (see e.g., Pigeon 75, Appendix Figure A5.5). Therefore, for cases where the AICc favoured the summed model for Both FI 2 and Both FI 8 trials, we used the first peak time and spread for analyses.

5.3.2.2 Changepoints

Response rates in individual peak trials typically transition from a low rate to a high rate at some point during each trial, and vice versa at a later time (Cheng & Westwood, 1993; Church, Meck, & Gibbon, 1994). The times at which such changepoints, termed *start* and *stop* times respectively, occur provide another measure of temporal discrimination. We obtained start and stop times using an algorithm described by Gallistel, Fairhurst, and Balsam (2004; see also Gallistel, King, Gottlieb, Balci, Papachristos, Szalecki, & Carbone, 2007; Taylor, Horvitz, & Balsam, 2007). Briefly, the algorithm steps through the cumulative record response by response, and for each response, computes the log of the odds that the response rate during the interval between origin of the record and the time of that response differs from the response rate between the time of that response and the end of the trial. If the log of the odds exceeds a specified decision criterion, a significant changepoint has been found. The algorithm then begins anew, now using this changepoint as the origin of the cumulative record. This continues until no more significant changepoints are found. For the present data, we used a binomial test to compare response rates before and after the putative changepoint, and an odds ratio of 20:1 ($p < .05$) as the decision criterion. Visual inspection of the cumulative records suggested that this criterion produced sensible changepoints. In most trials (generally, between 50 and 70%), the algorithm only detected one or two significant changepoints. Thus, we used the median first start time and median first stop time for analyses.

Trials with less than three responses were excluded from changepoint analyses. Additionally, because the changepoint algorithm only returned changepoints that exceeded the decision criterion, there were no stop times detected for Pigeon 74 in Both FI 8 trials when the reinforcer probability associated with the colour dimension was .8, and in Colour FI 2 trials when the reinforcer probability associated with the colour dimension was .9.

5.3 Results

5.3.1 Peak Response Functions

Figure 5.1 shows response rates across time in peak trials, averaged across all six pigeons, for each stimulus (columns) and *p*(food) for colour (*p*; rows). Appendix Figures A5.1 to A5.6 show peak response functions for individual pigeons. Although there was individual variability in temporal discrimination between pigeons, peak response functions followed the same general pattern across stimuli and conditions for all pigeons. Thus, the group-mean data in Figure 5.1 are representative of individual-pigeon data, and the patterns described here apply to both individual and group-mean data. As is evident from Figure 5.1, patterns of responding across time differed depending on the compound stimulus, hence, both the colour and flash-frequency dimensions controlled behaviour.

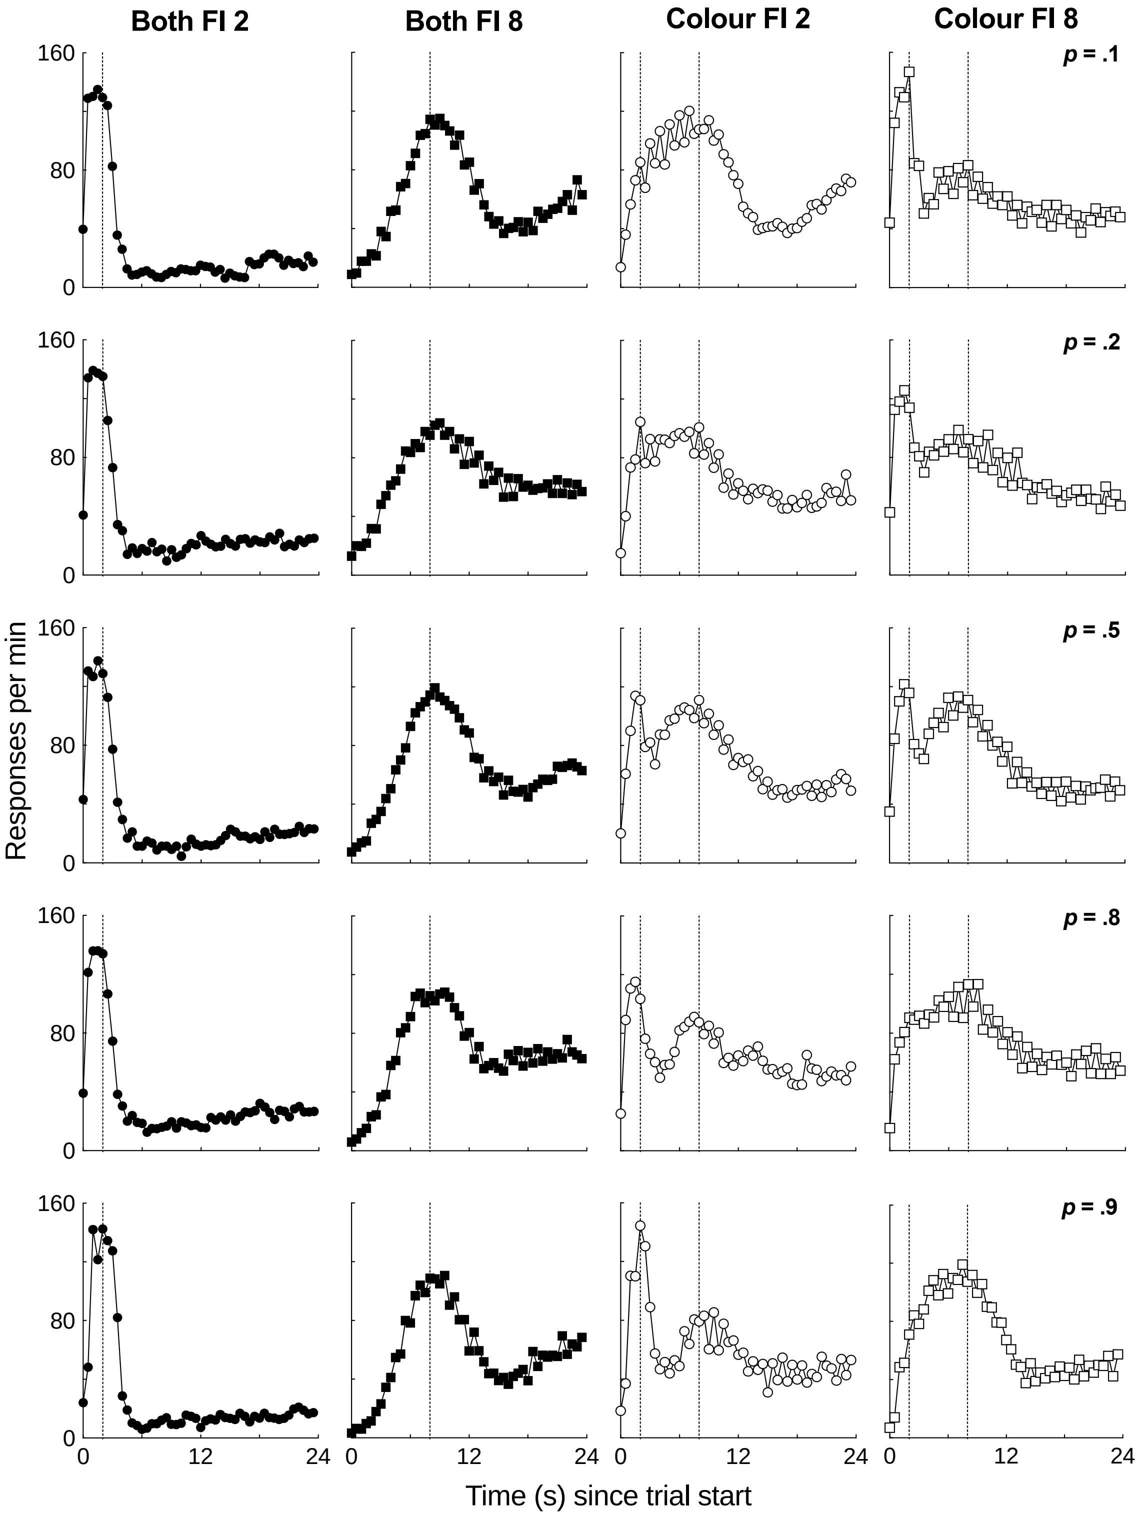

Figure 5.1. Mean (averaged across pigeons) response rates in peak trials. Data are separated according to the compound stimulus (columns) and the probability of reinforcer deliveries associated with colour stimuli (*p*; rows). The vertical lines denote the arranged FI durations (2 and 8 s).

5.3.1.1 Both FI 2 and Both FI 8 trials

When the colour and flash-frequency elements signalled the same FI schedule (Both FI 2 and Both FI 8 trials), response rates increased over time, reached a maximum at around the arranged reinforcer time, and then decreased thereafter (Figure 5.1, leftmost two columns). Response rates increased and decreased faster overall in Both FI 2 trials than in Both FI 8 trials. Additionally, patterns of responding at later times differed depending on whether the elements signalled an FI 2-s or 8-s schedule: In Both FI 2 trials, response rates remained low at later times, whereas in Both FI 8 trials, response rates were higher, and sometimes increased again after reaching a local minimum. These patterns of responding changed little as p(food) for colour increased.

Figure 5.2 shows peak times and normalised peak spreads (i.e., coefficients of variation: CVs) obtained from fits of a single (Equation 5.1) or summed (Equation 5.2) ramped Gaussian function to peak response functions for Both FI 2 and Both FI 8 trials. Peak times provide a measure of timing accuracy, and normalised peak spreads provide a measure of relative timing precision. Larger CVs indicate less precise temporal discrimination (i.e., larger peak spreads relative to peak time). For Both FI 2 and Both FI 8 trials, we used peak times and spreads from the model favoured by the AICc – either the single peak from the single model, or the first peak (which reflected temporal discrimination of the time to reinforcer delivery) from the summed model. Figure 5.2 also shows predicted response rates at the peak time (hereafter, *peak rates*), which were calculated using a Gaussian function (see Equation 5.1) with mean and standard deviation equal to the peak time and spread, respectively, from the curve-fitting procedure.

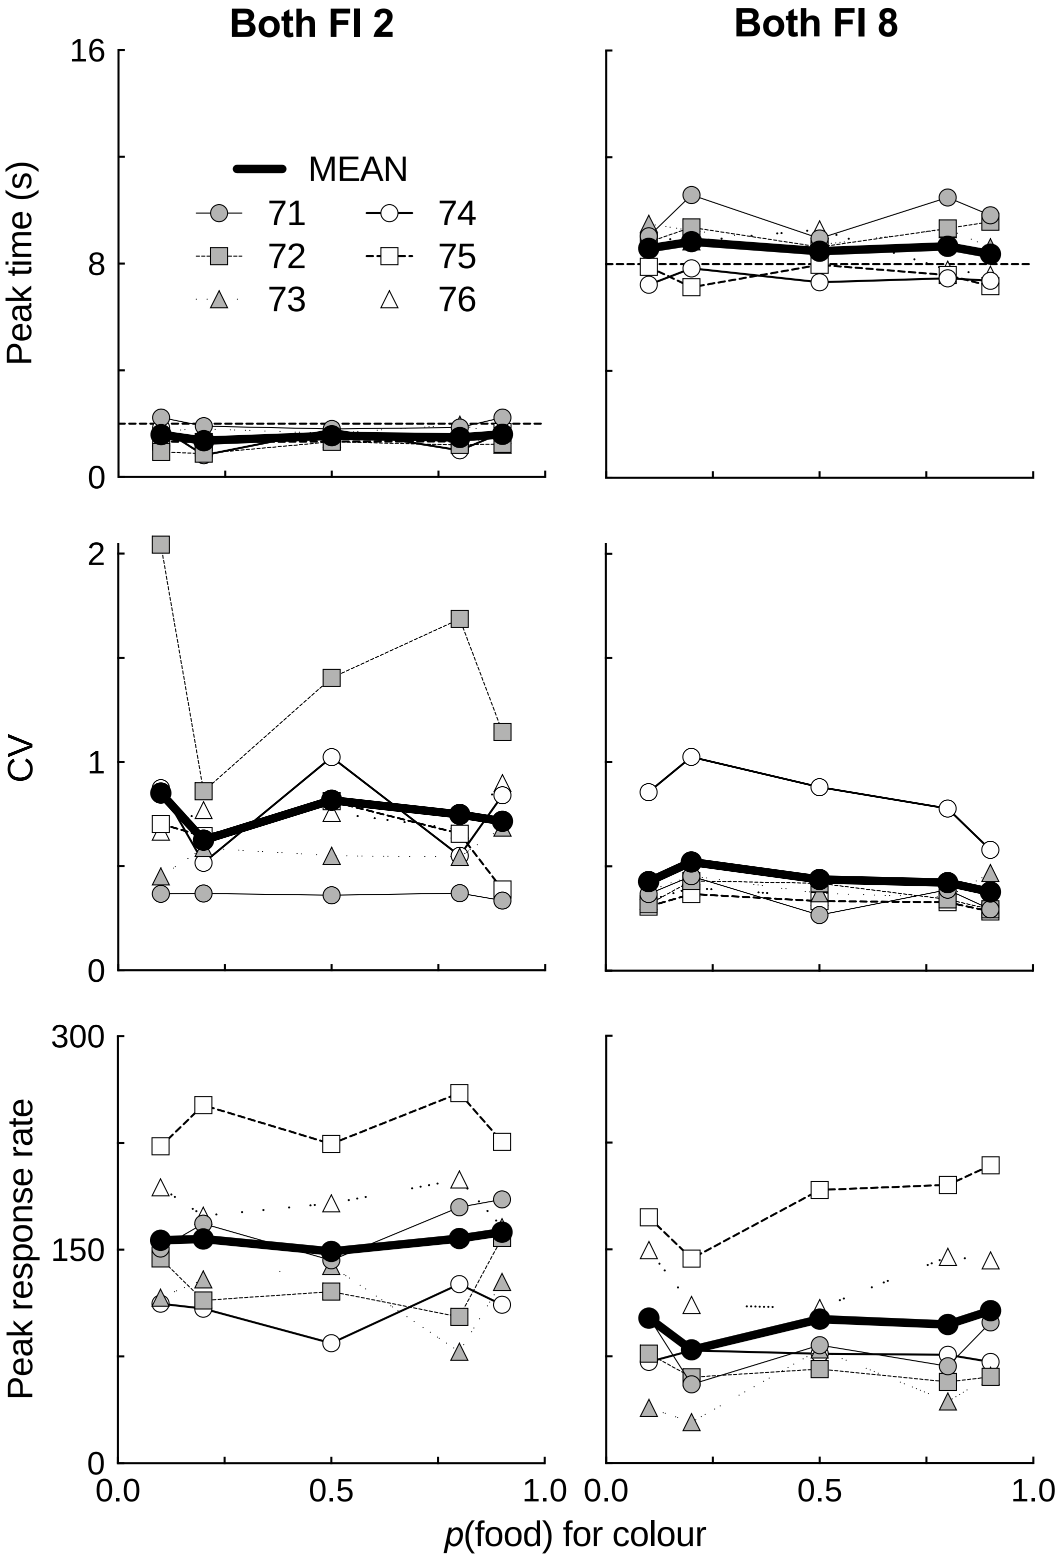

Figure 5.2. Peak times (top row), normalised peak spreads (coefficients of variation, CVs; middle row), and peak rates (bottom row) in Both FI 2 and Both FI 8 trials, plotted as a function of *p*(food) for colour. The white and grey symbols show individual-pigeon data, and the solid black symbols show group-mean data.

Consistent with the peak response functions in Figure 5.1, peak rates (Figure 5.2, bottom row) did not change systematically across conditions in Both FI 2 and Both FI 8 trials (two-tailed nonparametric trend tests [Elliffe & Elliffe, 2019; Kendall, 1955]: $\Sigma S = 8$ and 2 respectively, both $p > .4$). Peak times (Figure 5.2, top row) in Both FI 2 and Both FI 8 trials were generally close to the FI durations signalled by the compound-stimulus elements, and there was little change in peak times as *p*(food) for colour increased ($\Sigma S = 2$ and -8 for Both FI 2 and Both FI 8 trials, both $p > .4$). There was no systematic trend in CVs (Figure 5.2, middle row) for Both FI 2 trials ($\Sigma S = 0$, $p > .5$), but CVs appeared to decrease slightly in Both FI 8 trials, although this decrease was not statistically significant ($\Sigma S = -16$, $p > .05$). CVs differed significantly between Both FI 2 and Both FI 8 trials in some conditions; CVs were significantly larger in Both FI 2 trials when *p*(food) for colour was .1, .5, and .9 (binomial sign tests: all $p < .02$), but not when *p*(food) for colour was .2 or .8 ($p > .3$). The significant differences in CVs between Both FI 2 and Both FI 8 trials are not related to response rates at later times skewing the fits of Equations 5.1 or 5.2 in Both FI 2 trials, as both equations included a ramp to account for responding at such times. Instead, the differences in CVs suggest that, in general, the 2-s interval was timed somewhat less precisely relative to the 8-s interval. The reasons for this violation of scalar timing (Gibbon, 1977) are unclear, though it is worth noting that other researchers have found similar violations of scalar timing with short (< 2 s) fixed-interval durations (e.g., Grondin, 2012; Wearden & Lejeune, 2008).

5.3.1.2 Colour FI 2 and Colour FI 8 trials

The effects of relative reinforcer probabilities on divided stimulus control are evident in trials in which the colour and flash-frequency elements signalled different FI schedules (Colour FI 2 and Colour FI 8 trials). The rightmost two columns of Figure 5.1 show response rates across time in Colour FI 2 and Colour FI 8 peak trials, and Figures 5.3 and 5.4 show peak times, normalised peak spreads, and peak rates from fits of the summed Gaussian function (Equation 5.2) in Colour FI 2 and Colour FI 8 trials respectively. As p(food) for colour increased, a peak at 2 s became apparent in Colour FI 2 trials, and a peak at 8 s became apparent in Colour FI 8 trials. At the same time, peaks at around 8 s or 2 s became less apparent in those trials, respectively. More generally, as the relative reinforcer probability signalled by one stimulus dimension increased, a peak in responding near the reinforcer time signalled by that dimension developed, while a peak in responding near the reinforcer time signalled by the other dimension diminished. These patterns were reflected in peak rates (Figures 5.3 and 5.4, bottom rows); as p(food) for colour increased, peak rates for the 2-s interval increased in Colour FI 2 trials ($\Sigma S = 34$, $p < .001$) and decreased in Colour FI 8 trials ($\Sigma S = -36$, $p < .001$), whereas peak rates for the 8-s interval decreased in Colour FI 2 trials ($\Sigma S = -24$, $p < .02$) and increased in Colour FI 8 trials ($\Sigma S = 32$, $p < .001$).

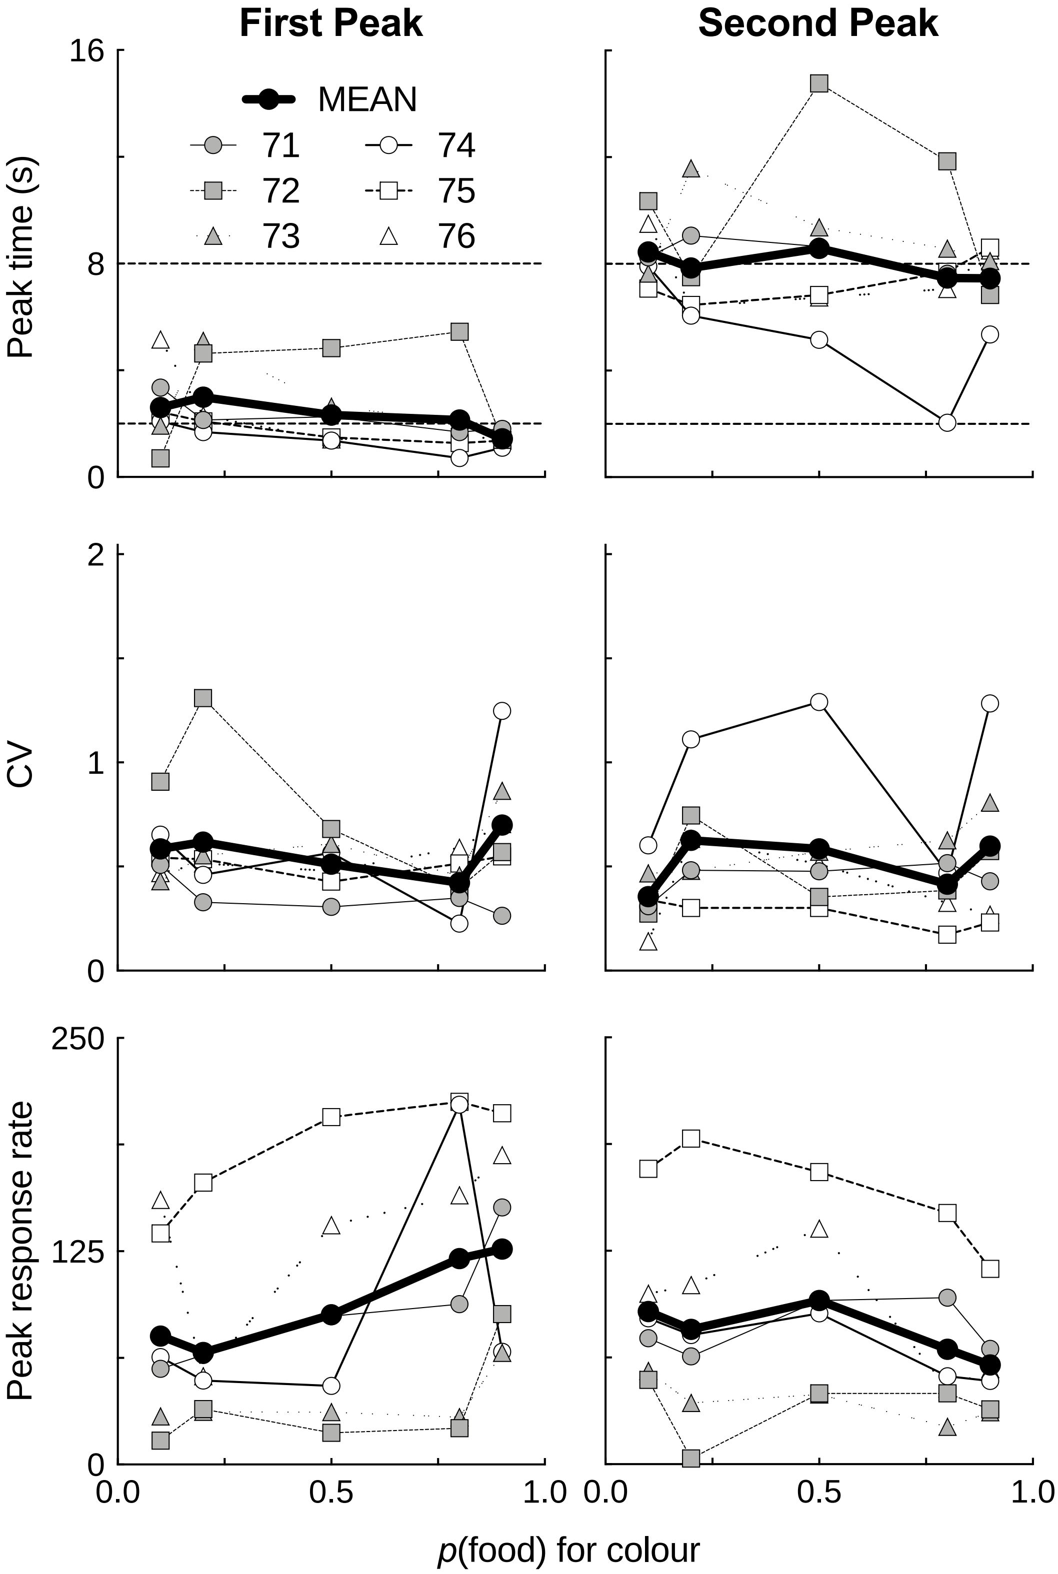

Figure 5.3. First- and second-peak times (top row), normalised peak spreads (coefficients of variation, CVs; middle row), and peak rates (bottom row) in Colour FI 2 trials, plotted as a function of *p*(food) for colour. The white and grey symbols show individual-pigeon data, and the solid black symbols show group-mean data.

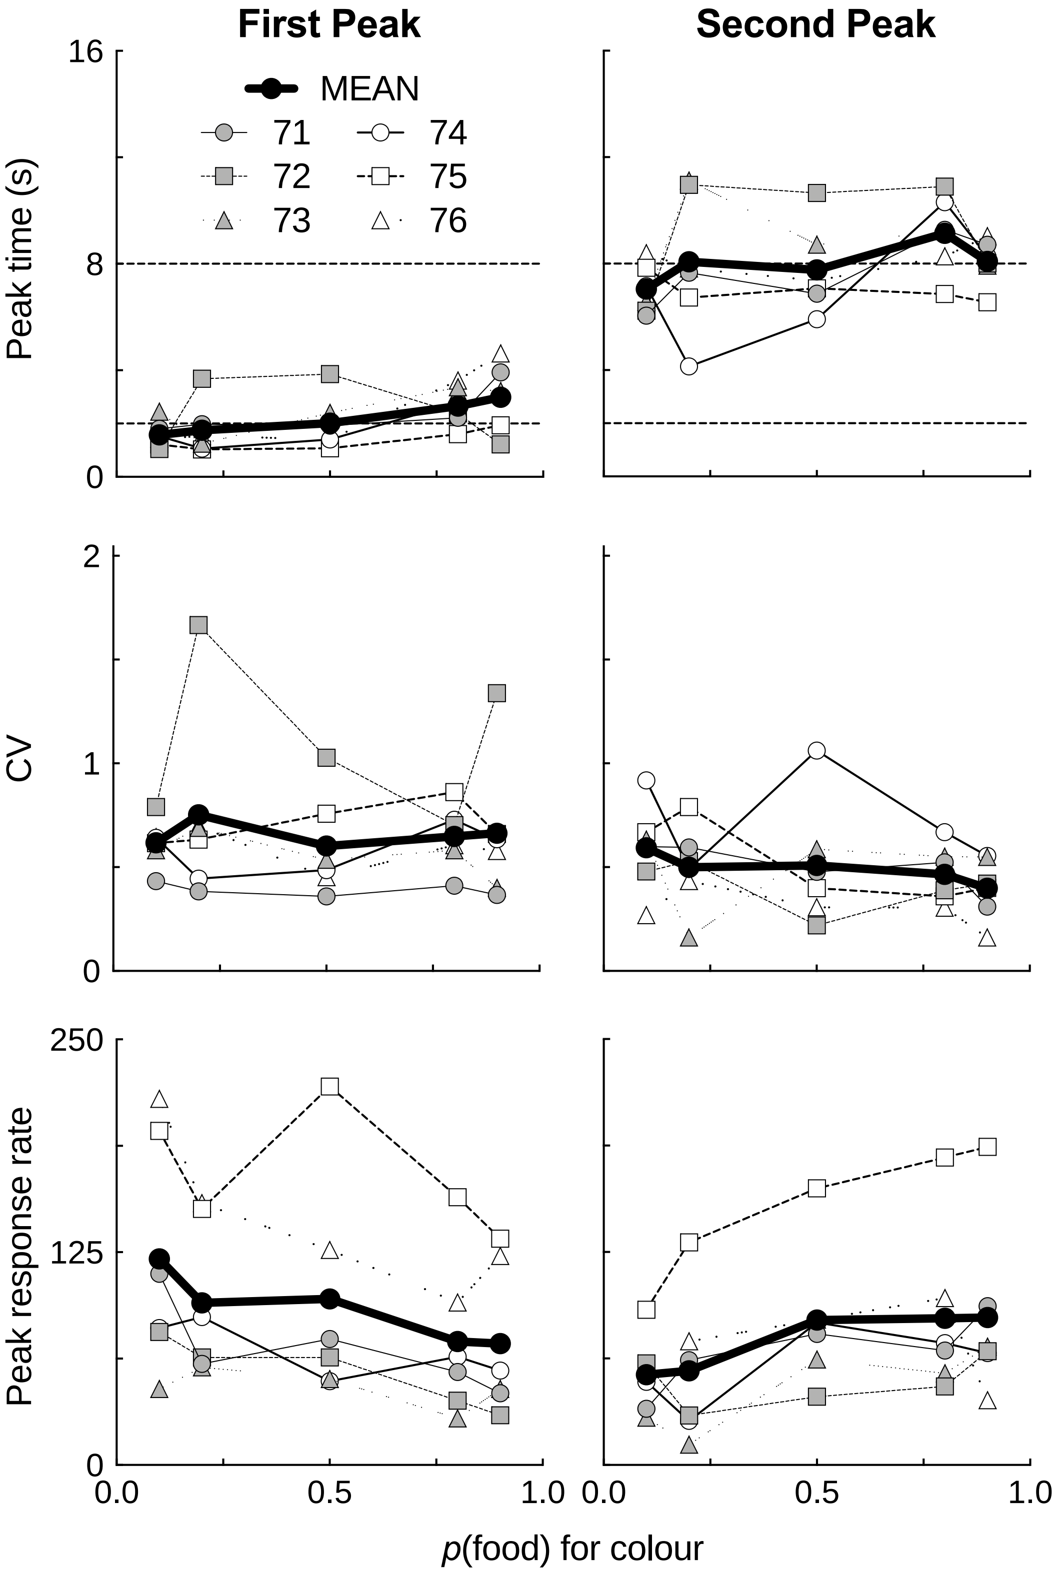

Figure 5.4. First- and second-peak times (top row), normalised peak spreads (coefficients of variation, CVs; middle row), and peak rates (bottom row) in Colour FI 8 trials, plotted as a function of *p*(food) for colour. The white and grey symbols show individual-pigeon data, and the solid black symbols show group-mean data.

Thus, the dimension associated with the relatively higher reinforcer probability appeared to exert stronger control over interval-timing behaviour in Colour FI 2 and Colour FI 8 trials. Nevertheless, some control by the other dimension was also evident in all conditions (Figure 5.1). At the beginning of peak trials, the slope of the increase in response rates was steeper in Colour FI 2 and Colour FI 8 trials than in Both FI 8 trials, even if a peak at 2 s was not apparent, suggesting that the element signalling the FI 2-s schedule exerted some control over responding. Likewise, response rates at later times in Colour FI 2 and Colour FI 8 trials were higher than response rates at later times in Both FI 2 trials, suggesting some control by the element signalling the FI 8-s schedule.

In addition to the effects of relative reinforcer probabilities on peak rates (Figures 5.1, 5.3, and 5.4), the accuracy and precision with which pigeons timed an interval may have also depended on the probability of reinforcer deliveries associated with the dimension signalling that interval. For example, when *p*(food) for colour was low, pigeons may have timed the interval signalled by the colour dimension less accurately and/or precisely. Indeed, unlike Both FI 2 and Both FI 8 trials (Figure 5.2), changes in *p*(food) for colour caused changes in peak times in Colour FI 2 and Colour FI 8 trials (Figures 5.3 and 5.4, top rows). In Colour FI 2 trials, first peak times decreased as *p*(food) for colour increased ($\Sigma S = -30$, $p < .002$), and became more similar to peak times in Both FI 2 trials. The opposite pattern was observed in Colour FI 8 trials, in which first peak times increased ($\Sigma S = 30$, $p < .002$), becoming less similar to peak times in Both FI 2 trials. Second peak times followed the same patterns, although changes in second peak times were less systematic and were not statistically significant ($\Sigma S = -8$ and 6 for Colour FI 2 and Colour FI 8 trials, both $p > .2$). Therefore, in general, as *p*(food) for colour increased, peak times shifted towards the interval signalled by the colour dimension and away from the interval signalled by the flash-frequency dimension, suggesting that changes in *p*(food) for colour affected the accuracy with which pigeons timed the 2- and 8-s intervals. This effect of *p*(food) for colour on timing accuracy was more pronounced for the 2-s interval.

In contrast to timing accuracy, relative reinforcer probabilities appeared to have little systematic effect on relative timing precision (CVs). The middle rows of Figures 5.3 and 5.4 show CVs in Colour FI 2 and Colour FI 8 trials. We tested for a possible effect of *p*(food) for colour using two-tailed nonparametric trend tests. Only the decreasing trend in second-peak CVs for Colour FI 8 trials was significant ($\Sigma S = -22$, $p < .02$). Because no other trend was observed either in first-peak CVs or second-peak CVs for other types of trial, the one significant trend probably represents a Type 1 error. Additionally, there were no significant differences in first- or second-peak CVs between Colour FI 2 or Colour FI 8 trials in all conditions (binomial sign tests: all $p > .1$). Therefore, unlike timing accuracy , (Figures 5.3 and 5.4, top rows), the relative precision with which pigeons timed the 2- and 8-s intervals was similar in Colour FI 2 and Colour FI 8 trials.

5.3.2 Changepoints

5.3.2.1 Both FI 2 and Both FI 8 trials

Figures 5.5 and 5.6 show, respectively, median first start and median first stop times, which represent the first times at which response rates transitioned from a low to a high rate (start times) and vice versa (stop times). In Both FI 2 and Both FI 8 trials, median start times were earlier than the arranged FI duration (Figure 5.5), and median stop times were around 2 s in Both FI 2 trials and later than 8 s in Both FI 8 trials (Figure 5.6). There was no systematic change in start times with changes in *p*(food) for colour in Both FI 2 and Both FI 8 trials (two-tailed nonparametric trend tests: $\Sigma S = -3$ and 13 respectively, both $p > .2$), and no systematic change in stop times for Both FI 2 trials ($\Sigma S = 0$, $p > .5$). However, a two-tailed nonparametric trend test on stop times in Both FI 8 trials was statistically significant ($\Sigma S =$ -

20, $p < .05$, $k = 5$ for all pigeons except Pigeon 74, with $k = 4$)[12]. Thus, median stop times decreased as *p*(food) for colour increased in Both FI 8 trials. However, this decrease was very slight (see Figure 5.6).

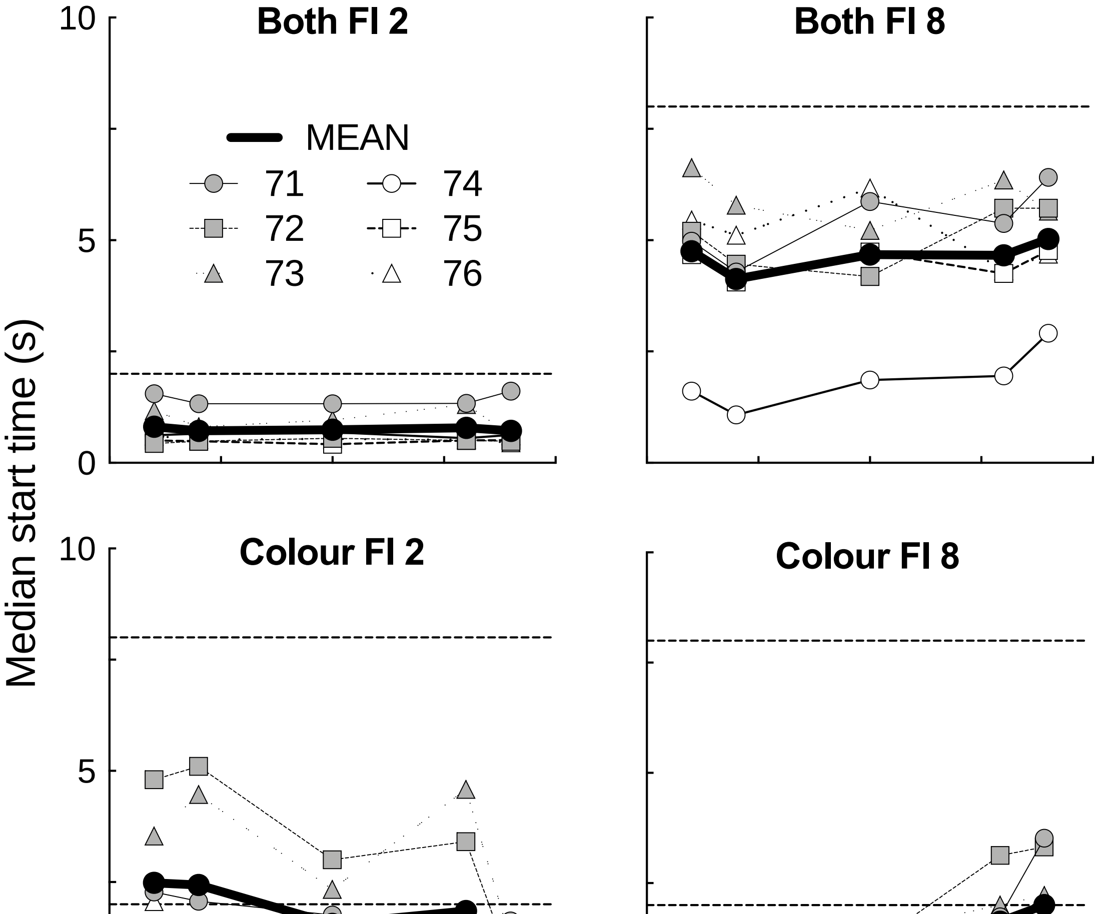

Figure 5.5. Median first start times for individual pigeons (white and grey symbols) and averaged across pigeons (solid black symbols). The dashed horizontal lines denote the arranged FI durations (2 and 8 s).

---

[12] The version of Kendall's (1955) trend test described by Elliffe and Elliffe (2019) can be used even when the number of data points differs between subjects. In this case, there are only four data points for Pigeon 74 because the algorithm used to obtain changepoints only returned those that exceeded a specified decision criterion. No stop times exceeded this criterion in Both FI 8 trials when *p*(food) for colour was .8 for Pigeon 74.

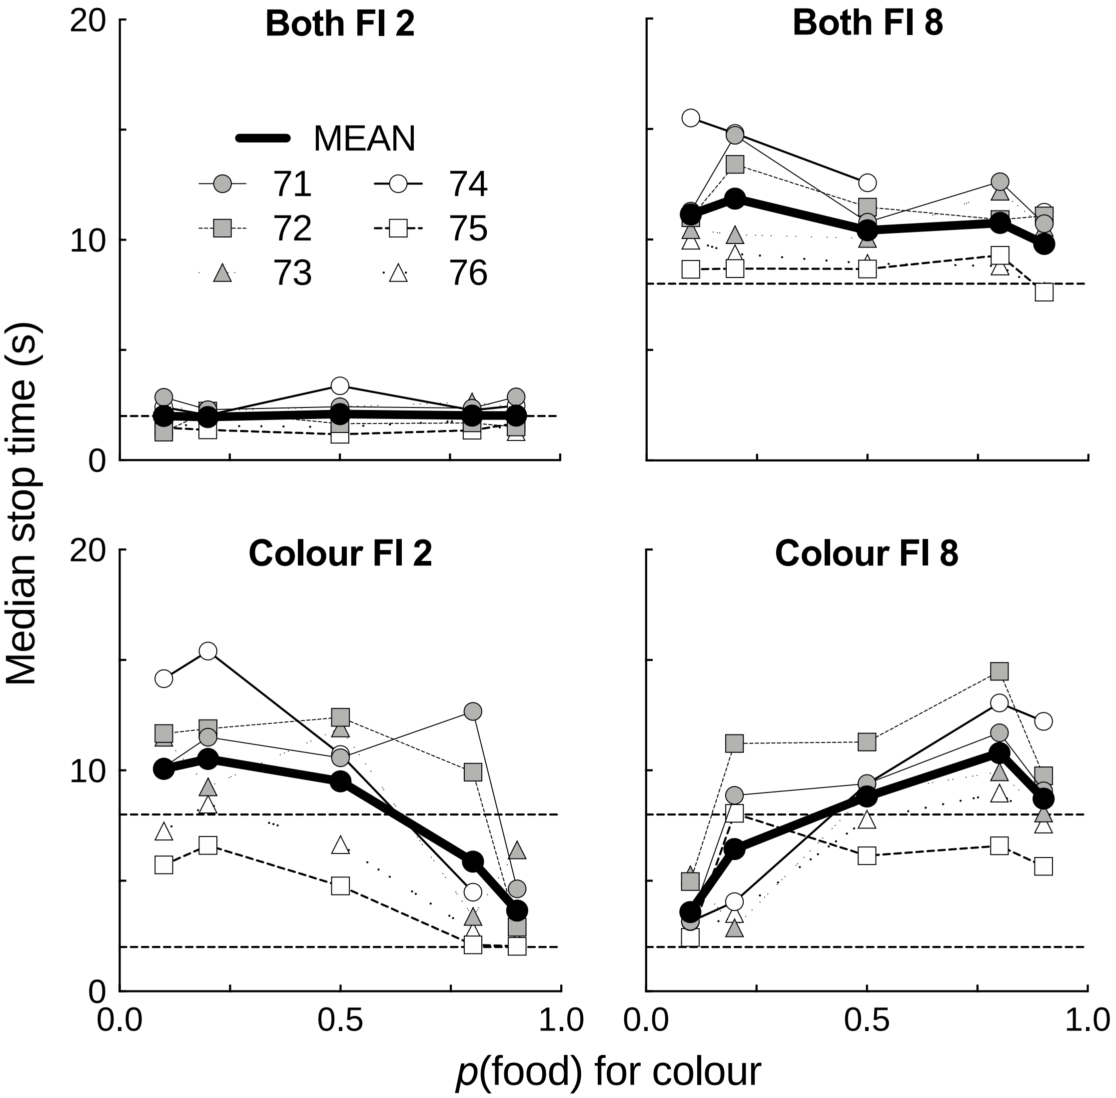

Figure 5.6. Median first stop times for individual pigeons (white and grey symbols) and averaged across pigeons (solid black symbols). The dashed horizontal lines denote the arranged FI durations (2 and 8 s).

5.3.2.2 Colour FI 2 and Colour FI 8 trials

In Colour FI 2 and Colour FI 8 trials, median start times were generally close to 2 s in all conditions (Figure 5.5). This suggests that start times depended strongly on the shortest interval to trial end; in Both FI 2, Colour FI 2, and Colour FI 8 trials, that interval was 2 s. In Both FI 8 trials, that interval was 8 s, and median start times were considerably longer than in all other trial types. There was also some effect of *p*(food) for colour; start times decreased in

Colour FI 2 trials ($\Sigma S = -44$, $p < .001$) and increased in Colour FI 8 trials ($\Sigma S = 44$, $p < .001$) as *p*(food) for colour increased. The same patterns were evident in stop times, except that the effects of *p*(food) for colour were much greater than for start times (Figure 5.6). In Colour FI 2 trials, stop times were similar to stop times in Both FI 8 trials when *p*(food) for colour was low, and decreased to become more similar to stop times in Both FI 2 trials as *p*(food) for colour increased ($\Sigma S = -28$, $p < .01$, $k = 5$ for all pigeons except Pigeon 74, with $k = 4$)[13]. The opposite pattern occurred in Colour FI 8 trials ($\Sigma S = 26$, $p < .01$).

These patterns in the median first start time and median first stop time (Figures 5.5 and 5.6) are consistent with the patterns of responding evident in Figure 5.1. In Colour FI 2 and Colour FI 8 trials, response rates immediately after trial start were similar to response rates in Both FI 2 trials, suggesting that subjects began by timing the 2-s interval. Thereafter, the time of peak responding depended on *p*(food) for colour, and hence first stop times depended on *p*(food) for colour; first stop times were early when there was a prominent peak in responding at 2 s (e.g., Colour FI 2 trials when *p*[food] for colour was .9), and later when there was a prominent peak in responding at 8 s (e.g., Colour FI 8 trials when *p*[food] for colour was .9).

