Managing Costly Generalisation Errors in the Food Domain

Camille Rioux, Connair Joseph Stephen Russell, Annie E Wertz

To cite this version:

Camille Rioux, Connair Joseph Stephen Russell, Annie E Wertz. Managing Costly Generalisation Errors in the Food Domain. 2025. ⟨hal-05341150⟩

HAL Id: hal-05341150

https://hal.inrae.fr/hal-05341150v1

Preprint submitted on 31 Oct 2025

Managing Costly Generalisation Errors in the Food Domain

Camille Rioux[1]*, Connair J. S. Russell[1]*, & Annie E. Wertz[1]

Authors' Affiliation:

[1]Max Planck Institute for Human Development, Max Planck Research Group Naturalistic Social Cognition, Lentzeallee 94, 14195 Berlin, Germany.

Corresponding Author:

russell@mpib-berlin.mpg.de

* Shared co-first authorship. Authors’ names are in alphabetical order.

Abstract

A bias towards certain kinds of information such as danger, has been observed in a number of domains, for instance learning and attention. Under Error Management Theory (EMT) it has been argued that these biases reflect the costs of making mistakes with this kind of information (Haselton & Nettle, 2006). Based on EMT principles we reasoned that such biases may also exist in generalisation, specifically for information where generalisation mistakes are costly (edibility and toxicity information). Experiment 1 tested this in adults (N = 88), and found that participants generalised toxicity information more, and edibility information less than control information. Experiment 2 tested this in 4- to 5-year-olds (N = 91) and did not find the same pattern, but observed the only effect of information type being in those children scoring high on food neophobia. In Experiment 3 results from Exp. 1 were compared with an additional categorisation condition (N = 30) for the same stimuli, and a different pattern was observed for categorisation compared to generalisation, indicating that participants were not relying primarily on category-based generalisation. Finally, Experiment 4 looked to test whether the findings of Exp. 1 are a function of negativity and positivity effects more broadly, and tested negative and positive information without the same direct costs to an individual (polluting and sustainability information). We found some indication of a broad negativity/positivity effect in generalisation, though this appears less strong than in Exp. 1, indicating a negativity bias effect cannot fully account for these results.

1. Introduction

Inductive inference, or generalisation of the known to the unknown, is a critical component of human cognition (Feeney & Heit, 2007; Murphy, 2002). It enables us to apply learned knowledge to make predictions about novel situations or objects. However, making inductive inferences is not error-free and some mistakes may be more costly than others. Consider the problem of identifying which foods are safe to eat, particularly for an omnivorous species like humans, who need to gather a wide variety of foods to ensure nutritional health while avoid ingesting something harmful (Drewnowski et al., 1997; Rozin & Todd, 2015; van Duyn & Pivonka, 2000). When confronted with a novel food candidate, an individual must decide whether or not it is edible based on their knowledge of edible foods in the environment. Food items of the same type may vary visually (e.g., ripening fruits), and different types of foods may look remarkably similar (e.g., edible redcurrants and toxic holly berries). This means that when one encounters a novel food candidate there is a degree of categorical uncertainty and, importantly, the consequences of making errors in generalisation are not uniform. For example, erroneously choosing to eat a novel fruit that is in fact toxic is generally more costly than erroneously choosing to forgo a novel fruit that is in fact edible. This asymmetry of errors is a core tenet of Error Management Theory (EMT; Haselton & Buss, 2000; Haselton & Nettle, 2006; Johnson et al., 2013) and here we explore the effects of this asymmetry on generalisation decisions in the food domain.

Generalisation is a complex process, and a number of factors may influence whether or not one generalises. Research on inductive inference has demonstrated that similarity is a key factor at play when deciding whether to generalise a specific property from a known exemplar to a novel item(Gelman & Markman, 1987; Murphy, 2002; Sloutsky & Fisher, 2004b). More formally, it has been argued that generalisation is a declining function of similarity, with generalisation declining as a novel item becomes less like a known exemplar (Shepard, 1958; Sloutsky, 2003). The ability to extrapolate known knowledge to novel items develops early in ontogeny, with infants as young as 9 months of age generalising nonobvious properties to novel items based on the visual appearance of a single known exemplar (Baldwin et al., 1993a; Welder et al., 2016). In certain domains, some features seem to be more important than others in judging to what extent two items are similar. For instance, in the food domain, colour similarity seems to play a key role in both children's and adults' decisions to generalise a property to a novel item (Lavin & Hall, 2002; Macario, 1991; Rumiati & Foroni, 2016, but see Rioux et al., 2018 for the importance of shape similarity as well), while shape similarity seems to be more important in the artefact domain (Landau et al., 1998; Welder et al., 2016).

Whilst the role of perceptual similarity appears clear, it has also been argued that category membership plays an important role in generalisation. Category-based induction is observed early in life, for example, 4-year-old children generalise a known property of an owl to a flamingo because both are from the same category “bird” (Gelman & Markman, 1986, 1987; Mandler & McDonough, 1993). Categorisation appears to develop as early in ontogeny as 3 months (Ferry et al., 2010; Quinn et al., 1993), young children are able to flexibly deploy different types of categories in the food domain (e.g., script, taxonomic and evaluative categories;Nguyen, 2020; Nguyen & Murphy, 2003), and selectively generalise a property depending of its type (e.g., incidental, functional or biological; Gelman, 1988; Nguyen, 2012; Thibaut et al., 2016). However, how such inductive processes function under categorical uncertainty—that is, when the category membership of the novel item is unclear—has received far less attention, with a few notable exceptions (Chen et al., 2014; Hayes et al., 2008; Hayes & Newell, 2009). Moreover, the question of whether humans engage in selective generalisation based on the type of information—that is, whether they generalise knowledge from known to novel instances differently depending on the content of that knowledge—does not appear to have been studied. In some areas, errors in generalisation decisions may have asymmetric costs. For example, incorrectly generalising edibility to something toxic is more costly than incorrectly generalising toxicity to something edible. In line with Error Management Theory (EMT; Haselton & Buss, 2000; Haselton & Nettle, 2006) when there are asymmetries in the costs of generalisation errors, one might expect to see selective generalisation of information. EMT posits that when the costs of false-positives (e.g., needlessly forgoing an edible fruit) and false-negatives (e.g., inadvertently eating a toxic fruit) have been asymmetric over phylogeny, cognitive systems should have arisen that reduce the rate of costly errors. These kinds of systems can evolve even if reducing costly errors results in more errors overall (e.g., one might prefer a smoke alarm that has greater sensitivity, and thus a greater number of overall errors, but reduced probability of a highly costly false-negative; Nesse, 2005).

It is not yet clear whether humans engage in selective generalisation depending on the type of information, but a growing body of evidence suggests that content selectivity is present in another crucial component of human cognition—learning (e.g., Barrett et al., 2016; Barrett & Broesch, 2012; Broesch et al., 2014; DeLoache & LoBue, 2009; Wertz & Wynn, 2019; Wertz & Wynn, 2014a). In fact, decades of evidence suggests that several species have selective learning mechanisms, such that certain evolutionarily relevant information and associations are learned more readily (e.g. Öhman & Mineka, 2001; Seligman, 1970; Seligman, 1971). Given the broad possibility space of what one could learn in any given environment, such selective learning systems could facilitate the learning of information most relevant to successful survival and reproduction. Indeed, evidence from experiments with fruit flies and computational models have shown environments in which preparedness can develop (Dunlap & Stephens, 2014; Lindström et al., 2016). In their seminal example of preparedness in animal learning, Garcia & Koelling (1966) demonstrated that rats learned to associate ingestion of a novel food with the experience of nausea in a single trial, but not with the experience of bright lights or novel sounds, even after many trials. These findings indicate that, at least in rats, certain ecologically recurrent associations (food causing nausea) are more easily learned. Similarly, monkeys have been shown to learn fear responses to snakes more readily than other non-dangerous stimuli like flowers (see Öhman & Mineka, 2001, 2003, for reviews).

Evidence for selective learning of different types of fitness-relevant information has also been found in humans. Human infants tend to selectively associate snakes, but not flowers or inoffensive animals such as frogs, with fear (DeLoache & LoBue, 2009; see LoBue & Rakison, 2013 for a review). Adults also appear to show preferential learning of fear associations for snakes and spiders (see Öhman & Mineka, 2001 for a review), including reduced habituation (Öhman et al., 1974) and non-conscious conditioning to fear-relevant but not fear-irrelevant stimuli (Öhman & Soares, 1998). Moreover, infants selectively associate edibility with plants compared to non-plants (Wertz & Wynn, 2014a). Children also appear to show preferential learning of danger information about animals, and their performance is similar across cultures(Barrett & Broesch, 2012). For example, children from Los Angeles, USA and from Indigenous Shuar communities in the Ecuadorian Amazon showed heightened recall of information about dangerous animals, compared to danger information about other stimuli (foods, artefacts), and other information (location, feature), and this advantage persisted with only minor attenuation a week later (Barrett et al., 2016). This preferential memory for danger information appears to reduce with age and is not present in adulthood (Broesch et al., 2014), suggesting differences across the lifespan in selective learning.

Given this evidence demonstrating that learning appears to be selective towards evolutionarily relevant information, the present study looks to examine whether this extends to how that information is used. That is, is there a selectivity in inductive inferences of evolutionarily relevant information? Specifically, we aimed to investigate how the type of information learned influences the generalisation of that information to novel stimuli, notably when the category membership of the novel stimulus is uncertain. Whilst the effect of information type on the degree of generalisation in humans has not been tested, some limited evidence suggests people do pay attention to such information type when generalising (Zhu & Murphy, 2013), and it has been found that rats would generalise a shock more broadly than a food reward (Hearst, 1962; Murray & Miller, 1952, though see Hoffman & Fleshler, 1963 for the opposite pattern), demonstrating, in rats at least, that generalisation is sensitive to what is being generalised. Our predictions were based on Error Management Theory, which, as outlined above, posits that selectivity in cognitive processes arises when there is an asymmetry in the cost of errors (Haselton et al., 2005; Haselton & Buss, 2000; Haselton & Nettle, 2006; Johnson et al., 2013). As a first step, we chose to examine selective generalisation about candidate food items. Decisions about what can (or cannot) be eaten are highly fitness-relevant and a single error when deciding whether or not to consume a novel food candidate can lead to severe negative consequences (e.g., poisoning; Hagen et al., 2009; Mithöfer & Boland, 2012). Importantly, the costs of different types of errors are asymmetric. Specifically, if you learn a food item is toxic you can make two generalisation errors: (1) you can incorrectly generalise this danger to another candidate food item that is in fact not dangerous (false-positive), or (2) you can fail to generalise this danger to another candidate food item that is in fact toxic (false-negative). The cost of the false-positive might be a missed dinner, but the cost of a false negative might be one’s life or a serious illness due to poisoning. Therefore, we predicted that participants would exhibit different generalisation patterns when presented with an ambiguous candidate food item depending on the type of information being generalised. In the present research we tested the generalisation of information about a candidate food's edibility or toxicity. Specifically, we predicted that, to avoid costly errors, humans will generalise toxicity information more broadly and edibility information more narrowly than neutral information.

To examine the effect of differing information on generalisation we tested adults and children in a series of four experiments. Participants performed a classical induction task in which they were asked questions about a series of 20 image triads (after e.g.,Gelman & Markman, 1986) of different fruits. Each triad consisted of two target images at the top and the test image below. Each test image was one stage of a five-stage transformation between the two target images (see Figure 1). Participants were told that one of the target images had a property (e.g., “This is edible”) and the other did not (e.g., “This is *not* edible”), and then were asked whether the test image also had that property (e.g., to select “This is edible" or "This is not edible”). Adults (Experiment 1) and children (Experiment 2) were asked to generalise three types of properties (across three independent conditions), edibility (positive condition), toxicity (negative condition) or seasonality (neutral condition). In these experiments, we predicted different generalisation patterns for edibility and toxicity information relative to neutral or control information. Specifically, we predicted that, to avoid costly errors, humans will generalise toxicity information more broadly and edibility information more narrowly than neutral information. We tested both children and adults in Experiments 1 and 2 as previous literature on content-specific learning mechanisms has revealed differences across the lifespan and mixed evidence in adult populations (e.g., Broesch et al., 2014; Öhman & Mineka, 2001). For the child age-group, we chose 4- to 5-year-olds because previous studies found that by this age children already selectively generalise properties based on their type (Gelman, 1988; Nguyen, 2012; Thibaut et al., 2016), and can use both visual similarity and category membership in inductive inference (Gelman & Davidson, 2013).

We ran two additional experiments to examine the role of categorisation (Experiment 3) and the possible effect of broader negativity or positivity biases (Experiment 4) in selective generalisation processes. In Experiment 3, adults were asked to categorise the same stimuli used in Experiments 1 and 2 based on name labels (e.g., "This is a Daxon/Sudeb") in order to measure their category membership judgements. This allowed us to examine whether participants form categories that guide their generalisation decisions or if they rely on similarity judgements alone. Finally, to test whether any effects of edibility or toxicity are a function of broader negativity or positivity biases, Experiment 4 looked the effects of negative and positive information more broadly on generalisation. In this experiment adults were asked to generalise three different properties (sustainability, polluting, storage instructions), similar in valence to those used in Experiment 1 (edibility, toxicity, season of growth), but which are not fitness-relevant properties and do not offer the same asymmetric costs and benefits as edibility and toxicity judgements.

2. Experiment 1

Experiment 1 examines whether type of information one learns about a stimulus influences how that information is generalised to similar stimuli. As outlined in more detail above, human learning systems seem to be particularly sensitive to information about danger (Barrett et al., 2016; Barrett & Broesch, 2012; Broesch et al., 2014; DeLoache & LoBue, 2009; LoBue et al., 2010; LoBue & DeLoache, 2008; New & German, 2015). In particular, it has been argued that biases in learning systems may reflect asymmetries in the costs of making errors about dangerous stimuli over evolutionary time (Haselton & Nettle, 2006). We reasoned that the selection pressures which give rise to these specialized learning systems may also have led to specialization in generalisation. We examined this in the domain of food because of the important role they play in human sustenance, the many risks they pose

(Belovsky & Schmitz, 2006; Hagen et al., 2009; Hanley et al., 2007; Keeler & Tu, 1983; Manners, 1996; Mithöfer & Boland, 2012; Wertz, 2019), and an apparent psychological sensitivity to information about this domain (Barrett et al., 2016; Krasnow et al., 2011; Wertz, 2019; Wertz & Wynn, 2019; Wertz & Wynn, 2014b, 2014a). Given the asymmetry in costs in generalisation, mistakes in these domains (i.e., generalising edibility to something toxic vs. generalising toxicity to something edible), we predicted that information about danger (i.e., toxicity) would be generalised more broadly, and information about edibility less broadly than neutral information. For the neutral seasonality information (whether the item grows in summer or not) was selected as a piece of information consistent with the domain of food that does not necessarily have the same cost of generalisation errors.

