Arbitrary Stimuli and Responses can be used to Investigate Planning Behaviours

Rachel Elizabeth Hawkins

A thesis submitted in fulfilment of the requirements for the degree of Master of Arts in Psychology, the University of Auckland, 2022.

Abstract

The ability for non-human animals to plan for the future – such that they act in the present in accordance with what will be best in the likely future – is an area of growing research interest. Evidence suggests that 'simple' operant behaviour is future-oriented, implying that behavioural control is fundamentally prospective. Evidence for future-oriented behaviour in non-human animals is often surrounded with debate as to whether genuine future planning abilities are demonstrated. Many demonstrations of non-human animal planning employ species-specific behaviours, limiting the generality of conclusions. To further our knowledge of the generality of planning across species, we investigated whether humans could learn to plan in procedures commonly used with non-human animals, with arbitrary stimuli and responses. In one experiment we replicated a procedure previously used with crows (Boeckle et al., 2020) and varied the arbitrariness of the stimuli used with human participants. We found that humans were successful in learning when trained in the same way as crows, however they performed worse than crows during testing trials. Stimuli that were more arbitrary resulted in slightly better performance during testing than stimuli that were less arbitrary. In another experiment we compared the performance of pigeons and humans on a task using arbitrary stimuli and responses. A response at the beginning of a trial was required for a different response to produce a reinforcer at the end of a trial. We found similar patterns of learning across both species. Both pigeons and humans required experience to learn the procedure, and both performed better when a correct response at the beginning of a trial caused an immediate stimulus change than when it did not. Both species learned to make responses with temporally and spatially distant consequences. Our findings suggest that we can study planning in humans and non-humans using arbitrary stimuli and responses, however the arbitrariness of the stimuli and the task used may be critical in determining whether experience generalizes to novel situations. These findings highlight a need to

understand how experience (both within and outside the experiment) underpins planning and prospective control.

Keywords: planning, future-oriented behaviour, prospective control, arbitrary stimuli

Acknowledgments

Firstly, I would like to thank my two supervisors – Dr Sarah Cowie and Associate Professor Alex Taylor. Thank you, Alex, for taking the time to assist us with designing our pigeon and human replications of the crow experiment. Even though the pandemic prevented us from running the pigeon experiment I appreciate your input in helping us design this, hopefully we can run it one day! Thank you, Sarah, for being such a wonderful supervisor. I wouldn't have been able to get through this process without all of your support, encouragement and advice, especially with all of the disruptions we faced throughout the year. I have learnt so much from you and feel very grateful that I got to work with you on this project, thank you.

Thank you to the members of the Operant Behaviour Lab and to my office mates for your support and encouragement throughout the year. Thank you to Dylan Taylor for helping with proofreading.

Completing a thesis in the middle of a pandemic was challenging to say the least, so a big thank you to all of the people who have supported me throughout this process. In particular my family and friends who have been by my side and offered words of encouragement when times were particularly challenging. I appreciate each and every one of you.

Table of Contents

List of Figures

Behaviour and The Future

The ability for nonhuman animals (henceforth, animals) to plan for the future – such that they can act in the present in accordance with an event that may occur at some later time – is an area of growing research interest. It is well known that humans can plan for the future, and we often engage in this behaviour in our day to day lives. Such behaviour requires learning about relations between events that are extended in time, and extrapolation about potential future conditions on the basis of past experience. These abilities are often considered higher-order, but research has challenged this assumption (e.g., Cowie, 2018, 2020; Cowie et al., 2011; Krageloh et al., 2005). Yet when behaviours that are considered 'complex' are studied in animals, evidence to suggest such behaviours reflect complex *processes* is often dismissed (e.g., de Mahy et al., 2021; Redshaw et al., 2017; Suddendorf et al., 2009). Rather, it is often argued that the complex behaviour is more likely a combination of simple behaviour and processes (Leslie, 2018).

One of the simplest learning processes is operant learning via reinforcement. The dominant view has been that reinforcers function by strengthening behaviour that has just occurred. Thorndike (1911) established the law of effect, stating that any behaviour followed by a favourable consequence (reinforcement) was more likely to occur again. Likewise, any behaviour followed by an unfavourable consequence (punishment) was less likely to occur again. This theory suggests that reinforcement functions retrospectively by strengthening the behaviour it follows, increasing its likelihood of reoccurring. If a pigeon pecks a key and receives a reinforcer, the pigeon will be more likely to peck that same key again. Within this view behaviour is retrospectively controlled by what has happened in the past. This view of response strengthening has generally been accepted within behaviour analysis without much question (Shahan, 2017).

Recent work, however, has questioned whether a response strengthening theory is the best way to conceptualise the function of reinforcement (Cowie, 2018, 2020; Davison, 2017; Shahan, 2017). The response strengthening theory is not always able to account for the performance and acquisition of behaviour (Shahan, 2017) and often does little to predict and explain simple operant behaviour (Cowie, 2020). Rather, a response signalling theory seems to function better to account for simple operant behaviour. This response signalling theory proposes that reinforcers function as discriminative stimuli, signalling to the organism what behaviour to do next (Cowie, 2020). In this case behaviour is prospective and controlled by what is likely to happen in the near future (Cowie, 2018). Recent studies have found promising evidence for the functionality of this theory when predicting and explaining simple operant behaviour.

A variety of results are inconsistent with a retrospective strengthening effect, and instead suggest control by what is likely to occur subsequently. For example, rats in a radial arm maze where arms are not rebaited will never return to an arm in which they have found food during that session (Olton & Samuelson, 1976). In procedures where two different behaviours might each produce reinforcers, pigeons (Cowie et al., 2011; Krageloh et al., 2005) and children (Cowie et al., 2021; Cowie, Virués-Ortega et al., 2021) will use more extended relations between reinforcers to choose in accordance with which behaviour is likely to produce the *next* reinforcer, even when this requires avoiding the most recently reinforced response.

These findings support a signalling view of reinforcement and therefore suggest that simple operant behaviour may be more complex than is typically assumed. Behaviour may not be as simple as repeating previously reinforced behaviours. Instead, animals can learn extended relations between events and use this information to respond in accordance with what is likely to produce food next. That is, simple operant behaviour seems to be

fundamentally future-oriented; behaviour is directed towards what may happen in the future, and reinforcers serve to signal what behaviour to do next. If even the simplest forms of behaviour are prospectively controlled, then 'higher-order' future-oriented behaviours like planning may be more ubiquitous across species than is typically assumed.

Memory

An important aspect of prospective behaviour is memory. For an animal to act in the present in accordance with the likely future they need to be able to remember what behaviour they have just done, and what events have happened in the past. Evidence has shown that animals can learn and remember when there is a delay between their behaviour and the consequence of that behaviour. A typical procedure used to examine memory in simple operant behaviour is a delayed matching-to-sample (DMTS) procedure. In this procedure the animal is presented with a sample stimulus and taught which comparison stimulus is a correct match. A delay is then imposed between the presentation of the sample stimulus and the presentation of comparison stimuli, during which the animal must remember what sample they were presented with in order to choose the correct comparison. This procedure has proved to be an effective way to measure and study memory, allowing researchers to investigate working memory in animals (Zentall & Smith, 2016).

A meta-analysis of DMTS tasks using data sets from 90 different studies and 25 different species calculated the performance half-life scores of animals on these tasks (Lind et al., 2015). This is where performance falls halfway between performance with no delay and chance performance, giving an accurate measure of the memory abilities of animals. Performance half-life scores were found to range from 2.4 s to 71 s with a median across all species of 27 s. This demonstrates that animals can use working memory and learn across a period of delay. They can remember information they have been given and use this when it becomes useful in the future.

Learning across a period of delay seems to be enhanced when behaviour produces an immediate stimulus change. Richards (1981) trained pigeons on a variable-interval (VI) 60 s schedule and then introduced a delay between the response that produced the reinforcer and reinforcer delivery. In a Signalled condition the pilot light within the pigeons' chamber remained illuminated during the delay. In an Unsignalled condition no stimulus change occurred. Decreases in responding were seen in both conditions when a delay was introduced, however responding decreased to a much greater extent in the Unsignalled condition than in the Signalled condition. These results show that pigeons can learn to respond when a period of delay is used, however responding seems to be enhanced when a signal is used during this delay.

Episodic memory is a form of memory that is involved when planning for the future. Episodic memory in humans consists of an individual being able to recall an event that has personally happened to them in the past (Raby & Clayton, 2012). Evidence has suggested that episodic memory is closely tied to future thinking and the ability for humans to imagine themselves in a future scenario (Raby & Clayton, 2012). In humans, episodic memory functions to allow old information to be used to solve problems in the present or the future (Raby & Clayton, 2012). As episodic memory involves subjective experiences, this can make it difficult to study in animals who cannot verbally communicate with us (Crystal, 2010). Therefore, the most effective way we can study episodic memory in animals is through their behaviour. Clayton et al. (2003) proposed studying 'episodic-like' memory in animals using three sets of behavioural criteria. To establish episodic-like memory in animals they suggest evidence that the animals can remember 'when' a particular event happened, as well as 'what' the event was and 'where' it occurred should be investigated (Clayton et al., 2003). Many studies investigating episodic-like memory in animals have adopted this 'what-when-where' (WWW) criteria to determine whether animals possess an episodic-like memory.

Caching tasks are often used to assess episodic-like memory in birds. Caching is a behaviour that is naturally present in birds in the wild; birds will store food in a location to later retrieve. Caching tasks often involve providing birds with two types of food, one that will degrade over time and one that will not degrade over time. The birds are then given the opportunity to cache these foods and retrieve them after either a short or a long period of delay. This tests their memory for what food was stored, where it was stored and when it was stored. If birds possess episodic-like memory, they should retrieve the degradable food after a short period of delay and the non-degradable food after a long period of delay.

Episodic-like memory has been examined in caching tasks with scrub-jays (Clayton & Dickinson, 1998) and black-capped chickadees (Feeney et al., 2009). Scrub-jays were given degradable (worms) and non-degradable (peanuts) food items to cache which they could recover after either a short or long delay. Chickadees foraged in an aviary environment where degradable (mealworms) and non-degradable (sunflower seeds) foods were stored in trays and replenished in the same locations after either a short or long delay. For both species, after short delays more time was spent recovering the degradable food and after long delays more time was spent recovering the non-degradable food. These results demonstrate that both scrub-jays and chickadees were using episodic-like memory to remember what type of food was stored, where it had been stored and how long ago it had been stored there.

Episodic-like memory has also been investigated in more species-general forms of behaviour. Babb and Crystal (2006) used rats in an eight-arm radial maze where four arms were accessible and baited with chow pellets, two of which had distinctive flavours. Rats were given the opportunity to investigate the maze before being removed and put back into the maze after a retention interval. After this interval, only the arms that had not initially been accessible were baited with pellets, and the distinctive flavoured pellets were only replenished after a long but not short retention interval. When rats were put back into the

maze, they typically avoided depleted locations, and were more likely to visit the distinctive flavoured arms after a long retention interval than after a short retention interval. This suggests the rats were able to remember information about the type of food available, the location of the food and the times it would be available, demonstrating episodic-like memory.

Episodic-like memory has also been investigated in the species-general behaviour of pigeons (Meyers-Manor et al., 2014). Two different coloured keylights were located on either side of an operant chamber. Depending on whether the session was carried out in the morning or the afternoon, one key would deliver a short duration of access to food after a high fixed-ratio (FR) value of keypecks, considered to be the non-optimal key. The other key would deliver a long duration of food after a low FR value of keypecks, considered to be the optimal key. The locations of the optimal and non-optimal keys reversed between morning and afternoon sessions. Responses followed the location of the optimal key based on the time of day. Responses to the key that had been optimal in the morning decreased when this key became non-optimal in the afternoon. This suggests the pigeons were able to remember WWW information about each type of key and adjusted their behaviour accordingly to respond to the optimal key.

Memory functions by guiding our behaviour towards the future, allowing organisms to predict the environment they are in (Osvath, 2016). Given how closely memory is tied towards future thinking in humans (Raby & Clayton, 2012), episodic-like memory in animals may indicate capacity for prospective behaviour.

Future-Oriented Behaviour

One more ‘complex’ form of prospective behaviour is the ability to plan. Planning is a behaviour that requires learning about relations between events that are extended in time, and then using this information to predict possible future events. In animals, planning is generally

defined as acting in the present in accordance with an event that may occur in the future – for example, planning for a future where you will be hungry even though you are not hungry in the present. Planning in animals generally involves performing a behaviour that has little to no value in the present but will come to pay off in the future. Different tasks have been used to investigate planning in animals such as tool use tasks, token exchange tasks and caching.

Tool Use Tasks

One type of task used to investigate future planning are tasks where tools available to the animal are not useful in the present but become useful in the near future. To do this, animals must obtain the tool before it becomes functionally useful. Within tool use tasks, animals are presented with an apparatus that is baited with something desirable, such as food. To access this food a specific tool must be used, which the animal must choose before the apparatus is available. As the tool is not useful in the present, any choice of the correct tool must be due to the anticipation of a future need of the tool. This suggests that the animal is planning for a future in which the tool is useful. Most tool use tasks use species that use tools in the wild. They often involve necessary training steps, such as teaching the animal which tools can and cannot be used on the apparatus and how food can be accessed from the apparatus. These training steps ensure the animal knows which tools are functional and how the apparatus works, so they can retrieve the food.

Brauer and Call (2015) investigated whether apes could produce multiple tools for future use. Apes were presented with an eight-arm apparatus baited with grapes. To retrieve the grapes a stick could be pushed through one of the arms making it available to eat. The nature of the apparatus meant that each stick tool could only be used once. Apes were provided with two pieces of wooden board which could be broken into pieces to make stick tools. A period of eight minutes commenced where the apes could not access the apparatus but could prepare tools to be used. After this period ended access to the apparatus was given

and tools prepared in advance or made during the access period could be used. When all arms of the apparatus were baited with grapes, the apes prepared more tools in advance compared to conditions where only one, or none, of the apparatus arms were baited. When tools were prepared in advance more grapes were collected than when tools were not prepared in advance. The tools were not useful to the apes while they were being made, suggesting that the apes were planning for future access to the apparatus by preparing tools in advance.

Corvids such as crows and ravens (Boeckle et al., 2020; Kabadayi & Osvath, 2017) also appear to be able to plan for future tool use. Boeckle et al. (2020) investigated whether crows could select an appropriate tool for future use. The crows were presented with an apparatus and after a waiting period of five minutes given a choice between the functional tool and four non-functional objects. Once they chose an item another waiting period of ten minutes ensued before they could use the chosen item on the apparatus. The crows performed significantly above chance choosing the correct tool, with one subject choosing correctly in 9 out of 10 trials. This suggests they were able to plan for the tool being useful in the future. Kabadayi and Osvath (2017) conducted a similar tool use task with ravens. The ravens were shown the apparatus and after a waiting period of one hour given the choice between the functional tool and other non-functional items. Once they had made their choice another short delay period ensued before they could use the item they had chosen on the apparatus. The ravens successfully chose the correct tool on an average of 11 out of 14 trials, suggesting that they were able to plan for the future use of the tool. Both experiments suggest that Corvids can choose tools that have no value in the present but will become useful in the near future.

Token Exchange Tasks

Token exchange tasks have also been commonly used to investigate planning in animals. In token exchange tasks animals are taught that a specific token is valuable as it can be exchanged for food in the near future. They are then presented with the valuable token and

several non-valuable tokens and asked to choose one or multiple tokens. After a waiting period, an experimenter will then exchange any valuable tokens the animal has chosen with food. Given that the token is not valuable in the present when it is chosen, any choice of a valuable token must reflect anticipation of the token being useful in the future.

Bourjade et al. (2014) examined whether apes could collect and transport tokens to exchange them for food. Apes were given access to a variety of tokens, some valuable and some non-valuable, which they could transport between two different rooms. They were then given the opportunity to exchange any valuable tokens they had transported with an experimenter for food. Two out of three species of apes collected and transported valuable tokens at a probability significantly higher than chance. This suggests that the apes were able to plan for the future in which the tokens would become valuable as they could be exchanged for food.

Kabadayi and Osvath (2017) showed ravens could also learn to plan for the future use of tokens. Ravens were presented with the token and distractor items, after which the experimenter appeared fifteen minutes later, and the ravens could exchange their valuable tokens for food. The ravens chose the correct token 143 out of 144 times, suggesting they were able to plan ahead to when the token would be valuable. These studies both suggest that non-human animals can choose an item in the present in order to use it in the future when it becomes valuable.

Caching and Food

Tasks involving food and caching, where food is stored and later retrieved, are commonly used among bird species to assess whether they can plan for the future. Caching studies typically use species who naturally cache their food in the wild. In these tasks animals will be given food with the opportunity to store and later retrieve it.

Raby et al. (2007) showed scrub-jays could plan for future hunger. The birds were taught that in one compartment they would receive breakfast in the morning and in another compartment, they would receive no breakfast. They were then given cacheable pine nuts in the evening which they could store in either compartment or eat immediately. The birds were able to plan for their future hunger by caching more pine nuts in the no-breakfast compartment than in the breakfast compartment. This suggests that the scrub-jays were able to anticipate that they would be hungry if placed in the no-breakfast compartment the next morning and therefore planned accordingly.

Feeney et al. (2011) showed black-capped chickadees could anticipate future needs for food. The birds were first given access to sunflower seeds which could be freely eaten for five minutes. After this a delay period of either five, ten or thirty minutes occurred, after which the birds were given highly preferred mealworms to eat. Birds who were given future access to mealworms ate significantly fewer sunflower seeds than birds who were not given access to mealworms. This suggests that the birds suppressed their current hunger in favour of a more preferred food in the future. These studies suggest that birds display future planning abilities in tasks that involve everyday behaviours such as caching and eating food.

Planning in Humans

The planning behaviours demonstrated by animals are also evident when humans perform similar experimental tasks. Miller et al. (2020) replicated the experiment conducted by Boeckle et al. (2020) with crows, however they used children as participants. The children were trained using the same apparatus and tool combinations as the crows. They were first shown a baited apparatus in location A and then waited in location B for a set period of delay before they were given the choice between a functional and a non-functional tool. After another delay period they were allowed back into location A where they could use the tool they had chosen on the apparatus to receive the reward. Children aged four and five

successfully chose the correct tool significantly above chance. This suggests that they were able to anticipate the future need of the tool when it became functional, and therefore planned accordingly by choosing it in advance. The similar patterns of behaviour by Boeckle et al.'s crows and Miller et al.'s children using a similar task provides further evidence that the planning engaged in by crows when using tools reflects the same sort of planning engaged in by humans.

Russell et al. (2010) also examined planning in children. In this experiment children aged 3 to 5 were taught to play a game where specific tools were needed. The game required a straw to be played and depending on what side of the table the game was played from, would also require a box. The experimenter played a round of the game with the child. During this round the child played from the side of the table that was made easily reachable for the children. Once the game had finished the children were told that they would be playing the game again tomorrow but from the other, unreachable side of the table. They were then presented with six tools, including the straw and a box, and asked which two items they wanted saved to use when they played the game the next day. It was found that 5-year-olds, but not 3- or 4-year-olds selected the correct two tools (the straw and the box) above chance. This suggests that the older children were thinking ahead to the game they would play tomorrow and selected tools that would be useful then, and again demonstrates similarity between tool use planning in humans and animals.

Planning with Less Specialised Behaviours

While most studies involving future planning tend to focus on more complex behaviours such as exchanging tokens and using tools, future planning in more simple behaviour such as that of pigeons pecking has also been investigated. Investigating planning using simple operant behaviour allows us to understand the extent to which planning is a species-general process. The stimuli and the behaviours used can be more arbitrary, and

therefore any evidence of planning is not reliant on a behaviour that the animal already demonstrates out in the wild, or subject to the influence of a behavioural history that is unknown to the experimenter.

Miyata and Fujita (2008) examined pigeons' ability to plan by using performance on a computerised maze task. Pigeons were presented with a red square shaped target stimulus and a blue square shaped goal stimulus. The target stimulus was surrounded by four small white dots on each side of the square. To move the target stimulus the pigeon had to peck at one of these dots, and the stimulus would be moved in the direction of the dot they had pecked. To solve the maze the pigeon had to move the target stimulus to the location of the goal stimulus. How far into the future the pigeons were planning was examined by suddenly changing the location of the goal stimulus. If the pigeon was planning ahead, we would expect to see them move the stimulus in the direction of the old goal location when this location was suddenly changed. When the goal location was changed, pigeons typically moved the target stimulus in the direction of the old goal location. The frequent movement of the target stimulus in the direction of the old goal location suggested that the pigeons were already planning to move the stimulus in that direction. This suggests that the pigeons were planning at least one step ahead where to move the stimulus.

Scarf and Colombo (2010) trained pigeons to respond to three stimuli (named Stimulus A, B and C) in a specific order by pecking the stimuli on a touch screen. They examined the pigeons' performance when stimuli order changed, and when stimuli were no longer visible. When Stimuli B and C switched positions after the pigeon responded to Stimulus A, responding to Stimulus B was slower than when the location of this stimulus remained the same. This suggests the pigeons were already planning where to respond next and had to inhibit this response and adjust to the new stimulus location. When Stimuli B and C were covered with a white square after the pigeon responded to Stimulus A, the pigeons

continued to follow the correct sequence of responding. This suggests they had already planned their next movement towards the stimulus and continued with this, even when the stimuli were no longer visible. These findings suggest that the pigeons were planning at least one step ahead as they continued to follow the correct sequence of responding when stimuli were not visible and adjusted their responding when stimuli changed locations.

Cowie and Davison (2021) investigated whether pigeons' behaviour could be controlled by consequences that were temporally and spatially distant. In this procedure the behaviour that produced a reinforcer was separated from the consequence by time and spatial location. Pigeons had to make an 'Investing' response at the beginning of a trial, followed by 25 'Outcome' responses in order to receive food. Pigeons were successful in making Investing responses early in a trial, after which responding switched exclusively to Outcome responses. This suggests that pigeons can make responses that are not temporally or spatially contiguous with food. Overall, these findings suggest that planning – or at least some of its fundamental building blocks – may be present even in the simple behaviour of stimuli pecking.

While the evidence discussed does seem to show the ability for animals to plan for the future, these conclusions can often be highly debated or critiqued. For example, Redshaw et al. (2017) have questioned the conclusions made by Kabadayi and Osvath (2017) that ravens could plan for the future. Redshaw et al. (2017) argued that while ravens preferred the correct items, this does not necessarily mean they were genuinely planning for the future. They propose that a possible explanation for this preference could be that the correct items were paired more with reward than distractor items, making them more attractive to the crows.

To counter debate about the planning abilities of animals Suddendorf and Corballis (2010) set out a specific set of criteria they believe should be used to properly demonstrate

future planning abilities in animals. These include ideas such as using single trials and novel problems to avoid repeated exposure and innate responses. They also propose a clear temporal and spatial separation be used between an action and the future consequence to avoid reliance on location and ensure long-term memory is used. Finally, they suggest that problems over a variety of domains are used to avoid a reliance on species-specific behaviours. These recommendations create a substantial challenge for the design of experiments to investigate planning. Many of the studies that provide more convincing evidence for planning have tended to use animals with species-specific skills such as caching or tool use. These behaviours are often innate and involve skills that the animals are already adept at, and hence the contribution of the animal's learning history to its ability to plan is unknown. Further, animals tend to require training. As seen with the tool use and token exchange tasks, animals need to be taught about the structure of the environment via experience – otherwise, there is no foreseeable future for which to plan. This makes it almost impossible to perform single-trial demonstrations of planning.

