Decomposing conditioned avoidance performance with computational models

Angelos-Miltiadis Krypotos[1,2], Geert Crombez[3], Ann Meulders[1,4], Nathalie Claes[1], & Johan W.S. Vlaeyen[1,4]

[1] Department of Health Psychology, KU Leuven

[2] Department of Clinical Psychology, Utrecht University

[3] Department of Experimental-Clinical and Heath Psychology, Ghent University

[4] Experimental Health Psychology, Maastricht University

Paper in press in Behaviour Research and Therapy

Author Note

The present work was supported by an FWO grant (Reg. # G071118N) awarded to JWSV and GC. AMK is supported by a senior post-doctoral grant from FWO (Reg. # 12X5320N) and a replication grant from NWO (Reg. # 401.18.056).

Correspondence concerning this article should be addressed to Angelos-Miltiadis Krypotos, Tiensestraat 102, 3000 Leuven, Belgium. E-mail: amkrypotos@gmail.com.

Abstract

Avoidance towards innocuous stimuli is a key characteristic across anxiety-related disorders and chronic pain. Insights into the relevant learning processes of avoidance are often gained via laboratory procedures, where individuals learn to avoid stimuli or movements that have been previously associated with an aversive stimulus. Typically, statistical analyses of data gathered with conditioned avoidance procedures include frequency data, for example, the number of times a participant has avoided an aversive stimulus. Here, we argue that further insights into the underlying processes of avoidance behavior could be unraveled using computational models of behavior. We then demonstrate how computational models could be used by reanalysing a previously published avoidance data set and interpreting the key findings. We conclude our article by listing some challenges in the direct application of computational modelling to avoidance data sets.

Keywords: pain, anxiety-related disorders, fear, escape, computational modeling

Decomposing conditioned avoidance performance with computational models

Excessive avoidance towards innocuous cues, situations, or movements is a central characteristic of anxiety-related disorders and chronic pain (American Psychiatric Association, 2013; Treede et al., 2015). Avoidance of aversive events is generally a protective strategy (e.g. avoid passing a red traffic light), fostering an organism's well-being and survival (i.e., *adaptive* avoidance). When avoidance is expressed in largely safe situations (e.g., flight phobics avoiding boarding planes), it may impede individuals' everyday functioning (i.e., *maladaptive* avoidance). To date, a major bulk of research in Clinical and Health Psychology focuses on unveiling the basic conditions under which adaptive and maladaptive avoidance is acquired, via mostly laboratory research (Arnaudova, Kindt, Fanselow, & Beckers, 2017; Beckers & Craske, 2017; Dymond, 2019; Meulders, Franssen, Fonteyne, & Vlaeyen, 2016; Pittig, Wong, Glück, & Boschet, 2020). It has been argued that by studying the basic processes of avoidance learning in experimental settings, more insight can be gained regarding both adaptive and maladaptive forms of avoidance (Krypotos, Vervliet, & Engelhard, 2018).

Avoidance learning procedures typically entail the learning of pairings between different stimuli/contexts with an aversive event and how the occurrence of such events can be prevented (Krypotos, Effting, Kindt, & Beckers, 2015; LeDoux, Moscarello, Sears, & Campese, 2017; Pittig et al., 2020). Such procedures have proven paramount in unveiling the basic learning processes (e.g., Mowrer, 1951), and the relevant neural substrates (LeDoux & Daw, 2018) of how defensive responses (e.g., escape or avoidance) arise. To date, avoidance learning procedures have been extended so as to include costs (e.g., Pittig & Dehler, 2019; Rattel, Miedl, Blechert, & Wilhelm, 2017), competing goals (e.g., Claes, Karos, Meulders, Crombez, & Vlaeyen, 2014), and virtual reality contexts (e.g., Glotzbach, Ewald, Andreatta, Pauli, & Mühlberger, 2012).