2.1. Method

Participants were 88 adults (43 females, 25 males, mean age = 33 years, range = 17 – 64, due to an error in the demographic survey 21 participants did not provide demographic information), divided evenly across three independent information conditions (edibility, toxicity, and seasonality), who were recruited via the online *Prolific.ac* platform. Data from two participants (in the edibility condition) were excluded for meeting our a priori exclusion criteria of giving the same response to at least 19 out of 20 questions (e.g. stating for 19 of 20 items “This is edible” for those in the edibility condition). Adult participation was restricted to those based in the United Kingdom and United States in order to limit the geographical range and thereby familiarity with the fruits from other regions that were used in the study. The study was approved by a local ethics committee and participants were compensated with 2.5 GBP (rate of 10 GBP/ hour).

The materials for the study included four image pairs of real fruits, and a spectrum of five morphs between each pair, created with *Abrosoft: Fanatmorph* software version 5.4.8.

We chose images of fruits that are not typically found in the UK and the US to ensure participants did not draw on prior knowledge to answer the task. There were originally nine different pairs of real fruit images and their respective morphs, which were piloted online with 30 English and American adults. The final four image pairs selected were those adults in the pilot rated least familiar and that had the most consistent similarity gradients across the morph sequence (i.e., the difference between morphs one and two were rated similarly to the difference between four and five, see Supplemental Material section 1.1). In addition, we chose this food category because fruits in their raw state can either be edible and nutritious, or they may contain secondary toxic components harmful to humans (e.g., Mithöfer & Boland, 2012). The morph sequences are shown in Figure 1 (left panel). The original images of the real fruits are shown at each end, and the five stages of the objects morphing into one another between. The morphs are at approximate degrees of similarity to the left real fruit picture of 83%, 67%, 50%, 33% and 17% (e.g., first row: a guava fruit transformed into a star fruit).

Image triads were created using *Affinity: Photo* software version 1.7.2.471, by presenting a final image pair at the top of the card (target pictures) alongside one level of their respective image morphs at the bottom of the card (test picture). This resulted in 20 triads (five morphs for each of the four pairs; see Fig. 1, right panel). The image triads were presented with the relevant information noted in text below each image. In each triad, one target had a property (the left top picture) and the other did not (the right top picture). Depending on the condition this was ``This is edible” / “This is not edible" (edibility condition, see Fig. 1 right panel), ``This is toxic” / “This is not toxic" (toxicity condition), or ``This grows in summer” / “This does not grow in summer" (seasonality condition). Which fruit of each pair was on the left varied across participants, with the item on the left always having the property; therefore, the fruit-property correspondence was counterbalanced across participants. The different pictures were never named, but referred to only as “this” to keep category membership ambiguous.

Figure 1. Stimuli used in all experiments. Left Panel depicts the four morph sequences, with (a) and (c) showing the original images and (b) showing the five-point morph sequence between each pair. Right Panel (d) depicts an example trial from the edibility condition; the images on the top of the triad are from (a) and (c) and the image on the bottom is from (b).

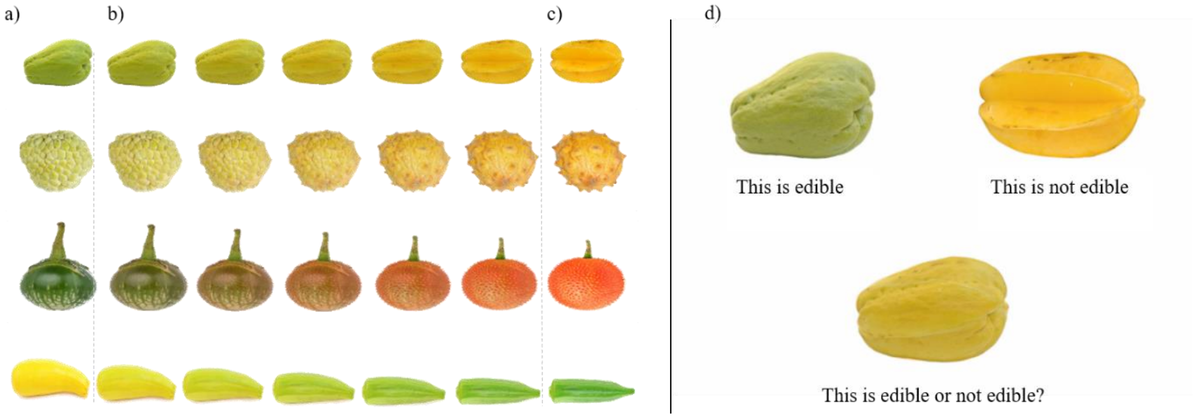

Participants completed the study online through the browser-based *Qualtrics* platform. To ensure relative parity across participants, participation was only possible on desktop browsers—not mobile phones or tablets. After reviewing the study information, and providing consent for participation and the use and storage of their data, participants were presented with the 20 image triads. The triads were presented in a pseudo-random, counterbalanced order such that no two triads from the same fruit pair would be presented back-to-back. This was done to prevent direct comparison between two morphs from the same morph sequence. For each of the 20 triads we recorded whether participants generalised the property to the test picture (morph picture) by asking explicitly whether the test picture had the property or not (see Fig. 1 right panel). A score of 1 was given when participants generalised the property to the test picture and a score 0 when they did not.

2.2. Results and Discussion

Statistical analyses were conducted using R 3. 6.1 (R Core Team, 2019), with main analyses completed using the lme4 package. Data and R scripts are available on the Open Science Framework page for this project: https://osf.io/smcy7/?view_only=7b12be167f484d5b8e3c626680f21fde.

A Binomial generalised linear mixed effect model (GLMM) with a logit link function was used to assess generalisation patterns across the 20 triads. Preliminary analysis found no improved model fit of different image pairs modelled as a random factor in the model. We also found no improved fit including sex and age as fixed factors in the model. Therefore, these variables were left out for further analysis. In the main analysis both participants and individual triads served as a random factor to account for shared variances within subjects and inter-triad variability. Similarity between the target picture that had the property and the test picture (83%, 67%, 50%, 33% and 17%), Condition (edibility, toxicity and seasonality information), as well as their interaction were modelled as fixed effects. The best fit model was selected according to Akaike's Information Criterion and was the full model including the main effects and the interaction between Similarity and Condition ($\chi2$ (2) = 18.51, $p < 0.001$, marginal R2 = 0.57, conditional R2 = 0.63).

In line with previous findings (Gelman & Markman, 1987; Sloutsky & Fisher, 2004b), effect of Similarity between the target picture that had the property and the test picture was significant (χ^2 (1) = 131.87, $p < 0.001$, see Table S4 in the SM section 2.1), with likelihood of property generalisation decreasing with similarity (see Fig. 2). This supports the established finding that similarity effects generalisation (e.g., Feeney & Heit, 2007; Murphy, 2002; Shepard, 1958; Sloutsky, 2003) and demonstrates our stimuli to successfully elicit this effect. As predicted, we also found a main effect of Condition (χ^2 (2) = 6.65, p = 0.036, see Table S4 in the SM) which was qualified by the interaction effect of Similarity*Condition (χ^2 (2) = 18.13, p < 0.001, see Fig. 2, see Table S4 in the SM), with the decrease of property generalisation with similarity in the seasonality condition being significantly steeper than in the edibility (z = 4.04, p < 0.001) and toxicity conditions (z = 2.60, p = 0.025).

To examine the interaction further we looked at generalisation rates at each point of our morph sequence (i.e., each level of similarity). This was done using Fisher's Exact Tests. Multiple comparisons were controlled for using the Bonferroni-Holm correction within each similarity level. In the *least* ambiguous morphs, we found a significant pairwise difference between the edibility and seasonality conditions (morph 1 (the most similar), p = 0.048, morph 5 (least similar), p = 0.024), with participants generalising edibility information at a higher rate in morph 1 and at a lower rate in morph 5 (see Fig. 2). For morph 5, a significant difference was also observed between the toxicity and seasonality conditions (p = 0.033), with participants generalising toxicity information at a lower rate (see Fig. 2). Finally, with the *most* ambiguous morph—morph 3, the halfway point in the morph sequence between the two original real food images (i.e., when similarity between the target that had the property and the test picture was 50%)— we found a significant pairwise difference between the edibility and toxicity conditions (morph 3, p = 0.003), with participants generalising toxicity information at a higher rate (see Fig. 2). As predicted, participants were more likely to generalise toxicity and less likely to generalise edibility to an ambiguous food candidate, yet this appears to only be the case in situations with greater uncertainty. That is, when told one item is toxic and another is not toxic, and then presented with a test item that is 50% similar to both, participants judge this item to be toxic at a rate significantly greater than chance (binomial P = 0.60 [95% CI: 0.51 – 0.69], p = 0.035). In the case of edibility, we see the opposite pattern, participants state this item to be edible at a rate significantly below chance (binomial P = 0.38 [95% CI: 0.29 – 0.48], p = .017), and in the neutral condition participants stated the item grew in summer at a rate no different from chance (binomial P = 0.49 [95% CI: 0.40 – 0.58], p = 0.93).

Figure 2: Generalisation scores in Experiment 1 at each point of the morph sequence (i.e., each level of Similarity), depending of Condition. Large dots show group means with 95% CI error bars,

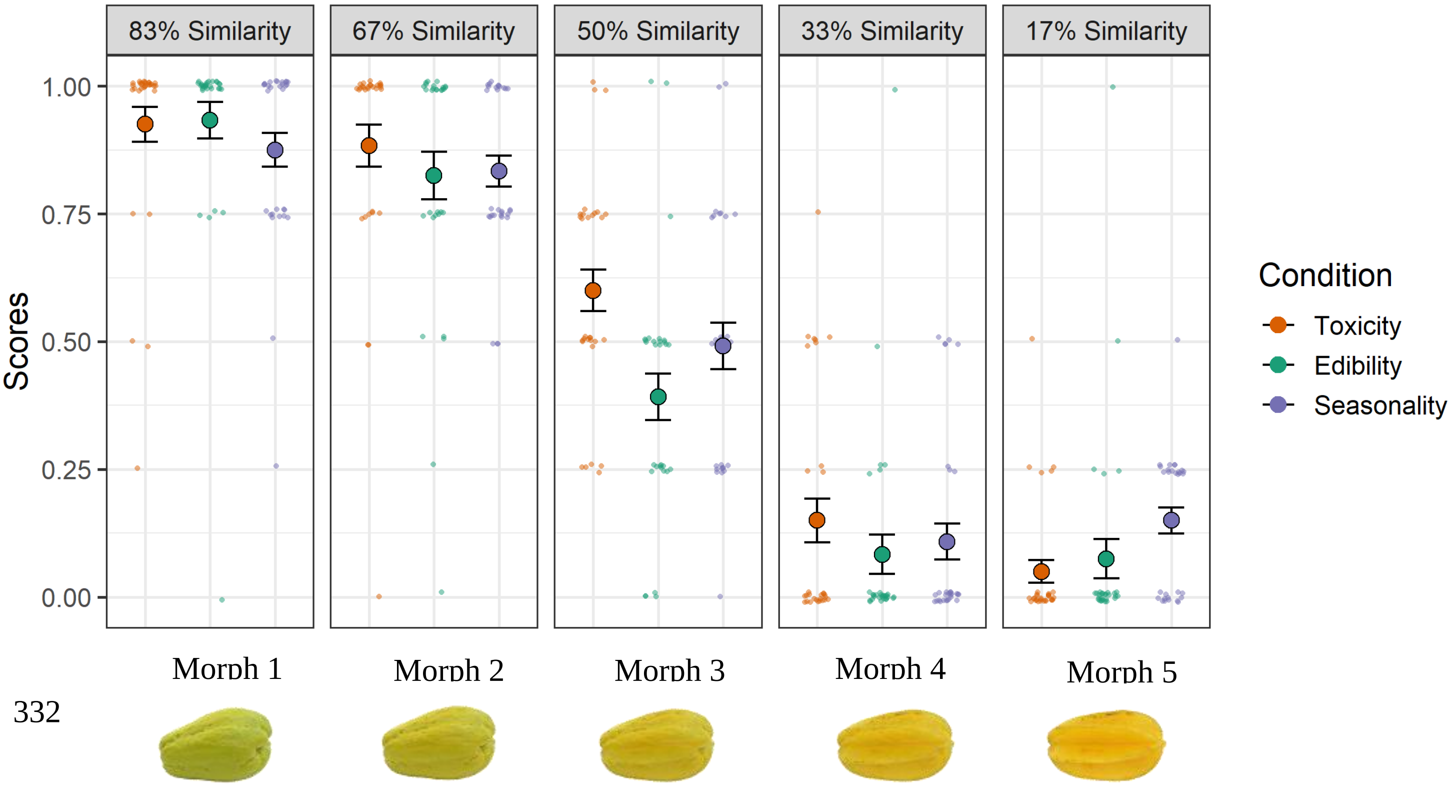

Finally, in order to examine participants' response stability, we analysed participants' response patterns across the different stimulus pairs. More than 90% participants exhibited a stable strategy (responding the same for 3 of the 4 image pairs) for the 4 least ambiguous morphs, but only 63% exhibited a stable strategy for the most ambiguous morph indicating greater choice uncertainty at this level of similarity (see SM Section 2.2 for details).

The results of Experiment 1 support the hypothesis that adults use selective generalisation strategies. Specifically, as predicted, adults overgeneralised toxicity information and undergeneralised edibility information, relative to a neutral information. Notably, our findings suggest that such selective generalisation strategies are deployed in situations with greater uncertainty. Participants generalised differently across the information conditions only for the most ambiguous morph.

