Linear integration and lexicographic models of choice: A cue weight learning perspective

Hrvoje Stojić*
Pompeu Fabra University

Henrik Olsson
Santa Fe Institute

Pantelis P. Analytis
University of Southern Denmark

Choosing between options characterized by multiple cues can be a daunting task. People may integrate all information at hand or just use lexicographic strategies that ignore most of it. Notably, integrative strategies require knowing exact cue weights, whereas lexicographic heuristics can operate by merely knowing the importance order of cues. Here we study how using integrative or lexicographic strategies interacts with learning about cues. In our choice-learning-estimation paradigm people first make choices, learning about cues from qualities of chosen options, and then estimate qualities of new options. We developed delta-elimination (DE), a new lexicographic strategy that generalizes previous heuristics to any type of environment, and compared it to the integrative weighted-additive (WADD) strategy. Our results show that participants learned cue weights regardless of the choice strategy employed. The group of people best described by the DE strategy learned cue weights, exactly like the group best described by the WADD. Still, there was an interaction between the adopted strategy and the cue weight learning process: the DE users learned cue weights slower than the WADD users. This work advances the study of lexicographic choice strategies, both empirically and theoretically, and deepens our understanding of strategy selection, in particular the interaction between the strategy used and learning the structure of the environment.

Keywords: decision making; cue weight learning; heuristics; strategy selection

Imagine you are an avid movie-goer and every couple of weeks you wind up in the local cinema looking for a good flick. There are several movies from which you can choose at each cinema visit, some that you would relish and others that you would loathe. Luckily, you have several useful pieces of information at your disposal to guide your choice: you can look at the average movie review ratings on a popular movie review website, the number of reviews, the director, the leading actor and the length of the movie. How do you go about picking the movie you will like the most, and how does this choice process interact with learning how useful various pieces of information are?

There are various ways in which people can use the available *cues* – pieces of information predictive of a variable of interest, a *criterion* variable – to choose among a number of different options. First, you could form an estimate of criterion value of each option by weighing and adding its cues and then choose the option that you gauge to have the highest value. This choice process summarizes the so-called weighted-additive strategy (WADD; Brehmer, 1994; Payne, Bettman, & Johnson, 1993). Alternatively, you could order cues on some measure of cue importance, compare the options on the first ranking cue and pick the one excelling on that cue (Take-the-best, see Gigerenzer & Goldstein, 1996;

Hogarth & Karelaia, 2005a) or eliminate options that are clearly inferior on that cue (Elimination-by-aspects, see Hogarth & Karelaia, 2005b; Tversky, 1972). If there is no difference or no *noticeable* difference on that cue (Δ -inference, see Luan, Schooler, & Gigerenzer, 2014), you could move to the next cue in the order and repeat these steps until you reach a decisive piece of information or are left with only one option, at which point you choose that option. In sequential search processes like these, choices are commonly based on the few decisive factors looked up first and, as a result, much of the information is ignored. Although the WADD strategy comes with the promise of accurate choices, it is also accompanied with a hefty mental effort price tag. Without compromising much on choice quality (or even at all, see Gigerenzer & Brighton, 2009; Hogarth & Karelaia, 2005a), lexicographic strategies can save a lot of cognitive effort, especially when facing demanding environments consisting of many options described on multiple dimensions.

Integrative strategies and lexicographic heuristics alike, rely on knowledge about the structure of the environment to attain their predictive potential. Yet the *type* of knowledge they require is fundamentally different. Integrative strategies require precise *values* of cue weight, representing (partial) correlations between each cue and the criterion value. By

contrast, lexicographic strategies need merely the *order* of cues, an estimate of how the cues rank in their ability to predict the criterion. Operating on cue orders has been argued to be one of the main advantages of lexicographic strategies (e.g. Gigerenzer & Brighton, 2009; Todd & Gigerenzer, 2000). The cue order approximates potentially complex cue-criterion relationships rather than trying to learn them precisely, such as with cue weights. Relying on cue order may prevent overfitting which tends to lead to worse performance when choosing among novel options. This argument has been a central-piece for explaining the surprising performance of the TTB heuristic (Gigerenzer & Brighton, 2009; Todd & Gigerenzer, 2000). Taken together, this implies an interaction between the strategy people use and the type of knowledge they acquire about the structure of the environment. People employing integrative strategies should learn cue weights, while those employing lexicographic strategies should learn cue order.

Hitherto, researchers have investigated the strategies that people should and do use (e.g., Gigerenzer & Goldstein, 1996; Hogarth & Karelaia, 2005a; Martignon & Hoffrage, 2002; Payne et al., 1993; Rieskamp & Otto, 2006). However, the interaction between choice strategies and learning about the environment has not been investigated so far. Our primary goal is to address this conceptual and empirical gap. We study the knowledge that people using either integrative strategies or lexicographic heuristics acquire about the structure of the environment—whether they learn cue weights or cue order. We achieve that by introducing a new choice-learning-estimation paradigm where participants first

complete a choice task, in which they simultaneously make choices and learn about cue weights or cue order through criterion value feedback. Using criterion value feedback allowed us to introduce a task where participants estimated the criterion values of new options. This estimation task immediately followed the choice task, and crucially, enabled us to identify whether participants learned the cue weights or only the cue order.

While there are no reasons to expect people using integrative strategies to learn cue order, there are several good reasons why people using lexicographic strategies might learn cue weights rather than only order. First, learning cue weights provides people with a more flexible handle on strategy selection—equipped with cue weights people can switch between integrative and lexicographic strategies at will. Hogarth and Karelaia (2007) showed that such flexibility can pay off: it is beneficial to adopt TTB at the beginning of a task, while cue weight knowledge is still poor, and switch to WADD later on when it is sufficiently developed. People then just need to learn to relate strategies with the environments where these strategies excel (Lieder & Griffiths, 2017; Rieskamp & Otto, 2006; Stojić, Olsson, & Speekenbrink, 2016). Second, cue weights can allow people to generalize to other types of tasks, beyond the choice task. In choice situations where criterion feedback is available cue weights allow predicting criterion values, something that is hard to achieve with cue order knowledge. Predicting the criterion value of an option has value of its own. For example, people might want to communicate their predictions to others, say quality of the movie (Analytis, Barkoczi, & Herzog, 2018). Benefits of this type of generalization might be large enough to make cue weight learning a default whenever possible. Finally, cue weight learning may provide an effective shortcut for learning cue orders, as the latter can be easily obtained from cue weights. A convincing theory for learning cue order directly is still lacking and most methods have been shown to be data hungry, improving strategy performance at a slow pace Dieckmann and Todd (2004); Todd and Dieckmann (2005). Algorithms for learning cue weights on a trial-by-trial basis have a long tradition (Rosenblatt, 1958; Rumelhart, Hinton, & Williams, 1986; Widrow & Hoff, 1960) and have been tested as models of human category and function learning (Gluck & Bower, 1988; Kelley & Busemeyer, 2008; Speekenbrink & Shanks, 2010), with convincing results about the match between the model and human behavior. Investigating the interplay between choice strategies and learning about the environment stands to improve our understanding of the nature of the strategy selection process. Obtaining a result that people using lexicographic strategies learn only cue orders would reveal a profound influence that choice strategies have on what people learn about the world. A result that people learn cue weights would on the other hand refine our view of strategy selection:

Hrvoje Stojić, Department of Economics and Business, Pompeu Fabra University; Henrik Olsson, Santa Fe Institute; Pantelis P. Analytis, Danish Institute of Advanced Studies and Department of Management and Marketing, University of Southern Denmark.

Early versions of the paper were presented at the Subjective Probability, Uncertainty and Decision Making conference in 2015 and the International Conference on Cognitive Modeling in 2016. The manuscript was made available as a PsyArXiv preprint at <https://psyarxiv.com/yj2zc/> (Stojic, Olsson, & P. Analytis, 2020). Project data, code and all related materials are publicly available at the Open Science Framework website: <https://osf.io/7yxrz/> (Stojic, Analytis, & Olsson, 2020).

We would like to thank Robin Hogarth, Gael Le Mens, Jörg Rieskamp, Thorsten Pachur, and Alexandros Gelastopoulos for their comments on previous versions of this manuscript. Hrvoje Stojić was supported by the Ministerio de Educación, Cultura y Deporte, Grant FPU12/05859.

*Correspondence should be addressed to Hrvoje Stojić, who is now at Max Planck UCL Centre for Computational Psychiatry and Ageing Research, University College London, 10-12 Russell Square, London, WC1B 5EH, United Kingdom. E-mail: h.stojic@ucl.ac.uk

people may choose strategies depending on the the weights they have learned and what these weights reveal about the choice setting.

Our second goal is to develop and test a lexicographic strategy that is well equipped for dealing with challenging choice environments. Most studies on inference, choice and learning have focused on problems with binary cues and two options (Gigerenzer & Goldstein, 1996; Nosofsky & Bergert, 2007; Rieskamp & Otto, 2006). In many real world settings, by contrast, people can choose among numerous options described by a number of cues, many of which are continuous (i.e. average movie rating or distance to the cinema). To increase the external validity in the study of lexicographic strategies we also need to devise heuristics that can cope with such challenging environments. Therefore, in our study we focus on an environment with three options described by continuous cues and criterion values, in which heuristics have not been put to test before.

What would be a good representative of lexicographic strategies that could perform competitively in settings with continuous cues and more than two options, and can match the ubiquitous character of WADD that can be applied across decision making settings? The TTB heuristic was developed for binary-valued cues (Gigerenzer, Hoffrage, & Kleinbölting, 1991) and collapses to a single-cue heuristic in environments with continuous cues. The Δ -inference heuristic effectively generalizes the principles of TTB to continuous environments by using minimal *noticeable* differences to control search (Luan et al., 2014). Still, the strategy can only be used to choose between two options. We propose *delta-elimination* (DE), a new lexicographic strategy that generalizes Δ -inference to any number of options by using an elimination mechanism, akin to that used in the elimination-by-aspects heuristic (Tversky, 1972). Since DE strategy nests Δ -inference and TTB, methods and insights produced in this study can be carried over to binary choice (or inference) problems with binary or continuous cues. We evaluate the DE strategy empirically as well, by comparing it to the WADD strategy in fitting choices of participants in our experiment. This development contributes to heuristic theories of choice and inference.

Choice strategies and learning the structure of the environment

Over the last decades, mathematical models and simulation studies have been used to evaluate the performance of different strategies across decision making contexts (e.g., Gigerenzer & Goldstein, 1996; Hogarth & Karelaia, 2005a). Depending on the structure of the (usually linear) environment (e.g., Hogarth & Karelaia, 2005b, 2007; Martignon & Hoffrage, 2002), and the cost of cognitive processing (e.g., Lieder & Griffiths, 2017; Payne et al., 1993) people might be well-advised to rely on integrative WADD or lexicographic

heuristic strategies. Experimental studies have followed up, assessing which of these strategies people actually use when they solve choice problems. In a typical experiment, participants are exposed to many trials with two or more options that are described by different cues, more or less predictive of the criterion value of the considered options. For instance, people may be asked to choose between two companies, that differ in their profitability (Rieskamp & Otto, 2006), described by different cue values on a number of informative dimensions (e.g., capital structure, financial resources, efficiency).

When it comes to knowledge about the structure of environment, a crucial input for integrative and lexicographic strategies respectively, simulation studies directly endowed the models with the required information (e.g. learning optimal cue weights or orders from data) and did not specify the exact learning process (e.g., Gigerenzer & Goldstein, 1996; Hogarth & Karelaia, 2005a, 2007; Şimşek & Buckmann, 2015). In a similar vein, most empirical studies explicitly provided information about cue weights or order (e.g., Bröder, 2000; Bröder & Schiffer, 2003a; Mata, Schooler, & Rieskamp, 2007; Newell & Shanks, 2003; Newell, Weston, & Shanks, 2003; Rieskamp, 2006; Rieskamp & Hoffrage, 2008; Rieskamp & Otto, 2006).¹ There is only a handful of studies that did not explicitly provide information about the cue weights or order. In these studies participants had to learn cue weights or order through, either (i) feedback about whether their choices were correct or not (Lee & Cummins, 2004; Newell, Rakow, Weston, & Shanks, 2004; Newell & Shanks, 2003; Nosofsky & Bergert, 2007; Pachur & Olsson, 2012; Rakow, Newell, Fayers, & Hersby, 2005), or (ii) feedback about both the choice correctness and the criterion values of the chosen option (Bröder & Schiffer, 2006a; Rakow, Hinvest, Jackson, & Palmer, 2004). Participants in these studies made better choices over time, a pattern that could be directly produced by learning either cue weights or cue orders. However, most of these studies did not examine the learning process beyond observing the raw learning effects, and focused on strategy selection instead (Bröder & Schiffer, 2006a; Lee & Cummins, 2004; Newell & Shanks, 2003; Nosofsky & Bergert, 2007; Pachur & Olsson, 2012).

A small subgroup of these studies has examined learning in greater detail, but focused exclusively on cue order learning. For example, in their experiments Rakow et al. (2004) and Newell et al. (2004) compared the cue orders that people reported and people's information search to several objective statistics according to which people might be ordering

¹Cue weights in many of these studies were represented with cue *validities*, a measure of cue goodness in choice environments where choice correctness information is available, rather than criterion value of the options. Cue validity is usually computed as $R/(R + W)$, where R and W are the number of correct and incorrect choices, respectively, based on the cue.

the cues – cue validity, success or discrimination rate ordering. They found that success measure corresponds better to people's self-reports and search patterns. Finally, [Rakow et al. \(2005\)](#) compared the learned cue order to expected information gain, an information theory measure, and simple cue-outcome correlations, alongside the usual cue validity ordering. They found that cue-outcome correlation ordering corresponded the best to the ordering reported in the experiments. None of these studies compared learning between different strategy users, how strategy selection might interact with what is being learned or interact with the learning process. They also did not postulate a mechanism of how cue weights or cue order might be learned.