Using arbitrary behaviours in planning experiments is an alternative approach that would overcome many of the concerns raised by Suddendorf and Corballis (2010) and others. Tasks employing arbitrary (non-species-specific) behaviours have the advantage of isolating the contribution of experience to the ability to plan; with arbitrary behaviours and stimuli, experimenters have increased control over an organism's relevant learning history. Indeed, the influence of experience on planning is largely ignored, but may be an important component of the ability to use past experience to navigate potential future conditions. Further, such tasks may be easily adapted for different species, facilitating comparisons across a variety of different animals, including those who do not engage in species-specific caching, tool use, or other behaviours which appear to recruit planning processes. Comparison between the 'signatures' of planning behaviour, and of the pattern of learning

that precedes such behaviour, across different species may well reveal similarities and differences that have gone previously unnoticed because of substantial procedural variations and will ultimately help to illuminate the ubiquity of planning and future-oriented behaviour.

The Current Experiments

One way we can enhance our knowledge of the generality of the ability to plan would be to investigate whether humans, who we already know can plan, can learn to do so when required to learn in the same way as animals, with arbitrary stimuli and responses. This would allow us to determine whether humans are able to acquire the ability to perform the task in the same way as animals when using similar training and experience levels, and hence whether the tasks do in fact reflect planning. We investigated this across three different experiments. In Chapter 1 and 2 we replicated the experiment conducted by Boeckle et al. (2020) using crows, replacing real tools and apparatuses with stimuli of varying levels of arbitrariness, with human participants. Experiment 1 allowed us to investigate the performance of humans on this task while Experiment 2 determined that participants' failure to choose correctly was not due to a procedural artefact. In Chapter 2 we replicated with human participants Cowie and Davison's (2021) procedure where a particular response had to be made at the beginning of a trial in order for another response to produce a reinforcer at the end of a trial. We assessed learning on this task and compared it with re-analysed data from Cowie and Davison's (2021) pigeons on the same task, to determine whether, and to what extent, the pattern of learning was similar across species. Overall, these three experiments will help us to examine whether we can investigate planning abilities in humans and animals using arbitrary stimuli and responses, potentially creating new avenues for understanding planning behaviour in non-human animals.

Chapter 1

Boeckle et al. (2020) and Miller et al. (2020) have shown that both crows and children can learn to plan for the future using a task where a specific tool is chosen in the present to be used on an apparatus in the near future. Both crows and children have experience with tools. Crows naturally use tools in the wild, and the 'tools' chosen for the children were objects they would have previously encountered such as pencils and paperclips. Therefore, an important question is whether this same task could be completed by animals for whom tool use is not part of their natural behavioural repertoire. Such demonstrations would require the use of arbitrary stimuli to demonstrate cross-species abilities. A failure of non-tool-users to complete this type of task would reflect the importance of tool-use as a prerequisite behaviour but may also reflect the difficulty in using arbitrary stimuli. We therefore asked whether adult humans could complete the same task as the crows and children using arbitrary stimuli, rather than objects that functioned as tools and apparatuses. We specifically designed the task as if it were an operant task to be completed by pigeons (unfortunately the pandemic prevented us from carrying out the work with the pigeons) as this requires no species-specific behaviour or stimuli. This creates an important foundation for subsequent work with animals and the use of arbitrary stimuli.

Method

Participants

Participants in Group Colour were 14 humans (6 males, 8 females) aged 20 to 31 years. Participants in Group Tool were 10 humans (8 males, 2 females) aged 19 to 35 years.

Apparatus

Experimental sessions and data collection were conducted online using Psytoolkit (Stoet, 2010, 2017). Participants were provided with a link that allowed them to access the experiment. A computer with a mouse or a touchscreen was required to participate.

Procedure

Participants played a computer game where their task was to choose stimuli presented on the screen in order to earn points. Some stimuli were 'tools' and other stimuli were 'apparatuses'; each 'tool' worked with a particular 'apparatus'. Across a series of levels participants were trained to use each tool with the corresponding apparatus by selecting a tool and then selecting the appropriate apparatus. Three 'tool' and 'apparatus' combinations were trained. Two groups were used, each with different stimuli. The stimuli used for Group Colour were coloured circles. Combination 1 consisted of a blue tool (T1) and a red apparatus (A1). Combination 2 was a pink tool (T2) and a green apparatus (A2), and Combination 3 was a purple tool (T3) and a yellow apparatus (A3). A grey circle was used as a distractor stimulus in Levels 7 and 8. The stimuli used for Group Tool were pictures of tools and materials. Combination 1 was a picture of a hammer (T1) and a picture of a nail (A1). Combination 2 was a picture of a saw (T2) and a picture of a plank of wood (A2) and Combination 3 was a picture of a screwdriver (T3) and a picture of a screw (A3). A picture of a pair of scissors was used as a distractor stimulus in Levels 7 and 8. The levels participants progressed through for both groups followed a similar sequence to the training used by Boeckle et al. (2020) with crows. Levels 1 to 6 were Training Levels and Levels 7 and 8 were Testing Levels.

Level 1

Level 1 mirrored the Tool Use Training undertaken by the crows (Boeckle et al., 2020). During this training, crows were taught how the first two tools worked with each

corresponding apparatus. For our participants, during each trial a tool stimulus and an apparatus stimulus were presented simultaneously on the screen, as illustrated in Figure 1.1 and Figure 1.2. To earn ten points the participant first had to click on the tool stimulus and then click on the apparatus stimulus. If the tool was clicked first, the message "nearly there…" appeared on the screen. If the correct sequence of responding was made a '10' appeared on the apparatus stimulus and the participant's total points score appeared at the bottom of the screen. If the apparatus stimulus was clicked first, the message "Mistake" appeared on the screen and the trial ended. If the participant took longer than 5 s to make a response the message "too slow" appeared on the screen and the trial ended. At the end of each trial there was a 3 s inter-trial interval (ITI) with the message "Get ready for the next trial". T1 and T2 stimuli were used during this level. The trial type and the position of the stimuli on the screen were randomised across trials. Participants played on the level until they chose correctly on three consecutive trials or until 100 trials had been completed. Each level began with a "click for the next level" message and ended with a "You've nailed the level" message and a display of the final points count for the level.

Level 2

Level 2 mirrored the Tool Selection Training undertaken by the crows (Boeckle et al., 2020). Crows were presented with one apparatus and both tools and taught to choose the correct tool for the available apparatus. For our participants, Level 2 was the same as Level 1 except one apparatus stimulus and two tool stimuli were presented each trial, as illustrated in Figure 1.1 and Figure 1.2. If either the apparatus or the incorrect tool was clicked first the message 'Mistake' appeared on the screen and the trial ended. There was no programmed consequence if the correct tool stimulus was clicked first. All other aspects of this level remained the same as Level 1.

Level 3

Level 3 mirrored the Apparatus Functionality Training undertaken by the crows (Boeckle et al., 2020). Crows were presented with one tool and both apparatuses and taught to use the tool on the correct apparatus. For our participants, Level 3 was the same as Level 2 except one tool stimulus and two apparatus stimuli were presented each trial, as illustrated in Figure 1.1 and Figure 1.2. If either of the apparatus stimuli was clicked first the message "Mistake" appeared on the screen. All other aspects of this level remained the same as Level 2.

Level 4

Level 4 mirrored the Mental Representation Training undertaken by the crows (Boeckle et al., 2020). Crows were presented with both tools and one apparatus and had to select the correct tool while the apparatus was no longer visible. For our participants, Level 4 was the same as Level 2 except the words "watch and wait" appeared at the bottom of the screen underneath the one apparatus stimulus and two tool stimuli for 5 s, as illustrated in Figure 1.1 and Figure 1.2. Clicks to the stimuli during this 5 s had no programmed consequence. After the 5 s had passed the apparatus stimulus disappeared from the screen. If the correct tool stimulus was clicked first, then the apparatus stimulus reappeared. All other aspects of this level remained the same as Level 2.

Level 5

Level 5 mimicked the Hook Training and Tool Transport Training undertaken by the crows (Boeckle et al., 2020). During this training crows were taught how to use the final tool with the corresponding apparatus. For our participants, Level 5 was the same as Level 1, except only T3 stimuli were used, as illustrated in Figure 1.1 and Figure 1.2. There was no programmed response after the tool was clicked first. All other aspects of this level remained the same as Level 1.

Level 6

Level 6 mimicked the Five Choice Functionality Training undertaken by the crows (Boeckle et al., 2020). Crows were presented with all three tools and one apparatus and taught to select the correct tool for the available apparatus. For our participants, one apparatus stimulus and all three tool stimuli were presented each trial, as illustrated in Figure 1.1 and Figure 1.2. T1, T2 and T3 stimuli were used during this level. If the apparatus stimulus or either of the two incorrect tool stimuli was clicked first the message "Mistake" appeared on the screen. All other aspects of this level remained the same as Level 2.

Level 7

Level 7 mimicked the Training Phase of the experiment proper, undertaken by the crows (Boeckle et al., 2020). Crows were trained on the temporal sequence that would occur during testing using the first tool-apparatus combination. For our participants, at the beginning of a trial A1 was presented on the screen for 5 s. The location of the apparatus stimulus was randomised across trials. After 5 s, A1 disappeared and a 2-minute delay period began with the message "select the matching card to get points". During this delay a Wisconsin Card Sorting Task (WCST) was used. For each trial of the WCST a sample card and three comparison cards were presented. Each comparison card matched the sample card on one aspect, either colour, number, or shape. The possible combinations of cards included colours of blue, yellow, red, or green, shapes of circles, stars, triangles, or crosses and either 1, 2, 3 or 4 objects. The correct aspect to match the card was changed after every 16 trials. If the correct card was chosen a '1' appeared on the screen next to the card as well as the participants total point count. If the incorrect card was clicked the message "Mistake" appeared on the screen and no points were earned. If the participant took longer than 10 s to respond the message "too slow!" appeared on the screen. Once the 2-minute delay period had passed, a low-value stimulus of a picture of a pigeon flashed up on the screen for 1 s. If this

was clicked the participant earned one point. If this was not clicked the message “missed a point!” appeared on the screen.

After the pigeon disappeared the Choice Phase began. All three tool stimuli, the low-value pigeon and a distractor stimulus were presented, as illustrated in Figure 1.1 and Figure 1.2. If the distractor stimulus was chosen the trial immediately ended. If the pigeon was chosen a ‘1’ point flashed up on the screen and the trial ended. If the participant took longer than 5 s to respond the message “too slow” appeared on the screen and the trial ended. If any of the three tools were chosen, then a second delay period of 5-minutes began. During this delay the same WCST was used. After five minutes had passed the chosen tool stimulus and the A1 stimulus were presented on the screen. If the participant had correctly chosen T1, a ‘10’ appeared on the apparatus stimulus after clicking T1 and then A1. If any other tool was chosen or the apparatus was clicked first the message ‘Mistake’ appeared on the screen. If the participant took longer than 5 s to respond the message “too slow” appeared on the screen. The location of the stimuli was randomised across trials. Level 7 ended after two trials had been completed.

Level 8

Level 8 mimicked the Testing Phase of the experiment proper, undertaken by the crows (Boeckle et al., 2020). During this phase, crows were tested with the same temporal sequence from the Training Phase using the final two tool-apparatus combinations. For our participants, Level 8 was the same as Level 7 expect for the trial types used. Level 8 consisted of two trials, the first trial was a T2 trial, and the second trial was a T3 trial. When the level was completed the message “You’ve nailed the experiment!” appeared on the screen followed by the participants points total for the level. All other aspects of this level remained the same as Level 7. At the end of the session participants exchanged their points for entries into a prize draw to win vouchers.

Figure 1.1

A Diagram of a Sample Trial During the Choice Phase from Each Level of the Procedure for Group Colour

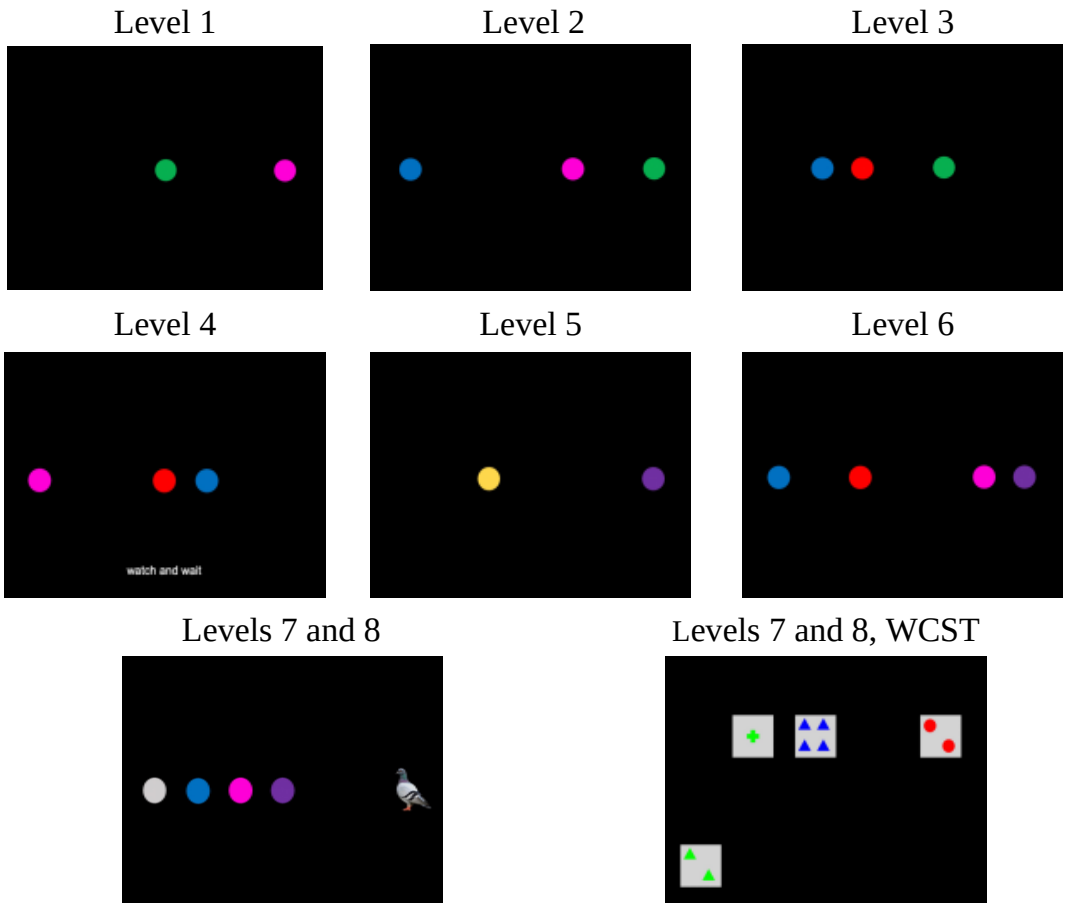

Figure 1.2

A Diagram of a Sample Trial During the Choice Phase from Each Level of the Procedure for Group Tool

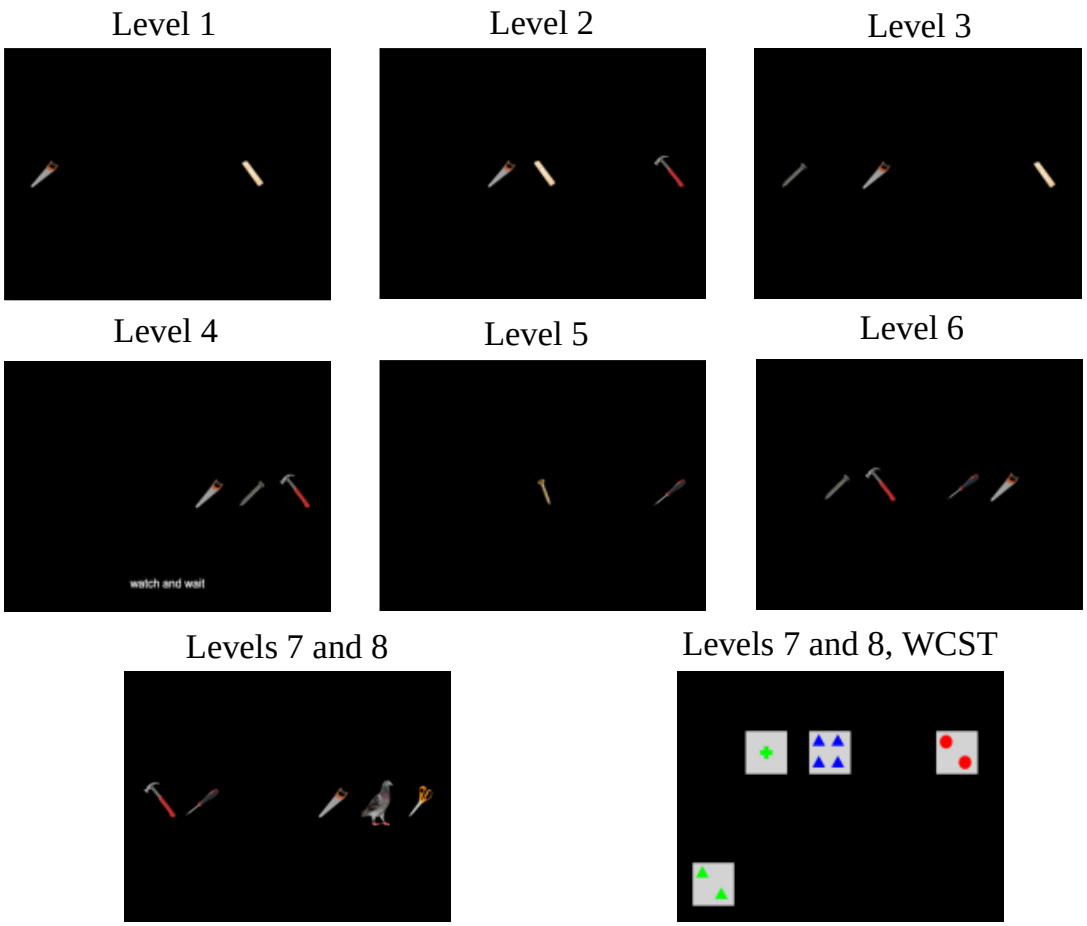

Results

We analysed performance of participants during Training (Levels 1 to 6) and Testing (Levels 7 and 8). Performance during Training tells us how participants were learning to make correct responses in the presence or absence of differing tools and apparatuses. Performance during Testing tells us whether the participants were able to use what they had learned during Training to choose the correct tool stimulus after a period of delay.

Training

Figure 1.3 shows the number of trials completed during Training for Groups 1 and 2. The white bars represent the mean number of trials completed across all participants. The coloured shapes represent the number of trials completed during each level for each individual participant. The pattern of behaviour across both groups was relatively similar. For Group Colour, the mean number of trials completed for the training criteria to be satisfied was higher in Levels 2 and 6 than in other levels. In both these levels extra tool stimuli were introduced, increasing to three total stimuli in Level 2 and four total stimuli in Level 6. For both groups, after the increase in mean number of trials for Level 2, the mean number of trials decreased as level number increased. The mean number of trials completed in Level 6 increased for Group Colour, but not for Group Tool. Level 5, where only one tool stimulus and one apparatus stimulus were present, required the lowest mean number of trials to reach the criterion for both groups. In Group Colour, only Participant 6 and Participant 14 reached the maximum number of trials that could be completed within a level (100) for Levels 2 to 6; the other participants completed these levels close to the mean, below 40 for Levels 2, 3 and 6, and below 20 for Levels 4 and 5. In Group Tool only Participant 7 reached the maximum number of trials in Levels 2 and 3; the other participants mostly completed these levels in under 20 trials. A one-way repeated measures ANOVA of the number of trials completed each level for Group Colour found a significant difference between the number of trials

required to complete Level 1 vs Level 2 ($p = .003$), Level 2 vs Level 3 ($p = .007$), Level 4 ($p = <.001$), Level 5 ($p = <.001$), and Level 6 ($p = .046$) and Level 5 vs Level 6 ($p = .049$). The same ANOVA for Group Tool found a significant difference between the number of trials required to complete Level 1 vs Level 4 ($p = .008$), Level 5 ($p = .004$), and Level 6 ($p = .010$), Level 2 vs Level 5 ($p = .035$) and Level 6 ($p = .037$), and Level 4 vs Level 5 ($p = .040$). A mixed ANOVA found a significant difference between the number of trials completed in Level 6 ($p = .038$) for Group Colour vs Group Tool.

Figure 1.4 shows the proportion of trials completed correctly during Training for Group Colour and Group Tool. A correct trial consisted of a tool being selected first, followed by the correct apparatus. An incorrect trial consisted of an incorrect tool being selected, or an incorrect apparatus being chosen, or a trial where no response was made after 5 s. The white bars represent the mean proportion of correct trials across participants for each level. The coloured shapes represent the proportion of correct trials for each individual participant during each level. For Group Colour, the mean proportion of correct trials decreased in both Levels 2 and 6, with a subsequent increase in proportion correct as level number increased between these two levels. For Group Tool, the mean proportion of correct trials increased as level number increased, beginning around 0.5 in Level 1, and steadily increasing to just below 1 in Level 6, but did not drop in Levels 2 and 6. While most individual participants' proportion of correct trials fell close to the mean, in Group Colour, eight participants completed all trials correctly in at least one level across Levels 3 to 6, and two participants (Participants 6 and 14) completed zero trials correctly across Levels 2 to 6. In Group Tool, all 10 participants completed at least one level with all correct trials across Levels 2 to 6. Only Participant 7 completed a level with no correct trials, in Levels 2 and 3. A one-way repeated measures ANOVA of the proportion of correct trials for Group Colour found a significant difference between Level 1 vs Level 2 ($p = <.001$), and Level 5 ($p = .014$),

Level 2 vs Level 3 (p = .004), Level 4 (p = <.001), Level 5 (p = <.001), and Level 6 (p = .001), Level 3 vs Level 5 (p = .020), Level 4 vs Level 6 (p = .007) and Level 5 vs Level 6 (p = .006). The same ANOVA for Group Tool found a significant difference between Level 1 vs Level 4 (p = .016), Level 5 (p = <.001), and Level 6 (p = <.001), and Level 2 vs Level 5 (p = .034) and Level 6 (p = .038). A mixed ANOVA found a significant difference between the proportion of trials completed correctly in Level 2 (p = .022) and Level 6 (p = .002) for Group Colour vs Group Tool.

Figure 1.5 shows the time (in seconds) each level took to complete during Training for Group Colour and Group Tool. The white bars represent the mean time it took to complete each level across participants. The coloured shapes represent the individual time it took each participant to complete each level. For Group Colour, the mean time taken to complete a level increased in Levels 2, 4 and 6, resulting in more time spent completing each level compared to the previous level. Level 4, where a 5 s waiting period occurred during each trial, took the longest to complete for Group Colour. Consistent with the higher mean number of trials and the lower mean proportion of correct trials in Levels 2 and 6, Group Colour took longer on average to complete Levels 2 and 6. For Group Tool, the mean time it took to complete each level was highest during Level 2 and decreased as level number increased after this level. Level 5, where only one tool stimulus and one apparatus stimulus were presented, took the shortest mean time to complete for Group Colour and Group Tool. For Group Colour, the individual time it took to complete each level varied, with most times close to or below the mean. An outlier for Participant 14 during Level 4 was removed from the figure as the time it took them to complete the level (1002 s) was triple that of the nearest data point. For Group Tool, the mean was representative of most of the individuals' performance. A one-way repeated measures ANOVA of the time it took to complete each level for Group Colour found a significant difference between Level 1 vs Level 2 (p = .014),

Level 2 vs Level 3 (p = .014), and Level 5 (p = .002) and Level 5 vs Level 6 (.036). The same ANOVA for Group Tool found a significant difference between Level 1 vs Level 5 (p = .027), and Level 6 (p = .045), Level 4 vs Level 5 (p = .019), and Level 6 (p = .027) and Level 5 vs Level 6 (p = .031). A mixed ANOVA found a significant difference between the time it took to complete Level 6 (p = .037) for Group Colour vs Group Tool.