Most often researchers use traditional analytic methods confined to the experimental design to make statistic inferences. Specifically, these entail the binomial categorization of avoidance and non-avoidance responses, and the corresponding calculations of the rate of (non-)avoidance. If, for example, a participant avoided in 5 out of the 10 trials, then the avoidance rate is 50% (Dymond, Roche, Forsyth, Whelan, & Rhoden, 2008; Krypotos & Engelhard, 2018; Vervliet & Indekeu, 2015). Alternatively, when using a continuous measure of avoidance, avoidance may be defined as the distance that the participant moves a virtual stimulus (e.g., an avatar) from a source of potential threat (Meulders et al., 2016; Mobbs et al., 2007). Despite the two approaches of data handling offer a straightforward way to measure avoidance responding, they are ambivalent about the involved learning processes. To illustrate, following Mowrer's 'two-factor theory' (Mowrer, 1951), the observed behavior could be attributed to current fear levels, with the participants emitting the observed behavior in order to *escape* the *currently* experienced fear, whereas according to Seligman and Johnston (1973), *avoidance* is the result of expecting an *impending* aversive event to occur. Arguably, our field could move forward faster if our analytic approaches were better able to disentangle the underlying processes involved during avoidance learning.

Assuming a well-executed experiment, one way to shed light on the underlying processes of avoidance is by using a *computational model* (Sharp & Eldar, 2019; Sutton & Barto, 2018). A computational model is here defined as the quantitative formalization of theories, hypotheses, and descriptions of narrative models (Schrater, Körding, & Blohm, 2019; see also Stafford, 2009; Thagard, 2018 for relevant discussions about computational models and their relations to philosophy of science). The goal of computational modeling is to link theoretical models of psychological processes (e.g., fear, expectancies) to the observed data (e.g., avoidance rates). As such, it is not a method for drawing statistical inferences of different distributions (e.g., in case of a *t*-test) (Lee, 2011). Computational models rely on translating the key latent processes of a theory (e.g., the *prediction* that a stimulus will be followed by an aversive event), into mathematical algorithms that link the latent processes with the observed behavior (e.g., frequency of avoidance).

To give an example, many *reinforcement computational models* characterize behavioral experiments using two key parameters (e.g., Ahn et al., 2014; Lindström, Selbing, Molapour, & Olsson, 2014). The first one, is *learning rate*, which refers to the impact that the discrepancy between an expected value (e.g., if I do press the button, I will not receive a painful stimulus) and the actual value (e.g., painful stimulus presentation) of an action may have on the value of the action the next time it is to be performed. A low learning rate means that many trials are considered in order to determine the current value of the action, whereas a high learning rate means that less trials are considered. The second key parameter is the *outcome impact/sensitivity*. This parameter reflects how much one values a reward (e.g., not receiving an electrocutaneous stimulus or receiving a monetary reward), with higher values meaning that a reward is highly liked and as such the emitted action will be repeated, or disliked a punishment (e.g., receiving an electrocutaneous stimulus or not receiving a monetary reward), with again higher values meaning that dislike is strong and as such the participant is expected to emit a different behavior. Naturally, the number and type of parameters may differ per model and it is up to the researcher to select, or build, theoretically principled models.

There are at least three important reasons why research in avoidance learning could benefit from computational modeling (see Marr, 1982 for extensive discussions). First, multiple computational models, each referring to the underlying latent processes of different narrative models, can be directly compared (Claeskens & Hjort, 2008). As such, multiple theories could be pitted with each other using a single data set. Second, by relying on computational formalization of theories, rather than only narrative accounts of them, researchers can make more specific predictions about the to-be-obtained results, rather than rely on a proposed pattern of results (e.g., one group will avoid less than the second one). This approach allows the easier modification, or falsification, of a model (see Meehl (1990);

Palminteri, Wyart, and Koechlin (2017); for extensive discussions). Model falsification is especially important given the recent replication crisis where non-replications, even when the same procedures are used as in the original studies, are often not regarded as direct evidence against a theory (Wagenmakers et al., 2016). Third, there are specific model variances (e.g., *hierarchical computational models*) that could help in the investigation of individual differences, an area of inquiry especially relevant in the field of fear and avoidance (Lonsdorf & Merz, 2017). Given these major advantages, computational models are widely used across scientific fields, such as neuroscience (Forstmann, Wagenmakers, & others, 2015) or computational psychiatry (Huys, Moutoussis, & Williams, 2011; Maia & Frank, 2011; Wang & Krystal, 2014). However, they are still to be applied systematically in our field. Here, we propose that computational models should also be used more routinely in avoidance learning.