5.3.3 Summary of Main Results

When two dimensions (colour and flash-frequency) of a compound stimulus signalled a 2-s or 8-s interval to response-contingent reinforcer delivery, pigeons timed the interval(s) signalled by the stimulus dimensions, indicating that both dimensions controlled behaviour (Figures 5.1 to 5.6). As the probability of reinforcer deliveries associated with the colour dimension increased, (1) a peak at the time of reinforcer delivery signalled by the colour dimension became apparent and a peak at the time of reinforcer delivery signalled by the

---

[13] No stop times exceeded the decision criterion in Colour FI 2 trials when *p*(food) for colour was .9 for Pigeon 74.

flash-frequency dimension diminished (Figures 5.1 to 5.4); (2) peak times shifted towards the interval signalled by the colour dimension in Colour FI 2 and Colour FI 8 trials (Figures 5.3 and 5.4); and (3) start and stop times shifted earlier in Colour FI 2 trials and later in Colour FI 8 trials (Figures 5.5 and 5.6). In contrast, relative reinforcer probabilities had no systematic effect on normalised peak spreads (CVs; Figures 5.3 and 5.4). Nevertheless, in combination, the changes seen in a variety of measures of temporal discrimination showed that as the relative reinforcer probability associated with the colour dimension increased, pigeons were more likely to time the interval signalled by the colour dimension and less likely to time the interval signalled by the flash-frequency dimension.

5.4 Discussion

The present experiment investigated whether relative reinforcer probabilities determine the division of control between the dimensions of compound stimuli when each dimension signalled a 2-s or 8-s interval to response-contingent reinforcer delivery in a multiple peak procedure. This experiment was a systematic replication of Davison and Elliffe (2010), who investigated the effects of relative reinforcer probabilities on divided control between stimulus dimensions that signalled the response alternative (left or right) that was more likely to deliver reinforcers in a DMTS task. Davison and Elliffe found that discriminative control by a dimension depended on the relative probability of reinforcer deliveries associated with that dimension; pigeons preferred the alternative signalled by the dimension associated with the higher reinforcer probability (see also e.g., Davison, 2018a; Shahan & Podlesnik, 2006, 2007). Similarly, we found that pigeons' response rates were higher and temporal discrimination was slightly more accurate for the interval signalled by the dimension associated with the higher reinforcer probability (Figures 5.1 to 5.4), thus suggesting stronger control by that dimension, compared with the dimension associated with the lower reinforcer probability, over interval timing. Therefore, the present findings extend previous research demonstrating that relative reinforcer probabilities divide stimulus control in space to the temporal domain.

In the peak procedure, strong control by a stimulus signalling a fixed interval to reinforcer delivery would be apparent from accurate and precise discrimination of the interval (i.e., peak times close to the arranged FI duration, and smaller peak spreads). Indeed, previous research suggests that when stimulus control shifts away from a stimulus signalling the time to reinforcer delivery and towards a stimulus that provides no such temporal information, temporal discrimination accuracy and/or precision worsens (Aum, Brown, & Hemmes, 2007; Cowie, Davison, & Elliffe, 2017; Gomes-Ng, Elliffe, & Cowie, 2018a; McMillan & Roberts, 2013; Sutton & Roberts, 2002). Thus, although relative reinforcer probabilities had strong and consistent effects on peak rates in Colour FI 2 and Colour FI 8 trials in the present experiment (Figures 5.3 and 5.4, bottom rows), peak rate alone is not a measure of temporal *discrimination* per se, and peak times and spreads therefore provide a better measure of the strength of *discriminative* stimulus control by the compound-stimulus dimensions.

Focusing on peak times and spreads (Figures 5.3 and 5.4, top and middle rows), the effects of relative reinforcer probabilities on divided stimulus control in the present experiment appear smaller than the effects reported by Davison and Elliffe (2010) and by other studies examining the relation between relative reinforcers and divided stimulus control in DMTS (e.g., Shahan & Podlesnik, 2006, 2007). Even when the probability of reinforcer deliveries associated with one dimension was lowest (.1), the pigeons still appeared to time the interval signalled by that dimension; first- and second-peak times in Colour FI 2 and Colour FI 8 trials were generally close to 2 and 8 s, respectively, in all conditions (Figures 5.2 to 5.4). Additionally, the effects of relative reinforcer probabilities on temporal discrimination appeared to be somewhat smaller for the 8-s interval, as a peak at 8 s was almost always apparent in Colour FI 2 and Colour FI 8 trials (Figure 5.1) and changes in timing accuracy (peak times; Figures 5.3 and 5.4) were less systematic than for the 2-s interval.

5.4.1 Joint Control by Compound Stimuli and by Elapsed Time

These smaller effects, relative to previous research (e.g., Davison & Elliffe, 2010), of relative reinforcer probabilities on divided stimulus control in the present experiment may be related to competition for control between the compound-stimulus dimensions and elapsed time itself. In the present procedure, changes in relative reinforcer probabilities associated with each stimulus dimension also changed the probability of reinforcer deliveries at two different times (2 s and 8 s) in Colour FI 2 and Colour FI 8 trials. As a result, changes in peak response functions across conditions may have been additionally driven by changes in the probability of reinforcer deliveries *at a time*, rather than solely by changes in the strength of control by the colour and flash-frequency dimensions. The similarities between response patterns at early times in Both FI 2, Colour FI 2, and Colour FI 8 trials, and between response patterns at later times in Both FI 8, Colour FI 2, and Colour FI 8 trials, provide some support for this suggestion – subjects appeared to behave *as if* an FI 2-s schedule was arranged at early times, and as if an FI 8-s schedule was arranged at later times, consistent with some control by elapsed time (i.e., by “what is likely to happen now”).

Furthermore, some evidence suggests that when only elapsed time controls behaviour in the peak procedure, changes in the probability of reinforcer deliveries at the end of the fixed interval do not necessarily cause changes in timing accuracy or precision, but they do cause changes in peak response rates (i.e., peak height). For example, Whitaker et al. (2008) found that a 30-s interval was timed with similar accuracy and precision when the probability of reinforcer deliveries varied from .1 to 1.0; the probability of reinforcer deliveries only determined rates of responding at 30 s (see also Bouton & Sunsay, 2003; Galtress & Kirkpatrick, 2009; S. Roberts, 1981). Thus, changes in *p*(food) for colour may have had small effects on temporal discrimination accuracy and precision, and larger effects on peak response rates, in the present experiment because elapsed time since trial start exerted strong control over behaviour.

Competition between control by the compound-stimulus dimensions and by elapsed time may also explain the slightly asymmetrical effects of relative reinforcers on discrimination of the 2- and 8-s intervals in Colour FI 2 and Colour FI 8 trials (Figures 5.1, 5.3, and 5.4). In these trials, there was a time-based change in the predictability of the next-reinforcer time (Whitaker, Lowe, & Wearden, 2003, 2008; see also Cowie, Davison, & Elliffe, 2011): Before 2 s had elapsed, the time of the next reinforcer (or trial end) was uncertain because it could occur at either 2 or 8 s, whereas after 2 s had elapsed, the next reinforcer (or trial end) would definitely occur at 8 s regardless of the compound stimulus. Thus, control by the compound stimulus may have been stronger at earlier times and control by elapsed time may have been stronger at later times, hence resulting in more systematic changes in discrimination (timing accuracy) of the 2-s, but not the 8-s, interval.

5.4.2 Processes Underlying Timing of Compound Stimuli

Additionally, or alternatively, the relatively small changes in temporal discrimination of the 2-s and 8-s intervals in Colour FI 2 and Colour FI 8 trials may be related to how the pigeons processed compound stimuli. According to a *configural*-processing approach, a compound stimulus is treated as a ‘whole’ stimulus, that is, as a separate entity from its constituent elements (e.g., Grings & Dawson, 1973; Pearce, 1987, 1994). This is contrasted with *elemental* processing, in which the compound stimulus is treated as the ‘sum’ of its constituent elements (e.g., Blough, 1975; Rescorla & Wagner, 1972). Clearly, divided control by two dimensions of a compound stimulus depends on elemental processing, as such division implies some degree of control by each individual dimension. Therefore, relative reinforcer probabilities can only determine the division of control between compound-stimulus dimensions to the extent that they are processed elementally. If our pigeons processed the compound stimuli configurally, then the multiple peak procedure arranged here would be similar to a multiple peak procedure in which four single-element stimuli (e.g., four colours) signalled a 2-s and/or 8-s interval to reinforcer delivery. That is, the four compound stimuli would have been discriminated to be four separate stimuli each signalling one (Both FI 2 and Both FI 8 trials) or two (Colour FI 2 and Colour FI 8 trials) FI schedules, rather than as stimuli comprised of common red, green, fast, and slow elements that each signalled one FI schedule. If so, changes in relative reinforcer probabilities would have appeared to cause changes in the probability of reinforcer deliveries at 2 and 8 s, rather than to cause changes in the probability of reinforcer deliveries associated with the colour and flash-frequency dimensions.

This configural-processing approach may also help to explain the differences between the present results and those of studies of temporal averaging. In studies of temporal averaging, separately trained light and tone stimuli, which each signal a different FI schedule, are occasionally presented simultaneously in peak trials. Like the present results, the extent of control by each stimulus in temporal-averaging studies appears to depend on relative reinforcer probabilities (Matell & Kurti, 2014). However, unlike our pigeons, Matell and Kurti’s (2014) subjects (rats) “averaged” the FI durations associated with the light and tone stimuli together, with each stimulus weighted by its relative reinforcer probability (see also De Corte & Matell, 2016a). Thus, the mechanisms underlying interval timing appear to differ depending on whether stimuli are trained separately or in compound. Some evidence suggests that a compound stimulus is more likely to be processed elementally if subjects previously experience discrimination trials with the individual elements that comprise the compound, whereas configural processing is more likely without such previous experience (e.g., M. F. Brown, 1987; D. A. Williams & Braker, 1999; D. A. Williams, Sagness, & McPhee, 1994). Additionally, elemental processing may be more likely with elements that are more perceptually separable (e.g., those from different stimulus modalities, such as auditory and visual stimuli in the temporal-averaging procedure) than those that are less separable (e.g., the colour and flash-frequency dimensions, which were unified into a single compound stimulus in the present experiment; Kehoe, Horne, Horne, & Macrae, 1994). Thus, the temporal-averaging procedure may be more likely to engender elemental processing than the present procedure. If so, then the absence of temporal averaging in the present experiment is unsurprising, because averaging the FI durations previously associated with individual elements of a compound stimulus requires elemental processing.

Matell and Kurti (2014) explained temporal averaging using Scalar Expectancy Theory (SET; Gibbon, 1977; Gibbon, Church, & Meck, 1984), a theory of timing which posits that subjects store a distribution of experienced intervals in memory, and the decision to respond depends on a comparison between the current time and a value sampled from the memory distribution. The memory distribution of intervals associated with a stimulus is formed as a result of the reinforcers obtained in the presence of a stimulus. When a stimulus signals a single interval to reinforcer delivery, it becomes associated with a single distribution of intervals. When multiple intervals are associated with the same stimulus, the stimulus may become associated with more than one distribution, provided that the intervals are different enough that the subject can discriminate between them (see Matell, Kim, & Hartshorne, 2014). Thus, in studies of temporal averaging (in which subjects obtain reinforcers in the presence of each individual stimulus), each stimulus becomes associated with a single memory distribution during training. In contrast, no distributions are associated with the compound stimulus because reinforcers are never delivered in its presence. The opposite was true in the present experiment; the pigeons obtained reinforcers in the presence of the compound stimuli, but never in the presence of the individual stimulus elements. This difference may be responsible for the difference between Matell and Kurti's and the present findings. When the individual elements of a compound stimulus are each associated with different memory distributions, values from each of the distributions may be sampled and then averaged together (Matell & Kurti, 2014). However, when the compound itself is associated with different memory distributions, the intervals may be timed simultaneously, resulting in bimodal peak response functions (see e.g., Church, Guilhardi, Keen, MacInnis, & Kirkpatrick, 2003; Leak & Gibbon, 1995; Whitaker et al., 2008).

Therefore, one fruitful avenue for future research would be to include some trials with single-element stimuli (i.e., just red, green, fast, and slow stimuli) in the present procedure. Such trials would provide insight into whether subjects learned the individual element-FI correlations during training with compound stimuli; if they did, then they should time the interval signalled by each element in single-element peak trials. Additionally, including single-element trials may increase the likelihood of elemental processing, and hence may result in temporal averaging in the presence of compound stimuli. Such a study would also help to elucidate whether temporal averaging is related to the prior training with individual stimulus elements, the *absence* of any training with the compound stimulus, or both. Further investigation of the species similarities and differences between rats and pigeons may also provide some useful insights into the discrepancies between the present data and temporal-averaging data. It is possible that when faced with compound stimuli comprised of elements that signal different FI durations, rats and pigeons use different timing strategies. Although temporal averaging has now been demonstrated several times in rats (e.g., Delamater & Nicolas, 2015; Kurti et al., 2013; Matell & Kurti, 2014; Swanton et al., 2009; Swanton & Matell, 2011), it remains to be seen whether the effect is robust in pigeons. One study (Cheng

& Roberts, 1991) provided inconclusive evidence of temporal averaging in pigeons, and some unpublished data from our lab provided similarly inconclusive evidence of temporal averaging in pigeons.

The difference in results between Matell and Kurti (2014) and the present experiment may also be related to stimulus-generalisation processes. Subjects hardly encounter the compound stimulus in temporal-averaging studies, whereas all trials involved compound stimuli in the present experiment. It is well established that behaviour generalises from previously trained stimuli to similar but novel stimuli (e.g., Guttman & Kalish, 1956; see Honig & Urcuioli, 1981 for a review). Such *stimulus generalisation* may explain temporal averaging; response-rate patterns learned in the presence of the visual and auditory elements may each generalise relatively equally to the visual-auditory compound stimulus, resulting in “averaged” peak response functions. In support of this, generalisation of responding from a previously trained element to a compound stimulus containing that element can occur (e.g., M. F. Brown, 1987; Grant & MacDonald, 1986). Furthermore, Kurti et al. (2013) found that when rats were injected with saline or amphetamine and trained in an FI 5-s or FI 20-s schedule respectively, administering an intermediate dose of amphetamine (i.e., a novel stimulus) resulted in temporal averaging (see also Gomes-Ng et al., 2018b; Nelson & Farthing, 1973). Thus, temporal averaging also appears to occur under conditions similar to those that are used to study stimulus generalisation. In contrast, no such generalisation would be expected in the present experiment, because the compound stimuli were not novel. Instead, pigeons in the present experiment would be expected to behave in accordance with the reinforcer contingencies (reinforcers at 2 and 8 s) signalled by the compound-stimulus dimensions.

5.4.3 Frequency of Experience versus Frequency of Reinforcers

The present experiment also shares similarities with a recent study on cue informativeness in the multiple peak procedure (Subramaniam & Kyonka, 2019). Subramaniam and Kyonka (2019) varied the correlation between two stimuli and two FI schedules such that both intervals were equally likely in the presence of each stimulus (low correlation) or one interval was more likely than the other for each stimulus (high correlation). They found that pigeons timed both intervals (i.e., bimodal peak response functions) when the correlation was low, whereas they timed only the interval signalled by the stimulus (i.e., unimodal functions) when the correlation was high. Additionally, as the stimulus-FI correlation increased, peak times shifted away from the unlikely interval and towards the likely interval, indicating less accurate timing of the unlikely interval, whereas peak spreads changed little with changes in the correlation. Thus, manipulations of stimulus-FI correlations and of relative reinforcer probabilities appear to have similar effects on temporal discrimination – response rates are higher and subjects are more likely to time accurately intervals that occur more frequently (Subramaniam & Kyonka, 2019), or that are reinforced more frequently (the present experiment). However, changes in timing accuracy were larger in Subramaniam and Kyonka than in the present experiment. One explanation for this difference is that Subramaniam and Kyonka's manipulation changed the frequency of experience with each interval in the presence of each stimulus, whereas our manipulation changed only the probability of a reinforcer delivery at the end of each interval. The frequency of experience with an interval may have stronger effects on temporal discrimination, perhaps because there are less opportunities to discriminate an interval as the frequency of that interval decreases. In support of this, some research suggests that temporal discrimination improves with training, suggesting that frequency of experience may play a role in temporal discrimination (see e.g., Balci, Gallistel, Allen, Frank, Gibson, & Brunner, 2009; Machado & Cevik, 1998).

5.5 Chapter Summary and Conclusions

In a recent paper, Davison (2018b) posed the question, “Do divided [stimulus control] effects occur in temporal discriminations?” (p. 136). The present experiment is the first step towards answering this question, and the present results tentatively suggest that the answer is “yes”, with the caveat that changes in accuracy were very small in the present experiment, and the effects of relative reinforcer probabilities on divided stimulus control appeared to be time-dependent; such effects were greater for the shorter (2 s) interval than for the longer (8 s) interval. Thus, the present findings extend previous research investigating the relation between relative reinforcers and divided stimulus control (e.g., Davison, 2018a; Davison & Elliffe, 2010; Shahan & Podlesnik, 2006, 2007; Experiments 1 and 2, Chapters 3 and 4) to temporal discriminations, and also suggest that there may be some differences between how relative reinforcers divide stimulus control in space and time.

Additionally, the present findings extend studies of temporal averaging, which examine the timing of compound stimuli comprised of previously trained visual and auditory elements in the multiple peak procedure (e.g., Delamater & Nicolas, 2015; Kurti et al., 2013; Matell & Kurti, 2014; Swanton et al., 2009; Swanton & Matell, 2011). Taken together, the present findings and those of temporal-averaging studies suggest that multiple stimuli that signal discrepant times to reinforcer delivery control behaviour jointly, but the processes that underlie timing under such conditions may differ depending on whether the compound stimulus itself is associated with different intervals or whether its constituent elements are separately associated with different intervals. This parallels the findings of Experiments 1 and 2 (Chapters 3 and 4), in which the type of training procedure (separately training elements versus training them in compound) modulated the behavioural expression of divided stimulus control.

The present experiment, like Experiments 1 and 2, helps to establish further the generality of the relation between relative reinforcer probabilities and divided stimulus control, and also suggests some of the limits of this generality. Specifically, when stimuli signal the time of the next reinforcer, the effects of relative reinforcer probabilities on divided stimulus control appear to be smaller compared with when stimuli signal the next-reinforcer location, and such effects may diminish over time. This may be related to the presence of an additional discriminative stimulus – elapsed time – that also exerts some control over behaviour (Figure 5.1). In contrast, no such additional stimulus controls behaviour when the compound-stimulus dimensions signal the next-reinforcer location. Thus, in addition to extending previous divided-stimulus-control and temporal-averaging research, the present findings also demonstrate that when reinforcer availability changes across time, non-temporal (in this case, visual compound stimuli) and temporal stimuli may compete for control over behaviour (see also McMillan, Spetch, Sturdy, & Roberts, 2017).

Appendix A5

Response rates in peak trials for individual pigeons, and fits of the single and summed Gaussian models (Equations 5.1 and 5.2) to peak-trial data

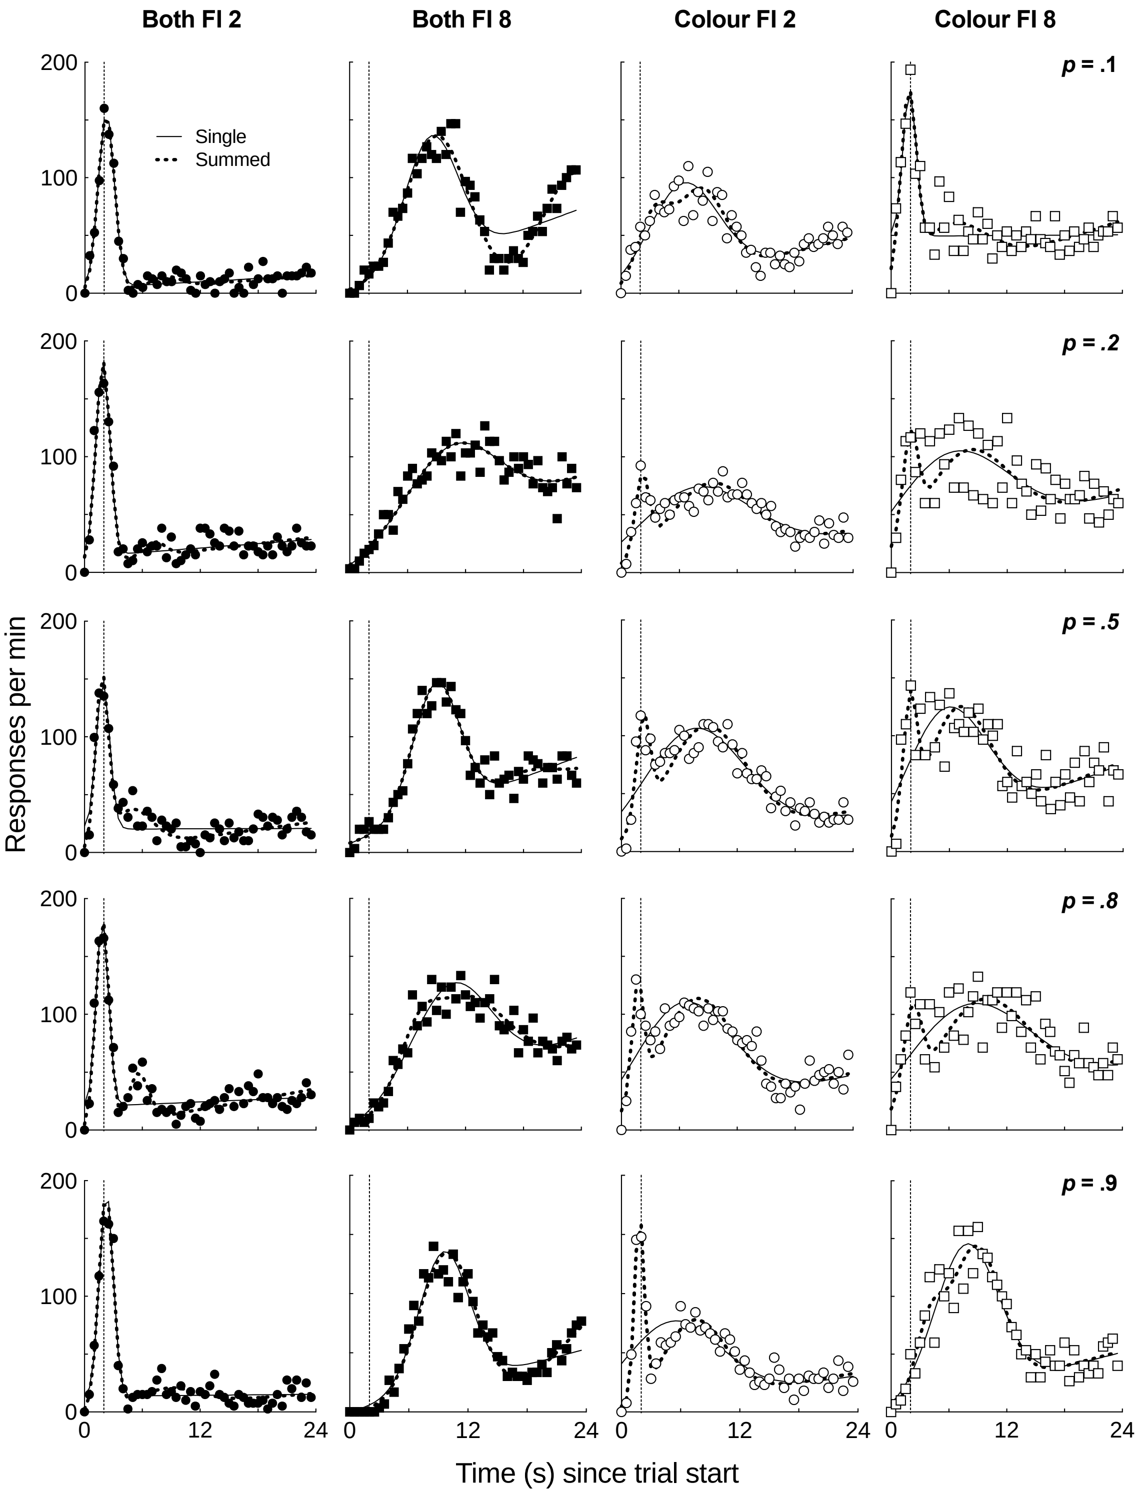

Figure A5.1. Responses per min across time in peak trials for Pigeon 71. The solid and dotted lines show fits of the single and summed Gaussian model, respectively.

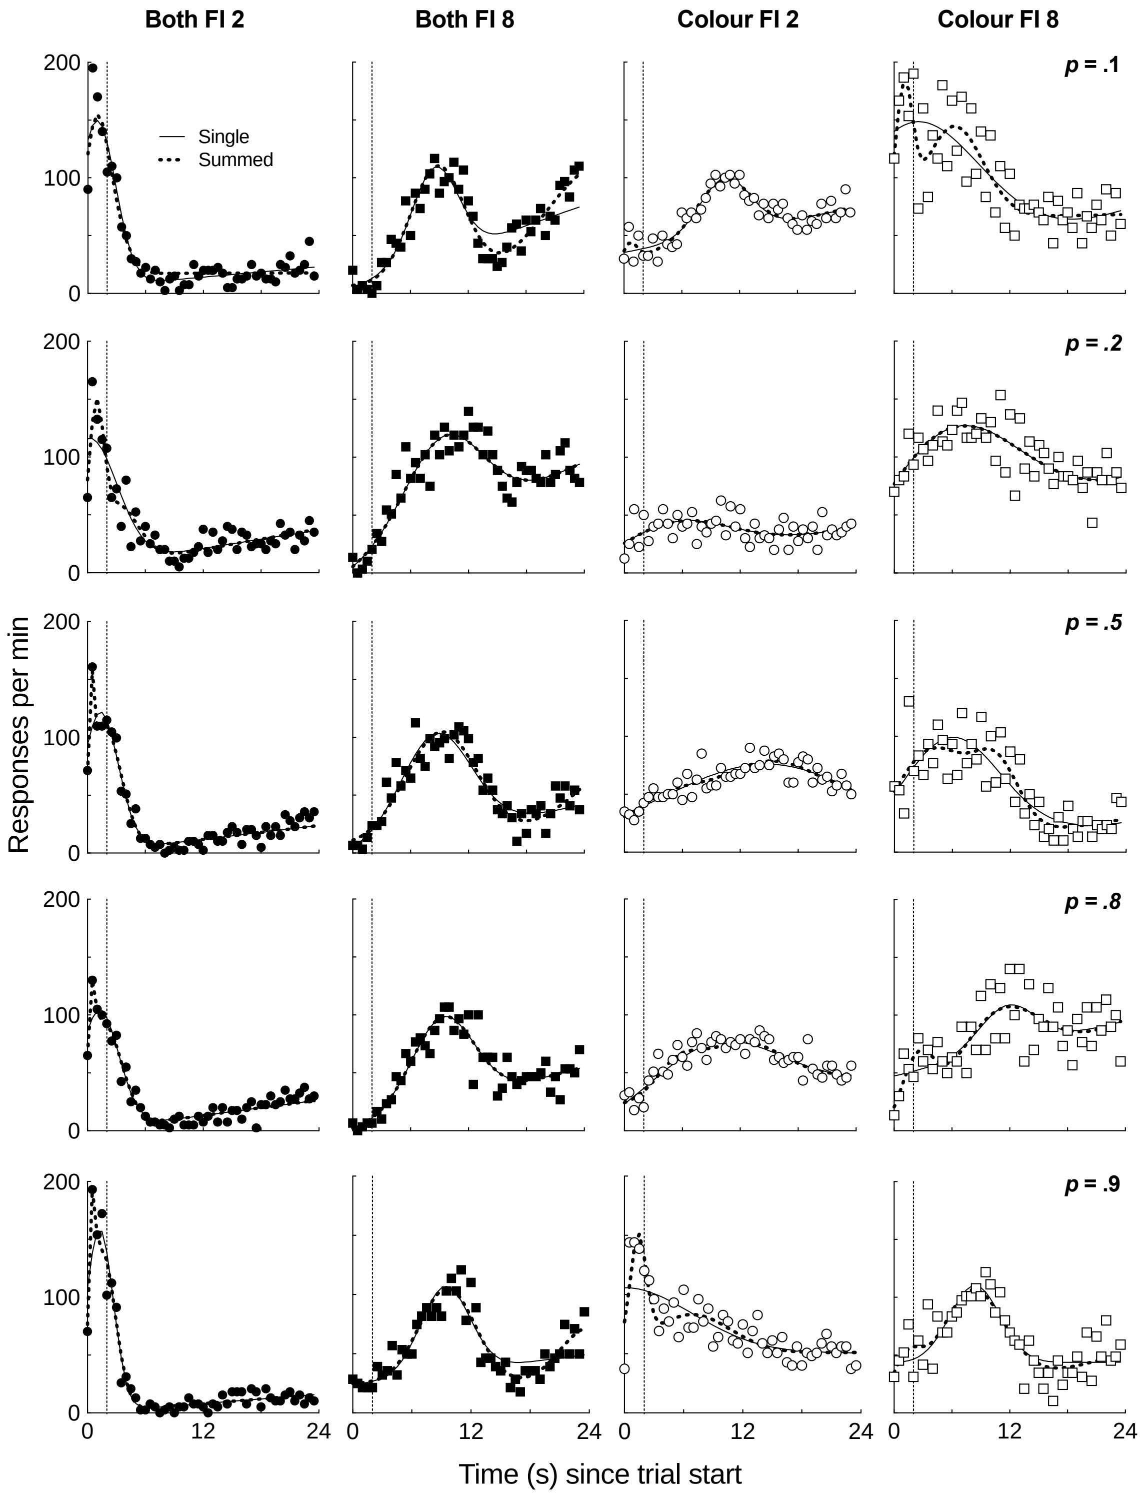

Figure A5.2. Responses per min across time in peak trials for Pigeon 72. The solid and dotted lines show fits of the single and summed Gaussian model, respectively.

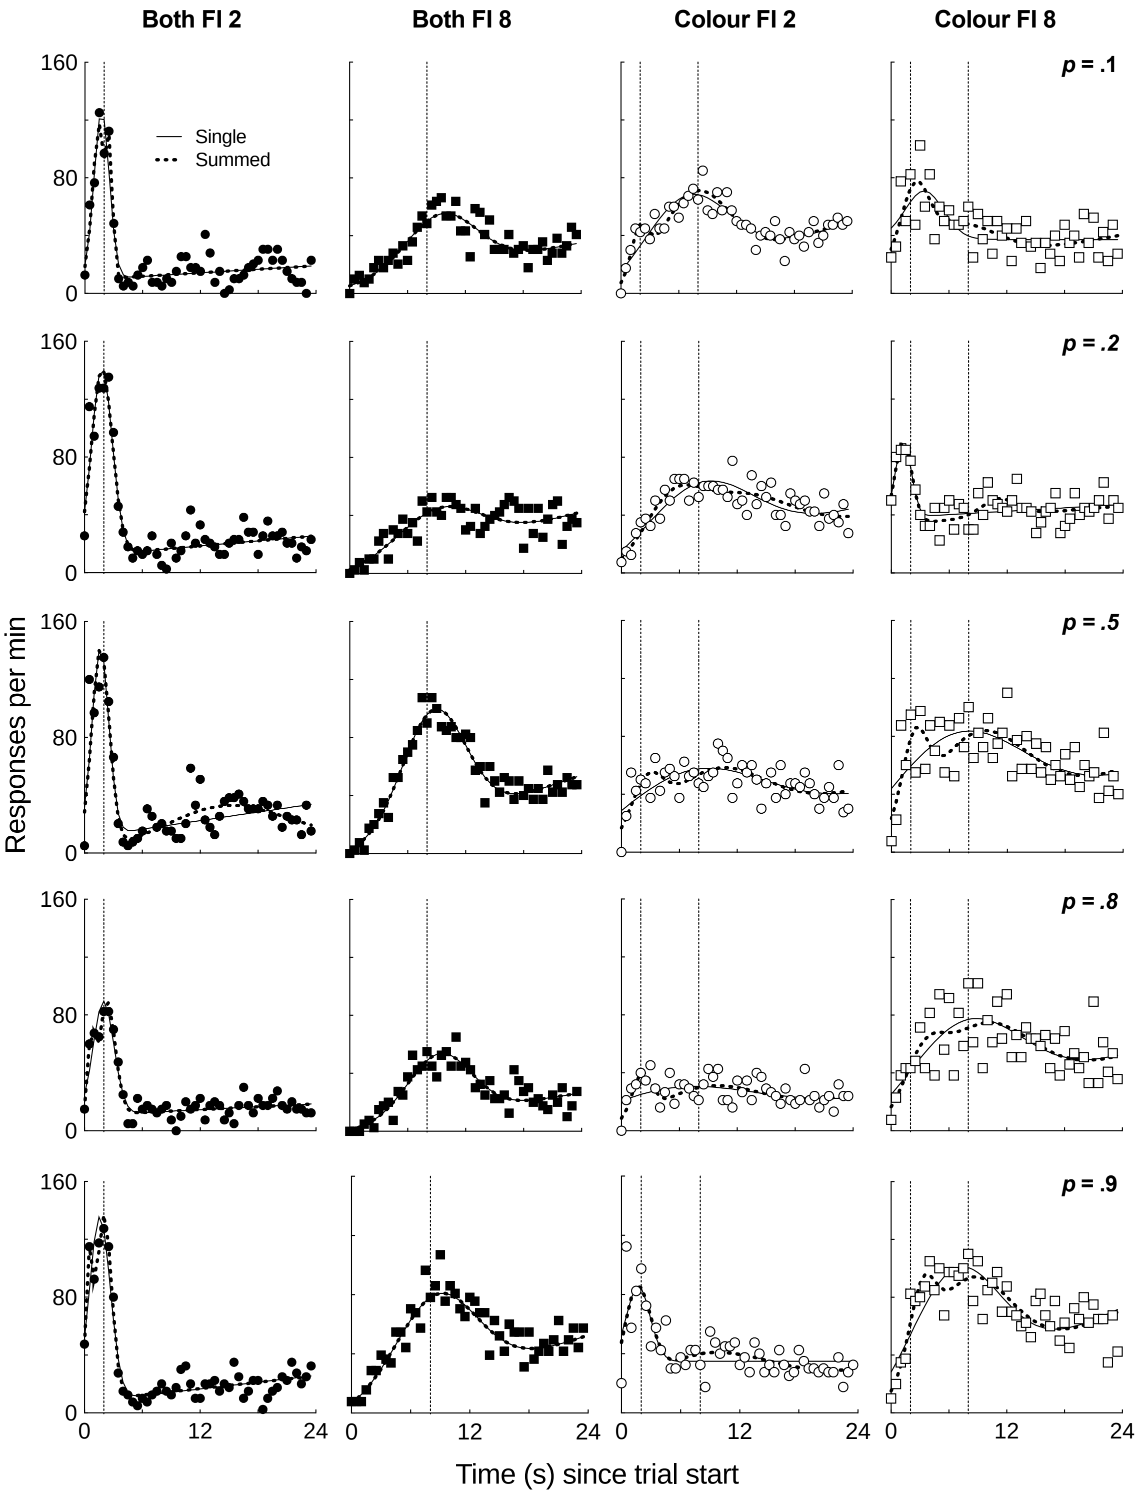

Figure A5.3. Responses per min across time in peak trials for Pigeon 73. The solid and dotted lines show fits of the single and summed Gaussian model, respectively.

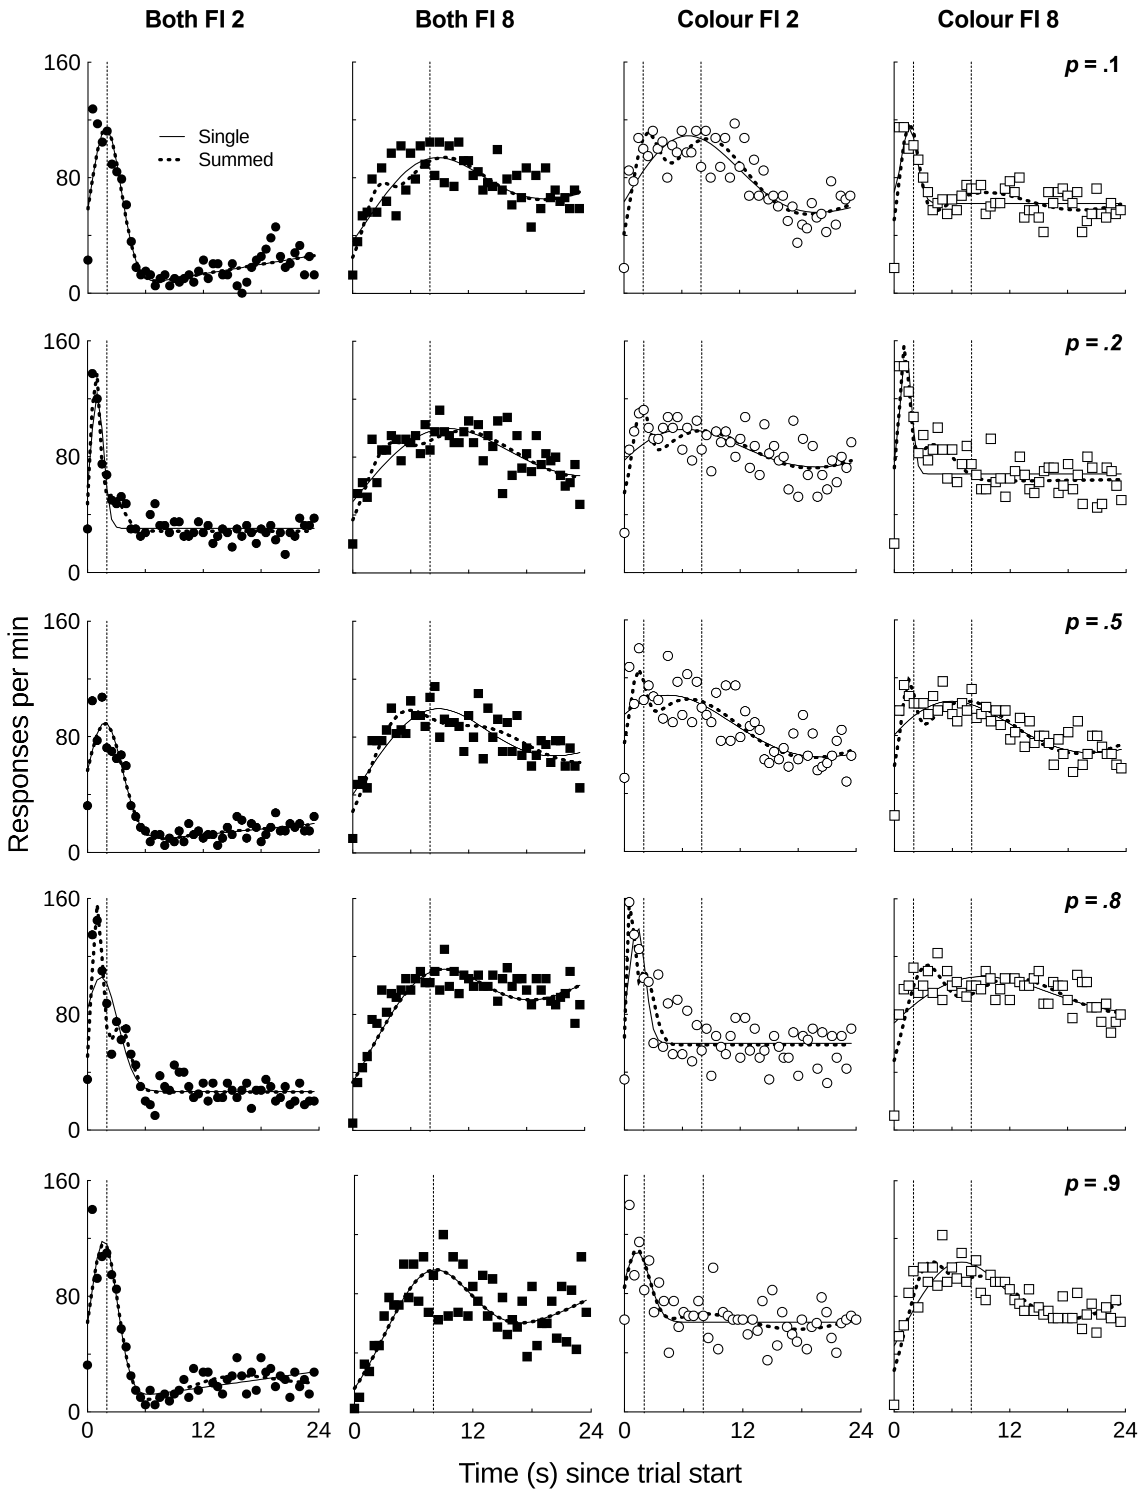

Figure A5.4. Responses per min across time in peak trials for Pigeon 74. The solid and dotted lines show fits of the single and summed Gaussian model, respectively.

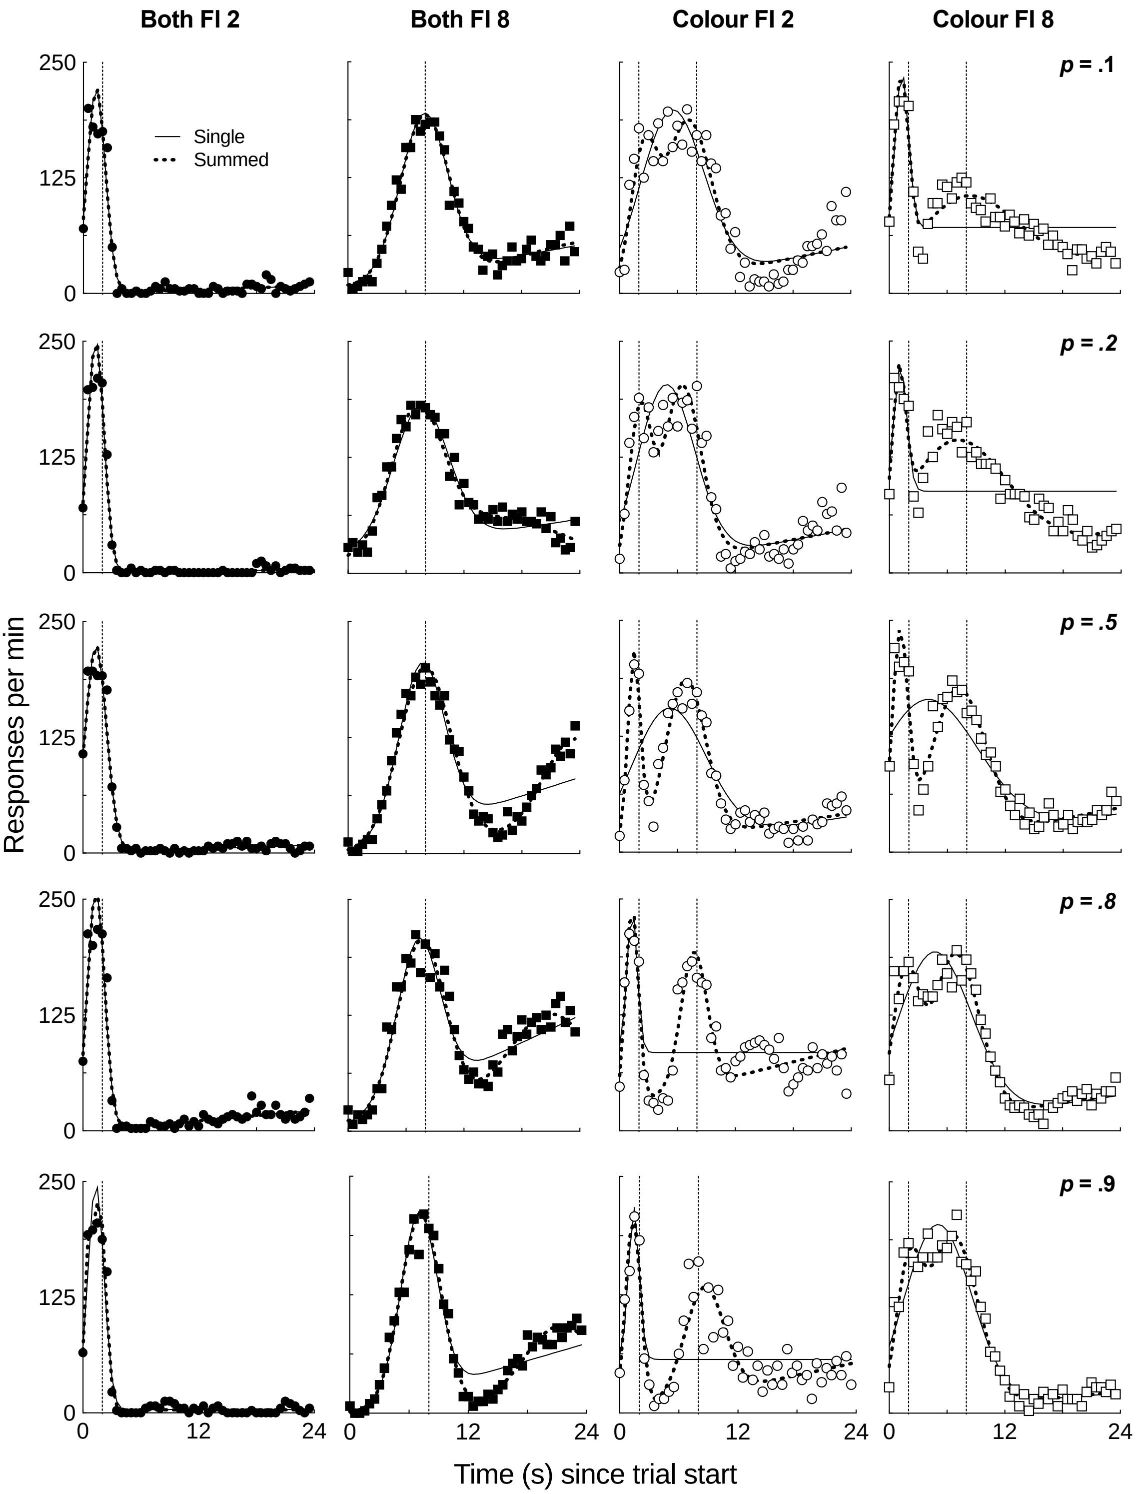

Figure A5.5. Responses per min across time in peak trials for Pigeon 75. The solid and dotted lines show fits of the single and summed Gaussian model, respectively.