We assessed the error rate for both groups to determine how many errors were made by each group on average across Training. Figure 1.6 shows the mean number of incorrect trials for Group Colour and Group Tool. Group Colour made 110 incorrect trials on average and Group Tool made 73 incorrect trials on average across all Training Levels. An independent samples t-test showed no statistical significance (p = .473) between the number of incorrect trials for Group Colour (*M* = 110, *SD* = 139) and Group Tool (*M* = 73, *SD* = 90).

Testing

Figure 1.7 shows the proportion of correct trials across participants during Testing for Group Colour. Group Tool made no correct trials during Testing and have therefore been excluded from this figure. Overall, the proportion of correct trials was low (~.2) in both levels, slightly increasing in Trial 2 of Level 8.

Figure 1.8 shows the frequency of choice for each item during the Choice Phase of Testing. During the Choice Phase in Levels 7 and 8, participants were presented with a choice between all three tool stimuli, a distractor stimulus, and a low-value stimulus. Across both groups the low-value stimulus was the most common choice. For Group Colour, the low-value stimulus was chosen most often across both levels and all trial numbers, except for Trial 1 of Level 7 where the distractor stimulus was the most chosen item. For Group Colour, in both trials of Level 7, the correct item (T1) was chosen only twice (by Participant 2 and 4 in Trial 1 and by Participant 2 and 3 in Trial 2). When T2 was the correct choice in Trial 1 of

Level 8 it was chosen twice (by Participant 10 and Participant 14), and when T3 was the correct choice in Trial 2 of Level 8 it was chosen three times (by Participants 3, 10 and 11). For Group Tool, no tools were chosen at any point during Testing. The low-value stimulus was the most chosen item across both levels and trial numbers. Trials where participants did not respond for more than 5 s, and therefore did not make a choice during the Choice Phase, were excluded from this figure. This included three trials from Group Colour, and six trials from Group Tool, mostly occurring during the first trial in Level 7 when the Choice Phase was first presented.

Figure 1.9 shows the proportion of correct WCST responses made during the delay periods of Testing for Group Colour and Group Tool. The white bars represent the mean proportion of correct WCST responses across participants for each level. The coloured circles represent the individual proportion of correct WCST responses for each individual participant in each level. Both groups showed a similar pattern of behaviour. Trial 1 of Level 7 (the first time the task was introduced) had the lowest mean proportion of correct responses for both groups. For Group Colour, the mean proportion of correct responses increased as trial number and level number increased, until Trial 2 of Level 8 where there was a slight decrease in proportion correct. For Group Tool, the mean proportion of correct responses increased as trial number and level number increased. The mean was representative of most of the individuals' performance across both groups. Participant 4 in Group Colour made no correct responses during Trial 2 of Level 8; but otherwise responded as the others did. A one-way repeated measures ANOVA of the proportion of correct WCST responses for Group Colour found a significant difference between Level 7, Trial 1 vs Level 7, Trial 2 ($p = .003$) and Level 8, Trial 1 ($p = <.001$). The same ANOVA for Group Tool found a significant difference between Level 7, Trial 1 vs Level 8, Trial 1 ($p = .005$), and Level 8, Trial 2 ($p = .006$), Level 8, Trial 1 vs Level 8, Trial 2 ($p = .025$) and Level 7, Trial 2 vs Level 8, Trial 2 (p

= .025). A mixed ANOVA found no significant difference between any of the four trials across both levels for Group Colour vs Group Tool.

Figure 1.3

Mean and Individual Number of Trials Completed During Training Levels for Group Colour and Group Tool

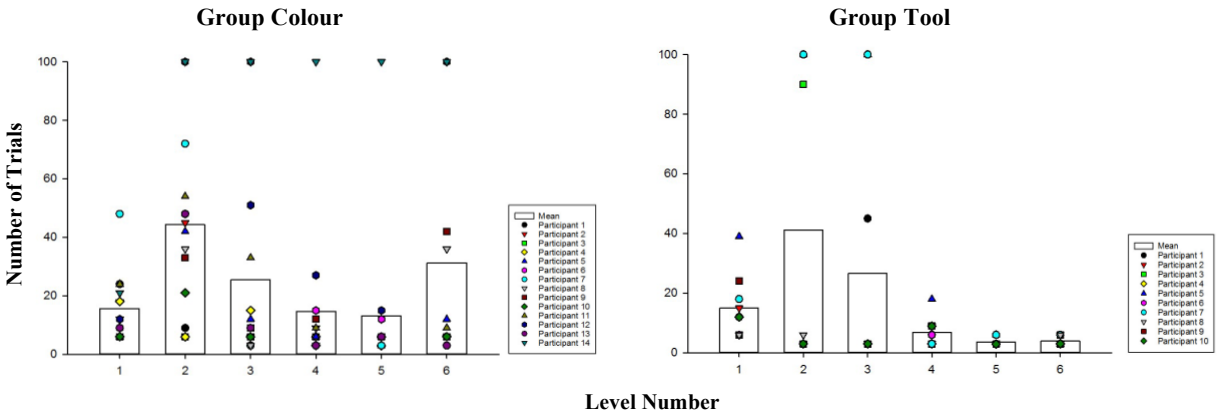

Figure 1.4

Mean and Individual Proportion of Correct Trials Completed During Training Levels for Group Colour and Group Tool

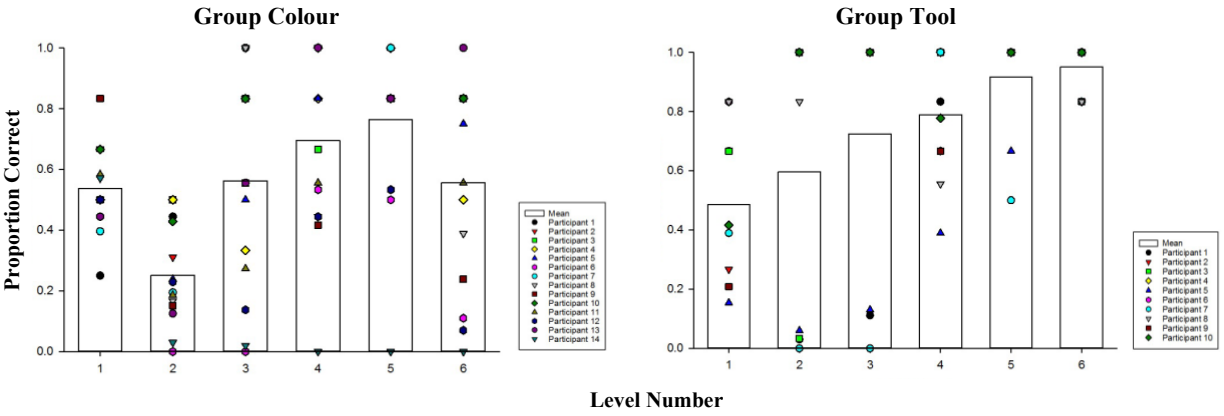

Figure 1.5

Mean and Individual Time (in Seconds) Each Level Took to Complete During Training Levels for Group Colour and Group Tool

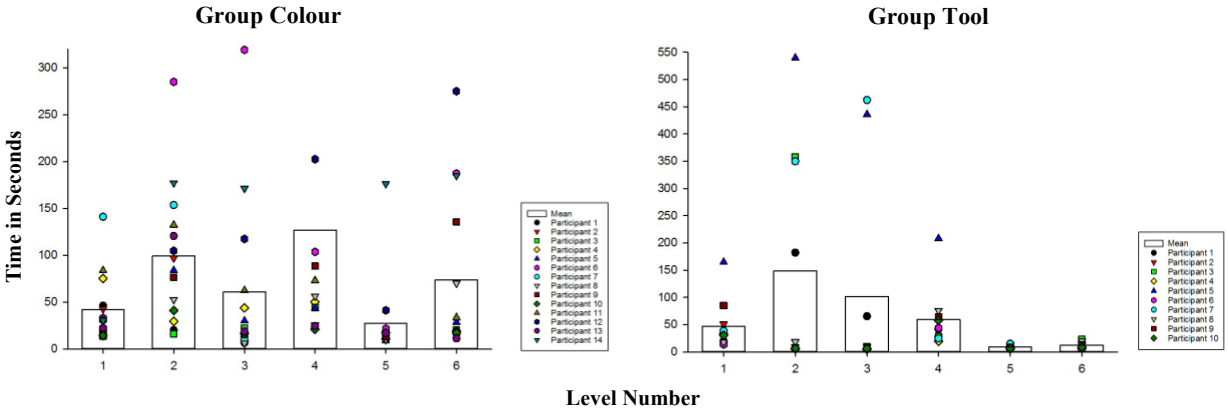

Figure 1.6

Mean Incorrect Trials Across all Training Levels for Group Colour and Group Tool

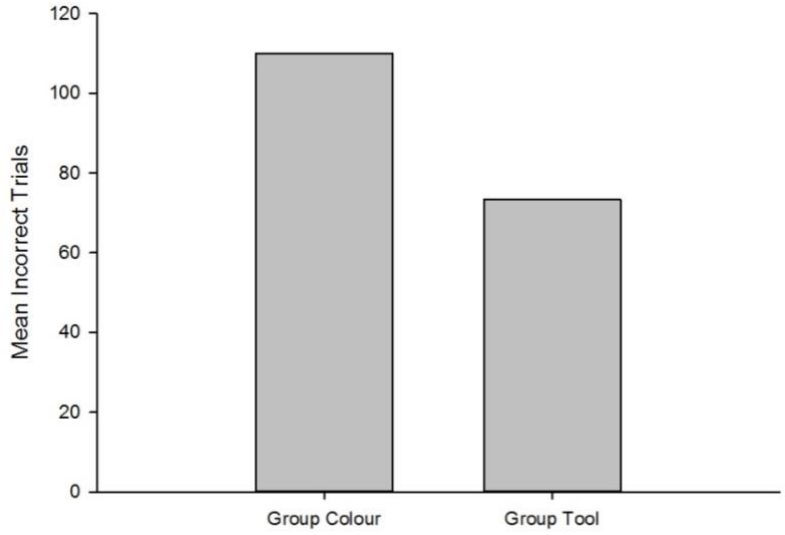

Figure 1.7

Proportion of Correct Trials During Testing Levels Across Participants for Group Colour

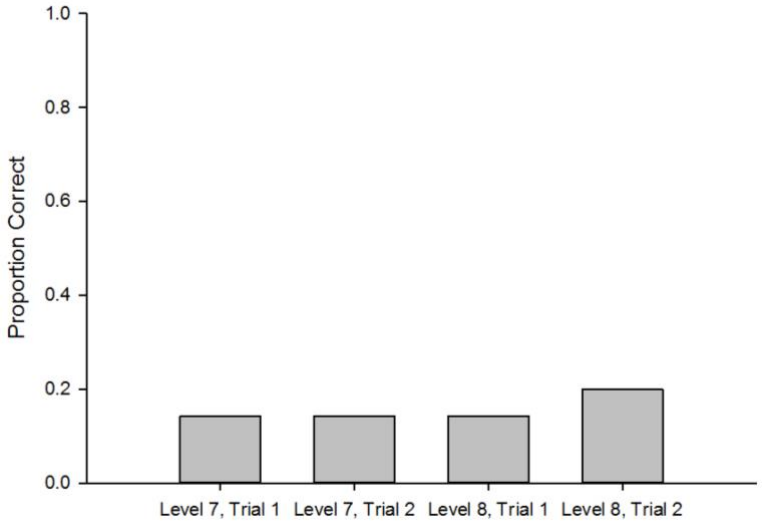

Figure 1.8

Number of Times Each Item was Chosen During Testing Trials Across Participants for Group Colour and Group Tool

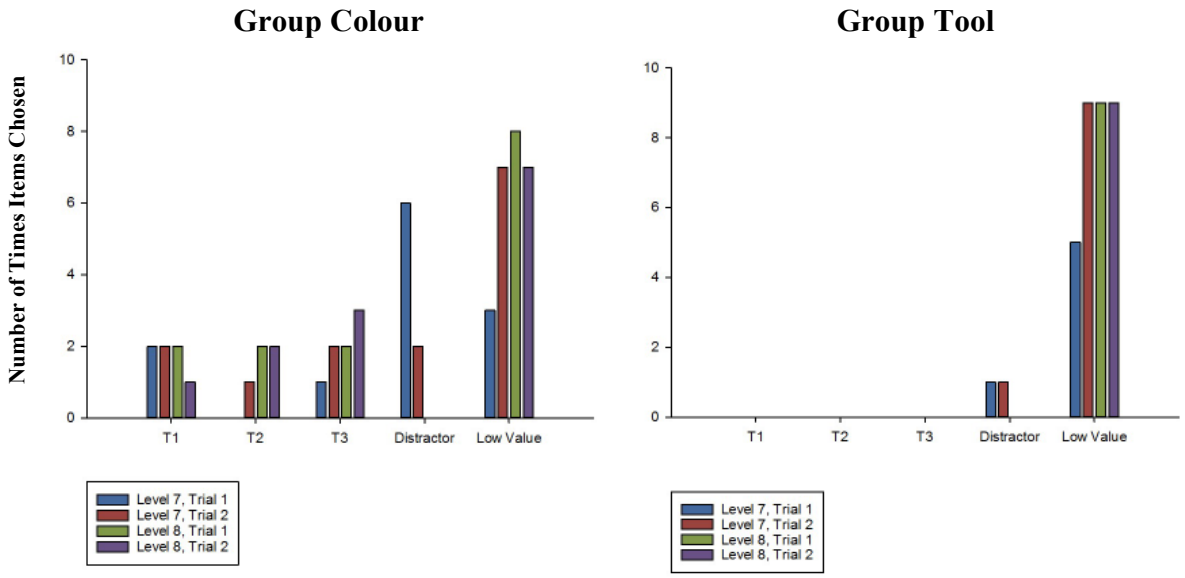

Figure 1.9

Mean and Individual Proportion of Correct WCST Responses During Testing Levels for Group Colour and Group Tool

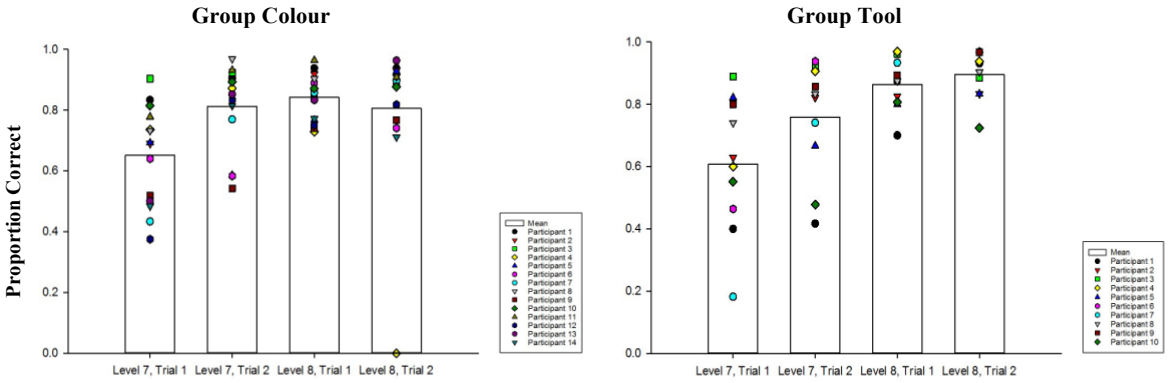

Discussion

We investigated the performance of humans on a planning task that had previously been used with animals. We replicated the Training and Testing conditions used by Boeckle et al. (2020) with crows and replaced the real tools and apparatuses with two types of stimuli that varied in their arbitrariness. We found that overall, the group that used less arbitrary stimuli (Group Tool) performed better during Training than the group that used arbitrary stimuli (Group Colour). The effect of added difficulty in new levels was particularly apparent for Group Colour; when an extra stimulus was added to a level, more trials were required to complete that level, more incorrect responses were made during that level, and more time was required to complete the level for Group Colour compared to Group Tool. Nevertheless, both groups were able to master the Training. While the proportion of correct choices during Testing was low across both groups, Group Colour performed better than Group Tool in Testing. Across both groups, the most common item chosen during Testing was the low-value stimulus, which provided a smaller more immediate reinforcer. Overall, we found that

the participants in our experiment performed significantly worse during Testing than Boeckle et al.'s (2020) crows, and Miller et al.'s (2020) children.

Our procedure replicated the Training and Testing used by Boeckle et al. (2020) with crows and aspects of Miller et al. (2020) with children. The crows used real tools and apparatuses and were trained across multiple levels on the functionality of the tool-apparatus combinations. They were then tested using a specific temporal sequence where delays were used between the presentation of the apparatus and the choice of an item, and between the Choice Phase and the ability to use the item on the available apparatus. For the sake of timing our Testing Levels used the delay periods and the number of trials used with children on the same procedure by Miller et al. during Testing (had we used the same timing as Boeckle et al., the experiment would have taken over 3 hours to complete). Boeckle et al. found that crows chose the correct tool during the Choice Phase of Testing significantly above chance. Miller et al. also found that children aged four to five could successfully choose the correct tool during Testing. In comparison, our participants were more likely to choose an incorrect stimulus than a correct one; that is, our adult human participants performed worse than both crows and children. Data for the performance of crows during Training was not provided, therefore we cannot directly compare the performance of our human participants with the crows during the Training Levels.

Our participants' failure to choose the correct tool in Testing cannot relate to the arbitrariness of the stimuli used. Despite performing better during Training, Group Tool surprisingly performed worse during Testing than Group Colour. While the proportion of correct trials for Group Colour was low, it was still higher than Group Tool who had no correct trials during Testing Levels. Thus, the use of familiar, non-arbitrary tool-apparatus combinations was not critical to (and indeed did not in fact facilitate) the ability to make a correct choice in Testing. This finding suggests that replication of the procedure with

arbitrary stimuli – a procedural feature critical to demonstrations of planning ability or lack thereof in 'simpler' organisms (e.g., pigeons) – is indeed feasible, at least in the sense that the arbitrariness of stimuli seems not to hinder task performance.

The arbitrariness of the stimuli may have indirect effects on what is learned, and hence on a participant's ability to choose correctly in Testing. One possible explanation for the difference in performance between Group Colour and Group Tool could be the difference in task difficulty between the groups, and its effect on the attention each group was paying to the task. The task of remembering three different colour combinations would have been more difficult for Group Colour than remembering the tool and material combinations for Group Tool. We can see this in the fact that Group Tool performed much better during Training. When tasks are more difficult, we tend to pay more attention to them (McDowd, 2007). Therefore, Group Colour may have been paying more attention to the task overall during both Training and Testing due to the added difficulty of remembering the colour combinations compared to Group Tool. Hence, the lack of attention from Group Tool during Training and Testing may have affected their performance, resulting in the incorrect stimulus being chosen more often. Regardless of why Group Colour performed better during Testing than Group Tool, this difference clearly shows that the arbitrariness of stimuli is not a factor that hampers performance.

Another explanation for why Group Colour performed better than Group Tool during Testing could be the difference in error rates during Training between the two groups. While there was no statistical significance between the error rates of the two groups, Group Colour made more errors on average during Training than Group Tool. Research has suggested that generalisation requires learning what is correct as well as learning what is not correct (Urcuioli & Nevin, 1975). Discrimination learning is also enhanced when both reinforcers and punishers are used, relative to when only reinforcers are used (e.g., see Brackbill &

O'Hara, 1958; Warden & Aylesworth, 1927) particularly when tasks are complex (Meyer & Offenbach, 1962). Therefore, the higher number of errors during Training for Group Colour would have allowed them more experience with what the incorrect responses were in comparison to Group Tool, facilitating better performance during Testing.

One explanation for why our results may have differed from those of the crows and children could be that we conducted the task using a computer-based program rather than using real tools and apparatuses. When the crows and the children completed this procedure, both used real tools and apparatuses that functioned together, and had to physically transport the tools to different rooms when choosing to use them (Boeckle et al., 2020; Miller et al., 2020). In comparison our participants used a computer mouse to click on objects on the screen to make their choice. This explanation, however, seems unlikely as our participants were able to learn during Training when trained in the same way as the crows were. Past behavioural research has also shown no difference in results when participants completed a tool-selection procedure in person compared to online (Casler et al., 2013). In Casler et al.'s experiment participants were presented with a novel tool and a familiar tool and asked to choose which tool they needed to complete a specific goal. The in-person participants physically handled the tools before making their choice whereas the online participants were shown videos of the tools and had to click a button on the screen to make their choice. The results showed no differences across the groups in how often they chose each type of tool. This suggests that human participants can learn to perform behavioural tasks to an equal level online when compared to the same in-person task. Of course, our participants saw only static images of the tools and apparatuses, but given they mastered the Training Levels, seeing the tool move does not appear to be central to being able to learn how it relates to an apparatus.

Unlike Boeckle et al.'s (2020) crows and Miller et al.'s (2020) children, our participants tended not to choose correctly during Testing. The most common item chosen

during Testing for both Group Colour and Group Tool was the low-value stimulus, which provided the participant with an immediate low-value reinforcer of one point. If a correct tool was chosen during the Choice Phase, then the participant would earn ten points after a 5-min delay. These two choices can be thought of as a choice between a 'smaller, sooner' reward (low-value stimulus) and a 'larger, later' reward (correct tool) (Rachlin, 1995). The choice of a smaller more immediate reward over a larger more valuable reward would suggest the participants were more impulsive and exerting less self-control (Rachlin, 1995). The participants in our experiment may have been impulsive in this particular situation, preferring to choose the immediate reward of one point instead of waiting five minutes to gain the later reward of ten points. It is also possible that our participants were more motivated to end the experiment, rather than to earn points, choosing the option that would end the experiment sooner, rather than the option that extended the experiment. Of course, given that impulsivity reduces across the lifespan (Green et al., 1994, 1999), and given humans are often more immune to the effects of delay than non-humans (e.g., Vanderveldt et al., 2016) we might have expected our adult participants to outperform the children and crows.

Another explanation for why so many incorrect responses were made by both groups in Testing could be that the participants were forgetting. The presentation of an apparatus stimulus at the beginning of a Testing trial was followed by a 2-min delay period before the Choice Phase began. During this delay period subjects completed a WCST. It is possible that by the time the subject reached the Choice Phase, they had forgotten that they had been presented with the apparatus at the beginning of the trial. When humans must hold multiple items in working memory longer delays lead to worse recall and more errors (Pertzov et al., 2017). Participants already had to remember three different matching pairs; therefore, the combination of the multiple stimuli and the delay length may have contributed to them forgetting. When children performed this same task using the same delay lengths as our

participants older children were successful in choosing the correct tool (Miller et al., 2020), suggesting that they were not forgetting. This difference between our participants and Miller et al.'s children could be due to the nature of the apparatus stimulus. It seems likely that an object that disappears from a screen would be easier to forget than a physical apparatus that is located in the room next to you, leading to a greater likelihood of forgetting in our participants.

One other potential explanation is that our procedure inadvertently encouraged participants to choose the low-value stimulus over the correct one. The low-value stimulus was selected from the same array as was used for a distractor task, a procedural decision that ensured our participants learned about the value of this stimulus. The distractor task employed specific instructions to click the low-value stimulus. Further, every click to the low-value stimulus was reinforced (whereas not every click to a tool or apparatus stimulus was reinforced; stimuli that are associated with higher overall rates of reinforcement tend to be more frequently chosen – see Baum, 1974; Davison & Elliffe, 2010; Shahan & Podlesnik, 2006). These factors, and the relative recency of this choice being the correct response may have overshadowed control by the more extended structure of the task (i.e., by the expectation that the apparatus would be presented again). Certainly, this sets our procedure apart from the procedures used by Boeckle et al. (2020) and Miller et al. (2020) who used a less preferred food item and a less preferred sticker, respectively. Future research may address this by using a low-value stimulus that is never explicitly the correct choice (e.g., a "+1 point" option.)