Example data set

In order to give an example of how computational modelling could be used in avoidance learning, we reanalyzed part of the data set originally reported in Meulders et al. (2016). In that study, participants learned to perform three different movement paths with a robotic arm to a target location, with each movement paths being associated with different probabilities of receiving a painful electrocutaneous stimulus as well as different effort to perform the movement (i.e., the associated cost).

Specifically, *Path 1* (i.e., the shortest path) was associated with a reinforcement rate of 100% but no resistance (i.e., no *costs*), *Path 2* (i.e., the middle path) was associated with a reinforcement rate of 50% and medium resistance, and *Path 3* (i.e., the longest path) was associated with a 0% reinforcement rate but the most resistance compared to the other paths. In terms of avoidance, participants could avoid the electrocutaneous stimulus 100% of the times by choosing *Path 3* , 50% of the times by choosing *Path 2*, and 0% of the times by choosing *Path 1.* Two groups were tested, an experimental group and a yoked control group, with each group undergoing through the following phases: 1) *practice*, 2) *acquisition*, 3) *extinction*, 4) *instructed movement under extinction*. For the purposes of our illustration, we focus only on the *acquisition* phase (48 trials), which was the most relevant phase for avoidance's acquisition. We also present the relevant data for the *extinction* phase in the supplemental material.

We chose to fit the data only in the experimental group ($N = 25$) as in yoked control group the presentation of the painful electrocutaneous stimulus was independent of participants' performance.

In the original study of Meulders et al. (2016), the authors quantified avoidance per phase as the mean maximum deviation of a virtual ball from its starting position (i.e., extreme left on the screen) to the target position (i.e., the extreme left, *Path 1*, middle *Path 2*, or extreme right parts of a screen, *Path 3*). The left panel of Figure 1 visualizes the mean distance from the initial starting position to one of the three paths. We observe that there is a gradual increase from choosing Path 1 and 2 to choosing Path 3.

For our modelling approach, we categorized responses in one of the three paths (i.e., Path 1, Path 2, Path 3). Then, we fitted different reinforcement computational learning models to the data. These models have been previously used in the literature (see Ouden et al., 2013; Ahn et al., 2014; Aylward et al., 2019; Haines, Vassileva, & Ahn, 2018; Seymour, Daw, Roiser, Dayan, & Dolan, 2012; Worthy, Pang, & Byrne, 2013) and are readily available at the `hBayesDM` (Ahn, Haines, & Zhang, 2017) package[1] for the `R` (R Core Team, 2018) or `Python` (Van Rossum & Drake, 2011) programming languages. Note, that although the used computational models were originally developed for n-bandit tasks (Sutton & Barto, 2018) or the Iowa Gambling Task (Buelow & Suhr, 2009), all models encompass basic reinforcement learning principles (e.g., the use of *prediction errors*). However, models differ

[1] We have modified the models slightly as they referred to experiments were participants had to choose among 4 options, instead of 3 as it is here.

on different levels. For example if they include a *lapse* (or noise) parameter that would account for random choices that do not follow the inferred values of each path, or the *decay* parameter that reflects the reduction in the weight that each choice gets the less it is chosen.[2] Our analysis at this stage is exploratory and used for didactic reasons. We do not have specific predictions about the pattern of results.