3. Experiment 2

Previous studies have demonstrated selective learning mechanisms for information about danger and edibility (Barrett et al., 2016; Barrett & Broesch, 2012; Wertz & Wynn, 2014a). In Experiment 1 we observed similar selectivity in generalisation, with toxicity (i.e., danger) information being generalised more broadly, and edibility more narrowly, than neutral information. However, the literature on specialized learning for danger information has revealed differences across the lifespan, with this effect diminishing with age and not being present in adulthood (Broesch et al., 2014). Further, induction abilities develop early in ontogeny with developmental changes (Baldwin et al., 1993; Gelman & Markman, 1986; Welder et al., 2016). Therefore, we aimed to examine whether there were developmental differences in the type of selective generalisation we found in Experiment 1. To examine if selective generalisation based on information type (toxicity, edibility) is also present early in life, and to compare any effect with that found in adults, we ran an adapted version of the task from Experiment 1 with preschool age children.

3.1. Method

Participants were 91 preschoolers between 4 and 5 years old (41 female, 50 male, mean age = 4 years, 11-months, range = 4 years, 0 months – 5 years, 11 months), divided evenly across three independent information conditions (edibility, toxicity, and seasonality). Children were recruited and tested at kindergartens, museums, or a research lab in or around Berlin, Germany. Due to the 2020 Covid-19 pandemic, in-person data collection had to conclude prematurely and the final 14 children were tested online with stimuli presented via video chat with the software *WebEx*. These children were recruited from the institution's participant database . Participants were compensated with a participation certificate. Data from 26 participants were excluded for meeting the a priori exclusion criteria of giving the same response to at least 19 out of 20 questions, missing 5 successive answers, failing the initial task example trials, or failing to complete the session due to inattentiveness. This study was reviewed and accepted by a local ethics committee. Written parental consent was obtained prior to testing, and children were first asked by the experimenter if they wanted to play a game. Upon providing verbal consent, the children who were tested in person sat at a table across from the experimenters in a quiet area. Two experimenters were present, one who ran the task, and the other who recorded the child’s responses. The in-person testing sessions were also filmed to allow checking of any discrepancies or missing reporting from the live sessions

The study began with a cover story, where the participant was told that they will be shown some foods from an alien planet and they will help sort them. Then, in a warm-up phase, children were presented with two practice trials. In these practice trials, children were shown triads of fruit images, like those used in the test phase, with two images at the top. Participants were told that one of the two images at the top had a property (the property differed depending on condition – see next paragraph for details) and asked whether the third image had the property or not. However, in the practice trials, the image at the bottom of the triad was identical to one of the two above so there was a clear *correct* and *incorrect* answer, which enabled us to assess children's understanding of the task. Data for participants who did not exhibit understanding in this phase was not included in the analysis.

For the test phase, children were presented with the 20 test triads (similar to Experiment 1). The experimenter showed the child the card and pointed at each of the two target images at the top of the triad and said whether it had or did not have a particular property. We used the same three properties as in Experiment 1: edibility, toxicity, and seasonality. The wording was adapted to be appropriate for preschool-aged children and the testing was conducted in German. In the edibility condition children were told "*Look! This is edible*" [German: "*Schau mal! Das hier ist essbar*"]. In the toxicity condition, "*Look! This is toxic*" [German: " *Schau mal! Das hier ist giftig*"] and in the seasonality condition "*Look! This grows in summer*" [German: " *Schau mal!Das hier wächst im Sommer*"]. For example, a participant in the edibility condition was shown an image triad, the experimenter pointed to one target image and said "Look! This is edible" [German: "*Schau mal! Das hier ist essbar*"], then pointed to the other and said "Look! This is not edible" [German: "*Das hier ist nicht essbar*"], and then asked the child if the test picture at the bottom was edible like the first target, or not edible like the second "Now it's your turn! Do you think this {pointing to the third picture} is edible like this {pointing to edible picture} or not edible like this {pointing to the not edible picture}? [German: *Jetzt bist du dran! Glaubst du das hier* {pointing to the third picture*} ist essbar wie das hier* {pointing to edible picture} *oder nicht essbar wie das hier* {pointing to not-edible picture}? In the online version of the task, participants saw the triads presented on shared PowerPoint slides, alongside a video stream of the experimenter, and the experimenter gestured to the foods with a red cursor.

In order to examine possible individual differences parents were also given a German translation of the Children Food Rejection Score (CFRS) questionnaire developed by Rioux et al. (2017) because previous work has shown that food rejection behaviours influence generalisation in the food domain (Rioux et al., 2018). Parents were also given a questionnaire about their child's familiarity with the fruits used in the present experiment and their answers confirmed that the fruits were unfamiliar to children (on average they already saw or ate only 1 fruit out of 8).

3.2. Results and Discussion

As with Experiment 1, the results from Experiment 2 were analysed using a Binomial GLMM with a logit link function. We found no improved model fit of including image pair, age or sex in the model. Due to the variation in testing location, we also included testing location as a random factor and found no improved model fit, and as such all these variables were excluded from further analysis. In order to examine the effect of neophobia we divided children into two groups—High-CFRS and Low-CFRS—using the average CFRS scores as the cut-off point (M = 35.72, SD = 5.94; no CFRS mean difference across experimental conditions). In the main analysis, Participant and Individual Triads served as random factors. Similarity between the left target picture and the test picture (83%, 67%, 50%, 33% and 17%), Condition (label, neutral property), CFRS group (High, Low) as well as their interaction were modelled as fixed effects. The best fit model was the full model according to Akaike’s Information Criterion ($\chi^2(11) = 103.07$, $p < 0.001$, marginal $R^2 = 0.46$, conditional $R^2 = 0.55$).

As in Experiment 1, and concurrent with previous research (e.g., Murphy, 2002; Shepard, 1958b; Sloutsky, 2003), we found a significant effect of Similarity (χ^2 (1) = 304.27, $p < 0.001$, see Fig. 3, see Table S5 in the SM section 2.1), demonstrating that as with adults, the likelihood of property generalisation is decreasing with similarity. We also found a significant Similarity*Condition interaction (χ^2 (2) = 6.32, $p = 0.042$, see Table S5 in the SM), which was qualified by a significant three-way interaction between

Similarity*Condition*CFRS group (χ^2 (2) = 17.66, p < 0.001, see Fig. 3, see Table S5 in the SM). In the High-CFRS group, the decrease of property generalisation with similarity in the seasonality condition was significantly steeper than in the toxicity condition (*z* = 2.59, *p* = 0.046). In the Low-CFRS group, the decrease of property generalisation with similarity in the seasonality condition was significantly more gradual than in the toxicity (*z* = 4.02, *p* < 0.001) and edibility conditions (*z* = 3.56, *p* = 0.0020).

As with the analysis of Experiment 1, we also conducted Fisher's Exact Tests to examine the 3-way interaction further. Specifically, in both CFRS groups we looked at generalisation rates at each point of our morph sequence, depending on the condition. The only significant differences were found in the two most similar morphs (i.e., morphs 1 and 2, see Fig. 3), with children high in food rejection scores (High CFRS group) generalising toxicity information at a higher rate overall, while children low in food rejection scores (Low CFRS group) generalised seasonality information at a higher rate overall (See SM section 2.3 for further details). Therefore, the three-way interaction noted above seems to be driven by the effect of information type for the morphs most similar to the target.

Figure 3: Generalisation scores in Experiment 2, at each point of the morph sequence (i.e., each level of similarity), depending on Condition and CFRS group. Large dots show group means with 95% CI error bars, smaller dots show individual participant means (minimally jittered for clarity). Example morph sequence for example triad after Fig. 1, where the real green fruit has the property.

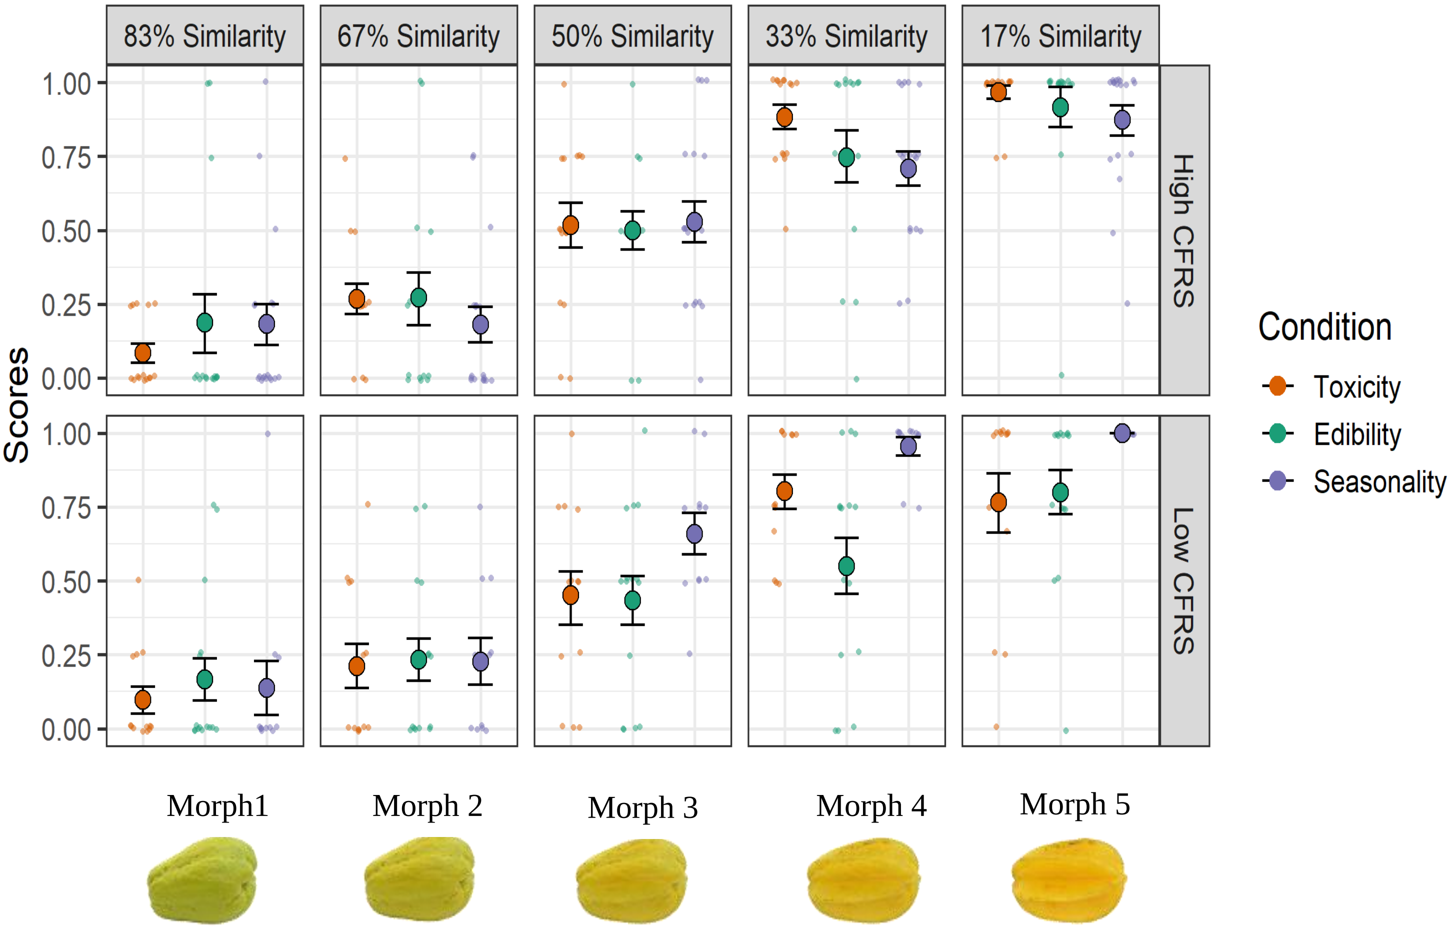

The results of Experiment 2 do not provide strong evidence that children use selective generalisation strategies consistent with minimizing costly errors. Whilst an effect of condition and similarity was observed, it did not follow the prediction of broader generalisation of toxicity information and narrower generalisation of edibility information relative to neutral information that was observed in adults in Experiment 1. Nevertheless, we found that children high in food neophobia appear to generalise toxicity at a higher rate. This suggests individual differences in the degree to which evolutionarily relevant information impacts generalisation, a finding that is consistent with the view that neophobia is a protective strategy against the risk of ingesting potentially toxic foods (Cashdan, 1994; Foinant et al., 2021; Lafraire et al., 2016; Rozin & Todd, 2015). Moreover, our results are in contrast with the lifespan findings from Broesch et al. (2014), who found preferential recall of danger information in children, but not in adults. Our results suggest that selectivity in generalisation seems to be limited to adulthood. Therefore, we explored the effect from Experiment 1 in adults in more detail in Experiments 3 and 4.

4. Experiment 3

The results from Experiment 1 suggest that, when confronted with ambiguous stimuli, adults use a cautious generalisation strategy in which they overgeneralise toxicity and undergeneralise edibility when asked about ambiguous novel candidate foods. One outstanding question is the role of categorisation in this generalisation process. Specifically, to what extent do participants form categories of the stimuli presented and use this categorical information in their responses? That is, are adult participants forming categorical judgements of the stimuli and using these to decide whether an item has or does not have a property, or are they relying solely on perceptual similarity? In the former case, the results from

Experiment 1 would mean that adults engage in selective generalisation when they are not able to assign category membership to ambiguous stimuli.

It is possible that participants in Experiment 1 formed categorical judgements about the target images, and then evaluated whether or not the test (morph) image had the shared property based on whether they considered the test morph to be a member of the relevant category. Alternatively, it could be the case that participants relied on similarity judgements alone to make this decision and did not form categorical judgements. In order to distinguish between these possibilities, we ran an additional experiment using a task similar to that in Experiment 1, but that tested whether or not participants engage in categorisation. Specifically, we examined categorisation by giving the stimuli from Experiment 1 labels (e.g., “This is a Daxon”) instead of properties (e.g., “This is edible”), thereby stating what a stimulus is, not a property it has. Such use of labels is an established way of inducing categorisation (e.g., Deng & Sloutsky, 2013; Fisher et al., 2011, 2015). The comparison of this new naming condition in Experiment 3 with the existing neutral seasonality condition of Experiment 1 allows us to examine the extent to which performance in the generalisation task (Experiment 1) is a function of categorisation. We used the neutral condition from Experiment 1 because the artificial name labels used in the Experiment 3 have no inherent costs associated (as is the case with edibility and toxicity information used in Experiment 1). Therefore, this provides a comparison of generalisation of neutral information with categorisation based upon neutral information. If participants in Experiment 1 did indeed base their decisions on categorisation, we would expect similar performance in generalisation of neutral information (seasonality information, Experiment 1) and categorisation of neutral information (Experiment 3).