Although Group Tool performed worse in Testing than Group Colour, during Training Group Tool performed better overall than Group Colour. As level number increased the average performance of Group Tool improved. Group Tool completed fewer trials per level, and more correct than incorrect trials, in less time relative to Group Colour. The main

difference between the two groups was the stimuli used to represent the tools and apparatuses. Group Colour used coloured circles, whereas Group Tool used pictures of tools and materials. Hence, Group Colour learned an arbitrary association between different colours, while Group Tool learned an association between objects that likely already existed in their repertoire. The time it takes to learn relations between stimuli can differ depending on the meaningfulness of the stimuli. Flemming et al. (2008) investigated matching performance in humans using either meaningful (stimuli that evoke an external object or concept – e.g., letters that form a word) or non-meaningful (e.g., a random string of letters) stimuli. Participants in Flemming et al.'s experiment who were presented with meaningful stimuli learned the task faster and completed the task with a higher number of correct trials than those who were presented with non-meaningful stimuli. Within our procedure the match between a hammer and a nail is more meaningful, as this evokes a concept that is commonly known, compared to a match between the colour blue and the colour red. Additionally, the association between two different colours may be a particularly challenging association to learn. Research with pigeons has demonstrated that when colours have to be matched together, the use of two different colours results in worse performance than when two of the same colour are used (Hogan et al., 1981). This could also explain why performance was worse for Group Colour who had to learn three different colour matching combinations, than for Group Tool who had to learn associations between more meaningful, less arbitrary stimuli.

Regardless of the arbitrariness of the stimuli, performance worsened when an extra stimulus was introduced. During Level 2 the number of stimuli presented during the Choice Phase rose to three, from two in the previous level. In Level 6 the number of stimuli presented during the Choice Phase rose to four, from two in the previous level. The addition of extra stimuli during these levels increased the ratio of wrong and right choices, increasing

the number of incorrect stimuli. More incorrect stimuli would have resulted in a higher chance of an incorrect stimulus being selected, particularly if a participant was choosing randomly. The extra stimuli may have also made it harder for the participants to remember which two stimuli matched due to retroactive interference. Retroactive interference occurs when new information interferes with old information and is worsened when items are unrelated and more competing items are present (Bower et al., 1994). For both groups the addition of extra stimuli led to worse performance, however this effect was seen to a greater extent for Group Colour. The extra stimuli that Group Colour were presented with existed within the same general category as the stimuli already present potentially leading to more interference between the options. For Group Tool, the extra stimuli existed within a different category than those already present leading to less interference between the options than Group Colour.

The main difference between our experiment and that done by Boeckle et al. (2020) was the stimuli used. We found that when arbitrary stimuli were used performance was worse during Training but slightly better during Testing. When less arbitrary stimuli were used performance was better during Training but worse during Testing. This may suggest that the stimuli used in these types of planning tasks need to fall in a middle ground between being too easy to remember (and therefore not facilitating learning about the structure of the task) and too hard to remember. For humans it seems that stimuli that are too easy to remember result in worse performance when the task is made more difficult, whereas stimuli that are too hard to remember result in relations between stimuli taking longer to learn but generalizing to a slightly greater extent. Future research could focus on using a similar task to our procedure, but testing this with stimuli that evoke differing levels of difficulty – for example, using stimuli that are familiar to humans but matching them with stimuli that would not normally be matched together. Overall, humans can learn when using a task that employs arbitrary

stimuli and behaviours, however this learning did not seem to generalise during Testing. Understanding why this learning did not generalise will be critical to future research, both in facilitating the design of an effective and fair test of planning in animals with species-general stimuli and in adding to knowledge about planning generally.

Chapter 2

In Chapter 1 we replicated an experiment previously used with crows (Boeckle et al., 2020) and children (Miller et al., 2020) with adult human participants, using arbitrary stimuli and responses. Surprisingly we found that our adult human participants performed worse than both crows and children during Testing. Our human participants frequently chose the low-value stimulus over the correct tool stimulus during the Choice Phase of Testing. One explanation for why participants may have chosen the low-value stimulus so frequently could have been that our procedure inadvertently encouraged participants to choose the low-value stimulus over the correct one. The low-value stimulus was selected from the same array as was used for a distractor task, so that the participants learned the value of this stimulus. Every click to this stimulus was reinforced, whereas correct tool or apparatus stimuli clicks were not always reinforced. This factor, combined with the relative recency of the low-value stimulus being the correct response may have overshadowed control by the more extended structure of the task. We therefore addressed this potential limitation by running the same experiment as in Chapter 1, but with a low-value stimulus option that had never explicitly been the correct choice (a "+1 point" stimulus). This allowed us to investigate whether the tendency to choose the low-value stimulus was due to its recent history of being the correct choice.

Method

Participants

Participants were 9 humans (6 males, 3 females) aged 19 to 36 years.

Apparatus

Experimental sessions and data collection were conducted online using Psytoolkit (Stoet, 2010, 2017). Participants were provided with a link that allowed them to access the experiment. A computer with a mouse or a touchscreen was required to participate.

Procedure

The procedure for this experiment replicated the procedure used for the Tool Experiment in Chapter 1. Participants transitioned through all six Training Levels (Levels 1 to 6) and both Testing Levels (Levels 7 and 8) using the same stimuli as Group Colour (coloured circles). The difference between this procedure and the procedure used in Chapter 1 was the low-value stimulus used. In Levels 7 and 8 of the present experiment the Choice Phase began immediately after the 2-min delay period and the low-value stimulus presented during the Choice Phase was a white square displaying the words "+1 point", rather than the pigeon picture used previously as a distractor task (as in Chapter 1). All other aspects of the procedure remained the same as the Tool Experiment in Chapter 1 (see the Method section of Chapter 1 for more detail).

Results and Discussion

As we were interested in the effects of the low-value stimulus on choice during Testing we analysed the choices made during the Choice Phase of Level 7 and 8 (Testing Levels). Figure 2.1 shows the frequency of choice for each item presented during the Choice Phase of the Testing Levels – all three tool stimuli, a distractor stimulus, and a low-value stimulus that had never previously been the correct choice. As with the Tool Experiment from Chapter 1, the low-value stimulus was the most common choice. Six participants chose the low-value stimulus in every trial in which they made a choice. One correct choice of T1 was made by a participant in Trial 1 of Level 7 and one correct choice of T3 was made by another participant in Trial 2 of Level 8. Trials where participants did not respond for more

than 5 s, and therefore did not make a choice during the Choice Phase, were excluded from this figure. Participants did not make a response in four instances, twice in Trial 1 of Level 7 and once in Trial 2 of Level 7 and Trial 2 of Level 8.

Figure 2.1

Number of Times Each Item was Chosen During Testing Trials Across all Participants

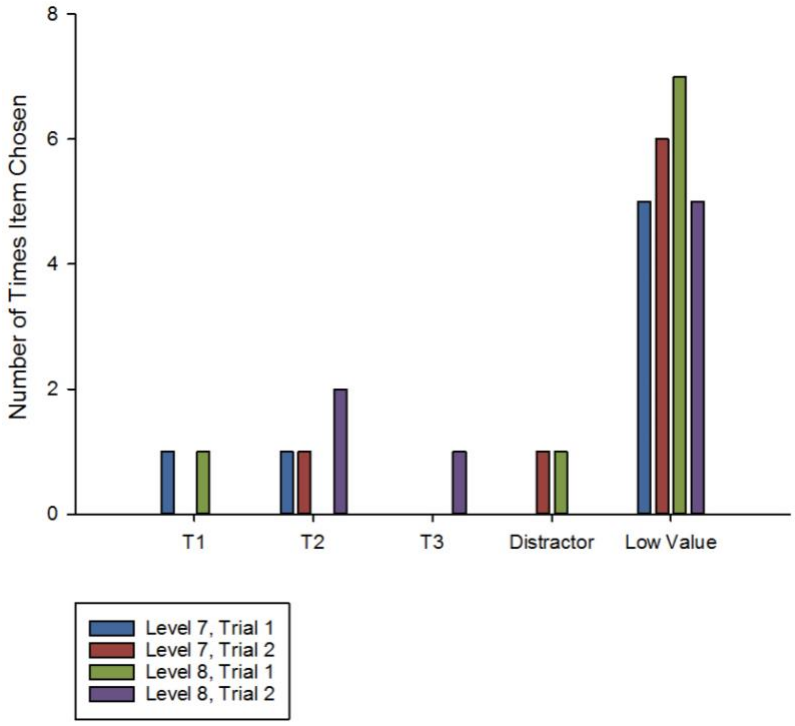

In this experiment, the low-value stimulus used had never explicitly been the correct choice. The first exposure the participants had to this stimulus during this experiment was during the Choice Phase of the first trial in Level 7. In the Tool Experiment in Chapter 1 the low-value stimulus was used during a distractor task where it was explicitly the correct choice directly before the participants chose an item in the Choice Phase. Given that the participants in this experiment continued to choose the low-value stimulus at a high rate, the stimulus' association with correct responding in the experiment from Chapter 1 appears not

to have been the cause of participants choosing this stimulus. We can conclude that the structure of the procedure and the low-value stimulus in Chapter 1 was not the cause of the frequency of choice of this stimulus. Rather, the participants must be choosing this stimulus for some other reason not related to our procedural set up (see the Discussion Section in Chapter 1 for a more in-depth review).

Chapter 3

Many demonstrations of planning use species-specific behaviour that often do not translate well into paradigms for other species and therefore do not facilitate cross-species comparisons of planning behaviours. Given debate around the definitions of what does or does not constitute prospection and planning, comparisons of the way different species approach planning tasks may be particularly useful for exploring the extent to which future-oriented behaviours are the same across species. Operant behaviour procedures do not rely on species-specific behaviour, and therefore can capture important elements of planning and future-oriented behaviour. These procedures are relatively easily adapted across different species, and hence can be used to explore differences and similarities in how species learn and engage with the task.

One potentially relevant procedure was that used by Cowie and Davison (2021), which was used to investigate whether pigeons could learn to make a response that was temporally and spatially distant from a reinforcer (see also Azrin et al., 1965). In Cowie and Davison's procedure a response needed to be made to one key at the beginning of a trial, followed by 25 responses to another key in order to produce a reinforcer. Part of the criteria for planning is that the planning behaviour is distant from its payoffs. For example, going to the supermarket today ensures I can cook dinner tonight. In Cowie and Davison's procedure the first response made in a trial was a planning response in this regard. In the present study, we compared signatures of learning on this task from human and pigeon participants. We re-

analysed data from the first two conditions of Cowie and Davison's study with pigeons (as they only looked at stable performance) and collected data from the same sort of procedure with humans choosing cards rather than pressing keys. This allowed us to make a direct comparison across different species, and to ask how future-oriented behaviour develops.

Method (Pigeons)

Subjects

The subjects were 5 pigeons numbered 31 to 35. They were housed in a shifted light/dark cycle where the room lights were turned off at 4 pm and turned on at midnight. The pigeons were maintained at 85% ± 15 g of their free-feeding weight. They were weighed after experimental sessions and fed mixed grain to maintain them at their set weight. They had access to water and grit at all times.

Apparatus

The pigeons were housed individually in their home cages which also served as experimental chambers. The home cages measured 375 mm high by 375 mm deep by 370 mm wide. Two wooden perches were mounted 100 mm from the wall and 20 mm from the floor. On the response panel, there were three 20 mm diameter translucent plastic keys set 100 mm apart centre to centre and 200 mm above the floor. Each key could be illuminated green, yellow, or blue. Responses to illuminated keys exceeding about 0.1 N were recorded. Beneath the centre key, 60 mm from the perch, was a magazine aperture measuring 40 mm high by 40 mm wide and 40 mm deep. When a reinforcer was delivered, key lights were extinguished, the aperture was illuminated, and the hopper containing wheat was raised for 3 s. All experimental events were programmed and recorded by a computer running MED PC® IV and located in another room. The subjects could see and hear other pigeons in the room during experimental sessions, but no person entered the room during this time.

Procedure

The pigeons had previous experience in a range of different procedures and required no pre-training. Experimental sessions ran everyday beginning at one am. Two keys were available for each trial: an Outcome key and a Plan key. Responses to the Outcome key produced either 3 s access to food provided the first response in that trial was made to the Plan key, or a 3 s blackout after 25 Outcome-key pecks if the first response in the trial was to the Outcome key. Pecks to the Plan key after the first peck in a trial had no programmed consequence. The left and right keys were always green and yellow at the start of each trial. In Unsignalled conditions, the keys remained the same colour throughout the trial; regardless of responses emitted; in Signalled conditions an effective peck to the Plan key would cause the Plan key to change colour to blue for the remainder of the trial. The location of the Outcome and Plan keys on the left or right was reversed across conditions. Sessions ran for 60 trials or 60 minutes, whichever came first. 64 sessions per condition were run.

Method (Humans)

Participants

The participants were 7 humans (5 females, 2 males) aged 21 to 29 years.

Apparatus

Experimental sessions were conducted, and data collected using the online software Psytoolkit (Stoet, 2010, 2017). Participants were provided with a link that allowed them to access the experiment. A computer with a mouse or a touchscreen was required to participate.

Procedure

During each trial the participant was presented with two playing cards representing two card decks – a Plan deck and an Outcome deck. In each trial, participants could choose to play cards from each deck in whatever sequence they liked; if the first card chosen was from

the Plan deck, the 12th card dealt from the Outcome deck would end the trial with 10 points. If the first card chosen was from the Outcome deck, then the 12th card dealt from the Outcome deck ended the trial with a message that showed '0 points'. The Plan card deck was located on the left; Plan cards had a white background with two blue squares. The Outcome card deck, located on the right, consisted of cards with a white background with four red stars. The participant was instructed to make a choice by clicking on one of the two cards. Once a choice had been made, a card dealt by the experimenter, with a white background with three purple triangles appeared on screen briefly, to signify the participant's response had been recorded. At the end of each trial, the participant's total points count was presented on the screen, followed by a 3-s inter-trial interval (ITI) before the next trial began. If a participant had not clicked on either of the cards for more than 5 s the two cards disappeared and the words "too slow" appeared on the screen, and the trial continued once this stimulus disappeared.

Each participant underwent two conditions: An Unsignalled condition in which the stimuli on screen remained the same whether or not the trial was going to end in point gain, and a Signalled condition in which the choice of a Plan card at the start of the trial would change Plan-card squares from blue to green squares for the rest of the trial. All participants experienced the Unsignalled condition first, followed by the Signalled condition. Each condition was run until a total of five correct trials in a row occurred. At the end of the session, participants were directed to a Qualtrics survey where they could enter a unique code to exchange their points for entries into a prize draw to win vouchers.

Results

To analyse performance for the pigeons we used data from the first session of the Unsignalled condition and the first session of the Signalled condition. The first session of the Unsignalled condition ended after 60 minutes for all pigeons, resulting in a variable number

of trials completed for each pigeon (*M* = 15, Range = 2-40). As Pigeons 31 to 33 completed a small number of trials (Range = 2-6) within the first session of the Unsignalled condition we also analysed data from the second session of the Unsignalled condition for these pigeons. In the first session of the Signalled condition all pigeons completed 60 trials. To analyse performance for the humans we used data from the Unsignalled condition and the Signalled condition. Within each condition, the number of trials a participant completed before reaching the criterion of five correct trials in a row varied (Unsignalled range: 7-31, Signalled range = 5-9). Most participants completed more trials in the Unsignalled condition (*M* = 15) than the Signalled condition (*M* = 6).

For both pigeons and humans, a successful trial (i.e., a trial ending in access to a reinforcer) involved starting with a Plan response, followed by a specified number of Outcome responses. A signature of efficient performance within this procedure would involve a tendency to begin a trial with a Plan response, and subsequent avoidance of Plan responses after the first response in a trial, resulting in completion of a trial with the minimum number of responses (26 for pigeons, 13 for humans). Thus, efficient responding would be evident in both choice at each response position within a trial, and in the number of responses per trial.

We assessed the overall pattern of choice within a trial by calculating the proportion of 'correct' responses in each trial. A correct response was defined as a Plan response for the first response, and an Outcome response for any responses thereafter. Figures 3.1 and 3.2 show these analyses for pigeons and humans respectively. The left panel shows trials from the Unsignalled condition, and the right panel shows trials from the Signalled condition. For pigeons (Figure 3.1), in the Unsignalled condition the proportion of correct responses began low and tended to increase as trial number increased, except for Pigeon 31 whose proportion correct decreased as trial number increased. Proportion correct at the end of the Unsignalled

condition tended to be around 0.8, except for Pigeon 31 whose proportion correct was below 0.2 at the end of the condition. Pigeons who completed more trials within the Unsignalled condition (Pigeons 33, 34 and 35) made more correct responses than incorrect responses across the condition. In the Signalled condition proportion correct began and remained high for all five pigeons across the condition, with most trials above a 0.8 proportion of correct responses. By the end of the Signalled condition all pigeons were at or close to a proportion correct of 1. Pigeons made more correct responses in a trial in the Signalled condition than the Unsignalled condition. A paired samples t-test showed statistical significance ($p = .044$) between the mean proportion of correct responses in the Unsignalled ($M = .629$, $SD = .207$) and Signalled condition ($M = .933$, $SD = .027$).

For the human participants (Figure 3.2), there was no consistent pattern across participants for the proportion of correct responses as trial number increased in the Unsignalled condition. The proportion of correct responses often started at or around 0.5 at the beginning of the condition. For some participants (Participants 1, 3 and 6) their proportion of correct responses increased as trial number increased, some (Participants 2, 5 and 7) decreased their proportion of correct responses as trial number increased, and one participant (Participant 4) remained relatively consistent across the condition. Proportion correct at the end of the Unsignalled condition tended to be around 0.5 for all participants. The only exception was Participant 1 who reached a proportion correct of 1 by the end of the Unsignalled condition. In the Signalled condition, proportion correct tended to start at 0.5 in Trial 1 and then increase as trial number increased. By the end of the Signalled condition, three participants (Participants 1, 2 and 7) had a proportion of correct responses of 1, three (Participants 3, 4 and 5) had a proportion around 0.5 and one (Participant 6) had a proportion of correct responses of 0.2. A paired samples t-test showed no statistical significance ($p =$

.136) between the mean proportion of correct responses in the Unsignalled (*M* = .523, *SD* = .126) and Signalled condition *(M* = .658, *SD* = .260).

Figures 3.3 and 3.4 show the proportion of first responses in a trial that were Plan responses across a condition for pigeons and humans, respectively. For pigeons (Figure 3.3), the proportion of trials beginning with a Plan response in the Unsignalled condition remained below 0.5, showing they were staring more trials with an Outcome response than a Plan response. Pigeon 32 made no first responses to the Plan key in the Unsignalled condition. All pigeons made a higher proportion of first responses to the Plan key in the Signalled condition compared to the Unsignalled condition. All pigeons, except for Pigeon 34, had a proportion of first responses to the Plan key above 0.6 in the Signalled condition. A paired samples t-test showed statistical significance (p = .005) between the proportion of first responses to the Plan key in the Unsignalled (*M* = .153, *SD* = .115) and Signalled condition (*M* = .640, *SD* = .140).

As with the pigeons, human participants had a higher proportion of first Plan responses in the Signalled condition compared to the Unsignalled condition (Figure 3.4), except for Participant 5, who had a higher proportion in the Unsignalled condition. In the Unsignalled condition all participants had a proportion of first Plan responses at or above 0.5, showing more trials were beginning with a Plan response than an Outcome response. In the Signalled condition, all proportion of first Plan responses fell around or above 0.8, with five participants (Participants 1, 2, 3, 6 and 7) starting every trial with a Plan response. A paired samples t-test showed statistical significance (p = .016) between the proportion of first Plan responses in the Unsignalled (*M* = .687, *SD* = .157) and Signalled condition (*M* = .944, *SD* = .096).

Figures 3.5 and 3.6 show the proportion of responses after the first response in a trial that were Plan responses (i.e., ineffective/unnecessary Plan responses) for the pigeons and

humans, respectively. For pigeons (Figure 3.5), in the Unsignalled condition the proportion of Plan responses tended to begin high and decreased as trial number increased, except for Pigeon 31 whose proportion of Plan responses increased as trial number increased. An exception was Pigeon 34 whose proportion of Plan responses began relatively low and did not systematically increase or decrease as trial count increased. In the Signalled condition, the proportion of Plan responses began low and remained relatively low across the condition for all pigeons. A paired samples t-test showed no statistical significance ($p = .051$) between the mean proportion of Plan responses in the Unsignalled ($M = .351$, $SD = .212$) and Signalled condition ($M = .054$, $SD = .029$).

For the human participants (Figure 3.6), in the Unsignalled condition the proportion of Plan responses tended to increase in earlier trials and then decrease, becoming stable at around 0.5 for most participants by the end of the condition. The only exception was Participant 1 for whom the number of Plan responses decreased to zero by the end of the condition. In the Signalled condition most participants proportion of Plan responses started around 0.5 at the beginning of the condition and remained relatively stable as trial number increased. Participant 1, 2 and 7's proportion of Plan responses decreased as trial number increased to zero Plan responses by the end of the condition. A paired samples t-test showed no statistical significance ($p = .173$) between the mean proportion of Plan responses in the Unsignalled ($M = .474$, $SD = .124$) and Signalled condition ($M = .351$, $SD = .263$).

Figures 3.7 and 3.8 show the total number of responses made within a trial for the pigeons and humans, respectively. Within a trial if a Plan response was made first the minimum number of responses required to produce a reinforcer was 26 for pigeons and 13 for humans. If an Outcome response was made first (so that the trial ended without a reinforcer) the minimum number of responses required to end the trial was 25 for pigeons and 12 for humans. For pigeons (Figure 3.7), in the Unsignalled condition response number in earlier

trials tended to be much higher than the minimum number of responses required within a trial. As trial number increased response number tended to decrease, except for Pigeon 31 whose response number increased near the end of the condition. Pigeons who completed more trials within the Unsignalled condition (Pigeons 33, 34 and 35) were more likely to complete a trial in fewer responses. In the Signalled condition the number of responses made by each pigeon tended to start and remain around the minimum number of responses required in a trial across the condition. A paired samples t-test showed no statistical significance ($p = .107$) between the mean number of responses made in a trial in the Unsignalled ($M = 55$, $SD = 29$) and Signalled condition ($M = 27$, $SD = 1$).

For the human participants (Figure 3.8), in the Unsignalled condition response number in earlier trials tended to be higher than the minimum number of responses required in a trial. As trial number increased, response number tended to decrease for all participants. In the Signalled condition most participants began the first few trials with response numbers slightly higher than the minimum. As trial number increased participants response number fell close to the minimum number of responses required, except for Participant 6 whose response number remained relatively high throughout the condition. A paired samples t-test showed no statistical significance ($p = .178$) between the mean number of responses made in a trial in the Unsignalled ($M = 37$, $SD = 18$) and Signalled condition ($M = 26$, $SD = 18$).

Figures 3.9 and 3.10 show the time (in seconds) a trial took to complete for pigeons and humans, respectively. For pigeons (Figure 3.9), in the Unsignalled condition the time a trial took to complete decreased as trial number increased, except for Pigeon 31 who increased the time a trial took to complete as trial number increased. Pigeons who completed fewer trials within the Unsignalled condition (Pigeons 31 and 32) took longer to complete a trial than pigeons who completed more trials (Pigeons 33, 34 and 35). By the end of the Unsignalled condition all pigeons were taking 200 s or less to complete a trial, except for

Pigeons 31 and 32 who were taking around 2000 s and 1200 s, respectively. In the Signalled condition, the time taken to complete each trial tended to be short at the beginning of the condition and would increase about halfway through the condition. Trial times would then generally decrease towards the end of the condition. By the end of the Signalled condition all pigeons were taking about 30 s to complete a trial. A paired samples t-test showed no statistical significance ($p = .124$) between the mean time trials took to complete in the Unsignalled ($M = 515$, $SD = 562$) and the Signalled condition ($M = 26$, $SD = 7$).