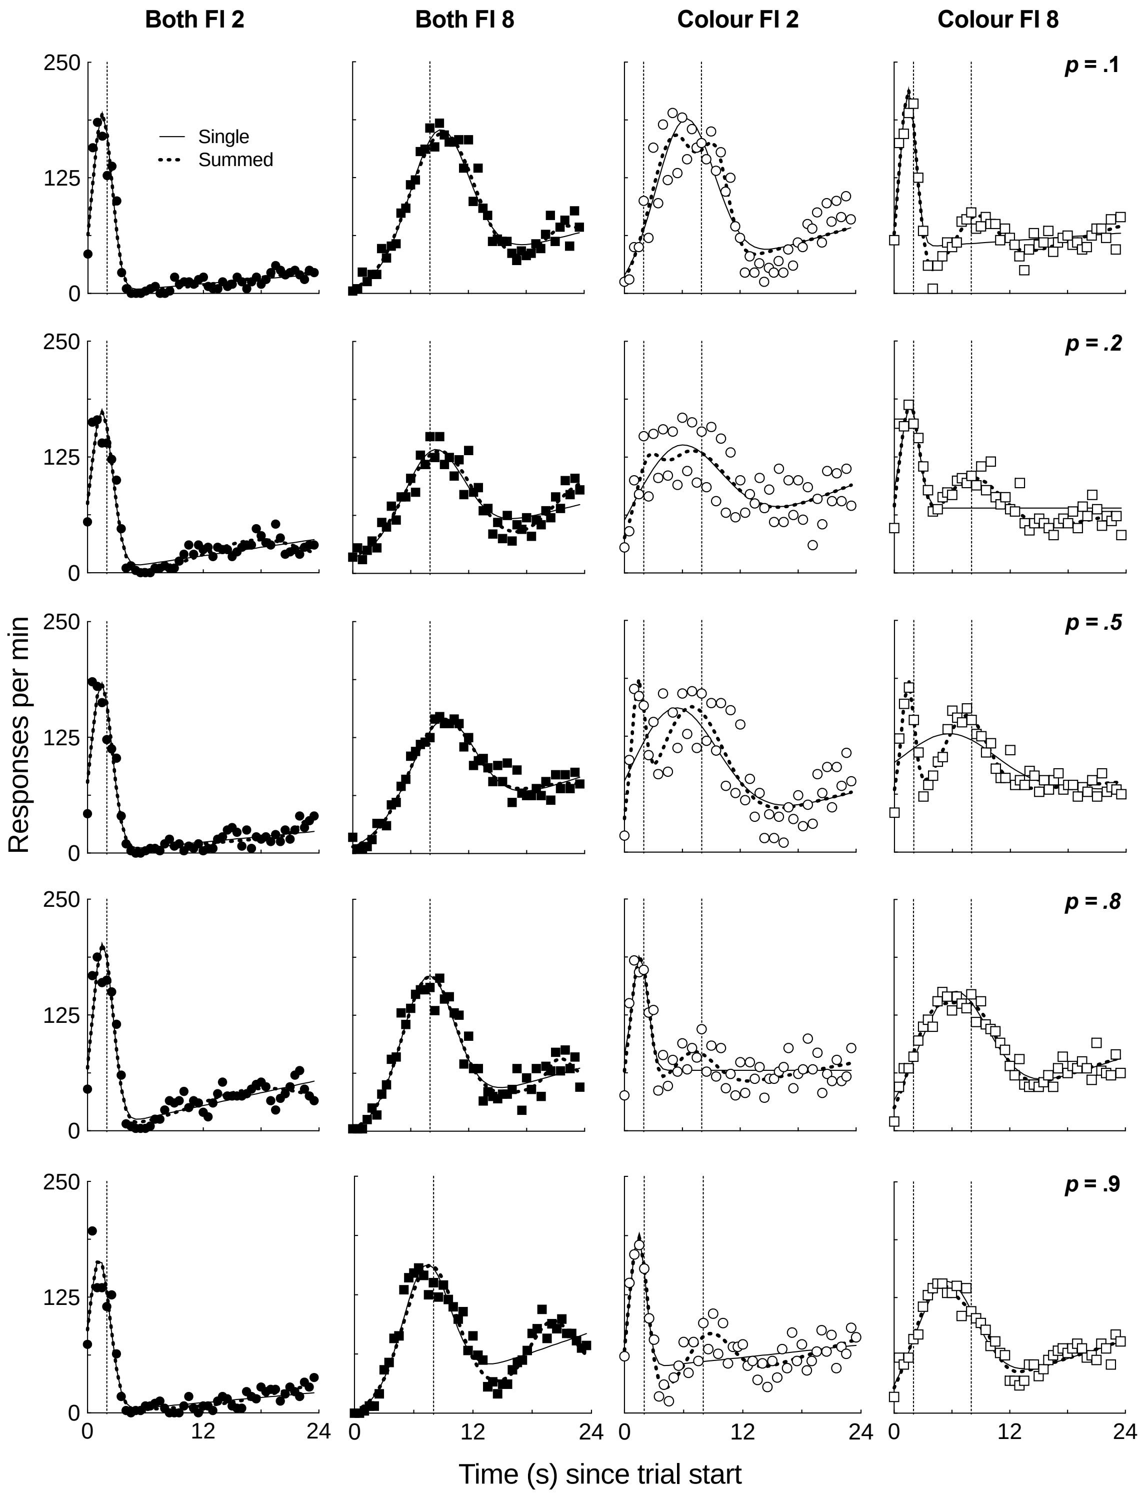

Figure A5.6. Responses per min across time in peak trials for Pigeon 76. The solid and dotted lines show fits of the single and summed Gaussian model, respectively.

Appendix B5

Parameter estimates, variance accounted for (VAC), and Akaike's Information Criterion corrected for small sample sizes (AICc) for fits of Equations 5.1 and 5.2

Table B5.1.

Parameter estimates, proportion of variance accounted for (VAC), and Akaike's Information Criterion (AICc) for single ramped Gaussian for Pigeon 71.

p(food for colour)	*a*	*M*	*SD*	*s*	*b*	*VAC*	*AICc*
Both FI 2							
.1	2.59	2.23	0.82	0.00	0.04	.96	**856.48**
.2	2.48	1.91	0.71	0.01	0.11	.91	**1209.42**
.5	1.84	1.92	0.69	0.00	0.17	.84	1145.15
.8	2.12	1.85	0.63	0.00	0.17	.87	1251.51
.9	2.59	2.27	0.70	0.00	0.11	.93	982.82
Both FI 8							
.1	4.81	8.47	2.78	0.02	0.00	.80	1499.51
.2	5.58	10.60	4.81	0.02	0.00	.86	**1619.48**
.5	4.16	8.97	2.40	0.02	0.05	.93	**1579.57**
.8	5.80	10.50	4.08	0.02	0.00	.89	**1612.46**
.9	4.58	9.55	2.56	0.01	0.00	.91	1307.73
Colour FI 2							
.1	6.12	6.47	3.55	0.02	0.00	.76	**1580.63**
.2	7.64	7.62	5.78	0.01	0.00	.57	1588.81
.5	10.48	7.47	5.20	0.01	0.00	.73	1705.20
.8	10.17	6.32	4.98	0.02	0.00	.67	1820.84
.9	7.17	5.10	4.89	0.01	0.00	.40	1520.32
Colour FI 8							
.1	1.38	1.91	0.70	0.00	0.31	.74	1498.09
.2	8.18	6.01	6.03	0.02	0.00	.32	1734.68
.5	6.97	5.61	4.18	0.02	0.00	.51	1709.79
.8	9.27	7.83	6.48	0.01	0.00	.49	1736.53
.9	6.65	7.81	3.30	0.01	0.00	.86	**1527.21**

Note: Bolded AICc indicates that this model is the best-fitting model for those data.

Table B5.2.

Parameter estimates, proportion of variance accounted for (VAC), and Akaike's Information Criterion (AICc) for summed ramped Gaussian for Pigeon 71.

p(food for colour)	a_1	M_1	SD_1	a_2	M_2	SD_2	*s*	*b*	*VAC*	*AICc*
Both FI 2										
.1	2.74	2.23	0.85	0.29	8.36	1.50	0.01	0.00	.96	856.77
.2	2.65	1/91	0.73	0.26	7.29	1.06	0.01	0.06	.92	1212.84
.5	1.95	1.81	0.65	1.40	4.91	2.10	0.01	0.00	.91	**1126.74**
.8	2.59	1.86	0.69	0.88	6.60	1.04	0.01	0.00	.93	**1232.81**
.9	2.91	2.24	0.75	0.97	8.64	3.12	0.01	0.00	.94	**978.32**
Both FI 8										
.1	7.11	9.05	3.33	6.03	23.63	3.67	0.00	0.00	.92	**1484.25**
.2	5.58	10.60	4.81	0.00	0.00	0.00	0.02	0.00	.86	1625.52
.5	4.40	8.98	2.45	0.68	18.05	3.03	0.02	0.05	.93	1584.89
.8	1.14	7.08	1.96	5.14	11.77	4.24	0.02	0.00	.92	1615.07
.9	5.99	9.84	2.90	16.87	32.47	8.36	0.00	0.00	.94	**1306.59**
Colour FI 2										
.1	2.00	3.36	1.71	3.94	8.25	2.55	0.02	0.00	.80	1580.79
.2	0.95	2.15	0.70	5.78	9.06	4.37	0.01	0.00	.83	**1574.06**
.5	1.26	2.26	0.69	8.28	8.65	4.13	0.01	0.00	.87	**1684.63**
.8	1.16	1.69	0.59	8.04	7.63	3.94	0.02	0.00	.85	**1802.49**
.9	1.50	1.81	0.48	4.48	7.38	3.17	0.01	0.00	.92	**1463.86**
Colour FI 8										
.1	1.85	1.81	0.79	2.50	6.06	3.62	0.02	0.00	.78	**1496.42**
.2	0.95	1.98	0.76	5.86	7.66	4.56	0.02	0.00	.44	**1728.00**
.5	1.08	1.93	0.69	5.28	6.88	3.28	0.02	0.00	.65	**1697.91**
.8	1.04	2.22	0.91	6.82	9.28	4.86	0.02	0.00	.64	**1729.16**
.9	1.27	3.93	1.44	5.26	8.71	2.70	0.01	0.00	.89	1527.84

Note: Bolded AICc indicates that this model is the best-fitting model for those data.

Table B5.3.

Parameter estimates, proportion of variance accounted for (VAC), and Akaike's Information Criterion (AICc) for single ramped Gaussian for Pigeon 72.

p(food for colour)	*a*	*M*	*SD*	*s*	*b*	*VAC*	*AICc*
Both FI 2							
.1	5.80	0.94	1.93	0.01	0.04	.91	**1146.12**
.2	6.70	0.16	2.89	0.01	0.04	.78	1352.52
.5	4.72	1.33	1.87	0.01	0.00	.91	**1068.47**
.8	4.37	1.21	2.03	0.01	0.00	.91	**1082.15**
.9	4.71	1.24	1.42	0.01	0.00	.92	**845.55**
Both FI 8							
.1	2.94	8.53	2.39	0.02	0.03	.76	1434.66
.2	5.10	9.38	4.05	0.02	0.00	.83	**1667.77**
.5	5.02	8.65	3.63	0.01	0.00	.83	**1346.63**
.8	3.83	9.34	3.21	0.01	0.01	.84	**1342.60**
.9	2.84	9.39	2.48	0.01	0.15	.75	1323.26
Colour FI 2							
.1	2.54	10.47	2.56	0.01	0.30	.84	**1862.71**
.2	4.54	4.92	6.09	0.01	0.00	.17	**1395.39**
.5	13.33	14.67	9.90	0.00	0.10	.72	**1769.30**
.8	8.59	9.44	6.86	0.01	0.00	.76	**1711.56**
.9	7.91	0.00	6.86	0.00	0.42	.56	1779.67
Colour FI 8							
.1	16.28	1.42	7.28	0.02	0.00	.61	1940.23
.2	12.00	5.70	7.24	0.02	0.00	.54	**1938.50**
.5	7.57	5.76	5.20	0.01	0.00	.68	**1392.28**
.8	1.73	11.73	2.92	0.01	0.30	.47	**1758.65**
.9	2.74	8.39	2.62	0.00	0.27	.67	**1453.87**

Note: Bolded AICc indicates that this model is the best-fitting model for those data.

Table B5.4.

Parameter estimates, proportion of variance accounted for (VAC), and Akaike's Information Criterion (AICc) for summed ramped Gaussian for Pigeon 72.

p(food for colour)	a_1	M_1	SD_1	a_2	M_2	SD_2	*s*	*b*	*VAC*	*AICc*
					Both FI 2					
.1	2.03	0.74	1.18	2.24	1.80	1.67	0.00	0.14	.92	1156.88
.2	1.83	0.89	0.76	2.32	3.01	2.15	0.01	0.02	.86	**1352.92**
.5	0.30	0.54	0.20	3.97	1.70	1.73	0.01	0.00	.97	1069.71
.8	0.31	0.68	0.17	3.70	1.57	1.92	0.01	0.00	.95	1085.35
.9	3.85	0.73	0.01	3.97	1.49	1.36	0.01	0.00	.96	846.29
					Both FI 8					
.1	4.52	8.84	2.81	10.10	26.14	6.06	0.00	0.04	.87	**1427.87**
.2	5.11	9.38	4.05	0.00	48.20	13.61	0.02	0.00	.83	1673.80
.5	5.79	9.23	3.77	9.08	30.70	6.86	0.00	0.04	.85	1351.68
.8	3.83	9.34	3.21	2.52	48.71	2.64	0.01	0.01	.84	1348.64
.9	3.59	9.60	2.83	2.11	23.92	2.86	0.00	0.15	.84	**1320.61**
					Colour FI 2					
.1	0.19	0.71	0.644	2.94	10.35	2.84	0.02	0.24	.84	1867.90
.2	4.14	4.63	6.07	0.41	7.50	5.59	0.01	0.00	.17	1401.43
.5	1.28	4.84	3.29	4.55	14.77	5.24	0.01	0.20	.74	1774.78
.8	0.96	5.45	2.16	3.95	11.85	4.56	0.01	0.18	.77	1716.77
.9	1.46	1.40	0.80	2.64	6.84	3.93	0.00	0.43	.76	**1774.58**
					Colour FI 8					
.1	1.35	1.05	0.83	3.72	6.24	2.98	0.00	0.40	.70	**1939.02**
.2	8.10	3.68	6.14	3.52	10.97	5.80	0.02	0.00	.54	1944.51
.5	5.24	3.86	3.96	1.96	10.66	2.33	0.01	0.00	.72	1392.86
.8	1.32	2.38	1.67	4.08	10.90	4.23	0.02	0.00	.52	1759.05
.9	1.00	1.21	1.63	4.71	8.01	3.37	0.01	0.00	.72	1454.81

Note: Bolded AICc indicates that this model is the best-fitting model for those data.

Table B5.5.

Parameter estimates, proportion of variance accounted for (VAC), and Akaike's Information Criterion (AICc) for single ramped Gaussian for Pigeon 73.

p(food for colour)	*a*	*M*	*SD*	*s*	*b*	*VAC*	*AICc*
			Both FI 2				
.1	1.91	1.74	0.79	0.00	0.08	.84	**974.32**
.2	2.88	1.81	1.07	0.00	0.10	.88	**1180.30**
.5	2.34	1.65	0.87	0.01	0.09	.77	1268.46
.8	1.76	1.97	1.07	0.00	0.09	.87	**991.92**
.9	2.89	1.58	1.09	0.01	0.06	.90	**1135.22**
			Both FI 8				
.1	2.95	9.51	3.64	0.01	0.03	.78	**1264.39**
.2	2.54	9.26	4.21	0.01	0.00	.64	**1247.87**
.5	5.40	8.72	3.24	0.02	0.00	.93	**1603.05**
.8	2.97	9.31	3.29	0.01	0.00	.79	**1107.73**
.9	5.20	8.64	4.05	0.02	0.00	.84	**1612.31**
			Colour FI 2				
.1	4.88	6.84	4.34	0.02	0.00	.73	**1529.82**
.2	5.40	8.17	5.44	0.02	0.00	.63	**1546.45**
.5	7.06	7.11	7.44	0.01	0.00	.33	**1555.96**
.8	4.34	5.16	8.27	0.01	0.00	.12	1106.08
.9	1.26	1.70	1.18	0.00	0.29	.55	**1459.62**
			Colour FI 8				
.1	1.40	3.42	2.00	0.00	0.31	.41	**1526.81**
.2	0.89	1.25	0.80	0.00	0.32	.59	**1577.99**
.5	9.91	6.51	6.98	0.02	0.00	.32	1870.18
.8	7.23	7.62	5.78	0.02	0.00	.38	**1710.85**
.9	7.66	6.62	4.60	0.03	0.00	.57	1943.43

Note: Bolded AICc indicates that this model is the best-fitting model for those data.

Table B5.6.

Parameter estimates, proportion of variance accounted for (VAC), and Akaike's Information Criterion (AICc) for summed ramped Gaussian for Pigeon 73.

p(food for colour)	a_1	M_1	SD_1	a_2	M_2	SD_2	*s*	*b*	*VAC*	*AICc*
					Both FI 2					
.1	1.52	1.51	0.68	0.39	2.66	0.23	0.00	0.08	.89	976.49
.2	0.47	1.64	1.10	2.41	1.84	1.06	0.00	0.10	.88	1186.34
.5	2.64	1.65	0.91	3.41	14.60	6.14	0.00	0.00	.81	**1262.48**
.8	0.34	0.72	0.30	1.37	2.33	0.84	0.00	0.09	.92	994.25
.9	0.66	0.48	0.36	2.08	2.03	0.78	0.01	0.07	.95	1136.54
					Both FI 8					
.1	2.96	9.52	3.65	2.57	78.24	30.36	0.01	0.03	.78	1270.42
.2	2.54	9.26	4.21	2.40	115.44	17.69	0.01	0.00	.64	1253.91
.5	5.40	8.72	3.24	2.37	112.66	19.93	0.02	0.00	.93	1609.08
.8	2.97	9.31	3.29	2.40	114.20	18.09	0.01	0.00	.79	1113.77
.9	5.20	8.64	4.05	2.38	108.29	22.33	0.02	0.00	.84	1618.34
					Colour FI 2					
.1	0.49	1.93	0.83	4.08	7.64	3.57	0.02	0.00	.79	1531.02
.2	1.83	5.10	2.83	4.17	11.58	5.55	0.01	0.00	.70	1549.38
.5	1.03	2.63	1.60	4.58	9.37	5.35	0.01	0.00	.39	1557.86
.8	0.56	2.12	0.97	2.47	8.57	5.36	0.01	0.00	.33	**1105.85**
.9	1.87	1.58	1.37	4.17	8.10	6.54	0.01	0.00	.60	1461.36
					Colour FI 8					
.1	1.33	2.45	1.43	2.41	6.56	4.16	0.01	0.10	.48	1528.59
.2	1.04	1.25	0.86	0.45	11.13	1.79	0.00	0.28	.66	1581.69
.5	1.36	2.41	1.29	6.62	8.71	5.10	0.02	0.00	.42	**1867.85**
.8	1.16	3.36	1.96	5.71	9.22	5.08	0.02	0.00	.40	1714.39
.9	1.21	3.22	1.29	6.32	7.92	4.35	0.02	0.00	.66	**1941.96**

Note: Bolded AICc indicates that this model is the best-fitting model for those data.

Table B5.7.

Parameter estimates, proportion of variance accounted for (VAC), and Akaike's Information Criterion (AICc) for single ramped Gaussian for Pigeon 74.

p(food for colour)	*a*	*M*	*SD*	*s*	*b*	*VAC*	*AICc*
Both FI 2							
.1	3.79	1.84	1.61	0.01	0.00	.86	**1056.42**
.2	1.25	1.00	0.64	0.00	0.26	.77	1330.42
.5	3.20	1.78	1.82	0.01	0.04	.88	**995.38**
.8	2.88	1.41	1.74	0.00	0.22	.74	1339.56
.9	3.27	1.67	1.40	0.01	0.05	.85	**1125.39**
Both FI 8							
.1	9.22	7.23	6.20	0.02	0.00	.55	**1861.43**
.2	13.31	7.85	8.05	0.02	0.00	.55	**1921.83**
.5	10.35	7.33	6.45	0.02	0.00	.58	**1894.34**
.8	9.23	7.47	5.81	0.03	0.00	.73	**1985.48**
.9	6.37	7.38	4.28	0.03	0.00	.55	**1793.94**
Colour FI 2							
.1	12.36	5.61	6.30	0.02	0.00	.63	1891.66
.2	15.09	3.53	8.55	0.02	0.00	.28	**1967.47**
.5	17.76	2.47	8.60	0.02	0.00	.54	**1979.11**
.8	1.38	1.32	0.82	0.00	0.50	.53	1830.53
.9	1.26	1.30	1.26	0.00	0.51	.46	**1796.50**
Colour FI 8							
.1	0.88	1.57	0.79	0.00	0.52	.47	**1807.96**
.2	0.98	1.20	0.56	0.00	0.57	.52	1889.85
.5	15.32	3.83	8.12	0.02	0.00	.53	**1958.57**
.8	18.73	5.76	10.36	0.02	0.00	.28	2059.65
.9	9.32	5.86	5.37	0.03	0.00	.61	**1875.58**

Note: Bolded AICc indicates that this model is the best-fitting model for those data.

Table B5.8.

Parameter estimates, proportion of variance accounted for (VAC), and Akaike's Information Criterion (AICc) for summed ramped Gaussian for Pigeon 74.

p(food for colour)	a_1	M_1	SD_1	a_2	M_2	SD_2	*s*	*b*	*VAC*	*AICc*
					Both FI 2					
.1	3.78	1.84	1.61	0.14	20.93	12.38	0.01	0.00	.86	1062.45
.2	0.96	0.82	0.42	0.59	2.55	1.14	0.00	0.24	.86	**1327.46**
.5	3.41	1.76	1.87	0.41	11.19	6.22	0.01	0.00	.88	1001.30
.8	1.46	1.00	0.55	1.00	3.37	1.05	0.00	0.22	.90	**1332.14**
.9	3.62	1.68	1.50	1.34	14.66	4.74	0.01	0.00	.87	1126.53
					Both FI 8					
.1	0.71	2.48	1.24	7.70	8.25	5.41	0.02	0.00	.58	1864.78
.2	1.25	3.11	1.93	10.76	9.71	7.02	0.02	0.00	.60	1925.15
.5	3.64	4.35	3.07	7.11	11.79	5.91	0.02	0.00	.67	1895.54
.8	0.00	1.72	6.01	9.23	7.47	5.81	0.03	0.00	.73	1991.52
.9	0.00	2.13	11.68	6.37	7.38	4.28	0.03	0.00	.55	1799.98
					Colour FI 2					
.1	1.78	2.08	1.36	8.51	7.91	4.76	0.02	0.00	.73	**1890.57**
.2	0.80	1.69	0.78	10.66	6.06	6.74	0.03	0.00	.42	1968.31
.5	0.75	1.36	0.77	12.27	5.15	6.65	0.02	0.00	.61	1981.98
.8	0.72	0.72	0.16	1.04	2.05	0.96	0.00	0.49	.66	**1829.10**
.9	1.90	1.10	1.37	7.03	5.35	6.87	0.02	0.00	.49	1800.97
					Colour FI 8					
.1	1.64	1.53	0.98	6.70	7.15	6.57	0.02	0.00	.57	1809.76
.2	0.86	1.07	0.47	1.10	4.15	2.04	0.00	0.52	.64	**1889.05**
.5	0.70	1.41	0.68	10.97	5.91	6.28	0.03	0.00	.64	1960.30
.8	2.70	2.79	2.03	10.33	10.32	6.90	0.02	0.00	.51	**2057.87**
.9	2.23	3.04	1.93	6.29	8.28	4.59	0.03	0.00	.71	1875.86

Note: Bolded AICc indicates that this model is the best-fitting model for those data.

Table B5.9.

Parameter estimates, proportion of variance accounted for (VAC), and Akaike's Information Criterion (AICc) for single ramped Gaussian for Pigeon 75.

p(food for colour)	*a*	*M*	*SD*	*s*	*b*	*VAC*	*AICc*
			Both FI 2				
.1	4.38	1.34	0.94	0.00	0.00	.93	**651.03**
.2	4.47	1.32	0.85	0.00	0.00	.96	**491.69**
.5	5.10	1.34	1.09	0.00	0.00	.97	**703.29**
.8	4.78	1.34	0.88	0.01	0.00	.94	**953.45**
.9	4.28	1.34	0.84	0.00	0.03	.95	587.43
			Both FI 8				
.1	8.77	7.90	2.43	0.02	0.06	.96	**1762.62**
.2	8.79	7.55	2.84	0.01	0.18	.93	1965.28
.5	8.72	7.66	2.34	0.03	0.02	.88	1929.70
.8	7.45	7.30	2.18	0.04	0.07	.89	2167.08
.9	7.29	7.03	1.83	0.03	0.01	.90	1769.19
			Colour FI 2				
.1	12.92	5.43	3.31	0.02	0.00	.80	1914.00
.2	12.16	4.77	3.00	0.02	0.00	.79	1848.17
.5	11.97	5.05	3.89	0.01	0.00	.61	1866.18
.8	1.69	1.29	0.52	0.00	0.70	.33	2256.22
.9	1.80	1.44	0.53	0.00	0.47	.42	1870.40
			Colour FI 8				
.1	2.36	1.32	0.67	0.00	0.59	.62	2022.02
.2	1.82	1.15	0.64	0.00	0.74	.30	2220.71
.5	17.95	3.71	5.44	0.01	0.00	.65	2012.64
.8	15.29	4.60	3.93	0.01	0.00	.86	1884.60
.9	14.69	5.01	3.52	0.01	0.00	.93	1676.91

Note: Bolded AICc indicates that this model is the best-fitting model for those data.

Table B5.10.

Parameter estimates, proportion of variance accounted for (VAC), and Akaike's Information Criterion (AICc) for summed ramped Gaussian for Pigeon 75.

p(food for colour)	a_1	M_1	SD_1	a_2	M_2	SD_2	*s*	*b*	*VAC*	*AICc*
Both FI 2										
.1	1.11	0.98	0.98	3.28	1.45	0.91	0.00	0.00	.93	656.87
.2	1.34	1.04	0.88	3.15	1.43	0.82	0.00	0.00	.96	497.49
.5	1.39	0.97	1.24	3.76	1.44	1.04	0.00	0.00	.97	709.18
.8	1.27	0.96	0.97	3.56	1.45	0.83	0.01	0.00	.94	959.14
.9	1.12	0.52	0.39	3.08	1.68	0.65	0.00	0.03	.98	**586.56**
Both FI 8										
.1	9.48	7.95	2.53	10.70	27.28	9.91	0.00	0.06	.96	1767.37
.2	7.89	7.14	2.63	8.63	14.89	7.78	0.00	0.06	.96	**1960.77**
.5	10.72	7.98	2.67	6.60	23.39	3.20	0.01	0.02	.97	**1893.58**
.8	10.21	7.59	2.50	11.65	20.97	4.76	0.00	0.08	.92	**2160.93**
.9	8.95	7.18	2.05	7.52	21.48	3.94	0.00	0.01	.97	**1730.14**
Colour FI 2										
.1	3.74	2.44	1.32	8.70	7.07	2.41	0.02	0.00	.86	**1898.72**
.2	3.85	2.08	1.12	7.76	6.46	1.95	0.02	0.00	.89	**1819.01**
.5	2.70	1.49	0.64	7.39	6.84	2.07	0.01	0.04	.94	**1785.85**
.8	2.91	1.27	0.66	4.09	7.70	1.33	0.02	0.23	.82	**2156.81**
.9	3.23	1.37	0.75	4.75	8.61	1.99	0.02	0.00	.81	**1756.48**
Colour FI 8										
.1	3.08	1.22	0.75	10.01	7.85	5.25	0.01	0.00	.89	**1967.15**
.2	2.05	1.03	0.65	14.54	6.73	5.32	0.02	0.00	.88	**2093.48**
.5	3.79	108	0.82	9.57	7.08	2.82	0.02	0.00	.94	**1953.20**
.8	4.57	1.61	1.39	9.32	6.86	2.47	0.01	0.00	.96	**1863.71**
.9	3.57	1.93	1.28	10.23	6.55	2.62	0.01	0.00	.97	**1667.40**

Note: Bolded AICc indicates that this model is the best-fitting model for those data.

Table B5.11.

Parameter estimates, proportion of variance accounted for (VAC), and Akaike's Information Criterion (AICc) for single ramped Gaussian for Pigeon 76.

p(food for colour)	*a*	*M*	*SD*	*s*	*b*	*VAC*	*AICc*
Both FI 2							
.1	4.04	1.50	1.00	0.01	0.00	.93	**922.49**
.2	4.06	1.47	1.12	0.01	0.00	.90	1190.97
.5	4.17	1.44	1.09	0.01	0.00	.90	**986.96**
.8	4.43	1.57	1.06	0.02	0.00	.89	**1451.42**
.9	3.82	1.23	1.10	0.01	0.00	.89	**922.58**
Both FI 8							
.1	9.48	9.01	3.04	0.02	0.03	.94	**1932.93**
.2	6.16	8.40	3.03	0.02	0.12	.83	1941.22
.5	7.52	9.30	3.30	0.03	0.04	.94	**2010.51**
.8	8.11	7.80	2.68	0.02	0.00	.91	**1839.76**
.9	6.46	7.26	2.36	0.03	0.00	.85	1864.29
Colour FI 2							
.1	10.69	6.28	3.02	0.03	0.00	.78	2006.88
.2	10.91	5.37	4.53	0.03	0.00	.42	**2195.22**
.5	13.44	5.01	4.55	0.02	0.00	.54	2088.93
.8	2.07	1.62	0.80	0.00	0.54	.68	2022.79
.9	2.26	1.38	0.76	0.01	0.38	.69	1970.95
Colour FI 8							
.1	2.71	1.46	0.76	0.01	0.40	.80	1893.55
.2	2.06	1.64	0.89	0.00	0.58	.64	1982.11
.5	5.84	5.68	4.69	0.00	0.59	.47	2180.10
.8	9.89	6.03	3.64	0.03	0.00	.90	**2009.28**
.9	8.16	5.25	3.13	0.03	0.00	.90	**1927.12**

Note: Bolded AICc indicates that this model is the best-fitting model for those data.

Table B5.12.

Parameter estimates, proportion of variance accounted for (VAC), and Akaike's Information Criterion (AICc) for summed ramped Gaussian for Pigeon 76.

P(food for colour)	a_1	M_1	SD_1	a_2	M_2	SD_2	s	b	VAC	$AICc$
Both FI 2										
.1	4.06	1.50	1.00	2.87	24.20	8.62	0.00	0.00	.93	926.55
.2	4.18	1.49	1.15	4.09	17.96	5.50	0.00	0.00	.92	**1183.36**
.5	4.18	1.44	1.09	0.47	23.26	1.28	0.01	0.00	.91	987.72
.8	4.48	1.58	1.07	4.32	17.63	6.32	0.01	0.00	.89	1452.32
.9	3.83	1.23	1.11	1.83	24.54	5.15	0.00	0.01	.89	922.61
Both FI 8										
.1	11.69	9.32	3.44	2.86	22.25	2.85	0.01	0.01	.96	1934.55
.2	8.19	8.86	3.53	5.13	24.41	3.77	0.01	0.11	.89	**1938.61**
.5	8.55	9.13	3.39	12.98	22.38	8.44	0.00	0.02	.94	2015.67
.8	8.49	7.88	2.76	0.65	21.40	1.19	0.02	0.00	.92	1841.86
.9	7.86	7.61	2.65	2.87	20.01	2.25	0.02	0.00	.91	**1854.47**
Colour FI 2										
.1	7.81	5.15	2.41	2.80	9.50	1.34	0.03	0.00	.83	**1998.12**
.2	1.28	2.31	1.19	8.89	6.41	4.06	0.03	0.00	.46	2196.49
.5	1.98	1.41	0.67	9.93	6.74	3.44	0.02	0.00	.69	**2063.14**
.8	3.11	1.60	0.94	2.14	7.06	2.29	0.02	0.22	.75	**2019.13**
.9	3.67	1.38	0.97	2.67	8.50	2.29	0.03	0.00	.84	**1950.97**
Colour FI 8										
.1	4.14	1.43	0.92	2.72	8.39	2.26	0.02	0.00	.91	**1872.83**
.2	3.44	1.53	1.07	5.14	7.82	3.39	0.02	0.10	.87	**1958.64**
.5	1.67	1.40	0.63	4.02	7.37	2.26	0.01	0.47	.89	**2153.81**
.8	4.37	3.62	2.19	5.16	8.28	2.52	0.03	0.00	.92	2012.65
.9	6.87	4.64	2.68	1.16	9.04	1.46	0.03	0.00	.92	1931.04

Note: Bolded AICc indicates that this model is the best-fitting model for those data.

CHAPTER 6

Experiment 4: Reinforcer Predictability and Divided Control by Temporal and Non-Temporal Stimuli[14]

6.1 Introduction

When elapsed time and other non-temporal stimuli both signal future reinforcer availability, they form a compound stimulus. As described in Chapter 2 of this thesis, previous research suggests that the division of control between temporal and non-temporal stimuli depends on the relative ability of each stimulus to predict future reinforcers. For example, Davison and Elliffe (2010) found that the division of control between the duration of presentation (a temporal stimulus) and the speed of keylight-colour alternation (a non-temporal stimulus) in a delayed matching-to-sample (DMTS) task depended on the relative reinforcer rate associated with each stimulus dimension.

Likewise, when reinforcer availability changes across time since a marker event and additional non-temporal stimuli signal future reinforcer availability, the extent of control by elapsed time and by non-temporal stimuli appears to depend on the relative reliability with which each stimulus signals future reinforcers. In Experiment 3 of the present thesis (Chapter 5), behaviour was jointly controlled by the compound-stimulus dimensions (colour and flash-frequency) and by elapsed time since trial start in a multiple peak procedure. Control by the compound-stimulus dimensions appeared to be weaker and control by elapsed time appeared to be stronger at times when elapsed time signalled the likely time of the next reinforcer with

---

[14] This experiment was conducted under Approval 001396 granted by the University of Auckland Animal Ethics Committee.
This chapter is an edited version of the publication Gomes-Ng, S., Elliffe, D., & Cowie, S. (2018a). Environment tracking and signal following in a reinforcer-ratio reversal procedure. *Behavioural Processes, 157,* 208-224. It has been edited to be consistent with the language used and overall aims of the present thesis. As per Elsevier's copyright policy, the author reserves the right to include the publication in this thesis.

greater certainty (see Figures 5.1 and 5.3; see also Aum, Brown, & Hemmes, 2007; McMillan & Roberts, 2013). Similarly, when the location of reinforcers changes at a fixed point in each experimental session (mid-session reversal, MSR, task), the strength of control by elapsed time and by additional non-temporal stimuli that signal the next-reinforcer location depends on the relative reliability of each stimulus; non-temporal stimuli exert stronger control at times when elapsed time signals the next-reinforcer location less reliably (e.g., McMillan, Sturdy, Pisklak, & Spetch, 2016; Rayburn-Reeves, Qadri, Brooks, Keller, & Cook, 2017; Smith, Beckmann, & Zentall, 2017; see also Cowie, Davison, Blumhardt, & Elliffe, 2016c).

However, in contrast, a recent study by Cowie, Davison, and Elliffe (2017) suggests that changes in the reliability of elapsed time do not always determine the division of control between elapsed time and non-temporal stimuli. In Cowie et al.'s experiment, the distribution of reinforcers between two response alternatives (the *reinforcer ratio*) changed across time after reinforcer deliveries: At early times after a reinforcer, the next reinforcer was more likely to occur on the not-just-reinforced alternative, whereas at later times, the next reinforcer was more likely to occur on the just-reinforced alternative. Because the change in reinforcer ratio depended on the last-reinforcer location, the reliability with which elapsed time signalled future reinforcer availability depended on the extent of behavioural control by last-reinforcer location. Such control declined across time (see Cowie, Davison, & Elliffe, 2011, 2017), and hence elapsed time signalled future reinforcer availability more reliably at early times and less reliably at later times after a reinforcer. Despite this, keylight-colour stimuli that signalled the definite location of the next reinforcer exerted similar, and in some cases slightly *weaker*, control over choice at later times than at early times after a reinforcer delivery. That is, the division of control between elapsed time and informative keylight-colour stimuli appeared not to depend on the reliability of elapsed time. In contrast, such division did depend on the reliability of keylight-colour stimuli; keylight colour exerted stronger control and elapsed time exerted weaker control when keylight colour signalled the next-reinforcer location more reliably. Therefore, the reliability of temporal and non-temporal stimuli had asymmetrical effects on divided stimulus control – the reliability of non-temporal stimuli (keylight colour), but not of elapsed time, determined the division of control between elapsed time and keylight-colour stimuli.

Cowie et al.'s (2017) findings appear at odds with other research (e.g., Aum et al., 2007; Rayburn-Reeves et al., 2017) showing that the division of control between elapsed time and non-temporal stimuli depends on the reliability with which each stimulus signals future reinforcer availability. This discrepancy may be related to several procedural differences. Cowie et al. arranged a gradual change in reinforcer availability, whereas the change in reinforcer availability was abrupt in other studies – in the peak procedure, the probability of a reinforcer delivery is 0 before the fixed interval and 1 after that interval has elapsed, and in the mid-session reversal task, the location of reinforcers changes abruptly at a fixed point in each session. Additionally, in Cowie et al.'s procedure, the change in reinforcer availability occurred after a reinforcer delivery, and depended on the location of the previous reinforcer. In contrast, in other studies, the change in reinforcer availability occurred after trial start (the peak procedure; e.g., Aum et al., 2007; Experiment 3 of this thesis) or session start (mid-session reversal task; e.g., Rayburn-Reeves et al., 2017; Smith et al., 2017), and hence did not depend on last-reinforcer location. Alternatively, it is possible that relative reinforcer predictability does not determine divided control between elapsed time and other non-temporal stimuli when reinforcer availability changes across time since a reinforcer delivery.

At present, little is known about how relative reinforcer predictability divides control between elapsed time and other non-temporal stimuli in environments in which reinforcer availability changes across seconds since a marker event. Furthermore, in previous studies demonstrating divided control between elapsed time and non-temporal stimuli (e.g., Cowie et al., 2017; Rayburn-Reeves et al., 2017; Experiment 3 of this thesis), changes in the reliability of elapsed time occurred as a result of the experimenter-arranged change in reinforcer availability. Hence, changes in the reliability of elapsed time were correlated with changes in elapsed time itself. For example, in the MSR task, elapsed time signals future reinforcer availability more reliably at early and late times, and less reliably towards the midpoint of the session, due to the arranged reversal in reinforcer location. Likewise, in Cowie et al., elapsed time signalled future reinforcer availability less reliably at later times. Such a correlation between changes in elapsed time and changes in the reliability of elapsed time makes it difficult to determine the extent to which the latter contributed to divided stimulus control.

Therefore, the present experiment was designed to examine more thoroughly how relative reinforcer predictability determines divided control between temporal and non-temporal stimuli in a concurrent schedule in which the reinforcer ratio changed across time after each reinforcer delivery. In the present experiment, reinforcer deliveries were more likely to occur on one alternative during the first 30 s after a reinforcer delivery, and thereafter reinforcer deliveries were more likely to occur on the other alternative (see Cowie, Elliffe, & Davison, 2013). Thus, as in Cowie et al. (2017), elapsed time since a reinforcer delivery signalled the likely location of the next reinforcer in the present experiment. However, unlike Cowie et al., the change in reinforcer ratio was abrupt and did not depend on the last-reinforcer location in the present experiment, hence making the change in reinforcer availability more similar to the abrupt change that occurs in the peak procedure or mid-session reversal task.

Across conditions, the probability of reinforcer deliveries on the locally richer key varied. As a result, elapsed time signalled the likely next-reinforcer location more reliably when the probability of reinforcer deliveries on the locally richer key was higher. Thus, unlike previous studies, changes in the reliability of elapsed time were not correlated with changes in elapsed time in the present experiment. In addition to manipulating the reliability of elapsed time, we also manipulated the reliability with which keylight-colour stimuli signalled the time or location of the next reinforcer. In one set of conditions, no keylight-colour stimuli were arranged, whereas in other conditions, keylight-colour stimuli signalled, in order of increasing reliability: (1) the location of the locally richer alternative; (2) whether the next reinforcer would occur before or after the reversal in reinforcer location, and hence by extension its likely location; or (3) the definite location of the next reinforcer, and hence by extension its likely time.

6.2 Method

6.2.1 Subjects and Apparatus

Six experimentally naïve pigeons, numbered 41 to 46, served as subjects. The pigeons were kept under the same housing and deprivation conditions as the pigeons in previous experiments in this thesis (see Experiment 1, Chapter 3). The experimental apparatus was similar to that used in previous experiments in this thesis, and differed only in the following aspects. The two wooden perches were perpendicular to each other; one perch was parallel to and 95 mm away from the right wall, and the other perch was parallel to and 95 mm away from the front wall. The operant panel was mounted on the right wall of the cage, 280 mm above the cage floor. The operant panel contained one row of four 20-mm diameter response keys, spaced 100 mm apart from centre to centre. Only the centre-left and centre-right keys (hereafter, left and right keys, respectively) were used; these could be illuminated yellow or red. The magazine aperture was 90 mm below the response keys. During a reinforcer delivery, the hopper was raised for 3 s.

6.2.2 Procedure

6.2.2.1 Pretraining

The pigeons were magazine trained and autoshaped using the same procedure that was used to pretrain Pigeons 71 to 76 in Experiment 1 (see Chapter 3). The experiment proper began once all pigeons were pecking on a variable-interval (VI) 45-s schedule.