The lack of a convincing theory for learning cue orders has been the subject of an important critique of the TTB heuristic. [Shanks and Lagnado \(2000\)](#) opened the critique by suggesting that the TTB heuristic is deceptively simple, hiding its complexity in the knowledge of cue orders. [Gigerenzer et al. \(1991\)](#) originally suggested that cue validities, and thereby order, might be learned by a simple frequency counting mechanism. Yet [Juslin and Persson \(2002\)](#) argued that much effort is needed to order cues as many events have to be stored in memory (see also [Dougherty, Franco-Watkins, & Thomas, 2008](#); [Newell, 2005](#)). There have been several lines of research grappling with this critique. First, [Dieckmann and Todd \(2004\)](#) studied a series of simple methods for learning cue order from data and found that although learning cue orders directly is in principle possible, numerous observations are required to improve strategy performance. More recently, Gigerenzer and colleagues (see [Gigerenzer, Hoffrage, & Goldstein, 2008](#); [Hutchinson & Gigerenzer, 2005](#)) pointed out that cue order knowledge is not necessarily driven by individual learning – it can be hard-coded by evolution, or transmitted socially. Short after, [Garcia-Retamero, Takezawa, and Gigerenzer \(2009\)](#) showed experimentally that social knowledge can indeed be beneficial – imitating the best decision maker or majority substantially accelerates individual learning. Regardless of the benefits of social knowledge transmission, knowledge about cue orders still has to be learned by some individuals (see [Rogers, 1988](#)). In a different line of research, [Katsikopoulos, Schooler, and Hertwig \(2010\)](#) have provided some evidence that cue order might not be that crucial for the accuracy of the TTB heuristic. When very small sample sizes were used to estimate cue order and directions (10 observations or fewer), the correctness of cue directions rather than order correlated highly with accuracy. The argument does not seem to hold for larger samples, however, and even in small samples some information about the cue order is needed for the TTB heuristic to achieve higher accuracy. In summary, these studies focused on identifying channels for acquiring cue order information other than individual learning ([Garcia-Retamero et al., 2009](#); [Gigerenzer et al., 2008](#), except for [Dieckmann](#)

and [Todd \(2004\)](#)) or argued order information is actually not needed ([Katsikopoulos et al., 2010](#)). All of them ignored possible interaction with strategy selection. Given that some studies did observe individual learning effects, the question of how individual learning of the cue information required for their choice strategies interacts with strategy selection still remains pertinent.

A key missing ingredient for studying the interaction between strategy selection and learning are plausible mechanisms for learning cue weights and/or cue order. We looked for inspiration in neighboring fields. In statistics, the least-mean-squares (LMS) algorithm has been a popular formalization for updating regression (i.e. cue) weights with each observation; the algorithm achieves that by leveraging the discrepancies between predicted and observed outcomes ([Rumelhart et al., 1986](#); [Widrow & Hoff, 1960](#), often called the Widrow-Hoff or Delta rule). Studies of human category and function learning have demonstrated that the LMS model is a plausible model of human learning (e.g., [Gluck & Bower, 1988](#); [Kelley & Busemeyer, 2008](#); [Speekenbrink & Shanks, 2010](#)). The settings in which these LMS models have been applied are similar to multi-cue choice tasks used to study integrative and lexicographic strategies. For instance, [Gluck and Bower \(1988\)](#) examined category learning in a task that presented stimuli described on four cues with binary cue values and provided correct/incorrect feedback. Participants had to learn the relationship between the cues and correct responses. Categorization task such as the one examined by [Gluck and Bower](#) and inference and choice tasks (i.e. city population task) that have been extensively used to test lexicographic heuristics ([Gigerenzer & Goldstein, 1996](#)) have the same zero-one error function (a categorization, inference or choice will be right or wrong). Further, [Kelley and Busemeyer \(2008\)](#) and [Speekenbrink and Shanks \(2010\)](#) examined how people learn functions when they are presented with a single stimulus per trial and they have access to multiple continuous-valued cues that they can use to predict a continuous criterion. In these tasks, the participants observed criterion value after each prediction, not unlike choice experiments where participants receive feedback about the criterion value on the chosen option (as in, for example [Bröder & Schiffer, 2003a, 2006a](#); [Hogarth & Karelaia, 2007](#)). In our study the LMS model learns cue weights based on observed cue and criterion values and cue order can be obtained as a byproduct of learning cue weights. Cue orders can be easily extracted by ordering absolute cue weight values. Thus, the LMS model provides an attractive alternative to learning cue orders directly ([Dieckmann & Todd, 2004](#)), allowing to easily go back and forth between integrative strategies that need cue weights and lexicographic strategies that require order, adaptively switching depending on each strategy utility ([Lieder & Griffiths, 2017](#); [Rieskamp & Otto, 2006](#); [Stojić et al., 2016](#)).

Here we built on previous theoretical and empirical work on the LMS model. Following function learning studies that examined mapping between multiple cues and continuous response (Kelley & Busemeyer, 2008; Speekenbrink & Shanks, 2010), participants in our experiment received criterion value of the chosen option as a feedback in a choice task (same as in Bröder & Schiffer, 2003a, 2006a; Hogarth & Karelaia, 2007; Rakow et al., 2004). An estimation task followed the choice task where participants predicted criterion value of new options that they had not seen during the choice task, and did not receive feedback after their predictions. The estimation task is essentially equivalent to a test phase commonly deployed in function learning tasks (Kelley & Busemeyer, 2008; Speekenbrink & Shanks, 2010). Introducing an estimation task after the choice task was a key innovation in our “choice-learning-estimation” paradigm. It allowed us to uncover the knowledge participants acquired during the choice task, and crucially, to identify whether participants learned cue weights or cue order. In addition, to independently cross-check the inferences drawn from the estimation task, the participants gave self-reports on the strength of the relationship between the cues and the criterion in regular intervals during the choice task. Predictions in the estimation task allowed us to fit the LMS model and obtain evolution of participants’ cue weights during the choice task. We then used the cue weights inferred by the LMS model to fit the WADD and DE models. Classifying participants as WADD or DE users then made it possible to compare the type and precision of knowledge acquired by the two groups.

The formal models

We describe the formal models with respect to their application to the choice-learning-estimation (from now on referred as CLE) experiment illustrated in Figure 1. In the first part of the CLE task participants repeatedly choose among three options; each varying along four continuous-valued cues and continuously along a criterion variable. Let $i = 1, \dots, N$ index the options and let $j = 1, \dots, M$ index the cues. By $x_{j,t}^i$ we denote the value of the i -th option for the j -th cue at trial t . Criterion value of the i -th option at each trial t , Y_t^i , is a simple linear combination of option’s cue values, $x_{j,t}^i$, and cue weights w_j

$$Y_t^i = \sum_{j=1}^M w_j x_{j,t}^i.$$

After each choice participants learn the criterion value, Y_t^i , of the option they selected. Their goal is to maximize the sum of criterion values. To achieve that they can use the feedback from their choices to learn about cue weights or cue order. At each step, the acquired knowledge feeds in the choice models and results in selecting one of the options. After the choice task decision makers get a surprise estimation

task where they repeatedly observe a single option and have to predict its criterion value.

Below we first introduce the LMS model with which we model the cue weight learning process. Key to this is the estimation task that gives us the observations required for fitting the LMS model. Next we describe the two choice models that we considered in this study, the WADD strategy and the heuristic DE strategy, with a focus on how these models employ cue weights inferred by the LMS model. We focus on the WADD and DE strategy as they can be taken as key representatives of the two qualitatively different classes of strategies – the integrative and lexicographic strategies – that we are interested in. Many of the other strategies lie in between and some can even be expressed as special cases. For instance, equal weighting strategy is a special case of the WADD model and TTB heuristic is a special case of the DE strategy. We describe baseline models that we use as benchmarks in each task in Appendix A, while parameter estimation and model selection procedures can be found in *Supplemental Material*.

LMS Model

We use the LMS model, originally used by Gluck and Bower (1988) for a category learning task, and subsequently for predicting continuous-valued outcomes in function learning tasks (e.g. Kelley & Busemeyer, 2008; Speekenbrink & Shanks, 2010). The model is a single-layer feed-forward network with one input node for each cue and a single output node. Cue weight learning is modeled as a dynamic process where weights are updated from trial to trial based on prediction error, that is difference between predicted and true criterion value.

Input nodes in each trial t are activated by cue values $x_{j,t}$, and the output node represents the predicted criterion value $\hat{Y}_t^i$:

$$\hat{Y}_t^i = \sum_{j=1}^M x_{j,t}^i u_{j,t}, \quad (1)$$

where $u_{j,t}$ denotes cue weight learned for cue j at trial t .

Weights are updated in every trial by the delta rule based on the difference between predicted criterion value, $\hat{Y}_t^i$ and true criterion value, Y_t^i , that participant receives as a feedback during the choice task for chosen option i at trial t :

$$u_{j,t+1} = u_{j,t} + \eta / t^\gamma (Y_t^i - \hat{Y}_t^i) x_{j,t}, \quad (2)$$

where $\eta \in (0, 1)$ is a learning rate and $\gamma \in (0, 1)$ specifies rate of decay of the learning rate. We initialize the weights to $u_{j,0} = 0$ for all j . If γ is set to zero, learning rates will be constant through all the trials, while for positive values learning rate will converge to zero over time. Thus, the LMS model provides us with cue weights for each trial, $u_{j,1}, \dots, u_{j,T}$. Note that the cue weight learning process is based only on the option for which participants receive feedback.

We allow for decaying learning rates as in previous experiments it has been found that the LMS model with a decay parameter can describe participants predictions better than without it (Kelley & Busemeyer, 2008). This version could also potentially capture a different learning mechanism based on Bayesian principles, where learning rate decreases with decrease in uncertainty about weight estimates (see, for example, Speekenbrink & Shanks, 2010). Hence, we have two versions of the LMS model – LMS without decay (where γ was set to zero) and LMS_d with decay.

A key innovation in our CLE paradigm is that LMS models start learning the cue weights on data from the choice task, by observing the cues and criterion values of chosen option. These cue weights are then frozen and used to make predictions in the estimation task (participants do not receive feedback here), which we use to fit the model and obtain the parameters (for more details see *Supplemental Material*). Resulting fits then allow us to generate cue weights evolving throughout the choice task trials that we can use as input for the choice models.

The weighted-additive model

The weighted-additive (WADD) decision making model (Payne, Bettman, & Johnson, 1988) takes into account information about all cues when making a decision. It is often considered a rational decision making and contrasted with lexicographic heuristic models such as TTB (Gigerenzer & Goldstein, 1996).

In our setup, the WADD model uses the knowledge about cue weights completely by combining the weights learned by the LMS model with cue values in a linear fashion to produce predicted value of each option i in trial t : $\hat{Y}_t^i = \sum_{j=1}^M x_{j,t}^i u_{j,t}$. The WADD model then deterministically decides by maximizing among the options: $C_t^{\max} = \arg\max_i \hat{Y}_t^i$, $i \in 1, \dots, N$.

To fit the model to choice data we need to assume additional error structure (Loomes, Moffatt, & Sugden, 2002; Rieskamp, 2008). We assume a so-called “trembling hand” error which postulates that participants choose the best option with probability $1 - \epsilon$, otherwise they mistakenly choose either of the three options with probability ϵ . The final probability that an option i is chosen, $P(C = i)$, taking into account the tremble error, ϵ , is given by:

$$P(C_t = i) = (1 - \epsilon) \times 1_{C_t^{\max}=i} + \frac{\epsilon}{N}, \quad (3)$$

where $1_{C_t^{\max}=i}$ is an indicator function that equals 1 if the maximizing option is option i , 0 otherwise.

The Delta-Elimination model

We consider delta-elimination (DE), a strategy that relies on the sequential consideration of cues, as in the TTB and Δ -inference heuristics (Gigerenzer & Goldstein, 1996; Luan et al., 2014) and is endowed with an elimination mechanism

as in elimination-by-aspects heuristic (Tversky, 1972). The strategy can be described formally as follows: let $i = 1, \dots, N$ index the options and let $j = 1, \dots, M$ index the cues from the most informative one to the least informative one. By $x_{j,t}^i$ we denote the value of the i -th option for the j -th cue at trial t and by $x_{j,t}^{\max}$ we denote the value of the best option on the j -th cue at trial t . Finally by Δ we denote a discriminatory threshold. We first eliminate all options such that $x_{1,t}^i < x_{1,t}^{\max} - \Delta$. If there are more than one options left, we proceed by eliminating all options such that $x_{2,t}^i < x_{2,t}^{\max} - \Delta$ and so on. This process continues until there is only one option left or until we run out of cues, where an option is selected at random from the non-eliminated options. We denote the chosen option as $C_t^{\max}$.

In our experiments we assume that in each trial the model orders the cues $i = 1, \dots, 4$ using the weights $\mathbf{u}_t$ provided by the LMS model: it first takes the absolute value of the weights, making the negative weights as predictive as the positive ones, and then ordering them from the largest weight to the lowest, producing a cue order r_t that serves as an estimate of the relative importance of different cues. Further, the values of cues with negative weights are multiplied with -1 . In sum, using only the order and the signs of the weights, the lexicographic model can perform well without knowing the exact cue weights.

The DE strategy effectively generalizes existing lexicographic strategies to settings with any type of the cues (binary, discrete or continuous) and any number of available options. The parameter Delta controls how frugal the strategy will be. For example, a small Delta value ensures that in large choice sets the model proceeds quickly in trimming down the choice set to the most promising options. The main building blocks of the strategy—sequential search and the use of perceptual thresholds—are common cognitive mechanisms in humans and have received ample empirical support in experimental studies (e.g. Bettman, Johnson, & Payne, 1990; Kohli & Jedidi, 2007). Thus, the strategy is both psychologically plausible and can provide a strong competitor to compensatory strategies in realistic choice problems with continuous cues and more than two options.

Similar to the WADD model, DE makes deterministic choices, so we add a trembling hand error term to arrive at final choice probability (Equation 3).

Method²

Participants

Seventy-eight participants (49 female, $M_{\text{age}} = 21.8$, age range: 17–54 years) took part in the experiment, recruited from the Pompeu Fabra University subject pool. They received a show-up fee of five Euros and a performance dependent bonus, an additional amount of 6.8 Euros on average. The experiment lasted for 43 minutes on average. We excluded from the analysis one participant that provided the same response in all trials in the estimation task, leaving us with 77 participants.

The choice-learning-estimation paradigm

The choice-learning-estimation (CLE) paradigm consisted of two parts: a choice task and an estimation task. In the choice task participants made choices and learned about the structure of the environment through outcome feedback. Estimation task that came after the choice task tested what they have learned during the choice task. This combination of the two tasks was a key innovation of the paradigm.