For the human participants (Figure 3.10), in the Unsignalled condition the time each trial took to complete tended to decrease as trial number increased for most participants. Participant 5 and Participant 7 both increased the time a trial took in the middle of the condition, but this time then decreased at the end of the condition. By the end of the Unsignalled condition all participants were taking less than 50 s to complete a trial, except for Participant 6 who was taking about 100 s to complete a trial. In the Signalled condition the time each trial took to complete began and remained consistently low as trial count increased. The only exception was Participant 6 who started the condition at a longer time and increased the time it took to complete each trial as trial count increased. By the end of the Signalled condition all participants were taking around 30 s to complete a trial, except for Participant 6 who was taking about 150 s to complete a trial. A paired samples t-test showed no statistical significance ($p = .197$) between the mean time trials took to complete for the Unsignalled ($M = 55$, $SD = 29$) and the Signalled condition ($M = 37$, $SD = 32$).

We assessed behaviour that followed an incorrect response (i.e., a first Outcome response) to see whether learning during the first response served as a cue for what to do next. Figures 3.11 and 3.12 show the proportion of Plan responses as a function of response number (up to 100 responses) for trials that began with an Outcome response (no-food or no-point trials) for pigeons and humans, respectively. For pigeons (Figure 3.11), in the

Unsignalled condition Plan responses occurred at all points within a trial. For Pigeons 31 and 32 the proportion of Plan responses tended to be highest around the beginning and middle of a trial. For Pigeons 33, 34 and 35 the proportion of Plan responses tended to be highest near the end of a trial. The proportion of Plan responses made tended to be lower than 0.5, except for Pigeons 31 and 32 who often made a proportion of Plan responses above 0.5, sometimes reaching a proportion as high as 1. In the Signalled condition all pigeons made a low proportion of Plan responses, often at or below 0.1. Plan responses always occurred at the beginning of a trial, generally within the first six responses in the Signalled condition. Pigeon 34 made no Plan responses in no-food trials in the Signalled condition.

For the human participants (Figure 3.12), within the Unsignalled condition Participants 1 and 5 made no Plan responses in trials that began with an Outcome response; therefore, their data has been excluded from this figure. Within the Signalled condition only Participants 4 and 5 started trials with an Outcome response. In the Unsignalled condition Plan responses occurred at all points within a trial across participants. The proportion of Plan responses made tended to be relatively high, often above 0.5 across participants. In the Signalled condition a high proportion of Plan responses was made every second response number, followed by no Plan responses suggesting Participants 4 and 5 were switching between making Plan and Outcome responses.

We also assessed behaviour that followed a correct response (i.e., a first Plan response) to see whether learning would be different following a correct response compared to an incorrect response (Figures 3.11 and 3.12). Figures 3.13 and 3.14 show the proportion of Plan responses as a function of response number (up to 100 responses) for trials that began with a Plan response (food or point trials) for pigeons and humans, respectively. For pigeons (Figure 3.13), in the Unsignalled condition Plan responses occurred at all points within a trial across pigeons. Pigeon 32 made no trials beginning with a Plan response during the

Unsignalled condition. When Plan responses occurred the proportion of Plan responses was often high. Later response numbers tended to have a higher proportion of Plan responses in the Unsignalled condition, often as high as 1, suggesting that longer trials involved responding mainly to the Plan key. In the Signalled condition the proportion of Plan responses tended to begin high and decreased as response number increased. Pigeons 31, 32 and 33 made no Plan responses at later response numbers in the Signalled condition, and Pigeons 34 and 35 made a small proportion of Plan responses at later response numbers.

For the human participants (Figure 3.14), in the Unsignalled condition Plan responses occurred at all points within a trial and the proportion of Plan responses tended to remain high across participants. In the Unsignalled condition later response numbers tended to have a high proportion of Plan responses, often as high as 1; longer trials involved responses mostly to the Plan card. In the Signalled condition a similar pattern emerged across participants of increasing and decreasing proportion of Plan responses with every second response number. This suggests that most participants were responding to the Plan card every second response in a trial. The proportion of Plan responses in the Signalled condition was often as high as 1, particularly towards the end of a trial. Participant 1 was an exception to this pattern, consistently making a low proportion of Plan responses across most response numbers in the Signalled condition.

Figure 3.1

Proportion of Correct Responses within a Trial for Each Individual Pigeon in the Unsignalled and Signalled Conditions

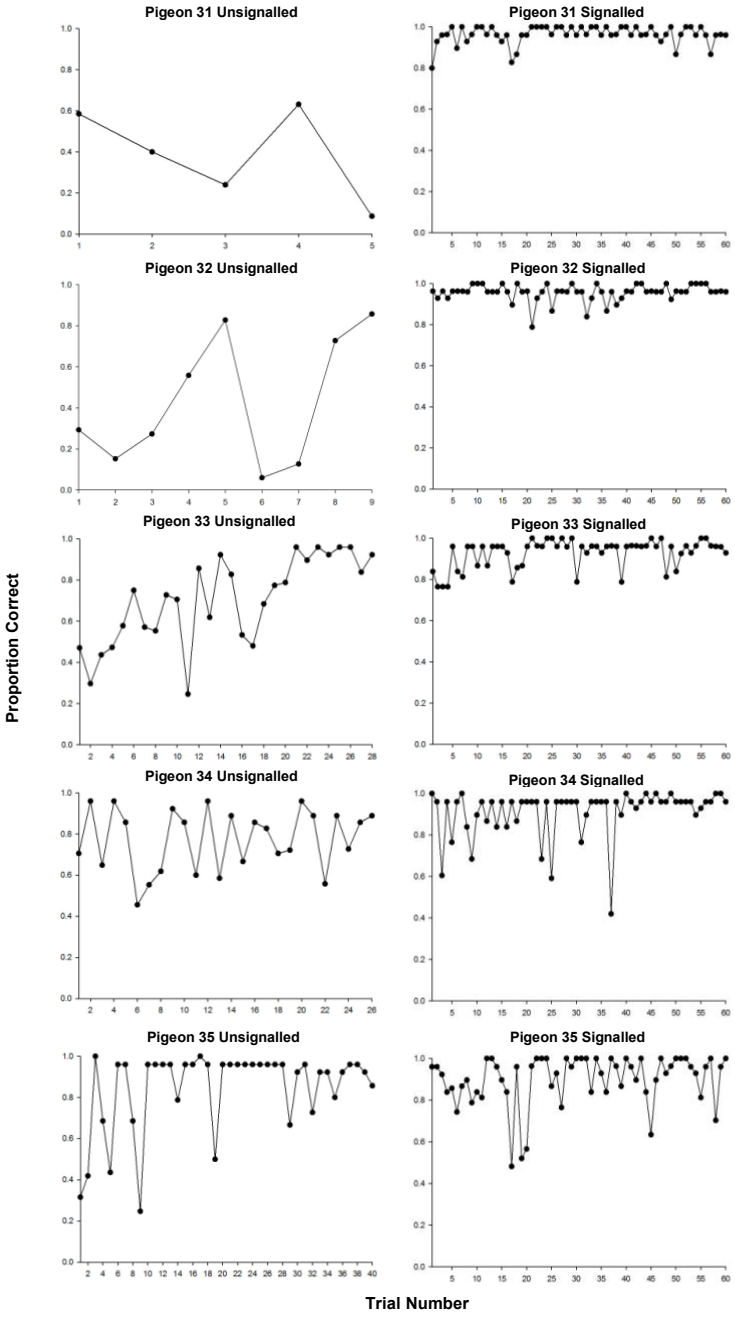

Figure 3.2

Proportion of Correct Responses within a Trial for Each Individual Human Participant in the Unsignalled and Signalled Conditions

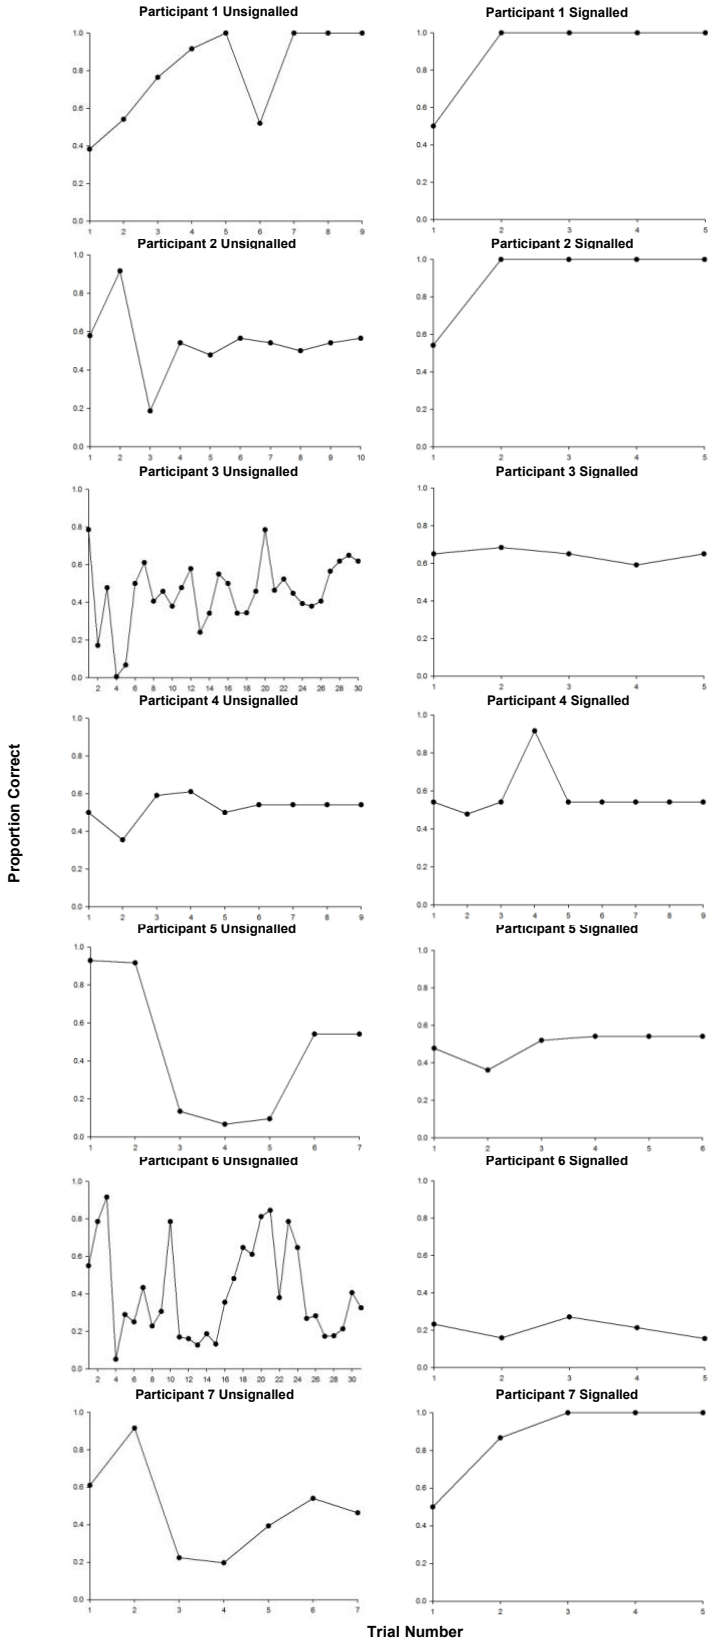

Figure 3.3

Proportion of First Responses to the Plan Key Across a Condition for Each Individual Pigeon in the Unsignalled and Signalled Conditions

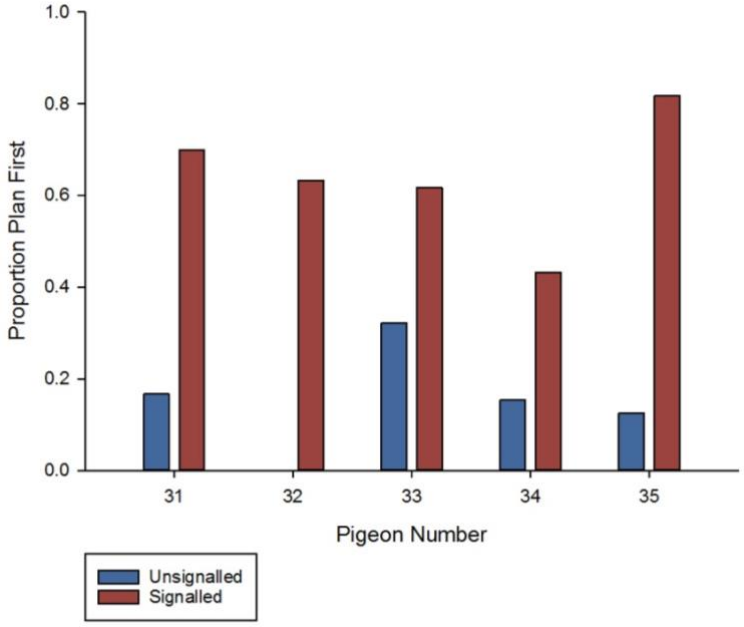

Figure 3.4

Proportion of First Responses to the Plan Key Across a Condition for Each Individual Human Participant in the Unsignalled and Signalled Conditions

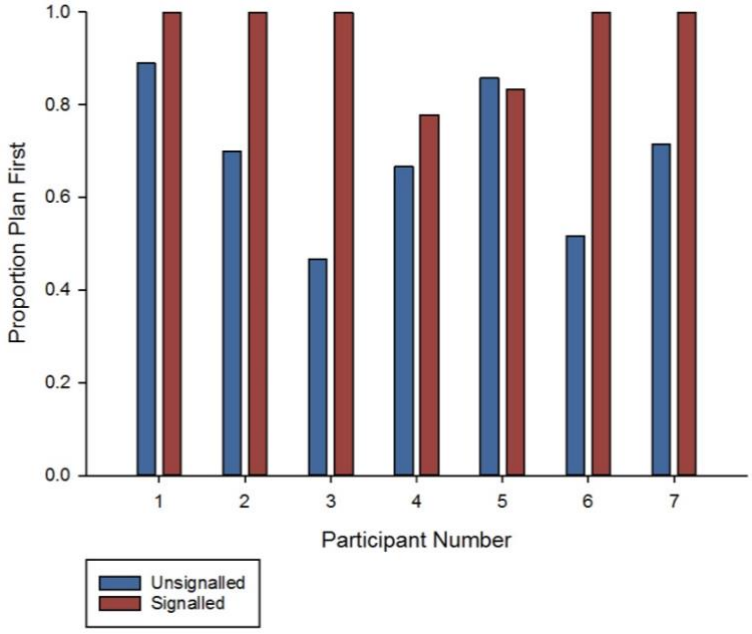

Figure 3.5

Proportion of Plan Responses Within a Trial After the First Response in a Trial for Each Individual Pigeon in the Unsignalled and Signalled Conditions

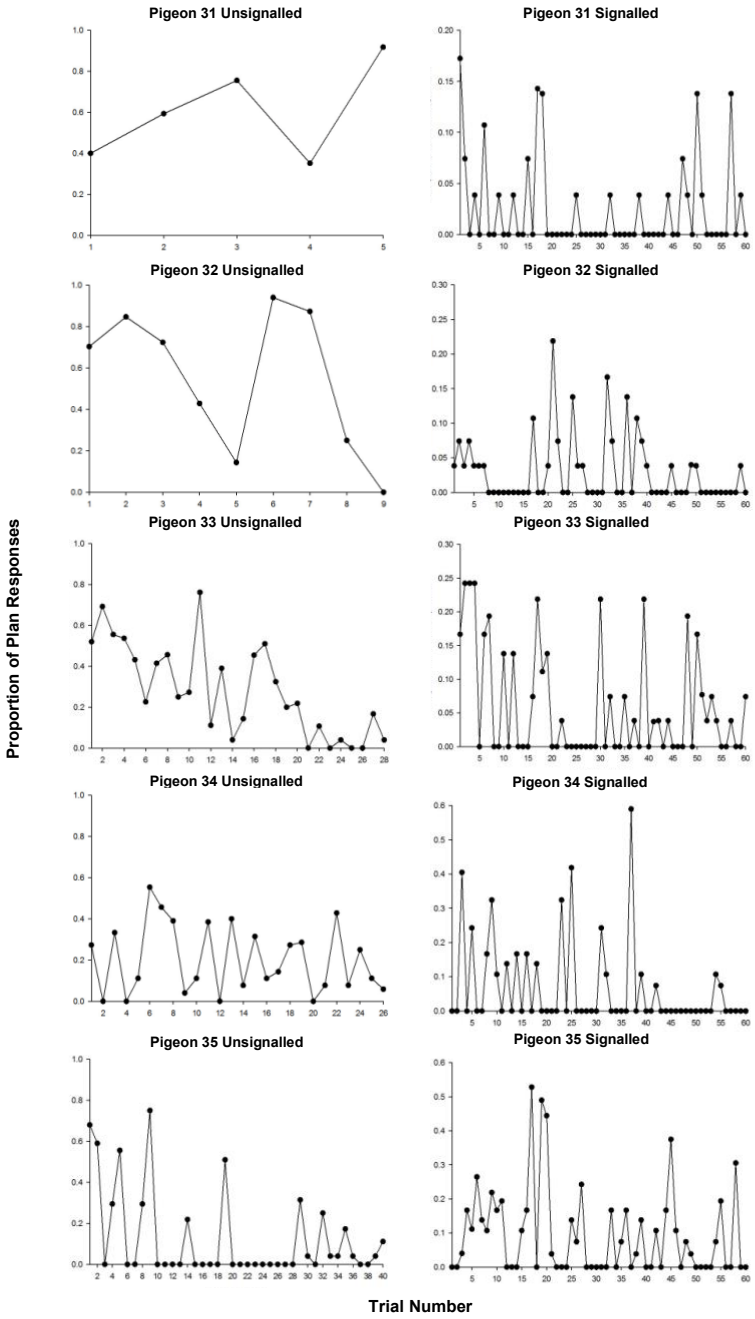

Figure 3.6

Proportion of Plan Responses Within a Trial After the First Response in a Trial for Each Individual Human Participant in the Unsignalled and Signalled Conditions

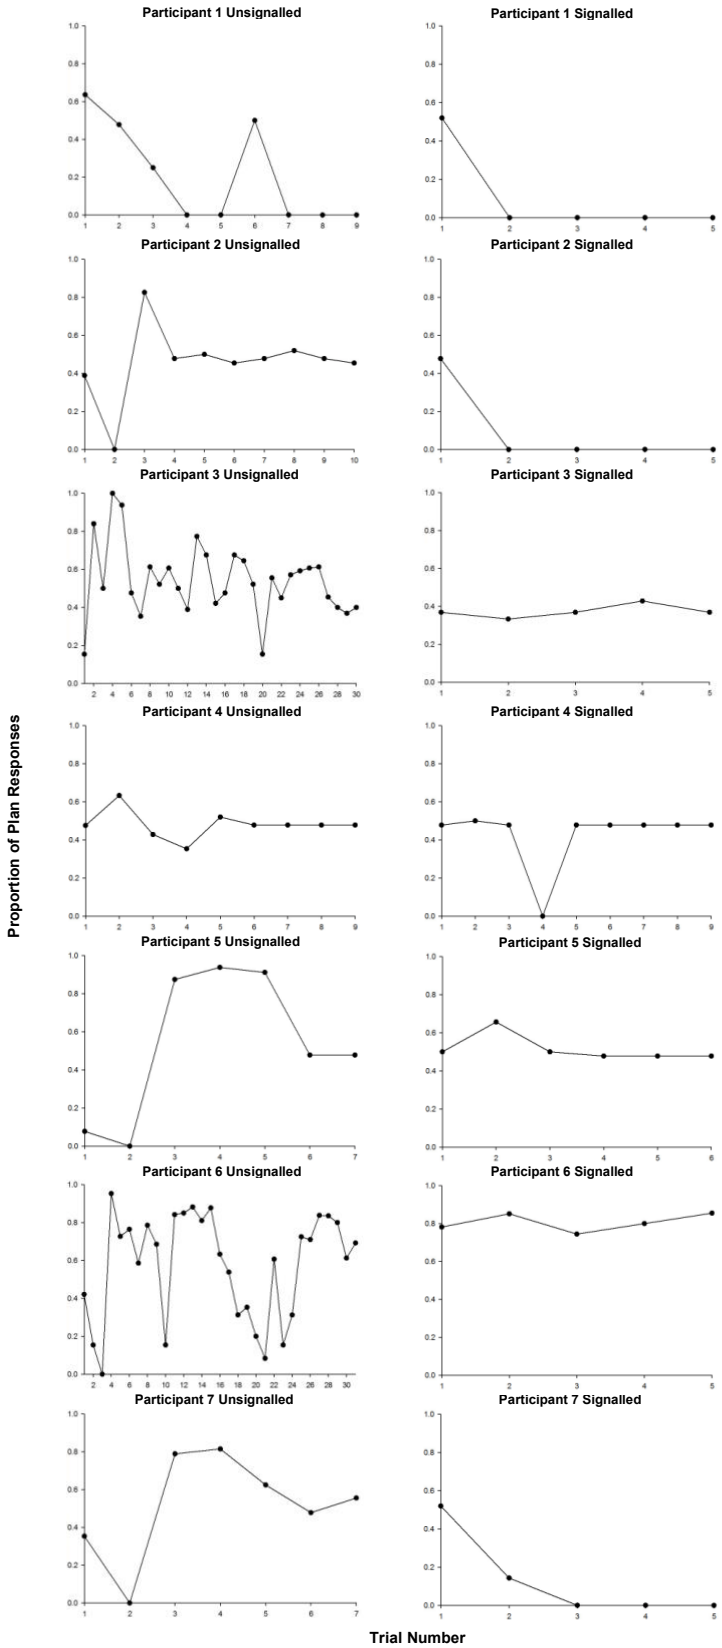

Figure 3.7

Total Number of Responses Within a Trial for Each Individual Pigeon in the Unsignalled and Signalled Conditions

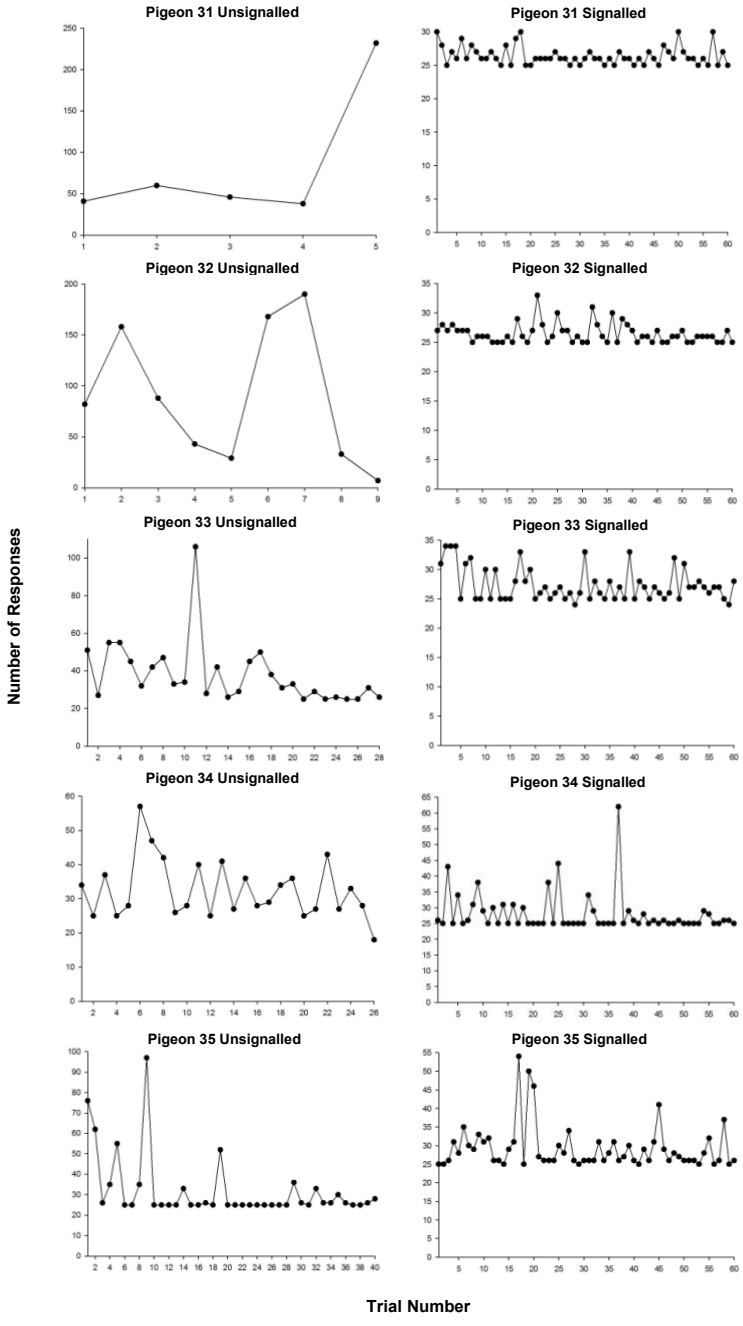

Figure 3.8

Total Number of Responses Within a Trial for Each Individual Human Participant in the Unsignalled and Signalled Conditions