4.1. Method

Participants were 30 adults (19 female, 11 male, mean age = 33 years, range = 20 – 59). This study was reviewed and approved by a local ethics committee. Recruitment and consent procedure were identical to Experiment 1. The stimulus pictures and procedure were identical to the Experiment 1, except that instead of assigning a property (edibility, toxicity, seasonality) to the target pictures, they were given a name. The names were novel labels based upon or taken from previous research using novel words and labels (e.g., Fusaro & Harris, 2013; Gaskell & Dumay, 2003). These words were selected to be plausible for the testing language (English) and were all disyllabic (See SM, section 1.2 for the list of names used).

In each trial of the task, an image triad was presented on the screen and one target fruit was given one name (e.g., Daxon for the left top picture), the other target fruit a different name (e.g., Sudeb for the right top picture). Participants were asked which of the two names the test fruit had (e.g., This is a Daxon or a Sudeb?). Randomisation and counterbalancing was the same as for the conditions used in Experiment 1. In order for data to be comparable with Experiment 1, a response was coded as 1 (yes) when participants categorised the target as the stimuli on the top left (e.g., Daxon) and a score 0 for right top target picture (e.g., Sudeb).

4.2. Results and Discussion

Categorisation scores were modelled in Binomial GLMM with a logit link function and were modelled alongside data from the seasonality condition from Experiment 1. As before, there was no model improvement from including sex and age as fixed factors and image pair as random factor, so these variables were excluded from further analysis. In the main analysis, Participant and Individual Triads served as random factors. Similarity between the left target picture and the test picture (83%, 67%, 50%, 33% and 17%), Condition (label, neutral property), as well as their interaction were modelled as fixed effects. The best fit model according to Akaike's Information Criterion was the full model including the main effects and the interaction between Similarity and Condition (χ^2 (1) = 52.27, $p < 0.001$, marginal $R^2 = 0.69$, conditional $R^2 = 0.76$).

As in Experiment 1, there was a significant effect of similarity (χ^2 (1) = 66.73, $p < 0.001$, see Table S4 in the SM section 2.1), indicating categorisation (naming condition, Experiment 3)/generalisation (neutral property, Experiment 1) declined with decreasing similarity (see Fig. 4). We also found a significant interaction effect of Similarity*Condition (χ^2 (1) = 34.56, $p < 0.001$, see Table S4 in the SM), indicating that, with decreasing similarity, the decrease of neutral property generalisation differed significantly from the decrease in categorisation ($z = 5.88$, $p < 0.001$). Fisher's Exact Tests at each stage of the morph spectrum demonstrate a statistically significant difference between the naming (Experiment 3) and seasonality information (Experiment 1) conditions in all morphs (all $p < 0.017$) except the most ambiguous morph, morph 3 ($p = .90$). When similarity between the target picture and the test picture was more than 50% (morph 1 and 2), the response rate was higher for naming (categorisation) relative to seasonality (generalisation, see Fig. 4). We observed the opposite pattern when similarity was lower than 50% (morph 4 and 5, see Fig. 4).

Figure 4: Generalisation scores in Experiment 1 seasonality condition and categorisation scores in Experiment 3, at each point of the morph sequence. Large dots show group means with 95% CI error bars, smaller dots show individual participant means. Morph images from an example triad after Fig. 1, the real green fruit has the property.

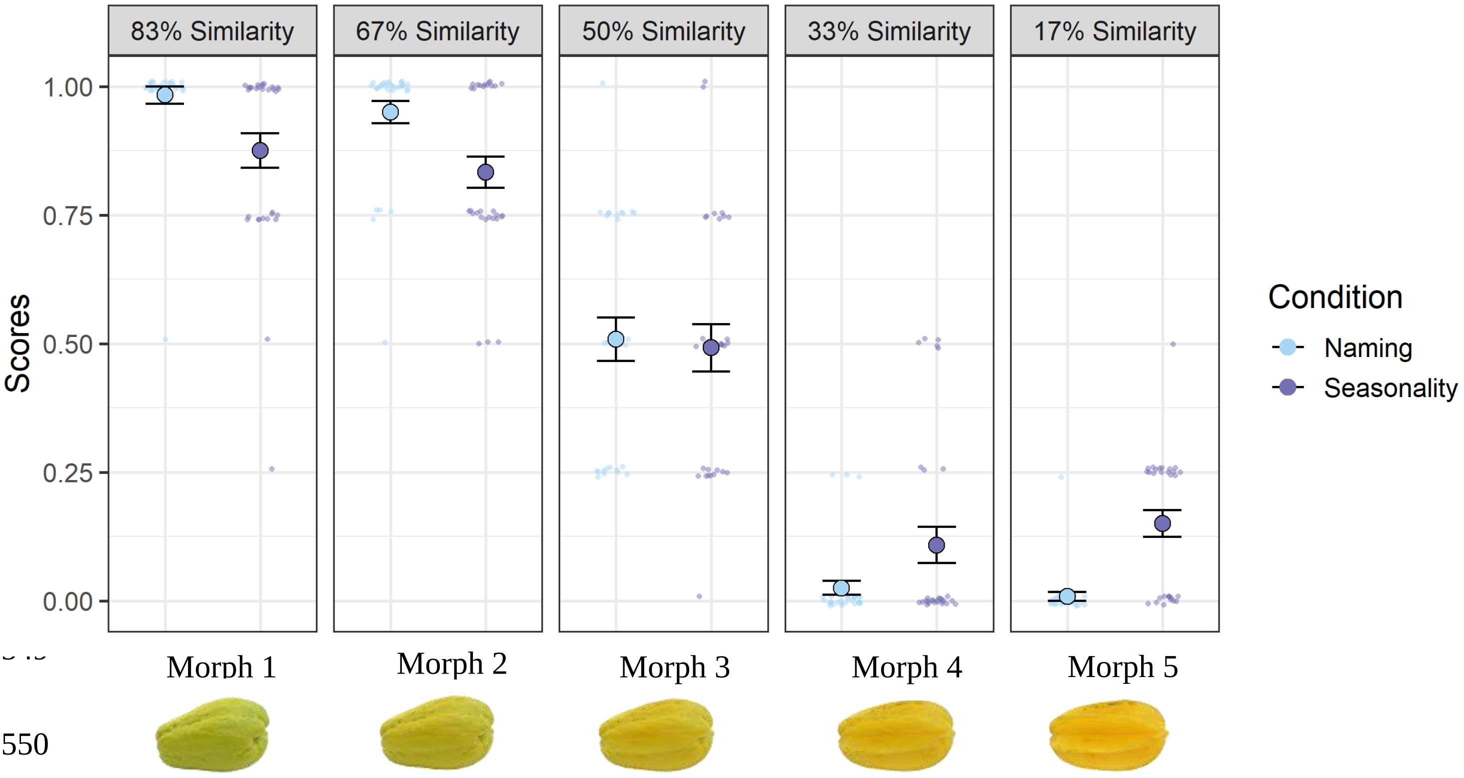

Finally, for all morphs, the percentage of stable responders was high (ranging from 73% for the most ambiguous morph and more than 96% for the other points of the morph spectrum), indicating that a majority of participants adopted a stable strategy regardless of the stimuli and the task (generalisation vs. categorisation), see SM Section 2.2 for further details.

In summary, when categorization is made explicit through naming, it differs from generalisation for our stimuli. This indicates that in generalisation, such as that observed in Experiment 1, participants are not generalising entirely based upon their categorical judgements. If this were the case, we would have observed the same results for both categorisation and generalisation, but we did not. Participants appear to show greater certainty when categorising the test pictures by labels compared to when generalising by a neutral information to them, even when the stimuli are the same. This suggest that participants are overall more cautious to generalise a property to a novel stimulus than to assign it a category, however under the highest degree of uncertainty generalisation and categorisation mechanisms appear to operate similarly

5. Experiment 4

In Experiment 1 we found that, under uncertainty, namely when adult participants did not know the category membership of the test picture, they generalised more broadly for toxicity (danger) information, and less broadly for edibility information, compared to seasonality (neutral) information. While these findings are consistent with the asymmetrical costs of generalisation errors in the food domain, one possible alternative explanation for these findings could be a broader negativity bias. This is a well-established asymmetry in the way adults and children process and use positive versus negative information in several psychological domains, such as decision-making, social learning, emotional processing (Baumeister et al., 2001; Bebbington et al., 2017; Stubbersfield et al., 2015; Vaish et al., 2008). If this was the case, and there is a broader negativity bias in generalisation, we would expect adults to generalise any kind of negative information more broadly than neutral information. However, if the selective generalisation we observed in Experiment 1 reflect the specific asymmetries in the costs of the two types of errors in generalising the edibility or toxicity of foods (false-positives and false-negatives; Haselton & Buss, 2000; Haselton & Nettle, 2006), we would predict different generalisation for information about danger than other negative information.

To test this, we ran a final experiment where we asked adults to generalise three novel properties with the same valence as in Experiment 1 (negative, positive, neutral) but where mistakes do not pose the same direct costs or benefits. In Experiment 1 we used the property “toxic”, which is both negative and offers direct costs to an individual if incorrectly generalised. In Experiment 4 we used the property “polluting” [this is polluting/ not polluting] which also has negative connotations, but does not have the same direct costs. That is, many more instances of incorrectly attributing the negative property "polluting" to food items would be required for the individual to experience adverse effects, if at all, from polluting foods. Likewise, the positive property “edible” used in Experiment 1, was replaced in Experiment 4 with the positive property of “sustainable” [this is sustainable/ not sustainable]. This property is positive but again without the same direct costs associated with errors in attributing it. A single instance of incorrectly generalising edibility to an inedible item is more costly than incorrectly generalising sustainability to something that is not sustainable. Finally, the neutral property of Experiment 1 “grows in summer” was replaced in Experiment 4 with another neutral property “stored in a cool dry place” [this is stored/ not stored in a cool dry place]. We replaced the seasonality control used in Experiment 1 to examine whether participants treated a different kind of neutral information in a similar way.

Adult participants once again evaluated triads of the same images used in the previous experiments, but with the properties polluting, sustainability, and storage instead of toxicity, edibility and seasonality. If the generalisation effects we observed in Experiment 1 are a function of a broader negativity/positivity effect we would expect to see a similar pattern in generalisation in Experiment 4, with an overgeneralisation of polluting information and undergeneralisation of sustainability information. However, if the effects in Experiment 1 are reflect the asymmetry in costs in generalising edibility and toxicity information in the food domain over evolutionary time, we would not expect to see this effect in Experiment 4.

5.1. Method

Participants were 90 adults (47 female, 41 male, mean age = 31 years, range = 17 – 77, 2 participants did not provide demographic information), divided evenly across three independent informational conditions (positive, negative, and neutral). As with previous experiments Experiment 4 was reviewed and approved by a local ethics committee, and the recruitment and consent procedure were the same as Experiments 1 and 3. The stimulus pictures and procedure were identical to the Experiment 1, except that the properties assigned to the target pictures were different. To determine the properties most relevant to food and with our desired valences, we ran a pilot experiment. We piloted two different properties per valence (positive, negative, neutral), which were online with 30 adults. The properties selected (i.e., polluting, sustainable, and stored in a cool dry place) were those adults rated most relevant to describe a food and that had the *a priori* expected valence (i.e., the property "polluting" was rated strongly negatively; see SM Section 1.3 for details).

5.2. Results and Discussion

As with previous experiments, results from Experiment 4 were analysed using a Binomial GLMM with a logit link function. We found no improved model fit of including image pair, age, or sex in the model. As such, these variables were again excluded from further analyses. In the main analysis Participants served as a random factor as well as Individual triads. Similarity between the left target picture and the test picture (83%, 67%, 50%, 33% and 17%), Condition (negative, positive, neutral), as well as their interaction were modelled as fixed effects. The best fit model according to Akaike's Information Criterion was the full model including the main effects and the interaction between Similarity and Condition (χ^2 (2) = 9.68, p = 0.0079 marginal R^2 = 0.74, conditional R^2 = 0.76).

Consistent with our previous findings, a significant main effect of Similarity was found, with the likelihood of generalisation decreasing as similarity decreased (χ^2 (1) = 291.09, p < 0.001, see Table S4 in the SM section 2.1). As in Experiment 1, we found a significant main effect of Condition (χ^2 (2) = 9.87, p = 0.0072, see Table S4 in the SM), which was qualified by an interaction effect of Similarity*Condition (χ^2 (2) = 8.39, p = 0.015, see Table S4 in the SM). The decline in generalisation due to decreasing similarity in the neutral condition was significantly steeper than in the positive condition (z = 2.42, p = 0.041) and negative condition (z = 2.83, p = 0.013). However, Fisher's Exact Tests at each point of our morph sequence (i.e., at each level of similarity) indicated that, in contrast to Experiment 1, no significant pairwise comparisons were found between the negative, positive, or neutral conditions. This indicates that there was no significant difference between generalisation rates of positive, negative and neutral information in Experiment 4. In particular, we compared generalisation rates where the morph test picture was most ambiguous (i.e. 50% similar to each of the original items) to the chance level. In contrast to Experiment 1, generalisation rates in neither experimental condition, negative (P = 0.58 [95% CI: 0.48 – 0.66], p = 0.12), nor positive (P = 0.42, [95% CI: 0.33-0.51], p = 0.082), differed significantly from chance. This indicates again that there was no significant difference between the generalisation of positive, negative and neutral information in Experiment 3.

Figure 5: Generalisation rates in Experiment 4 at each point of the morph sequence (i.e., each level of Similarity), depending of Condition. Large dots show group means with 95% CI error bars, smaller dots show individual participant means. Morph images from an example triad after Fig. 1, the real green fruit has the property.