In the choice task participants repeatedly chose between three options (Figure 1a). Options' criterion values were a function of four observable cues and participants observed only the criterion value of the option they chose. The goal was to maximize the sum of criterion values of the chosen options across trials. Criterion values of options, Y , were a simple linear combination of cue values, x , and cue weights w

$$Y = 4x_1 - 3x_2 + 2x_3 - x_4 + e, \quad (4)$$

where cue weights, $w = (4, -3, 2, -1)$, were kept fixed throughout the task and e was normally distributed error with mean 0 and standard deviation of 30. Positive weights indicate that the higher the cue values are, the higher the value of the option on average. We opted for some negative objective weights, as negative linear relationships are harder to learn and people seem to have a prior for positive linear relationships (Brehmer, 1994; Busemeyer, Byun, Delosh, & McDaniel, 1997). Hence, the negative weights would facilitate detecting patterns of true learning in the estimation task and insight questions.

For the choice task we generated 480 unique options by sampling cue values x from uniform distribution with range from 10 to 90, rounding them to the nearest integer, and computing the criterion values using Equation 4. Finally, we allocated options randomly across 160 trials, three options per trial. With this procedure we ensured that choosing any of the three options in the trial would yield on average equal amount of information with respect to learning the cue weights. Mean cue inter-correlation was -0.02 ($SD = 0.05$), ranging from a minimum of -0.08 to a maximum of 0.04 . To

reduce between-subject variance the stimuli were drawn once and all participants received the same set of stimuli. That is, the participants saw the same options in each trial.

The estimation task was comprised of 40 trials, where in each trial participants saw the cues characterizing a single option and their goal was to predict the criterion value of the presented option (Figure 1b). To prevent people from learning in the course the estimation task no feedback was provided about the guessed criterion values. We incentivised truthful reporting of predictions by computing the payoff as a difference between the prediction P and the option's criterion value Y : $200 - |P - Y|$.

Half of the 40 trials were *interpolations*, drawing the option's cue values from the same uniform distribution as in the choice task, $U(10, 90)$, and the remaining half were *extrapolations*, drawing cue values from two intervals at the extreme ends, $U(0, 10)$ and $U(90, 100)$, that had not been experienced in the choice task. We then computed the criterion values according to Equation 4. We included extrapolation trials as they can be useful for discriminating between the models describing the learning process (e.g. DeLosh, Busemeyer, & McDaniel, 1997). Same as in the choice task, here the stimuli was drawn once before the experiment and all participants received the same set of stimuli.

Insight questions

We divided the choice task into four blocks of 40 trials. After each block, participants answered insight questions where they reported their knowledge about the cue importance on simple Likert scale (Lagnado, Newell, Kahan, & Shanks, 2006; Rakow et al., 2004; Speekenbrink & Shanks, 2010). Insight questions provided us with another window into participants' learning during the choice task, complementary to the estimation task that participants completed only after the choice task. The scale used in these questions is arbitrary. Therefore, there is no clear-cut mapping to the objective cue weights. Nonetheless, these self-reports can inform us about participants' subjective cue order and evolution over time.

Lagnado et al. (2006) used similar insight questions in a category learning task and systematically compared experiment with and without insight questions during learning. They found that insight questions did not affect either learning or strategy selection. This implies that insight questions should not affect learning in our task as well.

²In this experiment we had four other conditions that we will analyze and report separately. We designed these conditions to tackle a different question than the one we ask here: whether the act of choosing *per se* confers an advantage for acquiring knowledge about the environment.

Figure 1. Screenshots from the choice-learning-estimation (CLE) paradigm. The CLE paradigm consists of two parts: a choice task followed by an estimation task. **(a)** In the first part participants encountered a choice task where in each trial they chose between three cheeses described by four cues – “Lactic”, “Acetic”, “Casein” and “Texture”, and observed the enjoyment units (i.e., criterion value) of the chosen cheese. The participants’ goal was to maximize the amount of enjoyment units collected throughout the 160 trials in the task. Unknown to them, enjoyment units were a linear function of the four cues. **(b)** In the second part participants completed an estimation task. Here we presented them with a single item in each trial, but now they had to predict the enjoyment units of the item. Participants were maximizing their earnings by making accurate predictions across 40 trials, matching the true criterion values of observed items. We did not provide them with feedback in this part.

Procedure

Participants completed the experiment on desktop computers, using a custom software written in Python, with the help of PsychoPy library (Peirce, 2007). The experiment was implemented in Spanish. We conducted the experiment in groups of about 15 people in BES laboratory at Universitat Pompeu Fabra.

We presented instructions about the experiment on the screen. We used a background story of a cheese loving consumer who in every trial faces a choice between three cheeses and whose goal is to maximize their enjoyment in cheese. We informed participants about the cues and the range of values they could take, and explicitly mentioned they could use this information in making their choices. We did not tell them anything else about the relationship between cue values and criterion values. We stressed that in each trial they would get three new cheeses that differ in their cue values. Participants had three practice trials where they could familiarize themselves with how the interface works. To motivate participants, we told them it is a difficult task, but that they will improve with practice.

During the choice task, every 40 trials participants answered insight questions. In the initial instructions we mentioned that we will interrupt the task on several occasions and ask them about their experience until that point, but without specifying details. When insight questions appeared during the choice task we gave the following instruction: “Please rate how strongly each feature relates to the amount of enjoyment (EU’s) you get from the cheese. Please rate each feature on a scale from -10 (highly negative) to 10 (highly positive).” We presented questions for all four cues on the

same screen, in the same order that was used to present the options.

We announced in the initial instructions that there will also be a second part of the experiment (the estimation task), saying only that it is much shorter than the first part and that performance will still count for their earnings. We did not inform them about the nature of the task, so the participants did not know that they would need to estimate the exact criterion values of the options. When they reached the second part of the experiment we eventually informed them about the details of the task, mentioned that we will present them a single cheese and ask them to estimate how tasty the cheese is going to be in terms of enjoyment units. We also mentioned that this time they will not be receiving any feedback.

We did not inform the participants about the exact number of trials in each task, instead we told them that it takes 60 minutes on average to complete the entire experiment. We carefully explained how we will calculate their earnings, illustrating the process with examples. We used enjoyment units (EU’s) as a token in the experiment, and we stated that EU’s will be transformed into Euro according to an exchange rate of 5000 EU’s for 1 Euro. Participants started with an initial fund of 5000 EU’s. During the choice task we displayed the total earnings on the screen at all times, but not in the estimation task. At the end of the experiment we showed the participants a breakdown of EU’s earned in each phase and the earned sum in Euro.

For each participant we randomized the order in which options were presented on the screen and the order of cues at the beginning of the experiment and then kept the ordering fixed throughout. Attribute labels, “Lactic”, “Acetic”, “Casein” and “Texture”, were also randomly attached to under-

lying cues for each participant to reduce the influence of any prior information participants might have had about them.

For the sake of transparency we list all the variables we recorded: age, gender and educational background of the participants, their choices and insight ratings in the choice task, predictions in the estimation task and response times in both tasks.

Data and code availability

The data, code used for analyzing the data and other project related files are publicly available at the Open Science Framework website: <https://osf.io/7yxrz> (Stojic, Analytis, & Olsson, 2020).

Results³

Participants are learning cue weights rather than order

Participant's performance in the estimation task is uniquely suitable for identifying whether they learned more than just the cue order during the choice task. Most participants were good at predicting the criterion values of options presented in the estimation task, as evidenced by the root mean square error (RMSE) and the correlation between each participant's predictions and criterion values (Figure 2a; LMS aspects of the figure will be addressed below). What would be a baseline performance in this task? Participants that learned only cue orders during the choice task, or did not learn how cues are related to criterion values at all, could of course make random guesses. They could, however, use a more sensible strategy that does not involve a sophisticated learning: keeping an estimate of a mean criterion value based on feedback in the choice task, and using this mean as a fixed prediction in the estimation task. We implemented this strategy in what we call a Baseline prediction model, one variant that perfectly computes the mean (BP_f) and another that learns the mean slowly trial-by-trial (BP_l ; see Appendix A for details on benchmark models). The mean RMSE between such baseline predictions (best fitting BP model) and criterion values (202, $SE = 0.54$) was larger than the observed mean RMSE (154, $SE = 4.59$) across participants (one-tailed Wilcoxon signed rank test: $\text{Median}_{\text{diff}} = -51.45$, $Z = 118$, $p \leq .0001$, $CI = [-Inf, -43.7]$). Another way to look at the estimation task performance is through correlations between estimates and criterion values, that speak more toward consistency, rather than how close predictions were to criterion values. The mean Spearman correlation (0.63, $SE = 0.07$) was substantially higher than zero (one-tailed Wilcoxon signed rank test: $\text{Median}_{\text{diff}} = 0.67$, $Z = 3000$, $p \leq .0001$, $CI = [0.61, Inf]$). Figure 2a also shows that individual variation in performance was substantial – while most participants were doing well, exhibiting high correlations and low RMSE's, 14 participants performed poorly in

the estimation task (worse or in par with the baseline prediction model), indicating that they did not acquire knowledge about the cue weights in the course of the choice task. Note that we designed the environment to be fairly predictable, with a relatively small error component. This can be seen from a high correlation between criterion values with and without error component for chosen items: mean across participants was 0.97 ($SE = 0.02$). Despite having an environment that in principle should be possible to learn, by the end of the choice task most participants were still far from perfect performance.

We obtained a more nuanced picture about participants' performance in the task by examining mean predictions across participants for each option (stimulus) in the estimation task (Figure 2b; the LMS aspects of the figure will be addressed below). Averaging the results across all options when computing RMSE and correlations can potentially obscure large differences between individual options. Figure 2b illustrates whether this was the case. We ordered options according to their criterion values and what is immediately obvious is that participants predictions were very close to criterion values for options with middle range criterion values, approximately from 0 to 200. Larger deviations occurred for options with more extreme values. Participants performed on average worse on items with extreme negative value than on items with extreme positive ones. This was likely due to having less experience with options with low criterion values. Participants were motivated to choose options with high criterion values in the choice task and over time they acquired a somewhat biased sample and had less observations about options with low values (Stojic, Schulz, P. Analytis, & Speekenbrink, 2020). Even in the edges of the range, however, participants' average predictions were going in the right direction. They were much better than those of the baseline prediction (BP) models, that we would expect as answers from participants knowing only the cue order, and they were relatively close to the true criterion values.

Insight questions provided us with another window into participants' learning in the choice task. Every 40 trials in the choice task we asked participants to rate the strength of each cue's effect on option value. Previous research using such questions showed that people have good insight into what they have learned (Lagnado et al., 2006; Rakow et al., 2004; Speekenbrink & Shanks, 2010). From mean ratings for all four cues we see that participants got the order and directions partially right on average (Figure 3a). Ordering and direction were roughly established already after the first block of trials and the pattern got clearer as the choice trials

³All the analysis and visualization of the results was implemented in R programming language, with the help of additional libraries (Auguie, 2012; R Core Team, 2017; RevolutionAnalytics, 2014; Sharpsteen & Bracken, 2015; Wickham, 2007, 2009; Wickham & Francois, 2015; Wuerz, 2013).

Figure 2. Participants learned cue weights rather than order. **(a)** Participants made good predictions in the estimation task, as indicated by the mean Spearman correlation and root mean square error (RMSE) between predictions and criterion values. Predictions were much better than what can be expected from knowing only the cue order – correlation of zero and the RMSE between the best fitting baseline prediction (BP) model and criterion values (continuous vertical line indicates the mean and shaded region the standard deviation). Points are individual results, while dashed horizontal and vertical lines are means across all participants. The different point shapes indicate the best fitting model for different participants – one of the BP or the least-mean-square models (LMS and LMS_d). **(b)** Mean predicted criterion values for individual options presented in the estimation task. Participants predicted well criterion values around the middle range (between 0 and 150), but overestimated criterion values at the lower end and underestimated them at the upper end. Options are sorted according to their criterion value (approximately diagonal black line), going from lowest to highest. Additional points denote mean predictions of the best fitting LMS model across participants (those fitted best by the BP model were excluded). Continuous horizontal line stands for the mean prediction for the best fitting BP model (line indicates the mean and shaded region the standard deviation). Standard errors of the means are smaller than the points.

progressed. Participants did not learn that the second cue had larger weight (although negative), than the third cue. Moreover, participants did not detect that the fourth cue was negative, instead they thought it was weak but positive. These results echo classical function learning findings – negative linear relationships are harder to learn than positive linear ones (Brehmer, 1994; DeLosh et al., 1997). Insight questions used an arbitrary scale that does not map directly to objective cue weights. Thus, it would be premature to infer whether participants knew cue weights or only cue orders on the basis of insight questions alone and without an estimation task. The ratings are informative about the learning process, however, regardless of the exact type of cue knowledge. The ratings changed over time, suggesting that slow, gradual learning characterizes both types of cue knowledge (Dieckmann & Todd, 2004; Gluck & Bower, 1988).

The estimation task results established that most participants learned cue weights—to perform well in this task participants had to know cue weights rather than just cue order. Insight questions ratings, in particular their adjustment over

time, indicate that cue knowledge evolved gradually during the choice task. Both the estimation task and the insight questions showed that participants’ knowledge was far from perfect. Still, taken together these results suggest that participants may have learned cue weights irrespective of the choice strategy that they adopted. Since it is possible that all or most participants used a single strategy, we have to postpone a stronger conclusion until we analyze what strategies participants adopted. These results also justify our next step: modeling a mechanism of how participants learn the cue weights.

Least-mean-squares learning as a mechanism for learning cue weights and order

We inferred participants’ subjective cue weights, and based on them the cue order as well, by fitting the LMS model to the estimation task while taking into account experiences in the choice task. This was a key innovation in our CLE paradigm: the LMS model learns weights from cues and criterion values that participants observed during

Figure 3. Participants knowledge of cue weights and order evolved during the choice task, but was still imperfect by the end of it. **(a)** Self-reports in insight questions that periodically appeared during the choice task (after trial 40, 80, 120 and 160) showed that participants acquired reasonably good knowledge of the cue order on average. The insight question scale captures how strongly each feature relates to the criterion values, on an arbitrary scale going from -10 (highly negative) to 10 (highly positive). The points indicate mean rating across participant while error bars represent standard errors of group means. To allow for easier discrimination, points are displaced horizontally and connected by lines. **(b)** Evolution of subjective cue weights inferred by the best fitting least-mean-squares model showed that participants learned the cue weights reasonably well, but did not reach the objective weights used for generating the stimuli. Lines are median cue weights across participants for each trial, smoothed with a moving average of ten trials, disregarding participants best fitted by one of the baseline models.

the choice task while following their choice strategies, but the evolution of cue weights was uncovered using the responses participants provided in the estimation task—these appropriately constrained the model.

