Demonstrating inferential reasoning through behavioural choices: an experimental analysis

Dániel Rivas-Blanco[1,2,*], Marco Vasconcelos[2], Friederike Range[1], Alex Kacelnik[3,†] and Tiago Monteiro[1,2,*,†]

[1]Domestication Lab, Konrad Lorenz Institute of Ethology, Department of Interdisciplinary Life Sciences, University of Veterinary Medicine Vienna, 1210 Vienna, Austria

[2]William James Center for Research, University of Aveiro, 3810-193 Aveiro, Portugal

[3]Department of Biology, University of Oxford, OX1 3PS Oxford, UK

[*]correspondence (daniel.rivas-blanco@vetmeduni.ac.at or tiago.monteiro@vetmeduni.ac.at)

[†]co-senior authors

Abstract

We critically examine theoretical and experimental evidence that animals' reason by exclusion, focusing on a widespread protocol known as the 2-Cups Task. We test the hypothesis that subjects solve inference by exclusion tasks by responding independently to two options, without implementing complex reasoning, as posited by the Sequential Choice Model. In the critical 2-Cups Task subjects are signalled that a bait is hidden in one of two cups and then shown that one is empty while the other's content is not revealed. Subjects are then allowed to choose and are considered successful if they choose the unrevealed cup. We emulate the procedure using arbitrary stimuli and starlings (Sturnus unicolor) as subjects. Birds learned in encounters with single stimuli that one coloured pecking key never yielded rewards ("empty cup") and another yielded rewards on half the trials ("resting cup"). When facing only one of the two stimuli, birds displayed differential behaviour, pecking at the resting cup and not at the empty one. From the first time they faced the stimuli together they all chose appropriately. We argue that primates in the 2-Cups Task can behave similarly, with correct choices not being evidence for inference by exclusion. Studies of inference by exclusion should report data on behaviour towards single cups, to consider the possibility that simple mechanisms explain subjects' choices. This contribution highlights the contrast between simple, generalisable, and falsifiable models of comparative cognition and alternatives that rely on high-level operations such as 'causal understanding', but without controls to exclude the former.

Keywords: Causal understanding; Inference by exclusion; Sequential Choice Model; Starling; 2-Cups Task.

Introduction

The quest to infer mental events from observable behaviour is a persistent source of inspiration and tension between different research programmes. Comparative cognition research benefits from a rich theoretical contrast between schools that avoid reliance on mental assumptions, such as Skinner's radical behaviourism (see Schneider, 1987), and at the other extreme, those that invoke high-level constructs such as 'causal understanding' as explanatory mechanisms for animal behaviour (e.g. Hanus & Call, 2011). Following Heyes (2012), we refer to the use of high-level cognitive constructs as explanatory tools as super-cognitivism, to highlight the contrast with cognitive approaches that model psychological processes algorithmically, using observable behaviour as both input and output. Scholars adept to either radical behaviorism or ultracognitivism often mutually criticise perceived excesses in anthropomorphism (Kennedy, 1992) or anthropodenialism (De Waal, 1999), perhaps with some justification in all cases.

We argue theoretically and show experimentally that some observations widely attributed to high-level constructs such as logical reasoning and/or causal understanding can be modelled precisely and productively in terms of observable behaviour, including cases in which the data are interpreted as justifying super-cognitive claims. We focus on the *2-Cups Task* paradigm (**Figure 1A**; Call, 2004; Call, 2006) as described in detail in a widely subscribed proposal by the *ManyPrimates* research consortium (ManyPrimates, 2024). We aim to show that an algorithmic model that is simple for agents to perform, empirically verifiable, and frugal in parameters can (and does) generate flexible behaviour in that and multiple alternative scenarios. Our proposal aims to implement objectives expressed (among others) by McNamara and Houston (2009) in behavioural ecology, Shettleworth (2010) in comparative cognition and Gigerenzer and collaborators (1999) in human decision research. Shettleworth (2010) expressed the problem we address succinctly:

> *"…demonstrations that complex human-like behavior arises from simple mechanisms rather than from 'higher' processes, such as insight or theory of mind, are often seen as uninteresting and 'killjoy', almost a denial of mental continuity between other species and humans"*.

Within this spirit, we propose a simple mechanism that can account for complex adaptive behaviour in many different situations, including the one at the centre of this study.

The *ManyPrimates 3A - Inference by Exclusion* proposal starts from the premise that individuals in nature need to make inferences about hidden information (ManyPrimates, 2024). We quote:

> "*Important events are often not observable. Consequently, individuals must make inferences about unobserved events. One method for making such inferences is inference by exclusion (IbE) - the ability to reconstruct the identity of an unknown alternative by eliminating known options*".

We believe that this premise is not justified theoretically and examine it experimentally. A useful theoretical starting point is Gibson's (1979) concept of affordances in visual perception. He argued that evolution shapes organisms to perceive and act upon what matters to them in their environment, writing that *to perceive visual stimuli is to perceive what they afford*. Gibson's ecological view of perception needs not be universally shared, but for present purposes it is enough to agree that if an event is important it must have observable consequences, and then species can evolve mechanisms to act adaptively towards them, without necessarily making mental inferences about events they cannot directly perceive or 'understand'.

The claimed need for inferential reasoning downplays the potential of simple learning processes evolutionarily designed to use reliable predictors to generate adaptive behaviour. To avoid confusions, we note that some authors use the phrase 'Inference by Exclusion' to address observable behaviour rather than mental processes (see discussion by Mody and Carey 2016), but here we follow its most widespread and conceptually ambitious use, in which making inferences and following logical reasoning is predicated to explain behaviour. The super-cognitivist handling of *Inference by Exclusion* does not just include logical reasoning, but implicitly involves psychological object permanence, representation of probabilities and causal understanding. We will argue that none of these is required to account for the reported data. To support our claims, we apply an alternative choice model to an

experimental analogue of the *2-Cups Task*, using spotless starlings (*Sturnus unicolor*) as subjects (**Figure 1B**), and suggest that our alternative framework may be considered for primates and other species. We do not predicate the absence of reasoning processes in any species, but object to some claims for their presence and show that there exist approaches to modelling cognition that are epistemically preferable for their frugality, their ability to generate testable predictions and their greater vulnerability to empirical falsification. We describe our model in the next section, then examine the *2-Cups Task* in some detail and present our experimental emulation.