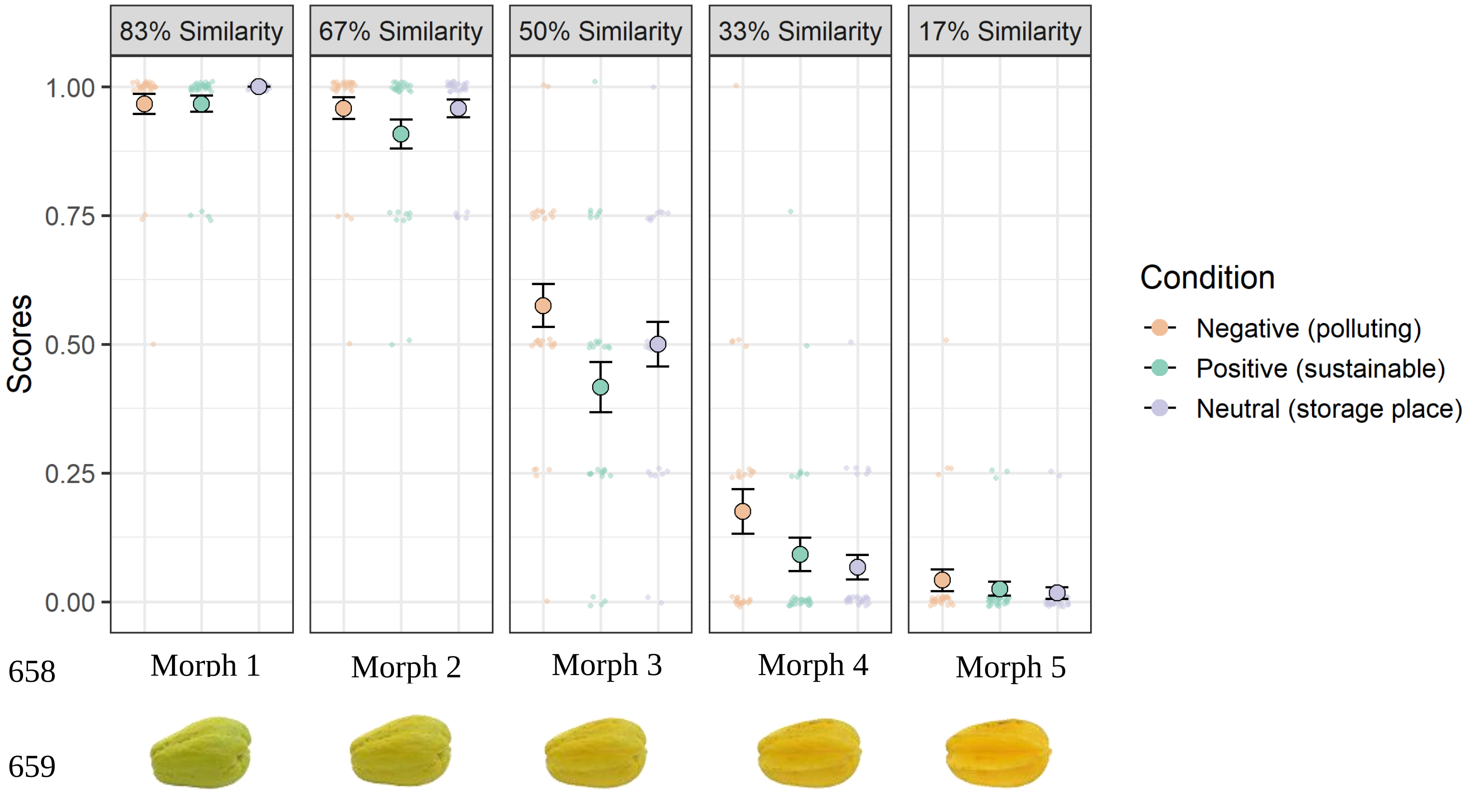

To examine the difference between Experiment 1 and Experiment 4 further we compared generalisation rates in cases where the test picture was most ambiguous (i.e. 50% similar to each of the original items) between the original toxicity condition from Experiment 1 and the negative (polluting) condition from Experiment 4, and also between the edibility condition and positive (sustainable) conditions from Experiments 1 and 4 respectively. This was done using Fishers Exact tests. We found no significant difference between the edibility and positive ($p = 0.69$), or toxicity and negative conditions ($p = 0.79$), indicating that the effect of edibility is not necessarily distinct from a general positivity effect, nor is toxicity distinct from a general negativity effect. However, we observed edibility and toxicity to be significantly different from chance in Experiment 1, and did not observe a similar pattern for positivity and negativity here. Further, we did find a significant difference between positivity, negativity, and neutral conditions in Experiment 4 while we did for toxicity and edibility in Experiment 1.

Finally, consistent with previous experiments, more than 93% of participants adopted a stable strategy for morphs at the beginning and end of the spectrum. For the most ambiguous morph, only 62% of participants adopted a stable strategy. This was significantly lower than for the other points of the morph spectrum ($p < .001$ for all; see SM Section 2.2 for further details), again indicating greater uncertainty in generalisation for the most ambiguous morph.

In summary, results of Experiment 4 give some indication that a general negativity/positivity effect is present in generalisation, however it is less pronounced than the effect for toxicity and edibility information we observed in Experiment 1, suggesting some level of specificity in generalisation.

6. General Discussion

When making inductive inferences under uncertainty, some generalisation errors are more costly than others. Generalising knowledge from one exemplar to a novel and ambiguous item, based on categories or similarity alone, is insufficient because the consequences of errors are not always uniform. Decisions about candidate food items provide a particularly clear example: Incorrectly deciding a toxic item is edible is a higher cost than incorrectly deciding that an edible item is toxic. In these cases, some selectivity is necessary to guide generalisation processes. Following the predictions of Error Management Theory (EMT; Haselton & Buss, 2000; Haselton & Nettle, 2006) and research showing that selective learning is present for certain fitness-relevant information (e.g., dangerous animals, foods, Barrett & Broesch, 2012; Wertz & Wynn, 2014a), the present work investigated how the type of information learned influences the generalisation of that information to novel items. Further, based on Broesch et al. (2014)'s finding that the selective learning for dangerous animals declined with age and was not present in adulthood, we also looked at age differences in selective generalisation processes by testing preschool-aged children and adults.

Our results demonstrate that both children and adults consistently generalise along a declining function of item similarity, consistent with the literature (e.g., Murphy, 2002; Shepard, 1958; Sloutsky, 2003), however adults were overall more cautious in generalising a property to a novel stimulus than assigning the same stimulus to a category. Moreover, as predicted, in Experiment 1, the findings reveal evidence of selective generalisation dependent on the information learned. Specifically, adults' generalisation decisions across the entire morph sequence show that generalisation of fitness-relevant information with asymmetric costs of errors (i.e., toxicity and edibility information about food) differs from generalisation of neutral content. Specifically, generalisation of toxicity information was less stringent (declined less rapidly with similarity) and edibility information more stringent (declined more rapidly with similarity), than neutral information. This is consistent with the error minimising principles of Error Management Theory (Haselton & Nettle, 2006), as this pattern appears to show a strategy that minimises the more costly error in each case. For edibility information, it would be more costly to incorrectly generalise edibility to entity that is not in fact edible, and we see reduced generalisation that would minimise this possibility. For toxicity information, the more costly error would be to fail to generalise toxicity information to an entity that is in fact toxic, and we see higher generalisation levels reflecting this. To our knowledge, this is the first demonstration of selective generalisation based on informational content.

This effect was driven by participants' generalisation decisions for the most ambiguous morph— the image that was halfway between the two target images in similarity. When confronted with the most ambiguous morph, and therefore under a state of categorical uncertainty (as shown by Experiment 3), it led participants generalise toxicity to a greater extent, and edibility to a lesser extent than neutral information. That this only appears in the most ambiguous morph highlights the role of uncertainty in selective generalisation. In fact, adults exhibited less certainty in their generalisation responses (Experiment 1) than in their categorisation responses (Experiment 3) in general, which could indicate that formation of categories requires less evidence than induction. This is consistent with the notion of minimizing costly errors— the mistakes involved in using information are possibly more costly than mistakes in learning that information That participants did not generalise (Experiment 1) or categorise (Experiment 3) neutral information differently from chance for the most ambiguous item, suggests that the differences in generalisation at this level of similarity in Experiment 1 are a function of the type of information to be generalised. More specifically, and consistent with the predictions drawn from EMT (Haselton & Buss, 2000;

Haselton & Nettle, 2006), generalisation patterns of fitness-relevant information (i.e., the toxicity and edibility of food) reflect the asymmetry in possible errors costs in those domains.

These findings are in line with previous research showing selective learning for danger and edibility information in infants and children (Barrett et al., 2016b; Barrett & Broesch, 2012b; Wertz & Wynn, 2014a), but differ in that selectivity in generalisation appears limited to adulthood. Given that selective learning is observed in younger age groups, we also examined selective generalisation with 4- to 5-year old children, who were given an adapted version of the generalisation task in Experiment 1. In contrast to the adults in Experiment 1, we did not find that children generalised toxicity to a greater extent and edibility to a lesser extent than neutral information. However, we found some hints of selective generalisation depending on type of information. Specifically, for the two most similar morphs, children in the high food neophobia group generalised toxicity to a higher degree than their counterparts, who generalised neutral information to a higher degree. These findings are consistent with the view that neophobia is a protective strategy against the risk of ingesting potentially toxic foods and suggest that neophobic individuals pay particular attention to negative information about foods (Cashdan, 1994; Foinant et al., 2021; Lafraire et al., 2016; Rozin & Todd, 2015). Notably, Foinant and colleagues presented children with a property generalisation task in which one target fruit was given a property (either positive or negative) and children were asked which other real fruits and vegetables also had the property. They found that neophobic children assign negative properties to more foods (e.g., “This food makes you throw up”) compared to individuals with less neophobic disposition, (Foinant et al., 2021). In addition, our findings are convergent with recent work showing that neophobic children show an attentional bias toward negatively perceived food stimuli, such as novel fruits and vegetables that are more likely than other types of foods to contain toxic compounds (Maratos et al., 2008) and may be more reluctant to approach those kinds of foods as infants (Rioux & Wertz, 2021).

Contrary to the findings of Broesch and colleagues (2014) showing more pronounced learning biases in children, we found evidence of selective generalisation in adults but not children. It has already been established that generalisation processes can unfold differently than other processes (Gelman & Markman, 1986, 1987; Sloutsky & Fisher, 2004a, 2004b; see also the results of Experiment 3 where we found a difference between categorisation and generalisation responses). It could be that differences between the developmental trajectories of selective learning of fitness-relevant information and selective generalisation are a function of the difference in the use of knowledge at different ages. That is, that selectivity in the process of acquiring knowledge is of higher value earlier in life than selectivity in generalisation of that knowledge, particularly when the overall structure of the generalisation processes is sufficient (e.g., the decrease in generalisation we observed over declining similarity between the target and test stimuli in Experiment 2). The importance of the human extended juvenile period for learning about our environment has been well established (e.g.,Crittenden et al., 2013; Joffe, 1997; Muthukrishna et al., 2018), with Gopnik et al. (2017) noting specifically that children learn better than adults across several domains.

Finally, Experiment 4 investigated to what extent the selective generalisation effect found in Experiment 1 may be a function of a broader valence effect. A general negativity bias is a well-established asymmetry in the way adults and children process and use positive versus negative information in several psychological domains, including decision-making, social learning, and emotional processing (Baumeister et al., 2001; Bebbington et al., 2017; Stubbersfield et al., 2015; Vaish et al., 2008). Therefore, in Experiment 4, we compared the generalisation of the fitness-relevant information used in Experiment 1 (toxicity and edibility information about foods) to information that was similarly valanced (negative, positive), but without the direct fitness risks of generalisation errors. Specifically, participants generalised negative and positive information about food (i.e., being polluting or sustainable) that does not pose the same kind of direct risks and benefits to an individual as decisions about edibility and toxicity information. For example, whilst costly on a broader scale, mistakenly deciding that something is sustainable when does not pose the same direct risk to an individual as thinking something is edible when it is not.

The results of Experiment 4 showed that, across the entire morph sequence, adults' generalisation decisions for valanced information (i.e., positive [sustainable] and negative [polluting] information) was *less* stringent than generalisation of neutral information (storage location). In addition, in contrast to Experiment 1, when participants were confronted with the most ambiguous morph items, participants did not show evidence of selective generalisation of positive and negative information. These results speak against a mere valence bias account for the findings of Experiment 1 and instead suggest that participants engaged in selective generalisation based on informational content in a manner consistent with the EMT prediction of minimizing of costly mistakes. However, we acknowledge that this interpretation must be taken with caution given that direct comparisons between generalisation of information with high vs cost of errors (*Toxic* vs. *Polluting* and *Edible* vs. *Sustainable*) revealed no significant differences.

6.1. Limitations and future directions

That the effect of information type on generalisation was only observed in the most ambiguous morph illustrates one limitation in the present design which warrants further investigation. We observed clear differences between how participants generalised to our most ambiguous morph (morph 3) and those on either side of it in the morph sequence (morphs 2 and 4). One outstanding question is the boundary condition for this difference, and whether this differs between information types. That is, at what degree of similarity to the target do these content effects appear? The current study limited the morph sequence to only five items, and thereby the ability to answer this question. This was done to limit the total number of questions presented to the child sample in Experiment 2. Future research could investigate the relationship between similarity and information type on generalisation by extending the morph sequence and systematically manipulating specific features of the exemplars. For example, previous research has shown that colour and shape play important roles in food categorisation (e.g., Rioux et al., 2018). In morphing the foods were items into one another, however the relative degrees of change of colour and shape along this morph sequence were not measured explicitly. Given the debate around the relative roles of these factors in food categorisation (e. g. Lavin & Hall, 2002; Macario, 1991; Rioux et al., 2018), further research could examine how similarity along each of these dimensions independently influences generalisation of different information types.

Finally, our results raised questions about the role of age in selective generalisation. As noted, we found an inverse pattern to Broesch et al. (2014), finding selectivity in generalisation was observed in adults but not children. We did not find an age effect within the sample, suggesting that the change occurs between the ages of our child sample, and adulthood. However, it should be noted that the effect of age within our adult sample was not a planned investigation so we may not have had enough power to examine this. This suggests the need for further research to examine the developmental trajectory of selectivity in generalisation, specifically the age at which the informational content effects observed here arise in later childhood or the teenage years. Alternatively, it could be the case that there is indeed an effect of informational content in children, but we did not observe it. For example, the difference between children and adults could be a function of differences in task competence, that is children’s ability to understand and complete the task rather than their performance on the task, a well-known issue in developmental research (e.g. O'Brien & Overton, 1982; Surian & Leslie, 1999). It would therefore be of value to replicate the present findings using alternative methods to examine this.