We fitted two versions of the LMS model to each individual's data – LMS and LMS_d where learning rate also decays over time (see **The formal models** Section for details on LMS models and **Supplemental Material** for details on fitting procedure). To select the best model we used the Bayesian information criterion (BIC, the smaller the better; Schwarz, 1978), as it takes into account differences in complexity of the models. Both versions fitted participants almost equally well on average (Table S1 in **Supplemental Material**), with slightly worse mean BIC score of 369 ($SE = 2.87$) for LMS than 365 ($SE = 2.85$) for LMS_d. For comparison, the best performance that can be achieved with these models when fitted to perfect predictions (i.e. equal to criterion values; see perfect LMS and LMS_d benchmark models in Appendix A), was a mean BIC score of 319 ($SE = 1.06$) and 306 ($SE = 0.98$) respectively. The LMS models outperformed the BP models, with a mean BIC score of 384 ($SE = 3.33$) for BP_l and 405 ($SE = 2.77$) for BP_f. Out of 14 participants that showed poor performance in the estimation task, 8 were best fitted by one of the BP models. Overall, the LMS model fits were closer

to the BP model fits than to the best performance that can be achieved using the LMS learning, suggesting that participants' learning was comparatively slow.

How far were the model predictions from participants' responses in the estimation task? The RMSE for LMS (100.4, $SE = 3.25$) was slightly worse than for LMS_d (91.5, $SE = 3.12$; two-tailed Wilcoxon signed rank test: Median_{diff} = 6.66, $Z = 2967$, $p \leq .0001$, $CI = [4.42, 8.89]$), both being substantially closer to participants predictions than the baseline prediction models, with RMSE of 157.1 ($SE = 4.98$) for BP_f and 123.2 ($SE = 4.32$) for BP_l. These differences were also reflected in the number of best fitting individuals per model. The largest number of participants was best fitted by the decay version (35, LMS_d), followed closely by non-decaying learning (31, LMS), and a much smaller subset by one of the BP models (11). When averaged only across participants best fitted by a model, the mean BIC weight of the model fit gives the mean probability with which that model was the best among the competing models (Wagenmakers & Farrell, 2004). In other words, it tells us how confident we can be the categorization of participants was correct (Bröder & Schiffer, 2003a). For example, a mean BIC weight close to 0 would mean negligible evidence for the model, while value close to 1 would mean that the evidence for best fitting model

was overwhelmingly better than for the competing models. The mean BIC weight was 0.75 ($SE = 0.02$) for the participants best fitted with LMS and 0.89 ($SE = 0.02$) for participants best fitted with LMS_d. This indicates that even though both LMS models fitted behavior equally well on *average*, there were substantial difference between the model fits at the individual level and we can be confident in our categorization of participants into LMS or LMS_d users.

Next, we made predictions for the estimation task options using the best fitting LMS models (Figure 2b). Mean predictions from the LMS models were close to the actual participants' predictions for most of the options. This was the case regardless of their criterion values, whether these were options with extreme values for which participants were less accurate or with middle values for which they were more accurate. This shows that the LMS models captured participants' predictions well. Model predictions were expectedly better for interpolation items (mean RMSE between participants' and model predictions: 77.6, $SE = 2.92$) than for extrapolation items (101.7, $SE = 3.97$; two-tailed Wilcoxon signed rank test: Median_{diff} = 23.07, $Z = 2645$, $p \leq .0001$, $CI = [16.96, 29.14]$).

We then examined in more detail how cue weight knowledge inferred by the best fitting LMS models evolved over time during the choice task (Figure 3b). On average participants' knowledge improved over time, but learning was slow and did not converge to the objective cue weights – distances from objective cue weights ranged from 0.91 to 2.09. Participants had easier time learning the positive weights than the negative weights, reflecting a well known finding in function learning literature (Brehmer, 1974). We also cross-checked the evolution of cue weights with patterns observed in the insight questions reports. In the LMS_d model the positive decay parameter values are compensated by faster learning rates (Table S1 in *Supplemental Material*). This results in fast learning in the beginning of the choice task and slower changes later on. By contrast, the LMS model without decay learns slowly but steadily across the task, and has relatively poor cue weight knowledge in the first stages of the choice task (Figure S2a in *Supplemental Material*). The insight question reports of participants best fitted with the LMS and LMS_d model showed that these qualitative differences in learning patterns between LMS and LMS_d model are accurately reflected in the insight questions as well (Figure S2b in *Supplemental Material*). This cross-task comparison provides additional evidence that the LMS models successfully captured the cue weight learning during the choice task.

Lexicographic heuristics need cue order rather than cue weights. Although the LMS model does not learn cue order directly, these can be easily derived from cue weights. Thus, we examined the evolution of cue orders by looking at the inferred cue weights. For most people this order of cue weights was established quite fast. Participants established

the first ranking cue in 21 trials ($SE = 4.02$) on average, that is, the rank stopped changing after that. The rank of other cue weights stopped changing later on, after 52 ($SE = 5.41$) trials for the second cue, 50 ($SE = 5.98$) for the third and 48 ($SE = 4.62$) for the fourth cue on average. Notably, individual variability on when ranks were established was substantial. The correctness of the established cue order varied - first ranking cue was correct in 0.99 ($SE = 0.01$) of cases, fourth in 0.57 ($SE = 0.06$), while second and third were confounded, being correct in 0.29 ($SE = 0.05$) and 0.29 ($SE = 0.05$) of cases. This pattern roughly corresponded to the insight questions results as well (Figure 3a).

The LMS model provided a good fit to the estimation task data. It also provided a reliable description of the evolution of cue weight knowledge during the choice task, as evidenced by cross-checks with insight question reports. Cue weights and order are crucial inputs for choice strategies. We next analyzed differences in knowledge of cue weights and order between participants using integrative and lexicographic strategy.

Differences in knowledge between WADD and DE strategy users

We examined the raw choice performance first and how it is related to performance in the estimation task. Cue weight knowledge is a crucial input for the choice strategies and the quality of choices depends on it. Therefore, we expected a high correlation between performances in the two tasks. We measured participants' performance in the choice task as the mean rank of options chosen in each block of 40 trials in the task. Rank was computed according to the expected values of the options present in a trial. Correct choices in all trials would yield mean rank of one, while random choices would yield a mean rank of two. The choice performance improved over time (Figure 4a; results related to WADD and DE perfect models are explained below), from 1.69 ($SE = 0.03$) in the first block, to 1.50 ($SE = 0.03$) in the fourth block (the lower the better; one-tailed Wilcoxon signed rank test: Median_{diff} = 0.21, $Z = 2125.5$, $p \leq .0001$, $CI = [0.15, Inf]$).⁴ We opted for the rank measure instead of percentage correct, because rank is more sensitive in choice sets with three options, where choosing a second best option is meaningfully better than the third best option. Nonetheless, we additionally examined the choice performance with a percentage correct measure that is somewhat

⁴The small decrease in choice performance in the last block probably occurred due to fluctuations in difficulty, although there was no significant difference between mean Euclidean norm of criterion values between 3rd and 4th block. We did not strictly control difficulty when constructing the stimuli, and since all participants received the same stimuli, this unlucky random draw was not canceled out.

Figure 4. (a) Participants' performance improved over time in the choice task. Performance was measured as the mean rank of options chosen in all the trials of a block. Correct choices (i.e. those with the highest expected value) would yield mean rank of 1 while random choices would yield a mean rank of 2. Thus, participant performance in each block represents the mean of individual mean ranks across the 40 block trials (i.e. 1-40, 41-80, 81-120, 121-160), and error bars represent the standard errors of the grand means. The shaded region represents the range of participant performance in the choice task, from the 10th to 90th percentile. The point marked with number five next to the y-axis illustrates the performance in the first five trials. We also illustrated the performance of the delta-elimination (DE) and weighted-additive (WADD) models fitted to perfect choices, using cue weights inferred by the least-mean-squares (LMS) model fitted to perfectly accurate predictions (dotted lines). (b) Participants who performed well in the estimation task tended to perform well in the choice task as well. Estimation task performance is expressed as a root mean squared error between predictions and criterion values (RMSE), and choice task performance is the mean rank of chosen options in the last block, to capture the stable part of the performance. Point shape denotes the best fitting choice model – WADD, DE or BC model (showing both BC model variants pooled together). Mean prediction from best fitting baseline prediction model (Mean BPM) and baseline choice model (Mean BCM; 4th block) are illustrated as thick gray lines and (one) standard deviation around the means is denoted by lighter grey area

stricter. There were only minor differences between measures and for brevity we do not report them here.

We fitted two baseline choice models, a random choice model (BC_f) where we fixed the probability of choosing any option to 1/3, and a more flexible variant (BC_m) where the probabilities of choosing different options were estimated from choice frequencies (see Appendix A). Participants did better than the baseline choice models already in the first five trials, with a mean rank of 1.71 ($SE = 0.04$; one-tailed Wilcoxon test: Median = 1.60, $Z = 231$, $p \leq .0001$, $CI = [-Inf, 1.70]$). This result concurs with previous findings, which suggest that people have a strong prior for positive linear relationships between cues and the criterion value (Brehmer, 1974; Bussemeyer et al., 1997). The initial bias we identified matches the function we used to construct the stimuli (Equation 4); the cues with positive weights were stronger predictors than the cues with negative weights. As a result, the initial guesses of many participants that cues were positively correlated with the criterion value were on average

correct.

Individual variability in choice performance was substantial; some participants got close to perfect choices, while others were still close to random by the fourth block, as indicated by the 10th and 90th percentile of performance (shaded region in Figure 4a). How was this variability related to variability in estimation task performance? Most of the people who performed well in the choice task also did well in the estimation task. This becomes immediately apparent by looking at individual performance on the choice task and performance on the estimation task jointly (Figure 4b) and is also supported by a strong Spearman rank correlation (0.77) between participants' performance in the two tasks. Further, many of the participants that performed close to the baseline or worse in the estimation task cluster close to random choice performance. It is likely that they were less motivated or attentive and consequently achieved low performances in both tasks.

While the WADD strategy requires cue weights, lexico-

Figure 5. Evolution of cue weights in the choice task, as inferred by the best fitting least-mean-squares model (LMS), shows that for both weighted-additive (WADD) and delta-elimination (DE) users cue weights were converging toward the objective ones, albeit at a slower pace for DE users. Weights are median cue weights across participants for each trial, smoothed with a moving average of ten trials.

graphic heuristic strategies such as DE can get away by only knowing the cue order (Hogarth & Karelaia, 2007; Katsikopoulos et al., 2010). In Figure 4a we illustrated the best possible choice performances of DE and WADD – when fitted to correct choices, based on cue weights inferred by LMS models fitted to perfect predictions (i.e. criterion values) in the estimation task (see benchmark choice models in Appendix A). These provide a performance benchmark – the best possible that can be achieved by taking learning in this environment into account. Perfect DE, using only cue order, achieved choice performance of 1.29 ($SE = 0.003$) in the 4th block, with whole performance curve being closer to most of participants’ performances across blocks. An optimally tuned WADD strategy that fully uses cue weight knowledge achieved almost perfect choice accuracy, with mean rank of 1.05 ($SE = 0.002$) in the 4th block. Some participants achieved accuracies that were higher than what was possible to achieve with perfect DE (mean rank < 1.29, $N = 17$, Figure 4b), suggesting that they adopted the more demanding WADD strategy. For participants with performance around and worse than the DE perfect performance level (mean rank > 1.29) it was less clear what strategy they adopted – they could be using DE or the WADD strategy but learning more slowly. To identify the strategy with greater certainty and examine the nature of their knowledge we modelled their choices, taking into account the evolution of each participant’s knowledge.

We fitted WADD and DE model to each individual’s data from the choice task (see The formal models Section for de-

tails on choice models and Supplemental Material for fitting procedure).⁵ This analysis assumes that participants used a single strategy across the choice task, a common practical assumption in studies of strategy selection (for a different approach see Rieskamp & Otto, 2006). Importantly, both models relied on the cue weight information supplied by the version of the LMS model that fitted a participant best in the estimation task.⁶ The WADD model used the weights directly while the DE model extracted cue order from cue weights. The WADD model fitted participants’ choices better, with mean BIC score of 302 ($SE = 4.94$) and BIC weight of 0.50 ($SE = 0.05$), but the DE was close behind, with mean BIC score of 306 ($SE = 4.61$) and BIC weight of 0.38 ($SE = 0.05$), both doing much better than the baseline choice model with BIC score of 352 ($SE = 0$; BC_f variant, better of the two) and BIC weight of 0.11 ($SE = 0.03$) (Table S1 in Supplemental Material). In absolute terms WADD and DE were predicting participants’ choices equally well, with average probability of 0.46 ($SE = 0.01$) and 0.46 ($SE = 0.01$).

⁵We examined how well we can differentiate between the two strategies in the choice task by using the performance of ideal DE and WADD strategies (see Figure 4a). On average choices between the strategies differed for 0.20 ($SE = 0.05$) of trials in the choice task, 32 ($SE = 0.41$) trials on average.

⁶Instead of removing participants who were best fitted with baseline prediction models rather than with one of the LMS models, we used their best fitting LMS model. We opted for this solution as these participants might have been heuristic users that potentially learned cue orders without learning weights.

For comparison, versions of WADD and DE that were fitted to perfect choices, relying on LMS based cue weights when fitted to perfect predictions in the estimation task (performance is illustrated in Figure 4a), achieved a mean BIC of 137 ($SE = 4.64$) and 221 ($SE = 2.35$), and predicted correct choices with average probability of 0.80 ($SE = 0.009$) and 0.66 ($SE = 0.005$), respectively. We also performed cross-validation, an alternative model selection procedure, as a robustness check and it showed similar model selection results (see *Supplemental Material*).