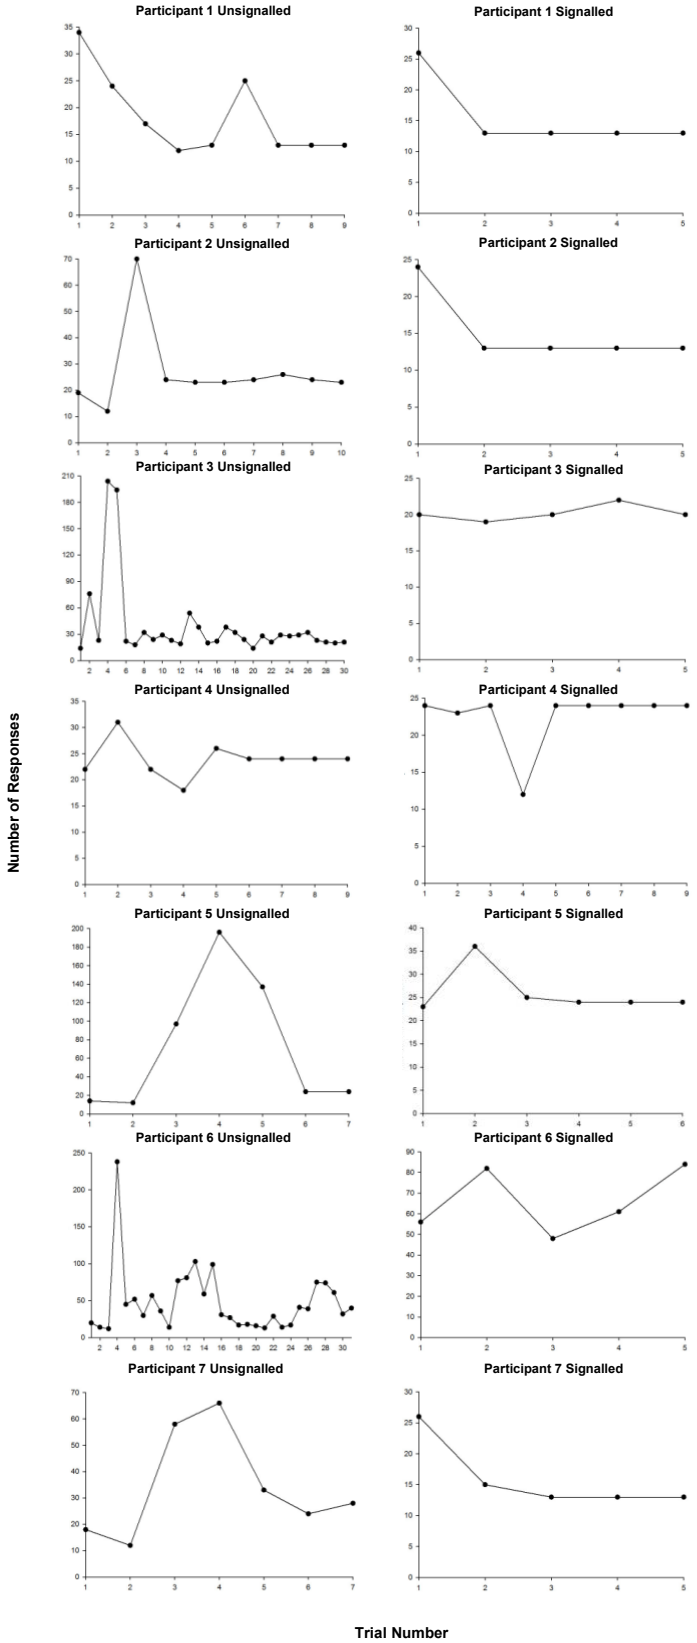

Figure 3.9

Time (in Seconds) Each Trial Took to Complete for Each Individual Pigeon in the Unsignalled and Signalled Conditions

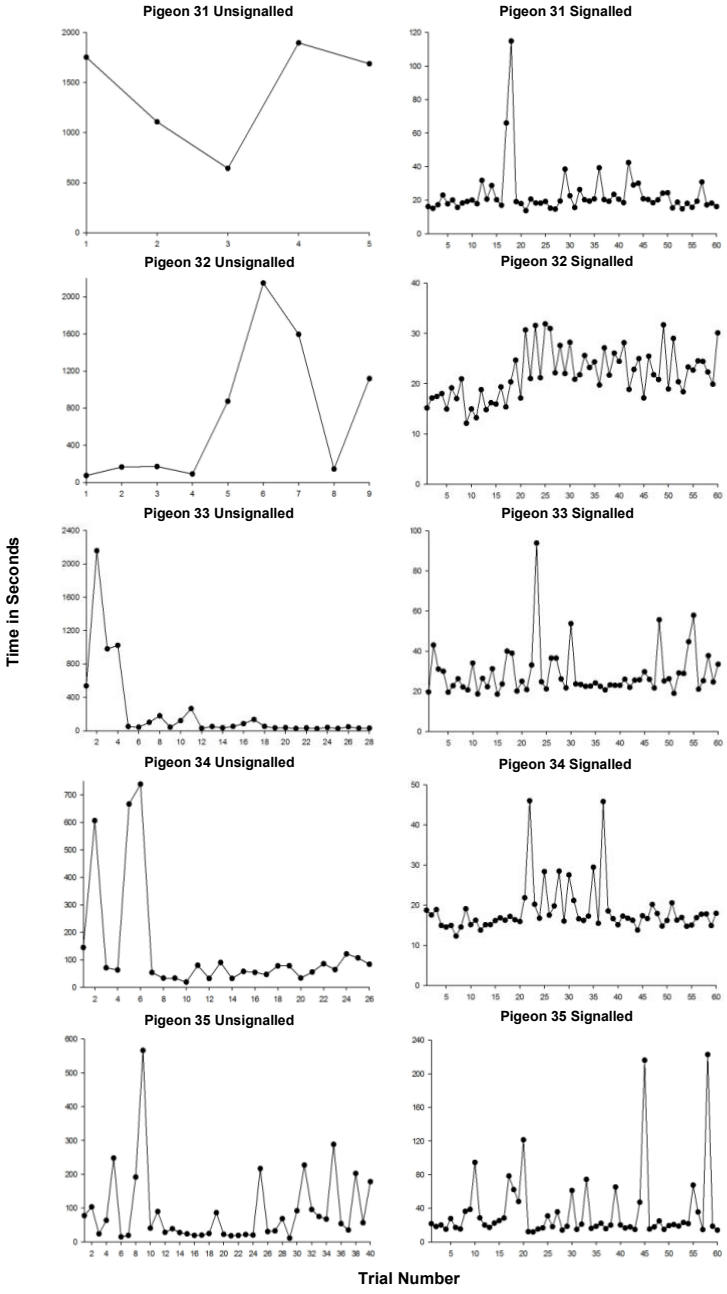

Figure 3.10

Time (in Seconds) Each Trial Took to Complete for Each Individual Human Participant in the Unsignalled and Signalled Conditions

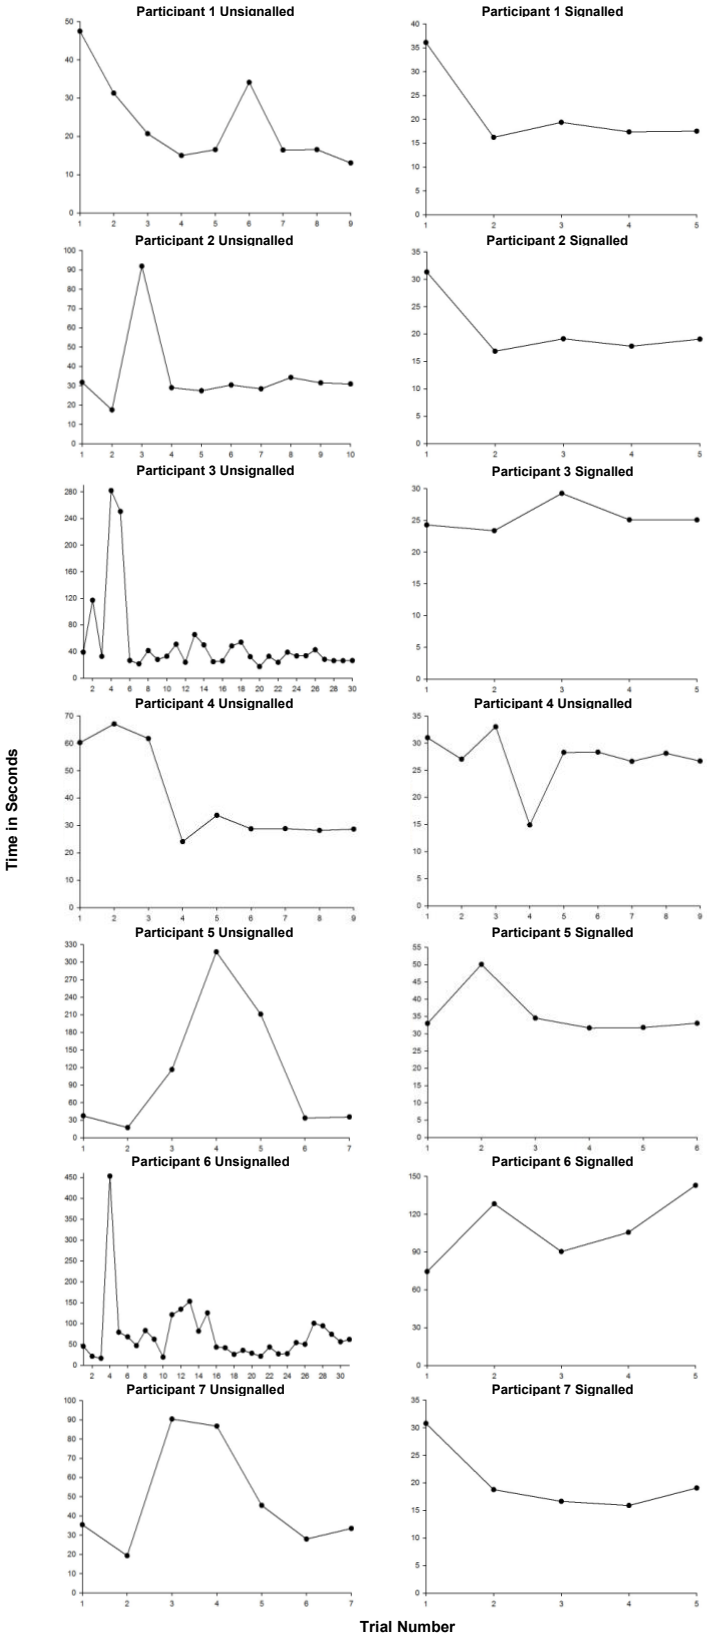

Figure 3.11

Proportion of Plan Responses as a Function of Response Number Across Trials Beginning with an Outcome Response for Each Individual Pigeon in the Unsignalled and Signalled Conditions

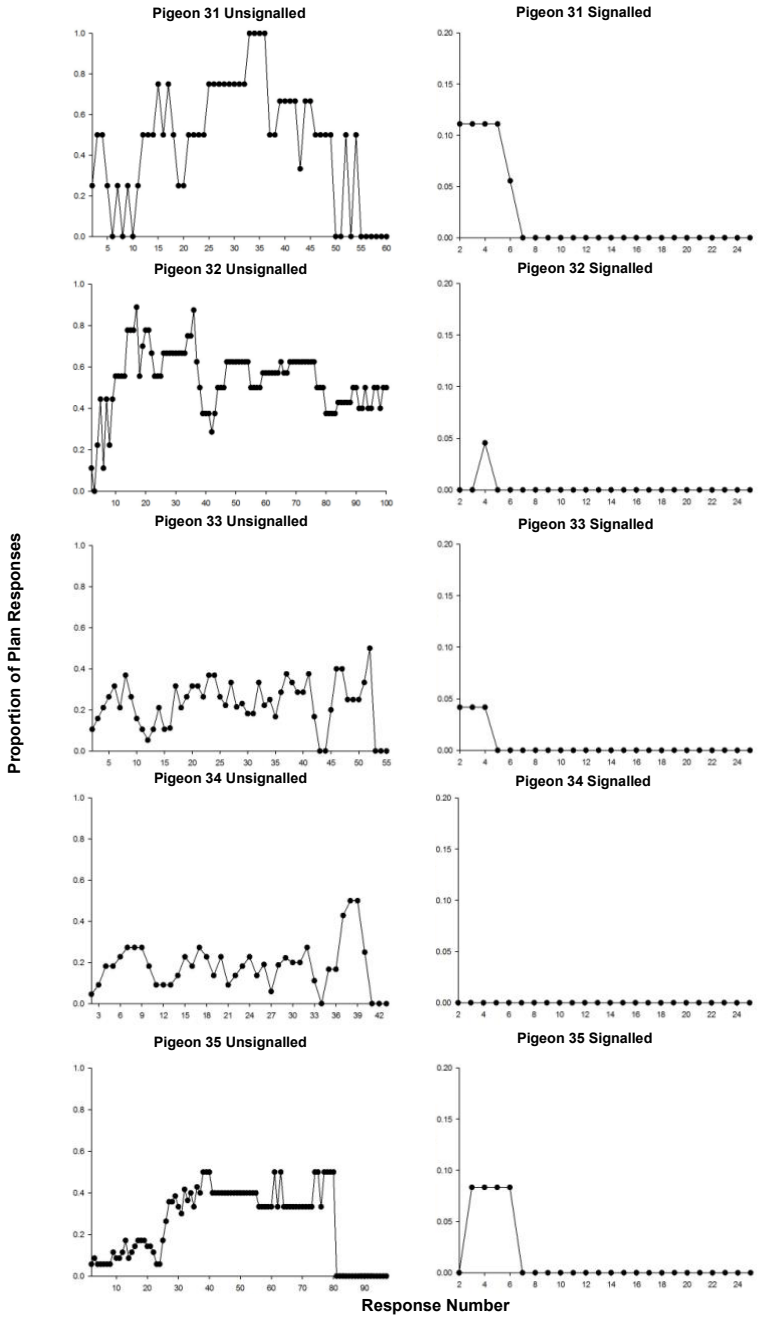

Figure 3.12

Proportion of Plan Responses as a Function of Response Number Across Trials Beginning with an Outcome Response for Each Individual Human Participant in the Unsignalled and Signalled Conditions

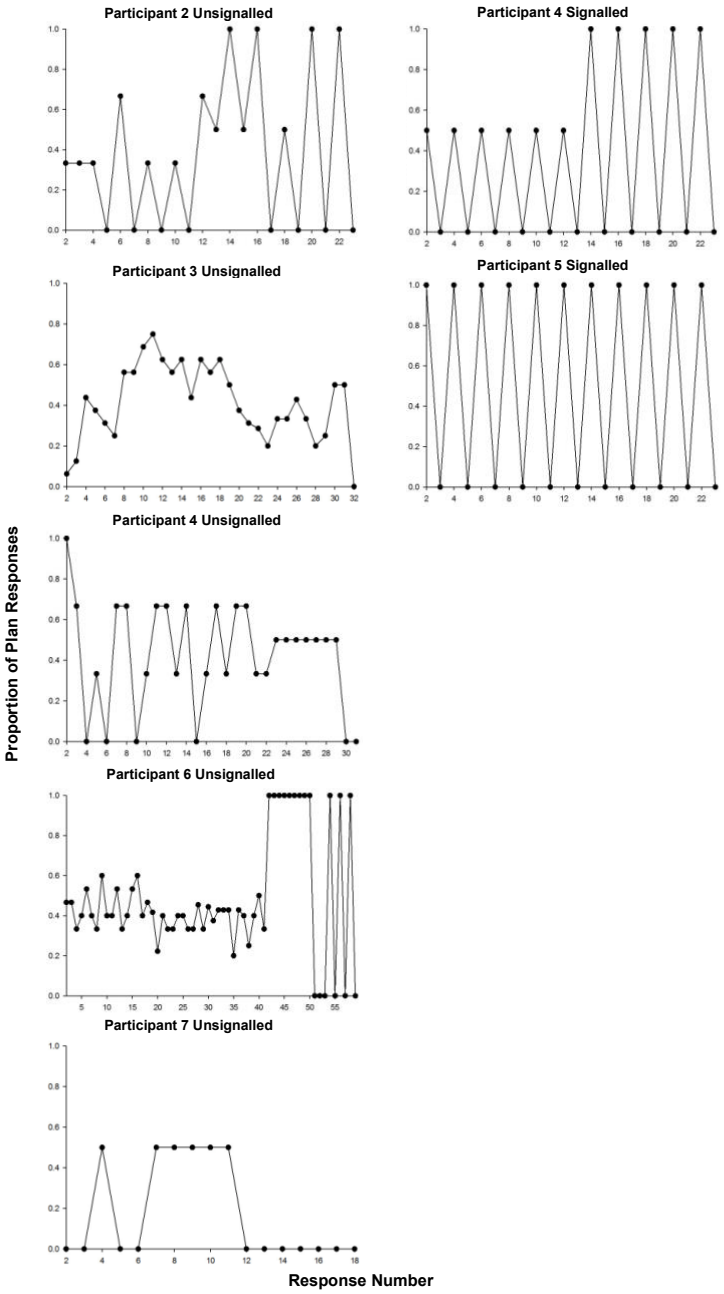

Figure 3.13

Proportion of Plan Responses as a Function of Response Number Across Trials Beginning with a Plan Response for Each Individual Pigeon in the Unsignalled and Signalled Conditions

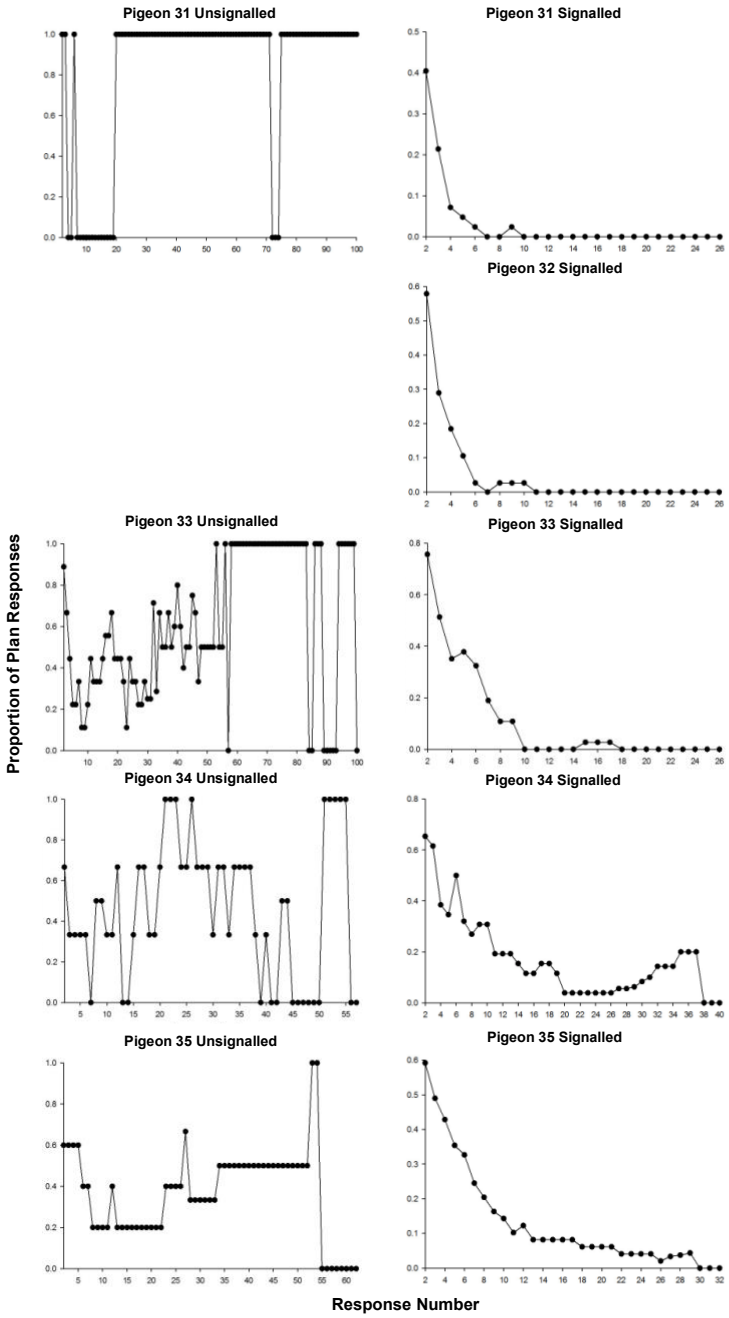

Figure 3.14

Proportion of Plan Responses as a Function of Response Number Across Trials Beginning with a Plan Response for Each Individual Human Participant in the Unsignalled and Signalled Conditions

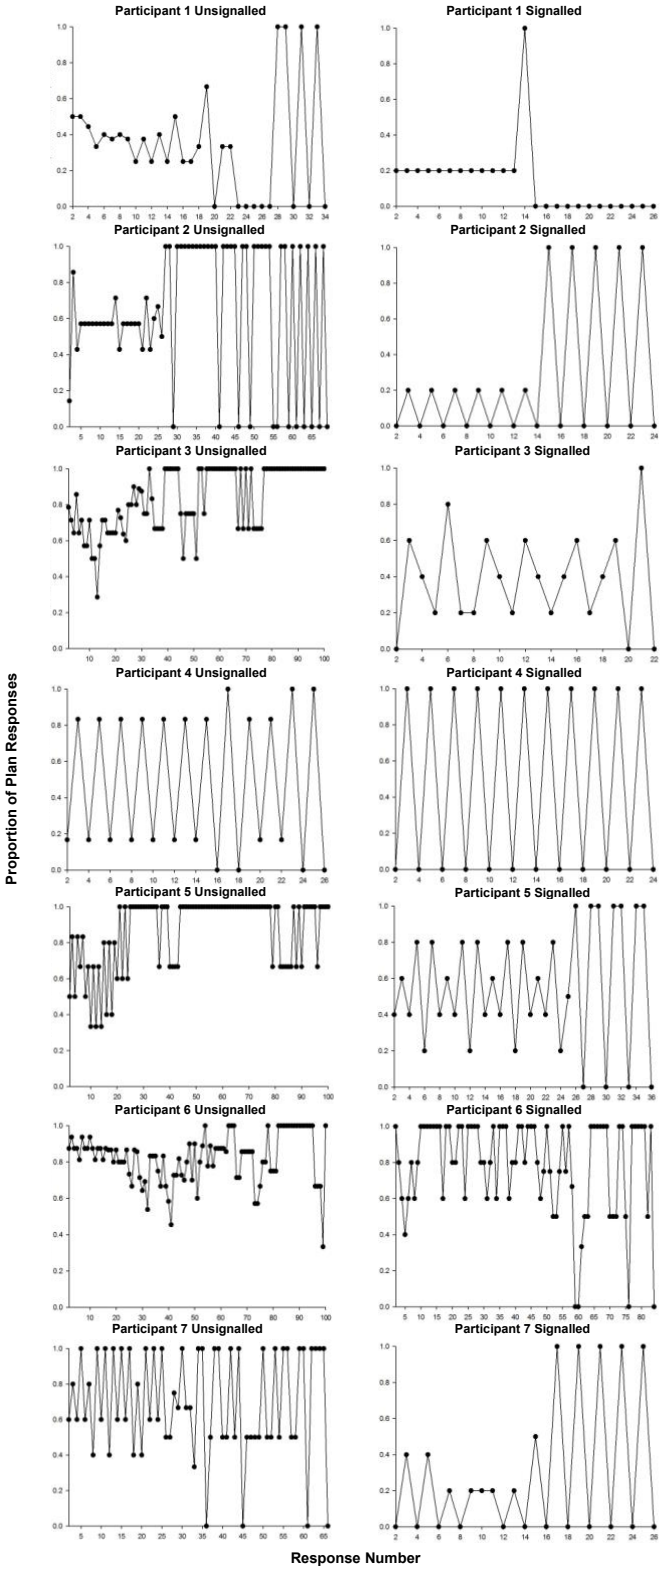

Discussion

We investigated how pigeons and humans learned to make responses with distant consequences using a task where a specific response had to be made at the beginning of a trial to receive a reinforcer at the end of the trial. This task used arbitrary stimuli and non-species-specific behaviour to examine how simple operant behaviour is controlled by potential future events. Given that pigeons can successfully learn using these types of tasks (e.g., Cowie et al., 2011; Cowie & Davison, 2021; Meyers-Manor et al., 2014) we compared their performance to the performance of humans to investigate whether learning would be similar across the two species. Overall, we found that both species demonstrated a similar *pattern* of learning and behaviour across the task. As each species gained experience with the procedure they made more correct responses, resulting in more Plan responses made first in a trial and less Plan responses made after the first response in a trial. When a correct response was signalled, the performance of both species was enhanced compared to when the same correct response was not signalled. Both species learned to begin a trial with a response that was temporally and spatially distant from the consequence.