7. Conclusions

The present work investigated the effect of different types of informational content on generalisation in adults and preschool-aged children. In line with the Error Management Theory (Haselton & Buss, 2000; Haselton & Nettle, 2006), we found evidence of selective generalisation. Specifically, under conditions of uncertainty, adults generalised danger information (i.e., toxicity) more broadly and edibility information more narrowly than neutral information. The present results also suggest a distinction between generalisation of fitness-relevant information with direct costs and benefits, and a generalisation of negative and positive information more broadly. Contrary to our predictions we did not find evidence of selective generalisation in young children, yet we did find some evidence of increased generalisation of danger information by children who scored higher on measures of food neophobia. Although previous work has found selectivity in other cognitive processes, such as learning and attentional biases (e.g., (Barrett et al., 2016; DeLoache & LoBue, 2009; LoBue & DeLoache, 2008; Soares et al., 2014; Wertz & Wynn, 2014a), to our knowledge, our findings constitute the first evidence that human generalisation systems reflect evolutionarily recurrent asymmetries in costs. When confronted with ambiguous candidate food items, adults use a cautious generalisation strategy consistent with minimizing costly errors.

8. References

Baldwin, D. A., Markman, E. M., & Melartin, R. L. (1993). Infants’ Ability to Draw Inferences about Nonobvious Object Properties: Evidence from Exploratory Play. *Child Development*, *64*(3), 711–728. https://doi.org/10.1111/j.1467-8624.1993.tb02938.x

Barrett, H. C., & Broesch, J. (2012). Prepared social learning about dangerous animals in children. *Evolution and Human Behavior*, *33*(5), 499–508. https://doi.org/10.1016/j.evolhumbehav.2012.01.003

Barrett, H. C., Peterson, C. D., & Frankenhuis, W. E. (2016). Mapping the Cultural Learnability Landscape of Danger. *Child Development*, *87*(3), 770–781. https://doi.org/10.1111/cdev.12495

Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad Is Stronger Than Good. *Review of General Psychology*, *5*(4), 323–370. https://doi.org/10.1037//1089-2680.5.4.323

Bebbington, K., MacLeod, C., Ellison, T. M., & Fay, N. (2017). The sky is falling: evidence of a negativity bias in the social transmission of information. *Evolution and Human Behavior*, *38*(1), 92–101. https://doi.org/10.1016/j.evolhumbehav.2016.07.004

Belovsky, G. E., & Schmitz, O. J. (2006). Plant Defenses and Optimal Foraging by Mammalian Herbivores. *Journal of Mammalogy*, *75*(4), 816–832. https://doi.org/10.2307/1382464

Broesch, J., Barrett, H. C., & Henrich, J. (2014). Adaptive Content Biases in Learning about Animals across the Life Course. *Human Nature*, *25*(2), 181–199. https://doi.org/10.1007/s12110-014-9196-1

Cashdan, E. (1994). A sensitive period for learning about food. *Human Nature*, *5*(3), 279–291. https://doi.org/10.1007/BF02692155

Chen, S. Y., Ross, B. H., & Murphy, G. L. (2014). Implicit and explicit processes in category-based induction: Is induction best when we don't think? *Journal of Experimental Psychology: General*, *143*(1), 227–246. https://doi.org/10.1037/a0032064

Crittenden, A. N., Conklin-Brittain, N. L., Zes, D. A., Schoeninger, M. J., & Marlowe, F. W. (2013). Juvenile foraging among the Hadza: Implications for human life history. *Evolution and Human Behavior, 34*(4), 299–304. https://doi.org/10.1016/J.EVOLHUMBEHAV.2013.04.004

DeLoache, J. S., & LoBue, V. (2009). The narrow fellow in the grass: Human infants associate snakes and fear. *Developmental Science, 12*(1), 201–207. https://doi.org/10.1111/j.1467-7687.2008.00753.x

Deng, W., & Sloutsky, V. M. (2013). The role of linguistic labels in inductive generalization. *Journal of Experimental Child Psychology*, *114*(3), 432–455. https://doi.org/10.1016/J.JECP.2012.10.011

Drewnowski, A., Henderson, S. A., Driscoll, A., & Rolls, B. J. (1997). The Dietary Variety Score: Assessing Diet Quality in Healthy Young and Older Adults. *Journal of the American Dietetic Association*, *97*(3), 266–271. https://doi.org/10.1016/S0002-8223(97)00070-9

Dunlap, A. S., & Stephens, D. W. (2014). Experimental evolution of prepared learning. *Proceedings of the National Academy of Sciences, 111*(32), 11750–11755. https://doi.org/10.1073/pnas.1404176111

Feeney, A., & Heit, E. (Eds.). (2007). *Inductive Reasoning: Experimental, Developmental, and Computational Approaches*. Cambridge University Press. https://doi.org/10.1017/CBO9780511619304

Ferry, A. L., Hespos, S. J., & Waxman, S. R. (2010). Categorization in 3- and 4-Month-Old Infants: An Advantage of Words Over Tones. *Child Development, 81*(2), 472. https://doi.org/10.1111/J.1467-8624.2009.01408.X

Fisher, A. v., Godwin, K. E., Matlen, B. J., & Unger, L. (2015). Development of Category-Based Induction and Semantic Knowledge. *Child Development*, *86*(1), 48–62. https://doi.org/10.1111/CDEV.12277

Fisher, A. v., Matlen, B. J., & Godwin, K. E. (2011). Semantic similarity of labels and inductive generalization: Taking a second look. *Cognition*, *118*(3), 432–438. https://doi.org/10.1016/J.COGNITION.2010.12.008

Foinant, D., Lafraire, J., & Thibaut, J. P. (2021). Strength or Nausea? Children's Reasoning About the Health Consequences of Food Consumption. *Frontiers in Psychology, 12*, 1119. https://doi.org/10.3389/FPSYG.2021.651889/BIBTEX

Garcia, J., & Koelling, R. A. (1966). Relation of cue to consequence in avoidance learning. *Psychonomic Society*, *4*(1), 123–124. https://doi.org/10.3758/BF03342209

Gelman, S. A. (1988). The development of induction within natural kind and artifact categories. *Cognitive Psychology, 20*(1), 65–95. https://doi.org/10.1016/0010-0285(88)90025-4

Gelman, S. A., & Davidson, N. S. (2013). Conceptual influences on category-based induction. *Cognitive Psychology, 66*(3), 327–353. https://doi.org/10.1016/J.COGPSYCH.2013.02.001

Gelman, S. A., & Markman, E. M. (1986). Categories and induction in young children. *Cognition*, *23*(3), 183–209. https://doi.org/10.1016/0010-0277(86)90034-X

Gelman, S. A., & Markman, E. M. (1987). *Young Children ' s Inductions from Natural Kinds : The Role of Categories and Appearances*. *58*(6), 1532–1541.

Gopnik, A., O'Grady, S., Lucas, C. G., Griffiths, T. L., Wente, A., Bridgers, S., Aboody, R., Fung, H., & Dahl, R. E. (2017). Changes in cognitive flexibility and hypothesis search across human life history from childhood to adolescence to adulthood. *Proceedings of the National Academy of Sciences of the United States of America*, *114*(30), 7892–7899. https://doi.org/10.1073/PNAS.1700811114/SUPPL_FILE/PNAS.1700811114.SAPP.PDF

Hagen, E. H., Sullivan, R. J., Schmidt, R., Morris, G., Kempter, R., & Hammerstein, P. (2009). Ecology and neurobiology of toxin avoidance and the paradox of drug reward. *Neuroscience*, *160*(1), 69–84. https://doi.org/10.1016/J.NEUROSCIENCE.2009.01.077

Hanley, M. E., Lamont, B. B., Fairbanks, M. M., & Rafferty, C. M. (2007). Plant structural traits and their role in anti-herbivore defence. *Perspectives in Plant Ecology, Evolution and Systematics*, *8*(4), 157–178. https://doi.org/10.1016/j.ppees.2007.01.001

Haselton, M. G., & Buss, D. M. (2000). Error management theory: A new perspective on biases in cross-sex mind reading. *Journal of Personality and Social Psychology*, *78*(1), 81–91. https://doi.org/10.1037//0022-3514.78.1.81

Haselton, M. G., & Nettle, D. (2006). The Paranoid Optimist : An Integrative Evolutionary Model of Cognitive Biases. *Personality and Social Psychology Review*, *10*(1), 47–66. https://doi.org/10.1207/s15327957pspr1001_3

Haselton, M. G., Nettle, D., & Murray, D. R. (2005). Evolution of Cognitive Bias. *The Handbook of Evolutionary Psychology*, 1–39. https://doi.org/10.1128/jb.177.6.1461-1469.1995

Hayes, B. K., Chen, T. H. J., & Jessamine Chen, T.-H. (2008). *Clinical expertise and reasoning with uncertain categories*. *15*(5), 1002–1007. https://doi.org/10.3758/PBR.15.5.1002

Hayes, B. K., & Newell, B. R. (2009). Induction with uncertain categories: When do people consider the category alternatives? *Memory & Cognition*, *37*(6), 730–743. https://doi.org/10.3758/MC.37.6.730

Hearst, E. (1962). Concurrent Generalization Gradients for Food-Controlled and Shock Controlled Behaviour. *Journal of the Experimental Analysis of Behavior*, *5*(1), 19–31. https://doi.org/10.1901/JEAB.1962.5-19

Hoffman, H. S., & Fleshler, M. (1963). Discrimination and stimulus generalization of approach, of avoidance, and of approach and avoidance during conflict. *Journal of Experimental Psychology*, *65*(3), 280–291. https://doi.org/10.1037/H0043630

Joffe, T. H. (1997). Social pressures have selected for an extended juvenile period in primates. *Journal of Human Evolution*, *32*(6), 593–605. https://doi.org/10.1006/JHEV.1997.0140

Johnson, D. D. P., Blumstein, D. T., Fowler, J. H., & Haselton, M. G. (2013). The evolution of error: error management, cognitive constraints, and adaptive decision-making biases. *Trends in Ecology & Evolution*, *28*(8), 474–481. https://doi.org/10.1016/J.TREE.2013.05.014

Keeler, R., & Tu, A. (Eds.). (1983). *Plant and Fungal Toxins: Handbook of Natural Toxins Vol. 1*. Marcel Decker.

Krasnow, M. M., Truxaw, D., Gaulin, S. J. C., New, J., Ozono, H., Uono, S., Ueno, T., & Minemoto, K. (2011). Cognitive adaptations for gathering-related navigation in humans. *Evolution and Human Behavior*, *32*(1), 1–12. https://doi.org/10.1016/j.evolhumbehav.2010.07.003

Lafraire, J., Rioux, C., Giboreau, A., & Picard, D. (2016). Food rejections in children: Cognitive and social/environmental factors involved in food neophobia and picky/fussy eating behavior. *Appetite*, *96*, 347–357. https://doi.org/10.1016/J.APPET.2015.09.008

Landau, B., Smith, L., & Jones, S. (1998). Object Shape, Object Function, and Object Name. *Journal of Memory and Language*, *38*(1), 1–27. https://doi.org/10.1006/JMLA.1997.2533

Lavin, T. A., & Hall, D. G. (2002). Domain effects in lexical development: Learning words for foods and toys. *Cognitive Development*, *16*(4), 929–950. https://doi.org/10.1016/S0885-2014(02)00070-9

Lindström, B., Selbing, I., & Olsson, A. (2016). Co-evolution of social learning and evolutionary preparedness in dangerous environments. *PLoS ONE*, *11*(8), 1–20. https://doi.org/10.1371/journal.pone.0160245

LoBue, V., & DeLoache, J. S. (2008). Detecting the Snake in the Grass. *Psychological Science*, *19*(3), 284. https://doi.org/10.1111/j.1467-9280.2008.02081.x

LoBue, V., & Rakison, D. H. (2013). What we fear most: A developmental advantage for threat-relevant stimuli. *Developmental Review*, *33*(4), 285–303. https://doi.org/10.1016/j.dr.2013.07.005

LoBue, V., Rakison, D. H., & DeLoache, J. S. (2010). Threat Perception Across the Life Span. *Current Directions in Psychological Science*, *19*(6), 375–379. https://doi.org/10.1177/0963721410388801

Macario, J. F. (1991). Young children’s use of color in classification: Foods and canonically colored objects. *Cognitive Development*, *6*(1), 17–46. https://doi.org/10.1016/0885-2014(91)90004-W

Mandler, J. M., & McDonough, L. (1993). Concept formation in infancy. *Cognitive Development*, *8*(3), 291–318. https://doi.org/10.1016/S0885-2014(93)80003-C

Manners, G. D. (1996). Plant Toxins. In B. R. Singh & A. T. Tu (Eds.), *Natural Toxins 2. Advances in Experimental Medicine and Biology* (Vol. 391, pp. 9–35). Springer. https://doi.org/10.1007/978-1-4613-0361-9_2

Mithöfer, A., & Boland, W. (2012). Plant Defense Against Herbivores: Chemical Aspects. *Http://Dx.Doi.Org/10.1146/Annurev-Arplant-042110-103854*, *63*, 431–450. https://doi.org/10.1146/ANNUREV-ARPLANT-042110-103854

Murphy, G. L. (2002). *The Big Book of Concepts*. MIT Press.