To decide which model fits the data best, we compared the Leave-out-One value (Gelman, Hwang, & Vehtari, 2014; Vehtari, Gelman, & Gabry, 2017) of each model with each other, and chose the model with the lowest value. Leave-out-One cross validation entails the training of a computational model on all but one observation in the data set. Then, this excluded observation is the test data set, where a predictive score is computed. After performing the same procedure for all observations in the data set, the predictive performance of the model is the sum of all the computed, predicted scores. The model with the *lowest* Leave-out-One value is the one with the best performance. Table 1 shows the Leave-out-One values for each model fitted in the present data.[3] The model with the lowest value is that described by Haines et al. (2018). The relevant model parameters are: *Reward learning rate*, *Punishment learning rate*, *Outcome frequency weight*, and *Perseverance weight*. Below, we present some of the most significant parameters of the model and their meaning, with the explanation of the parameters being available in Table 2.

According to the initial study of Meulders et al. (2016), the main conclusion drawn from the acquisition data was that participants in the experimental group chose Path 3. This conclusion, however, in not informative about how individuals chose between the different paths. For example, it is unclear whether participants preferred to explore the

[2] As the goal of this article was not to give an extensive review of each of the tested models, we refer to Ahn et al. (2017) for detailed description of each model and each model parameter.

[3] We have modified the models slightly as they referred to experiments were participants had to choose among 4 options, instead of 3 as it is here.

different options than exploit a single one.

This information could be acquired by the analyses of our computational model. We now present how the parameter estimates could provide a richer interpretation about the same data set. We start with the learning rate parameter. If a *learning rate* parameter (see also above) has a value closer to 1, it suggests that more weight is given to recent outcomes whereas for values closer to 0 more weight is placed on older outcomes. The selected model had two different learning rates, one for rewards — defined as not receiving the painful stimulus and putting the maximum effort — and punishments – defined as receiving a painful stimulus but putting the minimum resistance effort. The parameter value reveals that participants needed more trials to learn which stimulus was followed by rewards ($M = 0.05$) than punishments ($M = 0.26$). This parameter provides an assessment of whether outcome recency plays a role. In other words, individuals will learn more from recent outcomes rather than outcomes happening earlier in the experiment. The *perseverance decay* parameter indicates the tendency of the participants to repeat the same choice as the previous trial (positive values) or switch (negative values). The mean value of this parameter ($M = .81$) shows the tendency to repeat the previous trials, something in line with also the descriptive data (see Figure 1). This parameter is especially relevant in decision-making literature as how individuals balance between exploration-exploitation is a popular research topic across scientific fields (see Mehlhorn et al., 2015 for a recent review on the exploration-exploitation dilemma). The *outcome frequency weight* indicate the weight that participants put on the outcomes of each stimulus. The negative value ($M = 1.72$) indicates the preference to stimuli with low rewards, or in our case the stimuli with higher chances of not receiving a painful stimulus and maximum resistance. Lastly, the perseverance weight again indicates the switch or non-switch in choices, but now for each stimulus at the current trial. The negative value indicates that participants have the tendency to switch from recent trials, but the large standard deviations ($SD = 0.55$) make this parameter difficult to interpret in this specific data set. Collectively, the results provide a deeper insight of the learning processes in avoidance, in particular how participants choose between potential aversive or pleasant stimuli, whether they tend to explore or exploit more, and whether their performance was based on what was learned in the recent trials or not. Such insights would be difficult to obtain by using a standard model for drawing statistical inferences where the analyses are focused more on whether an individual avoided an aversive stimulus or not.

Collectively, the results provide a deeper insight of the learning processes in avoidance, in particular how participants choose between potential aversive or pleasant stimuli, whether they tend to explore or exploit more, and whether their performance was based on what was learned in the recent trials or not. Such insights would be difficult to obtain by using a standard model for drawing statistical inferences where the analyses are focused more on whether an individual avoided an aversive stimulus or not.

In follow-up studies, predictions can be formulated based on the meaning of each parameter. For example, it would be worthwhile to test how the value of the learning rate changes in case the value of the outcome changes. If a more aversive outcome would result in lower learning rates, or the opposite, could inform future studies and clinical therapists when tailoring interventions -– e.g., how many sessions are needed – to reduce avoidance in the clinic. It would be also informative to test how the perseverance decay parameter changes in presence of conflicting outcome, where each one of the different paths are now associated with a rewarding outcome (e.g., monetary outcome). Another relevant inquiry is how this parameter could be different if more individually tailored rewards were chosen rather than a one-size-fits-all approach. These examples serve more than little in showing that computational modeling makes room for many research questions both of scientific and clinical relevance.