Both pigeons and humans were able to learn to make a Plan response as the first response in a trial. Within this procedure, the Plan response was temporally separated from the reinforcer by the set number of Outcome responses, and spatially separated because the Outcome response produced the reinforcer at the end of a trial. Patterns of responding on this procedure are difficult to explain in terms of retrospective reinforcement processes (see, Cowie & Davison, 2021, for discussion). Instead, the data suggest that the behaviour of both species was under the control of likely future consequences. Animals can perform behaviours when consequences are distant (e.g., Brauer & Call, 2015; Kabadayi & Osvath, 2017; Raby et al., 2007), however these demonstrations often use complex, species-specific forms of behaviour such as preparing tools or storing food. Our results show that simple operant

responses can also be maintained by distant consequences. This could suggest that behaving in accordance with the likely future is a fundamental concept to behaviour as a whole (see also, Cowie, 2018, 2020; Shahan, 2017; Simon et al., 2020). Given that evidence was found for future-oriented behaviour in both pigeons and humans, simple operant procedures may be a useful way to investigate future-oriented behaviours such as planning in animals going forward.

Using the same procedure with both pigeons and humans allowed us to compare 'signatures' of learning and performance across the two species. When learning to Plan as the first response in a trial both species showed a similar pattern of learning. More Plan responses were made as the first response in a trial in the Signalled condition compared to the Unsignalled condition for both species (Figures 3.3 and 3.4). The human participants, however, had a higher proportion of first Plan responses in both conditions than the pigeons. Due to the smaller number of Outcome responses required for the humans compared to the pigeons, the human participants may have received reinforcers more often within a session leading to faster learning of making a Plan response as the first response in a trial.

Different patterns of responding emerged for the pigeons and humans when learning to choose the Outcome response after the first response in a trial. For pigeons, the proportion of Plan responses made after the first response in a trial tended to decrease (Figure 3.5) and the proportion of correct responses made within a trial tended to increase (Figure 3.1) as they gained experience with the procedure. Plan responses made after the first response in a trial tended to occur at all points within a trial in the Unsignalled condition for both food and no-food trials, but only at the beginning of a trial in the Signalled condition for both food and no-food trials (Figures 3.11 and 3.13). The human participants, however, seemed to develop a pattern of making alternate Plan and Outcome responses after the first response in a trial. This pattern of responding was seen across most participants in both point and no-point trials and

across both conditions (Figures 3.12 and 3.14). This may reflect a form of superstitious behaviour in our human participants (see Skinner, 1948). In humans, superstitious behaviour is often seen when patterns of responding that are thought to be causing or maintaining reinforcers are developed, when in fact these response patterns are not the cause of a reinforcer (see Ono, 1987; Rudski, 2001). This can be seen in most of our participants as they maintained a pattern of alternating Plan and Outcome responses despite this pattern not being the cause of a reinforcer at the end of a trial. The proportion of Plan responses made after the first response in a trial (Figure 3.6) and the overall proportion of correct responses in a trial (Figure 3.2) tended to remain stable across the experiment for the human participants as they stuck with their specific pattern of responding. While this pattern was seen across most participants, some (Participants 1, 2 and 7) did learn to make no unnecessary Plan responses by the end of the experiment (Figure 3.6). These slightly different patterns of learning seem to suggest that the pigeons were better at learning to respond efficiently with experience, whereas the humans learned a response pattern that resulted in a reinforcer and stuck with this throughout the experiment.

When first introduced to the procedure, both species made few first Plan responses and took longer to complete trials. As they gained experience with the procedure the number of first Plan responses increased (Figures 3.3 and 3.4), the proportion of subsequent Plan responses decreased (Figures 3.5 and 3.6) and trials took less time to complete (Figures 3.9 and 3.10), resulting in more trials ending in reinforcers. Both species learned to perform better with experience, generally making more correct responses later in a session compared to earlier in a session. This brings into question some of the criteria that is often set out to demonstrate planning abilities in animals. It has been suggested that to demonstrate genuine future planning, animals should be able to complete the task they have been set using single trials and novel problems (Suddendorf & Corballis, 2010). Our results, however,

demonstrated that even the human participants, whom we know can plan, took longer to learn than a single trial. While our human participants did learn faster than the pigeons, many were still learning by the end of the experiment, often continuing to make unnecessary responses. The use of single trials, or novel problems, within this experiment would have been too high of a bar even for our human participants to meet, suggesting animals may also struggle to meet this criterion. If single trials were used or novel problems were introduced, it would have been unlikely that either species would have performed appropriately.

Learning was still ongoing by the end of the experiment for both the pigeons and the human participants. By the end of the Signalled condition both the pigeons and the humans were still making unnecessary Plan responses after the first response in a trial (Figures 3.5 and 3.6), and some trials were still beginning with Outcome responses (Figures 3.3 and 3.4). This suggests a similar level of learning across the two species. While both the pigeons and humans improved on their performance relative to the beginning of the experiment, neither had managed to master the contingency by the end of the experiment. This suggests that the procedure used was equally challenging for both species. Given a similar amount of exposure to the procedure, both species were relatively similar in their performance and learned a similar amount about the task.

While learning was apparent in both conditions, we found that learning and performance were enhanced when a correct first Plan response was signalled. In the Signalled condition performance for both the pigeons and the humans was more consistent and fewer incorrect responses were made compared to the Unsignalled condition. The introduction of a signal between a response and a subsequent reinforcer has been shown to enhance behaviour (Richards, 1981). It is likely that this signal served to make remembering whether a correct Plan response had been made at the beginning of a trial easier. The effects of the signal on behaviour and memory could be explained through either a marking or a bridging hypothesis.

The marking hypothesis suggests that a salient stimulus occurring at the time of behaviour serves to mark the behaviour so that it is in memory when the consequence occurs (Lieberman et al., 1979). The bridging hypothesis suggests that when a stimulus is presented in the gap between behaviour and a reinforcer it bridges the temporal gap between the two (Kaplan & Hearst, 1982). Both hypothesises suggest that the signal functioned to help the pigeons and humans remember what response had been made and in turn enhanced performance.

It is also possible that the increased learning and performance in the Signalled condition may have been due to more exposure to the procedure. Both species experienced the Unsignalled condition first followed by the Signalled condition, which could explain why performance was better in the Signalled condition. While experience likely contributed to better learning, it does seem that the signal had an effect separate from experience. For the pigeons, the location of the Plan and Outcome keys reversed with each new condition. If learning was following experience alone, we may have expected worse performance for the pigeons if they only used experience from the last condition and simply kept doing what had previously been correct. Instead, we saw better learning and performance in this condition compared to the Unsignalled condition, despite the pigeons having to learn something new. Subsequent Signalled and Unsignalled conditions in Cowie and Davison (2021) showed that there was no improvement in performance in the Unsignalled condition across nine separate conditions. This clearly shows that the signal was having an effect on learning and performance separate from experience.

While the procedure used for both species was conceptually similar, the human participants had to make about half the number of Outcome responses (12) as the pigeons (25) potentially resulting in an easier task. Human participants often have lower motivation and limited attention in experiments in comparison to lab animals. Therefore, when human

participants are used in operant research rich schedules of reinforcement tend to be used (see Klapes et al., 2020; Podlesnick et al., 2022). Therefore, the reduction in Outcome responses for the humans would have resulted in a 'richer' schedule of reinforcement to ensure engagement throughout the procedure. Despite the difference in response number between the two species, by the end of the experiment both pigeons and humans were taking a similar amount of time to complete a trial (about 30 s). Therefore, by reducing the number of Outcome responses required for the humans we ensured a similar level of engagement with the task across both species.

Based on our findings, future research in this area could focus on using simple operant behaviour to further investigate the planning abilities of animals. Operant behaviour procedures tend to use arbitrary stimuli and non-species-specific behaviours, allowing us to investigate species-general behavioural processes. Given future-oriented behavioural control has been observed in a range of different species (e.g., apes (Brauer & Call, 2015), crows (Boeckle et al., 2020), pigeons (Cowie & Davison, 2021)), questions about future-oriented behaviour may be better tested with species-general procedures than with species-specific ones. Indeed, the use of such approaches may allow different components of future-oriented behaviour to be isolated and explored – as was done in the present procedure which investigated the ability of temporally and spatially distant events to control behaviour.

A further benefit of species-general procedures is that they permit comparisons of 'signatures' of learning and performance across different species. Overall, we found similar patterns of learning and behaviour across pigeons and humans using the same procedure, but a greater tendency for superstitious patterns of responding in humans. The behaviour of both species came under the control of likely future consequences and both species learned to make a response with a temporally and spatially distant consequence. Similarities in learning and behaviour between the two species suggest that this procedure was successful in

measuring the ability of both pigeons and humans to make responses with distant outcomes, and that both species learned in a similar way to navigate the procedure.

General Discussion

We assessed whether future-oriented behaviour could be studied in humans using arbitrary stimuli and responses. We did this using two experimental paradigms in which participants had to choose stimuli that would impact some future outcome, using stimuli that varied in their arbitrariness for both humans and pigeons. In Experiments 1 and 2 we used a tool-selection procedure that had previously been used to study planning in animals, however we made the stimuli and behaviours more arbitrary and used human participants. Experiment 1 showed that our participants were successful in learning with arbitrary stimuli and responses during Training, however this learning generally failed to transfer to Testing where participants performed worse than crows (Boeckle et al., 2020) and children (Miller et al., 2020). Experiment 2 showed that the poor performance during Testing where the low-value stimulus was the most common item chosen was not due to a procedural artefact. In Experiment 3 we explored similarities in how humans and pigeons learned to choose between arbitrary stimuli when consequences were distant, using a card game. We found similar patterns of learning across both species, with both taking time to learn the task, and both having performance facilitated by immediate stimulus changes. Thus, as a general conclusion, the sorts of arbitrary stimuli and responses used in laboratory tasks with 'simple' animals are a potentially valuable tool for species-general exploration of the ability to plan and learn about temporally extended relations between events. The use of arbitrary stimuli and arbitrary responses does not, in and of itself, limit one's ability to respond in accordance with likely future events.

Nevertheless, we found differing levels of behavioural control by the likely future in the Tool Experiment (Chapters 1 and 2) and Card Game Experiment (Chapter 3). In the Tool

Experiment, although participants completed the Training successfully, in a manner that mirrored crows' performance on the same sort of task (Boeckle et al., 2020), participants failed to choose the correct tool during Testing. In the Card Game Experiment, we found more convincing evidence of control by the likely future as the participants leant to make a response at the beginning of a trial to receive a reinforcer at the end of a trial. The two experiments were similar in the responses that participants had to make, and in the reinforcers for correct responses. Each experiment involved the participant clicking stimuli on a screen to try and work out how to earn points. Thus, the arbitrariness of the responses used is unlikely to have been a factor contributing *directly* to poor performance in the Tool Experiment.

Motivation may have been a factor affecting the differing levels of prospective control by our human participants. In both the Tool Experiment and the Card Game Experiment the participants worked to earn points for correct responses. In comparison to the crows (Boeckle et al., 2020) working for food rewards and the children working for a preferred sticker (Miller et al., 2020), points may be considered a less valuable reward, potentially reducing the motivation of our participants. While earning more points did increase the chance of the participants receiving a monetary reward, it did not guarantee this. Despite a possible difference in motivation levels, we did find evidence of learning across both experiments by our participants suggesting that lower motivation did not prevent them from learning throughout the experiments.

One key difference between the two experiments was the arbitrariness of the task participants were completing. In the Tool Experiment, participants were simply clicking objects on a screen and trying to work out how to earn points. Even when the stimuli were not arbitrary (Group Tool), the response to the stimuli was (rarely does one click on a hammer to make it work). In the Card Game Experiment, the participants were told they were playing a card game against an experimenter and had to earn points to try and beat the

experimenter. Although the cards themselves were comprised of shapes rather than more traditional suit-and-number combinations, the task of playing cards on a computer is less arbitrary as many individuals are likely to have had experience with playing different types of card games. Compared to clicking objects on a screen, clicking a 'card' in order to win a card game may have had more meaning – in that the combination of stimuli and behaviours related more closely to those in existing schemas – to the participants in this experiment. The meaningfulness of the task may explain why we found better evidence for future oriented behaviour in the Card Game Experiment compared to the Tool Experiment. With an existing schema, such as a card game, participants were likely to have had previous experience with its structure and function, allowing them to use existing knowledge to make sense of the unfamiliar task. With the colour version of the Tool experiment however, participants were unlikely to have had previous experience with clicking coloured circles (or pictures of tools in the tool version) on a screen, leading to the task seeming more random or unrelated to existing schemas. Therefore, it seems that existing general knowledge about a task and the meaningfulness of that task assisted in the ability for humans to be able to plan. This brings into question how we would use the same procedure to investigate planning with species-general behaviour and non-specific stimuli. Given that humans failed when the overall task was less familiar to them, arguments that animal behaviour does not demonstrate planning abilities because of familiarity with a task (e.g., Redshaw et al., 2017; Suddendorf & Corballis, 2010) seem redundant.

These findings raise the importance of biological constraints on learning, particularly when planning tasks involve the use of species-specific experience compared to species-general experience. Biological constraints can influence learning where combinations that are more relevant to an organism are more easily learned. For example, taste-aversion learning requires a large delay between taste and subsequent sickness and is better learned when the

association is between taste and sickness than taste and shock (Garcia & Koelling, 1966; Domjan, 2015). A long delay between food and sickness is biologically relevant as digestion takes time, therefore a short delay between taste and sickness would not be biologically appropriate, leading to better taste-aversion learning with a long delay. Species-specific schemas also play a role in learning as things that fit into one's species-specific schemas are more rapidly and easily learned. Dogs learning to go left or go right learn faster when the location of a tone is used than when different types of tone are used (Lawicka, 1968). Omnivorous noisy miner birds learn win-shift strategies better when a food reinforcer that depletes (nectar) is used compared to a food reinforcer that does not deplete (mealworms) (Sulikowski & Burke, 2007). These results were likely seen due to their biological relevance to the organism. The location of a noise is likely to be more relevant to a dog if they are out hunting for food in comparison to the type of noise, making location discriminations easier to learn. Likewise, noisy miners consume both nectar and mealworms as part of their diet making win-shift strategies easier to learn when one of these specific foods is used. Additionally, the relation between behaviours and subsequent reinforcers are important in the same way. As Sulikowski and Burke (2007) demonstrate, learning can be influenced by the type of reinforcers that are used particularly when these are relevant to the species. This suggests that the arbitrary stimuli that must necessarily be used with animals may make the task harder to learn.

Although arbitrary stimuli may make the task more difficult to learn, the results of both the Tool and Card Experiments (Chapters 1, 2 and 3) suggest that arbitrary stimuli and responses in and of themselves do not preclude successful planning behaviours. Certainly, the participants' failure to select the correct stimuli in the Testing Levels of the Tool Experiment (Chapters 1 and 2) was not a direct result of the use of arbitrary stimuli. This has important implications for research with non-human animals, allowing for further research with non-

species-specific behaviours – at least provided we are careful to account for the added difficulty such arbitrary stimuli may create. Further, it demonstrates the importance of understanding past experience – both species-specific and personal history (i.e., the experience that permits experiment stimuli and behaviours to fit into our existing schemas), and experiment-specific history (i.e., the experience that allows new schemas to develop) – on the ability to engage in successful 'planning' and learning. Future work might explore further the role of such history, such as investigating how differences in experience with a non-species-specific task may affect the ability to plan for the future.

The results from both the Tool Experiment and the Card Game Experiment call into question the often-difficult criterion that is set for animals to demonstrate the ability to plan. Consider the criteria set out by Tulving's (2005) famous spoon test. This test is based off a story of a young girl who dreams she is at a party with her friends eating pudding. She, however, has not brought a spoon with her and therefore cannot eat the pudding. When she goes to sleep the next night, she takes a spoon with her to bed. It is argued that this story demonstrates future planning as the girl has learned from a single experience and has mentally travelled in time to determine when she will need the spoon in the future. The spoon test often dictates the criteria required to demonstrate future planning and foresight in animals. If we consider our results in both the Card Game Experiment and the Tool Experiment in relation to this criterion, our human participants failed to meet the test as none of the participants learned to plan after a single experience. Further, in the Tool Experiment many did not select the correct tool in Testing even after *multiple* exposures. Much of the criteria set out by Suddendorf and Corballis (2010) to demonstrate planning in animals was also difficult for our human participants to meet. Participants were not able to learn using single trials, and the introduction of a novel problem during Testing in the Tool Experiment resulted in poor performance. Given that humans, who we already know have the ability to

plan, could not meet these criteria, it might suggest that the bar for what constitutes planning in non-human animals has been set too high. If humans performing very basic behaviours are unable to meet these criteria, then how can we expect that animals will be able to? Demonstrations of planning in animals can often be dismissed in the defence of human uniqueness (Leslie, 2018). Setting the criteria of planning in animals to a level that is considered appropriate for humans may mean that we will not be able to properly understand what animals are capable of. In fact, in light of our human participants' performance, it seems that even humans cannot reach these strict criteria, suggesting animals would also be unable to.

Our findings highlight the influence that experience and training seem to have on the ability to plan for the future. Most planning studies that use animals require training to demonstrate the ability to plan. These studies often rely on a combination of innate predispositions, previous experience, and training for the specific task at hand (Leslie, 2018). Often, training data are excluded from publications, and the use of species-specific behaviours like tool use means it is impossible to control for prior 'training' (relevant experience) in the animal's learning history outside the experiment. For example, the crows used by Boeckle et al. (2020) would likely have had previous experience with using tools in the wild, as well as being trained to use the specific tools required in the task. Much of the criteria set out to demonstrate planning abilities in animals does not give weight to this experience. If we consider Tulving's (2005) spoon test and Suddendorf and Corballis's (2010) criteria, both Boeckle et al.'s crows and Miller et al.'s (2020) children would be considered to have failed to demonstrate an ability to plan as training and multiple trials were used for both. Our human participants also showed better evidence of prospective behaviour when they had experience with the task and were more familiar with the task overall (e.g., the Card Experiment relative to the Tool Experiment). Given that humans required experience to

plan this may suggest that experience and training are important in the ability to plan for the future. Without knowledge of the structure of the environment or what is likely to occur in the future, the ability to plan for the future may be hampered. This highlights the importance of using species-general behaviour and arbitrary stimuli as training and experience can be provided for tasks that implement these (as seen with our human participants) allowing for experience with the structure of the environment without relying on innate predispositions or species-specific behaviour.

Overall, we found that the use of arbitrary stimuli and responses does not limit the ability of human participants to respond in accordance with the likely future. Planning for the future seemed to be enhanced when the task more closely mirrored tasks encountered outside the experiment, hence longer-term experience could be used, suggesting that experience may be an important factor that influences prospective behaviour. This has implications for the way planning is studied in animals, suggesting that experience and training should be considered when demonstrating the ability for animals to plan. The use of non-species-specific behaviours and arbitrary stimuli in planning tasks provide a promising avenue to demonstrate prospective behaviour and planning in animals.