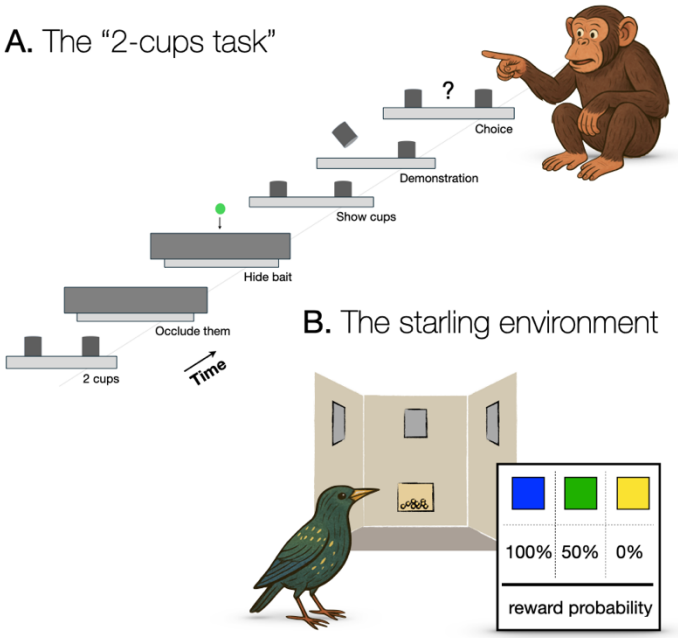

Figure 1. The 2-cups task. A. Example trial in the 2-cups protocol where the subject is shown the content of the empty cup before being allowed to make a choice. **B.** Emulation of the 2-cups task with starlings using a 3-key operant panel and 3 coloured stimuli signalling *Sure Reward* (100%), sure *No Reward* (0%), or *Uncertain* (50%). Animal images were generated using ChatGPT (https://chat.openai.com).

Logical reasoning or independent responding

When a goal-directed agent faces multiple mutually exclusive options, it is tempting to model the agent as reasoning about the options' relative merits respect to its goal, then making a choice determined by these deliberations. This implies that preference is constructed when facing the choice, fits our subjective experience, and dominates the field of human decision making. However, the agent may be alternatively conceptualised as possessing a pattern of behaviour towards each option when

encountering it on its own and resorting to the same responses when facing more than one simultaneously. In this case, the behavioural candidates compete for a mutually exclusive *Behavioural Final Common Path* (McFarland & Sibly 1975), so that only one action is realised. When this happens, preference is pre-determined and expresses cognitive contents that exist before the subject makes a choice. Such a mechanism makes it superfluous to argue that complex logical operations cause the choice. The critical question is then: is behaviour when facing single options sufficient to predict behaviour when two options are present? If the answer is yes, then we should remain agnostic as to what mental operations happen at choice time. Curiously, this simple control is not reported in studies that conclude that choices are supported by logical reasoning at the time of choice.

Patterns of responding to single options can be innate or acquired through previous learning in sequential encounters. This distinction is interesting but not crucial to the present argument. What is critical is whether options faced in choices are treated in parallel or compared by some rational process. Parallel processing of options has been tested successfully in diverse protocols (Kacelnik et al., 2011; Monteiro et al., 2020; Shapiro et al., 2008). The Sequential Choice Model (SCM) includes hypotheses about the environment of evolutionary adaptation, about individual learning and about the cognitive process of action competition in choice encounters. This is a summary:

(1) *Environment:* sequential encounters with single targets are probably more frequent and evolutionarily influential than choices between multiple options.

(2) *Acquisition:* in sequential encounters, experience (or natural selection) leads to differential behaviour (pointing, grasping, pecking, touching, approach-withdrawal, sniffing, etc.) towards recognisable stimuli. For instance, a frog may innately snap at a worm and ignore a leaf, or an infant primate may gesture towards a ripe fruit and ignore an unripe one.

(3) *Choice:* if 2 or more targets are faced simultaneously, their option-specific actions compete for expression. The hypothetical frog and infant primate would 'choose' appropriately without any comparative process.

These three hypotheses have different status:

The first addresses ecological relevance and is not critical to the model's verification. It is plausible, because if potential targets have independent distributions, facing them sequentially should be more frequent and evolutionarily influential than simultaneously. Early applications of optimality theory to behavioural ecology addressed patch exploitation, diet choice and copulation duration as sequential decisions of the kind 'stay or go away' (Charnov 1976 a, b; Parker & Stuart, 1976; see also Freidin et al 2009). While leks or fruit clusters exemplify simultaneous encounters, the set of scenarios in which sequential encounters prevail is rich enough to suspect their primacy.

The second hypothesis rests on empirical observations (e.g. Aw et al, 2012; Macías et al., 2021; Mazur, 2010; Monteiro et al., 2020; Vasconcelos et al., 2010; Vasconcelos et al., 2013) and should be verified in every empirical test. It is not unique to this model, but this model needs it. One corroboration of this hypothesis is seen in laboratory tests of SCM, because options encountered sequentially lead to the acquisition of differential latencies to respond (faster response times to better options). This may seem intuitively trivial, but it is in fact puzzling: under the assumption of reward rate maximisation that underlies classical foraging models there is no reason why taking any opportunity within the acceptable set should take longer than minimally possible, but models that include these latencies yield better empirical fitting than those ignoring them (Reboreda & Kacelnik, 1991; Shapiro et al., 2008). Another corroboration is observed in classical (Pavlovian) conditioning. When discriminable stimuli reliably predict different outcomes, animals show differential anticipatory behaviour even if the outcomes are not contingent on the behaviours (e.g., Mazur, 2023; Lovibond & Westbrook, 2024). This matters, because response latencies or other forms of anticipatory behaviours may be sufficient to explain how animals act when facing a choice.

The third element is in fact the hypothesis under test: if choices are predictable from pre-existing behaviour, then no novel cognitive operation (such as logical reasoning) is necessary at choice time.