6.2.2.2 Experimental Sessions

Sessions began at 2:00 a.m. daily for all pigeons, and ran for 60 min or until 60 reinforcers had been delivered, whichever occurred first. Each session began with the illumination of the left and right keylights, signalling the availability of a VI schedule on each key. Reinforcers were scheduled dependently (Stubbs & Pliskoff, 1969), and were arranged by interrogating a probability gate set at $p = .022$ every second (i.e., a single exponential VI 45-s schedule). The minimum arranged inter-reinforcer interval (IRI) was 1 s, and the maximum arranged IRI was 453 s. No changeover delay (Herrnstein, 1961) was arranged.

After a reinforcer was arranged by the VI schedule, it was allocated to the left or right key with a probability that depended on the time since the previous reinforcer, or since session start for the first reinforcer of the session. One key, the *higher-to-lower* (H→L) key, delivered reinforcers with a higher probability during the first 30 s after the previous reinforcer (or session start), and with a lower probability after 30 s. The reverse was arranged on the other key (the *lower-to-higher*, L→H, key). For example, in Condition 1, reinforcers arranged within the first 30 s after the previous reinforcer were allocated to the L→H key with a probability of .1, and reinforcers arranged after 30 s were allocated to the L→H key with a probability of .9. The middle panel of Figure 6.1 shows the arranged proportion of reinforcers on the L→H key across time (dashed line) since the most recent reinforcer for

Condition 1. The number of reinforcers arranged before and after the reversal was roughly equal (see Cowie et al., 2013).

Thus, if the IRI was longer than 30 s, the reinforcer ratio reversed to its reciprocal at 30 s after the previous reinforcer. Occasionally, a reinforcer arranged on a key before the reversal was not obtained before the reversal occurred. To ensure that reinforcers arranged on the locally richer key were always obtained from the locally richer key and likewise that reinforcers arranged on the locally leaner key were always obtained from the locally leaner key, the location of any arranged but not obtained reinforcers changed to the other key at the time of the reinforcer-ratio reversal. This maintained the abrupt change in the reinforcer ratio.

Table 6.1 shows the sequence of conditions. The location of the H→L key alternated between conditions. The probability of reinforcer deliveries on the locally richer key also varied across conditions. In Conditions 1 to 6, the locally richer key was nine times more likely to deliver reinforcers than the locally leaner key (9x-Richer conditions; Figure 6.1, middle panel). In Conditions 7 to 10, the locally richer key was three times more likely to deliver reinforcers (3x-Richer conditions; Figure 6.1, top panel), and in Conditions 11 to 13, all reinforcers were arranged on the locally richer key (All-on-Richer conditions; Figure 6.1, bottom panel). As the probability of reinforcer deliveries on the locally richer key increased, elapsed time since a reinforcer signalled the likely location of the next reinforcer with greater certainty.

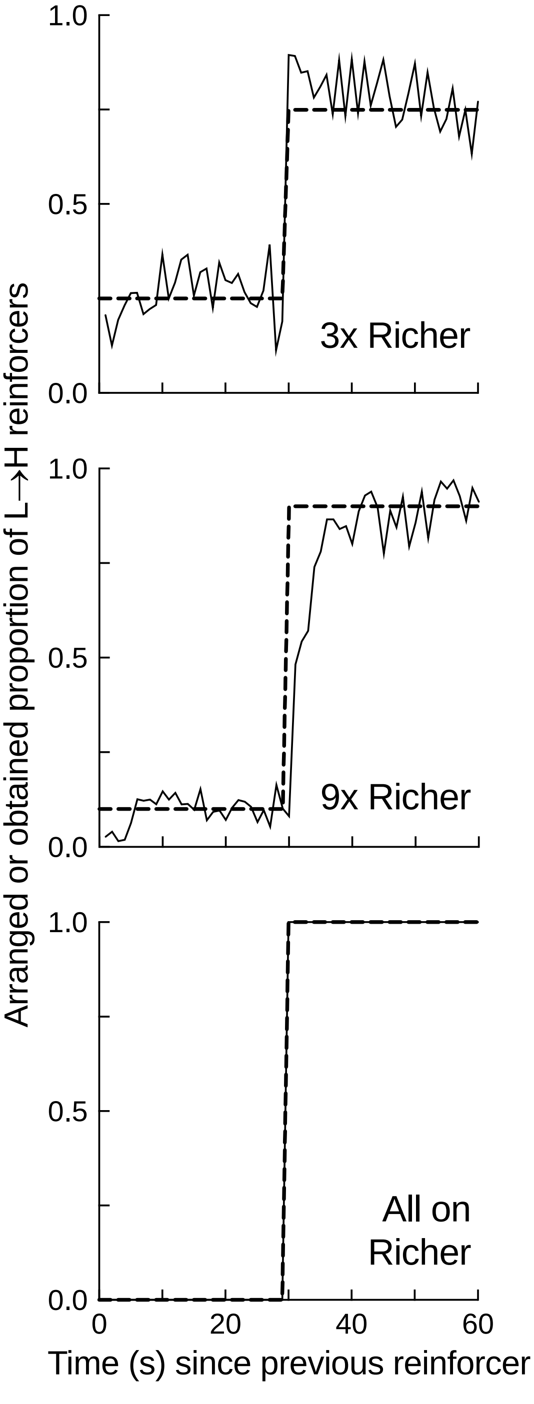

Figure 6.1. Arranged (dashed lines) and mean (group sum) obtained (solid lines) proportion of reinforcers on the L→H key as a function of time, in 1-s bins, since the previous reinforcer. Obtained data are from No-Signal conditions (Conditions 7, 1, and 11).

Table 6.1.

Condition number, location of H→L key, probability of reinforcers on the locally richer key, and stimulus properties for all conditions. The reinforcer-ratio reversal occurred at 30 s in all conditions.

Condition	H→L key	*p*(reinforcer on locally richer)	Condition type	Description of stimuli
1	Left	.9	No Signal	Both keys yellow
2	Right	.9	No Signal	Both keys yellow
3	Left	.9	Signal Richer	Both keys yellow before reversal, red after reversal
4	Right	.9	Signal Time	Both keys yellow if next reinforcer before reversal, red if next reinforcer after reversal
5	Left	.9	Signal Location	Both keys yellow if next reinforcer on left, red if next reinforcer on right
6	Right	.9	Signal Location	Both keys yellow if next reinforcer on left, red if next reinforcer on right
7	Left	.75	No Signal	Both keys yellow
8	Right	.75	Signal Richer	Both keys yellow before reversal, red after reversal
9	Left	.75	Signal Time	Both keys yellow if next reinforcer before reversal, red if next reinforcer after reversal
10	Right	.75	Signal Location	Both keys yellow if next reinforcer on left, red if next reinforcer on right
11	Left	1.0	No Signal	Both keys yellow
12	Right	1.0	Signal Richer	Both keys yellow before reversal, red after reversal
13	Left	1.0	Signal Time and Location	Both keys yellow if next reinforcer before reversal (i.e., on left), red if next reinforcer after reversal (i.e., on right)

Across conditions (Table 6.1), the colour of the keys provided either no information about the next reinforcer ('No-Signal' conditions), or signalled the locally richer key ('Signal-Richer' conditions), the time of the next reinforcer ('Signal-Time' conditions), or the location of the next reinforcer ('Signal-Location' conditions). In No-Signal conditions, both keys were lit yellow at all times, and hence only elapsed time since the previous reinforcer signalled the current reinforcer ratio. In Signal-Richer conditions, both keys were lit yellow before the reinforcer-ratio reversal and were lit red after the reversal. Hence, both elapsed time and keylight colour signalled the current reinforcer ratio. In Signal-Time conditions, both keys were lit yellow if the next reinforcer was arranged before the reversal, and were lit red if the next reinforcer was arranged after the reversal. In Signal-Location conditions, both keys were lit yellow if the next reinforcer was arranged on the left key, and were lit red if the next reinforcer was arranged on the right key. Although the location and time of reinforcers arranged but not obtained before the reversal changed when the reinforcer ratio reversed, the colour of the keys did not change in Signal-Time and Signal-Location conditions. In All-on-Richer conditions, stimuli that signalled whether the next reinforcer occurred before or after the reversal also signalled the location of that reinforcer (the H→L key before the reversal, and the L→H key after the reversal). Hence, we ran only three All-on-Richer conditions (see Table 6.1).

In order to signal the time and/or location of the next reinforcer in Signal-Time and Signal-Location conditions, the time and/or location of the next reinforcer was pre-determined using a method that was essentially the same as the method used to arrange reinforcers in No-Signal conditions. The first 10 IRIs were computed and stored in a list before each session. In Signal-Time conditions, the location of the next reinforcer was determined when its arranged time was reached, in the same way as in No-Signal conditions. In Signal-Location conditions, the location of the first 10 reinforcers was also pre-determined. The location of each arranged reinforcer depended on its arranged time; if the IRI was shorter than 30 s, the reinforcer location was chosen according to the reinforcer ratio before the reversal, whereas if the IRI was longer than 30 s, the reinforcer location was chosen according to the reinforcer ratio after the reversal. During sessions, a counter kept track of the time (Signal-Time conditions) or time and location (Signal-Location conditions) of the next reinforcer. After a reinforcer delivery, the counter was incremented, and a new IRI (and location, in Signal-Location conditions) replaced the previous interval. The counter was reset to zero whenever the end of the list was reached.

Conditions 1 and 2 lasted for 33 sessions each, and all subsequent conditions lasted for 30 sessions. Responding for all pigeons was stable, determined by visual inspection, by the 15th session of each condition. Hence, the last 15 sessions of each condition were used for analyses.

6.2.3 Data Analysis

Responses and reinforcers were aggregated according to the time at which they occurred, in 1-s bins, since the most recent reinforcer delivery. Events in each bin were then separated according to their location (the H→L or L→H key), and further separated according to keylight colour (see Table 6.1) in conditions in which keylight-colour stimuli signalled the time or location of the next reinforcer (Signal-Time and Signal-Location conditions). We used these response and reinforcer counts to calculate the proportion of responses made to and the proportion of reinforcers obtained on the L→H key in each time bin (*local choice* and *local reinforcer ratios*, respectively). We also calculated response rates on each key by dividing the number of responses made to each key in each bin by the number of times that bin was reached, and then multiplying by 60 to yield responses per min.

Because time bins with low response or reinforcer counts do not provide enough data to determine local choice or obtained reinforcer ratios reliably, such bins were excluded from analyses. For conditions in which both keys remained yellow (No-Signal conditions) or in which the colour of the keys changed at the time of the reversal (Signal-Richer conditions), bins with fewer than 20 responses or 5 reinforcers were excluded from individual-pigeon data. For group-sum data, bins with fewer than 100 responses or 20 reinforcers were excluded from analyses. For conditions in which keylight-colour stimuli signalled the next-reinforcer time or location (Signal-Time and Signal-Location conditions), data were separated according to keylight colour, and hence bins with fewer than 10 responses for individual-pigeon data or 50 responses for group-sum data were excluded from analyses. These criteria generally excluded the first one to four time bins after a reinforcer, likely because of the post-reinforcer pause, and because reinforcers were rarely arranged in the first few seconds after a reinforcer had been obtained.

6.3 Results

Although there were differences in overall levels of preference and in the degree of changes in choice between pigeons, response rates and local choice across time followed similar patterns for all pigeons. Thus, patterns of responding and of choice in group-sum data were generally representative of such patterns in individual-pigeon data. Therefore, group data are presented here, and individual-pigeon data are shown in Appendix Figures A6.1 to A6.4. In each type of condition (No Signal, Signal Richer, Signal Time, and Signal Location; see Table 6.1), responding followed a similar pattern, hence, the data were replicable across conditions.

6.3.1 Obtained versus Arranged Proportions of Reinforcers

In the present procedure, the arranged reinforcer ratio depended on elapsed time since the previous reinforcer (Figure 6.1): If less than 30 s had elapsed since the previous reinforcer, reinforcer deliveries were more likely on the H→L key, whereas if more than 30 s had elapsed, reinforcer deliveries were more likely on the L→H key. In all conditions and for all pigeons, obtained reinforcer ratios approximated arranged reinforcer ratios. The solid lines in Figure 6.1 show the mean (group sum) proportion of reinforcers obtained from the L→H key in each 1-s time bin since the most recent reinforcer for 3x-Richer, 9x-Richer, and All-on-Richer conditions in which both keys remained yellow (Conditions 7, 1, and 11, respectively). The obtained reinforcer ratios shown in Figure 6.1 are representative of obtained local reinforcer ratios in all other conditions, and for all pigeons.

Any reinforcers arranged but not obtained before the reinforcer-ratio reversal were reallocated to the other alternative at the time of the reversal. If such reallocation happened often, it would have caused an abrupt increase in the obtained reinforcer rate at the time of the reversal, and hence may have affected behaviour. To determine whether the reinforcer reallocation could have affected behaviour, we analysed the percentage of reallocated reinforcers for each pigeon in each condition, and also obtained reinforcer rates across time. In general, fewer than 10% of reinforcers arranged before the reversal were reallocated. Thus, reinforcers arranged before the reversal were almost always obtained before the reversal, and hence the reinforcer-location reallocation probably did not affect behaviour. The only exception was Pigeon 45, for whom there was an abrupt increase in obtained reinforcers just after the reversal in Conditions 5, 7, and 11. Even so, the reinforcer reallocation maintained the step-change in the reinforcer ratio, and therefore did not affect the *proportion* of reinforcers obtained from a key (i.e., it did not affect the reliability with which elapsed time signalled the likely next-reinforcer location).

6.3.2 Choice Without Keylight-Colour Stimuli

In conditions in which the keys were always yellow (No-Signal conditions), only elapsed time since a reinforcer signalled the likely location of the next reinforcer. Hence, patterns of responding in such conditions provide an indication of control by elapsed time in the absence of additional keylight-colour stimuli. If elapsed time strongly controlled behaviour, choice should favour the likely location of the next reinforcer as signalled by elapsed time – that is, local choice should follow the same pattern as the local reinforcer ratio.

The left panels of Figure 6.2 show response rates on the H→L and L→H keys (top panel) and the proportion of responses made to the L→H key (bottom panel) in No-Signal conditions. Response rates on the H→L key were high immediately after a reinforcer and decreased thereafter, whereas response rates on the L→H key were initially low and increased gradually over time. These response-rate patterns were reflected in local-choice analyses; choice strongly favoured the H→L key immediately after a reinforcer and shifted gradually towards and stabilised to favour the L→H key thereafter. Thus, choice followed a similar pattern to the local reinforcer ratio, indicating control by the time-based change in the reinforcer ratio. Such control by elapsed time was imperfect, however; the shift in choice towards the L→H key was gradual despite the abrupt step-change in the reinforcer ratio, and choice was consistently less extreme than reinforcer ratios.

To quantify more precisely control by elapsed time in No-Signal conditions, we fit a sigmoidal function (Cowie, Davison, Blumhardt, & Elliffe, 2016b) to individual-pigeon data using Microsoft® Excel Solver:

$$\frac{B_{L\to H}}{B_{L\to H}+B_{H\to L}} = Y_0 + \frac{\alpha}{1+e^{-\left(\frac{t-X_0}{B}\right)}} \qquad \text{(Equation 6.1)}$$

The relevant parameters in Equation 6.1 are X_0 and B. X_0 represents the "preference-reversal" time – the time at which choice reached halfway between the minimum (Y_0) and maximum ($Y_0 + \alpha$) of the function (i.e., the vertical midpoint of the function). B represents the speed of change in choice at the preference-reversal time; larger values indicate a slower change. In the present procedure, strong control by elapsed time would be evident as X_0 values close to 30 and small B values, because the reinforcer ratio reversed abruptly at 30 s (Cowie et al., 2016b). It should be noted that we are not attributing any theoretical significance to Equation 6.1. Equation 6.1 simply provides a useful way to quantify the preference-reversal time and speed of change in choice.

The right panels of Figure 6.2 shows median X_o and B values from fits of Equation 6.1 to individual-pigeon data. The variance accounted for (VAC) for fits of Equation 6.1 are shown in Appendix Table B6.1. In general, VACs were above .90, with only a few exceptions. Thus, the obtained X_o and B values provide a good indication of the preference-reversal time and the abruptness of the change in preference. The preference-reversal time (X_o; Figure 6.2, top-right panel) was always earlier than the reinforcer-ratio reversal time (30 s), indicating imperfect control by elapsed time. B values (Figure 6.2, bottom-right panel) ranged from about 3 to 9, reflecting the gradual change in choice across time.

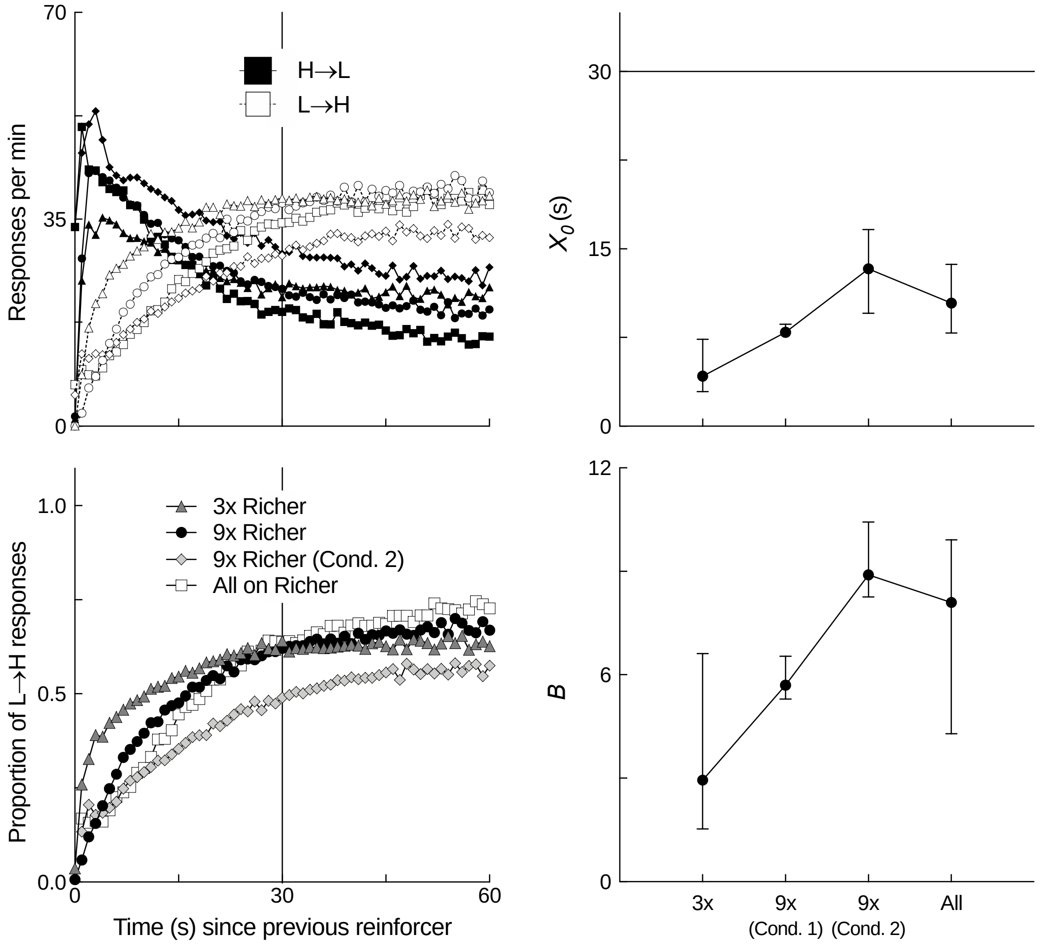

Figure 6.2. Response rates (top-left panel), the proportion of responses to the L→H key (bottom-left panel), and median obtained X_0 and B values (estimates of the time at which preference reversed and the speed of change in preference respectively; top-right and bottom-right panels) in No-Signal conditions. For the response-rate graph, filled symbols represent response rates on the H→L key and unfilled symbols represent response rates on the L→H key, and the symbol shapes represent the different conditions (3x Richer, 9x Richer, and All on Richer). The solid vertical or horizontal line at 30 s denotes the reinforcer-ratio reversal time. In the right panels, the error bars represent the interquartile range across pigeons.

In general, as the likely location of reinforcers at a particular time became more certain in No-Signal conditions, choice favoured that location more strongly: The proportion of responses to the L→H key before the reversal was lower, and the proportion of responses to the L→H key after the reversal was higher (Figure 6.2, bottom-left panel). Also, choice favoured the H→L key for longer, and hence reversed to favour the L→H key later. The latter finding suggests that control by elapsed time was slightly stronger when it signalled the location of the next reinforcer more reliably, as the preference-reversal time (X_o) occurred closer to the arranged reversal time. However, the median preference-reversal time was never later than 15 s (Figure 6.2, top-right panel). Additionally, the speed of change in choice was *slower* – that is, choice changed more gradually, rather than more abruptly – as the probability of reinforcer deliveries on the locally richer key increased (B values; Figure 6.2, bottom-right panel). Therefore, although the reliability with which elapsed time signalled the likely location of the next reinforcer appeared to have some effect on control by elapsed time, any such effects of reliability were small.

There was one exception to these effects of reliability on control by elapsed time. Preference for the H→L key was strongest and choice changed the slowest in the second 9x-Richer condition (Condition 2; Figure 6.2). This exception was probably related to a right-key bias exhibited by most subjects, rather than to the probability of reinforcer deliveries on the locally richer key. In Condition 2, the H→L key was the right key, whereas it was the left key in the other No-Signal conditions (see Table 6.1). A bias towards the right key in Condition 2 would have increased, and also prolonged the period of, preference for the H→L (right) key before the reversal and attenuated preference for the L→H (left) key after the reversal, compared with the other No-Signal conditions.

6.3.3 Choice With Keylight-Colour Stimuli

In Signal-Richer, Signal-Time, and Signal-Location conditions, keylight-colour stimuli provided information about the likely or definite time or location of the next reinforcer. These keylight-colour stimuli did not affect the local reinforcer ratio (Figure 6.1). Hence, if elapsed time controlled choice exclusively in conditions with keylight-colour stimuli, patterns of responding should be identical to patterns of responding in No-Signal conditions (Figure 6.2). However, if keylight colour also (or exclusively) controlled choice, then patterns of responding should differ between conditions with and without keylight-colour stimuli, and any such differences should be larger as control by keylight colour increases and/or control by elapsed time decreases. Furthermore, if control by elapsed time is stronger when elapsed time signals the location of future reinforcers more reliably, then patterns of responding in conditions with keylight-colour stimuli should become more similar to No-Signal conditions as the probability of reinforcers on the locally richer key increases.

6.3.3.1 Signal-Richer Conditions

In Signal-Richer conditions, keylight colour signalled the location of the locally richer key. Hence, the colour of the keys provided the same information as elapsed time. Figure 6.3 shows response rates (top panel) and the proportion of responses made on the L→H key (bottom panel) in Signal-Richer conditions. In these conditions, choice followed the local reinforcer ratio closely: Before the reversal, choice strongly favoured the H→L key, and choice reversed abruptly to favour the L→H key at the reinforcer-ratio reversal time. Thus, the addition of keylight-colour stimuli signalling the locally richer key enhanced control by the abrupt change in reinforcer ratio relative to No-Signal conditions (compare Figures 6.2 and 6.3), indicating that keylight-colour stimuli exerted strong control over behaviour. There also appeared to be some control by elapsed time before the reversal, although such control was much weaker than in No-Signal conditions (Figure 6.2); before the reversal in Signal-Richer conditions (Figure 6.3), response rates on the H→L key decreased slightly across time, and hence there was a slight shift in choice towards the L→H key. There were no systematic differences in the extent of change in choice between conditions arranging different probabilities of reinforcer deliveries on the locally richer key, suggesting that the strength of control by keylight-colour stimuli and by elapsed time did not depend on the reliability with which elapsed time signalled the likely next-reinforcer location.

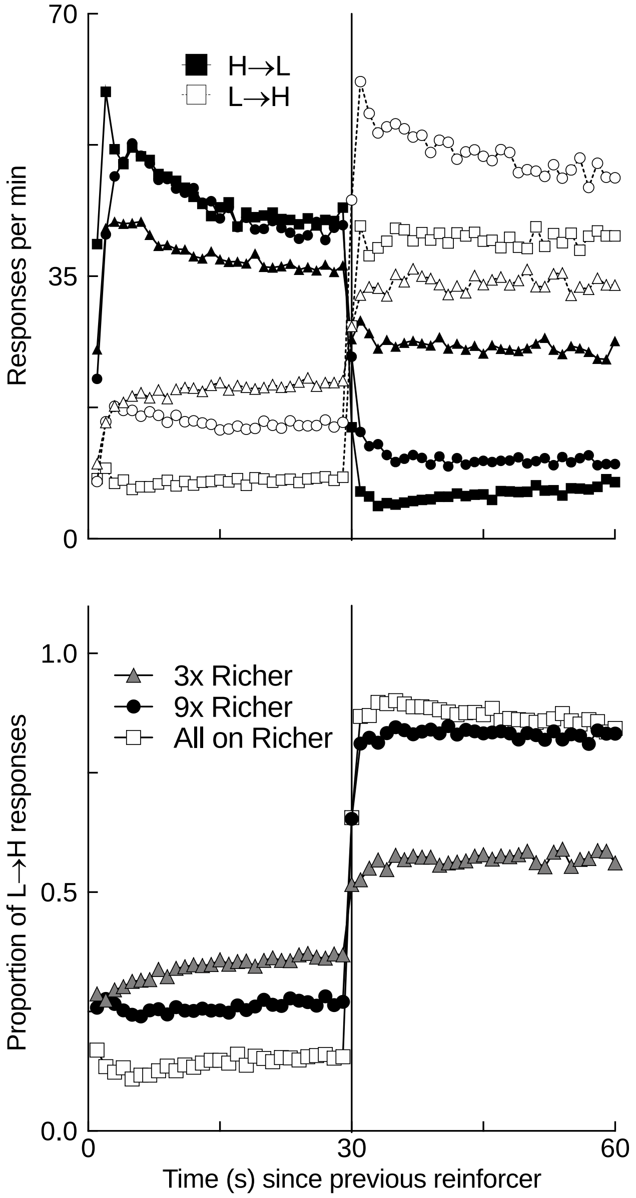

Figure 6.3. Response rates (top panel) and the proportion of responses to the L→H key (bottom panel) in Signal-Richer conditions. For the response-rate graph, filled symbols represent response rates on the H→L key and unfilled symbols represent response rates on the L→H key. The symbol shapes represent the different conditions (3x Richer, 9x Richer, and All on Richer). The solid vertical line denotes the reinforcer-ratio reversal.

6.3.3.2 Signal-Time Conditions

In Signal-Time conditions, keylight colour signalled whether the next reinforcer would occur before or after the reversal, and hence also signalled its likely location by extension – the H→L key if the next reinforcer was arranged before the reversal, and the

L→H key if the next reinforcer was arranged after the reversal. If the keylight-colour stimuli strongly controlled choice in these conditions, absolute response rates on both keys should increase systematically over time, but relative response rates (i.e., choice) should remain stable. This is because systematic changes in absolute response rates over time imply control by the time of the next reinforcer; response rates increase as the time of that reinforcer approaches, and increase more rapidly and are higher overall when the reinforcer occurs sooner than when it occurs later (Ferster & Skinner, 1957; see also Experiment 3, Figure 5.1). Hence, response rates should follow this pattern if responding was controlled by the likely time of the next reinforcer as signalled by keylight colour in Signal-Time conditions. In contrast, relative response rates should remain stable because changes in choice across time imply changes in the likely location of reinforcers across time, and no such changes were signalled by the keylight-colour stimuli. Instead, keylight colour signalled the likely next-reinforcer location with a *fixed* probability. For example, when keylight colour signalled that the next reinforcer would occur after the reversal in the 3x-Richer condition, the probability of that reinforcer occurring on the L→H key was .75, and this probability never changed. Therefore, if choice was controlled only by keylight colour, the proportion of responses to the likely next-reinforcer location should approximately equal the probability of the next reinforcer occurring on that location signalled by keylight colour, regardless of elapsed time.

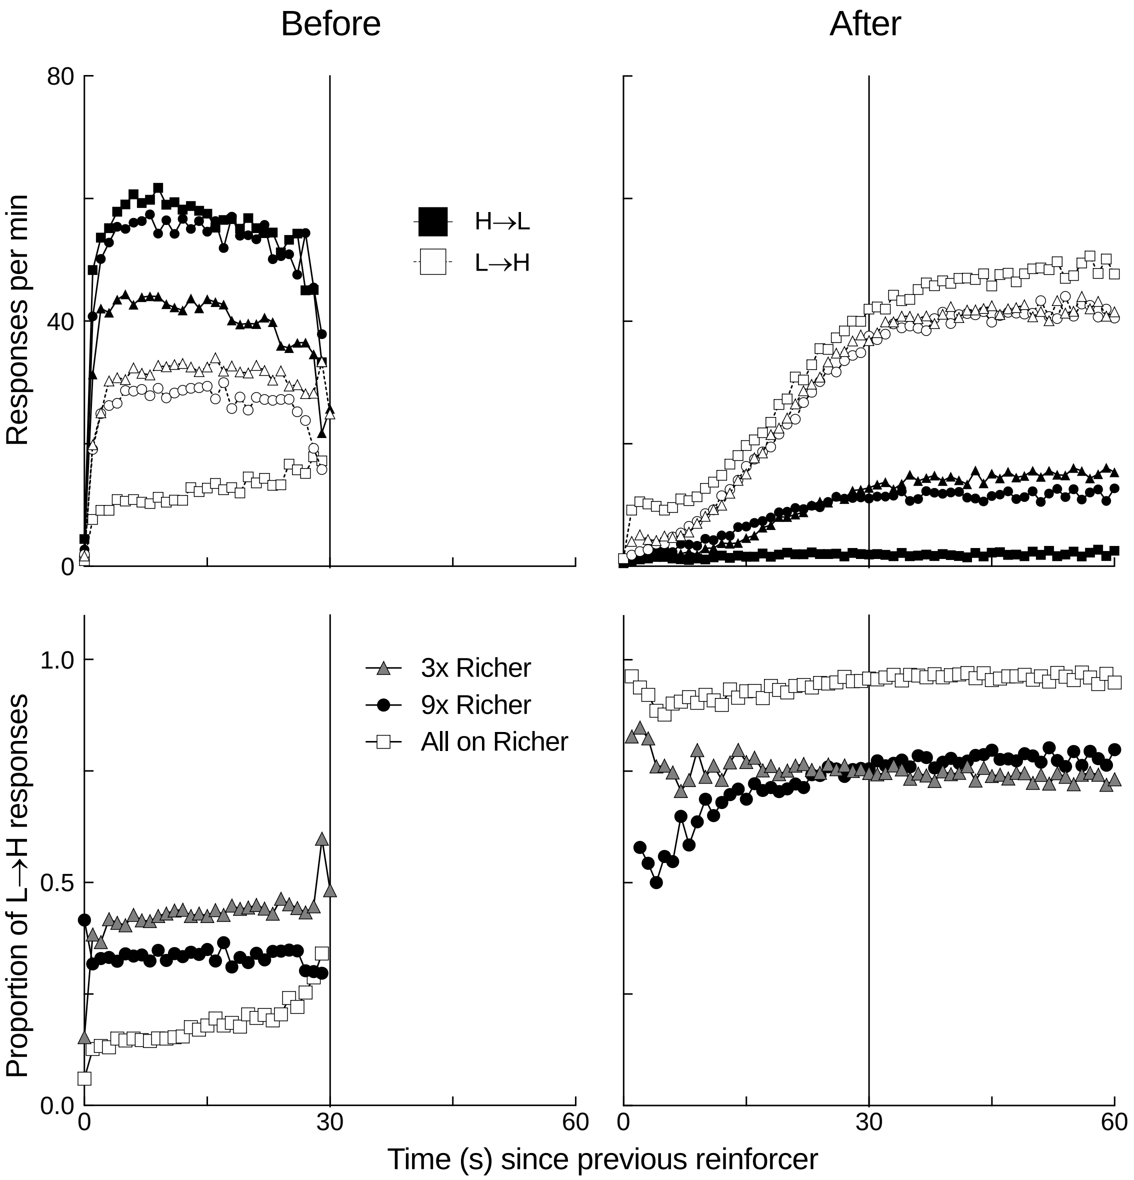

Figure 6.4. Response rates (top panels) and the proportion of responses to the L→H key (bottom panels) in Signal-Time conditions. Left panels show data from when keylight colour signalled that the next reinforcer would occur before the reversal, and the right panels show data from when keylight colour signalled that the next reinforcer would occur after the reversal. For the response-rate graph, filled symbols represent response rates on the H→L key and unfilled symbols represent response rates on the L→H key. The symbol shapes represent the different conditions (3x Richer, 9x Richer, and All on Richer). The solid vertical line denotes the reinforcer-ratio reversal.

Figure 6.4 shows response rates on both keys (top panels) and the proportion of responses made to the L→H key (bottom panels) in Signal-Time conditions. The left panels show response rates and choice when the keylight colour signalled that the next reinforcer would occur before the reversal, and the right panels show data from when the keylight colour signalled that the next reinforcer would occur after the reversal. Patterns of responding appeared to differ between the 9x-Richer condition (Condition 4) and the other Signal-Time conditions (Conditions 9 and 13), but these differences were most likely related to a strong right-key bias exhibited by some pigeons (e.g., Pigeon 45, see Appendix Figure A6.3). In the 9x-Richer condition, the H→L key was the right key, whereas in the other Signal-Time conditions, the H→L key was the left key (see Table 6.1).

Responding was strongly controlled by keylight-colour stimuli in all Signal-Time conditions (Figure 6.4). When the next reinforcer was signalled to occur before the reversal, response rates increased rapidly within the first few seconds after a reinforcer delivery and remained high thereafter, and response rates were higher on the H→L key than on the L→H key (Figure 6.4, top-left panel). When the next reinforcer was signalled to occur after the reversal, response rates increased gradually before the reversal and then remained high after the reversal, and were higher on the L→H key than on the H→L key (Figure 6.4, top-right panel). Thus, response rates on both keys were higher during the time in which the next reinforcer was signalled to occur than at other times, and choice favoured the likely next-reinforcer location signalled by keylight colour.

When the keylight-colour stimuli signalled that the next reinforcer would occur before the reversal, preference for the likely next-reinforcer location was more extreme when the probability of reinforcer deliveries on the locally richer key was higher (Figure 6.4, bottom-left panel). However, this was not the case when keylight-colour stimuli signalled that the next reinforcer would occur after the reversal (Figure 6.4, bottom-right panel). Some of this discrepancy was probably related to the strong right-key bias, which would have enhanced preference for the H→L key before the reversal and attenuated preference for the L→H key after the reversal in the 9x-Richer condition compared with the other conditions.

In Signal-Time conditions, some control by elapsed time was also evident (Figure 6.4, bottom panels). Even though the keylight-colour stimuli did not signal a change in the likely location of the next reinforcer, choice changed slightly over time in accordance with the local reinforcer ratio. The magnitude of changes in choice depended on the time of the next reinforcer; changes in choice across time were larger when the next reinforcer was signalled to occur after the reversal than when it was signalled to occur before the reversal (compare Figure 6.4, bottom-left and bottom-right panels). Additionally, the direction of change in choice differed depending on the next-reinforcer time. When the next reinforcer was signalled to occur before the reversal, choice shifted slightly towards the L→H key over time, whereas when the next reinforcer was signalled to occur after the reversal, choice shifted towards the H→L key during the first 5 s after a reinforcer delivery, and thereafter shifted towards the L→H key.

To assess the effects of stimulus reliability (i.e., the probability of reinforcer deliveries on the locally richer key) on control by elapsed time and by keylight-colour stimuli in Signal-Time conditions, we focused only on the 3x-Richer and All-on-Richer conditions. This was because the H→L key was the same in both of those conditions, and so any differences in responding were related to the probability of reinforcer deliveries on the locally richer key and not to the location of the H→L key. The extent of changes in choice across time was similar in both 3x-Richer and All-on-Richer conditions, although changes in choice did appear to be slightly larger in the All-on-Richer condition (Figure 6.4, bottom panels). Nevertheless, any such difference between conditions was very small, and hence as in No-

Signal and Signal-Richer conditions, the reliability with which elapsed time signalled the likely location of the next reinforcer had little, if any, effect on control by elapsed time.

6.3.3.3 Signal-Location Conditions

In Signal-Location conditions, keylight colour signalled the definite location of the next reinforcer, and by extension its likely time – before the reversal if the next reinforcer was arranged on the H→L key, and after the reversal if the next reinforcer was arranged on the L→H key. Hence, as in Signal-Time conditions, absolute response rates should increase systematically over time and choice should remain stable across time if keylight-colour stimuli strongly controlled behaviour in Signal-Location conditions. Figure 6.5 shows response rates (top panels) and the proportion of responses made to the L→H key (bottom panels) in Signal-Location conditions, separated based on whether the keylight colour signalled that the next reinforcer would occur on the H→L key (left panels) or on the L→H key (right panels). Because keylight colour signalled both the time and location of the next reinforcer in the All-on-Richer condition, the data plotted in Figure 6.5 for this condition are the same as the data plotted in Figure 6.4.

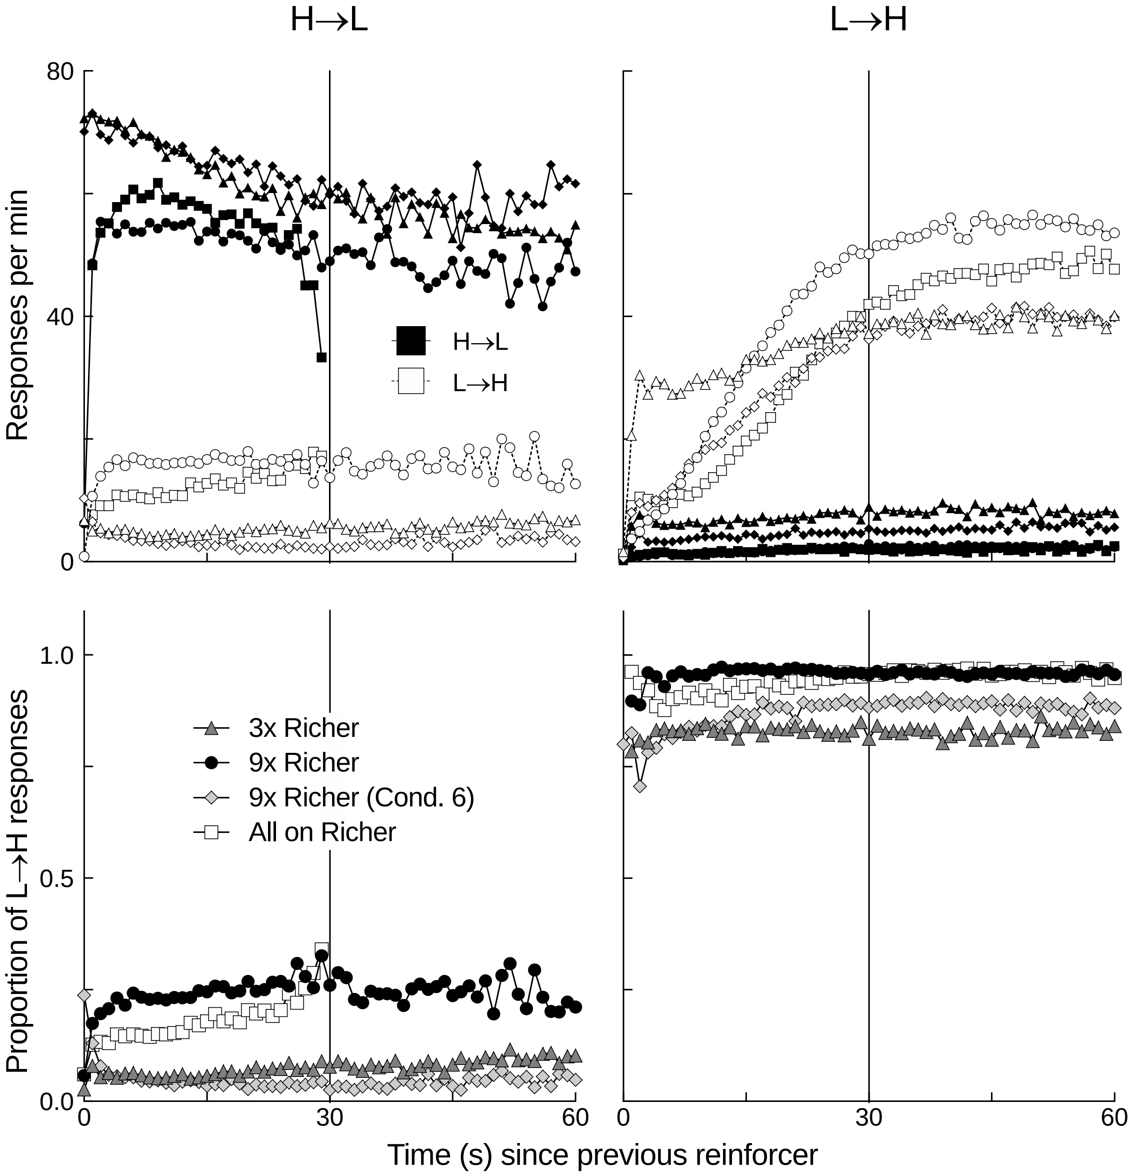

Figure 6.5. Response rates (top panels) and the proportion of responses to the L→H key (bottom panels) in Signal-Location conditions. Left panels show data from when keylight colour signalled that the next reinforcer would occur on the H→L key, and the right panels show data from when keylight colour signalled that the next reinforcer would on the L→H key. For the response-rate graph, filled symbols represent response rates on the H→L key and unfilled symbols represent response rates on the L→H key. The symbol shapes represent the different conditions (3x Richer, 9x Richer, and All on Richer). The solid vertical line denotes the reinforcer-ratio reversal.

As in Signal-Time conditions, responding was strongly controlled by keylight colour in Signal-Location conditions. Response rates were high immediately after a reinforcer when the next reinforcer was signalled to occur on the H→L key and hence was likely to occur earlier (Figure 6.5, top-left panel), whereas response rates increased gradually when the next reinforcer was signalled to occur on the L→H key and hence was likely to occur later (Figure 6.5, top-right panel). Choice strongly favoured the definite location of the next reinforcer regardless of elapsed time in Signal-Location conditions (Figure 6.5, bottom panels). Such preference was stronger than preference for the likely next-reinforcer location in Signal-Time conditions (Figure 6.4, bottom panels). This difference in the strength of preference between Signal-Time and Signal-Location conditions is consistent with the information provided by keylight-colour stimuli. Keylight colour signalled the location of the next reinforcer more definitively in Signal-Location conditions than in Signal-Time conditions, resulting in stronger preference for the next-reinforcer location in the former conditions.

Levels of preference in Signal-Location conditions were modulated by a right-key bias (Figure 6.5; see also Appendix Figure A6.5). When the next-reinforcer location was the right key (the H→L key in Conditions 6 and 10, and the L→H key in Conditions 5 and 13; see Table 6.1), preference for the next-reinforcer location was stronger than when the next-reinforcer location was the left key. These effects of a right-key bias are clearest in the two 9x-Richer conditions (Conditions 5 and 6), which were identical except that the H→L key was the left key in Condition 5 and the right key in Condition 6.

Preference in Signal-Location conditions was also modulated by the local reinforcer ratio (Figure 6.5, bottom panels). Choice never favoured the next-reinforcer location exclusively, and preference was slightly more extreme when the probability of reinforcer deliveries on the locally richer key was higher, even though the keylight-colour stimuli always signalled the definite location of the next reinforcer. Additionally, choice shifted slightly towards the L→H key in the first few seconds after a reinforcer delivery, therefore following the same pattern as the local reinforcer ratio. Thus, in addition to strong control by the keylight-colour stimuli, elapsed time also exerted a small amount of control over choice in Signal-Location conditions.