Murray, E. J., & Miller, N. E. (1952). Displacement; steeper gradient of generalization of avoidance than of approach with age of habit controlled. *Journal of Experimental Psychology*, *43*(3), 222–226. https://doi.org/10.1037/H0062579

Muthukrishna, M., Doebeli, M., Chudek, M., & Henrich, J. (2018). The Cultural Brain Hypothesis: How culture drives brain expansion, sociality, and life history. *PLOS Computational Biology*, *14*(11), e1006504. https://doi.org/10.1371/JOURNAL.PCBI.1006504

Nesse, R. M. (2005). Natural selection and the regulation of defenses. *Evolution and Human Behavior, 26*(1), 88–105. https://doi.org/10.1016/j.evolhumbehav.2004.08.002

New, J. J., & German, T. C. (2015). Spiders at the cocktail party: An ancestral threat that surmounts inattentional blindness. *Evolution and Human Behavior*, *36*(3), 165–173. https://doi.org/10.1016/j.evolhumbehav.2014.08.004

Nguyen, S. P. (2012). Inductive Selectivity in Children’s Cross-Classified Concepts. *Child Development*, *83*(5), 1748–1761. https://doi.org/10.1111/J.1467-8624.2012.01812.X

Nguyen, S. P. (2020). From foods to artifacts: Children's evaluative and taxonomic categorization across multiple domains. *Cognitive Development*, *56*, 100894. https://doi.org/10.1016/J.COGDEV.2020.100894

Nguyen, S. P., & Murphy, G. L. (2003). An Apple is More Than Just a Fruit: Cross-Classification in Children's Concepts. *Child Development*, *74*(6), 1783–1806. https://doi.org/10.1046/J.1467-8624.2003.00638.X

O'Brien, D. P., & Overton, W. F. (1982). Conditional reasoning and the competence-performance issue: A developmental analysis of a training task. *Journal of Experimental Child Psychology*, *34*(2), 274–290. https://doi.org/10.1016/0022-0965(82)90046-7

Öhman, A., Eriksson, A., Fredriksson, M., Hugdahl, K., & Olofsson, C. (1974). Habituation of the electrodermal orienting reaction to potentially phobic and supposedly neutral stimuli in normal human subjects. *Biological Psychology*, *2*(2), 85–93. https://doi.org/10.1016/0301-0511(74)90017-9

Öhman, A., & Mineka, S. (2001). Fears, phobias, and preparedness: Toward an evolved module of fear and fear learning. *Psychological Review*, *108*(3), 483–522. https://doi.org/10.1037/0033-295X.108.3.483

Öhman, A., & Mineka, S. (2003). The Malicious Serpent: Snakes as a Prototypical Stimulus for an Evolved Module of Fear. *Current Directions in Psychological Science*, *12*(1), 5–9. https://doi.org/10.1111/1467-8721.01211

Öhman, A., & Soares, J. J. F. (1998). Emotional Conditioning to Masked Stimuli: Expectancies for Aversive Outcomes Following Nonrecognized Fear-Relevant Stimuli. *Journal of Experimental Psychology: General*, *127*(1), 69–82. https://doi.org/10.1037/0096-3445.127.1.69

Quinn, P. C., Eimas, P. D., & Rosenkrantz, S. L. (1993). Evidence for representations of perceptually similar natural categories by 3-month-old and 4-month-old infants. *Perception, 22*(4), 463–475. https://doi.org/10.1068/P220463

Rioux, C., Lafraire, J., & Picard, D. (2018). Food rejection and the development of food category-based induction in 2–6 years old children. *Journal of Cognitive Psychology, 30*(1), 5–17. https://doi.org/10.1080/20445911.2017.1367688

Rozin, P., & Todd, P. M. (2015). The Evolutionary Psychology of Food Intake and Choice. *The Handbook of Evolutionary Psychology*, 1–23. https://doi.org/10.1002/9781119125563.EVPSYCH106

Rumiati, R. I., & Foroni, F. (2016). We are what we eat: How food is represented in our mind/brain. *Psychonomic Bulletin and Review*, *23*(4), 1043–1054. https://doi.org/10.3758/s13423-015-0908-2

Seligman, M. E. (1970). On the generality of the laws of learning. *Psychological Review*, *77*(5), 406–418. https://doi.org/10.1037/h0029790

Seligman, M. E. P. (1971). Phobias and preparedness. *Behavior Therapy*, *2*(3), 307–320. https://doi.org/10.1016/S0005-7894(71)80064-3

Shepard, R. N. (1958). Toward a Universal Law of Generalization for Psychological Science. *Science*, *13*.

Sloutsky, V. M. (2003). The role of similarity in the development of categorization. *Trends in Cognitive Sciences*, *7*(6), 246–251. https://doi.org/10.1016/S1364-6613(03)00109-8

Sloutsky, V. M., & Fisher, A. v. (2004a). Induction and categorization in young children: A similarity-based model. *Journal of Experimental Psychology: General*, *133*(2), 166–188. https://doi.org/10.1037/0096-3445.133.2.166

Sloutsky, V. M., & Fisher, A. v. (2004b). When development and learning decrease memory: Evidence against category-based induction in children. *Psychological Science*, *15*(8), 553–558. https://doi.org/10.1111/j.0956-7976.2004.00718.x

Soares, S. C., Lindström, B., Esteves, F., & Öhman, A. (2014). The hidden snake in the grass: Superior detection of snakes in challenging attentional conditions. *PLoS ONE*, *9*(12), 1–26. https://doi.org/10.1371/journal.pone.0114724

Stubbersfield, J. M., Tehrani, J. J., & Flynn, E. G. (2015). Serial killers, spiders and cybersex: Social and survival information bias in the transmission of urban legends. *British Journal of Psychology*, *106*(2), 288–307. https://doi.org/10.1111/bjop.12073

Surian, L., & Leslie, A. M. (1999). Competence and performance in false belief understanding: A comparison of autistic and normal 3-year-old children. *British Journal of Developmental Psychology*, *17*(1), 141–155. https://doi.org/10.1348/026151099165203

Thibaut, J. P., Nguyen, S. P., & Murphy, G. L. (2016). Body and soul: Do children distinguish between foods when generalizing biological and psychological properties? *Early Education and Development*, *27*(8), 1250–1262. https://doi.org/10.1080/10409289.2016.1146551

Vaish, A., Grossmann, T., & Woodward, A. (2008). Not All Emotions Are Created Equal: The Negativity Bias in Social-Emotional Development. *Psychological Bulletin*, *134*(3), 383–403. https://doi.org/10.1037/0033-2909.134.3.383

van Duyn, M. A. S., & Pivonka, E. (2000). Overview of the health benefits of fruit and vegetable consumption for the dietetics professional: selected literature. *Journal of the American Dietetic Association*, *100*(12), 1511–1521. https://doi.org/10.1016/S0002-8223(00)00420-X

Welder, A. N., Graham, S. A., & Development, C. (2016). *The Influence of Shape Similarity and Shared Labels on Infants ' Inductive Inferences about Nonobvious Object Properties*. *72*(6), 1653–1673.

Wertz, A. E. (2019). How Plants Shape the Mind. *Trends in Cognitive Sciences*, *23*(7), 528–531. https://doi.org/10.1016/j.tics.2019.04.009

Wertz, A. E. A. E., & Wynn, K. (2019). Can I eat that too? 18-month-olds generalize social information about edibility to similar looking plants. *Appetite*. https://doi.org/10.1016/j.appet.2019.02.013

Wertz, A. E., & Wynn, K. (2014a). Selective social learning of plant edibility in 6- and 18-month-old infants. *Psychological Science*, *25*(4), 874–882. https://doi.org/10.1177/0956797613516145.Selective

Wertz, A. E., & Wynn, K. (2014b). Thyme to touch: Infants possess strategies that protect them from dangers posed by plants. *Cognition*, *130*(1), 44–49. https://doi.org/10.1016/j.cognition.2013.09.002.Thyme

Zhu, J., & Murphy, G. L. (2013). Influence of Emotionally Charged Information on Category-Based Induction. *PLOS ONE*, *8*(1), e54286. https://doi.org/10.1371/JOURNAL.PONE.0054286

Supplementary Material

Contents

1. Methods

1.1. Selection procedure for the image stimuli used in all Experiments

There were originally nine different pairs of real fruit images and their respective morphs, which were piloted online with 30 English or American adults. For each pair, participants were asked to rate their familiarity with each of the two real fruits on a 7-point Likert scale from 1 (“Not familiar at all’) to 7 (‘Very familiar’). Table S1 shows the familiarity ratings.

Participants were also asked to rate how similar on a 7-point Likert scale from 1 (“Not similar at all’) to 7 (‘Very similar’) each real fruit of the pair was with each of the five morph of the sequence (e.g., in the first pair, how similar was the real guava with morph 1, …, morph 5 and how similar was the real star fruit with morph 1, …, morph 5). Figure S1 show the similarity ratings within each pair. For each pair, on the x axis, ticks 1-5 represent similarity ratings between Image 1 of the pair and Morph 1-5 and ticks 6-10 represent similarity ratings between Image 2 of the pair and Morph 1-5.

The final four image pairs selected were those adults rated least familiar (see Pair 1-4 in Table S1) and that had the most consistent similarity gradients across the morph sequence: the difference in similarity between morphs one and two were roughly equivalent to the difference between four and five. In addition, in the final four selected pairs, for each morph, the difference in similarity between one of the real fruit (e.g. guava) and the morph was roughly equivalent to the difference between the other real fruit of the pair and the morph (e.g., star fruit, see Fig. S1, Pair 1 – Morph 1, the far left and far right points should be on the same horizontal line).

Table S1: Familiarity ratings for each food item

Pair 1	Guava	1.80
	Star Fruit	2.67
Pair 2	Yellow zucchini	3.33
	Okra	3.53
Pair 3	Honey melon	1.53
	Custard apple	1.47
Pair 4	Green eggplant	1.47
	Bitter melon	2.17
Pair 5	Jack fruit	1.63
	Swede	5.60
Pair 6	Butternut	6.00
	Barbary fig	2.33
Pair 7	Dragon fruit	2.87
	Papaya	2.77
Pair 8	Cacao fruit	2.00
	Bitter cucumber	1.57
Pair 9	Pomegranate	4.67
	Snake fruit	1.33

Figure S1: Similarity ratings for each pair.