We categorized participants as users of a certain strategy by selecting the strategy with the lowest BIC value for that participant. The largest group of participants were categorized as WADD users (38), whereas a sizable group were categorized as DE users (29). A smaller group were classified as BC users (10, two baseline choice models pooled together), five of which were also best fitted with one of the baseline prediction models in the estimation task. Mean BIC weight of the model fit, when averaged only across participants best fitted with a model, tells us how confident we can be in categorization results. Participants were classified with a good degree of confidence, as indicated by the mean BIC weight of 0.93 ($SE = 0.01$) for WADD, 0.94 ($SE = 0.02$) for DE. Visualizing the classification results jointly with choice and estimation task performance revealed that the classification yielded reasonable results (Figure 4b). Analysis of raw choice performance suggested that participants that achieved choice accuracy higher than what was possible to achieve with perfect DE (mean rank < 1.29, Figure 4b) adopted the WADD strategy. Using model classification results we can confirm that indeed most of these decision makers were categorized as WADD users (14 out of 17). On average participants categorized as DE users achieved choice performance of 1.5 ($SE = 0.03$). There were also quite a few WADD users with relatively poor choice performance (6 out of 38). Most of them also performed poorly in the estimation task, suggesting that they possessed less accurate cue weight knowledge.

There were five participants that were classified as BC users and obviously had choice performance well below the random performance line. These cases are indicative of some limitations of the CLE paradigm. Since LMS models were fitted to the estimation task data, predictions in the choice data are not constrained. For most people LMS modeling worked fine, but in these few cases LMS produced unrealistically high or low weights in the first half of the choice phase, to which choice models dealt with by increasing error parameter almost to maximum, consequently leading to being best fitted with one of the BC models. Instead of trying out tailor-made fits for these participants we opted for leaving them as they are, to be as transparent as possible and illustrate the limits of our approach. There were also several participants who had performance close to random, but were

not best fitted with the BC models. For some of them the WADD and DE strategies adopted very high error parameter values and were only marginally better than the BC models. For others the WADD and DE were much better than the BC models – cue knowledge of these participants was somewhat skewed and they were producing choices consistent with their knowledge, even if these were wrong choices, that is, close to random. Despite these limitations the modeling results were sensible for most of the participants and produced valuable insights to otherwise impervious cognitive processes.

We next examined the differences in cue weight knowledge between participants classified as DE or WADD strategy users. Broad patterns were visible from average estimation task performance (Figure 4b). Most importantly, DE users on average had estimation task performance that was substantially below the baseline prediction model performance, showing that DE users knew cue weights, rather than just the cue order. We obtained more detailed picture by examining cue weights inferred by the LMS models and insight question reports, separately for each strategy. Most interesting were again the DE users. Their cue weight evolution showed a slower convergence toward objective weights than of that by the WADD users, and by the end of the choice task their level of knowledge was somewhat worse than that of WADD users (Figure 5). Based on the same LMS inferred cue weights we examined correctness of cue orders for participants in the two groups. Correct cue orders are particularly important for DE users to make accurate choices. For both strategy users the first rank cue was on average correct (0.97 ($SE = 0.03$) for DE and 1.00 ($SE = 0$) for WADD), while for all other cues DE users learned the correct order in smaller proportion (0.24 ($SE = 0.08$) for the second, 0.21 ($SE = 0.08$) for the third and 0.48 ($SE = 0.09$) for the fourth ranking cue) than WADD users (0.39 ($SE = 0.08$) for the second, 0.42 ($SE = 0.08$) for the third and 0.76 ($SE = 0.07$) for the fourth ranking cue). Overall, fewer DE users learned the correct cue order for all four cues (0.14, $SE = 0.06$), in comparison to WADD users (0.39, $SE = 0.08$; two-tailed proportion test, $\chi^2(1) = 4.15$, $p = 0.042$, $CI = [-0.49, -0.03]$). Insight questions of DE users on average pointed in the same direction: reports were almost the same as those of the WADD users, with the exception of slower learning of the weight for the second cue and wrong direction for the fourth cue (Figure S3a in *Supplemental Material*). Insight questions and LMS weights of BC users showed that they indeed had much poorer knowledge about the cue weights (Figure S3b and S3c in *Supplemental Material*).

In summary, we found both the WADD strategy and the DE strategy users among experiment participants; importantly, the DE group also learned cue weights rather than exclusively cue orders. However, we found differences in cue weight learning between strategy users. The cue weight

Figure 6. (a) The choice accuracy of the delta-elimination heuristic as a function of the discrimination threshold (Δ) value, obtained in simulations on our choice task using objective cue weights. The choice accuracy of the strategy peaks for threshold value of 19 (vertical dashed line) and deteriorates rapidly afterwards, reaching random choice performance for values above 70 (horizontal dashed line). (b) The total number of elements considered, i.e. individual cue values inspected until decision was made, and the estimated threshold parameters for all participants. Shape of the point marks the model fitting the participant's choices best: weighted-additive (WADD), delta-elimination (DE) or one of the baseline choice (BC) models. All the participants who were best fitted by DE had threshold values between 0 and 29, in the region or close to the region where the effort/accuracy trade-off applies. Overall, they considered much less information than people relying on the WADD strategy. Dashed lines mark the optimal Delta value in terms of the achieved option rank and the associated number of elements inspected.

evolution and insight questions of DE users showed a slower cue weight learning. This result suggests an interaction between the two processes, even if the choice strategy does not affect the type of knowledge that people acquire.

Delta-elimination and the effort/accuracy trade-off

Our second goal was developing a new lexicographic heuristic – the DE strategy – and testing it empirically. This new heuristic is well equipped to tackle challenging choice environments with more than two options and continuous-valued cues, in which heuristic strategies have rarely been put into test before. The DE strategy described well choices of a sizable group of participants. Here we discuss the prescriptive and descriptive performance of the strategy.

To examine the effects of the discrimination threshold (Δ) on choice accuracy and search through the cues we simulated the DE strategy on our choice task. We used objective cue weight values and varied the discrimination threshold from 0 to 100. Choice accuracy and the number of informational elements considered by the DE strategy vary with the threshold value in a different way (Figure 6a and Figure S1a in *Supplemental Material*), producing an interesting effort/accuracy trade-off. The DE subsumes the TTB and

single-attribute heuristics as a special case where $\Delta = 0$. At this point effort is at the lowest point, with the least amount of cues inspected. However, the accuracy is compromised and can be improved with somewhat larger threshold levels. The choice accuracy peaks at threshold value of 19, looking up only 4.8 cue values on average. For larger threshold values the accuracy of the strategy deteriorates, as it starts to rely increasingly on less important pieces of information to make a choice. The region between $\Delta = 0$ and $\Delta = 19$ is the Pareto efficient frontier in terms of the effort/accuracy trade-off – it is impossible to further increase the accuracy of the strategy without also increasing the implied effort. This is the region that we should expect to see populated by the participants relying on the DE rule. Overall, the strategy can capture different points on the effort/accuracy continuum, from the immediate choice implied by the TTB heuristic for continuous cues, to a more balanced trade-off implied by somewhat larger threshold levels.

The DE strategy fitted best 29 participants (from the remaining participants 38 were best modeled with WADD and 10 with the one of the baseline choice models). The mean threshold value for DE users was $\Delta = 15.3$ ($SE = 1.85$), close to the theoretically optimal value of 19 for this environment.

The maximum observed value was 28.7 and values were generally distributed close to the optimal value (participants marked with a circle, Figure 6b, see also Figure S1b in *Supplemental Material* where distribution is easier to discern). Slightly over half of the participants (15) had Delta values within the Pareto efficient effort/accuracy frontier ($\Delta < 19$) and the remaining slightly above it (from 19 to 29), which also corresponds to a relatively flat region on the performance function. Although none of the uncovered participant threshold values was exactly $\Delta = 0$, a cluster of eight participants had Delta values that were very close to it, with $\Delta < 5$. People belonging to that group were very near to the absolute minimum of 3 pieces of information that would be consumed by TTB in this environment—the Delta was sufficiently small that most cases were decided already on the first cue. Looking at participants best fitted with WADD, they were fitted with higher Delta values (mean of 35.8, $SE = 3.95$), with one subset of participants having much higher values than participants best fitted with DE. All participants best fitted with one of the BC models, reassuringly had Delta values larger than 70, same as in our simulations (Figure 6a; see also Figure S1b in *Supplemental Material*). How much information did people best described by the DE consume as compared to the more information hungry WADD strategy? Our results suggest that they inspected on average 4.6 ($SE = 0.20$) pieces of information and at most 6.1 pieces of information (Figure 6b), whereas people using the WADD strategy would consider all 12 of them.

Discussion

For most choices people make in their every-day life they can choose among several options and have access to numerous informative cues. People can integrate all the information at hand or they can use lexicographic heuristic strategies that omit most of it. Yet these two popular types of choice strategies require different forms of knowledge about the informational value of different cues. People using integrative strategies need to learn exact cue weights whereas people relying on lexicographic heuristics merely need to learn cue order. Here, we introduced the CLE paradigm that allowed us to investigate the interplay between strategy selection and learning about the structure of the environment. We used participants' responses in an estimation task deployed after they made choices and their periodical self-reports on cue importance while they were making choices. We found that almost all participants exhibited knowledge that goes beyond mere cue order – they learned cue weight values. We then postulated least-mean-squares learning as a mechanism of how cue weights are learned and uncovered how the weights evolved during the choice task while the participants were receiving feedback on chosen options. Feeding the obtained cue weights into the integrative WADD strategy and newly developed DE lexicographic strategy, we found that the WADD

strategy described the largest group of participants. The DE heuristic best described a sizable group of participants, and more importantly, these participants learned cue weights as well. Nonetheless, the group that used the DE strategy performed worse in the estimation task and the insight questions showed weaker discrimination between the cues.

The adaptive appeal of learning cue weights

Why would people using lexicographic heuristics such as the DE strategy learn cue weights when they only need cue order? Firstly, it has been argued and empirically demonstrated that people select their strategies adaptively, and they are more likely to choose a strategy in environments where this strategy excels (e.g. Lieder & Griffiths, 2017; Rieskamp & Otto, 2006). For example, lexicographic strategies are known to perform well in settings where there is one very informational cue and there are high inter-correlations among the cues (e.g. Katsikopoulos, 2011; Şimşek & Buckmann, 2015). In such settings people tend to use lexicographic heuristics more often. However, the structure of the environment, and consequently the best-performing strategy, might be unknown beforehand (Stojić et al., 2016). By learning cue weights people can switch between the different types of strategies as deemed appropriate, whereas by learning only cue order they would be locked into choosing among lexicographic strategies. Hogarth and Karelaia (2007), for instance, showed theoretically that it would be beneficial to adopt the TTB when people have little experience and their knowledge about the structure of the environment is poor, and switch to the WADD later on as they become more experienced. If TTB users learned only cue order such adaptive switch could not be executed.

Second, learning the cue weights has been shown to be an efficient way to learn cue orders. Todd and Dieckmann (2005) designed a simulation in which they compared how fast a series of learning mechanisms learn cue order from choice feedback. The competing approaches included a model that directly learned cue validities, models that tallied how many times a cue led to correct decisions and reordered cues accordingly, simpler ordering rules that after each decision moved cues that led to correct decisions earlier in the cue order, and an associative LMS model that learned orders from weights. They found that the LMS model was the fastest cue order learner, despite learning orders indirectly. However, they argued that the LMS algorithm implies a high cognitive cost and memory load. One caveat of these cue order learning models is that they lack the mechanism for learning cue direction (i.e. whether positive or negative difference in cue values leads to a correct choice), instead assuming this information is known. By contrast, the LMS model simultaneously learns cue directions and takes them into account in computing cue order. More importantly, Todd and Dieckmann only investigated how these models learn when they

are given feedback about whether choices were correct. This is somewhat limiting as people often do not experience options they do not choose, and therefore they can rarely learn whether their choices were correct. It is possible to extend the Todd and Dieckmann analyses in scenarios like the one we studied, where the models are provided with feedback about the criterion value of chosen option. The LMS model can operate well with either type of feedback—cue weights can be learned either from correctness or criterion value feedback. The models directly learning cue-orders studied by Todd and Dieckmann could leverage stored cue and criterion values of chosen options to simulate choices in memory. Thus, the fundamental learning events would be simulated choices, as opposed to finding out whether actual choices were correct. For instance, 10 choices imply an equal number of stored criterion values. These could be used to simulate 45 choices among 2 options and 120 choices among 3 options, a substantial increase in learning sample when compared to learning from choice feedback. Such an approach would also lead to orders of magnitude more computation for all models directly learning cue-orders, whereas the LMS approach would require exactly the same amount of computation. Thus, the LMS approach becomes the most cost-efficient approach in the setting with criterion value feedback, while there is no reason to expect it will be learning slower than models learning directly cue-orders.

Finally, in real life people do not tackle only choice problems and they might use the knowledge gained in other ways. One way the cue weight knowledge can be used is for predicting the quality of options, as in our estimation task. This is available in choice situations where feedback comes in terms of criterion values. This is always the case when people have only their experience to teach them, as opposed to having information about outcomes of actions that were not taken or an authority informing them whether their choice was correct or not. Ability to predict qualities of options has its own value; for example, people often have to communicate their predictions, say quality of a restaurant or price of a stock a year from now. Hence, by learning cue weights people can seamlessly transition between choice and estimation tasks. Another way is for tackling choices in other domains. Richer (generative) models of the environment can be used to form hierarchical representations that facilitate learning in related other problems. For example, Lake, Salakhutdinov, and Tenenbaum (2015) showed that this mechanism might underlie human ability to successfully generalize their experience to new tasks based only on few observations, and illustrated how the same hierarchical representation can be used in multiple ways – action, imagination and explanation (for more general discussion, see also Lake, Ullman, Tenenbaum, & Gershman, 2017). The benefits of generalization might be large enough to outweigh any potential costs (i.e. cognitive or memory costs). Thus, learning the structure of the

choice environment in a detailed manner might be people's default, either hard-wired by evolution or learned through experience.

Interactions between strategy selection and cue weight learning

Even though a choice of a particular strategy did not influence whether participants learned cue weights or cue order, we found differences in cue weight knowledge between strategy users, suggesting an interaction between strategy selection and cue weight learning. On average WADD users acquired better knowledge of cue weights than DE users, as evidenced by their estimation task performance, insight questions reports and LMS inferred cue weights.