Concluding remarks

We argue that computational modelling could be a helpful tool towards better understanding adaptive and maladaptive avoidance. As we have shown above, by analyzing avoidance data using theory-driven computational models, it is possible to draw conclusions about the latent processes that govern avoidance, something that following traditional statistical analyses would typically need several experiments, and several experimental manipulations. To showcase our arguments for computational modeling as well as for encouraging its broader use, we have reanalysed a published avoidance data set and by drawing inferences according to the values of each model parameter. We conclude our manuscript by sketching some avenues for future research as well as highlighting some challenges when using computational models.

Here, we have used computational models already presented in the literature. Although these models are based on reinforcement learning principles (Sutton & Barto, 2018), also present in avoidance learning theories (see Krypotos et al., 2015 for a review), we still miss the direct translation of avoidance learning models (e.g., models by Lovibond, 2006; Vlaeyen, Crombez, & Linton, 2016) to computational models that could be directly compared with each other. We hope that our work could inspire such translation in the future, something that could potentially transform past theories or even generate new ones.

One challenge when building original computational models is that mathematical and programming skills are required, which are not typically acquired during the standard curriculum of Psychology. Although nowadays there are programming platforms that allow the relatively easy implementation of complex models (e.g., Stan Development Team, 2018), challenges remain. These challenges stress once more the need for more interdisciplinary collaboration (Ledford, 2015).

Two words of caution should be mentioned. First, whenever using a computational model, researchers should follow a standard, principled workflow in order to ensure that the conclusions follow logically the results of the model, and all the relevant model assumptions have been met. Such workflows have been presented in the literature (e.g., Lee et al., 2019; Wilson & Collins, 2019). From this literature, we would like to highlight *3* suggestions. First, and prior to seeing the actual data, it is a good practice to simulate data according to the model parameters and subsequently fit the data to the model. Unless the parameter values are recovered, it is doubtful whether the model is appropriate for the real data. Second, after a model has been chosen from competitive models, the researcher should check whether predictions made by this data set now fit the real data. This can be done by the so-called post-prediction checks, where the model parameters are used for generating synthetic data and then these data are compared to the real data. To illustrate this point, on Figure 2, we plotted the real data against the data predicted by the model for two participants. For the first participant (left panel) the model does a poor job as it predicts the correct choice only 18 out of the 48 trials. On the other hand, for the second participant the model predicts correctly 44 out of 48 trials. Given that the two participants had different performance patterns – with participant 24 switching much more than participant 13 – then that should inform the experimenters that the model maybe not good enough whenever frequent switches are present. Lastly, and despite our enthusiasm for computational models, a computational model can never substitute a good experiment; a poorly designed experiment does not allow researchers to draw stable inferences of the underlying processes, no matter how principled the computational model might be. For example, before the interested reader may attempt to fit the data in an archived data set, we suggest to first think carefully whether the described models could be meaningful for the collected data or not. No matter how sophisticated a computational model could be, there is no escape from poor data quality.

To conclude, we argue that scientific progress in the field of avoidance learning could accelerate by using computational models for statistical analysis, rather than relying on statistical inferences based on traditional models. Although this change in the field comes with challenges, we are convinced that such challenges may not stand as obstacles towards better understanding how *adaptive* and *maladaptive* avoidance is learned and reduced, issues of high relevance for the individuals and the society as a whole.

References

Ahn, W.-Y., Haines, N., & Zhang, L. (2017). Revealing neurocomputational mechanisms of reinforcement learning and decision-making with the hBayesDM package. *Computational Psychiatry*, *1*, 24–57. Retrieved from https://doi.org/10.1162/CPSY_a_00002

Ahn, W.-Y., Vasilev, G., Lee, S.-H., Busemeyer, J. R., Kruschke, J. K., Bechara, A., & Vassileva, J. (2014). Decision-making in stimulant and opiate addicts in protracted abstinence: Evidence from computational modeling with pure users. *Frontiers in Psychology*, *5*, 849.