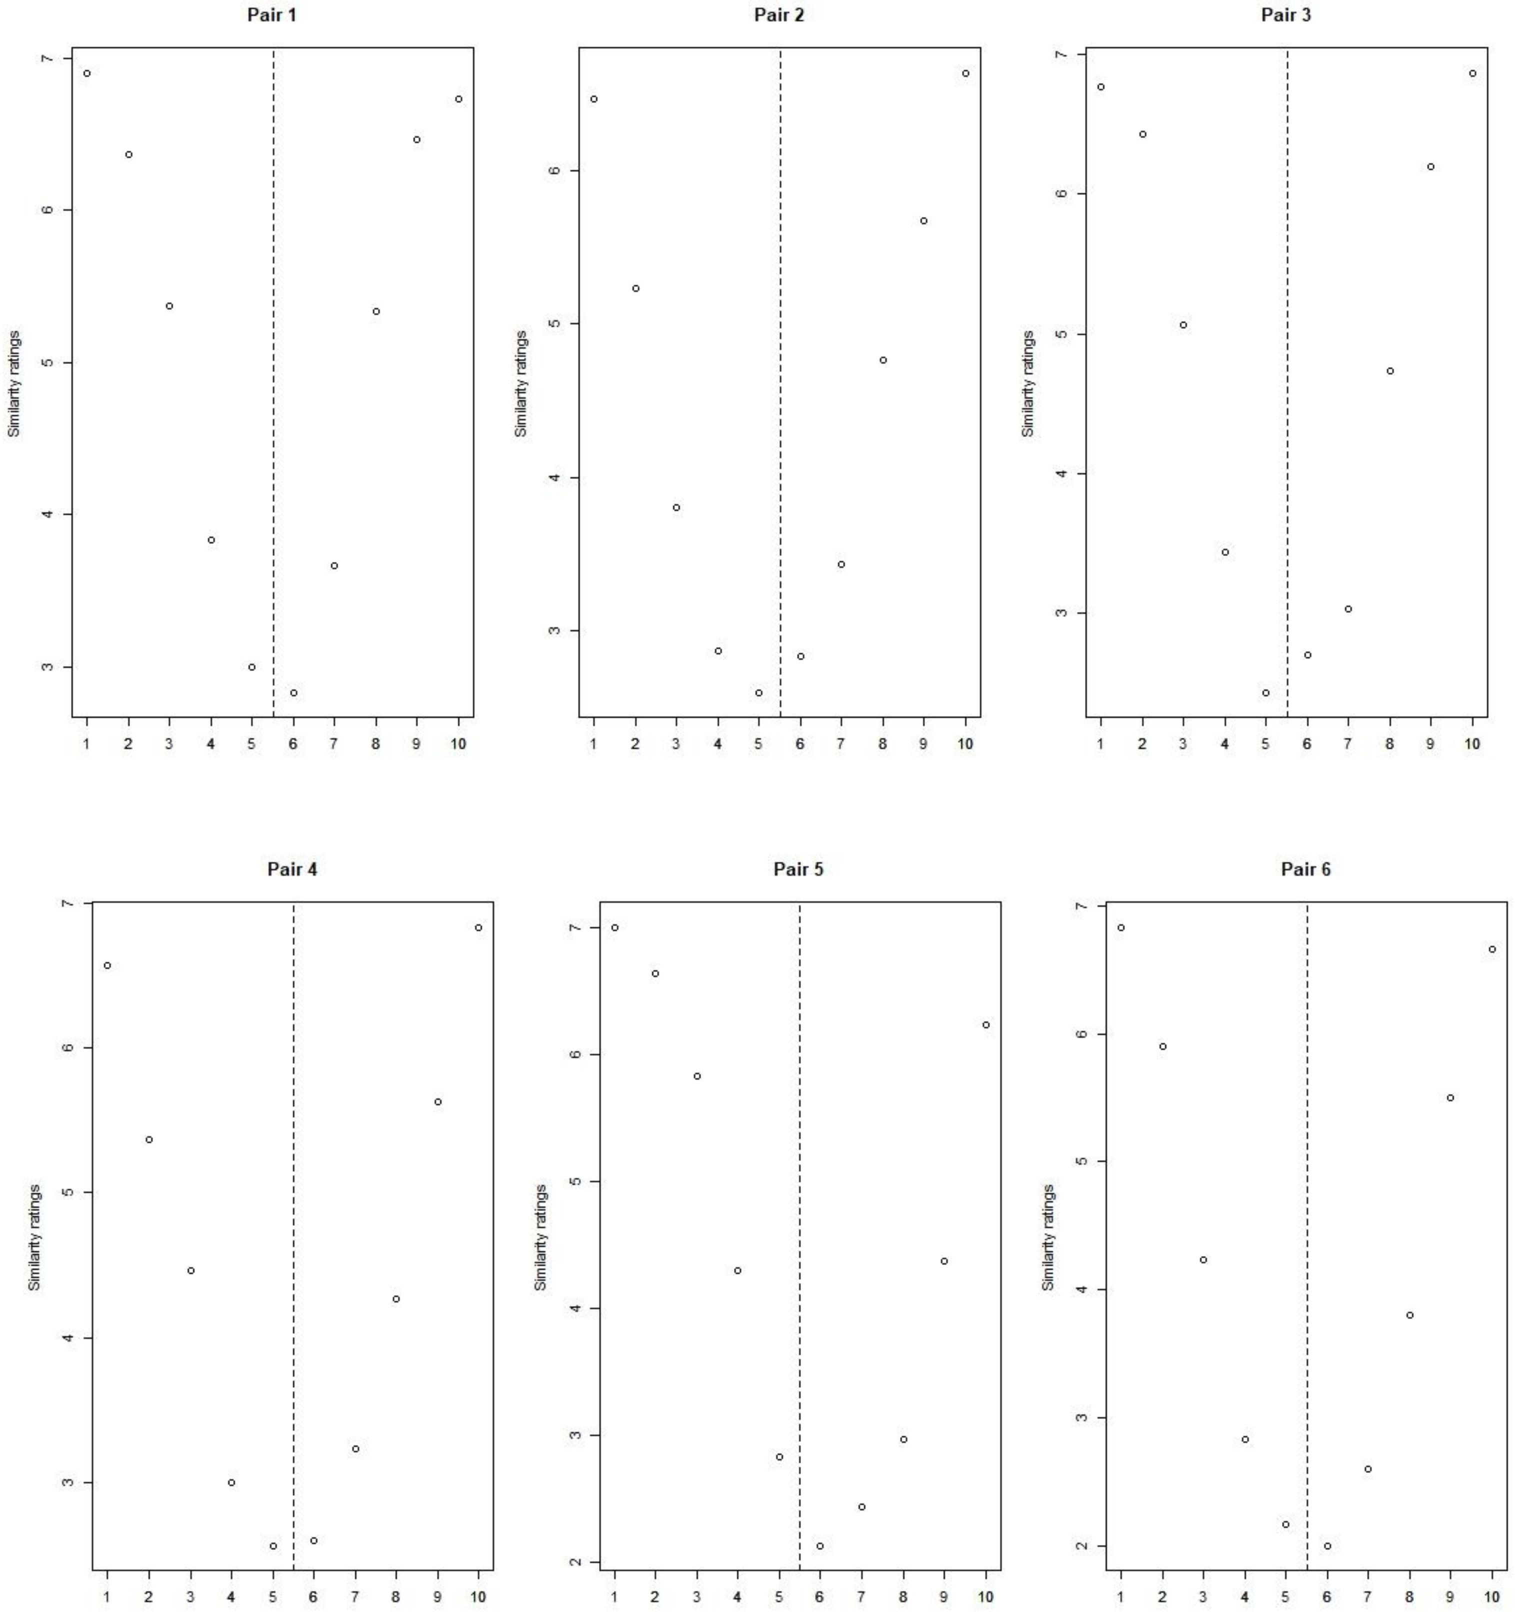

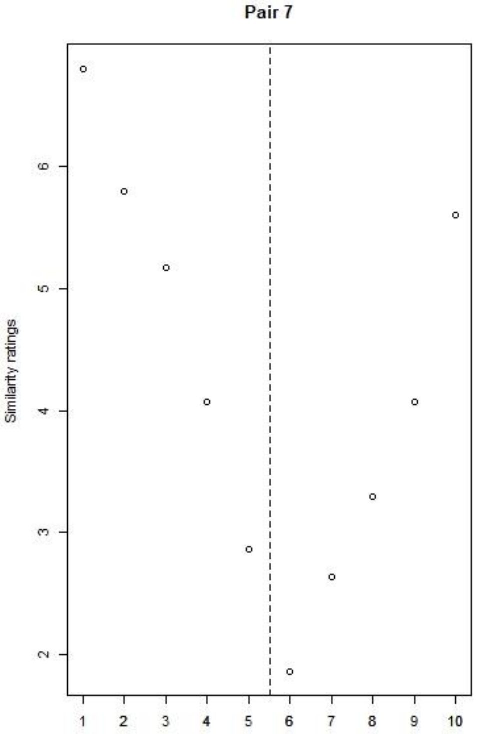

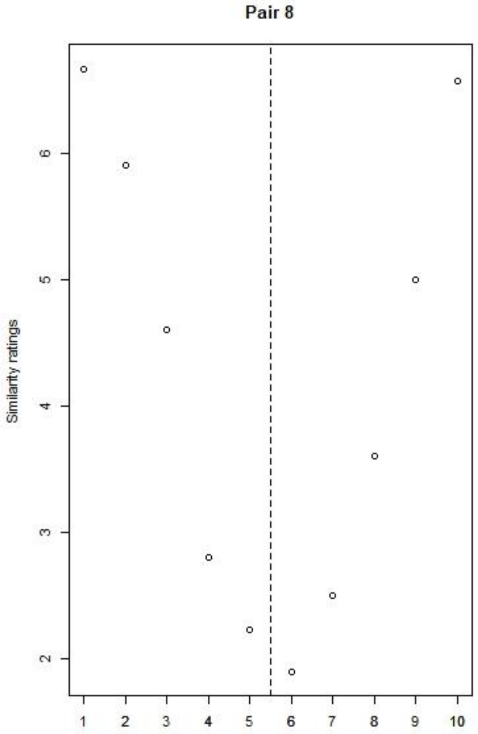

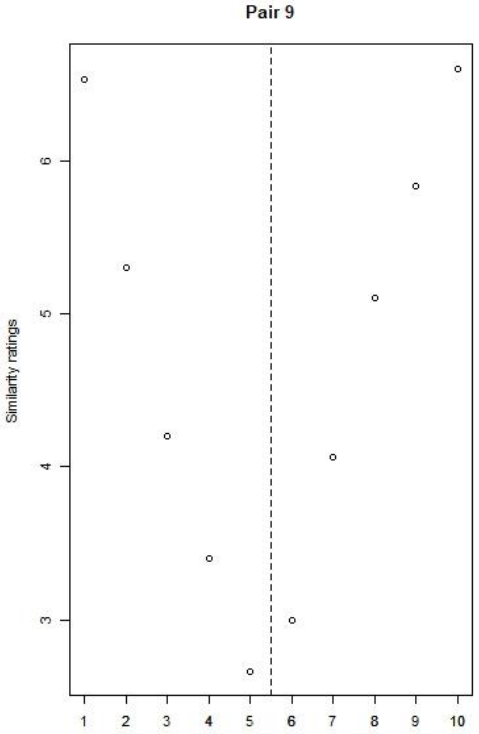

1.2. List of word stimuli used in Experiment 2 (Naming condition)

In Experiment 2, in each trial of the task, an image triad was presented on the screen and one target fruit was given one name (e.g. Daxon for the left top picture), the other target fruit a different name (e.g. Sudeb for the right top picture). Participants were asked which of the two names the test fruit had (e.g., This is a Daxon or a Sudeb?).

The names were novel labels based upon or taken from previous research using novel words and labels (e.g. Fusaro & Harris, 2013; Gaskell & Dumay, 2003). These words were selected to be plausible for the testing language (English) and were all disyllabic (See

Table S2 for the list of names used). A different pair of word was used for each of four pairs, and we counterbalanced across participants which fruit of the pair had name 1 or 2.

Table S2: List of names used in Experiment 2

Pair	Name 1	Name 2
1	Graidesh	Mefop
2	Daxon	Sudeb
3	Tander	Nasket
4	Toma	Fepich

1.3. List of properties used in Experiment 4 (Negative bias condition)

In Experiment 4, the properties assigned to the target pictures were different than in Experiment 1 and 2. We aimed to have properties that had the same valence as in Experiment 1 and 2 (positive, negative and neutral) but without the same immediate cost of ingestion.

There were originally two different properties per valence (positive, negative, neutral), which were piloted online with 30 English or American adults (see Table S3 for the list of properties). Participants were asked to rate the valence of each property on a 7-

point Likert scale from 1 (“Extremely negative”) to 7 (“Extremely positive”). There were also asked to rate the relevance of each property to describe a food item on a 7-point Likert scale from 1 (“Extremely irrelevant”) to 7 (“Extremely relevant”).

The properties selected (i.e., polluting, sustainable, and stored in a cool dry place) were those adults rated most relevant to describe a food item and that had the a priori expected valence (i.e. the property “polluting” was rated strongly negatively; see Table S3), with the constraint that the property “polluting” should be matched with the property “sustainable” and the property “environmentally damaging” should be matched with the property “environmentally friendly”.

Table S3: List of properties used in Experiment 4

Valence	**Property**	**Valence rating**	**Relevance rating**
Negative	**Polluting**	1.67	5.67
	Environmentally damaging	1.53	5.57
Positive	**Sustainable**	6.10	6.13
	Environmentally friendly	6.17	6.17
Neutral	**Stored in a cool dry place**	5.00	5.27
	Growing on a tree	5.37	4.57

Note. In bold the final properties selected for Experiment 4

2. Results

2.1. Complete ANOVA results for linear mixed-effect model for participants' scores in all experiments

Statistical analyses were conducted using R 3. 6.1 (R Core Team, 2019), with main analyses completed using the lme4 package. Data and R scripts are available on the Open Science Framework page for this project (XX). In all experiments Participants' genrelisation/categorization scores have been analyzed, using a binomial generalised linear mixed effect model (GLMM) with a logit link function.

In the main analysis of Experiments 1, 3 and 4 (adult participants) Participants and Individual triads served as a random factor. Similarity between the target picture that had the property and the test picture, Condition as well as their interaction were modelled as fixed effects. Main and interaction effects of the full models are presented in Table S4.

Table S4: Complete ANOVA results for linear mixed-effect model for adult participants' scores in Experiments 1,3 and 4

Effects	**Experiment 1** Condition (edibility, toxicity and seasonality information)			**Experiment 3** Condition (label and seasonality information)			**Experiment 4** Condition (negative, positive and neutral information)		
	χ2	df	p	χ2	df	p	χ2	df	p
Similarity	131.87	1	< .001 ***	66.73	1	< .001 ***	291.09	1	< .001 ***
Condition	6.65	2	.036 *	.039	1	.84	9.87	2	.0072 **
Similarity* Condition	18.13	2	< .001 ***	34.56	1	< .001 ***	8.39	1	.015 *

Note. χ2-values for effects using Type II Wald chi-square tests. P-values < .001 are marked with ***, p values < .01 with ** and p values < .05 with *.

In the main analysis of Experiments 2 (child participants) Participants and Individual triads served as a random factor. Similarity between the target picture that had the property and the test picture, Condition, CFRS group as well as their interaction were modelled as fixed effects. Main and interaction effects of the full models are presented in Table S5.

Table S5: Complete ANOVA results for linear mixed-effect model for child participants' scores in Experiment 2

Effects	**Experiment 2**		
	Condition (edibility, toxicity and seasonality information)		
	χ2	***f***	***p***
Similarity	304.27	1	< .001 ***
Condition	.88	2	.64
CFRS group	.68	1	.41
Similarity*Condition	6.32	2	.042 *
Similarity*CFRS group	.86	1	.35
Condition*CFRS group	5.94	2	.051
Similarity*Condition*CFRS group	17.66	2	< .001 ***

Note. χ2-values for effects using Type II Wald chi-square tests. P-values < .001 are marked with ***, p values < .01 with ** and p values < .05 with *.

2.2. Participants' response stability

In order to examine participants' response stability, in each experiment we analyzed participants' response patterns across the four different stimulus pairs. As participants were presented with morph sequences for four different food pairs and their respective morphs, we were able to examine response stability across the stimuli sets for each point of the morph spectrum. For each morph, participants were considered to have a stable strategy if they responded consistently across the four different food pairs (having the same response for three or four out of four food pairs), or unstable, if they responded differently (different responses in two out of four food pairs).

In Experiment 1, more than 90% of participants adopted a stable strategy for the four morphs at the beginning and end of the spectrum (i.e. the most and least similar morphs). For the most ambiguous morph (i.e., when similairty between the target was 50%), only 63% of participants adopted a stable strategy. This was significantly lower than for the other points of the morph spectrum (p < .0001 for all). This pattern was consistent across conditions. However, whilst lower than in the less ambiguous morphs, the stability rate is still high and statistically greater than chance (Binomial P = 0.63 [95% CI: 0.54 – 0.72, p = 0.0044).

In Experiment 2, more than 84% of children adopted a stable strategy for the two morphs at the beginning and two morphs at the end of the spectrum. For the most ambiguous morph, only 60% of children adopted a stable strategy. As with adults, this was significantly lower than for the other points of the morph spectrum (p < .001 for all). We observed similar pattern in across conditions and CFRS group.

In Experiment 3, for all morphs the percentage of stable responders was high (73% for the most ambiguous morph and more than 96% for the other points of the morph spectrum), indicating that a majority of participants adopted a stable strategy regardless of the stimuli and the task (generalization vs. categorization).

In Experiment 4, more than 93% of participants adopted a stable strategy for morphs at the beginning and end of the spectrum. For the most ambiguous morph, only 62% of participants adopted a stable strategy. This was significantly lower than for the other points of the morph spectrum (p < .001 for all). We observed similar pattern in across conditions.

2.3. Influence of children's food neophobia on generalization responses

The results from Experiment 2 were analyzed using a Binomial GLMM with a logit link function.

We found a significant effect of Similarity (χ^2 (1) = 304.27, $p < 0.001$), demonstrating that as with adults, the likelihood of property generalization is decreasing with similarity. We also found a significant Similarity*Condition interaction (χ^2 (2) = 6.32, $p = 0.042$), which was qualified by a significant three-way interaction between Similarity*Condition*CFRS group (χ^2 (2) = 17.66, $p < 0.001$). In the High-CFRS group, the decrease of property generalization with similarity in the seasonality condition was significantly steeper than in the toxicity condition ($z = 2.59$, $p = 0.046$). In the Low-CFRS group, the decrease of property generalization with similarity in the seasonality condition was significantly more gradual than in the toxicity ($z = 4.02$, $p < 0.001$) and edibility conditions ($z = 3.56$, $p = 0.0020$).

We also conducted Fisher's Exact Tests to examine the 3-way interaction further. Specifically, in both CFRS groups we looked at generalization rates at each point of our morph sequence, depending on the condition. The only significant differences were found in the two most similar morphs (i.e., morphs 1 and 2, see Fig. 3). For Morph 1, children with high CFRS scores generalized toxicity at a significantly higher rate ($p = 0.013$) than children with low CFRS scores. In the Low-CFRS group only, a difference was also found between the seasonality condition and both the toxicity and edibility conditions, respectively $p = 0.013$ and $p = 0.047$, with seasonality being generalized at a higher rate. For Morph 2 children in the Low-CFRS group generalizing seasonality at a significantly higher rate ($p = 0.011$) than children with high CFRS scores. A significant difference was also found between the seasonality and edibility conditions in the Low-CFRS group ($p < 0.001$), with seasonality being generalized at a higher rate.

Therefore, the three-way interaction noted above seems to be driven by the effect of information type for the morphs most similar to the target.