One plausible explanation for this difference among strategies is that cue weight knowledge influences the adoption of a choice strategy (Stojic, Olsson, & Analytis, 2016). In a theoretical analysis of strategy performance in choice tasks, Hogarth and Karelaia (2007) showed that people should use a single variable heuristic while their knowledge of cue weights is imperfect and switch to the integrative strategy once they have learned the cue weights sufficiently well. Our results are consistent with this perspective. The DE users were learning at a slower pace and relied on that strategy because their knowledge was still imprecise. Given more time and observations their cue weight knowledge would improve and the linear integrative strategy would become more appealing in terms of maximizing utility. People potentially develop this strategy switching rule through their experiences with many environments. Instead of using a fixed switching rule, this could be a continuous optimization process. Along with learning of the cue weights there could be a simultaneous learning process of how effective the strategy is (Lieder & Griffiths, 2017; Rieskamp & Otto, 2006; Stojic et al., 2016). Early in the learning process one of the strategies might yield better cost-accuracy trade-off, while another strategy could take over later when the knowledge improves. When coupled with individual variability in learning rates, commonly found in cue weight learning tasks, the causal link between strategy performance and cue weight knowledge could give rise to systematic differences between strategies observed in our data. What's more, this mechanism could explain variability in strategy adoption in empirical choice studies that allow learning, providing a potential explanation to an otherwise elusive phenomenon (Bröder, 2012; Bröder & Schiffer, 2006b; Newell, 2005).

However, it is also possible that the adopted strategy influenced how well the participants learned the cue weights in our task. For instance, using the DE strategy might result in paying attention to a subset of cues in each trial. This influence could be potentially captured with attention weighted updates of cue weights in the learning process of the LMS model (Leong, Radulescu, Daniel, DeWoskin, & Niv, 2017;

Mackintosh, 1975). The resulting bias in cue weight learning could explain why DE users were learning slower and consequently showed poorer performance in the estimation task. This mechanism would not be mutually exclusive with variability in learning rates influencing strategy selection, both mechanisms could play a role. With our data we cannot exclude either possibility. Evaluating an LMS model where cue weight updates are made dependent on cues examined by the DE strategy in each trial, would yield insights into these interactions. Collecting data on attention to cues, through eye- or mouse-tracking, would be indispensable to appropriately constrain and evaluate such a model. Studying the impact of the selected strategy on cue weight learning is an important direction for future research.

A recent study by Parpart, Schulz, Speekenbrink, and Love (2017) also focused on interactions between strategy selection and learning the structure of the environment. They formulated active learning versions of WADD and TTB, assuming that people using the former will learn cue weights and the latter cue order. Then they investigated how people choose between two options, while they also strive to learn cue weights or order. Notably, in contrast to our approach, they tied strategies with the type of knowledge they require. They found that an active learning version of WADD described most of the participants. Our results show that even people that use lexicographic strategies such as TTB may preferentially learn cue weights, rather than order. Participants categorized as WADD users by Parpart et al. (2017) may be people using lexicographic strategies who are learning the cue order by learning the cue weights first. In sum, we believe that our cue weight learning perspective and the active learning perspective of Parpart et al. (2017) provide converging evidence that most people learn cue weights rather than directly learning order.

From accidental errors to adaptive exploration

To fit the WADD and DE models to human data we assumed a trembling hand error. It is also well-known that parameters accounting for human errors improve the descriptive value of choice models (e.g. Loomes et al., 2002; Rieskamp, 2008). A trembling hand type of error introduces a free parameter ϵ , which assumes that participants choose the best option with probability $1 - \epsilon$, otherwise they mistakenly choose either of the three options with probability ϵ . This type of error has often been used in context such as ours (e.g. Bergert & Nosofsky, 2007; Bröder & Schiffer, 2003a, 2003b; Nosofsky & Bergert, 2007; Rieskamp, 2006; Rieskamp & Otto, 2006). In environments where there is no learning and the cue weights are already known, such as those used in most of the empirical comparisons of integrative and lexicographic strategies, a trembling hand parameter is undoubtedly sub-optimal and can be safely interpreted as lapse in judgment. By contrast, in environments where

people have to learn the structure of the environment, such randomness in the choice process may reflect sensible and adaptive behavior. People may occasionally choose at random to *explore* the structure of the environment to gain information, that they could later on *exploit* to make better decisions. This is the so-called exploration-exploitation trade-off, a signature of reinforcement learning problems (Niv, 2009; Sutton & Barto, 2018). In fact, the trembling hand parameter corresponds to the exploration parameter in the classical ϵ -greedy model in the reinforcement learning literature (Sutton & Barto, 2018). From the reinforcement learning point of view, the somewhat large ϵ parameter that we found in our modeling results might indicate a large exploration tendency, rather than an error in executing a choice strategy. Furthermore, choice tasks where options are characterized with multiple cues and cue-criterion relations have to be learned from experience is closely related to contextual multi-armed bandit tasks, a class of reinforcement learning problems that has been recently under focus (e.g. Schulz, Konstantinidis, & Speekenbrink, 2018; Stojic, Analytis, & Speekenbrink, 2015; Stojic, Schulz, et al., 2020; Wu, Schulz, Speekenbrink, Nelson, & Meder, 2018). The link between multi-cue choice and reinforcement learning has been seldom made and there is potential for transferring insights from one field into the other in the future studies. For example, reinforcement learning models usually do not assume that individuals choose among multiple choice strategies, and heuristic strategies are rarely considered among the strategies that people may use.

Lexicographic heuristics for continuous multi-alternative environments

Up to now, the descriptive appeal of lexicographic heuristic strategies has been primarily assessed in choice environments with two options and binary cues describing them (e.g. Gigerenzer & Goldstein, 1996; Newell & Shanks, 2003). Although this setting has taken us a long way in the descriptive study of choice mechanisms, it is restrictive in the sense that it corresponds only to a subset of environments people experience in every-day life. Characteristically, when people have access to continuous cues, the TTB heuristic essentially reduces to a single-cue strategy and its prescriptive appeal is somewhat reduced (Hogarth & Karelaia, 2005b). Assuming that people use only one cue in contexts with many continuous cues and many options is a strong, but potentially unrealistic assumption. Nonetheless, the informational and computational load of integrative strategies also increases in these contexts, making them much less appealing as well. Can lexicographic heuristic strategies provide good descriptive models in these settings? What type of strategies should we consider?

The Δ -inference strategy effectively generalizes TTB to continuous environments with the help of an additional discrimination parameter (Luan et al., 2014). Luan and col-

leagues have showed that the strategy can achieve competitive performance in continuous environments. We equipped Δ -inference with an elimination mechanism, as implemented in the elimination-by-aspects strategy (Tversky, 1972). The resulting delta-elimination strategy can cope well with environments with more than two options and continuous cues. Our results show that DE described a sizable group of participants, even in an environment where an integrative strategy was bound to perform better. The strategy substantially reduced the pieces of information that people need to consider and the total number of cognitive operations performed. This is one of the few experiments where lexicographic heuristics have been tested and shown to capture people's behavior in scenarios with continuous cues where people choose among more than two options. Our work also corroborates and expands recent work by Luan, Reb, and Gigerenzer (2019) showing that Δ -inference strategy describes well a large group of participants in binary choice problems. Based on previous results showing that people can adaptively switch decision strategies depending on the structure of the environment (e.g. Rieskamp & Otto, 2006), we expect that in settings that are favorable to lexicographic strategies, such as environments with a skewed (J-shaped) distribution of cue importance or validities, DE would describe an even larger group of participants. Given our finding that DE users learned cue weights despite not needing them, we expect the same to happen in environments more favorable to lexicographic strategies.

Limits of the cue weight learning approach

We offered an account of how cue weight learning interacts with choice strategies in one linear environment. This raises several questions about the generalizability of our results and conclusions. First, the objective weights in our environment were fixed at 4, -3, 2, and -1. That is, there is a large dispersion in cue weights, and all cues make independent contributions. In such environments the WADD strategy is expected to thrive and therefore learning cue weights might be particularly beneficial. In settings where the cue dispersion is low and where there are high inter-correlations between cues, by contrast, lexicographic heuristic strategies are known to perform as well, and therefore, knowing cue weights would not confer a clear advantage. However, the DE users in our current experiment did not have a use of knowing cue weights and still learned them. In low cue dispersion or high cue inter-correlation environments we would expect to see more DE users, but we would expect them to still learn the cue weights rather than order. Earlier in the discussion we provided several arguments for people to learn cue weights even in settings where it is less beneficial to do so from a choice strategy point of view. Nonetheless, future experiments in settings with low cue dispersion and high inter-correlations among cues would provide a good robust-

ness check of our results.

Further, our cue weight modeling approach is unlikely to work in more complex environments where cues are nonlinearly related to criterion. The LMS model provides a good approximation of what people know in linear environments, where there is a perfect match between the objective structure and the model (Kelley & Busemeyer, 2008; Speekenbrink & Shanks, 2010). The LMS model belongs to a family of rule-based learning strategies, well known in category and function learning literature (e.g. Busemeyer et al., 1997). In nonlinear environments, by contrast, cue weights are not necessarily a good summary representation of the environment and people are more likely to learn in qualitatively different manner, better described with exemplar-based learning (e.g. Hoffmann, von Helversen, & Rieskamp, 2016). Could people still rely on heuristic strategies? While knowledge in exemplar-based learning is contained in memories of encountered options there is an attention weight associated with each cue, presumably also learned through trial-and-error (Kruschke, 1992). These attention weights do not necessarily represent objective cue weights that determine relations between cues and criterion, but they do reflect their contribution and could be used as order information in lexicographic heuristics. Finally, even in these settings some people may learn using rule-based strategies (e.g. Hoffmann et al., 2016; Olsson, Enkvist, & Juslin, 2006), and consequently, they may rely on heuristics for choosing among different options. Our understanding of how heuristics work in nonlinear environments, as well as how choices interact with learning of the structure of the environment, is still developing and constitutes an exciting venue for future research.

In our choice task participants got criterion values as feedback. Several studies found that with binary feedback people achieve high choice performance levels, implying that they successfully learned the cue validities (e.g. Lee & Cummins, 2004; Nosofsky & Bergert, 2007). Learning the structure of the environment with binary feedback proceeds differently. While cue weights can be learned (Gluck & Bower, 1988), they do not necessarily correspond to the objective cue weights that determine the criterion values. This means that in a choice task with binary feedback people should have a hard time providing good estimates in the estimation task. However, Pachur and Olsson (2012) found that participants show a good performance in an estimation task that followed a choice task with binary feedback. This could possibly be a consequence of the comparison of two objects fostering learning of how changes in cue values are associated with changes in the criterion, rather than learning the relations between cue values and the criterion (Klayman, 1988a, 1988b). An interesting direction for future studies would be to examine in greater detail what underlies this performance by applying the same cue weight learning approach as done here.

There is scope for further development of the CLE

paradigm we have used here. We noted that the LMS model fits of some participants produced unrealistic cue weights in the first half of the choice task, causing issues with fitting choice models. This is possible because the predictions the LMS makes for the choice task are not constrained. A potential remedy is to intersperse some estimation task trials in the choice task (e.g. Gershman & Niv, 2015). Lagnado et al. (2006) found that self-reports similar to our insight questions do not affect learning and strategy use, suggesting that such estimation task trials might not interfere either. Here we took a conservative approach, leaving the systematic examination of whether estimation task trials change how people go about their choices for future studies.

In introducing insight questions during the choice task we relied on results by Lagnado et al. (2006). They examined learning and insight in a category learning task, which used four cues with binary cue values and correct/incorrect feedback. This task is structurally similar to the pairwise comparison tasks that have been extensively used to test integrative and lexicographic strategies (Gigerenzer & Goldstein, 1996). In one of their experiments, exactly like in our study, they asked insight questions after each learning block. In another experiment they asked insight questions only at the end. They considered similar integrative and heuristic strategies with us and did not find any differences between experiments in people's strategy adoption, or in their learning performance overall. Given that the task in Lagnado et al. (2006) was not exactly the same as ours, there is a small possibility that insight questions played a role in participants adopting cue weight learning independent of strategy use. Future studies using the CLE paradigm could also include control groups without insight questions to replicate the results from Lagnado et al. (2006) in our setting.

Conclusion

This work contributes to our understanding of the strategy selection, in particular the interaction between choice strategy and learning of the structure of the environment. Our results showed that participants learned cue weights regardless of the choice strategy employed. Adopting a lexicographic strategy does not affect the type of knowledge about cues that participants acquired. The fact that lexicographic strategies require less knowledge to operate – the cue order – does not imply that people learn only the absolutely necessary knowledge. However, there was an interaction between adopted strategy and the cue weight learning process. Either individual variation in cue weight learning lead to differences in strategy selection, or strategy might have influenced how well participants learned the cue weights, or both; our design does not allow us to identify the direction. This work also contributes to development of heuristic strategies, both theoretically and empirically. The Delta-elimination strategy generalizes previous lexicographic strategies and can deal

with choice environments with more than two options and continuous cues. Our findings show that it is an attractive strategy that substantial number of people adopted, thus providing empirical support for this new lexicographic heuristic. Our choice-learning-estimation paradigm is a small methodological contribution as well. Combining choice and estimation task allowed us to examine the extent and type of knowledge about the environment people acquired while they made choices. Our hope is that more studies will take the integrative approach and study strategy selection jointly with learning, as outlined here.