Elapsed time appeared to exert slightly stronger control over choice in Signal-Location conditions in which the probability of reinforcer deliveries on the locally richer key was higher. Comparing conditions in which the H→L key was the same – for example, the 3x-Richer condition and the second 9x-Richer condition (Condition 6, see Table 6.1) – changes in choice appeared slightly larger when the probability of reinforcer deliveries on the locally richer key was higher (Figure 6.5, bottom panels). This was particularly true when the next reinforcer was arranged on the L→H key. Nevertheless, as in the other conditions with keylight-colour stimuli (Signal-Richer and Signal-Time conditions), these effects of the reliability of elapsed time on divided control between elapsed time and keylight colour were small.

6.3.4 Summary of Main Results

In the present experiment, elapsed time since a reinforcer delivery signalled the likely location of the next reinforcer: The H→L key during the first 30 s after a reinforcer, and the L→H key thereafter (Figure 6.1). When no keylight-colour stimuli were present (No-Signal conditions), choice was controlled, albeit imperfectly, by elapsed time (Figure 6.2). When keylight-colour stimuli signalled the likely location of the next reinforcer as well (Signal-Richer conditions), such stimuli controlled choice almost exclusively (Figure 6.3). When keylight-colour stimuli signalled the time or location of the next reinforcer (Signal-Time and Signal-Location conditions), such stimuli also controlled choice strongly, but weak control by elapsed time was also evident (Figures 6.4 and 6.5). Keylight-colour stimuli that signalled future reinforcer availability with more certainty (Signal-Location conditions) exerted stronger control over behaviour than stimuli that signalled future reinforcer availability with less certainty (Signal-Time conditions). In contrast, the reliability with which elapsed time signalled future reinforcer availability had virtually no effect on the division of control between elapsed time and keylight colour.

6.4 Discussion

The present experiment investigated divided control between elapsed time and keylight-colour stimuli in a concurrent schedule in which the likely location of future reinforcers reversed abruptly at a fixed time since the previous reinforcer (Figure 6.1). We found that stimulus control was divided between elapsed time and keylight colour, but the latter strongly overshadowed the former (Figures 6.2 to 6.5). The effects of relative stimulus reliability were asymmetrical: The reliability of elapsed time had negligible effects on divided stimulus control, whereas more reliable keylight colour stimuli appeared to exert stronger control over behaviour.

Thus, our findings are inconsistent with the findings from another reversal-learning paradigm, the mid-session reversal task (MSR). In the MSR, elapsed time signals the likely location of the next reinforcer more reliably at early and late times during each experimental session, and less reliably at times closer to the reversal in reinforcer location (i.e., around the midpoint of the session; Rayburn-Reeves et al., 2017). Rayburn-Reeves et al. (2017) showed that divided control between elapsed time since session start and a keylight-colour stimulus that signalled the location of the next reinforcer depended on the reliability of elapsed time; elapsed time exerted stronger control and keylight colour exerted weaker control at times when elapsed time was more reliable (see also Smith et al., 2017). Likewise, when elapsed time is made less reliable by varying the location of the reversal during each session, control by elapsed time decreases and control by other discriminative stimuli that signal future reinforcer availability increases in the MSR task (e.g., McMillan et al., 2016; Rayburn-Reeves, Molet, & Zentall, 2011).

In contrast, our findings are similar to those of Cowie et al. (2017), who also investigated divided control between elapsed time and keylight-colour stimuli when reinforcer deliveries were followed by a change in reinforcer ratio in a concurrent schedule. Like us, Cowie et al. found that keylight-colour stimuli always exerted strong control over choice regardless of the reliability of elapsed time, and keylight-colour stimuli that signalled the location of the next reinforcer more definitively exerted stronger control. Therefore, both Cowie et al.'s and the present findings suggest that divided stimulus control may not depend on the reliability with which elapsed time signals future reinforcer availability when changes in reinforcer availability occur across seconds after a reinforcer delivery. Compared with Cowie et al., the present results provide stronger evidence for this suggestion, as we dissociated changes in the reliability of elapsed time from changes in elapsed time itself, whereas changes in the reliability of elapsed time were correlated with changes in elapsed time in Cowie et al.'s procedure.

6.4.1 Effects of Reliability of Elapsed Time on Divided Stimulus Control

Why might the reliability of elapsed time affect divided control between elapsed time and other non-temporal stimuli in the MSR task, but not in the present experiment or in Cowie et al.'s (2017) experiment? One possibility is that the effects of the reliability of elapsed time on divided stimulus control depend on the time-scale on which changes in reinforcer availability occur. In the MSR, the change in reinforcer location occurs on a larger time-scale – across the session – whereas in the reinforcer-ratio reversal task arranged here (Figure 6.1) and in Cowie et al., the change in reinforcer availability occurs on a smaller time-scale – across seconds after a reinforcer delivery. Perhaps the time-scale modulates the extent to which the reliability of elapsed time determines the division of control between elapsed time and other stimuli.

Alternatively, the reliability of elapsed time may have had little effect in the present and Cowie et al.’s (2017) experiments, and a larger effect in the MSR task, because of differences in the *discriminability* of changes in reliability. In order for changes in stimulus reliability to shift divided stimulus control, subjects must discriminate changes in reliability (see e.g., Davison & Nevin, 1999). In other words, subjects must be able to detect changes in reliability after they occur. When the likely location of future reinforcers changes across time in a concurrent schedule, such discrimination requires that subjects correctly attribute obtained reinforcers to the response that produced them, and discriminate accurately the time at which reinforcers were obtained (see Cowie, Davison, & Elliffe, 2014, 2016d; Davison & Nevin, 1999). Although response-reinforcer discriminability is probably high when reinforcers are obtained from separate response alternatives (as in the present procedure), discrimination of the time at which reinforcers were obtained is probably relatively inaccurate and imprecise (see Cowie et al., 2014). As a result of this poor temporal discrimination, subjects are unable to discriminate changes in the reliability of elapsed time (e.g., in the probability of reinforcer deliveries on the locally richer key), and hence such changes have little effect on control by elapsed time and by other non-temporal stimuli. In support of such a discriminability account, any effects of reliability on control by elapsed time were greatest at early times after a reinforcer in the present experiment. Because the error in temporal discrimination is proportional to the duration to-be-timed (scalar property of timing; Gibbon, 1977), the discriminability of elapsed time would be higher immediately after a reinforcer, and hence any changes in reliability would be easier to discriminate, compared with later times after a reinforcer.

However, similarly poor discrimination by elapsed time would be expected in the present experiment, Cowie et al.’s (2017) experiment, and the MSR task. Therefore, on its own, poor temporal discrimination cannot explain why changes in the reliability of elapsed time affect divided stimulus control in the MSR but not in concurrent schedules that arrange time-based changes in the reinforcer ratio after reinforcer deliveries. Indeed, it appears that changes in reinforcer availability were *harder* to discriminate in the present and Cowie et al.’s procedures than in the MSR task; compared with the present No-Signal conditions (Figure 6.2) and with Cowie et al.’s conditions without keylight-colour stimuli, deviations in choice from reinforcer ratios appear to be smaller in the MSR task. This difference may arise due to several key procedural factors. Both the present experiment and Cowie et al.’s experiment arranged VI schedules, whereas every correct response is followed by a reinforcer delivery in the MSR task. Choice tends to follow changes in contingency more closely when reinforcer rates are higher (Bizo & White, 1994; Cowie, Bizo, & White, 2016a; Cowie et al., 2016b; Davison & Baum, 2010; Mazur, 1997), and so choice may follow changes in reinforcer availability more closely in the MSR than in concurrent VI VI schedules because the overall reinforcer rate is higher in the former than in the latter.

Furthermore, in the present experiment and in Cowie et al.’s (2017) experiment, the reinforcer ratio never favoured one key exclusively except in our All-on-Richer conditions. In contrast, in the MSR, reinforcers are only available at one location at a time. This may explain why choice follows the change in reinforcer ratio more closely in the MSR. However, this line of reasoning implies that choice should have followed the reinforcer ratio more closely in our All-on-Richer conditions. Therefore, differences in the probability of reinforcers on the locally richer key cannot explain completely why it was harder for subjects to discriminate the reinforcer-ratio reversal in the present procedure compared with the MSR task. Another possible procedural difference that may contribute to discriminability of the reversal is trial-length variability. In the MSR, the number of trials before and after the reversal remains the same across sessions, and the reversal occurs in every session. This was not the case in the present experiment; although 30 s always elapsed before the reversal, the number of seconds after the reversal varied depending on the VI schedule, and the reversal occurred only if more than 30 s had elapsed since the previous reinforcer. Likewise, in Cowie et al.'s experiment, the time between reinforcer deliveries varied substantially depending on the VI schedule. Temporal discrimination may be more precise and accurate, and hence control by elapsed time may be stronger, when events before and after a change in reinforcer availability remain relatively constant, and when the change in reinforcer availability is experienced more regularly. If so, then choice would follow the arranged contingency more closely in the MSR task than in concurrent VI VI schedules in which changes in reinforcer availability occur during variable-length inter-reinforcer intervals (see also Cowie et al., 2014).

Finally, one other notable difference between the MSR task and concurrent schedules in which the reinforcer ratio changes across time since a reinforcer (e.g., Cowie et al., 2017; the present procedure, Figure 6.1) is that the change in reinforcer availability is contingent on number of trials, not on elapsed time, in the MSR task. As a result, number of trials since session start serves as an additional discriminative stimulus signalling the likely location of the next reinforcer. Choice will follow the change in reinforcer location more closely to the extent that this additional stimulus controls behaviour; indeed, time and number have been shown control behaviour jointly, although elapsed time typically exerts stronger control (Davison & Cowie, 2019; Fetterman, 1993; MacDonald & Roberts, 2018; Meck, Church, & Gibbon, 1985; W. A. Roberts & Mitchell, 1994).

Along similar lines, the location of and consequence following each response provides information about the next-reinforcer location in the MSR task; if a response is reinforced, there is a high probability of the same response producing a reinforcer in the next trial, whereas if a response is not reinforced, there is a high probability of the other response producing a reinforcer. Hence, choice may be additionally controlled by the location or outcome of the previous response in the MSR task (see e.g., Santos, Soares, Vasconcelos, & Machado, 2019). In contrast, no such information about future reinforcer availability was signalled by the location or outcome of the previous response in the present procedure. Although last-reinforcer location was a discriminative stimulus in Cowie et al.'s (2017) experiment, it signalled a time-based change in reinforcer ratio, and hence provided less certain information about future reinforcer availability compared with the MSR. Thus, the presence of additional discriminative stimuli signalling the next-reinforcer location – number of trials, the location of the previous response, and the consequence of that response – in the MSR task may serve to enhance the discriminability of, and hence control by, the change in reinforcer location.

6.4.2 Effects of Reliability of Non-Temporal Stimuli on Divided Stimulus Control

Although the relative reliability with which elapsed time signalled future reinforcer availability had little effect on divided stimulus control in the present experiment, the division of control between elapsed time and keylight colour did differ between Signal-Richer, Signal-Time, and Signal-Location conditions (Figures 6.3 to 6.5). When keylight-colour stimuli signalled the location of the locally richer key (Signal-Richer conditions), keylight colour controlled choice nearly exclusively. In contrast, when keylight colour signalled the time or location of the next reinforcer (Signal-Time and Signal-Location conditions), stimulus control was divided between elapsed time and keylight colour. Additionally, elapsed time appeared to exert stronger control over choice in Signal-Time conditions than in Signal-Location conditions, as changes in choice across time were larger in the former than in the latter conditions.

These differences in divided control between elapsed time and keylight colour in Signal-Richer, Signal-Time, and Signal-Location conditions were likely related to the type and reliability of the information provided by keylight-colour stimuli. In Signal-Richer conditions, elapsed time and keylight colour were redundant relevant cues, as both signalled the same information – the location of the locally richer key. Previous research suggests that stimulus control is more likely to be selective than divided in the presence of redundant relevant cues (e.g., Blaisdell, Denniston, & Miller, 1998; Boutros, Davison, & Elliffe, 2009, 2011; Cowie et al., 2017; Egger & Miller, 1962, 1963; McCausland, Menzer, Dempsey, & Birkimer, 1967; McLinn & Stephens, 2006, 2010), and that stimuli that provide more reliable information about future reinforcer availability are more likely to control behaviour (e.g., Blaisdell, Schroeder, & Fast, 2018; Cheng, Shettleworth, Huttenlocher, & Rieser, 2007; Cowie et al., 2017; Davison & Elliffe, 2010; Du, McMillan, Madan, Spetch, & Mou, 2017; Lechelt & Spetch, 1997; McLinn & Stephens, 2006, 2010; Rayburn-Reeves et al., 2017; Rubi & Stephens, 2016a, 2016b; Shahan & Podlesnik, 2006, 2007). In the present experiment, keylight-colour stimuli were probably more reliable than elapsed time, because the error in discriminating colour is likely smaller than the error in discriminating time since a reinforcer delivery (see Cowie et al., 2014, 2016c). Therefore, it is unsurprising that keylight colour controlled choice exclusively in Signal-Richer conditions here.

In contrast to Signal-Richer conditions, elapsed time and keylight colour were not redundant relevant cues in Signal-Time and Signal-Location conditions. Thus, obtained reinforcers were differential with respect to elapsed time, and were also differential with respect to keylight colour. Given previous research demonstrating that stimulus control is divided when multiple stimuli differentially signal reinforcer deliveries (e.g., Blough, 1969; Davison & Elliffe, 2010; Leith & Maki, 1975; Shahan & Podlesnik, 2006), it is unsurprising that stimulus control was divided between elapsed time and keylight colour in Signal-Time and Signal-Location conditions. As in Signal-Richer conditions, keylight-colour stimuli signalled future reinforcer availability more reliably than elapsed time, and so keylight colour exerted stronger control over choice in Signal-Time and Signal-Location conditions. Additionally, elapsed time probably exerted weaker control over choice in Signal-Location conditions than in Signal-Time conditions because keylight colour provided more reliable information about the next reinforcer in Signal-Location conditions: Whereas keylight colour signalled the definite location of the next reinforcer in Signal-Location conditions, no such definitive information about the next reinforcer was signalled by keylight colour in Signal-Time conditions. Therefore, relative stimulus reliability appears to have similar effects on divided control between multiple non-temporal stimuli, and on divided control between elapsed time and other non-temporal stimuli (e.g., keylight colour), at least when the reliability of the non-temporal stimuli is manipulated.

Divided stimulus control also appeared to differ *within* each condition. This was clearest in Signal-Time conditions, in which changes in choice were larger when keylight colour signalled that the next reinforcer was arranged after the reversal compared with when the next reinforcer was arranged before the reversal (Figure 6.4). That is, some control by elapsed time – by what was *likely* to happen at the current time – was evident when the next reinforcer was signalled to occur after the reversal, but not when it was signalled to occur before the reversal. This result was not due to differences in the duration of keylight-colour stimuli signalling that the next reinforcer would occur before or after the reversal; although the keylight-colour stimulus signalling that the next reinforcer was arranged after the reversal was presented for longer because it signalled a longer time to the next reinforcer, changes in choice only occurred in the first few seconds after a reinforcer delivery.

Instead, the difference in local choice depending on keylight colour in Signal-Time conditions appears to be related to differences in the correlation between keylight colour and changes in the reinforcer ratio, and also to the congruency between the information provided by elapsed time and by keylight colour. When keylight colour signalled that the next reinforcer would occur before the reversal, such stimuli were not correlated with a change in reinforcer ratio because reinforcers were obtained before the change occurred, whereas keylight-colour stimuli that signalled that the next reinforcer would occur after the reversal were always correlated with the reinforcer-ratio reversal. Hence, elapsed time only signalled a change in reinforcer availability when the keylight-colour signalled that the next reinforcer would occur after the reversal. This may explain why control by elapsed time appeared to be stronger when keylight colour signalled that the next reinforcer would occur after the reversal; systematic changes in choice across time are more likely to occur if a time-based change in reinforcer availability occurs (e.g., Cowie et al., 2011; see Cowie & Davison, 2016 for a review).

Additionally, at early times after a reinforcer delivery, elapsed time always signalled that the next reinforcer was likely to occur on the H→L key. Thus, when keylight colour signalled that the next reinforcer would occur before the reversal and probably on the H→L key in Signal-Time conditions, keylight colour provided the same information as elapsed time. However, when keylight colour signalled that the next reinforcer would occur after the reversal and hence probably on the L→H key, this information was incongruent with elapsed time at early times after a reinforcer delivery. That is, when both elapsed time and keylight colour provided congruent information, control by elapsed time was weaker, whereas when they provided incongruent information, some control by elapsed time was evident.

Changes in choice in Signal-Location conditions were much smaller than in Signal-Time conditions, and typically occurred only within the first 5 or so seconds after a reinforcer delivery (Figure 6.5). Nevertheless, there was a slight difference in the extent of change in choice within Signal-Location conditions as well; changes in choice across time were more apparent when the next reinforcer was signalled to occur on the L→H key than when it was signalled to occur on the H→L key, particularly in the 3x-Richer and 9x-Richer conditions. Reinforcers arranged on the L→H key were more likely to occur after the reversal, and hence keylight-colour stimuli that signalled that the next reinforcer would occur on the L→H key were more likely to be correlated with a change in the reinforcer ratio than stimuli that signalled that the next reinforcer would occur on the H→L key. Also, when the next reinforcer was arranged on the L→H key, the keylight-colour stimulus conflicted with elapsed time at early times after a reinforcer, because elapsed time signalled that the H→L key was more likely to deliver reinforcers at early times. Thus, as in Signal-Time conditions, the extent of control by elapsed time at early times appeared to depend on the correlation between keylight colour and the change in reinforcer ratio, and on the congruency of the information provided by elapsed time and keylight colour. Similarly, when Cowie et al. (2017) signalled the location of the next reinforcer (as in our Signal-Location conditions), control by elapsed time was greater at times when keylight colour provided information that was incongruent with elapsed time.

6.4.3 Other Factors Contributing to Divided Stimulus Control

One other factor that may have contributed to the changes in choice across time in Signal-Time and Signal-Location conditions is the pigeons' previous experience in No-Signal conditions. All pigeons experienced conditions without keylight-colour stimuli (No-Signal conditions) before conditions with keylight-colour stimuli (see Table 6.1). In a similar reinforcer-ratio reversal procedure to that arranged here, Davison and Cowie (2019) found persistently strong control by elapsed time, even if the reversal was contingent on number of responses instead of on time since the previous reinforcer, and even if the reliability of elapsed time was degraded. Davison and Cowie suggested that their pigeons' extensive experience discriminating time-based reinforcer-ratio reversals may have contributed to such pervasive control by elapsed time, much like pre-training with one element of a compound stimulus can block conditioning to the other element(s) of the compound (e.g., Kamin, 1969; Mackintosh, 1971). A similar explanation may apply to the present experiment. Here, the pigeons' initial experience in No-Signal conditions may have established a reinforcer history for discriminating the time-based reinforcer-ratio reversal, thereafter resulting in continued control (albeit weak) by elapsed time even when keylight colour stimuli signalled reinforcer availability more reliably.

Compared with the present experiment, changes in choice across time were larger in Cowie et al. (2017), suggesting that elapsed time exerted stronger control in their experiment than in our experiment. In support of this, local choice in conditions with and without keylight-colour stimuli was more similar in Cowie et al.'s experiment than in the present experiment. This stronger control by elapsed time in Cowie et al. may be related to the discriminative properties of last-reinforcer location, as the change in reinforcer ratio depended on last-reinforcer location in Cowie et al.'s procedure, whereas no such dependence on last-reinforcer location was arranged in the present procedure. Thus, taken together, the present results and Cowie et al.'s results suggest that when elapsed time and other non-temporal discriminative stimuli signal future reinforcer availability, control by elapsed time is stronger and control by non-temporal stimuli is weaker when (1) the non-temporal stimuli signal future reinforcer availability less reliably, (2) the non-temporal stimuli are more strongly correlated with a change in reinforcer availability, (3) the information provided by non-temporal stimuli is incongruent with the information signalled by elapsed time, (4) subjects have previous experience discriminating time-based changes in reinforcer availability (see also Davison & Cowie, 2019), and (5) the change in reinforcer availability across time depends on last-reinforcer location.

6.5 Chapter Summary and Conclusions

Previous research suggests that when changes in reinforcer availability occur over time, elapsed time may come to control behaviour, in conjunction with other non-temporal discriminative stimuli (e.g., Cowie et al., 2011, 2014, 2017; Rayburn-Reeves et al., 2017; Smith et al., 2017; Experiment 3, Chapter 5, of this thesis). The present experiment asked how relative stimulus reliability divides control between elapsed time and a non-temporal stimulus (keylight colour). Our findings suggest that when reinforcer availability changes across time since a reinforcer delivery in a concurrent schedule, the reliability with which non-temporal stimuli signal future reinforcer availability affects the division of control between elapsed time and non-temporal stimuli, whereas the reliability of elapsed time does not. More generally, whereas changes in the ability of a non-temporal stimulus to predict reinforcer availability result in concomitant changes in divided stimulus control, changes in the predictive ability of elapsed time itself may not.

The present results also demonstrate that when non-temporal stimuli provide information that conflicts with elapsed time, control by elapsed time is stronger compared with when non-temporal stimuli provide information that is congruent with elapsed time (see also Cowie et al., 2017). Thus, these findings contribute to a growing body of evidence showing that elapsed time exerts pervasive control over behaviour, even if other reliable non-temporal stimuli signal future reinforcer availability (e.g., Cowie et al., 2017; Davison & Cowie, 2019; Rayburn-Reeves et al., 2011, 2017; see McMillan, Spetch, Sturdy, & Roberts, 2017 for a review). In relation to previous studies that have investigated reversal learning (e.g., in the MSR), our results suggest that procedural factors, such as the overall reinforcer rate, trial-length variability, and the presence of additional discriminative stimuli (e.g., number of trials, location and outcome of the previous response), may affect the discriminability of the time-based change in reinforcer availability, and hence whether or not the reliability of elapsed time influences divided stimulus control.

Tasks that involve changes in reinforcer availability across time, such as that arranged here, are likely to be more ecologically valid than those that ask subjects to discriminate stimulus duration, such as the DMTS task arranged by Davison and Elliffe (2010). It is unlikely that human and non-human animals would be asked to report the duration of a stimulus in the natural world, whereas changes in contingency across time probably occur frequently (Buhusi & Meck, 2005; Davison & Cowie, 2019; Marshall & Kirkpatrick, 2015). For example, for a foraging animal, the probability of food at different foraging patches may vary depending on the time of day (see e.g., Biebach, Gordijn, & Krebs, 1989; Daan & Koene, 1981; Thorpe & Wilkie, 2002; Wilkie, 1995; Wilkie & Willson, 1992), just as the probability of reinforcers on the higher-to-lower key varied across time in the present experiment. Additionally, in the wild, time-based changes in food availability may be correlated with additional, non-temporal discriminative stimuli, such as footprints that signal the presence of predators or prey. The present results suggest that under such conditions, both elapsed time and informative non-temporal stimuli control behaviour, with the latter exerting much stronger control regardless of the reliability of elapsed time. Therefore, the present experiment provides insight into some of the variables that likely contribute to divided stimulus control in the natural world.

Appendix A6

Proportion of responses made on the lower-to-higher (L→H) key for individual pigeons in each condition

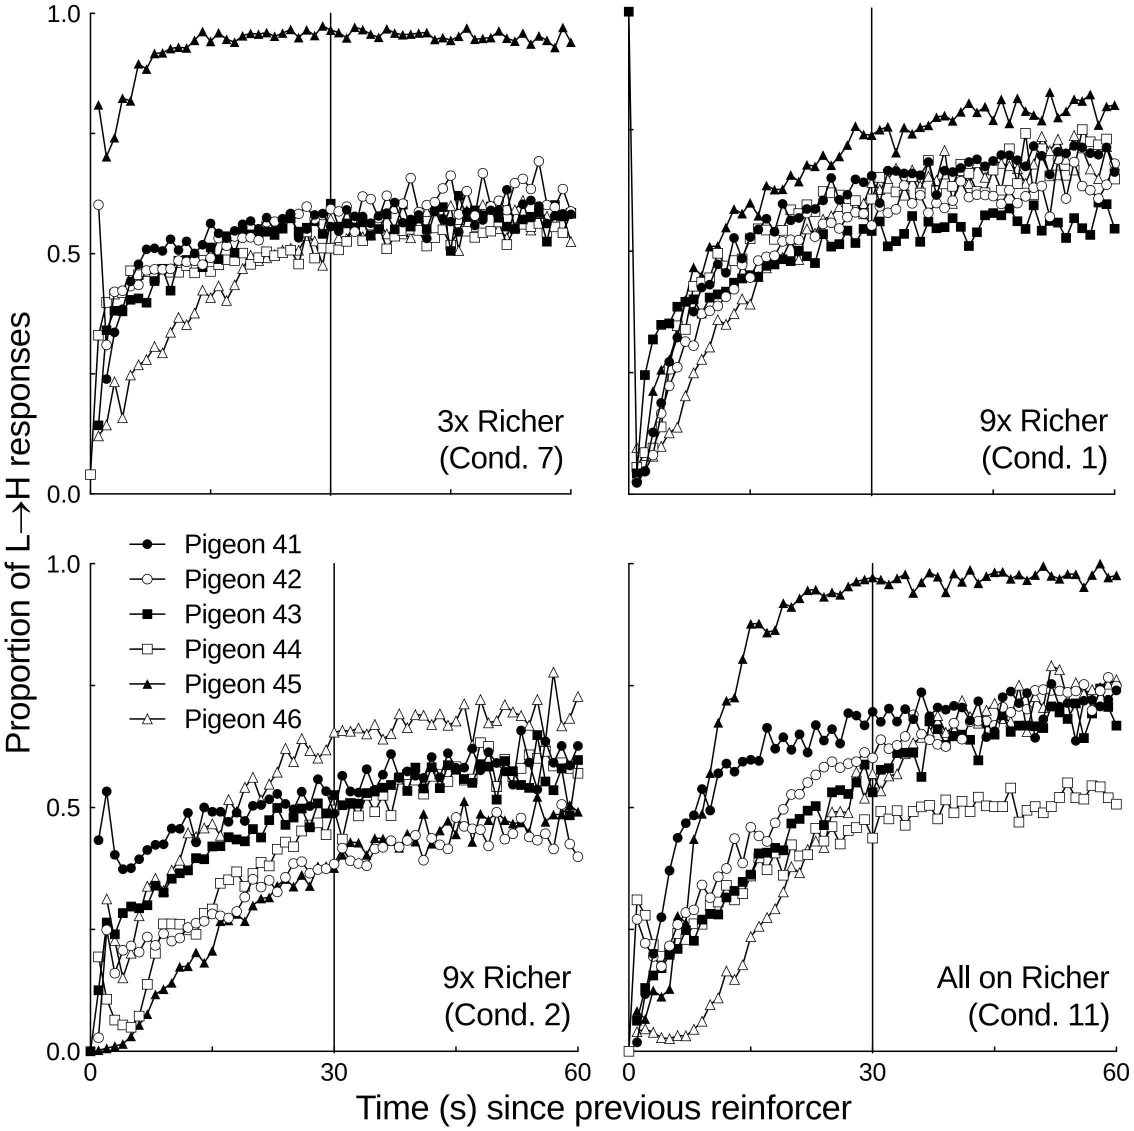

Figure A6.1. Proportion of responses made to the lower-to-higher (L→H) key across time since a reinforcer, for individual pigeons in No-Signal conditions. The solid vertical line denotes the reinforcer-ratio reversal time.

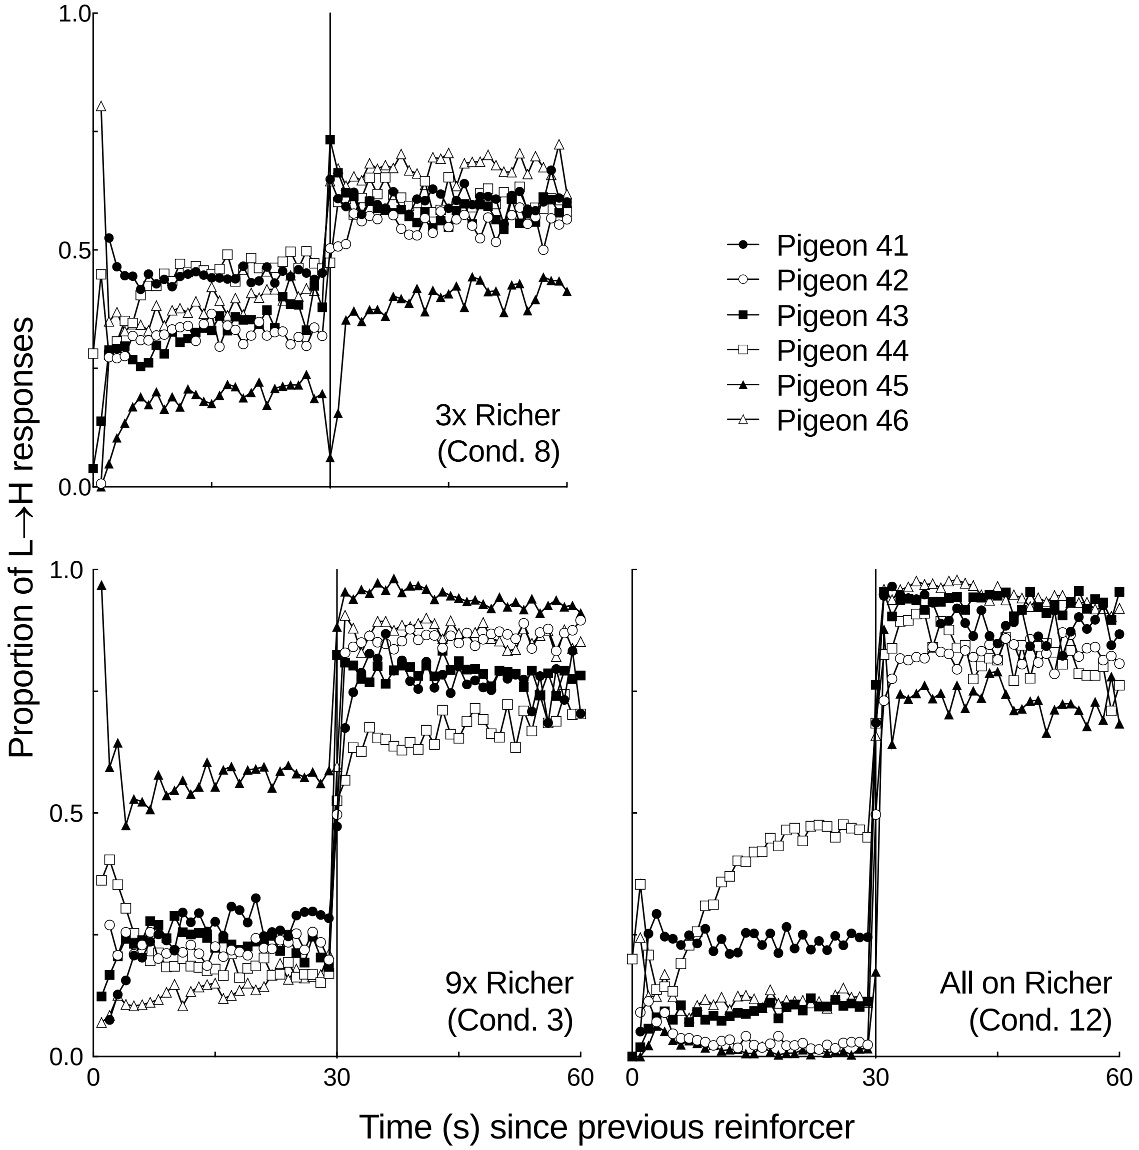

Figure A6.2. Proportion of responses made to the lower-to-higher (L→H) key across time since a reinforcer, for individual pigeons in Signal-Richer conditions. The solid vertical line denotes the reinforcer-ratio reversal time.

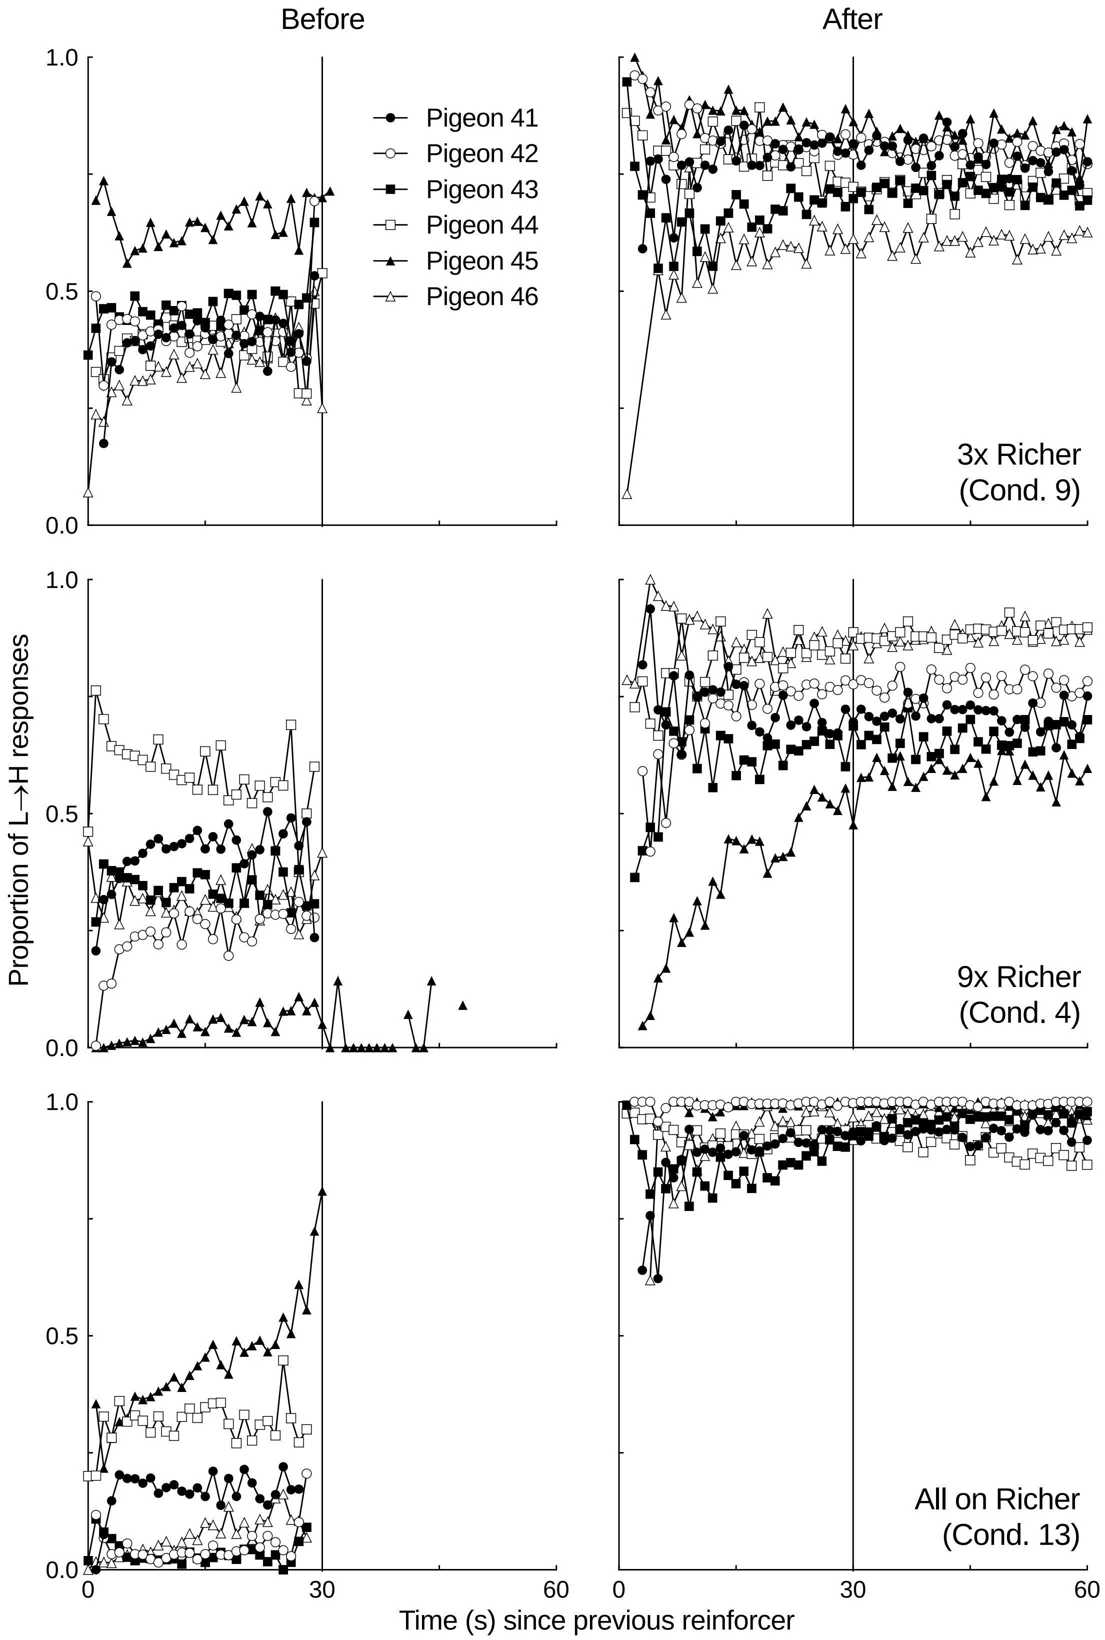

Figure A6.3. Proportion of responses made to the lower-to-higher (L→H) key across time since a reinforcer, for individual pigeons in Signal-Time conditions. Data are separated according to whether keylight colour signalled that the next reinforcer would occur before or after the reversal. The solid vertical line denotes the reinforcer-ratio reversal time.

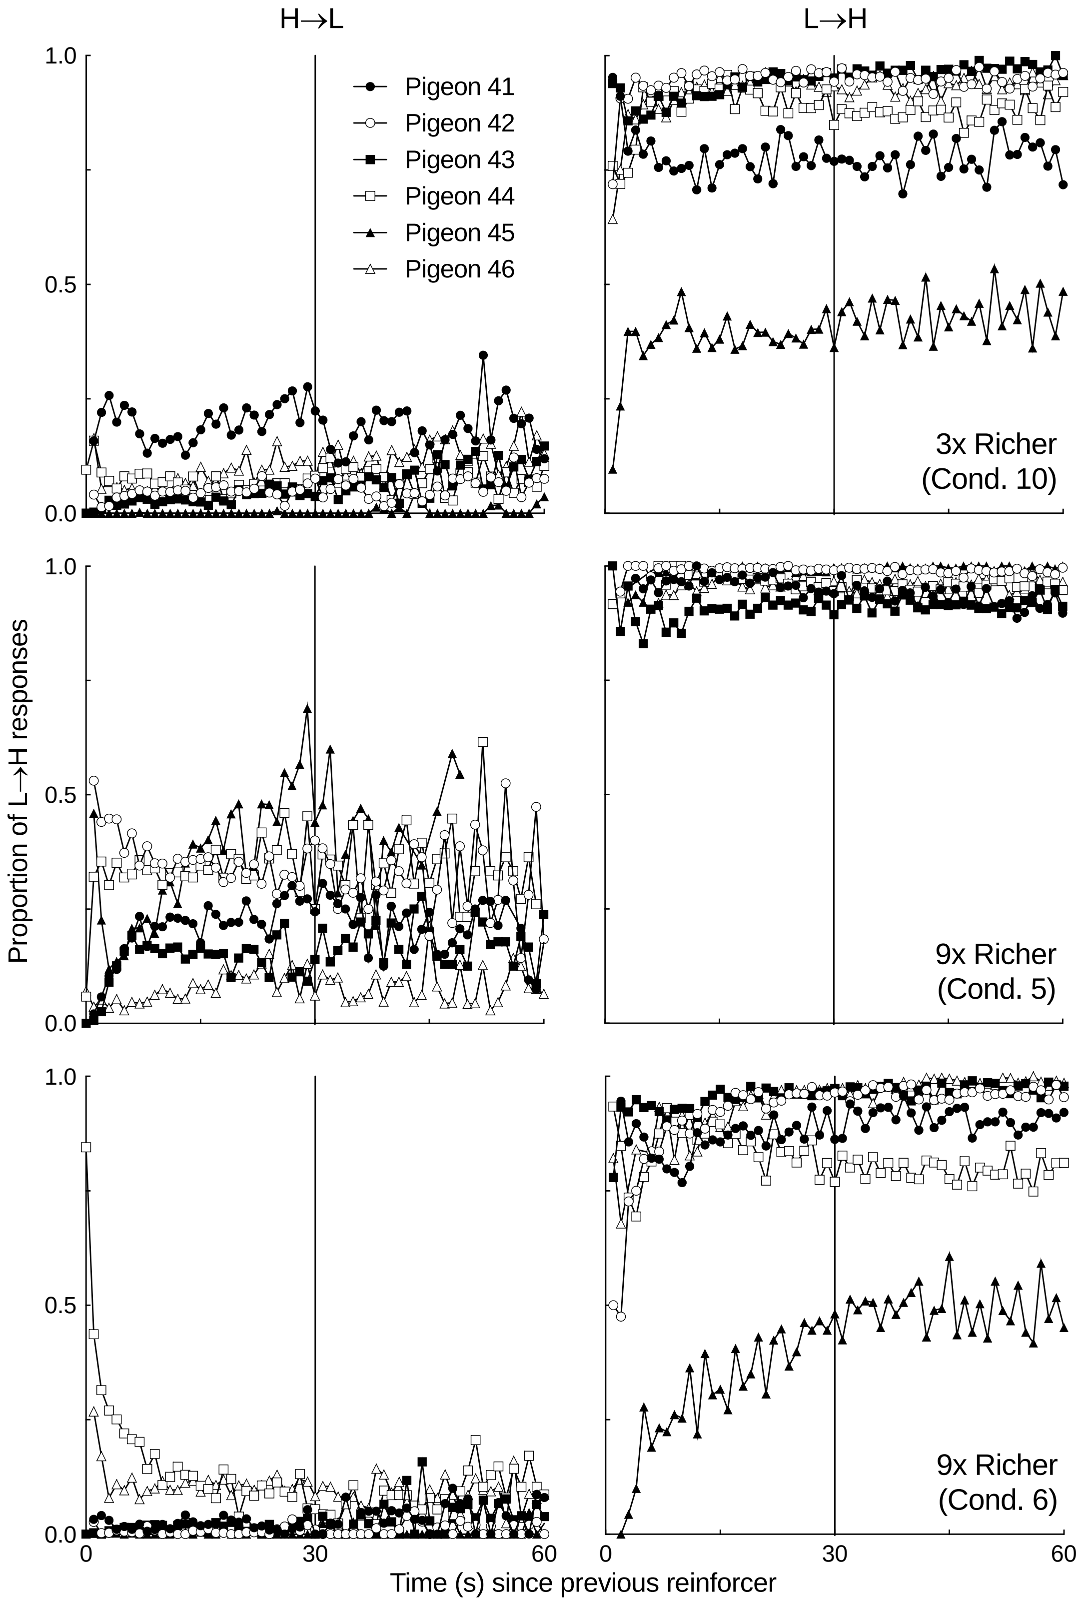

Figure A6.4. Proportion of responses made to the lower-to-higher (L→H) key across time since a reinforcer, for individual pigeons in Signal-Location conditions. Data are separated according to whether keylight colour signalled that the next reinforcer would on the H→L or L→H key. The solid vertical line denotes the reinforcer-ratio reversal time.