If the last two hypotheses are confirmed in a choice experiment, then the experiment does not demonstrate inferential reasoning or other high-level cognitive hypotheses such as causal understanding. Mody and Carey (2016) and Rescorla (2009) also question the evidence for current claims for inferential reasoning, but our rationale is more radical, because we argue that animals facing multiple options may not even cognise the task as a choice, but rather treat options independently. Of course, in this scheme actual behaviour in a choice may be slightly different from that when there is a single target for mechanical or attentional reasons, but this is not a justification for inferring reasoning.

The *2-Cups Task*

The search for evidence of thinking in non-human animals can be traced to the ancient Greeks. In the 3rd century BCE, Chrysippus discussed a form of reasoning called "disjunctive syllogism" or *modus tollendo ponens*, which is summarised as "Either p or q. Not p. Therefore, q". This deceptively simple operation assumes competence to handle the relation 'OR', implying that out of two world states only one can be true (see Mody & Carey, 2016). It also requires understanding the concept of negation, 'NOT'. Together, given OR, if one is negated, the other must be true. Five centuries later, Sextus Empiricus wrote that Chrysippus had used a thought experiment to claim that a dog could demonstrate logical reasoning without language, through its behaviour. In a simplified version, an imaginary dog chases a hare that gets out of sight where the road bifurcates. The dog sniffs one side and, detecting no hare smell, runs straight into the other, without sniffing. The failure to seek information in the latter is taken as evidence of inferential reasoning by exclusion. This assumes that the dog understands that there is exactly one hare, and then the probabilities of it being in each of the roads add to unity: if it is zero in one, it must be one in the other. At behavioural level the dog's actions are consistent with such logic, but reasoning is not required. If in its past whenever the dog run into a sniffed-but-odourless road it missed dinner, but running into un-sniffed roads had mixed outcomes, it will waste no time in entering an uncertain, un-sniffed, roads but pause in front of odourless roads. In a road fork, this results in a 'choice' without the

need for logical reasoning; the dog promptly runs into the unsniffed road without using OR or NOT, does not estimate probabilities, and does not assume that there is exactly one hare in the world.

In the *2-Cups Task* a subject familiar with receiving food after appropriate behaviour (such as pointing) faces opaque cups. There are variants of this protocol with more than 2 cups (e.g., Mody & Carey, 2016; Call, 2022), but we focus on the simplest 2-cups case because the argument would not differ. An experimenter hides one bait without identifying in which cup and then briefly lifts one of the cups, showing its content (or absence of it), leaving the other resting. For each cup, the subject witnesses one out of three scenes that signal that the cup is baited, empty, or unknown. The subject then points to either cup and receives its content (i.e. food or nothing, see **Figure 1A**). The three displays are observable events that signal important outcomes. In the super-cognitive interpretation if the subject points at the resting cup when the lifted box was empty, it is concluded that it has cognitively performed the *Inference by Exclusion* logical operation. In the SCM interpretation, the animal would treat each of the three scenes differently if met in isolation, for instance pointing fast to a manifestly full cup, slower to an uncertain resting one, and not at all to a manifestly empty one. If it just aims to act in the same way when two cups are present, it would show the correct 'choice' without logical reasoning.

Correct behaviour in the *2-Cups* protocol can be generated by multiple processes. M. Rescorla (2009) used a Bayesian analysis to argue that "*results supposedly indicative of non-linguistic deduction can be explained without citing logical reasoning or logically structured mental states",* and Lauffer et al (2016) used a 2-alternative forced-choice discrimination task to demonstrate that behaviour consistent with inference by exclusion can be generated by the subject's acquisition of response equivalence, an associative rather than reasoning process. Both are algorithmically efficient explanations that omit 'logical reasoning' or 'causal understanding'. We test experimentally whether animals can solve such tasks without even framing the problem as a choice.

An emulation of the *2-Cups Task*

Our experimental procedure reproduces the critical contingencies between stimuli and behavioural outcomes in the *2-Cups Task* (**Figure 1B**). Using a discrete trials procedure, starlings were pre-trained to peck at four different stimuli encountered sequentially as single options. During pre-training the 5th peck at a single key caused food to be delivered. Once they pecked reliably, they entered a *Sequential Training* phase. Now, three stimuli indicated different food contingencies, labelled *Sure Reward*, *No-Reward*, or *Uncertain*. Pecking them was followed by food, no food, *or* food /no food respectively. A response to a central attention key was followed by one of the three stimuli appearing for 2s in either of two side keys. This emulated the hiding of a bait between the 2 cups. During these 2s, pecks were recorded (as 'Pre-GO Pecks') but had no programmed consequence. Once the 2s lapsed, the stimulus was switched off and a 'GO' stimulus (always white) was presented in the same side key. The first peck at a key in its GO state realised the trial's food contingency (**Figure 2A**, top). After 10 days in this phase, the *Simultaneous Testing* phase followed. This phase included choice trials, that were as before but with both side keys showing two of the three stimuli simultaneously for 2s and then turning white to enter the GO state. A single peck to a key in the GO state caused the offset of the alternative and the realisation of the respective food contingency (**Figure 2B**, top). Pre-GO pecks were recorded but had no effect. Choices were determined by the first peck to a key in the GO state. This emulates the choice between cups in the *2-Cups* protocol. We describe starlings' behaviour in sequential encounters and in choices, examining whether responding in single encounters is sufficient to predict choices in simultaneous trials.

Our procedure emulates the contingencies in the basic *2-Cups* protocol, but with arbitrary stimuli and controlled acquisition instead of food containers. Animals in *2-Cups* experiments do not have a pre-training phase, but they are often experienced adults that had opportunities to learn that a cup recently seen to be full is likely to contain food and one that has not been seen to be either full or empty may lead to food or no food. Fortunately, this difference in training history between the *2-Cups* experiments and our emulation is not an obstacle for the present argument. If reports

of *2-Cups* experiments included the behaviour of subjects in single cup presentations, our claim that *1-Cup* encounters would predict choices in simultaneous trials would be put to test.

Results

The hypothesis under test is that choice (pecking at one of the GO stimuli in a choice trial, equivalent to primates pointing after the demonstration or different choice actions in other species) is pre-determined before pairs of options become available, implying that preference does not depend on cognitive operations that include the logical disjunction 'OR'. We first examine behaviour during the *Sequential Training* and then in the *Simultaneous Testing* phases (**Figure 2**).