American Psychiatric Association. (2013). *Diagnostic and statistical manual of mental disorders (dsm-5)*. American Psychiatric Pub.

Arnaudova, I., Kindt, M., Fanselow, M., & Beckers, T. (2017). Pathways towards the proliferation of avoidance in anxiety and implications for treatment. *Behaviour Research and Therapy*, *96*, 3–13.

Aylward, J., Valton, V., Ahn, W.-Y., Bond, R. L., Dayan, P., Roiser, J. P., & Robinson, O. J. (2019). Altered learning under uncertainty in unmedicated mood and anxiety disorders. *Nature Human Behaviour*, *3*, 1116–1123.

Beckers, T., & Craske, M. G. (2017). Avoidance and decision making in anxiety: An introduction to the special issue. *Behaviour Research and Therapy*, *96*, 1.

Buelow, M. T., & Suhr, J. A. (2009). Construct validity of the Iowa gambling task. *Neuropsychology Review*, *19*, 102–114.

Claes, N., Karos, K., Meulders, A., Crombez, G., & Vlaeyen, J. W. (2014). Competing goals attenuate avoidance behavior in the context of pain. *The*

Journal of Pain, *15*, 1120–1129.

Claeskens, G., & Hjort, N. L. (2008). *Model selection and model averaging.* Cambridge University Press.

Dymond, S. (2019). Overcoming avoidance in anxiety disorders: The contributions of pavlovian and operant avoidance extinction methods. *Neuroscience & Biobehavioral Reviews*, *98*, 61–70.

Dymond, S., Roche, B., Forsyth, J. P., Whelan, R., & Rhoden, J. (2008). Derived avoidance learning: Transformation of avoidance response functions in accordance with same and opposite relational frames. *The Psychological Record*, *58*, 269–286.

Forstmann, B. U., Wagenmakers, E.-J., & others. (2015). *An introduction to model-based cognitive neuroscience*. Springer.

Gelman, A., Hwang, J., & Vehtari, A. (2014). Understanding predictive information criteria for Bayesian models. *Statistics and Computing*, *24*, 997–1016.

Glotzbach, E., Ewald, H., Andreatta, M., Pauli, P., & Mühlberger, A. (2012). Contextual fear conditioning predicts subsequent avoidance behaviour in a virtual reality environment. *Cognition & Emotion*, *26*, 1256–1272.

Haines, N., Vassileva, J., & Ahn, W.-Y. (2018). The outcome-representation learning model: A novel reinforcement learning model of the Iowa Gambling Task. *Cognitive Science*, *42*, 2534–2561.

Huys, Q. J., Moutoussis, M., & Williams, J. (2011). Are computational models of any use to psychiatry? *Neural Networks*, *24*, 544–551.

Krypotos, A.-M., Effting, M., Kindt, M., & Beckers, T. (2015). Avoidance learning: A review of theoretical models and recent developments. *Frontiers in Behavioral*

Neuroscience, *9*, 189.

Krypotos, A.-M., & Engelhard, I. M. (2018). Testing a novelty-based extinction procedure for the reduction of conditioned avoidance. *Journal of Behavior Therapy and Experimental Psychiatry*, *60*, 22–28.

Krypotos, A.-M., Vervliet, B., & Engelhard, I. M. (2018). The validity of human avoidance paradigms. *Behaviour Research and Therapy*, *111*, 99–105.

Ledford, H. (2015). Team science. *Nature*, *525*, 308–311.

LeDoux, J., & Daw, N. D. (2018). Surviving threats: Neural circuit and computational implications of a new taxonomy of defensive behaviour. *Nature Reviews Neuroscience*, *19*, 269–282.