Appendix B6

Variance accounted for (VAC) for fits of Equation 6.1 to individual-pigeon data for No-Signal conditions

Table B6.1

Variance accounted for (VAC) for fits of Equation 6.1 to individual-pigeon data for No-Signal conditions.

Pigeon	Cond. 1 (9x Richer)	Cond. 2 (9x Richer)	Cond. 7 (3x Richer)	Cond. 11 (All on Richer)
41	.94	.81	.89	.93
42	.95	.91	.77	.98
43	.80	.88	.74	.97
44	.92	.95	.84	.94
45	.96	.97	.90	.99
46	.98	.96	.96	.99

CHAPTER 7

General Discussion

> This area [divided stimulus control] is new to the experimental analysis of behavior. It was first investigated and quantified by Shahan and Podlesnik (2006), with further papers by Shahan and Podlesnik (2007) and Davison and Elliffe (2010). The general finding is that discrimination (measured by log *d*; Davison & Tustin, 1978) between two stimuli located on one dimension is controlled by the frequency of reinforcers on this dimension compared with the frequency of reinforcers for discriminating between two stimuli on another dimension. Another way to put this is that the matching law determines the relative attention to (or control by) stimuli on one dimension versus on another dimension. **This result cries out for future research … We know a great deal about how reinforcer-frequency differences bias signal detection and conditional discrimination, but not nearly enough about how reinforcer frequency divides stimulus control.**
>
> Davison (2018b, p. 136, emphasis added)

This passage from a recent article by Davison (2018b) highlights the problem that this thesis aimed to address. Presently, few studies have investigated how relative reinforcer frequencies determine the division of control between multiple discriminative stimuli, and those that have arrange the same general delayed matching-to-sample (DMTS) procedure in which the dimensions of compound sample stimuli signal the likely location of future reinforcers (see Figure 2.1; e.g., Davison, 2018a; Davison & Elliffe, 2010; Podlesnik, Thrailkill, & Shahan, 2012; Shahan & Podlesnik, 2006, 2007). The present thesis extended such research, and asked whether relative reinforcers determine divided stimulus control in four experiments, each of which arranged a procedure in which the effects of relative reinforcers on divided stimulus control have not yet been investigated. Thus, this thesis examined the generality of the relation between relative reinforcers and divided stimulus control in several diverse situations. Table 7.1 provides a summary of the four experiments in this thesis, and their novel contributions.

Table 7.1.

Summary of main aims, findings, and novel contributions of the experiments in the present thesis.

	Main aim(s)	Main findings and novel contributions
Experiment 1 (Chapter 3; Gomes-Ng et al., in press)	- Investigated whether relative reinforcer rates determine divided stimulus control when stimulus-response-reinforcer relations are trained separately	- Relative reinforcers determined divided stimulus control between separately trained stimuli if they were spatially separated (Experiment 1B), but not if they were combined into a unified compound stimulus (Experiment 1A) - Behavioural expression of relation between relative reinforcers and divided stimulus control was evident in analyses of sample, not comparison, choice
Experiment 2 (Chapter 4)	- Investigated further the effects of spatial separation of compound-stimulus dimensions on behavioural expression of relation between relative reinforcers and divided stimulus control - Investigated whether Davison and Nevin's (1999) model could describe relation between relative reinforcers and divided control between spatially separated stimuli	- Compound-stimulus dimensions exerted discriminative control over comparison choice, suggesting that the results of Experiment 1B were related to the training procedure - Highly selective stimulus control, due to a change in the contingencies associated with one stimulus dimension from Experiment 1 to Experiment 2 - Relative reinforcer rates determined divided stimulus control between spatially separated stimuli that were trained together, but such effects were small, probably due to highly selective stimulus control - Davison and Nevin's (1999) model described the data well, suggesting that divided stimulus control depends on differential reinforcers, nondifferential reinforcers, and reinforcer-generalisation processes
Experiment 3 (Chapter 5)	- Investigated whether relative reinforcer probabilities determine divided control between stimuli signalling the time (2 or 8 s) to reinforcer delivery	- Relative reinforcer probabilities determined divided control between stimuli signalling the time of the next reinforcer - Elapsed time since trial start also controlled behaviour, and appeared to influence divided stimulus control; effects of relative reinforcer probabilities on divided stimulus control were stronger at earlier times than at later times - Elapsed time and compound-stimulus dimensions may compete for control over behaviour
Experiment 4 (Chapter 6; Gomes-Ng et al., 2018a)	- Investigated further the division of control between elapsed time and visual discriminative stimuli - Investigated whether reliability of elapsed time and visual stimuli determines divided stimulus control	- Stimulus control was divided between elapsed time and visual stimuli signalling next-reinforcer time or location, but the latter always exerted stronger control - Division of control depended on elapsed time, reliability of visual stimuli, and degree of conflict between elapsed time and visual stimuli - Changes in reliability of elapsed time had little effect on divided stimulus control, likely due to imperfect discrimination of elapsed time

7.1 Here or There? Relative Reinforcers and Divided Stimulus Control in Space

Experiments 1 and 2 (Chapters 3 and 4) asked whether relative reinforcer rates determine the division of control between two dimensions (colour and flash-frequency) of a compound stimulus when those dimensions signalled the location of future reinforcers in a symbolic DMTS task. Collectively, these experiments investigated how the training procedure (separately training stimuli versus training them together) and spatial configuration of the compound stimulus (spatially separated versus unified) modulate the effects of relative reinforcer rates on divided stimulus control. In conjunction with previous research (e.g., Davison & Elliffe, 2010; Shahan & Podlesnik, 2006, 2007), the results of Experiments 1 and 2 suggested that the type of training procedure and spatial configuration jointly determine whether a relation between relative reinforcers and divided stimulus control is observed, and may also impact the behavioural expression of that relation. Specifically, when stimuli were trained separately and then presented together, a systematic relation between relative reinforcer rates and divided stimulus control was observed only if the stimuli were spatially separated (Experiment 1). In contrast, relative reinforcer rates appeared to determine divided stimulus control when stimuli were trained together, regardless of stimulus configuration (Experiment 2, and previous research).

Additionally, the type of training procedure influenced the behavioural expression of this relation between relative reinforcers and divided stimulus control. In Experiment 1B, pigeons preferred to peck the stimulus dimension associated with the higher reinforcer rate during sample-stimulus presentations, but that dimension did not exert stronger *discriminative* control over behaviour. In contrast, the opposite was true in Experiment 2. Thus, the results of Experiments 1 and 2 show that responses made directly to a stimulus may provide some indication of which aspects of that stimulus control behaviour under some conditions, whereas this is not the case under other conditions (see also e.g., M. F. Brown, Cook, Lamb, & Riley, 1984; Castro & Wasserman, 2014, 2016, 2017; Dittrich, Rose, Buschmann, Bourdonnais, & Güntürkün, 2010; Furrow & LoLordo, 1975). The location of responses to a stimulus may be correlated with stimulus control when differential responding to different aspects of a stimulus facilitates discriminative performance (as in Experiment 1; see Chapter 3), whereas such a correlation may not eventuate if responding consistently to a specific aspect of a stimulus does not facilitate performance (as in Experiment 2; see Chapter 4).

The transition from Experiment 1 to Experiment 2 involved a change in the contingencies associated with one stimulus dimension, while the contingencies associated with the other dimension remained unchanged. The unchanged dimension exerted near-exclusive control over behaviour in Experiment 2, suggesting that the pigeons may not have discriminated the change in contingencies associated with the other dimension. Indeed, discriminating the change in contingency in Experiment 2 was probably difficult because every trial involved a compound stimulus comprising the unchanged and changed dimensions, and the dimension to match was unsignalled. As a result, obtained reinforcers would have appeared consistent with the previously learned contingencies according to the unchanged dimension, and inconsistent with the previously learned contingencies according to the changed dimension. Such apparent (i.e., discriminated by the subject; Davison & Nevin, 1999) differential and nondifferential reinforcers according to the unchanged and changed dimensions, respectively, may have then maintained the selective control by the unchanged dimension throughout Experiment 2 (see also Ryan, Hemmes, & Brown, 2011). Thus, the results of Experiment 2 highlight the importance of *contingency discriminability* (Davison & Nevin, 1999) – reinforcers will only divide stimulus control to the extent that they are discriminated to be differential with respect to more than one stimulus (see also Davison, 2018a; Davison & Elliffe, 2010).

Therefore, an unintended consequence of the transition from Experiment 1 to Experiment 2 was that it provided insight into the effects of relative reinforcer rates when stimulus control is highly selective. Although the pigeons' previous experience in Experiment 1 was responsible for the highly selective stimulus control in Experiment 2, this is not the only reason why stimulus control may be highly selective. Other reasons, some of which were outlined in Chapter 2, include the sensory capacities or phylogenetic history of the subject, inherent biases, and stimulus characteristics such as intensity, disparity, and spatial or temporal configuration. The results of Experiment 2 suggest that when stimulus control is highly selective, the effects of relative reinforcer rates on divided stimulus control are much attenuated; although changes in relative reinforcer rates did shift stimulus control in Experiment 2, such shifts were generally small, and the unchanged dimension continued to exert much stronger control.

7.2 Now or Then? Relative Reinforcers and Divided Stimulus Control in Time

Experiment 3 (Chapter 5) asked whether relative reinforcers determine divided stimulus control in a multiple peak procedure in which the dimensions of a compound stimulus (colour and flash-frequency) signalled the time (2 or 8 s) to probabilistic response-contingent reinforcer delivery. The pigeons were more likely to time the interval signalled by the dimension associated with the higher reinforcer probability, suggesting that relative reinforcer probabilities divide stimulus control in time. However, compared with previous research examining the effects of relative reinforcer probabilities on divided stimulus control in DMTS (e.g., Davison & Elliffe, 2010; Shahan & Podlesnik, 2006), the effects of relative reinforcer probabilities on divided stimulus control appeared to be smaller in Experiment 3. Thus, Experiment 3 is the first to extend previous research demonstrating the relation between relative reinforcers and divided stimulus control to temporal discriminations, and also suggests that the extent to which relative reinforcers determine divided stimulus control may differ in spatial and temporal discriminations.

Experiment 3 also shares similarities with studies of temporal averaging, which investigate divided control between separately trained stimuli in the multiple peak procedure (e.g., Matell & Kurti, 2014). However, whereas pigeons in Experiment 3 timed the interval(s) signalled by the compound-stimulus elements, rats in temporal averaging studies appear to "average" the intervals together. Thus, like Experiments 1 and 2, the results of Experiment 3 and of studies of temporal averaging suggest that whether stimulus-response-reinforcer relations are trained together (Experiments 2 and 3; also e.g., Davison & Elliffe, 2010; Shahan & Podlesnik, 2006) or apart (Experiment 1, and studies of temporal averaging) modulates the behavioural expression of divided stimulus control.

Furthermore, the extent of control by the compound-stimulus dimensions, and the effects of relative reinforcer probabilities on divided stimulus control, depended on elapsed time in Experiment 3. At earlier times in each trial, the stimulus dimensions appeared to control behaviour, and relative reinforcer probabilities divided such control. However, as time progressed, control by the stimulus dimensions appeared to weaken, and hence elapsed time appeared to exert primary control over behaviour at later times. This time-dependent change in the stimuli controlling behaviour (compound-stimulus dimensions and elapsed time) was consistent with a change in the certainty of the time of the next reinforcer – after 2 s had elapsed in a trial, the next reinforcer (if one was arranged) would definitely occur at 8 s, regardless of the compound stimulus. That is, elapsed time exerted stronger control over behaviour when it signalled the next-reinforcer time with more certainty. Therefore, Experiment 3 also demonstrated that when reinforcer availability changes across time, elapsed time and other informative non-temporal stimuli (e.g., compound visual stimuli) may compete for control over behaviour, and the degree of control by elapsed time and other stimuli may change across time depending on their relative abilities to predict future reinforcer availability.

7.3 Here or There, and Now or Then?

The fourth experiment in this thesis (Chapter 6) took a slightly different approach to the first three. Unlike the first three experiments, in which reinforcer availability varied across spatial locations (Experiments 1 and 2) or across time (Experiment 3), reinforcer availability varied across both space *and* time in Experiment 4. Two other points of difference between Experiment 4 and Experiments 1 to 3 are that Experiment 4 investigated (1) divided control between elapsed time and a visual stimulus, rather than divided control between two dimensions of a compound visual stimulus, and (2) the effects of relative *reliabilities*, rather than relative reinforcer probabilities.

In Experiment 4, the availability of reinforcers on two response alternatives changed across time since the previous reinforcer delivery, and additional keylight-colour stimuli signalled the time or location of future reinforcers. Across conditions, the reliability with which elapsed time or keylight colour signalled future reinforcer availability varied. Shifts in divided stimulus control were driven solely by changes in the reliability of keylight-colour stimuli; keylight-colour stimuli exerted stronger control over choice when they signalled future reinforcer availability more reliably, whereas changes in the reliability of elapsed time had little to no effect on control by elapsed time and by keylight colour. The latter result appeared to be related to the pigeons' imperfect discrimination of the reinforcer-ratio reversal (see also Cowie, Davison, & Elliffe, 2014, 2016b, 2016c, 2016d; Cowie, Elliffe, & Davison, 2013). Thus, like Experiment 2, Experiment 4 demonstrated that the effects of reinforcer contingencies on divided stimulus control depend on subjects’ ability to discriminate those contingencies; if such discrimination is imperfect, then changes in reinforcer contingencies will likely have smaller effects on divided stimulus control.

Although changes in the reliability with which keylight-colour stimuli signalled reinforcer availability produced changes in divided stimulus control in Experiment 4, such changes were small overall; in all conditions, control by keylight colour was always much stronger than control by elapsed time. This was most likely because keylight-colour stimuli always provided more reliable information about the next reinforcer than elapsed time. Thus, the results of Experiment 4 are similar to those of Experiment 2, in which stimulus control was highly selective and the effects of relative reinforcer rates on divided stimulus control were attenuated.

Additionally, like Experiment 3, the division of control between elapsed time and non-temporal stimuli changed across time in Experiment 4 – control by elapsed time was strongest immediately after a reinforcer delivery, and declined over time. This time-based change in divided stimulus control was probably partly related to the discriminability of elapsed time, as shorter durations are generally timed more accurately and precisely than longer durations (Gibbon, 1977). The divided control between elapsed time and keylight colour at early times also appeared to be related to the congruency of the information provided by both stimuli; control by elapsed time was stronger when the stimuli provided incongruent information than when they provided congruent information. Thus, in summary, Experiment 4 suggested that the effects of relative reinforcers on divided control between temporal and non-temporal stimuli depend jointly on contingency discriminability, the overall level of selectivity of stimulus control, and on the congruency of the information provided by the stimuli.

7.4 Theoretical Relevance

The four experiments in this thesis thus join a growing body of evidence demonstrating that relative reinforcer predictability determines divided stimulus control (see Chapter 2; e.g., Davison, 2018a; Davison & Elliffe, 2010; Matell & Kurti, 2014; Podlesnik et al., 2012; Shahan & Podlesnik, 2006, 2007). The experiments in this thesis extend such previous research by showing that whether stimuli are trained separately or together can impact the behavioural expression of such divided stimulus control. Furthermore, the experiments in this thesis also suggest that the effects of relative reinforcer predictability on divided stimulus control may be attenuated, or even eliminated, (1) by procedural variables such as stimulus configuration, type of training procedure, or changes in previously learned contingencies; (2) if stimulus control is highly selective overall; (3) when time-based changes in reinforcer availability occur; and (4) if subjects are unable to discriminate accurately the contingencies. In addition to these insights into the conditions under which relative reinforcers do and do not determine divided stimulus control, the experiments in this thesis are also relevant to several other theoretical issues in the experimental analysis of behaviour.

7.4.1 Divided Stimulus Control in Discrimination Learning

All basic behavioural experiments arrange situations in which multiple stimuli may control behaviour. Even those experiments that arrange a single, apparently simple stimulus – such as a coloured keylight – may engender divided stimulus control, because all stimuli are multi-dimensional, and other stimuli present in the experimental chamber may exert additional control[15]. Furthermore, all events occur in time, and hence elapsed time since important events (e.g., trial start or reinforcer delivery) may also control behaviour, even if

---

[15] I recall, from my time as an undergraduate student in a second-year psychology paper on Learning and Behaviour at The University of Auckland, the lecturer stating that “even the pile of pigeon poo in the corner of the operant cage may exert some control”. Though this is highly unlikely, it conveys the sentiment well.

other stimuli signal reinforcer availability more reliably. Indeed, a growing body of evidence suggests that control by elapsed time is pervasive in a range of conditional-discrimination tasks (see McMillan, Spetch, Sturdy, & Roberts, 2017 for a review).

The findings in the present thesis, along with those of Davison and Elliffe (2010) and Shahan and Podlesnik (2006, 2007), suggest that if other temporal or non-temporal stimuli are additionally correlated with reinforcers, then they may exert control over behaviour as well. Thus, as the examples provided in Chapter 1 of this thesis illustrate, the experimenter-arranged stimuli may not control behaviour exclusively, and in some cases, may not even exert the strongest control over behaviour (Ray, 1972). For example, in Experiment 1, behaviour was strongly controlled by the location of the stimulus elements, rather than by their identities. Studies of divided stimulus control thus highlight the importance of considering which stimuli in the environment are actually controlling a subject's behaviour. Such considerations will improve our understanding of how behaviour is learned, maintained, and generalised to novel contexts; of individual-subject and inter-species similarities and differences in behaviour; and of why discrimination learning is sometimes successful and other times unsuccessful (see also Pinto, Fortes, & Machado, 2017; Pinto & Machado, 2017 for similar discussions).

7.4.2 Stimulus Control and Attention

The term *attention* refers to the mechanisms by which some information in the environment is processed and other information is ignored (James, 1890). Skinner (1953) described attention as the "controlling relation" (p. 123) between an antecedent stimulus and behaviour. In other words, a stimulus will only control a subject's behaviour if that subject attended to it (see also Dinsmoor, 1985; Nevin, Davison, & Shahan, 2005; Ray, 1969, 1972; Reynolds, 1961). Thus, the generally accepted view is that "any observation of stimulus control is [also] an observation of attention[16]" (Johnson & Cumming, 1968, p. 157). As a result, studies of divided stimulus control can also provide insight into the processes that may govern the allocation of attention between stimuli (Shahan, 2013).

Dinsmoor (1985) speculated that "the processes involved in attention are not as readily accessible to observation as the more peripheral adjustments, but it is my hope and my working hypothesis that they obey similar principles" (p. 365). The experiments in this thesis, along with Davison and Elliffe's (2010), Podlesnik et al.'s (2012), and Shahan and Podlesnik's (2006, 2007) findings, suggest that subjects allocate more attention to stimuli associated with higher relative reinforcer rates, in the same way as they allocate more responses to alternatives associated with higher reinforcer rates in concurrent-choice schedules. Thus, in support of Dinsmoor's speculation, attending appears to be governed by the same variables that govern overt operant behaviour. Attending, then, may be considered to be an operant – a behaviour (albeit unobservable[17]) that is controlled by the consequences that follow (Skinner, 1937; Staddon & Cerutti, 2003).

7.4.3 Behavioural Measures of Divided Stimulus Control

Experiments 1 to 3 highlighted some of the complications associated with measuring divided stimulus control. Such complications have been acknowledged since the study of divided stimulus control began in behaviour analysis. For example, as described in Chapter 2, early research arranged go/no-go discrimination training with compound stimuli, and then

---

[16] It is important to note that the opposite statement – that the absence of stimulus control implies the absence of attention – is not necessarily true, because attention to a stimulus may not always be expressed behaviourally (see e.g., Born, Snow, & Herbert, 1969; Broomfield, McHugh, & Reed, 2008; Farthing & Hearst, 1970; Reed, Broomfield, McHugh, McCausland, & Leader, 2009; Wilkie & Masson, 1976).

[17] It is worth noting that the observing-response procedure, in which subjects emit a response (e.g., pecking a separate key) that produces a stimulus which signals the current reinforcer contingencies (Wyckoff, 1952), is one attempt to measure more directly the operant of 'attending' (see also Dinsmoor, 1985). However, most observing paradigms arrange single-element stimuli, and hence do not examine the division of attention between *multiple* stimuli (i.e., divided stimulus control). For this reason, further discussion of the relation between observing, attending, and stimulus control is beyond the scope of this thesis.

assessed divided stimulus control by presenting each compound-stimulus element individually and measuring overall response rates to each (e.g., Reynolds, 1961). However, subsequent studies suggested that post-discrimination tests with individual elements did not always reflect divided stimulus control during discrimination training with compound stimuli (e.g., Born, Snow, & Herbert, 1969; Farthing & Hearst, 1970; Wilkie & Masson, 1976).

Experiments 1 and 2 of this thesis emphasise the importance of choosing appropriate measures of divided stimulus control in the DMTS task, particularly when the compound-stimulus elements are spatially separated. In Experiment 1, control by the compound-stimulus dimensions was evident in analyses of sample choice, rather than in the more typical measure of comparison choice. Although responses directly to a stimulus may not necessarily always constitute control by that stimulus, this was not the case in Experiment 1, because comparison choice was controlled by the location of the last-pecked sample key. Hence, behaviour during sample-stimulus presentations, rather than during the comparison-choice phase, provided the best measure of divided stimulus control in Experiment 1. In contrast, the reverse was true in Experiment 2. Thus, appropriate measures of divided stimulus control depend on the task arranged, and the most conventional measure (e.g., log *d* or matching accuracy in DMTS) is not necessarily always the best measure.

Experiment 3 showed that divided stimulus control is much harder to quantify in temporal discriminations than in DMTS. The most obvious measures of temporal discrimination are peak times and spreads, which provide indications of timing accuracy and precision, respectively. However, the division of stimulus control may change across time, as was the case in both Experiments 3 and 4 of this thesis, and hence one single measure may fail to provide a complete picture of divided stimulus control. Indeed, peak times and spreads are typically obtained by fitting a Gaussian equation (see e.g., Equation 5.1) to data aggregated across many peak trials, and such data aggregation may mask patterns of responding in individual test trials (e.g., Cheng & Westwood, 1993; Church, Meck, & Gibbon, 1994). For example, in Experiment 3, start and/or stop times changed across different experimental conditions, but peak times remained relatively stable. Due to the complexities of temporal-discrimination data, a range of measures – such as those used in Experiment 3 – should be used to assess divided stimulus control in the multiple peak procedure.

In summary, Experiments 1 to 3 of this thesis demonstrated that using a range of behavioural measures to assess divided stimulus control can provide valuable insight into the exact nature of divided stimulus control. Indeed, behaviour itself is multi-dimensional, and hence stimuli may exert control over one or more dimensions of behaviour. This was shown clearly in Experiments 1 to 3, in which several different measures of behaviour (e.g., sample choice, comparison choice, peak times, normalised peak spreads, changepoints) were used to examine closely the effects of relative reinforcer frequencies on divided stimulus control. Therefore, researchers should consider the idiosyncrasies of the task arranged and the multi-dimensional nature of operant behaviour when deciding how best to measure divided stimulus control.

7.4.4 Theories of Divided Stimulus Control

The literature review of this thesis provided a brief overview of theories of divided stimulus control, which assume that the division of control between multiple stimuli depends jointly on the subject's sensory capacities, stimulus characteristics, and reinforcer predictability. Sensory capacities and stimulus characteristics determine the initial associability of each stimulus in respondent paradigms (Mackintosh, 1975) or overall bias towards some stimuli over others in operant paradigms (Davison & Elliffe, 2010; Shahan & Podlesnik, 2006, 2007), and relative reinforcer predictability modulates these initial levels of control by each stimulus. The experiments reported in this thesis are largely consistent with these theories of divided stimulus control. Because the experiments in this thesis used operant paradigms, the remainder of this section focuses on Shahan and Podlesnik's (2006, 2007) and Davison and Elliffe's (2010) models.

Shahan and Podlesnik's (2006, 2007) and Davison and Elliffe's (2010) models describe the relation between relative reinforcer rates and divided control between stimuli that signal the location of reinforcers. According to Shahan and Podlesnik, the effects of relative reinforcer rates on divided stimulus control can be described by a linear equation with the slope parameter representing *sensitivity* to changes in relative reinforcer rates (Lobb & Davison, 1975) and the y-intercept parameter representing inherent bias in stimulus control towards one dimension over the other. This model has been shown to describe the relation between relative reinforcer rates and divided stimulus control in DMTS (Podlesnik et al., 2012; Shahan & Podlesnik, 2006, 2007), and the data from Experiments 1 and 2 of this thesis also appear amenable to fits of Shahan and Podlesnik's model; in both experiments, the relation between relative reinforcer rates and log d_x was roughly linear, and each pigeon also displayed a bias in favour of one dimension over the other (Figures 3.3, 3.4, and 4.3).

However, although sensitivity to reinforcement captures the extent to which changes in relative reinforcer rates cause changes in divided stimulus control, it provides no underlying *mechanism* for sensitivity to reinforcement. To illustrate, if we had fit a linear equation to the data from Experiment 2, such fits would have produced relatively low sensitivity to reinforcement values, as changes in relative reinforcer rates had smaller-than-proportional changes in log d_x values (see Figure 4.3). But *why* did changes in relative reinforcer rates cause smaller-than-proportional changes in divided stimulus control in Experiment 2? Shahan and Podlesnik's model cannot answer this question. In contrast, the fits of Davison and Elliffe's (2010) model to data from Experiment 2 suggested that relative reinforcer rates had small effects on divided stimulus control because of a high degree of reinforcer generalisation (i.e., low log d_{sb} values; Table 4.2) between the elements along one stimulus dimension, compared with the other dimension. Thus, Davison and Elliffe's model proposes a mechanism – differential and nondifferential reinforcers, and reinforcer generalisation – to explain how and why relative reinforcer rates determine divided stimulus control. Davison and Elliffe's model therefore provides a more comprehensive understanding of how relative reinforcers divide stimulus control than Shahan and Podlesnik's model.

7.5 Applied Relevance

In addition to the theoretical issues outlines above, findings from the experiments in this thesis also have applied relevance. The complexity of natural environments means that most behaviours are likely to be controlled by multiple stimuli. The general introduction of this thesis (Chapter 1) provided several examples of such behaviours. Additionally, natural environments are constantly changing, and so divided stimulus control must shift dynamically – towards relevant stimuli and away from irrelevant stimuli – in response to changes in the environment. When multiple stimuli provide important information about how, when, or where to behave, the failure of behaviour to come under the control of those multiple relevant stimuli can have far-reaching detrimental consequences. For example, as outlined in Chapter 1, selective stimulus control has been linked to linguistic and social-communicative deficits, and can prevent learning of new behaviours (see Ploog, 2010 for a review). Thus, it is important that we understand how contingencies in the environment determine control by multiple stimuli, and how changes in such contingencies can produce shifts in divided stimulus control.

Selective stimulus control (termed *stimulus overselectivity* in the applied literature) has been found to be more likely in several populations, including preschool- and kindergarten-aged children (e.g., Bickel, Stella, & Etzel, 1984; Eimas, 1969), older adults aged 60 to 89 (e.g., Kelly, Leader, & Reed, 2016; McHugh & Reed, 2007; McHugh, Simpson, & Reed, 2010), individuals with learning disabilities (e.g., Allen & Fuqua, 1985; Dube et al., 2010; Matthews, Shute, & Rees, 2001), individuals with traumatic brain injury (e.g., Wayland & Taplin, 1982), and typically developing individuals under high cognitive load (e.g., Reed & Gibson, 2005; Reed, Petrina, & McHugh, 2011; Reed, Reynolds, & Fermandel, 2012). Stimulus overselectivity is especially prevalent in individuals with autism spectrum disorder (ASD; Ploog, 2010), and is thought to underpin several key symptoms of ASD, such as social-communicative deficits and the reduced ability to respond to stimuli in the environment (Lovaas, Koegel, & Schreibman, 1979; Ploog, 2010). Applied behaviour analysts often work with such populations, hence, stimulus overselectivity can pose a challenge to the success of behavioural interventions, as irrelevant stimuli, or only a small subset of relevant stimuli, may exert complete control over behaviour (e.g., Koegel & Rincover, 1976; Rincover & Koegel, 1975).

Therefore, research examining how reinforcer contingencies determine divided stimulus control may inform the development of new interventions, or the improvement of existing interventions, to remediate overselectivity. For example, recent research suggests that *revaluing* the overselected stimulus by reducing its reinforcer rate in separate training sessions, or by differentially reinforcing observing of underselected stimuli, may be a promising intervention to remediate stimulus overselectivity in applied settings (e.g., Broomfield, McHugh, & Reed, 2010; Dube & McIlvane, 1999; Kelly, Leader, & Reed, 2015; Reed, Broomfield, McHugh, McCausland, & Leader, 2009; Reed et al., 2012). The findings in this thesis provide further evidence that stimuli that are more highly correlated with reinforcers are more likely to exert stronger control over behaviour, and hence that manipulations of reinforcer contingencies may be used to reduce stimulus overselectivity, and more generally, to shift divided stimulus control away from some stimuli and towards other stimuli. Such manipulations of reinforcer contingencies may be more practical than manipulating stimulus characteristics; some stimuli in natural environments cannot be changed or removed, and so their characteristics cannot be manipulated to shift divided stimulus control. Under such conditions, it may be possible to manipulate the reinforcer rates associated with the different stimuli to shift divided stimulus control.

However, each experiment in this thesis also suggests that there are conditions under which reinforcer contingencies may have little to no effect on divided stimulus control. Some of these conditions – such as stimulus configuration or previously learned contingencies (Experiments 1 and 2), highly selective stimulus control (Experiment 2), imperfect contingency discriminability (Experiments 2 and 4), and time-based changes in reinforcer availability (Experiment 3 and 4) – are likely to occur in the natural world. Thus, future research should investigate the extent to which changes in reinforcer contingencies affect divided stimulus control under such conditions, and how the effects of relative reinforcers may be maximised. For example, in applied settings, changes in relative reinforcer rates may have larger effects on divided stimulus control if they are combined with verbal prompts, explicit requirements to observe (e.g., to point or look at) each stimulus, or salient changes in the physical appearance (e.g., intensity) of the stimuli (e.g., Dube et al., 2010; Dube & McIlvane, 1999; Walpole, Roscoe, & Dube, 2007).

Given the close relation between stimulus control and attention (Dinsmoor, 1985; Skinner, 1953), studies of divided stimulus control may also be informative with respect to the treatment of disorders that are characterised by deficits in attention, such as attention deficit hyperactivity disorder (ADHD). Additionally, individuals with substance-abuse disorders or addictions appear to allocate more attention to drug- or alcohol-related stimuli than to other non-drug or non-alcohol-related stimuli (e.g., Ehrman, Robbins, Bromwell, Lankford, Monterosso, & O’Brien, 2002; Field & Cox, 2008; Harris, Donohue, Ilse, Schoenfeld, Heinze, & Woldorff, 2018; Johnsen, Laberg, Cox, Vaksdal, & Hugdahl, 1994). Thus, research examining the variables that determine divided stimulus control may provide insight into the improvement or development of interventions aimed at treating ADHD or addiction and substance abuse disorders (Shahan, 2013; Shahan & Podlesnik, 2008) – for example, manipulating reinforcer contingencies may shift attention towards task-relevant stimuli and away from distractor or drug- or alcohol-related stimuli.

7.6 Future Directions

As noted by Davison (2018b) in the quoted passage at the beginning of this chapter, divided stimulus control is a ripe area for future research in behaviour analysis, and is an area of research with clear theoretical and applied relevance. The experiments in this thesis raise new questions about the effects of relative reinforcers on divided stimulus control. First, the experiments in this thesis represent a small selection of the diverse scenarios in which divided stimulus control may occur. Future research should continue to investigate the generality of the relation between relative reinforcers and divided stimulus control in a range of procedures beyond the typical DMTS task. Additionally, future studies could examine whether the relation between relative reinforcer rates and divided stimulus control extends to other dimensions of reinforcers (such as type, magnitude, or delay). Such research will provide further insight into the generality of the effects of relative reinforcers on divided stimulus control.

A procedural limitation of both Experiments 1 and 2 was the configuration of the sample and comparison keys on the operant panel; the sample keys were the top-left and top-right keys, and the left and right comparison keys were positioned directly below the sample keys. It is possible that the results of both experiments were idiosyncratic – they may have been related to this operant-panel configuration, rather than to the type of training procedure or stimulus configuration per se. A systematic replication of Experiments 1 and 2, with a different operant-panel configuration (e.g., with the sample stimulus presented on a single centred display), may help to clarify whether the results of those experiments were indeed idiosyncratic. Additionally, reinforcers were arranged independently in Experiment 2, and this may have served to maintain the carryover effects from the pigeons' previous experience in Experiment 1. Hence, a replication of Experiment 2 with dependently scheduled reinforcers (see Davison, 2018a; Davison & Elliffe, 2010) may enhance control by the reinforcer contingencies associated with each dimension. Alternatively, replicating Experiment 2 with naïve pigeons would eliminate any effects of previous experience on divided stimulus control, and hence provide greater insight into how spatial configuration and relative reinforcer rates *jointly* determine divided stimulus control in the DMTS task – for example, does greater spatial separation attenuate the effects of relative reinforcers on divided stimulus control? More generally, how do stimulus characteristics and relative reinforcer predictability interact to determine divided stimulus control?

The excellent fits of Davison and Nevin's (1999) model to the data from Experiment 2 (see also Davison & Elliffe, 2010) raise several interesting questions about quantitative models of divided stimulus control. For example, are the effects of explicitly arranged nondifferential reinforcers on a stimulus dimension similar to the effects of increasing reinforcers arranged on the other dimension in the DMTS task? That is, if we were to arrange nondifferential reinforcers on the colour dimension, would this have the same effect as increasing differential reinforcers on the flash-frequency dimension? Can the Davison-Nevin model describe such effects of explicitly arranged differential and nondifferential reinforcers on divided stimulus control? Investigating the answers to these questions will shed further light on how differential reinforcers, nondifferential reinforcers, and reinforcer generalisation between stimuli and between responses contributes to divided stimulus control in space.

Furthermore, can Davison and Nevin's (1999) model be extended to temporal discriminations, to account for the results of Experiment 3? Cowie et al. (2014, 2016d) have developed a model describing control by time-based changes in the reinforcer ratio in concurrent schedules that closely follows the logic of Davison and Nevin's model. In Cowie et al.'s model, discrimination of time-based changes in the reinforcer ratio in concurrent schedules is influenced by both reinforcer generalisation between alternatives and by reinforcer generalisation across time bins (i.e., a reinforcer obtained at one time – e.g., at 2 s – may generalise to surrounding times). Future research might apply Cowie et al.'s model to data from a multiple peak procedure by replacing the parameter that describes reinforcer generalisation between alternatives with a parameter describing reinforcer generalisation between stimuli.

Experiment 3 is the first to extend Shahan and Podlesnik's (2006, 2007) and Davison and Elliffe's (2010) findings to the temporal domain. Thus, it serves as a valuable starting point for future research examining the effects of relative reinforcers on divided stimulus control in time. Some questions worth addressing in future research include: Would similar results to Experiment 3 be obtained with a "delayed" multiple peak procedure similar to DMTS, in which the compound stimulus is removed before the start of the FI schedule? Do the effects of relative reinforcer probabilities on divided stimulus control depend on the absolute or relative values of the intervals signalled by the compound-stimulus dimensions? To what extent do frequency of experience with an interval and frequency of reinforcers at the end of the interval contribute to divided stimulus control? How do the effects of reinforcers on divided stimulus control change as temporal discrimination becomes less accurate or precise? Further investigation of how reinforcers influence temporal discrimination when multiple stimuli signal different intervals to reinforcer delivery will also help to elucidate the timing mechanisms that may be responsible for the results of

Experiment 3 and of studies of temporal averaging, as obtained reinforcers may play some role in determining these mechanisms.

Together, Experiments 3 and 4 raise questions about the variables that determine the extent of control by elapsed time and by other non-temporal stimuli. In Experiment 3, control by elapsed time appeared to be stronger when it signalled the time of the next reinforcer with more certainty (after 2 s had elapsed). If so, then manipulations that reduce this certainty, or that increase the reliability with which non-temporal stimuli signal reinforcer availability, should reduce control by elapsed time. However, the results of Experiment 4 suggested that the division of control between elapsed time and other non-temporal stimuli depends on subjects' ability to discriminate elapsed time, and may also be sensitive to procedural factors, such as the type of task arranged (e.g., a reinforcer-ratio reversal task vs. mid-session reversal task), whether changes in reinforcer availability occur on a larger or smaller time-scale (e.g., across a session vs. across seconds), the overall reinforcer rate, and trial-length variability. Future research should examine how these procedural factors modulate the division of control between elapsed time and non-temporal stimuli. In addition, control by elapsed time appeared to be stronger when elapsed time and non-temporal stimuli provided conflicting information in Experiment 4, and in Cowie, Davison and Elliffe (2017). However, neither of those experiments manipulated the degree of conflict explicitly. Thus, a future study could explicitly manipulate the degree of conflict between elapsed time and a non-temporal stimulus, to investigate the conditions under which control by elapsed time is likely to be pervasive.

At present, studies directly comparing divided stimulus control in spatial and temporal discrimination tasks are scarce, as research on divided stimulus control in space and time remains relatively separate. Thus, a question that remains concerns the similarities and differences in the processes that govern divided stimulus control in comparable spatial and temporal discrimination tasks. There may be some differences in the processes underlying behaviour in the presence of multiple stimuli that signal the location or time of future reinforcers, and depending on whether the stimuli are trained together or apart. Future research should arrange equivalent spatial and temporal discrimination tasks with the same subjects in order to examine more thoroughly the similarities and differences between divided stimulus control in space and time. Experiment 3 provides a useful starting point for such future research – the multiple peak procedure in Experiment 3 was an analogue of Davison and Elliffe's (2010) DMTS task used to study divided stimulus control. A future study might replicate Experiment 3 and compare the results with an equivalent DMTS task using the same subjects and experimental parameters (e.g., stimulus dimensions, range of relative reinforcer probabilities, order of conditions), in order to gain further insight divided stimulus control in space and time.

Finally, and to anchor the present thesis among recent research examining the effects of reinforcers on behaviour (see Cowie & Davison, 2016 for a review), investigating the *local* effects of reinforcer deliveries on divided stimulus control will shed further light on how reinforcers divide stimulus control. Recent concurrent-choice research suggests that reinforcer deliveries act primarily as *signals* that provide information about the likely time and/or location of future reinforcers. Such research may be extended to divided stimulus control – for example, if a reinforcer delivery signals more reinforcers for a colour discrimination in the DMTS task, will this shift divided stimulus control towards the colour dimension and away from other dimensions? Does divided stimulus control change dynamically, across time, in accordance with changes in the reinforcer contingencies signalled by each stimulus? For example, if reinforcers for a colour discrimination are more likely at early times and reinforcers for a flash-frequency discrimination are more likely at later times, will divided stimulus control shift across time in accordance with these contingencies? Both Experiments 3 and 4 showed that time-based changes in divided stimulus control can occur, and so we might expect divided stimulus control to follow time-based changes in contingency. Extending choice research on the signalling properties of reinforcers to divided stimulus control will help to reveal the mechanisms by which reinforcer deliveries divide stimulus control.

7.7 Conclusion

Control by multiple stimuli is crucial to successful learning, maintenance, and generalisation of behaviours. Previous research has shown that such divided stimulus control depends on a range of factors (see Chapter 2), including the subject’s sensory capacities and phylogenetic history; stimulus characteristics such as intensity, disparity, and spatial or temporal configuration; and the ability of each stimulus to predict future reinforcer availability. The present thesis provided further evidence that stimuli that are better predictors of future reinforcers exert stronger control over behaviour. However, in each experiment in this thesis, other factors – such as stimulus configuration, type of training procedure, orienting behaviours, changes in previously learned contingencies, changes in reinforcer availability across time, and contingency discriminability – also modulated the effects of reinforcer predictability on divided stimulus control (see Table 7.1). Thus, in addition to establishing the generality of the relation between relative reinforcers and divided stimulus control in novel procedures, this thesis also highlights some of the potential limits of this generality, and suggests that the effects of relative reinforcers on divided stimulus control may sometimes be modulated by procedural variables.

In closing, we return to the quote at the beginning of this thesis, in which Dinsmoor (1995a) emphasises the importance of understanding how antecedent stimuli control behaviour:

> …antecedent stimuli exert a pervasive influence on operant, as well as respondent, behavior. They serve as the crucial link between current behavior and past reinforcement. They are the direct and immediate determinants of what the individual does at any given moment.
>
> Dinsmoor (1995a, p. 67)

Although it has been 14 years since the publication of Dinsmoor's quote, many of the complexities of stimulus control remain under-investigated, and in some cases, *un*-investigated, in the experimental analysis of behaviour (see Cowie, 2018; Davison, 2018b for brief discussions). In contrast, much research has examined closely the complexities of consequence (e.g., reinforcer) control on behaviour (see Cowie, 2018; Cowie & Davison, 2016 for reviews). Thus, what is needed is a similarly thorough research programme examining the processes that govern complex stimulus control. In conjunction with recent research on consequence control, a close examination of the complexities of stimulus control will enable us to explain better, to predict more accurately, and to modify more efficiently and effectively, behaviour in both the laboratory and in the natural world. The present thesis is a step towards accomplishing this goal, and provides a platform for further research examining the role that reinforcer contingencies play in the division of control between multiple stimuli in space and time.

REFERENCES

Allan, R. W. (1993). Control of pecking response topography by stimulus-reinforcer and response-reinforcer contingencies. In H. P. Ziegler & H-J. Bischof (Eds.), *Vision, brain, and behavior in birds* (pp. 285-300). Cambridge, MA, USA: MIT Press.

Allen, K. D., & Fuqua, R. W. (1985). Eliminating selective stimulus control: A comparison of two procedures for teaching mentally retarded children to respond to compound stimuli. *Journal of Experimental Child Psychology, 39,* 55-71. https://doi.org/10.1016/0022-0965(85)90029-3

Anderson, B. A. (2013). A value-driven mechanism of attentional selection. *Journal of Vision, 13,* 1-16. https://doi.org/10.1167/13.3.7

Anderson, B. A. (2016). The attention habit: How reward learning shapes attentional selection. *Annals of the New York Academy of Sciences, 1369,* 24-39. https://doi.org/10.1111/nyas.12957

Anderson, B. A., Laurent, P. A., & Yantis, S. (2011). Value-driven attentional capture. *Proceedings of the National Academy of Sciences of the United States of America, 108,* 10367-10371. https://doi.org/10.1073/pnas.1104047108

Anderson, B. A., Laurent, P. A., & Yantis, S. (2013). Reward predictions bias attentional selection. *Frontiers in Human Neuroscience, 7:262.* https://doi.org/10.3389/fnhum.2013.00262

Aum, A., Brown, B. L., & Hemmes, N. S. (2007). The effect of intruded events on peak time: The role of reinforcement history during the intruded event. *Behavioural Processes, 74,* 187-197. https://doi.org/10.1016/j.beproc.2006.11.001

Aust, U., & Huber, L. (2002). Target-defining features in a “people-present/people-absent” discrimination task by pigeons. *Animal Learning & Behavior, 30,* 165-176. https://doi.org/10.3758/BF03192918

Awh, E., Armstrong, K. M., & Moore, T. (2006). Visual and oculomotor selection: Links, causes and implications for spatial attention. *Trends in Cognitive Science, 10,* 124-130. https://doi.org/10.1016/j.tics.2006.01.001

Baetu, I., Baker, A. G., Darredeau, C., & Murphy, R. A. (2005). A comparative approach to cue competition with one and two strong predictors. *Learning & Behavior, 33,* 160-171. https://doi.org/10.3758/BF03196060

Baker, T. W. (1968). Properties of compound conditioned stimuli and their components. *Psychological Bulletin, 70,* 611-625. https://doi.org/10.1037/h0026744

Balci, F., Gallistel, C. R., Allen, B. D., Frank, K. M., Gibson, J. M., & Brunner, D. (2009). Acquisition of peak responding: What is learned?. *Behavioural Processes, 80,* 67-75. https://doi.org/10.1016/j.beproc.2008.09.010

Baron, M. R. (1965). The stimulus, stimulus control and stimulus generalization. In D. I. Mostofsky (Ed.), *Stimulus generalization* (pp. 62-71). Stanford, CA, USA: Stanford University Press.