Sequential Training **phase**

At the start of sequential training all starlings reliably pecked at all keys. **Figure 2A** shows pecking at the Pre-GO stimuli during the 2s Pre-GO interval and response latencies towards the GO key in the first four sequential sessions and in the last one, just before switching to the *Simultaneous Testing* phase. In summary, the birds rapidly extinguished responding at the *No Reward* option and maintained responding when there was a chance of food. These trends were statistically reliable. For the *Sure Reward* option, Pre-GO pecks persisted (effect of session: $p = 0.114$), and response latency to GO decreased ($p < 0.001$); for the *Uncertain Reward* Pre-GO pecks decreased ($p = 0.031$) and response time increased ($p < 0.001$) when compared with the *Sure Reward* option (the model's intercept), and for the *No Reward* stimulus there was a fast decrement in Pre-GO pecks and an increase in GO response latency ($p < 0.001$ for both). See **Table S1 — Models 1.1 and 2.1—** for statistical details. Consistently, response times to the first peck at the GO stimulus remained very short for the two rewarded stimuli and increased for the unrewarded one.

By the last session of the *Sequential Training* phase both Pre-GO pecking and response times to peck at the GO key were significantly different between the *Sure*

Reward stimulus and the *No Reward* one (p<0.001 for both pecks and response times), with no significant differences between the two rewarded alternatives (pecks: p=0.086, response times: p=0.14; see **Table S1 — Models 3 and 4—** for details and **Supplementary Figures S1 and S2** for individual results). In terms of the *2-Cups Task* emulation, that would be equivalent to subjects discriminating behaviourally between the unrevealed uncertain cup and the revealed empty one (in this emulation the *Uncertain* and *No Reward* arbitrary stimuli) when a single cup is presented.

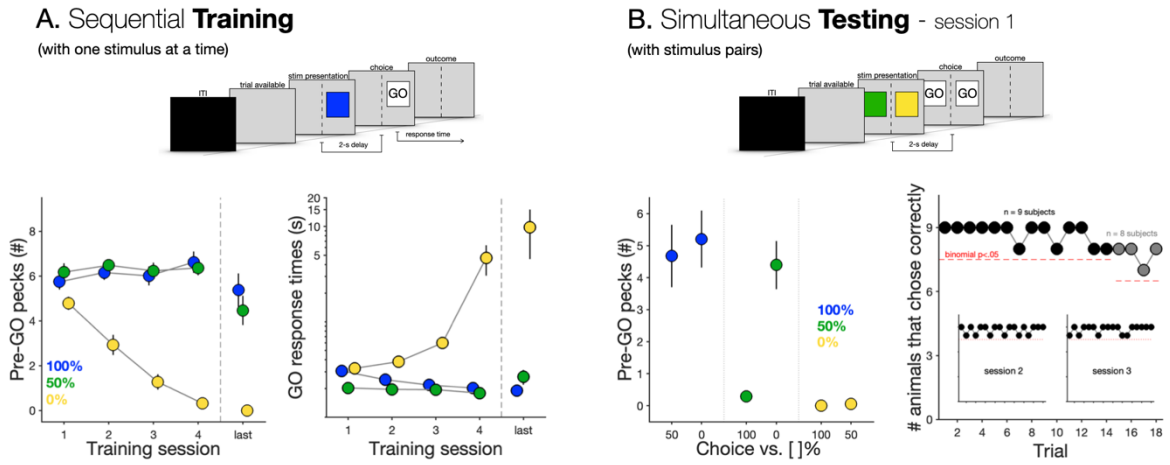

Figure 2. Starlings' performance in the 2-cups task analogue. A. Example of a sequential trial (top). Average (±SE) Pre-GO pecks (bottom left) and response times (bottom right) for each alternative (*Sure Reward* - 100%; *No Reward* - 0% and *Uncertain* - 50%) during the first 4 (and last) sessions of sequential training. **B.** Example of simultaneous choice trial (top). Average (±SE) Pre-GO pecks for choice trials for all simultaneous choice trial types: *Sure Reward* vs. *Uncertain*, *Sure Reward* vs. *No Reward*, *Uncertain* vs. *Sure Reward*, *Uncertain* vs. *No Reward*, *No Reward* vs. *Sure Reward* and *No Reward* vs. *Uncertain*, respectively in blue, green and yellow (bottom left). Number of animals choosing *Uncertain* (trial-by-trial) in choices between it and *No Reward* for the first testing session (see Figure Supplementary S5 for the remaining choice pairings). Black and grey markers depict trials with 9 and 8 subjects, respectively, and red dashed line binomial significance. Insets show the same data for testing sessions 2 and 3 (n=9 subjects).

Simultaneous Testing **phase**

In the final phase, subjects faced a pseudo-random sequence of 'sequential' and 'choice' trials (see *Methods* section, "*Simultaneous Testing"* phase subsection). We focus on preference during Pre-GO between *No Reward* and *Uncertain*, as this

emulates the critical choice between a lifted empty cup and a resting cup, the condition used as evidence of inference by exclusion.

Pre-GO pecks in choice trials

Figure 2B (bottom left; see also **Supplementary Figure S3** for individual data) shows Pre-GO pecks in the first choice session, in three panels. The left panel shows Pre-GO pecks at the *Sure Reward* (100%) stimulus when paired with either the *Uncertain* (50%, left) or *No Reward* (0%, right) alternatives. In both cases pecking was like that in sequential trials (p=0.766 and p=0.064, respectively; see **Table S1 —Models 6 and 7—** for details). The middle panel shows Pre-GO pecks at the *Uncertain* (50%) stimulus when paired with the *Sure Reward* (100%, left) or *No Reward* (0%, right) alternatives. In this case, Pre-GO pecks depended on the alternative. When paired with *Sure Reward* (100%) pecking dropped substantially (p<0.001; see **Table S1 — Models 8—** for details), while when paired with *No Reward* (0%) pecking was maintained (p=0.239; see **Table S1 —Models 9—** for details). This is evidence for response competition: the animal cannot be in two places at once, and discrimination ability acquired in sequential encounters is not fully captured by number of pecks in sequential encounters, probably due to a ceiling effect. Finally, the rightmost panel shows Pre-GO pecks at the *No Reward* (0%) option when paired with either of the other two. In both cases there was very little pecking.