References

- Analytis, P. P., Barkoczi, D., & Herzog, S. M. (2018). Social learning strategies for matters of taste. *Nature Human Behaviour*, 2, 415.
- Auguie, B. (2012). *gridExtra: functions in Grid graphics*. Retrieved from <http://cran.r-project.org/package=gridExtra>
- Bergert, F. B., & Nosofsky, R. M. (2007). A response-time approach to comparing generalized rational and take-the-best models of decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33, 107–129. doi: 10.1037/0278-7393.33.1.107
- Bettman, J. R., Johnson, E. J., & Payne, J. W. (1990). A componential analysis of cognitive effort in choice. *Organizational Behavior and Human Decision Processes*, 45, 111–139.
- Brehmer, B. (1974). Hypotheses about relations between scaled variables in the learning of probabilistic inference tasks. *Organizational Behavior and Human Performance*, 11, 1–27. doi: 10.1016/0030-5073(74)90002-6
- Brehmer, B. (1994). The psychology of linear judgement models. *Acta Psychologica*, 87, 137–154.
- Bröder, A. (2000). Assessing the empirical validity of the "take-the-best" heuristic as a model of human probabilistic inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 1332–1346.
- Bröder, A. (2012). The quest for take the best - Insights and outlooks from experimental research. In P. M. Todd, G. Gigerenzer, & the ABC Research Group (Eds.), *Ecological rationality: Intelligence in the world* (pp. 216–240). New York, NY, US: Oxford University Press.
- Bröder, A., & Schiffer, S. (2003a). Bayesian strategy assessment in multi-attribute decision making. *Journal of Behavioral Decision Making*, 16, 193–213. doi: 10.1002/bdm.442
- Bröder, A., & Schiffer, S. (2003b). Take The Best versus simultaneous feature matching: Probabilistic inferences from memory and effects of representation format. *Journal of Experimental Psychology: General*, 132, 277–293. doi: 10.1037/0096-3445.132.2.277
- Bröder, A., & Schiffer, S. (2006a). Adaptive flexibility and maladaptive routines in selecting fast and frugal decision strategies. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32, 904–918. doi: 10.1037/0278-7393.32.4.904
- Bröder, A., & Schiffer, S. (2006b). Stimulus format and working memory in fast and frugal strategy selection. *Journal of*

- Behavioral Decision Making*, 19, 361–380. doi: 10.1002/bdm.533
- Busemeyer, J. R., Byun, E., Delosh, E. L., & McDaniel, M. A. (1997). Learning functional relations based on experience with input-output pairs by humans and artificial neural networks. In K. Lamberts & D. R. Shanks (Eds.), *Knowledge, concepts and categories. studies in cognition*. (pp. 408–437). Cambridge, MA, US: MIT Press.
- DeLosh, E. L., Busemeyer, J. R., & McDaniel, M. A. (1997). Extrapolation: the sine qua non for abstraction in function learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 968–986.
- Dieckmann, A., & Todd, P. M. (2004). Simple ways to construct search orders. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 26).
- Dougherty, M. R. P., Franco-Watkins, A. M., & Thomas, R. (2008). Psychological plausibility of the theory of probabilistic mental models and the fast and frugal heuristics. *Psychological Review*, 115, 199–213. doi: 10.1037/0033-295X.115.1.199
- Garcia-Retamero, R., Takezawa, M., & Gigerenzer, G. (2009). Does imitation benefit cue order learning? *Experimental Psychology*, 56, 307–320.
- Gershman, S. J., & Niv, Y. (2015). Novelty and Inductive Generalization in Human Reinforcement Learning. *Topics in Cognitive Science*, 7, 391–415. doi: 10.1111/tops.12138
- Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science*, 1, 107–143.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103, 650–669.
- Gigerenzer, G., Hoffrage, U., & Goldstein, D. G. (2008). Fast and frugal heuristics are plausible models of cognition: reply to Dougherty, Franco-Watkins, and Thomas (2008). *Psychological Review*, 115, 230–239. doi: 10.1037/0033-295X.115.1.230
- Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: a Brunswikian theory of confidence. *Psychological Review*, 98, 506–28.
- Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, 117, 227–247.
- Hoffmann, J. A., von Helversen, B., & Rieskamp, J. (2016). Similar task features shape judgment and categorization processes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42, 1193–1217. doi: 10.1037/xlm0000241
- Hogarth, R. M., & Karelaia, N. (2005a). Ignoring information in binary choice with continuous variables: When is less “more”? *Journal of Mathematical Psychology*, 49, 115–124.
- Hogarth, R. M., & Karelaia, N. (2005b). Simple Models for Multiattribute Choice with Many Alternatives: When It Does and Does Not Pay to Face Trade-offs with Binary Attributes. *Management Science*, 51, 1860–1872. doi: 10.1287/mnsc.1050.0448
- Hogarth, R. M., & Karelaia, N. (2007). Heuristic and linear models of judgment: matching rules and environments. *Psychological Review*, 114, 733–758. doi: 10.1037/0033-295X.114.3.733
- Hutchinson, J. M., & Gigerenzer, G. (2005). Simple heuristics and rules of thumb: Where psychologists and behavioural biologists might meet. *Behavioural Processes*, 69, 97–124.
- Juslin, P., & Persson, M. (2002). PROBABILITIES from EXemplars (PROBEX): A lazy algorithm for probabilistic inference from generic knowledge. *Cognitive Science*, 26, 563–607.
- Katsikopoulos, K. V. (2011). Psychological heuristics for making inferences: Definition, performance, and the emerging theory and practice. *Decision Analysis*, 8, 10–29.
- Katsikopoulos, K. V., Schooler, L. J., & Hertwig, R. (2010). The robust beauty of ordinary information. *Psychological Review*, 117, 1259–1266. doi: 10.1037/a0020418
- Kelley, H., & Busemeyer, J. R. (2008). A comparison of models for learning how to dynamically integrate multiple cues in order to forecast continuous criteria. *Journal of Mathematical Psychology*, 52, 218–240. doi: 10.1016/j.jmp.2008.01.009
- Klayman, J. (1988a). Cue discovery in probabilistic environments: Uncertainty and experimentation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 317.
- Klayman, J. (1988b). On the how and why (not) of learning from outcomes. In *Advances in psychology* (Vol. 54, pp. 115–162). Elsevier.
- Kohli, R., & Jedidi, K. (2007). Representation and inference of lexicographic preference models and their variants. *Marketing Science*, 26, 380–399.
- Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, 99, 22–44. doi: 10.1037/0033-295X.99.1.22
- Lagnado, D. A., Newell, B. R., Kahan, S., & Shanks, D. R. (2006). Insight and strategy in multiple-cue learning. *Journal of Experimental Psychology: General*, 135, 162–183. doi: 10.1037/0096-3445.135.2.162
- Lake, B. M., Salakhutdinov, R., & Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, 350, 1332–1338.
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, e253. doi: 10.1017/S0140525X16001837
- Lee, M. D., & Cummins, T. D. R. (2004). Evidence accumulation in decision making: unifying the “take the best” and the “rational” models. *Psychonomic Bulletin & Review*, 11, 343–352.
- Leong, Y. C., Radulescu, A., Daniel, R., DeWoskin, V., & Niv, Y. (2017). Dynamic interaction between reinforcement learning and attention in multidimensional environments. *Neuron*, 93, 451–463.
- Lieder, F., & Griffiths, T. L. (2017). Strategy selection as rational metareasoning. *Psychological Review*, 124, 762.
- Loomes, G., Moffatt, P. G., & Sugden, R. (2002). A microeconomic test of alternative stochastic theories of risky choice. *Journal of Risk and Uncertainty*, 24, 103–130.
- Luan, S., Reb, J., & Gigerenzer, G. (2019). Ecological rationality: Fast-and-frugal heuristics for managerial decision making under uncertainty. *Academy of Management Journal*(ja).
- Luan, S., Schooler, L. J., & Gigerenzer, G. (2014). From perception to preference and on to inference: An approach-avoidance analysis of thresholds. *Psychological Review*, 121, 501.

- Mackintosh, N. J. (1975). A theory of attention: variations in the associability of stimuli with reinforcement. *Psychological Review*, 82(4), 276.
- Martignon, L., & Hoffrage, U. (2002). Fast, frugal, and fit: Simple heuristics for paired comparison. *Theory and Decision*, 52, 29–71.
- Mata, R., Schooler, L. J., & Rieskamp, J. (2007). The aging decision maker: cognitive aging and the adaptive selection of decision strategies. *Psychology and Aging*, 22, 796–810. doi: 10.1037/0882-7974.22.4.796
- Newell, B. R. (2005). Re-visions of rationality? *Trends in Cognitive Sciences*, 9, 11–15. doi: 10.1016/j.tics.2004.11.005
- Newell, B. R., Rakow, T., Weston, N. J., & Shanks, D. R. (2004). Search strategies in decision making: The success of "success". *Journal of Behavioral Decision Making*, 17, 117–137.
- Newell, B. R., & Shanks, D. R. (2003). Take the best or look at the rest? factors influencing "one-reason" decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 53–65.
- Newell, B. R., Weston, N. J., & Shanks, D. R. (2003). Empirical tests of a fast-and-frugal heuristic: Not everyone "takes-the-best". *Organizational Behavior and Human Decision Processes*, 91, 82–96.
- Niv, Y. (2009). Reinforcement learning in the brain. *Journal of Mathematical Psychology*, 53, 139–154. doi: 10.1016/j.jmp.2008.12.005
- Nosofsky, R. M., & Bergert, F. B. (2007). Limitations of Exemplar Models of Multi-Attribute Probabilistic Inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33, 999–1019. doi: 10.1037/0278-7393.33.6.999
- Olsson, A.-C., Enkvist, T., & Juslin, P. (2006). Go with the flow: How to master a nonlinear multiple-cue judgment task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32, 1371.
- Pachur, T., & Olsson, H. (2012). Type of learning task impacts performance and strategy selection in decision making. *Cognitive Psychology*, 65, 207–240. doi: 10.1016/j.cogpsych.2012.03.003
- Parpart, P., Schulz, E., Speekenbrink, M., & Love, B. C. (2017). Active learning reveals underlying decision strategies. *bioRxiv*, 239558.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 534–552.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. New York, NY, US: Cambridge University Press.
- Peirce, J. W. (2007). PsychoPy - Psychophysics software in Python. *Journal of Neuroscience Methods*, 162, 8–13. doi: 10.1016/j.jneumeth.2006.11.017
- R Core Team. (2017). R: A Language and Environment for Statistical Computing [Computer software manual]. Vienna, Austria. Retrieved from <https://www.r-project.org/>
- Rakow, T., Hinveest, N., Jackson, E., & Palmer, M. (2004). Simple heuristics from the adaptive toolbox: Can we perform the requisite learning? *Thinking & Reasoning*, 10, 1–29.
- Rakow, T., Newell, B. R., Fayers, K., & Hersby, M. (2005). Evaluating three criteria for establishing cue-search hierarchies in inferential judgment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31, 1088–1104.
- RevolutionAnalytics. (2014). *doMC: Foreach parallel adaptor for the multicore package*. Retrieved from <http://cran.r-project.org/package=doMC>
- Rieskamp, J. (2006). Perspectives of probabilistic inferences: Reinforcement learning and an adaptive network compared. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32, 1355–1370. doi: 10.1037/0278-7393.32.6.1355
- Rieskamp, J. (2008). The probabilistic nature of preferential choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34, 1446.
- Rieskamp, J., & Hoffrage, U. (2008). Inferences under time pressure: how opportunity costs affect strategy selection. *Acta Psychologica*, 127, 258–76. doi: 10.1016/j.actpsy.2007.05.004
- Rieskamp, J., & Otto, P. E. (2006). SSL: a theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, 135, 207–236. doi: 10.1037/0096-3445.135.2.207
- Rogers, A. R. (1988). Does biology constrain culture? *American Anthropologist*, 90, 819–831.
- Rosenblatt, F. (1958). The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological Review*, 65, 386.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing, exploitation in the microstructure of cognition - vol. 1: Foundations* (pp. 154–172). Cambridge, MA, US: Oxford University Press.
- Schulz, E., Konstantinidis, E., & Speekenbrink, M. (2018). Putting bandits into context: How function learning supports decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 44, 927–943. doi: 10.1037/xlm0000463
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6, 461–464.
- Shanks, D. R., & Lagnado, D. (2000). Sub-optimal reasons for rejecting optimality. *Behavioral and Brain Sciences*, 23, 761–762.
- Sharpsteen, C., & Bracken, C. (2015). *tikzDevice: R Graphics Output in LaTeX Format*. Retrieved from <http://cran.r-project.org/package=tikzDevice>
- Şimşek, Ö., & Buckmann, M. (2015). Learning from small samples: An analysis of simple decision heuristics. In *Advances in neural information processing systems* (pp. 3159–3167).
- Speekenbrink, M., & Shanks, D. R. (2010). Learning in a changing environment. *Journal of Experimental Psychology: General*, 139, 266–298. doi: 10.1037/a0018620
- Stojic, H., Analytis, P. P., & Olsson, H. (2020, Apr). *Project files for "linear integration and lexicographic models of choice: A cue weight learning perspective"*. Open Science Framework. Retrieved from osf.io/7yxrz
- Stojic, H., Analytis, P. P., & Speekenbrink, M. (2015). Hu-

- man behavior in contextual multi-armed bandit problems. In *Proceedings of the Thirty-Seventh Annual Conference of the Cognitive Science Society* (pp. 2290–2295).
- Stojic, H., Olsson, H., & Analytis, P. P. (2016). Explaining inter-individual variability in strategy selection: A cue weight learning approach. In *Proceedings of the 14th International Conference on Cognitive Modeling* (pp. 144–150).
- Stojic, H., Olsson, H., & P. Analytis, P. (2020, April). *Linear integration and lexicographic models of choice: A cue weight learning perspective*. PsyArXiv. Retrieved from psyarxiv.com/yj2zc doi: 10.31234/osf.io/yj2zc
- Stojić, H., Olsson, H., & Speekenbrink, M. (2016). *Not everything looks like a nail: Learning to select appropriate decision strategies in multiple environments*. Retrieved from <https://osf.io/fma3p/> doi: 10.17605/OSF.IO/FMA3P
- Stojic, H., Schulz, E., P. Analytis, P., & Speekenbrink, M. (2020). It's new, but is it good? How generalization and uncertainty guide the exploration of novel options. *Journal of Experimental Psychology: General*. doi: <http://dx.doi.org/10.1037/xge0000749>
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
- Todd, P. M., & Dieckmann, A. (2005). Heuristics for ordering cue search in decision making. In *Advances in neural information processing systems* (pp. 1393–1400).
- Todd, P. M., & Gigerenzer, G. (2000). Précis of Simple heuristics that make us smart. *Behavioral and Brain Sciences*, 23, 727–780.
- Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, 79, 281.
- Wagenmakers, E.-J., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*, 11, 192–196. doi: 10.3758/BF03206482
- Wickham, H. (2007). Reshaping Data with the reshape Package. *Journal of Statistical Software*, 21(12), 1–20. Retrieved from <http://www.jstatsoft.org/v21/i12/>
- Wickham, H. (2009). *ggplot2: elegant graphics for data analysis*. Springer New York. Retrieved from <http://had.co.nz/ggplot2/book>
- Wickham, H., & Francois, R. (2015). *dplyr: A Grammar of Data Manipulation*. Retrieved from <http://cran.r-project.org/package=dplyr>
- Widrow, B., & Hoff, M. E. (1960). *Adaptive switching circuits* (Tech. Rep.). Stanford, CA, US: Stanford Univ Ca Stanford Electronics Labs.
- Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2018). Generalization guides human exploration in vast decision spaces. *Nature Human Behaviour*, 2, 915.
- Wuertz, D. (2013). *fOptions: Basics of Option Valuation*. Retrieved from <http://cran.r-project.org/package=fOptions>

Appendix

Benchmark prediction models

To provide a benchmark to participants' performance in the estimation task and a comparison to the LMS and LMS_d models, we use two baseline prediction (BP) models and perfect LMS and LMS_d based learners. Baseline models provide a lower bound on the performance, the worst we would expect to see, while perfect learners provide an upper bound on the performance, the best we would expect to see.