Barthold, C. H., & Egel, A. L. (2001). Stimulus overselectivity and generative language instruction for students with autism: An issue that needs to be revisited. *The Behavior Analyst Today, 2,* 33-39. https://doi.org/10.1037/h0099910

Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior, 22,* 231-242. https://doi.org/10.1901/jeab.1974.22-231

Berg, M. E., & Grace, R. C. (2011). Categorization of multidimensional stimuli by pigeons. *Journal of the Experimental Analysis of Behavior, 95,* 305-326. https://doi.org/10.1901/jeab.2010.94-305

Bickel, W. K., & Etzel, B. C. (1985). The quantal nature of controlling stimulus-response relations as measured in tests of stimulus generalization. *Journal of the Experimental Analysis of Behavior, 44,* 245-270. https://doi.org/10.1901/jeab.1985.44-245

Bickel, W. K., Stella, E., & Etzel, B. C. (1984). A reevaluation of stimulus overselectivity: Restricted stimulus control or stimulus control hierarchies. *Journal of Autism and Developmental Disorders, 14,* 137-157. https://doi.org/10.1007/BF02409657

Biebach, H., Gordijn, M., & Krebs, J. R. (1989). Time-and-place learning by garden warblers, Sylvia borin. *Animal Behaviour, 37,* 353-360. https://doi.org/10.1016/0003-3472(89)90083-3

Birnie-Selwyn, B., & Guerin, B. (1997). Teaching children to spell: Decreasing consonant cluster errors by eliminating selective stimulus control. *Journal of Applied Behavior Analysis, 30,* 69-91. https://doi.org/10.1901/jaba.1997.30-69

Bizo, L. A., & White, K. G. (1994). The behavioral theory of timing: Reinforcer rate determines pacemaker rate. *Journal of the Experimental Analysis of Behavior, 61,* 19-33. https://doi.org/10.1901/jeab.1994.61-19

Blaisdell, A. P., Denniston, J. C., & Miller, R. R. (1998). Temporal encoding as a determinant of overshadowing. *Journal of Experimental Psychology: Animal Behavior Processes, 24,* 72-83. https://doi.org/10.1037/0097-7403.24.1.72

Blaisdell, A. P., Schroeder, J. E., & Fast, C. D. (2018). Spatial integration during performance in pigeons. *Behavioural Processes, 154,* 73-80. https://doi.org/10.1016/j.beproc.2017.12.012

Blough, D. S. (1959). Delayed matching in the pigeon. *Journal of the Experimental Analysis of Behavior, 2,* 151-160. https://doi.org/10.1901/jeab.1959.2-151

Blough, D. S. (1969). Attention shifts in a maintained discrimination. *Science, 166,* 125-126. https://doi.org/10.1126/science.166.3901.125

Blough, D. S. (1975). Steady state data and a quantitative model of operant generalization and discrimination. *Journal of Experimental Psychology: Animal Behavior Processes, 1,* 3-21. https://doi.org/10.1037/0097-7403.1.1.3

Born, D. G., & Peterson, J. L. (1969). Stimulus control acquired by components of two color-form compound stimuli. *Journal of the Experimental Analysis of Behavior, 12,* 437-442. https://doi.org/10.1901/jeab.1969.12-437

Born, D. G., & Snow, M. E. (1970). Stimulus control by relevant, irrelevant, and redundant components of complex stimuli as assessed by two testing methods. *The Psychological Record, 20,* 311-319. https://doi.org/10.1007/BF03393947

Born, D. G., Snow, M. E., & Herbert, E. W. (1969). Conditional discrimination learning in the pigeon. *Journal of the Experimental Analysis of Behavior, 12,* 119-125. https://doi.org/10.1901/jeab.1969.12-119

Bouton, M. E., & Sunsay, C. (2003). Importance of trials versus accumulating time across trials in partially reinforced appetitive conditioning. *Journal of Experimental Psychology: Animal Behavior Processes, 29,* 62-77. https://doi.org/10.1037/0097-7403.29.1.62

Boutros, N., Davison, M., & Elliffe, D. (2009). Conditional reinforcers and informative stimuli in a constant environment. *Journal of the Experimental Analysis of Behavior, 91,* 41-60. https://doi.org/10.1901/jeab.2009.91-41

Boutros, N., Davison, M., & Elliffe, D. (2011). Contingent stimuli signal subsequent reinforcer ratios. *Journal of the Experimental Analysis of Behavior, 96,* 39-61. https://doi.org/10.1901/jeab.2011.96-39

Brodbeck, D. R. (1994). Memory for spatial and local cues: A comparison of a storing and a nonstoring species. *Animal Learning & Behavior, 22,* 119-133. https://doi.org/10.3758/BF03199912

Brodbeck, D. R., & Shettleworth, S. J. (1995). Matching location and color of a compound stimulus: Comparison of a food-storing and a nonstoring bird species. *Journal of Experimental Psychology: Animal Behavior Processes, 21,* 64-77. https://doi.org/10.1037/0097-7403.21.1.64

Broomfield, L., McHugh, L., & Reed, P. (2008). Re-emergence of under-selected stimuli, after the extinction of over-selected stimuli in an automated match to samples procedure. *Research in Developmental Disabilities, 29,* 503-512. https://doi.org/10.1016/j.ridd.2007.09.001

Broomfield, L., McHugh, L., & Reed, P. (2010). Factors impacting emergence of behavioral control by underselected stimuli in humans after reduction of control by overselected stimuli. *Journal of the Experimental Analysis of Behavior, 94,* 125-133. https://doi.org/10.1901/jeab.2010.94-125

Brown, B. L., Hemmes, N. S., & De Vaca, S. C. (1992). Effects of intratrial stimulus change on fixed-interval performance: The roles of clock and memory processes. *Animal Learning & Behavior, 20,* 83-93. https://doi.org/10.3758/BF03199948

Brown, M. F. (1987). Dissociation of stimulus compounds by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes, 13,* 80-91. https://doi.org/10.1037/0097-7403.13.1.80

Brown, M. F. (1991). A search for the locus of information overload in pigeon compound matching-to-sample performance. *Bulletin of the Psychonomic Society, 29,* 337-340. https://doi.org/10.3758/BF03333937

Brown, M. F., Cook, R. G., Lamb, M. R., & Riley, D. A. (1984). The relation between response and attentional shifts in pigeon compound matching-to-sample performance. *Animal Learning & Behavior, 12,* 41-49. https://doi.org/10.3758/BF03199811

Brown, M. F., & Morrison, S. K. (1990). Element and compound matching-to-sample performance in pigeons: The roles of information load and training history. *Journal of Experimental Psychology: Animal Behavior Processes, 16,* 185-192. https://doi.org/10.1037/0097-7403.16.2.185

Brown, P. L., & Jenkins, H. M. (1968). Auto-shaping of the pigeon’s key-peck. *Journal of the Experimental Analysis of Behavior, 11,* 1-8. https://doi.org/10.1901/jeab.1968.11-1

Buhusi, C. V., & Meck, W. H. (2005). What makes us tick? Functional and neural mechanisms of interval timing. *Nature Reviews Neuroscience, 6,* 755-765. https://doi.org/10.1038/nrn1764

Burnham, K., & Anderson, D. (2002). *Model selection and multimodel inference: A practical information-theoretic approach.* New York, NY, USA: Springer-Verlag.

Bushnell, M. C., & Weiss, S. J. (1977). The effect of reinforcement differences on choice and response distribution during stimulus compounding. *Journal of the Experimental Analysis of Behavior, 27,* 351-362. https://doi.org/10.1901/jeab.1977.27-351

Castro, L., & Wasserman, E. A. (2014). Pigeons’ tracking of relevant attributes in categorization learning. *Journal of Experimental Psychology: Animal Learning and Cognition, 40,* 195-211. https://doi.org/10.1037/xan0000022

Castro, L., & Wasserman, E. A. (2016). Attentional shifts in categorization learning: Perseveration but not learned irrelevance. *Behavioural Processes, 123,* 63-73. https://doi.org/10.1016/j.beproc.2015.11.001

Castro, L., & Wasserman, E. A. (2017). Feature predictiveness and selective attention in pigeons’ categorization learning. *Journal of Experimental Psychology: Animal Learning and Cognition, 43,* 231-242. https://doi.org/10.1037/xan0000146

Catania, A. C. (1970). Reinforcement schedules and psychophysical judgements: A study of some temporal properties of behavior. In W. N. Schoenfeld (Ed.), *The theory of reinforcement schedules* (pp. 1-42). New York, NY, USA: Appleton-Century-Crofts.

Chase, S. (1968). Selectivity in multidimensional stimulus control. *Journal of Comparative and Physiological Psychology, 66,* 787-792. https://doi.org/10.1037/h0026542

Chase, S., & Heinemann, E. G. (1972). Choices based on redundant information: An analysis of two-dimensional stimulus control. *Journal of Experimental Psychology, 92,* 161-175. https://doi.org/10.1037/h0032083

Chatlosh, D. L., & Wasserman, E. A. (1987). Delayed temporal discrimination in pigeons: A comparison of two procedures. *Journal of the Experimental Analysis of Behavior, 47,* 299-309. https://doi.org/10.1901/jeab.1987.47-299

Chelazzi, L., Perlato, A., Santandrea, E., & Della Libera, C. (2013). Rewards teach visual selective attention. *Vision Research, 85,* 58-72. https://doi.org/10.1016/j.visres.2012.12.005

Chen, X., McNamara, T. P., Kelly, J. W., & Wolbers, T. (2017). Cue combination in human spatial navigation. *Cognitive Psychology, 95,* 105-144. https://doi.org/10.1016/j.cogpsych.2017.04.003

Cheng, K., & Roberts, W. A. (1991). Three psychophysical principles of timing in pigeons. *Learning and Motivation, 22,* 112-128. https://doi.org/10.1016/0023-9690(91)90019-5

Cheng, K., Shettleworth, S. J., Huttenlocher, J., & Rieser, J. J. (2007). Bayesian integration of spatial information. *Psychological Bulletin, 133,* 625-637. https://doi.org/10.1037/0033-2909.133.4.625

Cheng, K., & Westwood, R. (1993). Analysis of single trials in pigeons' timing performance. *Journal of Experimental Psychology: Animal Behavior Processes, 19,* 56-67. https://doi.org/10.1037/0097-7403.19.1.56

Chiang, H-M., & Carter, M. (2008). Spontaneity of communication in individuals with autism. *Journal of Autism and Developmental Disorders, 38,* 693-705. https://doi.org/10.1007/s10803-007-0436-7

Chudasama, Y., & Muir, J. L. (1997). A behavioural analysis of the delayed non-matching to position task: The effects of scopolamine, lesions of the fornix and of the prelimbic region on mediating behaviours by rats. *Psychopharmacology, 134,* 73-82. https://doi.org/10.1007/s002130050427

Church, R. M., Guilhardi, P., Keen, R., MacInnis, M., & Kirkpatrick, K. (2003). Simultaneous temporal processing. In H. Helfrich (Ed.), *Time and mind II: Information processing perspectives* (pp. 3-19). Gottingen, Germany: Hogrefe & Huber Publishers.

Church, R. M., Meck, W. H., & Gibbon, J. (1994). Application of scalar timing theory to individual trials. *Journal of Experimental Psychology: Animal Behavior Processes, 20,* 135-155. https://doi.org/10.1037/0097-7403.20.2.135

Clayton, N. S., & Krebs, J. R. (1994). Memory for spatial and object-specific cues in food-storing and non-storing birds. *Journal of Comparative Physiology A, 174,* 371-379. https://doi.org/10.1007/BF00240218

Cohen, J. S., & Hachey, G. (1977). Utilization and dominance of spatially separated cues as a function of water deprivation. *Animal Learning & Behavior, 5,* 103-109. https://doi.org/10.3758/BF03209139

Colavita, F. B. (1974). Human sensory dominance. *Perception & Psychophysics, 16,* 409-412. https://doi.org/10.3758/BF03203962

Cole, R. P., Barnet, R. C., & Miller, R. R. (1995). Effect of relative stimulus validity: Learning or performance deficit?. *Journal of Experimental Psychology: Animal Behavior Processes, 21,* 293-303. https://doi.org/10.1037/0097-7403.21.4.293

Colombo, M., & Graziano, M. (1994). Effects of auditory and visual interference on auditory-visual delayed matching-to-sample in monkeys (Macaca fascicularis). *Behavioral Neuroscience, 108,* 636-639. https://doi.org/10.1037/0735-7044.108.3.636

Cook, R. G., Riley, D. A., & Brown, M. F. (1992). Spatial and configural factors in compound stimulus processing by pigeons. *Animal Learning & Behavior, 20,* 41-55. https://doi.org/10.3758/BF03199945

Cook, R. G., & Wixted, J. T. (1997). Same-different texture discrimination in pigeons: Testing competing models of discrimination and stimulus integration. *Journal of*

Experimental Psychology: Animal Behavior Processes, 23, 401-416. https://doi.org/10.1037/0097-7403.23.4.401

Couvillon, P. A., Tennant, W. A., & Bitterman, M. E. (1976). Intradimensional vs. extradimensional transfer in the discriminative learning of goldfish and pigeons. *Animal Learning & Behavior, 4,* 197-203. https://doi.org/10.3758/BF03214035

Cowie, S. (2018). Behavioral time travel: Control by past, present, and potential events. *Behavior Analysis: Research and Practice, 18,* 174-183. https://doi.org/10.1037/bar0000122

Cowie, S., Bizo, L. A., & White, K. G. (2016a). Reinforcer distributions affect timing in the free-operant psychophysical choice procedure. *Learning and Motivation, 53,* 24-35. https://doi.org/10.1016/j.lmot.2015.10.003

Cowie, S., & Davison, M. (2016). Control by reinforcers across time and space: A review of recent choice research. *Journal of the Experimental Analysis of Behavior, 105,* 246-269. https://doi.org/10.1002/jeab.200

Cowie, S., Davison, M., Blumhardt, L., & Elliffe, D. (2016b). Does overall reinforcer rate affect discrimination of time-based contingencies?. *Journal of the Experimental Analysis of Behavior, 105,* 393-408. https://doi.org/10.1002/jeab.204

Cowie, S., Davison, M., Blumhardt, L., & Elliffe, D. (2016c). Learning in a changing environment: Effects of the discriminability of visual stimuli and of time. *Learning and Motivation, 56,* 1-14. https://doi.org/10.1016/j.lmot.2016.08.001

Cowie, S., Davison, M., & Elliffe, D. (2011). Reinforcement? Food signals the time and location of future food. *Journal of the Experimental Analysis of Behavior, 96,* 63-86. https://doi.org/10.1901/jeab.2011.96-63

Cowie, S., Davison, M., & Elliffe, D. (2014). A model for food and stimulus changes that signal time-based contingency changes. *Journal of the Experimental Analysis of Behavior, 102,* 289-310. https://doi.org/10.1002/jeab.105

Cowie, S., Davison, M., & Elliffe, D. (2016d). A model for discriminating reinforcers in time and space. *Behavioural Processes, 127,* 62-73. https://doi.org/10.1016/j.beproc.2016.03.010

Cowie, S., Davison, M., & Elliffe, D. (2017). Control by past and present stimuli depends on the discriminated reinforcer differential. *Journal of the Experimental Analysis of Behavior, 108,* 184-203. https://doi.org/10.1002/jeab.268

Cowie, S., Elliffe, D., & Davison, M. (2013). Concurrent schedules: Discriminating reinforcer-ratio reversals at a fixed time after the previous reinforcer. *Journal of the Experimental Analysis of Behavior, 100,* 117-134. https://doi.org/10.1002/jeab.43

Cox, J. K., & D’Amato, M. R. (1982). Matching to compound samples by monkeys (Cebus apella): Shared attention or generalization decrement?. *Journal of Experimental Psychology: Animal Behavior Processes, 8,* 209-225. https://doi.org/10.1037/0097-7403.8.3.209

Cuell, S. F., Good, M. A., Dopson, J. C., Pearce, J. M., & Horne, M. R. (2012). Changes in attention to relevant and irrelevant stimuli during spatial learning. *Journal of Experimental Psychology: Animal Behavior Processes, 38,* 244-254. https://doi.org/10.1037/a0028491

Daan, S., & Koene, P. (1981). On the timing of foraging flights by oystercatchers, haematopus ostralegus, on tidal mudflats. *Netherlands Journal of Sea Research, 15,* 1-22. https://doi.org/10.1016/0077-7579(81)90002-8

Dall, S. R. X., Giraldeau, L-A., Olsson, O., McNamara, J. M., & Stephens, D. W. (2005). Information and its use by animals in evolutionary ecology. *Trends in Ecology & Evolution, 20,* 187-193. https://doi.org/10.1016/j.tree.2005.01.010

D'Amato, M. R., & Fazzaro, J. (1966). Attention and cue-producing behavior in the monkey. *Journal of the Experimental Analysis of Behavior, 9,* 469-473. https://doi.org/10.1901/jeab.1966.9-469

Davison, M. (2018a). Divided stimulus control: Which key did you peck, or what color was it?. *Journal of the Experimental Analysis of Behavior, 109,* 107-124. https://doi.org/10.1002/jeab.295

Davison, M. (2018b). The experimental analysis of behavior: Whence and thence?. *Behavior Analysis: Research and Practice, 18,* 134-143. https://doi.org/10.1037/bar0000101

Davison, M., & Baum, W. M. (2010). Stimulus effects on local preference: Stimulus-response contingencies, stimulus-food pairing, and stimulus-food correlation. *Journal of the Experimental Analysis of Behavior, 93,* 45-59. https://doi.org/10.1901/jeab.2010.93-45

Davison, M. C. (1976). Preference for fixed-interval schedules: Effects of unequal initial links. *Journal of the Experimental Analysis of Behavior, 25,* 371-376. https://doi.org/10.1901/jeab.1976.25-371

Davison, M., & Cowie, S. (2019). Timing or counting? Control by contingency reversals at fixed times or numbers of responses. *Journal of Experimental Psychology: Animal Learning and Cognition, 45,* 222-241. https://doi.org/10.1037/xan0000201

Davison, M. C., & Tustin, R. D. (1978). The relation between the generalized matching law and signal-detection theory. *Journal of the Experimental Analysis of Behavior, 29,* 331-336. https://doi.org/10.1901/jeab.1978.29-331

Davison, M., & Elliffe, D. (2010). Divided stimulus control: A replication and a quantitative model. *Journal of the Experimental Analysis of Behavior, 94,* 13-23. https://doi.org/10.1901/jeab.2010.94-13

Davison, M., & Ferguson, A. (1978). The effects of different component response requirements in multiple and concurrent schedules. *Journal of the Experimental Analysis of Behavior, 29,* 283-295. https://doi.org/10.1901/jeab.1978.29-283

Davison, M., & Jenkins, P. E. (1985). Stimulus discriminability, contingency discriminability, and schedule performance. *Animal Learning & Behavior, 13,* 77-84. https://doi.org/10.3758/BF03213368

Davison, M., & McCarthy, D. (1980). Reinforcement for errors in a signal-detection procedure. *Journal of the Experimental Analysis of Behavior, 34,* 35-47. https://doi.org/10.1901/jeab.1980.34-35

Davison, M., & Nevin, J. A. (1999). Stimuli, reinforcers, and behavior: An integration. *Journal of the Experimental Analysis of Behavior, 71,* 439-482. https://doi.org/10.1901/jeab.1999.71-439

De Corte, B. J., & Matell, M. S. (2016a). Interval timing, temporal averaging, and cue integration. *Current Opinion in Behavioral Sciences, 8,* 60-66. https://doi.org/10.1016/j.cobeha.2016.02.004

De Corte, B. J., & Matell, M. S. (2016b). Temporal averaging across multiple response options: Insight into the mechanisms underlying integration. *Animal Cognition, 19,* 329-342. https://doi.org/10.1007/s10071-015-0935-4

Delamater, A. R., & Nicolas, D-M. (2015). Temporal averaging across stimuli signaling the same or different reinforcing outcomes in the peak procedure. *International Journal of Comparative Psychology, 28*. https://escholarship.org/uc/item/0gc1q7zs

De Lillo, C., Spinozzi, G., Palumbo, M., & Giustino, G. (2011). Attention allocation modulates the processing of hierarchical visual patterns: A comparative analysis of capuchin monkeys (Cebus apella) and humans. *Journal of Experimental Psychology: Animal Behavior Processes, 37,* 341-352. https://doi.org/10.1037/a0022989

Della Libera, C., & Chelazzi, L. (2006). Visual selective attention and the effects of monetary rewards. *Psychological Science, 17,* 222-227. https://doi.org/10.1111/j.1467-9280.2006.01689.x

Della Libera, C., & Chelazzi, L. (2009). Learning to attend and to ignore is a matter of gains and losses. *Psychological Science, 20,* 778-784. https://doi.org/10.1111/j.1467-9280.2009.02360.x

Denniston, J. C., Miller, R. R., & Matute, H. (1996). Biological significance as a determinant of cue competition. *Psychological Science, 7,* 325-331. https://doi.org/10.1111/j.1467-9280.1996.tb00383.x

Denton, S. E., & Kruschke, J. K. (2006). Attention and salience in associative blocking. *Learning & Behavior, 34,* 285-304. https://doi.org/10.3758/BF03192884

Dias, R., Robbins, T. W., & Roberts, A. C. (1996). Primate analogue of the Wisconson card sorting test: Effects of excitotoxic lesions of the prefrontal cortex in the marmoset. *Behavioral Neuroscience, 110,* 872-886. https://doi.org/10.1037/0735-7044.110.5.872

Dickinson, A., Hall, G., & Mackintosh, N. J. (1976). Surprise and the attenuation of blocking. *Journal of Experimental Psychology: Animal Behavior Processes, 2,* 313-322. https://doi.org/10.1037/0097-7403.2.4.313

Dinsmoor, J. A. (1985). The role of observing and attention in establishing stimulus control. *Journal of the Experimental Analysis of Behavior, 43,* 365-381. https://doi.org/10.1901/jeab.1985.43-365

Dinsmoor, J. A. (1995a). Stimulus control: Part I. *The Behavior Analyst, 18,* 51-68. https://doi.org/10.1007/BF03392691

Dinsmoor, J. A. (1995b). Stimulus control: Part II. *The Behavior Analyst, 18,* 253-269. https://doi.org/10.1007/BF03392712

Dittrich, L., Rose, J., Buschmann, J-U. F., Bourdonnais, M., & Güntürkün, O. (2010). Peck tracking: A method for localizing critical features within complex pictures for pigeons. *Animal Cognition, 13,* 133-143. https://doi.org/10.1007/s10071-009-0252-x

Doran, J., & Holland, J. G. (1979). Control by stimulus features during fading. *Journal of the Experimental Analysis of Behavior, 31,* 177-187. https://doi.org/10.1901/jeab.1979.31-177

Du, Y., McMillan, N., Madan, C. R., Spetch, M. L., & Mou, W. (2017). Cue integration in spatial search for jointly learned landmarks but not for separately learned landmarks. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 43,* 1857-1871. https://doi.org/10.1037/xlm0000416

Dube, W. V., Dickson, C. A., Balsamo, L. M., O'Donnell, K. L., Tomanari, G. Y., Farren, K. M., …McIlvane, W. J. (2010). Observing behavior and atypically restricted stimulus control. *Journal of the Experimental Analysis of Behavior, 94,* 297-313. https://doi.org/10.1901/jeab.2010.94-297

Dube, W. V., McDonald, S. J., McIlvane, W. J., & Mackay, H. A. (1991). Constructed-response matching to sample and spelling instruction. *Journal of Applied Behavior Analysis, 24,* 305-317. https://doi.org/10.1901/jaba.1991.24-305

Dube, W. V., & McIlvane, W. J. (1997). Reinforcer frequency and restricted stimulus control. *Journal of the Experimental Analysis of Behavior, 68,* 303-316. https://doi.org/10.1901/jeab.1997.68-303

Dube, W. V., & McIlvane, W. J. (1999). Reduction of stimulus overselectivity with nonverbal differential observing responses. *Journal of Applied Behavior Analysis, 32,* 25-33. https://doi.org/10.1901/jaba.1999.32-25

Duncan, J. H., & Slotnick, B. M. (1990). Comparison of visual and olfactory reversal learning in pigeons. *Chemical Senses, 15,* 59-73. https://doi.org/10.1093/chemse/15.1.59

Dunlap, G., Koegel, R. L., & Burke, J. C. (1981). Educational implications of stimulus overselectivity in autistic children. *Remedial and Special Education, 2,* 37-49. https://doi.org/10.1177/074193258100200309

Edhouse, W. V., & White, K. G. (1988). Sources of proactive interference in animal memory. *Journal of Experimental Psychology: Animal Behavior Processes, 14,* 56-70. https://doi.org/10.1037/0097-7403.14.1.56

Egger, M. D., & Miller, N. E. (1962). Secondary reinforcement in rats as a function of information value and reliability of the stimulus. *Journal of Experimental Psychology, 64,* 97-104. https://doi.org/10.1037/h0040364

Egger, M. D., & Miller, N. E. (1963). When is a reward reinforcing? An experimental study of the information hypothesis. *Journal of Comparative and Physiological Psychology, 56,* 132-137. https://doi.org/10.1037/h0040744

Ehrman, R. N., Robbins, S. J., Bromwell, M. A., Lankford, M. E.,Monterosso, J. R., & O'Brien, C. P. (2002). Comparing attentional bias to smoking cues in current smokers, former smokers, and non-smokers using a dot-probe task. *Drug and Alcohol Dependence, 67,* 185-191. https://doi.org/10.1016/S0376-8716(02)00065-0

Eimas, P. D. (1969). Multiple-cue discrimination learning in children. *The Psychological Record, 19,* 417-424. https://doi.org/10.1007/BF03393868

Elliffe, D., & Elliffe, M. (2019). Rank-permutation tests for behavior anaysis, and a test for trend allowing unequal data numbers for each subject. *Journal of the Experimental Analysis of Behavior, 111,* 342-358. https://doi.org/10.1002/jeab.502

Emmerton, J. (1998). Numerosity differences and effects of stimulus density on pigeons' discrimination performance. *Animal Learning & Behavior, 26,* 243-256. https://doi.org/10.3758/BF03199218

Escobar, M., Matute, H., & Miller, R. R. (2001). Cues trained apart compete for behavioral control in rats: Convergence with the associative interference literature. *Journal of Experimental Psychology: General, 130,* 97-115. https://doi.org/10.1037/0096-3445.130.1.97

Fairhurst, S., Gallistel, C. R., & Gibbon, J. (2003). Temporal landmarks: Proximity prevails. *Animal Cognition, 6,* 113-120. https://doi.org/10.1007/s10071-003-0169-8

Fantino, E., & Davison, M. (1983). Choice: Some quantitative relations. *Journal of the Experimental Analysis of Behavior, 40,* 1-13. https://doi.org/10.1901/jeab.1983.40-1

Farthing, G. W. (1971). Discrimination of compound stimuli involving the presence or absence of a distinctive visual feature. *Journal of the Experimental Analysis of Behavior, 16,* 327-336. https://doi.org/10.1901/jeab.1971.16-237

Farthing, G. W. (1972). Overshadowing in the discrimination of successive compound stimuli. *Psychonomic Science, 28,* 29-32. https://doi.org/10.3758/BF03328652

Farthing, G. W., & Hearst, E. (1970). Attention in the pigeon: Testing with compounds or elements. *Learning and Motivation, 1,* 65-78. https://doi.org/10.1016.0023-9690(70)90129-3

Feldman, J. M. (1971). Added cue control as a function of reinforcement predictability. *Journal of Experimental Psychology, 91,* 318-325. https://doi.org/10.1037/h0031888

Feldman, J. M. (1975). Blocking as a function of added cue intensity. *Animal Learning & Behavior, 3,* 98-102. https://doi.org/10.3758/BF03209108

Ferster, C. B., & Skinner, B. F. (1957). *Schedules of reinforcement*. New York, NY, USA: Appleton-Century-Crofts.

Fetterman, J. G. (1993). Numerosity discrimination: Both time and number matter. *Journal of Experimental Psychology: Animal Behavior Processes, 19,* 149-164. https://doi.org/10.1037/0097-7403.19.2.149

Fetterman, J. G., & Killeen, P. R. (2011). Interactions of numerical and temporal stimulus characteristics on the control of response location by brief flashes of light. *Learning & Behavior, 39,* 191-201. https://doi.org/10.3758/s13420-011-0016-8

Field, M., & Cox, W. M. (2008). Attentional bias in addictive behaviors: A review of its development, causes, and consequences. *Drug and Alcohol Dependence, 97,* 1-20. https://doi.org/10.1016/j.drugalcdep.2008.03.030

Fields, L. (1978). Fading and errorless transfer in successive discriminations. *Journal of the Experimental Analysis of Behavior, 30,* 123-128. https://doi.org/10.1901/jeab.1978.30-123

Fields, L. (1979). Acquisition of stimulus control while introducing new stimuli in fading. *Journal of the Experimental Analysis of Behavior, 32,* 121-127. https://doi.org/10.1901/jeab.1979.32-121

Fields, L., Bruno, V., & Keller, K. (1976). The stages of acquisition in stimulus fading. *Journal of the Experimental Analysis of Behavior, 26,* 295-300. https://doi.org/10.1901/jeab.1976.26-295

Fox, A. E., & Kyonka, E. G. E. (2016). Effects of signaling on temporal control of behavior in response-initiated fixed intervals. *Journal of the Experimental Analysis of Behavior, 106,* 210-224. https://doi.org/10.1002/jeab.226

Fremouw, T., Herbranson, W. T., & Shimp, C. P. (1998). Priming of attention to local or global levels of visual analysis. *Journal of Experimental Psychology: Animal Behavior Processes, 24,* 278-290. https://doi.org/10.1037/0097-7403.24.3.278

Furrow, D. R., & LoLordo, V. M. (1975). Stimulus control in a discrimination based on a distinctive feature. *Journal of the Experimental Analysis of Behavior, 23,* 217-222. https://doi.org/10.1901/jeab.1975.23-217

Gadbois, S., & Reeve, C. (2014). Canine olfaction: Scent, sign, and situation. In A. Horowitz (Ed.), *Domestic Dog Cognition and Behavior* (pp. 3-29). Berlin, Germany: Springer. https://doi.org/10.1007/978-3-642-53994-7_1

Gallistel, C. R. (2003). Conditioning from an information processing perspective. *Behavioural Processes, 62,* 89-101. https://doi.org/10.1016/S0376-6357(03)00019-6

Gallistel, C. R., Fairhurst, S., & Balsam, P. (2004). The learning curve: Implications of a quantitative analysis. *Proceedings of the National Academy of Sciences of the United States of America, 101,* 13124-13131. https://doi.org/10.1073/pnas.0404965101

Gallistel, C. R., King, A. P., Gottlieb, D., Balci, F., Papachristos, E. B., Szalecki, M., & Carbone, K. S. (2007). Is matching innate?. *Journal of the Experimental Analysis of Behavior, 87,* 161-199. https://doi.org/10.1901/jeab.2007.92-05

Galloway, C. (1967). Modification of a response bias through differential amount of reinforcement. *Journal of the Experimental Analysis of Behavior, 10,* 375-382. https://doi.org/10.1901/jeab.1967.10-375

Galloway, C., & Petre, R. D. (1968). An analysis of the functional equivalence of stimulus class members. *Journal of Experimental Child Psychology, 6,* 384-390. https://doi.org/10.1016/0022-0965(68)90119-7

Galtress, T., & Kirkpatrick, K. (2009). Reward value effects on timing in the peak procedure. *Learning and Motivation, 40,* 109-131. https://doi.org/10.1016/j.lmot.2008.05.004

Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review, 84,* 279-325. https://doi.org/10.1037/0033-295X.84.3.279

Gibbon, J., Church, R. M., & Meck, W. H. (1984). Scalar timing in memory. *Annals of the New York Academy of Sciences, 423,* 52-77. https://doi.org/10.1111.j.1749-6632.1984.tb23417.x

Gliner, C. R., Pick, A. D., Pick, H. L., & Hales, J. J. (1969). A developmental investigation of visual and haptic preferences for shape and texture. *Monographs of the Society for Research in Child Development, 34,* 1-40.

Godfrey, R., & Davison, M. (1998). Effects of varying sample- and choice-stimulus disparity on symbolic matching-to-sample performance. *Journal of the Experimental Analysis of Behavior, 69,* 311-326. https://doi.org/10.1901/jeab.1998.69-311

Gomes-Ng, S., Elliffe, D., & Cowie, S. (2018a). Environment tracking and signal following in a reinforcer-ratio reversal procedure. *Behavioural Processes, 157,* 208-224. https://doi.org/10.1016/j.beproc.2018.10.001

Gomes-Ng, S., Elliffe, D., & Cowie, S. (2018b). Generalization of response patterns in a multiple peak procedure. *Behavioural Processes, 157,* 361-371. https://doi.org/10.1016/j.beproc.2018.07.012

Gomes-Ng, S., Elliffe, D., & Cowie, S. (in press). Relative reinforcer rates determine pigeons' attention allocation when separately trained stimuli are presented together. *Learning & Behavior.* https://doi.org/10.3758/s13420-019-00381-w

Grant, D. S. (1981). Stimulus control of information processing in pigeon short-term memory. *Learning and Motivation, 12,* 19-39. https://doi.org/10.1016/0023-9690(81)90023-0

Grant, D. S., & MacDonald, S. E. (1986). Matching to element and compound samples in pigeons: The role of sample coding. *Journal of Experimental Psychology: Animal Behavior Processes, 12,* 160-171. https://doi.org/10.1037/0097-7403.12.2.160

Grant, D. S., & Roberts, W. A. (1973). Trace interaction in pigeon short-term memory. *Journal of Experimental Psychology, 101,* 21-29. https://doi.org/10.1037/h0035532

Gray, V. A. (1976). Stimulus control of differential-reinforcement-of-low-rate responding. *Journal of the Experimental Analysis of Behavior, 25,* 199-207. https://doi.org/10.1901/jeab.1976.25-199

Grings, W. W., & Dawson, M. E. (1973). Complex variables in conditioning. In W. Prokasy (Ed.), *Electrodermal activity in psychological research* (pp. 203-154). New York, NY, USA: Academic Press.

Grondin, S. (2012). Violation of the scalar property for time perception between 1 and 2 seconds: Evidence from interval discrimination, reproduction, and categorization. *Journal of Experimental Psychology: Human Perception and Performance, 38,* 880-890. https://doi.org/10.1037/a0027188

Guttman, N., & Kalish, H. I. (1956). Discriminability and stimulus generalization. *Journal of Experimental Psychology, 51,* 79-88. https://doi.org/10.1037/h0046219

Hale, G. A., & Green, R. Z. (1976). Children's attention to stimulus components with variation in relative salience of components and degree of stimulus integration. *Journal of Experimental Child Psychology, 21,* 446-459. https://doi.org/10.1016/0022-0965(76)90073-4

Hall, G., & Channell, S. (1985). A comparison of intradimensional and extradimensional shift learning in pigeons. *Behavioural Processes, 10,* 295-295. https://doi.org/10.1016/0376-6357(85)90075-0

Hall, G., Mackintosh, N. J., Goodall, G., & Dal Martello, M. (1977). Loss of control by a less valid or by a less salient stimulus compounded with a better predictor of reinforcement. *Learning and Motivation, 8,* 145-158. https://doi.org/10.1016/0023-9690(77)90001-7

Harris, J. A., Donohue, S. E., Ilse, A., Schoenfeld, M. A., Heinze, H-J., & Woldorff, M. G. (2018). EEG measures of brain activity reveal that smoking-related images capture the attention of smokers outside of awareness. *Neuropsychologia, 111,* 324-333. https://doi.org/10.1016/j.neuropsychologia.2018.02.005

Hautus, M. J. (1995). Corrections for extreme proportions and their biasing effects on estimated values of d'. *Behavior Research Methods, Instruments, & Computers, 27,* 46-51. https://doi.org/10.3758/BF03203619

Hearst, E. (1988). The feature-positive effect in pigeons: Conditionality, overall predictiveness, and type of feature. *Bulletin of the Psychonomic Society, 26,* 73-76. https://doi.org/10.3758/BF03334866

Heinemann, E. G., Chase, S., & Mandell, C. (1968). Discriminative control of "attention". *Science, 160,* 553-554. https://doi.org/10.1126/science.160.3827.553

Herborn, K., Alexander, L., & Arnold, K. E. (2011). Colour cues or spatial cues? Context-dependent preferences in the European greenfinch (Carduelis chloris). *Animal Cognition, 14,* 269-277. https://doi.org/10.1007/s10071-010-0360-7

Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior, 4,* 267-272. https://doi.org/10.1901/jeab.1961.4-267

Heyman, G. M., Grisanzio, K. A., & Liang, V. (2016). Introducing a method for calculating the allocation of attention in a cognitive "two-armed bandit" procedure: Probability matching gives way to maximizing. *Frontiers in Psychology, 7:223.* https://doi.org/10.3389/fpsyg.2016.00223

Hinson, J. M., Cannon, C. B., & Tennison, L. R. (1999). Orthogonal stimulus variation and attention in dimensional contrast. *Animal Learning & Behavior, 27,* 181-189. https://doi.org/10.3758/BF03199674

Hodos, W. (1993). The visual capabilities of birds. In H. P. Ziegler & H-J. Bischof, *Vision, brain, and behavior in birds* (pp. 63-76). Cambridge, MA, USA: MIT Press.

Honig, W. K., & Urcuioli, P. J. (1981). The legacy of Guttman and Kalish (1956): 25 years of research on stimulus generalization. *Journal of the Experimental Analysis of Behavior, 36,* 405-445. https://doi.org/10.1901/jeab.1981.36-405

House, B. J., & Zeaman, D. (1963). Miniature experiments in the discrimination learning of retardates. In L. P. Lipsitt & C. C. Spiker (Eds.), *Advances in child development and behaviour* (pp. 313-374). New York, NY, USA: Academic Press.

Huguenin, N. H. (1987). Assessment of attention to complex cues in young children: Manipulating prior reinforcement histories of stimulus components. *Journal of Experimental Child Psychology, 44,* 283-303. https://doi.org/10.1016/0022-0965(87)90035-X

Huguenin, N. H., & Touchette, P. E. (1980). Visual attention in retarded adults: Combining stimuli which control incompatible behavior. *Journal of the Experimental Analysis of Behavior, 33,* 77-86. https://doi.org/10.1901/jeab.1980.33-77

Hunter, W. S. (1913). The delayed reaction in animals and children. *Behavioral Monographs, 2,* 1-86.

James, W. (1890). *The principles of psychology*. New York, NY, USA: Henry Holt and Company.

Jenkins, H. M., & Sainsbury, R. S. (1969). The development of stimulus control through differential reinforcement. In N. J. Mackintosh & W. K. Honig (Eds.), *Fundamental issues in associative learning* (pp. 123-167). Halifax, Canada: Dalhousie University Press.

Jenkins, H. M., & Sainsbury, R. S. (1970). Discrimination learning with the distinctive feature on positive or negative trials. In D. Mostofsky (Ed.), *Attention: Contemporary theory and analysis* (pp. 239-275). New York, NY, USA: Appleton-Century-Crofts.

Jennings, J. W., & Keefer, L. H. (1969). Olfactory learning set in two varieties of domestic rat. *Psychological Reports, 24,* 3-15. https://doi.org/10.2466.pr0.1969.24.1.3

Johnsen, B. H., Laberg, J. C., Cox, W. M., Vaksdal, A., & Hugdahl, K. (1994). Alcoholic subjects’ attentional bias in the processing of alcohol-related words. *Psychology of Addictive Behaviors, 8,* 111-115. https://doi.org/10.1037/0893-164X.8.2.111

Johnson, D. F. (1970). Determiners of selective stimulus control in the pigeon. *Journal of Comparative and Physiological Psychology, 70,* 298-307. https://doi.org/10.1037/h0020524

Johnson, D. F., & Cumming, W. W. (1968). Some determiners of attention. *Journal of the Experimental Analysis of Behavior, 11,* 157-166. https://doi.org/10.1901/jeab.1968.11-157

Kalafut, K. L., & Church, R. M. (2017). The quantification of behavior in the presence of compound stimuli. *Journal of Experimental Psychology: Animal Learning and Cognition, 43,* 96-108. https://doi.org/10.1037/xan0000128

Kamin, L. J. (1969). Predictability, surprise, attention, and conditioning. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior* (pp. 279-296). New York, NY, USA: Appleton-Century-Crofts.

Kehoe, E. J., & Gormezano, I. (1980). Configuration and combination laws in conditioning with compound stimuli. *Psychological Bulletin, 87,* 351-378. https://doi.org/10.1037/0033-2909.87.2.351

Kehoe, E. J., Horne, A. J., Horne, P. S., & Macrae, M. (1994). Summation and configuration between and within sensory modalities in classical conditioning of the rabbit. *Animal Learning & Behavior, 22,* 19-26. https://doi.org/10.3758/BF03199952

Kelly, M. P., Leader, G., & Reed, P. (2015). Stimulus over-selectivity and extinction-induced recovery of performance as a product of intellectual impairment and autism severity. *Journal of Autism and Developmental Disorders, 45,* 3098-3106. https://doi.org/10.1007/s10803-015-2466-x

Kelly, M. P., Leader, G., & Reed, P. (2016). Factors producing over-selectivity in older individuals. *Age, 38:63*. https://doi.org/10.1007/s11357-016-0026-x

Kendall, M. G. (1955). *Rank correlation methods*. London, United Kingdom: Charles Griffin.

Kirk, E. C., & Kay, R. F. (2004). The evolution of high visual acuity in the Anthropoidea. In C. F. Ross & R. F. Kay (Eds.), *Anthropoid origins* (pp. 539-602). Boston, MA, USA: Springer. https://doi.org/10.1007/978-1-4419-8873-7_20

Klein, S. B., Weston, D., McGee-Davis, T., & Cohen, L. (1984). The relative contributions of predictiveness and salience in flavor aversion learning. *Learning and Motivation, 15,* 188-202. https://doi.org/10.1016/0023-9690(84)90029-8

Klosterhalfen, S., Fischer, W., & Bitterman, M. E. (1978). Modification of attention in honey bees. *Science, 201,* 1241-1243. https://doi.org/10.1126/science.694513

Koegel, R. L., & Rincover, A. (1976). Some detrimental effects of using extra stimuli to guide learning in normal and autistic children. *Journal of Abnormal Child Psychology, 4,* 59-71. https://doi.org/10.1007/BF00917605

Kraemer, P. J. (1993). Attention to uninformative features expressed by pigeons in a duration matching-to-sample task. *Behavioural Processes, 30,* 291-297. https://doi.org/10.1016/0376-6357(93)90141-D

Kraemer, P. J., Mazmanian, D. S., & Roberts, W. A. (1987). Simultaneous processing of visual and spatial stimuli in pigeons. *Animal Learning & Behavior, 15,* 417-422. https://doi.org/10.3758/BF03205050

Kraemer, P. J., & Roberts, W. A. (1985). Short-term memory for simultaneously presented visual and auditory signals in the pigeon. *Journal of Experimental Psychology: Animal Behavior Processes, 11,* 137-151. https://doi.org/10.1037/0097-7403.11.2.137

Kraemer, P. J., & Roberts, W. A. (1987). Restricted processing of simultaneously presented brightness and pattern stimuli in pigeons. *Animal Learning & Behavior, 15,* 15-24. https://doi.org/10.3758/BF03204899

Krebs, R. M., Boehler, C. N., & Woldorff, M. G. (2010). The influence of reward associations on conflict processing in the Stroop task. *Cognition, 117,* 341-347. https://doi.org/10.1016/j.cognition.2010.08.018

Kurti, A. N., Swanton, D. N., & Matell, M. S. (2013). The potential link between temporal averaging and drug-taking behavior. In V. Arstila & D. Lloyd (Eds.), *Subjective time: The philosophy, psychology, and neuroscience of temporality* (pp. 599-620). Cambridge, MA, USA: MIT Press.