In summary, Pre-GO pecking, although inconsequential in terms of reward, implies that discriminating behaviour exists before the animal is allowed to choose, in the sequential phase (for Pre-GO pecking in sequential trials during the *Simultaneous Testing* phase see **Supplementary Figure S4**). In terms of the *2-Cups Task* emulation, a direct comparison between the experiments could be explored if subjects' behaviours such as gazing or pointing were reported during the demonstration or just before consequential (instrumental) choices.

Choice in paired trials

Successful subjects in the *2-Cups Task* prefer the resting cup when paired with a lifted but empty one. **Figure 2B** (bottom right) shows that starlings were almost totally

error-free from the very first choice trial (this persisted throughout this session and the following ones). Mean proportion of correct choices in the first choice session was 0.968, ranging from 0.951 to 0.980 when subtracting and adding the standard error, respectively, this being substantially better than random ($p<0.001$; see **Table S1 —Model 5—** for details). The birds could have achieved this without logical inferences, by treating each stimulus as a competing opportunity to respond. The remaining pairings, although less relevant to the arguments, are presented in detail in the Methods section and in **Supplementary Figure S5**.

Discussion

We highlight how simple cognitive processes can lead to rich and adaptive behaviour, so that the observation of such behaviour does not evidence complex cognition. In particular, we show that differential behaviour towards opportunities typically met in sequential encounters generates correct choices when multiple opportunities are met simultaneously. The mechanism we explore with starlings shows that they can behave as if they reasoned by exclusion, when in fact choice is a consequence of each stimulus in the set eliciting pre-existent behaviour. Behaving as if making inferences by exclusion does not require forming estimates of probabilities or connecting alternatives with the OR disjunction and is less demanding computationally than logical reasoning. In addition, it seems ecologically more relevant, since encounter rates with single potential targets in nature are probably more frequent than simultaneous encounters requiring a choice. It remains possible that many species have evolved mechanisms to respond appropriately to recognisable stimuli but have no dedicated mechanisms evolved for choosing between simultaneously present targets.

We provide a novel test of the Sequential Choice Model, which in essence states that behaviour in sequential encounters can explain choices in a variety of experimental protocols (Kacelnik et al., 2011; Monteiro et al., 2020; Shapiro et al., 2008). Although some may consider this to be a 'killjoy' approach (Shettleworth 2010) given the widespread appeal of attributing high-level cognition, the opposite view is

compelling: simple mechanisms of behaviour can be readily implemented in natural or artificial agents and are sufficient to produce rich, flexible, and adaptive behavioural repertoires, and they should be preferred for their parsimony and simplicity. This is encouraging not only because it advances understanding of animal behaviour, but also because it potentially simplifies the design of intelligent artificial agents.

There are important caveats regarding the extent to which our experiment constitutes an informative emulation of the *2-Cups* protocol. We explicitly trained the starlings with sequential encounters before giving them choices, while this is not the case in tests of the *2-Cups* task. We believe that this is not a fatal difference for two main lines of argument. First, all tested animals in *2-Cups* protocols have rich histories, in which they must have acquired the relevant competences. In a study interpreted as demonstrating causal understanding in chimpanzees, Hanus and Call (2011) wrote:

> "*Recurrent experience during ontogeny definitely plays an important role in the acquisition of such cognitive skills—in human and nonhuman primates alike. Individuals need to interact repeatedly with various causal regularities within their physical world in order to be capable of transferring this specific knowledge to new situations.*"

This surely applies to how a subject in a *2-Cups* experiment learns to treat a cup lifted and shown to hide a bait differently from one that is resting on the surface. Regardless of the process of acquisition, it seems safe to expect that a subject that can solve the *2-Cups* problem as if employing inferential reasoning will respond differently to the same single cups if it were tested in sequential presentations. This is all that is needed for our emulation to be informative. We focused our analysis on the very first choices faced by the starlings. The reason for this is that we wanted to exclude the possibility that their behaviour had been shaped by having been rewarded for particular choices. While their sequential training was intensive, on their first simultaneous presentations the birds had no other way of choosing correctly (as they did) than relying on biases acquired in sequential encounters, when no inferential reasoning by exclusion is relevant. Our most conclusive argument however is that doubts about the relevance

of the emulation have a clearcut empirical solution by adding a control to the *2-Cups* protocol: testing what subjects do when a single cup, baited or unbaited, is presented, either briefly lifted or left resting. If, as we expect, their behaviour is predictive of what they do when 2 cups are present, then there would be no need to invoke a constructive cognitive operation at the time of choice. We look forward to reports of such controls.

Methods

Subjects

The starlings (*Sturnus unicolor*; n=9) were housed in pairs in indoor cages (1350 × 784 × 800 mm (l × w × h)) and kept under a 12L:12D cycle with gradual transitions at dawn and dusk. They were confined individually to one section of their cage for tests, but after each session, they had 4h of ad-libitum food and social interaction with the cage-partner. By the end of the testing season (early summer), the birds were released into the wild.

The experimental protocols were approved by the University's Animal Welfare Body (B003/2020) and by the Portuguese General Directorate of Food and Veterinary Medicine (0421/000/000/ 2021), and their capture and holding were approved by licences 514/2021/CAPT and 07/2021/DET respectively (both from the Portuguese Department of Wildlife and Conservation).

Apparatus

The experimental cages were provided with two side areas that could be split with sliding panels from a shared middle section (see Vasconcelos et al., 2015). Each side area had a working panel with three sections attached, each with a centrally placed response key (**Figure 1B**). The central section of the side panels also had a food hopper connected to a pellet dispenser (Campden Instruments) containing 20 mg *BioServ* precision dustless pellets. Contingencies were controlled by custom

software running on the Microsoft WINDOWS operating system and attached to an Animal Behaviour Environment Test System (Campden Instruments) via WhiskerServer (Cardinal & Aitken, 2010).