Simplest baseline model would be a random guessing model. However, there is a more sensible strategy that still does not involve any learning of cue-criterion relationship. They could keep an estimate of a mean criterion value based on feedback in the choice task, and use this mean as a fixed prediction in the estimation task. It is a simple strategy that requires minimal learning and can provide better performance in the estimation task than random guesses. We formulate this strategy as one baseline prediction model that perfectly computes the mean (BP_f) and another that learns the mean slowly trial-by-trial (BP_l), potentially arriving at an imperfect estimate of the mean. These two strategies are the best that participants that learn only cue order during the choice task could do.

More precisely, the BP_f model estimates the mean by taking the arithmetic average across all trials in the choice task $\hat{P} = \frac{1}{T_c} \sum_{t=1}^{T_c} Y_t$, while the BP_l model updates the mean estimate after each observation according to the delta rule

$$\hat{P}_{t+1} = \hat{P}_t + \eta(Y_t - \hat{P}_t) \quad (5)$$

where $\eta \in (0, 1)$ is a learning rate parameter. $T_c = 160$ is the number of trials in the choice task.

The perfect LMS and LMS_d models provide an upper bound on prediction performance. They are fitted to predictions of a perfect learner, that is, a simulated participant that managed to produce completely accurate predictions in the estimation task. These models still learn the cue weights based on choices and feedback of each individual participant data from the choice task. However, they have learning rates that allow them to learn as much as possible from the feedback, trying to achieve perfectly accurate predictions in the estimation task.

For details on parameter estimation and model selection please see *Supplemental Material*.

Benchmark choice models

As for the estimation task, we provide benchmark models for the choice task, two baseline choice (BC) models and perfect WADD and DE models. Again, baseline models provide a lower bound on the performance, the worst we would expect to see, while perfect learners provide an upper bound on the performance, the best we would expect to see.

Baseline choice models consist of only basic information processing and do not require any knowledge about

cues. The BC_f model captures simple random choice, where the probability of choosing any option is equal and fixed at $1/3$. The BC_m model is more flexible. The probabilities of choosing options are also fixed, but rather than set to being equal they were estimated from the empirical frequencies of choosing them. This is the so-called multinomial model according to which participants choose randomly, but allowing for the possibility that participants have a response bias for either left, center or right option.

Perfect WADD and DE models are fitted to choices of a perfect decision maker, that is, a simulated participant that

always made correct choices in the choice task, selecting options with highest criterion value in every trial. These models use cue weights learned by the best fitting perfect LMS model. Relying on the best possible cue weights for each individual, these models have tremble error (WADD and DE) and Δ (DE) parameter that allow them to make choices as close as possible to the perfectly accurate choices.

For details on parameter estimation and model selection please see *Supplemental Material*.

Supplemental Material for “Linear integration and lexicographic models of choice: A cue weight learning perspective”

Hrvoje Stojić*
Pompeu Fabra University

Henrik Olsson
Santa Fe Institute

Pantelis P. Analytis
University of Southern Denmark

Hrvoje Stojić, Department of Economics and Business, Pompeu Fabra University; Henrik Olsson, Santa Fe Institute; Pantelis P. Analytis, Danish Institute of Advanced Studies and Department of Management and Marketing, University of Southern Denmark.

*Correspondence should be addressed to Hrvoje Stojić, who is now at Max Planck UCL Centre for Computational Psychiatry and Ageing Research, University College London, 10-12 Russell Square, London, WC1B 5EH, United Kingdom. E-mail: h.stojic@ucl.ac.uk

Additional details on method

Parameter estimation in the estimation task

We fit each prediction model (LMS, LMS_d, BP_f and BP_l) to individual participant's predictions in the estimation task by minimizing the sum of squared error (SSE) between participant's predictions (P_t) and model predictions ($\hat{P}_t$):

$$\text{SSE} = \sum_{t=1}^{T_e} (P_t - \hat{P}_t)^2 \quad (1)$$

where P_t is participant's prediction and $\hat{P}_t$ is prediction of the model in trial t , while $T_e = 40$ is the number of trials in the estimation task.

We fit the perfect LMS and LMS_d models by using criterion values Y_t instead of individual predictions P_t in Equation 1, essentially using a simulated participant that managed to make perfectly accurate predictions in the estimation task. Although there is a single set of criterion values in the estimation task across all participants, choice trajectories of participants differed and cue weight learning in the choice task could yield slightly different cue weights by the end of it, resulting in variation in fits of perfect LMS and LMS_d across participants.

To optimize the parameters we use the Nelder–Mead simplex algorithm implemented in the `optim` function in R (R Core Team, 2017). We use multiple starting points that we choose by first evaluating a uniform distributed grid of initial parameters, performing K-means clustering on them and choosing the best parameter set in each cluster of points.

Parameter estimation in the choice task

We fit each choice model (WADD, DE, BC_f and BC_m) to individual participant's choices in the choice task by minimizing the log likelihood of the data given the choice probabilities predicted by the model. The likelihood of the data set, L , is given by

$$L = \prod_{t=1}^{T_c} P(M_t = C_t) \quad (2)$$

where $T_c = 160$ is the number of trials in the choice task, and $P(M_t = C_t)$ is probability of model making the same choice as participant made in trial t . We use the same algorithm and procedure as in the modeling of the estimation task to optimize the parameters.

We fit the perfect WADD and DE models by using choices corresponding to option i with maximum criterion value $\text{argmax}_i Y_t^i$ in each trial instead of participants' choices C_t in Equation 2. In other words, we use a simulated participant that managed to make perfectly accurate choices in the choice task.

Model Selection

As a measure of model performance we use Bayesian information criterion (BIC, the smaller the better; Schwarz, 1978). The BIC score in the choice task is a monotonic transformation of the likelihood of the choice data

$$\text{BIC} = -2\ln(L) + N_{\text{par}}\ln(T_c) \quad (3)$$

where $\ln(L)$ is the log transformation of the choice data likelihood, N_{par} is the number of free parameters in the model, and $T_c = 160$ is the number of trials in the choice task.

The BIC score in the estimation task is a similar transformation of the SSE of the estimation data

$$\text{BIC} = T_e\ln(\text{SSE}/T_e) + N_{\text{par}}\ln(T_e) \quad (4)$$

where $T_e = 40$ is the number of trials in the estimation task.

While BIC scores take into account differences in complexity of the models by penalizing models with more parameters, differences between the models can be hard to interpret. Hence, for model comparison we use BIC weights, a simple transformation of BIC scores that can be directly interpreted as conditional probabilities that each model is the best one among the models compared (Wagenmakers & Farrell, 2004).

First we subtract BIC score of the best model in the comparison set, $\text{BIC}^{\min}$, from BIC score of each model i

$$\Delta\text{BIC}_i = \exp[-0.5(\text{BIC} - \text{BIC}^{\min})] \quad (5)$$

Then we normalize the relative BIC scores to obtain the BIC weights

$$\text{BIC}w_i = \frac{\Delta\text{BIC}_i}{\sum_{k=1}^K \Delta\text{BIC}_k} \quad (6)$$

where $K \in \{\text{WADD}, \text{DE}, \text{BC}_f, \text{BC}_m\}$ in the choice task and $K \in \{\text{LMS}, \text{LMS}_d, \text{BP}_f, \text{BP}_l\}$ in the estimation task. This makes the weights dependent on models that are being compared, stressing the relative aspect of the model comparison.

Additional details on results

Cross-validation model selection results

We also performed cross-validation, an alternative model selection procedure, as a robustness check. Although BIC is widely applied and is supposed to approximate cross-validation in theory, it can be informative to compare our BIC based results to those obtained by cross-validation.

We estimated parameters of the choice models on the first three blocks of the choice task and tested their predictions on the fourth block. The results were similar to those obtained with model selection based on BIC. Mean log-likelihood deviance (computed for test trials) for DE was 78 while for WADD it was 74. Log-likelihood weights showed that the probability that DE was the best model was 0.30 while for WADD it was 0.54. These results show that the balance shifted further in favour of WADD, but not by much.

When we used cross-validation results to classify participants into strategy users we obtained 23 DE users, 44 WADD users and 10 baseline choice model users. These results are again similar to BIC based results, with the balance shifted more toward WADD. The confidence with which we classified participants was somewhat lower, as indicated by the mean log-likelihood weights (computed only for participants best fitted for a model) – 0.73 for DE and 0.82 for WADD. This was to be expected, due to having less training data for fitting the models and generalization nature of the cross-validation.

Table S1

Results of modeling predictions in the estimation task and choices in the choice task. Choice models used the cue weights inferred by the best fitting LMS model for each participant.

Model	N_{par}	BIC	BICw	N_{best}	η	γ	Δ	ϵ
<i>Predictions</i>								
BP_f	1	405 (2.77)	0.03 (0.02)	3 (0.66)	—	—	—	—
BP_l	1	384 (3.33)	0.12 (0.03)	8 (0.79)	1.9e-2 (1.0e-2)	—	—	—
LMS	1	369 (2.87)	0.35 (0.04)	31 (0.75)	1.2e-5 (1.8e-6)	—	—	—
LMS_d	2	365 (2.85)	0.50 (0.04)	35 (0.89)	1.7e-4 (2.4e-5)	0.62 (0.04)	—	—
Perfect LMS	1	319 (1.06)	0.05 (0.02)	2 (0.72)	3.6e-5 (6.0e-7)	—	—	—
Perfect LMS_d	2	305 (0.98)	0.95 (0.02)	75 (0.96)	3.0e-4 (1.9e-5)	0.50 (0.01)	—	—
<i>Choices</i>								
BC_f	0	352 (0)	0.11 (0.03)	9 (0.82)	—	—	—	—
BC_m	2	360 (0.26)	0.01 (0.01)	1 (0.60)	—	—	—	—
DE	2	306 (4.61)	0.38 (0.05)	29 (0.94)	—	—	34.0 (3.09)	0.45 (0.02)
WADD	1	302 (4.94)	0.50 (0.05)	38 (0.93)	—	—	—	0.62 (0.02)
Perfect DE	2	221 (2.35)	0.01 (0.001)	0 (—)	—	—	22.2 (0.26)	0.30 (0.01)
Perfect WADD	1	137 (4.64)	0.99 (0.001)	77 (—)	—	—	—	0.16 (0.01)

Note. BP = Baseline prediction model with estimated or learned mean (suffixes f and l respectively); LMS = Least-mean-squares network model with or without additional decay of the learning rate (suffix d); BC = Random choice model with fixed or estimated choice probabilities (suffixes f or m respectively); WADD = Weighted-additive model; DE = Delta-elimination model; N_{par} = Number of free parameters of the model; BIC = Bayesian information criterion; BICw = Bayesian information criterion weights; N_{best} = Number of participants best fitted by the model according to BICw; η and γ = Parameters in the LMS models; ϵ and Δ = parameters in the WADD and DE models. In all columns except for the first two we display the mean and standard error of the mean in parenthesis below it, with exception for N_{best} where in parenthesis we show mean BICw for those participants).

Figure S1. (a) The number of elements inspected by the delta-elimination (DE) heuristic as a function of the discrimination threshold (Δ) value, obtained in simulations on our choice task using objective cue weights. The number of elements associated with the optimal threshold value of 19 (vertical dashed line), is 4.8 (horizontal dashed line). This is much fewer number of elements than the maximum of 12 for the weighted-additive strategy. *(b)* The model evidence, i.e. how well the fitted DE strategy predicts choices (the lower the better), and the estimated discrimination threshold (Δ) parameter for all participants. The shape of the point marks the model fitting the participant's choices best: weighted-additive (WADD), DE or one of the baseline choice (BC) models.

Figure S2. (a) Participants best fitted with an LMS model with decaying learning rate (LMS_d , $N = 35$) exhibited a faster evolution of subjective cue weights at the beginning and slower changes later on than those best fitted by the model without decay (LMS, $N = 31$). We illustrate median cue weight across participants for each trial, smoothed with a moving average of ten trials. (b) This pattern matched the differences in insight question ratings during the choice task when we divided participants into those best fitted with the LMS and LMS_d model. Participants best fitted with LMS_d model had more accurate ratings already in the first instance when the insight questions were asked and smaller changes later on, while participants best fitted with the LMS model exhibited more gradual changes throughout the choice task. In both panels participants best fitted with one of the baseline prediction models (BP_f or BP_l) were not included.

Figure S3. (a) Insight questions data of both weighted-additive (WADD) strategy users (participants best fitted by WADD model) and delta-elimination (DE) users indicates they learned the cue order fairly well. (b) Insight questions data for baseline choice (BC) model users are mixed (pooled across both model variants). By the end of the choice task they seem to have gained some knowledge about the structure of the task. In (a) and (b) the points at trials 40, 80, 120 and 160 in the choice task are mean ratings from the arbitrary insight question scale, how strongly each feature relates to the criterion values, on a scale from -10 (highly negative) to 10 (highly positive). To allow for easier discrimination, ratings are displaced horizontally and connected by lines. Error bars represent standard errors of group means. (c) Evolution of cue weights in the choice task, as inferred by the best fitting least-mean-squares model (LMS or LMS_d), shown for participants best fitted with one of the BC models. They exhibit very limited knowledge by the end of the choice task. Weights are median cue weights across participants for each trial, smoothed with a moving average of ten trials.

References

- R Core Team. (2017). R: A Language and Environment for Statistical Computing [Computer software manual]. Vienna, Austria. Retrieved from <https://www.r-project.org/>
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6, 461–464.
- Wagenmakers, E.-J., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*, 11, 192–196. doi: 10.3758/BF03206482