Appendix A

Table 1

One-way Repeated Measures ANOVA of the Number of Trials Completed During Training Levels for Group Colour

Level Number		Mean Difference	Std. Error	p-value
Level 1	Level 2	-28.643*	7.863	.003
	Level 3	-9.786	9.814	.337
	Level 4	1.000	7.277	.893
	Level 5	2.500	7.136	.732
	Level 6	-15.643	11.335	.191
Level 2	Level 3	18.857*	5.904	.007
	Level 4	29.643*	6.940	<.001
	Level 5	31.143*	7.031	<.001
	Level 6	13.000	7.942	.126
Level 3	Level 4	10.786	6.202	.106
	Level 5	12.286	6.595	.085
	Level 6	-5.857	5.255	.285
Level 4	Level 5	1.500	1.123	.205
	Level 6	-16.643*	7.551	.046
Level 5	Level 6	-18.143*	8.345	.049

Table 2

One-way Repeated Measures ANOVA of the Number of Trials Completed During Training Levels for Group Tool

Level Number		Mean Difference	Std. Error	p-value
Level 1	Level 2	-26.100	14.602	.108
	Level 3	-11.600	11.037	.321
	Level 4	8.100*	2.410	.008
	Level 5	11.400*	2.960	.004
	Level 6	11.100*	3.407	.010
Level 2	Level 3	14.500	9.719	.170
	Level 4	34.200	15.250	.052
	Level 5	37.500*	15.128	.035
	Level 6	37.200*	15.228	.037
Level 3	Level 4	19.700	12.423	.147
	Level 5	23.000	12.537	.100
	Level 6	22.700	12.856	.111
Level 4	Level 5	3.300*	1.375	.040
	Level 6	3.000	1.673	.107
Level 5	Level 6	-.300	.539	.591

Table 3

Mixed ANOVA of the Number of Trials Completed During Training Levels for Group Colour vs Group Tool

Level Number		Mean Difference	Std. Error	p-value
Level 1 (Group Colour)	Level 1 (Group Tool)	.643	4.543	.889
Level 2 (Group Colour)	Level 2 (Group Tool)	3.186	16.110	.845
Level 3 (Group Colour)	Level 3 (Group Tool)	-1.171	15.370	.940
Level 4 (Group Colour)	Level 4 (Group Tool)	7.743	8.191	.355
Level 5 (Group Colour)	Level 5 (Group Tool)	9.543	8.032	.247
Level 6 (Group Colour)	Level 6 (Group Tool)	27.386*	12.429	.038

Table 4

One-way Repeated Measures ANOVA of the Proportion of Correct Trials Completed During Training Levels for Group Colour

Level Number		Mean Difference	Std. Error	p-value
Level 1	Level 2	.285*	.063	<.001
	Level 3	-.026	.104	.803
	Level 4	-.159	.104	.152
	Level 5	-.228*	.080	.014
	Level 6	-.020	.108	.856
Level 2	Level 3	-.311*	.090	.004
	Level 4	-.444*	.069	<.001
	Level 5	-.513*	.060	<.001
	Level 6	-.305*	.076	.001
Level 3	Level 4	-.132	.085	.144
	Level 5	-.201*	.076	.020
	Level 6	.006	.083	.939
Level 4	Level 5	-.069	.055	.227
	Level 6	.139*	.043	.007
Level 5	Level 6	.208*	.064	.006

Table 5

One-way Repeated Measures ANOVA of the Proportion of Correct Trials Completed During Training Levels for Group Tool

Level Number		Mean Difference	Std. Error	p-value
Level 1	Level 2	-.110	.173	.542
	Level 3	-.239	.145	.133
	Level 4	-.305*	.103	.016
	Level 5	-.432*	.077	<.001
	Level 6	-.464*	.092	<.001
Level 2	Level 3	-.129	.095	.208
	Level 4	-.195	.172	.285
	Level 5	-.322*	.129	.034
	Level 6	-.354*	.145	.038
Level 3	Level 4	-.066	.149	.669
	Level 5	-.193	.103	.095
	Level 6	-.225	.141	.145
Level 4	Level 5	-.127	.085	.171
	Level 6	-.159	.078	.071
Level 5	Level 6	.032	.054	.570

Table 6

Mixed ANOVA of the Proportion of Correct Trials Completed During Training Levels for Group Colour vs Group Tool

Level Number		Mean Difference	Std. Error	p-value
Level 1 (Group Colour)	Level 1 (Group Tool)	.052	.080	.522
Level 2 (Group Colour)	Level 2 (Group Tool)	-.345*	.140	.022
Level 3 (Group Colour)	Level 3 (Group Tool)	-.162	.165	.339
Level 4 (Group Colour)	Level 4 (Group Tool)	-.094	.114	.417
Level 5 (Group Colour)	Level 5 (Group Tool)	-.152	.097	.130
Level 6 (Group Colour)	Level 6 (Group Tool)	-.394*	.110	.002

Table 7

One-way Repeated Measures ANOVA of the Time (in Seconds) Each Level Took to Complete During Training Levels for Group Colour

Level Number		Mean Difference	Std. Error	p-value
Level 1	Level 2	-56.040*	19.764	.014
	Level 3	-18.746	26.323	.489
	Level 4	-85.002	70.334	.248
	Level 5	14.728	15.997	.374
	Level 6	-31.766	27.011	.261
Level 2	Level 3	37.293*	13.211	.014
	Level 4	-28.963	64.016	.658
	Level 5	70.767*	18.715	.002
	Level 6	24.273	21.164	.272
Level 3	Level 4	-66.256	61.610	.302
	Level 5	33.474	21.294	.140
	Level 6	-13.020	18.196	.487
Level 4	Level 5	99.730	56.923	.103
	Level 6	53.236	59.538	.387
Level 5	Level 6	-46.494*	19.823	.036

Table 8

One-way Repeated Measures ANOVA of the Time (in Seconds) Each Level Took to Complete During Training Levels for Group Tool

Level Number		Mean Difference	Std. Error	p-value
Level 1	Level 2	-101.228	56.513	.107
	Level 3	-54.480	51.031	.314
	Level 4	-11.945	9.205	.227
	Level 5	38.207*	14.457	.027
	Level 6	34.301*	14.733	.045
Level 2	Level 3	46.749	38.981	.261
	Level 4	89.283	55.278	.141
	Level 5	139.435	62.367	.052
	Level 6	135.530	62.098	.057
Level 3	Level 4	42.535	51.441	.430
	Level 5	92.686	57.589	.142
	Level 6	88.781	57.891	.159
Level 4	Level 5	50.151*	17.545	.019
	Level 6	46.246*	17.512	.027
Level 5	Level 6	-3.905*	1.529	.031

Table 9

Mixed ANOVA of the Time (in Seconds) Each Level Took to Complete During Training Levels for Group Colour vs Group Tool

Level Number		Mean Difference	Std. Error	p-value
Level 1 (Group Colour)	Level 1 (Group Tool)	-5.044	16.795	.767
Level 2 (Group Colour)	Level 2 (Group Tool)	-48.947	57.540	.404
Level 3 (Group Colour)	Level 3 (Group Tool)	-40.777	56.296	.476
Level 4 (Group Colour)	Level 4 (Group Tool)	68.013	82.911	.421
Level 5 (Group Colour)	Level 5 (Group Tool)	18.435	13.924	.199
Level 6 (Group Colour)	Level 6 (Group Tool)	61.024*	27.411	.037

Table 10

Independent Samples t-test for the Number of Incorrect Trials Across Training Levels for Group Colour vs Group Tool

Variable	F	Levene's Test p-value	*t*	*df*	p-value
Number of Incorrect Trials	.439	.515	.731	22	.473

Table 11

One-way Repeated Measures ANOVA of the Proportion of Correct WCST Responses During Testing Levels for Group Colour

Level and Trial Number		Mean Difference	Std. Error	p-value
Level 7, Trial 1	Level 7, Trial 2	-.160*	.045	.003
	Level 8, Trial 1	-.191*	.038	<.001
	Level 8, Trial 2	-.156	.079	.071
Level 7, Trial 2	Level 8, Trial 1	-.031	.033	.374
	Level 8, Trial 2	.004	.074	.955
Level 8, Trial 1	Level 8, Trial 2	.035	.057	.551

Table 12

One-way Repeated Measures ANOVA of the Proportion of Correct WCST Responses During Testing Levels for Group Tool

Level and Trial Number		Mean Difference	Std. Error	p-value
Level 7, Trial 1	Level 7, Trial 2	-.152	.073	.068
	Level 8, Trial 1	-.258*	.070	.005
	Level 8, Trial 2	-.288*	.080	.006
Level 7, Trial 2	Level 8, Trial 1	-.106*	.039	.025
	Level 8, Trial 2	-.136*	.051	.025
Level 8, Trial 1	Level 8, Trial 2	-.030	.029	.326

Table 13

Mixed ANOVA of the Proportion of Correct WCST Responses Completed During Testing Levels for Group Colour vs Group Tool

Level Number		Mean Difference	Std. Error	p-value
Level 7, Trial 1 (Group Colour)	Level 7, Trial 1 (Group Tool)	.044	.078	.575
Level 7, Trial 2 (Group Colour)	Level 7, Trial 2 (Group Tool)	.054	.066	.422
Level 8, Trial 1 (Group Colour)	Level 8, Trial 1 (Group Tool)	-.022	.033	.517
Level 8, Trial 2 (Group Colour)	Level 8, Trial 2 (Group Tool)	-.089	.081	.280

Appendix B

Table 1

Paired Samples t-test for the Pigeon Subjects in the Unsignalled Compared to the Signalled Condition

Variable	*M*	*SD*	*t*	*df*	p-value
Proportion Plan First	-.486661	.19058	-5.709	4	.005
Proportion Plan Responses	.29681	.24024	2.763	4	.051
Proportion Correct	-.30384	.23421	-2.901	4	.044
Number of Responses	27.41460	29.58366	2.072	4	.107
Time Each Trial Took	489.62152	564.49637	1.939	4	.124

Table 2

Paired Samples t-test for the Human Participants in the Unsignalled Compared to the Signalled Condition

Variable	*M*	*SD*	*t*	*df*	p-value
Proportion Plan First	-.25733	.20491	-3.323	6	.016
Proportion Plan Responses	.12327	.21099	1.546	6	.173
Proportion Correct	-.13506	.20757	-1.722	6	.136
Number of Responses	11.44847	19.85937	1.525	6	.178
Time Each Trial Took	17.83459	32.55797	1.449	6	.197

References

Azrin, N. H., Hake, D. F., Holz, W. C., & Hutchinson, R. P. (1965). Motivational aspects of escape from punishment 1. *Journal of the Experimental Analysis of Behavior*, *8*(1), 31-44. https://doi.org/10.1901/jeab.1965.8-31

Babb, S. J., & Crystal, J. D. (2006). Episodic-like memory in the rat. *Current Biology*, *16*(13), 1317-1321. https://doi.org/10.1016/j.cub.2006.05.025

Baum, W. M. (1974). Choice in free-ranging wild pigeons. *Science*, *185*(4145), 78-79. https://doi.org/10.1126/science.185.4145.78

Boeckle, M., Schiestl, M., Frohwieser, A., Gruber, R., Miller, R., Suddendorf, T., Gray, R. D., Taylor, A. H., & Clayton, N. S. (2020). New Caledonian crows plan for specific future tool use. *Proc. R. Soc. B*, *287*(1938), 1-7. https://doi.org/10.1098/rsbp.2020.1490

Bourjade, M., Call, J., Pelé, M., Maumy, M., & Dufour, V. (2014). Bonobos and orangutans, but not chimpanzees, flexibly plan for the future in a token-exchange task. *Animal Cognition*, *17*(6), 1329-1340. https://doi.org/10.1007/s10071-014-0768-6

Bower, G. H., Thompson-Schill, S., & Tulving, E. (1994). Reducing retroactive interference: An interference analysis. *Journal of Experimental Psychology*, *20*(1), 51-66. https://doi.org/10.1037/0278-7393.20.1.51

Brackbill, Y., & O'Hara, J. (1958). The relative effectiveness of reward and punishment for discrimination learning in children. *The Journal of Comparative and Physiological Psychology*, *51*(6), 747-751. https://doi.org/10.1037/h0038282

Brauer, J., & Call, J. (2015). Apes produce tools for future use. *American Journal of Primatology*, *77*(3), 254-263. https://doi.org/10.1002/ajp.22341

Casler, K., Bickel, L., & Hackett, E. (2013). Separate but equal? A comparison of participants and data gathered via Amazon's MTurk, social media, and face-to-face behavioral testing. *Computers in Human Behavior*, *29*(6), 2156-2160. https://doi.org/10.1016/j.chb.2013.05.009

Clayton, N. S., Bussey, T. J., & Dickinson, A. (2003). Can animals recall the past and plan for the future?. *Nature Reviews Neuroscience*, *4*(8), 685-691. https://doi.org/10.1038/nrn1180

Clayton, N. S., & Dickinson, A. (1998). Episodic-like memory during cache recovery by scrub jays. *Nature*, *395*(6699), 272-274. https://doi.org/10.1038/26216

Cowie, S. (2018). Behavioral time travel: Control by past, present, and potential events. *Behavior Analysis: Research and Practice*, *18*(2), 174-183. https://doi.org/10.1037/bar0000122

Cowie, S. (2020). Some weaknesses of a response-strength account of reinforcer effects. *European Journal of Behavior Analysis*, *21*(2), 348-363. https://doi.org/10.1080/15021149.2019.1685247

Cowie, S., & Davison, M. (2021). Pigeons prefer to invest early for future reinforcers. *Journal of the Experimental Analysis of Behavior*, *115*(3), 650-666. https://doi.org/10.1002/jeab.687

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*(1), 63-86. https://doi.org/10.1901/jeab.2011.96-63

Cowie, S., Virués-Ortega, J., McCormack, J., Hogg, P., & Podlesnik, C. A. (2021). Extending a misallocation model to children's choice behavior. *Journal of Experimental*

Psychology: Animal Learning and Cognition, 47(3), 317-325. https://doi.org/10.1037/xan0000299

Cowie, S., Zhai, E., & Elliffe, D. (2021). Surprise! Pigeons and humans respond similarly to unexpected reinforcers. *Conductual, 9*(1), 45-56. http://conductual.com/articulos/Surprise.%20Pigeons%20and%20humans%20respond%20similarly%20to%20unexpected%20reinforcers.pdf

Crystal, J. D. (2010). Episodic-like memory in animals. *Behavioural Brain Research, 215*(2), 235-243. https://doi.org/10.1016/j.bbr.2010.03.005

Davison, M. (2017). Killeen and Jacobs (2016) are not wrong. *Behavior Analyst, 40*(1), 57-64. https://doi.org/10.1007/s40614-017-0118-5

Davison, M., & Elliffe, D. (2010). Divided stimulus control: A replication and a quantitative model. *Journal of the Experimental Analysis of Behavior, 94*(1), 13-23. https://doi.org/10.1901/jeab.2010.94-13

de Mahy, D., Esteve, N. A., & Santariello, A. (2021). New test, old problems: Comment on 'New Caledonian crows plan for specific future tool use'. *Proc. R. Soc. B, 288*(1958), 1-3. https://doi.org/10.1098/rspb.2021.0186

Domjan, M. (2015). The Garcia-Koelling selective association effect: A historical and personal perspective. *International Journal of Comparative Psychology, 28*(1), 1-14. https://doi.org/10.46867/ijcp.2015.28.01.08

Feeney, M. C., Roberts, W. A., & Sherry, D. F. (2009). Memory for what, where and when in the black-capped chickadee (Poecile atricapillus). *Animal Cognition, 12*(6), 767-777. https://doi.org/10.1007/s10071-009-0236-x

Feeney, M. C., Roberts, W. A., & Sherry, D. F. (2011). Black-capped chickadees (Poecile atricapillus) anticipate future outcomes of foraging choices. *Journal of Experimental Psychology*, *37*(1), 30-40. https://doi.org/10.1037/a0019908

Flemming, T. M., Beran, M. J., Thompson, R. K. R., Kleider, H. M., & Washburn, D. A. (2008). What meaning means for same and different: Analogical reasoning in humans (Homo sapiens), chimpanzees (Pan troglodytes), and rhesus monkeys (Macaca mulatta). *Journal of Comparative Psychology*, *122*(2), 176-185. https://doi.org/10.1037/0735-7036.122.2.176

Garcia, J., & Koelling, R. A. (1966). Relation of cue to consequence in avoidance learning. *Psychonomic Science*, *4*, 123-124. https://doi.org/10.3758/BF03342209

Green, L., Fry, A. F., & Myerson, J. (1994). Discounting of delayed rewards: A life-span comparison. *Psychological Science*, *5*(1), 33-36. https://doi.org/10.1111/j.1467-9208.1994.tb00610.x

Green, L., Myerson, J., & Ostaszewski, P. (1999). Discounting of delayed rewards across the life span: Age differences in individual discounting functions. *Behavioural Processes*, *46*(1), 89-96. https://doi.org/10.1016/S037-6357(99)00021-2

Hogan, D. E., Edwards, C. A., & Zentall, T. R. (1981). The role of identity in the learning and memory of a matching-to-sample problem by pigeons. *Bird Behavior*, *3*(1-2), 27-36. https://doi.org/10.3727/015613881791560892

Kabadayi, C., & Osvath, M. (2017). Ravens parallel great apes in flexible planning for tool-use and bartering. *Science*, *357*(6347), 202-204. https://doi.org/10.1126/science.aam8138

Kaplan, P. S., & Hearst, E. (1982). Bridging temporal gaps between CS and US in autoshaping: Insertion of other stimuli before, during and after CS. *Journal of Experimental Psychology, 8*(2), 187-203. https://doi.org/10.1037/0097-7403.8.2.187

Klapes, B., Calvin, O. L., & McDowell, J. J. (2020). A discriminated rapid-acquisition laboratory procedure for human continuous choice. *Journal of the Experimental Analysis of Behavior, 114*(1), 142-159. https://doi.org/10.1002/jeab.612

Krageloh, C. U., Davison, M., & Elliffe, D. (2005). Local preference in concurrent schedules: The effects of reinforcer sequences. *Journal of the Experimental Analysis of Behavior, 84*(1), 37-64. https://doi.org/10.1901/jeab.2005.114-04

Lawicka, W. (1968). Differing effectiveness of auditory quality and location cues in two forms of differentiation learning. *Acta Biologiae Experimentalias, 29*(1), 83-92. https://rcin.org.pl/ibd/Content/4775/WA488_4271_P180-T29-z1_ABE.pdf#page=85

Leslie, J. C. (2018). Testing the limits of behavior analysis: A review of Frans de Waal's Are we smart enough to know how smart animals are?. *Journal of the Experimental Analysis of Behavior, 110*(3), 569-585. https://doi.org/10.1002/jeab.482

Lieberman, D. A., McIntosh, D. C., & Thomas, G. V. (1979). Learning when reward is delayed: A marking hypothesis. *Journal of Experimental Psychology: Animal behavior processes, 5*(3), 224-242. https://doi.org/10.1037/0097-7403.5.3.224

Lind, J., Enquist, M., Ghirlanda, S. (2015). Animal memory: a review of delayed matching-to sample data. *Behavioural Processes 117*, 52-58. https://doi.org/10.1016/j.beproc.2014.11.019

McDowd, J. M. (2007). An overview of attention: Behavior and brain. *Journal of Neurologic Physical Therapy, 31*(3), 98-103. https://doi.org/10.1097/NPT.0b013e31814d7874

Meyer, W. J., & Offenbach, S. I. (1962). Effectiveness of reward and punishment as a function of task complexity. *The Journal of Comparative and Physiological Psychology, 55*(4), 532-534. https://doi.org/10.1037/h0049119

Meyers-Manor, J. E., Overmier, B., Hatfield, D. W., & Croswell, J. (2014). Not so bird brained: Pigeons show what-where-when memory both as time of day and how long ago. *Journal of Experimental Psychology, 40*(2), 225-240. https://doi.org/10.1037/xan0000016

Miller, R., Frohnweiser, A., Ding, N., Troisi, C. A., Schiestl, M., Gruber, R., Taylor, A. H., Jelbert, S. A., Boeckle, M., & Clayton, N. S. (2020). A novel test of flexible planning in relation to executive function and language in young children. *R. Soc. Open Sci, 7*(4), 1-13. https://doi.org/10.1098/rsos.192015

Miyata, H., & Fujita, K. (2008). Pigeons (Columba livia) plan future moves on computerised maze tasks. *Animal Cognition, 11*(3), 505-516. https://doi.org/10.1007/s10071-008-0141-8

Olton, D. S., & Samuelson, R. J. (1976). Remembrance of places passed: Spatial memory in rats. *Journal of Experimental Psychology, 2*(2), 97-116. https://doi.org/10.1037/0097-7403.2.2.97

Ono, K. (1987). Superstitious behavior in humans. *Journal of the Experimental Analysis of Behavior, 47*(3), 261-271. https://doi.org/10.1901/jeab.1987.47-261

Osvath, M. (2016). Putting flexible animal prospection into context: Escaping the theoretical box. *Wiley Interdisciplinary Reviews, 7*(1), 5-18. https://doi.org/10.1002/wcs.1372

Pertzov, Y., Manohar, S., & Husain, M. (2017). Rapid forgetting results from competition over time between items in visual working memory. *Journal of Experimental Psychology, 43*(4), 528-536. https://doi.org/10.1037/xlm0000328

Podlesnik, C. A., Ritchey, C. M., Kuroda, T., & Cowie, S. (2022). A quantitative analysis of the effects of alternative reinforcement rate and magnitude on resurgence. *Behavioural Processes, 198*, 1-17. https://doi.org/10.1016/j.beproc.2022.104641

Raby, C. R., Alexis, D. M., Dickinson, A., & Clayton, N.S. (2007). Planning for the future by western scrub-jays. *Nature, 445*(7130), 919-921. https://doi.org/10.1038/nature05575

Raby, C. R., & Clayton, N. S. (2012). Episodic memory and planning. In T. K. Shackelford & J. Vonk (Eds.), *The Oxford Handbook of Comparative Evolutionary Psychology* (pp. 217-235). https://doi.org/10.1093/oxfordhb/9780199738182.013.0012

Rachlin, H. (1995). Self-control: Beyond commitment. *The Behavioral and Brain Sciences, 18*(1), 109-121. https://doi.org/10.1017/S0140525X00037602

Redshaw, J., Taylor, A. H., & Suddendorf, T. (2017). Flexible planning in ravens?. *Trends in Cognitive Sciences, 21*(11), 821-822. https://doi.org/10.1016/j.tics.2017.09.001

Richards, R. W. (1981). A comparison of signaled and unsignaled delay of reinforcement. *Journal of the Experimental Analysis of Behavior, 35*(2), 145-152. https://doi.org/10.1901/jeab.1981.35-145

Rudski, J. (2001). Competition, superstition and the illusion of control. *Current Psychological Research and Reviews, 20*(1), 68-84. https://doi.org/10.1007/s12144-001-1004-5

Russell, J., Alexis, D., & Clayton, N. S. (2010). Episodic future thinking in 3- to 5-year-old children: The ability to think of what will be needed from a different point of view. *Cognition, 114*(1), 56-71. https://doi.org/10.1016/j.cognition.2009.08.013

Scarf, D., & Colombo, M. (2010). The formation and execution of sequential plans in pigeons (Columba livia). *Behavioural Processes, 83*(2), 179-182. https://doi.org/j.beproc.2009.12.004

Shahan, T. A. (2017). Moving beyond reinforcement and response strength. *Behavior Analyst, 40*(1), 107-121. https://doi.org/10.1007/s40614-017-0092-y

Shahan, T. A., & Podlesnik, C. A. (2006). Divided attention performance and the matching law. *Learning & Behavior, 34*(3), 255-261. https://doi.org/10.3758/BF03192881

Simon, C., Bernardy, J. L., & Cowie, S. (2020). On the "strength" of behavior. *Perspectives on Behavior Science, 43*(4), 677-696. https://doi.org/10.1007/s40614-020-00269-5

Skinner, B. F. (1948). 'Superstition' in the pigeon. *Journal of Experimental Psychology, 38*(2), 168-172. https://doi.org/10.1037/h0055873

Stoet, G. (2010). PsyToolkit: A software package for programming psychological experiments using Linux. *Behavior Research Methods, 42*(4), 1096-1104. https://doi.org/10.3758/BRM.42.4.1096

Stoet, G. (2017). PsyToolkit: A novel web-based method for running online questionnaires and reaction-time experiments. *Teaching of Psychology, 44*(1), 24-31. https://doi.org/10.1177/0098628316677643

Suddendorf, T., & Corballis, M. C. (2010). Behavioural evidence for mental time travel in nonhuman animals. *Behavioural Brain Research, 215*(2), 292-298. https://doi.org/10.1016/j.bbr.2009.11.044

Suddendorf, T., Corballis, M. C., & Collier-Baker, E. (2009). How great is great ape foresight?. *Animal Cognition, 12*(5), 751-754. https://doi.org/10.1007/s10071-009-0253-9

Sulikowski, D., & Burke, D. (2007). Food-specific spatial memory biases in an omnivorous bird. *Biology Letters, 3*(3), 245-248. https://doi.org/10.1098/rsbl.2007.0122

Thorndike, E. L. (1911). *Animal intelligence: Experimental studies*. New York, NY: Macmillan. https://doi.org/10.5962/bhl.title.55072

Tulving, E. (2005). Episodic memory and autonoesis: Uniquely human?. In H. Terrace & J. Metcalfe (Eds.), *The Missing Link in Cognition: Origins of Self-Reflective Consciousness* (pp. 3-56). https://doi.org/10.1093/acprof:oso/9780195161564.001.0001

Urcuioli, P. J., & Nevin, J. A. (1975). Transfer of hue matching in pigeons. *Journal of the Experimental Analysis of Behavior, 24*(2), 149-155. https://doi.org/10.1901/jeab.1975.24-149

Vanderveldt, A., Oliveira, L., & Green, L. (2016). Delay discounting: Pigeon, rat, human – does it matter?. *Journal of Experimental Psychology: Animal learning and cognition, 42*(2), 141-162. https://doi.org/10.1037/xan0000097

Warden, C. J., & Aylesworth, M. (1927). The relative value of reward and punishment in the formation of a visual discrimination habit in the white rat. *The Journal of Comparative Psychology, 7*(2), 117-127. https://doi.org/10.1037/h0073058

Zentall, T. R., & Smith, A. P. (2016). Delayed matching-to-sample: A tool to assess memory and other cognitive processes in pigeons. *Behavioural Processes, 123*, 26-42. https://doi.org/10.1016/j.beproc.2015.07.002