Trials and Experimental phases

The experiment had two phases: S*equential Training* and *Simultaneous Testing*. During the *Sequential Training* phase only sequential trials took place, while in the *Simultaneous Testing* phase sequential and choice trials were interspersed. In sequential trials, a pre-GO stimulus was presented for 2s, followed in the same location by a GO stimulus, always white, that stayed on until pecked once (**Figure 2A**, top). In choice trials two pre-GO stimuli were simultaneously presented (**Figure 2B,** top), each followed by a GO stimulus. The stimuli were square 4 by 4 LED light panels lit in green, blue, yellow or white (**Figure 1B**). The three colours were used as pre-GO stimuli, and white served as a common GO stimulus that, when pecked, realised the contingency. The three discriminative contingencies were *Sure Reward* (rewarded in 100% of the trials), *No Reward* (rewarded in 0% of the trials), and *Uncertain* (rewarded in 50% of the trials). During the *Simultaneous Testing* phase pre-GO stimuli were combined in their three possible pairs plus a fourth double presentation of the *Uncertain* stimulus in both left and right key (see *Simultaneous Testing* phase for details).

General procedure

At the beginning of each trial, the centre key of the apparatus displayed a blinking (700 milliseconds ON, 300 milliseconds OFF) white square, signalling that a trial was available. Blinking continued until the subject pecked, which turned off the centre key and initiated a trial. After this, either one (sequential trial) or both (simultaneous choice trial) side keys were lit and the 2s Pre-GO period started. In sequential trials (**Figure 2A**, top), the unused side key remained off and disabled. In choice trials (**Figure 2B**, top), both side keys lit up and remained enabled. Depending on the trial type, the

stimulus was either an unblinking green square, or a blinking (700 milliseconds ON, 300 milliseconds OFF) blue or yellow square (see the "*Sequential Training phase*" sub-section). During the stimuli (2 seconds; referred to as the "Pre-GO period" hereafter), pecks to each side key were recorded, but had no other effect.

At the end of the Pre-GO period, illuminated side key/keys (one or two in sequential or choice trials, respectively) switched to white (the GO signal) and remained on until pecked. Food was delivered or not depending on the corresponding Pre-Go stimulus. After this, both keys were turned off again for a 30-s inter-trial-interval (ITI).

Trials terminated if the subject failed to peck at the GO key within 1 minute. If the stimulus had been *Sure Reward* or *Uncertain*, the same trial was repeated after an ITI up to a maximum of 3 presentations (correction procedure).

Pre-training

All subjects had experience with the apparatus in tasks with different contingencies (a task on 'suboptimal/paradoxical choice', and another one using the mid-session reversal procedure). In the *Pre-training* phase, pecking at the white central key was followed by the onset in a side key of one of the stimuli that would later serve in the present study. All keys were assigned a fixed ratio (FR) reward schedule. Over two days, each stimulus was presented for three 50-trial sessions, divided in blocks of 25 trials with a 15-minute break and with a 60-s ITI. The fixed ratio increased in consecutive sessions, with FR1, FR3 and FR5 in the first, second and third sessions respectively. One individual needed an extra session due to apparatus malfunction. Before starting the Sequential Training phase, subjects had an additional FR5 session consisting of two blocks of 35 trials for each of the 4 stimuli. At the end of pre-training, all subjects were pecking reliably and similarly to the three coloured and white stimuli.

Sequential Training phase

In the Sequential Training phase, only one alternative was presented in each trial, while the other key remained dark. The side of the stimulus presented was pseudo-randomized so that it appeared the same number of times on both sides and not more than twice consecutively on the same side.

One of three coloured stimuli was presented in each trial, each associated with a specific probability of obtaining a food reward (**Figure 1B;** see also Monteiro et al., 2013). For half of the subjects the *Sure Reward* stimulus was yellow and the *No Reward* stimulus was blue, and the opposite for the other half of the animals. For the *Uncertain* stimulus, the key stayed green for the 2 seconds of Pre-GO time.

The subjects were trained daily, with up to 4 sessions per day, each comprising 108 trials[1].

All subjects shifted to the *Simultaneous Testing* phase after 10 consecutive days in sequential training.

Simultaneous Testing phase

During Simultaneous Testing, sequential trials were interspersed with five kinds of simultaneous choice trials, in which two alternatives were presented, one in each side key. Simultaneous test trials displayed two of the stimuli used in the *Sequential Training* phase, but, crucially, the possibility of getting the reward was now dependent on what the other stimulus was displaying (i.e., if the *Sure Reward* stimulus was used for one of the sides, the other side was not rewarded).

There were 4 different choice trial types that emulated the possible contingencies in a 2-Cups procedure, as follows:

[1] For 5 animals, the first training day was comprised of 3 sessions (as, by that point, it was unknown whether the animals were able to complete 4 sessions within a day). Due to this, analyses using the first 4 sessions of those animals include the first session of the second testing day.

1. *Uncertain* vs. *No Reward*: This choice emulates trials in the *2-Cups* procedure in which one cup is left resting and the other is lifted briefly to show that it is empty. Correct choice with these contingencies is interpreted as showing Inference by Exclusion.
2. *Sure Reward* vs. *No Reward*: This emulates a choice between a cup shown to be baited and another shown to be empty. This trial type came in two variants with stimuli either presented simultaneously (synchronous) or sequentially (asynchronous), in order to control for recency effects. In the latter variant the presentation of the second stimulus started after the end of the presentation time of the first one. The apparatus remained unresponsive until the end of the second presentation.
3. *Sure Reward* vs. *Uncertain*: This emulates a choice between a cup shown to be baited and another left resting so that its contents are unknown. In contrast with pre-training, in these trials Uncertain was programmed to be unrewarded, but this is irrelevant for the analysis of the initial simultaneous trials.
4. *Uncertain* vs. *Uncertain*: The purpose of these trials was to control for side preferences when both contingencies are equated.

All trial types were presented in a pseudo-randomized manner, with no more than two consecutive trials of the same kind or where the expected correct response was on the same side. In the case of *Sure Reward* vs. *No Reward* (asynchronous variant), the order and side of stimuli were balanced. Pseudo-randomization rules applied equally to sequential and choice trials. Sessions in this phase had 90 choice trials (18 per pairwise combination) and 18 sequential trials (6 per pairwise combination). This phase lasted for 10 consecutive days[2].

[2] An experimental error in *Uncertain* vs. *Uncertain* trial type during the first testing session of the first 5 animals being tested led to an early abortion of this session for the corresponding animals (leading to the early dropout in subjects seen in the last few trials of Figures 2 and S5). Due to this error, for the 5 animals affected, the data used for analyses and graphs not involving the *Uncertain* vs. *Uncertain* trial type starts at session 1, while for the graph involving this trial type (last panel of figure S5), data from session 2 was used instead. All corrupted data was removed from the dataset provided.