Lamb, M. R. (1988). Selective attention: Effects of cuing on the processing of different types of compound stimuli. *Journal of Experimental Psychology: Animal Behavior Processes, 14,* 96-104. https://doi.org/10.1037/0097-7403.14.1.96

Lamb, M. R. (1991). Attention in humans and animals: Is there a capacity limitation at the time of encoding?. *Journal of Experimental Psychology: Animal Behavior Processes, 17,* 45-54. https://doi.org/10.1037/0097-7403.17.1.45

Lamb, M. R., & Riley, D. A. (1981). Effects of element arrangement on the processing of compound stimuli in pigeons (Columba livia). *Journal of Experimental Psychology: Animal Behavior Processes, 7,* 45-48. https://doi.org/10.1037/0097-7403.7.1.45

Lander, D. G., & Irwin, R. J. (1968). Multiple schedules: Effects of the distribution of reinforcements between components on the distribution of responses between components. *Journal of the Experimental Analysis of Behavior, 11,* 517-524. https://doi.org/10.1901/jeab.1968.11-517

Langley, C. M., & Riley, D. A. (1993). Limited capacity information processing and pigeon matching-to-sample: Testing alternative hypotheses. *Animal Learning & Behavior, 21,* 226-232. https://doi.org/10.3758/BF03197986

Lashley, K. S. (1938). The mechanism of vision: XV. Preliminary studies of the rat's capacity for detail vision. *The Journal of General Psychology, 18,* 123-193. https://doi.org/10.1080/00221309.1938.9709894

Lea, S. E. G., Pothos, E. M., Wills, A. J., Leaver, L. A., Ryan, C. M. E., & Meier, C. (2018). Multiple feature use in pigeons' category discrimination: The influence of stimulus set structure and the salience of stimulus differences. *Journal of Experimental Psychology: Animal Learning and Cognition, 44,* 114-127. https://doi.org/10.1037/xan0000169

Lea, S. E. G., Wills, A. J., Leaver, L. A., Ryan, C. M. E., Bryant, C. M. L., & Millar, L. (2009). A comparative analysis of the categorization of multidimensional stimuli: II. Strategic information search in humans (Homo sapiens) but not in pigeons (Columba livia). *Journal of Comparative Psychology, 123,* 406-420. https://doi.org/10.1037/a0016851

Leader, G., Loughnane, A., McMoreland, C., & Reed, P. (2009). The effect of stimulus salience on over-selectivity. *Journal of Autism and Developmental Disorders, 39,* 330-338. https://doi.org/10.1007/s10803-008-0626-y

Leak, T. M., & Gibbon, J. (1995). Simultaneous timing of multiple intervals: Implications of the scalar property. *Journal of Experimental Psychology: Animal Behavior Processes, 21,* 3-19. https://doi.org/10.1037/0097-7403.21.1.3

Lechelt, D. P., & Spetch, M. L. (1997). Pigeons' use of landmarks for spatial search in a laboratory arena and in digitized images of the arena. *Learning and Motivation, 28,* 424-445. https://doi.org/10.1006/lmot.1997.0975

Legge, E. L. G., Madan, C. R., Spetch, M. L., & Ludvig, E. A. (2016). Multiple cue use and integration in pigeons (Columba livia). *Animal Cognition, 19,* 581-591. https://doi.org/10.1007/s10071-016-0963-8

Leith, C. R., & Maki, W. S. (1975). Attention shifts during matching-to-sample performance in pigeons. *Animal Learning & Behavior, 3,* 85-89. https://doi.org/10.3758/BF03209105

Leith, C. R., & Maki, W. S. (1977). Effects of compound configuration on stimulus selection in the pigeon. *Journal of Experimental Psychology: Animal Behavior Processes, 3,* 229-239. https://doi.org/10.1037/0097-7403.3.3.229

Le Pelley, M. E., Beesley, T., & Griffiths, O. (2014). Relative salience versus relative validity: Cue salience influences blocking in human associative learning. *Journal of Experimental Psychology: Animal Learning and Cognition, 40,* 116-132. https://doi.org/10.1037/xan0000006

Lionello, K. M., & Urcuioli, P. J. (1998). Control by sample location in pigeons' matching to sample. *Journal of the Experimental Analysis of Behavior, 70,* 235-251. https://doi.org/10.1901/jeab.1998.70-235

Lobb, B., & Davison, M. C. (1975). Performance in concurrent interval schedules: A systematic replication. *Journal of the Experimental Analysis of Behavior, 24,* 191-197. https://doi.org/10.1901/jeab.1975.24-191

Lobb, B., & Davison, M. C. (1977). Multiple and concurrent schedule performance: Independence from concurrent and successive schedule contexts. *Journal of the*

Experimental Analysis of Behavior, 28, 27-39. https://doi.org/10.1901/jeab.1977.28-27

Lovaas, O. I., Koegel, R. L., & Schreibman, L. (1979). Stimulus overselectivity in autism: A review of research. *Psychological Bulletin, 86,* 1236-1254. https://doi.org/10.1037/0033-2909.86.6.1236

Lovaas, O. I., Schreibman, L., Koegel, R., & Rehm, R. (1971). Selective responding by autistic children to multiple sensory input. *Journal of Abnormal Psychology, 77,* 211-222. https://doi.org/10.1037/h0031015

Lovejoy, E., & Russell, D. G. (1967). Suppression of learning about a hard cue by the presence of an easy cue. *Psychonomic Science, 8,* 365-366. https://doi.org/10.3758/BF03332242

MacDonald, H., & Roberts, W. A. (2018). Cognitive flexibility and dual processing in pigeons: Temporal and contextual control of midsession reversal. *Journal of Experimental Psychology: Animal Learning and Cognition, 44,* 149-161. https://doi.org/10.1037/xan0000161

Machado, A., & Cevik, M. (1998). Acquisition and extinction under periodic reinforcement. *Behavioural Processes, 44,* 237-262. https://doi.org/10.1016/S0376-6357(98)00052-7

Mackintosh, N. J. (1971). An analysis of overshadowing and blocking. *Quarterly Journal of Experimental Psychology, 23,* 118-125. https://doi.org/10.1080/00334447143000121

Mackintosh, N. J. (1974). *The psychology of animal learning*. Oxford, United Kingdom: Academic Press.

Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review, 82,* 276-298. https://doi.org/10.1037/h0076778

Mackintosh, N. J. (1976). Overshadowing and stimulus intensity. *Animal Learning & Behavior, 4,* 186-192. https://doi.org/10.3758/BF03214033

Mackintosh, N. J., Bygrave, D. J., & Picton, B. M. B. (1977). Locus of the effect of a surprising reinforcer in the attenuation of blocking. *Quarterly Journal of Experimental Psychology, 29,* 327-336. https://doi.org/10.1080/14640747708400608

Mackintosh, N. J., Dickinson, A., & Cotton, M. M. (1980). Surprise and blocking: Effects of the number of compound trials. *Animal Learning & Behavior, 8,* 387-391. https://doi.org/10.3758/BF03199622

Mackintosh, N. J., & Little, L. (1969). Intradimensional and extradimensional shift learning by pigeons. *Psychonomic Science, 14,* 5-6. https://doi.org/10.3758/BF03336395

Mackintosh, N. J., & Turner, C. (1971). Blocking as a function of novelty of CS and predictability of US. *Quarterly Journal of Experimental Psychology, 23,* 359-366. https://doi.org/10.1080/14640747109400245

Maki, W. S., & Leith, C. R. (1973). Shared attention in pigeons. *Journal of the Experimental Analysis of Behavior, 19,* 345-349. https://doi.org/10.1901/jeab.1973.19-345

Maki, W. S., & Leuin, T. C. (1972). Information-processing by pigeons. *Science, 176,* 535-536. https://doi.org/10.1126/science.176.4034.535

Marshall, A. T., & Kirkpatrick, K. (2015). Everywhere and everything: The power and ubiquity of time. *International Journal of Comparative Psychology, 28.* https://escholarship.org/uc/item/8hg831n3

Matell, M. S., De Corte, B. J., Kerrigan, T., & DeLussey, C. M. (2016). Temporal averaging in response to change. *Timing & Time Perception, 4,* 223-247. https://doi.org/10.1163/22134468-00002068

Matell, M. S., Kim, J. S., & Hartshorne, L. (2014). Timing in a variable interval procedure: Evidence for a memory singularity. *Behavioural Processes, 101,* 49-57. https://doi.org/10.1016/j.beproc.2013.08.010

Matell, M. S., & Kurti, A. N. (2014). Reinforcement probability modulates temporal memory selection and integration processes. *Acta Psychologica, 147,* 80-91. https://doi.org/10.1016.j.actpsy.2013.06.006

Matthews, B., Shute, R., & Rees, R. (2001). An analysis of stimulus overselectivity in adults with autism. *Journal of Intellectual & Developmental Disability, 2,* 161-176. https://doi.org/10.1080/13668250020054477

Matute, H., & Pineño, O. (1998). Stimulus competition in the absence of compound conditioning. *Animal Learning & Behavior, 26,* 3-14. https://doi.org/10.3758/BF03199157

Mazur, J. E. (1997). Effects of rate of reinforcement and rate of change on choice behaviour in transition. *Quarterly Journal of Experimental Psychology, 50,* 111-128. https://doi.org/10.1080/713932646

McCausland, D. F., Menzer, G. W., Dempsey, T. K., & Birkimer, J. C. (1967). Response-contingent and non-contingent informative and redundant secondary reinforcers. *Psychonomic Science, 8,* 293-294. https://doi.org/10.3758/BF03331668

McHugh, L., Simpson, A., & Reed, P. (2010). Mindfulness as a potential intervention for stimulus over-selectivity in older adults. *Research in Developmental Disabilities, 31,* 178-184. https://doi.org/10.1016/j.ridd.2009.08.009

McHugh, L., & Reed, P. (2007). Age trends in stimulus overselectivity. *Journal of the Experimental Analysis of Behavior, 88,* 369-380. https://doi.org/10.1901/jeab.2007.88-369

McLean, A. P., & White, K. G. (1983). Temporal constraint on choice: Sensitivity and bias in multiple schedules. *Journal of the Experimental Analysis of Behavior, 39,* 405-426. https://doi.org/10.1901/jeab.1983.39-405

McLinn, C. M., & Stephens, D. W. (2006). What makes information valuable: Signal reliability and environmental uncertainty. *Animal Behaviour, 71,* 1119-1129. https://doi.org/10.1016/j.anbehav.2005.09.006

McLinn, C. M., & Stephens, D. W. (2010). An experimental analysis of receiver economics: Cost, reliability and uncertainty interact to determine a signal's value. *Oikos, 119,* 254-263. https://doi.org/10.1111/j.1600-0706.2009.17756.x

McMillan, N., & Roberts, W. A. (2013). Interval timing under variations in the relative validity of temporal cues. *Journal of Experimental Psychology: Animal Behavior Processes, 39,* 334-341. https://doi.org/10.1037/a0032470

McMillan, N., Spetch, M. L., Sturdy, C. B., & Roberts, W. A. (2017). It's all a matter of time: Interval timing and competition for stimulus control. *Comparative Cognition & Behavior Reviews, 12,* 83-103. https://doi.org/10.3819/ccbr.2017.120007

McMillan, N., Sturdy, C. B., Pisklak, J. M., & Spetch, M. L. (2016). Pigeons perform poorly on a midsession reversal task without rigid temporal regularity. *Animal Cognition, 19,* 855-859. https://doi.org/10.1007/s10071-016-0962-9

McSweeney, F. K., Farmer, V. A., Dougan, J. D., & Whipple, J. E. (1986). The generalized matching law as a description of multiple-schedule responding. *Journal of the Experimental Analysis of Behavior, 45,* 83-101. https://doi.org/10.1901/jeab.1986.45-83

Meck, W. H., Church, R. M., & Gibbon, J. (1985). Temporal integration in duration and number discrimination. *Journal of Experimental Psychology: Animal Behavior Processes, 11,* 591-597. https://doi.org/10.1037/0097-7403.11.4.591

Meltzer, D., & Robertson, K. (1989). Intradimensional and extradimensional shifts in conditional discrimination. *Behavioural Processes, 20,* 49-59. https://doi.org/10.1016/0376-6357(89)90011-9

Meltzer, M., & Masaki, M. A. (1973). Measures of stimulus control and stimulus dominance. *Bulletin of the Psychonomic Society, 1,* 28-30. https://doi.org/10.3758/BF03333328

Miles, C. M., & Jenkins, H. M. (1973). Overshadowing in operant conditioning as a function of discriminability. *Learning and Motivation, 4,* 11-27. https://doi.org/10.1016/0023-9690(73)90036-2

Miller, R. R., & Matute, H. (1996). Biological significance in forward and backward blocking: Resolution of a discrepancy between animal conditioning and human causal judgment. *Journal of Experimental Psychology: General, 125,* 370-386. https://doi.org/10.1037/0096-3445.125.4.370

Miranda-Dukoski, L., Davison, M., & Elliffe, D. (2014). Choice, time and food: Continuous cyclical changes in food probability between reinforcers. *Journal of the Experimental Analysis of Behavior, 101,* 406-421. https://doi.org/10.1002/jeab.79

Morse, W. H., & Skinner, B. F. (1958). Some factors involved in the stimulus control of operant behavior. *Journal of the Experimental Analysis of Behavior, 1,* 103-107. https://doi.org/10.1901/jeab.1958.1-103

Munoz, N. E., & Blumstein, D. T. (2012). Multisensory perception in uncertain environments. *Behavioral Ecology, 23,* 457-462. https://doi.org/10.1093/beheco/arr220

Nelson, T. D., & Farthing, G. W. (1973). Generalization of fixed-interval response patterning in pigeons. *Animal Learning & Behavior, 1,* 183-188. https://doi.org/10.3758/BF03199071

Nevin, J. A. (1974). Response strength in multiple schedules. *Journal of the Experimental Analysis of Behavior, 21,* 389-408. https://doi.org/10.1901/jeab.1974.21-389

Nevin, J. A., Cate, H., & Alsop, B. (1993). Effects of differences between stimuli, responses, and reinforcer rates on conditional discrimination performance. *Journal of the Experimental Analysis of Behavior, 59,* 147-161. https://doi.org/10.1901/jeab.1993.59-147

Nevin, J. A., Davison, M., & Shahan, T. A. (2005). A theory of attending and reinforcement in conditional discriminations. *Journal of the Experimental Analysis of Behavior, 84,* 281-303. https://doi.org/10.1901/jeab.2005.97-04

Nevin, J. A., Jenkins, P., Whittaker, S., & Yarensky, P. (1982). Reinforcement contingencies and signal detection. *Journal of the Experimental Analysis of Behavior, 37,* 65-79. https://doi.org/10.1901/jeab.1982.37-65

Nevin, J. A., Olson, K., Mandell, C., & Yarensky, P. (1975). Differential reinforcement and signal detection. *Journal of the Experimental Analysis of Behavior, 24,* 355-367. https://doi.org/10.1901/jeab.1975.24-355

Newman, F. L., & Benefield, R. L. (1968). Stimulus control, cue utilization, and attention: Effects of discrimination training. *Journal of Comparative and Physiological Psychology, 66,* 101-104. https://doi.org/10.1037/h0025988

Nicholls, E., Ryan, C. M. E., Bryant, C. M. L., & Lea, S. E. G. (2011). Labeling and family resemblance in the discrimination of polymorphous categories by pigeons. *Animal Cognition, 14,* 21-34. https://doi.org/10.1007/s10071-010-0339-4

Oberling, P., Bristol, A. S., Matute, H., & Miller, R. R. (2000). Biological significance attenuates overshadowing, relative validity, and degraded contingency effects. *Animal Learning & Behavior*, *28,* 172-186. https://doi.org/10.3758/BF03200252

Olson, D. J. (1991). Species differences in spatial memory among Clark’s nutcrackers, scrub jays, and pigeons. *Journal of Experimental Psychology: Animal Behavior Processes, 17,* 363-376. https://doi.org/10.1037/0097-7403.17.4.363

Oscar-Berman, M., & Bonner, R. T. (1985). Matching- and delayed matching-to-sample performance as measures of visual processing, selective attention, and memory in aging and alcoholic individuals. *Neuropsychologia, 23,* 639-651. https://doi.org/10.1016/0028-3932(85)90065-X

Pavlov, I. P. (1927). *Conditioned reflexes*. London, United Kingdom: Oxford University Press.

Pearce, J. M. (1987). A model for stimulus generalization in Pavlovian conditioning. *Psychological Review, 94,* 61-73. https://doi.org/10.1037/0033-295X.94.1.61

Pearce, J. M. (1994). Similarity and discrimination: A selective review and a connectionist model. *Psychological Review, 101,* 587-607. https://doi.org/10.1037/0033-295X.101.4.587

Perez, W. F., Endemann, P., Pessôa, C. V. B. B., & Tomanari, G. Y. (2015). Assessing stimulus control in a discrimination task with compound stimuli: Evaluating testing procedures and tracking eye fixations. *The Psychological Record, 65,* 83-88. https://doi.org/10.1007/s40732-014-0092-1

Perone, M. (2019). How I learned to stop worrying and love replication failures. *Perspectives on Behavior Science, 42,* 91-108. https://doi.org/10.1007/s40614-018-0153-x

Pinto, C., Fortes, I., & Machado, A. (2017). Joint stimulus control in a temporal discrimination task. *Animal Cognition, 20,* 1129-1136. https://doi.org/10.1007/s10071-017-1130-6

Pinto, C., & Machado, A. (2017). Unraveling sources of stimulus control in a temporal discrimination task. *Learning & Behavior, 45,* 20-28. https://doi.org/10.3758/s13420-016-0233-2

Ploog, B. O. (2010). Stimulus overselectivity four decades later: A review of the literature and its implications for current research in autism spectrum disorder. *Journal of Autism and Developmental Disorders, 40,* 1332-1349. https://doi.org/10.1007/s10803-010-0990-2

Podlesnik, C. A., Thrailkill, E., & Shahan, T. A. (2012). Differential reinforcement and resistance to change of divided-attention performance. *Learning & Behavior, 40,* 158-169. https://doi.org/10.3758/s13420-011-0052-4

Posner, M. I., Nissen, M. J., & Klein, R. M. (1976). Visual dominance: An information-processing account of its origins and significance. *Psychological Review, 83,* 157-171. https://doi.org/10.1037/0033-295X.83.2.157

Rand, J. F. (1977). Behaviors observed during S- in a simple discrimination learning task. *Journal of the Experimental Analysis of Behavior, 27,* 103-117. https://doi.org/10.1901/27-103

Randich, A., Klein, R. M., & LoLordo, V. M. (1978). Visual dominance in the pigeon. *Journal of the Experimental Analysis of Behavior, 30,* 129-137. https://doi.org/10.1901/jeab.1978.30-129

Ray, B. A. (1969). Selective attention: The effects of combining stimuli which control incompatible behavior. *Journal of the Experimental Analysis of Behavior, 12,* 539-550. https://doi.org/10.1901/jeab.1969.12-539

Ray, B. A. (1972). Strategy in studies of attention: A commentary on D. I. Mostofsky’s Attention: Contemporary theory and analysis. *Journal of the Experimental Analysis of Behavior, 17,* 293-297. https://doi.org/10.1901/jeab.1972.17-293

Rayburn-Reeves, R. M., & Cook, R. G. (2016). The organization of behavior over time: Insights from mid-session reversal. *Comparative Cognition and Behavior Reviews, 11,* 103-125. https://doi.org/10.3819/ccbr.2016.110006

Rayburn-Reeves, R. M., Molet, M., & Zentall, T. R. (2011). Simultaneous discrimination reversal learning in pigeons and humans: Anticipatory and perseverative errors. *Learning & Behavior, 39,* 125-147. https://doi.org/10.3758/s13420-010-0011-5

Rayburn-Reeves, R. M., Moore, M. K., Smith, T. E., Crafton, D. A., & Marden, K. L. (2018). Spatial midsession reversal learning in rats: Effects of egocentric cue use and memory. *Behavioural Processes, 152,* 10-17. https://doi.org/10.1016/j.beproc.2018.03.005

Rayburn-Reeves, R. M., Qadri, M. A. J., Brooks, D. I., Keller, A. M., & Cook, R. G. (2017). Dynamic cue use in pigeon mid-session reversal. *Behavioural Processes, 137,* 53-63. https://doi.org/10.1016/j.beproc.2016.09.002

Reed, P., Broomfield, L., McHugh, L., McCausland, A., & Leader, G. (2009). Extinction of over-selected stimuli causes emergence of under-selected cues in higher-functioning children with Autistic Spectrum Disorders. *Journal of Autism and Developmental Disorders, 39,* 290-298. https://doi.org/10.1007/s10803-008-0629-8

Reed, P., & Gibson, E. (2005). The effect of concurrent task load on stimulus over-selectivity. *Journal of Autism and Developmental Disorders, 35,* 601-614. https://doi.org/10.1007/s10803-005-0004-y

Reed, P., Petrina, N., & McHugh, L. (2011). Over-selectivity as a learned response. *Research in Developmental Disabilities, 32,* 201-206. https://doi.org/10.1016/j.ridd.2010.09-016

Reed, P., Reynolds, G., & Fermandel, L. (2012). Revaluation manipulations produce emergence of underselected stimuli following simultaneous discrimination in humans. *The Quarterly Journal of Experimental Psychology, 65,* 1345-1360. https://doi.org/10.1080/17470218.2012.656663

Reinhold, D. B., & Perkins, C. C. (1955). Stimulus generalization following different methods of training. *Journal of Experimental Psychology, 49,* 423-427. https://doi.org/10.1037/h0046761

Rescorla, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology, 66,* 1-5. https://doi.org/10.1037/h0025984

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64-99). New York, NY, USA: Appleton-Century-Crofts.

Reynolds, G. S. (1961). Attention in the pigeon. *Journal of the Experimental Analysis of Behavior, 4,* 203-208. https://doi.org/10.1901/jeab.1961.4-203

Reynolds, G. S. (1963). On some determinants of choice in pigeons. *Journal of the Experimental Analysis of Behavior, 6,* 53-59. https://doi.org/10.1901/jeab.1963.6-53

Reynolds, G. S., & Limpo, A. J. (1969). Attention and generalization during a conditional discrimination. *Journal of the Experimental Analysis of Behavior, 12,* 911-916. https://doi.org/10.1901/jeab.1969.12-911

Riley, D. A. (1984). Do pigeons decompose stimulus compounds?. In H. L. Roitblat, T. G. Bever, & H. S. Terrace (Eds.), *Animal cognition: Proceedings of the Harry Frank Guggenheim conference June 2-4, 1982* (pp. 333-352). Hillsdale, NJ, USA: Lawrence Erlbaum Associates, Inc.

Riley, D. A., & Leith, C. R. (1976). Multidimensional psychophysics and selective attention in animals. *Psychological Bulletin, 83,* 138-160. https://doi.org/ 10.1037/0033-2909.83.1.138

Rincover, A., & Koegel, R. L. (1975). Setting generality and stimulus control in autistic children. *Journal of Applied Behavior Analysis, 8,* 235-246. https://doi.org/10.1901/jaba.1975.8-235

Roberts, A. C., Robbins, T. W., & Everitt, B. J. (1988). The effects of intradimensional and extradimensional shifts on visual discrimination learning in humans and non-human primates. *The Quarterly Journal of Experimental Psychology Section B: Comparative and Physiological Psychology, 40,* 321-341.

Roberts, S. (1981). Isolation of an internal clock. *Journal of Experimental Psychology: Animal Behavior Processes, 7,* 242-268. https://doi.org/10.1037/0097-7403.7.3.242

Roberts, W. A., & Boisvert, M. J. (1998). Using the peak procedure to measure timing and counting processes in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes, 24,* 416-430. https://doi.org/10.1037/0097-7403.24.4.416

Roberts, W. A., Coughlin, R., & Roberts, S. (2000). Pigeons flexibly time or count on cue. *Psychological Science, 11,* 218-222. https://doi.org/10.1111/1467-9280.00244

Roberts, W. A., & Mitchell, S. (1994). Can a pigeon simultaneously process temporal and numerical information?. *Journal of Experimental Psychology: Animal Behavior Processes, 20,* 66-78. https://doi.org/10.1037/0097-7403.20.1.66

Rodrigo, T., Gimeno, E., Ayguasanosa, M., & Chamizo, V. D. (2014). Navigation with two landmarks in rats (Rattus norvegicus): The role of landmark salience. *Journal of Comparative Psychology, 128,* 378-386. https://doi.org/10.1037/a0036544

Rubi, T. L., & Stephens, D. W. (2016a). Does multimodality per se improve receiver performance? An explicit comparison of multimodal versus unimodal complex signals in a learned signal following task. *Behavioral Ecology and Sociobiology, 70,* 409-416. https://doi.org/10.1007/s00265-016-2061-y

Rubi, T. L., & Stephens, D. W. (2016b). Should receivers follow multiple signal components? An economic perspective. *Behavioral Ecology, 27,* 36-44. https://doi.org/10.1093/beheco/arv121

Ryan, C. S., Hemmes, N. S., & Brown, B. L. (2011). Effects of conditioning history on selective stimulus control by elements of compound discriminative stimuli. *Behavioural Processes, 87,* 291-301. https://doi.org/10.1016/j.beproc.2011.06.001

Sainsbury, R. (1971). The “feature positive effect” and simultaneous discrimination learning. *Journal of Experimental Child Psychology, 11,* 347-356. https://doi.org/10.1016/0022-0965(71)90039-7

Sali, A. W., Anderson, B. A., & Yantis, S. (2014). The role of reward prediction in the control of attention. *Journal of Experimental Psychology: Human Perception and Performance, 40,* 1654-1664. https://doi.org/10.1037/a0037267

Santos, C., Soares, C., Vasconcelos, M., & Machado, A. (2019). The effect of reinforcement probability on time discrimination in the midsession reversal task. *Journal of the Experimental Analysis of Behavior, 111,* 371-386. https://doi.org/10.1002/jeab.513

Schmidt, K. A., Dall. S. R. X., & Van Gils, J. A. (2010). The ecology of information: An overview on the ecological significance of making informed decisions. *Oikos, 119,* 304-316. https://doi.org/10.1111/j.1600-0706.2009.17573.x

Schreibman, L., & Lovaas, O. I. (1973). Overselective response to social stimuli by autistic children. *Journal of Abnormal Child Psychology, 1,* 152-168. https://doi.org/10.1007/BF00916110

Segal, M., & Harrison, J. M. (1978). The control of responding by auditory stimuli: Interactions between different dimensions of the stimuli. *Journal of the Experimental Analysis of Behavior, 30,* 97-106. https://doi.org/10.1901/jeab.1978.30-97

Seligman, M. E. (1970). On the generality of the laws of learning. *Psychological Review, 77,* 406-418. https://doi.org/10.1037/h0029790

Seraganian, P., & vom Saal, W. (1969). Blocking the development of stimulus control when stimuli indicate periods of nonreinforcement. *Journal of the Experimental Analysis of Behavior, 12,* 767-772. https://doi.org/10.1901/jeab.1969.12-767

Shahan, T. A. (2010). Conditioned reinforcement and response strength. *Journal of the Experimental Analysis of Behavior, 93,* 269-289. https://doi.org/10.1901/jeab.2010.93-269

Shahan, T. A. (2013). Attention and conditioned reinforcement. In G. J. Madden, W. V. Dube, T. D. Hackenberg, & K. A. Lattal (Eds.), *APA handbook of behavior analysis, Vol 1. Methods and principles* (pp. 387-410). Washington DC, VA, USA: American Psychological Association. https://doi.org/10.1037/13937-017

Shahan, T. A., & Cunningham, P. (2015). Conditioned reinforcement and information theory reconsidered. *Journal of the Experimental Analysis of Behavior, 103,* 405-418. https://doi.org/10.1002/jeab.142

Shahan, T. A., & Podlesnik, C. A. (2006). Divided attention performance and the matching law. *Learning & Behavior, 34,* 255-261. https://doi.org/10.3758/BF03192881

Shahan, T. A., & Podlesnik, C. A. (2007). Divided attention and the matching law: Sample duration affects sensitivity to reinforcement allocation. *Learning & Behavior, 35,* 141-148. https://doi.org/10.3758/BF03193049

Shahan, T. A., & Podlesnik, C. A. (2008). Quantitative analyses of observing and attending. *Behavioural Processes, 78,* 145-157. https://doi.org/10.1016/j.beproc.2008.01.012

Shepp, B. E., & Eimas, P. D. (1964). Intradimensional and extradimensional shifts in the rat. *Journal of Comparative and Physiological Psychology, 57,* 357-361. https://doi.org/10.1037/h0043967

Shepp, B. E., & Schrier, A. M. (1969). Consecutive intradimensional and extradimensional shifts in monkeys. *Journal of Comparative and Physiological Psychology, 67,* 199-203. https://doi.org/10.1037/h0026785

Sherry, D., & Duff, S. (1996). Behavioural and neural bases of orientation in food-storing birds. *Journal of Experimental Biology, 199,* 165-172.

Sherry, D. F., & Schacter, D. L. (1987). The evolution of multiple memory systems. *Psychological Review, 94,* 439-454. https://doi.org/10.1037/0033-295X.94.4.439

Shettleworth, S. J., & Hampton, R. R. (1998). Adaptive specializations of spatial cognition in food-storing birds? Approaches to testing a comparative hypothesis. In R. P. Balda, I. M. Pepperberg, & A. C. Kamil (Eds.), *Animal cognition in nature* (pp. 65-98). Cambridge, MA, USA: Academic Press. https://doi.org/10.1016/B978-012077030-4/50055-6

Shettleworth, S. J., & Westwood, R. P. (2002). Divided attention, memory, and spatial discrimination in food-storing and nonstoring birds, black-capped chickadees (Poecile atricapilla) and dark-eyed juncos (Junco hyemalis). *Journal of Experimental Psychology: Animal Behavior Processes, 28,* 227-241. https://doi.org/10.1037/0097-7403.28.3.227

Shomstein, S., & Johnson, J. (2013). Shaping attention with reward: Effects of reward on space- and object-based selection. *Psychological Science, 24,* 2369-2378. https://doi.org/10.1177/0956797613490743

Sidman, M. (1969). Generalization gradients and stimulus control in delayed matching-to-sample. *Journal of the Experimental Analysis of Behavior, 12,* 745-757. https://doi.org/10.1901/jeab.1969.12-745

Singh, N. N., & Beale, I. L. (1978). Attentional changes during discrimination learning by retarded children. *Journal of the Experimental Analysis of Behavior, 29,* 527-533. https://doi.org/10.1901/jeab.1978.29-527

Skinner, B. F. (1932). Drive and reflex strength. *The Journal of General Psychology, 6,* 22-37. https://doi.org/10.1080/00221309.1932.9711852

Skinner, B. F. (1937). Two types of conditioned reflex: A reply to Konorski and Miller. *Journal of General Psychology, 1,* 272-279. https://doi.org/10.1080/00221309.1937.9917951

Skinner, B. F. (1938). *The behavior of organisms*. New York, NY, USA: Appleton-Century Co.

Skinner, B. F. (1948). ‘Superstition’ in the pigeon. *Journal of Experimental Psychology, 38,* 168-172. https://doi.org/10.1037/h0055873

Skinner, B. F. (1953). *Science and human behavior*. New York, NY, USA: The Free Press.

Slotnick, B. M., & Katz, H. M. (1974). Olfactory learning-set formation in rats. *Science, 195,* 796-798. https://doi.org/10.1126/science.185.4153.796

Smith, A. P., Beckmann, J. S., & Zentall, T. R. (2017). Mechanisms of midsession reversal accuracy: Memory for preceding events and timing. *Journal of Experimental Psychology: Animal Learning and Cognition, 43,* 62-71. https://doi.org/10.1037/xan0000124

Spetch, M. L. (1995). Overshadowing in landmark learning: Touch-screen studies with pigeons and humans. *Journal of Experimental Psychology: Animal Behavior Processes, 21,* 166-181. https://doi.org/10.1037/0097-7403.21.2.166

Spetch, M. L., & Wilkie, D. M. (1994). Pigeons’ use of landmarks presented in digitized images. *Learning and Motivation, 25,* 245-275. https://doi.org/10.1006/lmot.1994.1014

Staddon, J. E. R., & Cerutti, D. T. (2003). Operant conditioning. *Annual Review of Psychology, 54,* 115-144. https://doi.org/10.1146/annurev.psych.54.101601.145124

Stickney, K. J., & Donahoe, J. W. (1983). Attenuation of blocking by a change in US locus. *Animal Learning & Behavior, 11,* 60-66. https://doi.org/10.3758/BF03212308

Stubbs, D. A., & Pliskoff, S. S. (1969). Concurrent responding with fixed relative rate of reinforcement. *Journal of the Experimental Analysis of Behavior, 12,* 887-895. https://doi.org/10.1901/jeab.1969.12-887

Subramaniam, S., & Kyonka, E. (2019). Selective attention in pigeon temporal discrimination. *The Quarterly Journal of Experimental Psychology, 72,* 298-310. https://doi.org/10.1080/17470218.2017.1360921

Sutherland, N. S., & Andelman, L. (1967). Learning with one and two cues. *Psychonomic Science, 7,* 107-108. https://doi.org/10.3758/BF03328487

Sutton, J. E., & Roberts, W. A. (2002). The effect of nontemporal information processing on time estimation in pigeons. *Learning and Motivation, 33,* 124-140. https://doi.org/10.1006/lmot.2001.1103

Swanton, D. N., Gooch, C. M., & Matell, M. S. (2009). Averaging of temporal memories by rats. *Journal of Experimental Psychology: Animal Behavior Processes, 35,* 434-439. https://doi.org/10.1037/a0014021

Swanton, D. N., & Matell, M. S. (2011). Stimulus compounding in interval timing: The modality-duration relationship of the anchor durations results in qualitatively different response patterns to the compound cue. *Journal of Experimental Psychology: Animal Behavior Processes, 37,* 94-107. https://doi.org/10.1037/a0020200

Taylor, K. M., Horvitz, J. C., & Balsam, P. D. (2007). Amphetamine affects the start of responding in the peak interval timing task. *Behavioural Processes, 74,* 168-175. https://doi.org/10.1016/j.beproc.2006.11.005

Terrace, H. S. (1966). Stimulus control. In W. K. Honig (Ed.), *Operant behavior* (pp. 271-344). New York, NY, USA: Appleton-Century-Crofts.

Thesen, A., Steen, J. B., & Doving, K. B. (1993). Behaviour of dogs during olfactory tracking. *Journal of Experimental Biology, 180,* 247-251.

Thomas, D. R., Burr, D. E., & Eck, K. O. (1970). Stimulus selection in animal discrimination learning: An alternative interpretation. *Journal of Experimental Psychology, 86,* 53-62. https://doi.org/10.1037/h0030023

Thomas, D. R., Cook, S. C., & Terrones, J. P. (1990). Conditional discrimination learning by pigeons: The role of simultaneous versus successive stimulus presentations. *Journal of Experimental Psychology: Animal Behavior Processes, 16,* 390-401. https://doi.org/10.1037/0097-7403.16.4.390

Thorpe, C. M., & Wilkie, D. M. (2002). Unequal interval time-place learning. *Behavioural Processes, 58,* 157-166. https://doi.org/10.1016/S0376-6357(02)00030-X

Uengoer, M., & Lachnit, H. (2012). Modulation of attention in discrimination learning: The roles of stimulus relevance and stimulus-outcome correlation. *Learning & Behavior, 40,* 117-127. https://doi.org/10.3758/s13420-011-0049-z

Urcuioli, P. J. (1984). Overshadowing in matching-to-sample: Reduction in sample-stimulus control by differential sample behaviors. *Animal Learning & Behavior, 12,* 256-264. https://doi.org/10.3758/BF03199966

Urcuioli, P. J. (1985). On the role of differential sample behaviors in matching-to-sample. *Journal of Experimental Psychology: Animal Behavior Processes, 11,* 502-519. https://doi.org/10.1037/0097-7403.11.4.502

Urcuioli, P. J., & Honig, W. K. (1980). Control of choice in conditional discriminations by sample-specific behaviors. *Journal of Experimental Psychology: Animal Behavior Processes, 6,* 251-277. https://doi.org/10.1037/0097-7403-6.3.251

von Fersen, L., & Lea, S. E. G. (1990). Category discrimination by pigeons using five polymorphous features. *Journal of the Experimental Analysis of Behavior, 54,* 69-84. https://doi.org/10.1901/jeab.1990.54-69

Vyazovska, O. V., Navarro, V. M., & Wasserman, E. A. (2018). Pigeons deploy selective attention to efficiently learn a stagewise multidimensional visual discrimination task. *Journal of Experimental Psychology: Animal Learning and Cognition, 44,* 162-167. https://doi.org/10.1037/xan0000168

Vyazovska, O. V., Teng, Y., & Wasserman, E. A. (2014). Attentional tradeoffs in the pigeon. *Journal of the Experimental Analysis of Behavior, 101,* 337-354. https://doi.org/10.1002/jeab.82

Wagner, A. R. (1969). Stimulus validity and stimulus seletion. In W. K. Honig & N. J. Mackintosh (Eds.), *Fundamental issues in associative learning* (pp. 90-122). Halifax, Canada: Dalhousie University Press.

Wagner, A. R., Logan, F. A., Haberlandt, K., & Price, T. (1968). Stimulus selection in animal discrimination learning. *Journal of Experimental Psychology, 76,* 171-180. https://doi.org/10.1037/h0025414

Walpole, C. W., Roscoe, E. M., & Dube, W. V. (2007). Use of a differential observing response to expand restricted stimulus control. *Journal of Applied Behavior Analysis, 40,* 707-712. https://doi.org/10.1901/jaba.2007.707-712

Wasserman, E. A. (1974). Stimulus-reinforcer predictiveness and selective discrimination learning in pigeons. *Journal of Experimental Psychology, 103,* 284-297. https://doi.org/10.1037/h0036872

Wasserman, E. A., & Anderson, P. A. (1974). Differential autoshaping to common and distinctive elements of positive and negative discriminative stimuli. *Journal of the Experimental Analysis of Behavior, 22,* 491-496. https://doi.org/10.1901/jeab.1974.22-491

Wasserman, E. A., & Miller, R. R. (1997). What's elementary about associative learning?. *Annual Review of Psychology, 48,* 573-607. https://doi.org/10.1146/annurev.psych.48.1.573

Wayland, S., & Taplin, J. E. (1982). Nonverbal categorization in fluent and nonfluent anomic aphasics. *Brain and Language, 16,* 87-108. https://doi.org/10.1016/0093-934X(82)90074-8

Wearden, J. H., & Lejeune, H. (2008). Scalar properties in human timing: Conformity and violations. *The Quarterly Journal of Experimental Psychology, 61,* 569-587. https://doi.org/10.1080/17470210701282576

Weaver, J. E., Dorrance, B. R., & Zentall, T. R. (1999). Present/absent sample matching in pigeons: Is comparison choice controlled by the same stimulus or by differential sample responding?. *Animal Learning & Behavior, 27,* 288-294. https://doi.org/10.3758/BF03199727

Werner, C. W., Tiemann, I., Cnotka, J., & Rehkämper, G. (2005). Do chickens (Gallus gallus f. domestica) decompose visual figures?. *Animal Cognition, 8,* 129-140. https://doi.org/10.1007/s10071-004-0229-8

Whitaker, S., Lowe, C. F., & Wearden, J. H. (2003). Multiple-interval timing in rats: Performance on two-valued mixed fixed-interval schedules. *Journal of Experimental Psychology: Animal Behavior Processes, 29,* 277-291. https://doi.org/10.1037/0097-7403.29.4.277

Whitaker, S., Lowe, C. F., & Wearden, J. H. (2008). When to respond? And how much?: Temporal control and response output on mixed-fixed-interval schedules with unequally probable components. *Behavioural Processes, 77,* 33-42. https://doi.org/10.1016/j.beproc.2007.06.001

White, K. G. (1986). Conjoint control of performance in conditional discriminations by successive and simultaneous stimuli. *Journal of the Experimental Analysis of Behavior, 45,* 161-174. https://doi.org/10.1901/jeab.1986.45-161

Wilkie, D. M. (1973). Behavioral interactions and stimulus control during conditional discriminations. *Journal of the Experimental Analysis of Behavior, 20,* 483-487. https://doi.org/10.1901/jeab.1973.20-483

Wilkie, D. M. (1995). Time-place learning. *Current Directions in Psychological Science, 4,* 85-89. https://doi.org/10.1111/1467-8721.ep10772332

Wilkie, D. M., & Masson, M. E. (1976). Attention in the pigeon: A reevaluation. *Journal of the Experimental Analysis of Behavior, 26,* 207-212. https://doi.org/10.1901/jeab.1976.26-207

Wilkie, D. M., & Willson, R. J. (1992). Time-place learning by pigeons, Columba livia. *Journal of the Experimental Analysis of Behavior, 57,* 145-158. https://doi.org/10.1901/jeab.1992.57-145

Williams, B. A. (1984). Relative stimulus validity in conditional discrimination. *Animal Learning & Behavior, 12,* 117-121. https://doi.org/10.3758/BF03213129

Williams, B. A. (1994). Blocking despite changes in reinforcer identity. *Animal Learning & Behavior, 22,* 442-457. https://doi.org/10.3758/BF03209164

Williams, D. A., & Braker, D. S. (1999). Influence of past experience on the coding of compound stimuli. *Journal of Experimental Psychology: Animal Behavior Processes, 25,* 461-474. https://doi.org/10.1037/0097-7403.25.4.461

Williams, D. A., Sagness, K. E., & McPhee, J. E. (1994). Configural and elemental strategies in predictive learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 694-709. https://doi.org/10.1037/0278-7393.20.3.694

Williams, D. I. (1967). Constant irrelevant cue learning in the pigeon. *Animal Behaviour, 15,* 229-230. https://doi.org/10.1016/0003-3472(67)90004-8

Wolfgang, A. K., & Richardson, J. (1973). Effect of type of stimulus component on cue selection. *The American Journal of Psychology, 86*, 635-642. https://doi.org/10.2307/1421949

Wyckoff, L. B. (1952). The role of observing responses in discrimination learning. *Psychological Review, 59,* 431-442. https://doi.org/10.1037/h0053932

Yokoyama, K., Dailey, D., & Chase, S. (2006). Processing of conflicting and redundant stimulus information by pigeons. *Learning & Behavior, 34,* 241-247. https://doi.org/10.3758/BF03192879

Zentall, T. R. (2005a). Configural/holistic processing or differential element versus compound similarity. *Animal Cognition, 8,* 141-142. https://doi.org/10.1007/s10071-004-0241-z

Zentall, T. R. (2005b). Selective and divided attention in animals. *Behavioural Processes, 69,* 1-15. https://doi.org/10.1016/j.beproc.2005.01.004

Zentall, T. R., Sherburne, L. M., & Zhang, Z. (1997). Shared attention in pigeons: Retrieval failure does not account for the element superiority effect. *Learning and Motivation, 28,* 248-267. https://doi.org/10.1006/lmot.1996.0965

Zuckerman, D. C. (1973). Steady state responding based upon simple and compound stimuli. *Journal of the Experimental Analysis of Behavior, 20,* 209-218. https://doi.org/10.1901/jeab.1973.20-209