Statistics

We analysed number of pecks in the Pre-GO periods and both proportion of choices between stimuli and response time to the GO stimulus. The main topic of interest is how behaviour in the sequential pre-training phase relates to that in the initial simultaneous choices, when rules for logical reasoning about interdependencies such as 'if the reward is not here it must be there' are implausible.

Acquisition of discrimination

To examine acquisition of responding as a function of early experience, response times to the GO stimulus during the first day of the *Sequential Training* phase were analysed through a Linear Mixed Model (LMM) with the logarithm of the response times as the response variable, the subject as a random intercept, and the identity of the stimulus (*Sure Reward*, *No Reward*, or *Uncertain* stimuli) as both a fixed effect and a random slope. To calculate any changes in response times during the 4 sessions of the first day, one model (see **Table S1 — Model 1.1**) included the interaction between the stimulus and the session number (z-transformed, added as a random slope as well) and another without this variable (see **Table S1 — Model 1.2**). The same model structure as this latter one was used to analyze any differences in response times between stimuli in the last day of the *Sequential Training* Phase (see **Table S1 — Model 3**).

Since the amount of interactions with the keys before the GO stimulus (*i* pecks) was also dependent on the animal's acquisition of the stimuli (with more pecks for preferred stimuli), a parallel set of Generalized Mixed Models (GLMM) with a Poisson distribution were ran with the same model structure but with the number of Pre-GO pecks as a response variable (see **Table S1 — Models 2.1, 2.2, and 4**). Due to the distribution of the data, both models accounted for zero-inflation.

Choices in simultaneous trials

In order to ascertain whether the animals significantly preferred the rewarded option in the *Uncertain* vs. *No Reward* trials (a behaviour that would be consistent with *Inference by Exclusion*), the outcome (success or failure) of the trials was analysed

using a GLMM with a binomial distribution and a logit link function, with subject as a random intercept. As the purpose of this model was to calculate the difference between the observed results and chance (a probability of 0.5), no fixed effects were added to the model, and calculations were made based on the intercept (as per Nagel et al., 2024; see **Table S1 — Model 5**).

To assess any potential changes in the amount of Pre-GO pecks for each stimulus between the *Sequential Training* and the *Simultaneous testing* phases (and thus, any changes in the animals' response as soon as more than one alternative was available), an additional set of GLMMs with a Poisson distribution was run using data from the last session of the *Sequential Training* phase and the choice trials of the first session of *Simultaneous Testing* phase. A different model was fit for the stimuli in each of the potential choices (*Sure Reward* vs. *Uncertain*, *Sure Reward* vs. *No Reward*, *Uncertain* vs. *Sure Reward*, and *Uncertain* vs. *No Reward*). All of these models had the phase (*Sequential Training* or *Simultaneous Testing*) as both a fixed effect and a random slope, and the identity of the subjects as the random intercept (see **Table S1 — Models 6 to 9**). Due to general lack of Pre-GO pecks for the *No Reward* stimulus, no models were fitted for this.

All statistical analyses were carried out in R with the packages lme4, lmerTest, and car (Bates et al., 2014; Brooks et al., 2017; Fox et al., 2012; Kuznetsova et al., 2015; R Core Team, 2021).

Funding

D.R.B., F.R. and T.M. were supported by the Austrian Science Fund (FWF) Grant DOI 10.55776/P37052 and 10.55776/P33928. D.R.B. was funded by a Marietta Blau Grant from the Austrian Agency for International Cooperation in Education, Science and Research (MMC-2023-07030). MV was supported by grant 2022.04861.PTDC from the Portuguese Foundation for Science and Technology (FCT). AK is grateful for the support of the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2002/1 "Science of

Intelligence" – project number 390523135. FCT (https://www.fct.pt/en/) supported this work, through a grant to T.M. (CDL-CTTRI- 249-SGRH/2022), and multiannual funding to the WJCR in the context of the project UID/04810/2020, DOI: 10.54499/UIDB/04810/2020 and 10.54499/UIDP/04810/2020.

Data availability

All data and code will be made available upon acceptance or at reasonable request to the corresponding authors.

Author contributions

D.R.B. conceptualization, data curation, formal analysis, investigation, methodology, funding acquisition, writing— original draft, writing—review and editing.
M.V. supervision, funding acquisition, writing—review and editing.
F.R. conceptualization, supervision, writing—review and editing.
A.K. formal analysis, writing— original draft, writing—review and editing.
T.M. conceptualization, supervision, formal analysis, investigation, methodology, writing— original draft, writing—review and editing.

Acknowledgements

TM would like to thank Jeff Stevens and Robert Hampton for fruitful discussions on earlier versions of the manuscript.

Conflict of interests

The authors declare no conflict of interests.

References

Aw, J., Monteiro, T., Vasconcelos, M., & Kacelnik, A. (2012). Cognitive mechanisms of risky choice: Is there an evaluation cost?. *Behavioural processes*, *89*(2), 95-103.

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2014). Fitting Linear Mixed-Effects Models using lme4. arXiv [stat.CO]

Brooks, M. E., Kristensen, K., Van Benthem, K. J., Magnusson, A., Berg, C. W., Nielsen, A., Skaug, H. J., Mächler, M. & Bolker, B. M. (2017). glmmTMB balances speed and flexibility among packages for zero-inflated generalized linear mixed modeling.

Brunton, B. W., Botvinick, M. M., & Brody, C. D. (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science*, *340*(6128), 95-98.

Call, J. (2004). Inferences about the location of food in the great apes (Pan paniscus, Pan troglodytes, Gorilla gorilla, and Pongo pygmaeus). J. Comp. Psychol. *118*, 232–241.

Call, J. (2006). Inferences by exclusion in the great apes: the effect of age and species. *Animal cognition*, *9*(4), 393-403.

Call, J. (2022). The “avoid the empty cup” hypothesis does not explain great apes’(Gorilla gorilla, Pan paniscus, Pan troglodytes, Pongo abelii) responses in two three-cup one-item inference by exclusion tasks. *Journal of Comparative Psychology*, *136*(3), 172.

Cardinal, R.N., & Aitken, M.R.F. (2010). Whisker: A client—server high-performance multimedia research control system. Behav. Res. Methods *42*, 1059–1071.

Charnov, E. L. (1976a). Optimal foraging: attack strategy of a mantid. *The American Naturalist*, *110*(971), 141-151.

Charnov, E. L. (1976b). Optimal foraging, the marginal value theorem. *Theoretical population biology*, *9*(2), 129-136.

De Waal, F. B. (1999). Anthropomorphism and anthropodenial: consistency in our thinking about humans and other animals. *Philosophical topics*, *27*(1), 255-280.

Fox, J., Weisberg, S., Adler, D., Bates, D., Baud-Bovy, G., Ellison, S., Firth, D., Friendly, M., Gorjanc, G., Graves, S., et al. (2012). Package "car." Vienna: R Foundation for Statistical Computing *16*.

Freidin, E., Aw, J., & Kacelnik, A. (2009). Sequential and simultaneous choices: Testing the diet selection and sequential choice models. *Behavioural processes*, *80*(3), 218-223.

Gibson, J. J. *The Ecological Approach to Visual Perception: Classic Edition* (Psychology Press, 1979).

Gigerenzer, G., Todd, P. M., & ABC Research Group, T. (2000). *Simple heuristics that make us smart*. Oxford University Press.

Hanus, D., & Call, J. (2011). Chimpanzee problem-solving: contrasting the use of causal and arbitrary cues. *Animal Cognition*, *14*, 871-878.

Kacelnik, A., Vasconcelos, M., Monteiro, T., & Aw, J. (2011). Darwin's "tug-of-war" vs. starlings' "horse-racing": how adaptations for sequential encounters drive simultaneous choice. *Behavioral Ecology and Sociobiology*, *65*(3), 547-558.

Kennedy, J. S. (1992). *The new anthropomorphism*. Cambridge University Press.

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2015). Package 'lmertest'. *R package version*, *2*(0), 734.

Lauffer, M. C., Castro, L., & Wasserman, E. A. (2017). Chrysippus's pigeon: Exclusion-based responding in an avian model. *Journal of Experimental Psychology: Animal Learning and Cognition*, *43*(2), 139.

Lichtenstein S, Slovic P. *The Construction of Preference*. Cambridge University Press; 2006.

Lovibond, P. F., & Westbrook, R. F. (2024). The role of expectancy in Pavlovian conditioning. *Psychological Review*.

Macías, A., González, V. V., Machado, A., & Vasconcelos, M. (2021). The functional equivalence of two variants of the suboptimal choice task: choice proportion and response latency as measures of value. *Animal cognition*, *24*(1), 85-98.

ManyPrimates. (2024). ManyPrimates 3A - Inference by Exclusion. https://doi.org/10.17605/OSF.IO/6W2RP

Mazur, J. E. (2010). Distributed versus exclusive preference in discrete-trial choice. *Journal of Experimental Psychology: Animal Behavior Processes*, *36*(3), 321.

Mazur, J.E., & Odum, A.L. (2023). Learning and Behavior (9th ed.). Routledge. https://doi.org/10.4324/9781003215950

McFarland, D. J., & Sibly, R. M. (1975). The behavioural final common path. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, *270*(907), 265-293.

McNamara, J. M., & Houston, A. I. (2009). Integrating function and mechanism. *Trends in ecology & evolution*, *24*(12), 670-675.

Mody, S., & Carey, S. (2016). The emergence of reasoning by the disjunctive syllogism in early childhood. *Cognition*, *154*, 40-48.

Monteiro, T., Vasconcelos, M., & Kacelnik, A. (2013). Starlings uphold principles of economic rationality for delay and probability of reward. *Proceedings of the Royal Society B: Biological Sciences*, *280*(1756), 20122386.

Monteiro, T., Vasconcelos, M., & Kacelnik, A. (2020). Choosing fast and simply: Construction of preferences by starlings through parallel option valuation. *PLoS Biology*, *18*(8), e3000841.

Nagel, R., Ruxton, G. D., & Morrissey, M. B. (2024). Classical tests, linear models and their extensions for the analysis of 2× 2 contingency tables. *Methods in Ecology and Evolution*, *15*(5), 843-855.

Parker, G. A., & Stuart, R. A. (1976). Animal behavior as a strategy optimizer: evolution of resource assessment strategies and optimal emigration thresholds. *The American Naturalist*, *110*(976), 1055-1076.

R Core Team (2021). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. 2012. Preprint.

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural computation*, *20*(4), 873-922.

Reboreda, J. C., & Kacelnik, A. (1991). Risk sensitivity in starlings: variability in food amount and food delay. *Behavioral Ecology*, *2*(4), 301-308.

Rescorla, M. (2009). Chrysippus' dog as a case study in non-linguistic cognition. *The philosophy of animal minds*, 52-71.

Schneider, S. M., & Morris, E. K. (1987). A history of the term radical behaviorism: From Watson to Skinner. *The Behavior Analyst*, *10*, 27-39.

Seeley, T. D. (2011). *Honeybee democracy*. Princeton University Press.

Shapiro, M. S., Siller, S., & Kacelnik, A. (2008). Simultaneous and sequential choice as a function of reward delay and magnitude: normative, descriptive and process-based models tested in the European starling (Sturnus vulgaris). *Journal of Experimental Psychology: Animal Behavior Processes*, *34*(1), 75.

Shettleworth, S. J. (2010). Clever animals and killjoy explanations in comparative psychology. *Trends in cognitive sciences*, *14*(11), 477-481.

Slovic, P. (1995). The construction of preference. *American psychologist*, *50*(5), 364.

Vasconcelos, M., Monteiro, T., Aw, J., & Kacelnik, A. (2010). Choice in multi-alternative environments: a trial-by-trial implementation of the sequential choice model. *Behavioural processes*, *84*(1), 435-439.

Vasconcelos, M., Monteiro, T., & Kacelnik, A. (2013). Context-dependent preferences in starlings: linking ecology, foraging and choice. *PLoS One*, *8*(5), e64934.

Vasconcelos, M., Monteiro, T., & Kacelnik, A. (2015). Irrational choice and the value of information. *Scientific reports*, *5*(1), 13874.