LeDoux, J. E., Moscarello, J., Sears, R., & Campese, V. (2017). The birth, death and resurrection of avoidance: A reconceptualization of a troubled paradigm. *Molecular Psychiatry*, *22*, 24–36.

Lee, M. D. (2011). How cognitive modeling can benefit from hierarchical Bayesian models. *Journal of Mathematical Psychology*, *55*, 1–7.

Lee, M. D., Criss, A. H., Devezer, B., Donkin, C., Etz, A., Leite, F. P., … others. (2019). Robust modeling in cognitive science. *Computational Brain & Behavior*, 1–13.

Lindström, B., Selbing, I., Molapour, T., & Olsson, A. (2014). Racial bias shapes social reinforcement learning. *Psychological Science*, *25*, 711–719.

Lonsdorf, T. B., & Merz, C. J. (2017). More than just noise: Inter-individual differences in fear acquisition, extinction and return of fear in humans-biological, experiential, temperamental factors, and methodological pitfalls. *Neuroscience &*

Biobehavioral Reviews, *80*, 703–728.

Lovibond, P. (2006). Fear and avoidance: An integrated expectancy model.

Maia, T. V., & Frank, M. J. (2011). From reinforcement learning models to psychiatric and neurological disorders. *Nature Neuroscience*, *14*, 154–162.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: WH Freeman.

Meehl, P. E. (1990). Why summaries of research on psychological theories are often uninterpretable. *Psychological Reports*, *66*, 195–244.

Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M. D., Morgan, K., Braithwaite, V. A., … Gonzalez, C. (2015). Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, *2*, 191.

Meulders, A., Franssen, M., Fonteyne, R., & Vlaeyen, J. W. (2016). Acquisition and extinction of operant pain-related avoidance behavior using a 3 degrees-of-freedom robotic arm. *Pain*, *157*, 1094–1104.

Mobbs, D., Petrovic, P., Marchant, J. L., Hassabis, D., Weiskopf, N., Seymour, B., … Frith, C. D. (2007). When fear is near: Threat imminence elicits prefrontal-periaqueductal gray shifts in humans. *Science*, *317*, 1079–1083.

Mowrer, O. H. (1951). Two-factor learning theory: Summary and comment. *Psychological Review*, *58*, 350–354.

Ouden, H. E. den, Daw, N. D., Fernandez, G., Elshout, J. A., Rijpkema, M., Hoogman, M., … Cools, R. (2013). Dissociable effects of dopamine and serotonin on reversal learning. *Neuron*, *80*, 1090–1100.

Palminteri, S., Wyart, V., & Koechlin, E. (2017). The importance of falsification in computational cognitive modeling. *Trends in Cognitive Sciences*, *21*, 425–433.

Pittig, A., & Dehler, J. (2019). Same fear responses, less avoidance: Rewards competing with aversive outcomes do not buffer fear acquisition, but attenuate avoidance to accelerate subsequent fear extinction. *Behaviour Research and Therapy*, *112*, 1–11.

Pittig, A., Wong, A. H., Glück, V. M., & Boschet, J. M. (2020). Avoidance and its bi-directional relationship with conditioned fear: Mechanisms, moderators, and clinical implications. *Behaviour Research and Therapy*, 103550.

Rattel, J. A., Miedl, S. F., Blechert, J., & Wilhelm, F. H. (2017). Higher threat avoidance costs reduce avoidance behaviour which in turn promotes fear extinction in humans. *Behaviour Research and Therapy*, *96*, 37–46.

R Core Team. (2018). *R: A language and environment for statistical computing.* Vienna, Austria: R Foundation for Statistical Computing. Retrieved from https://www.R-project.org/

Schrater, P. R., Körding, K. P., & Blohm, G. (2019). Modeling in neuroscience as a decision process.

Seligman, M. E. P., & Johnston, J. C. (1973). A cognitive theory of avoidance learning. In D. B. McGuigan F. J.and Lumsden (Ed.), *Contemporary approaches to conditioning and learning* (pp. 69–110). Washington, DC: Winston & Sons Inc.

Seymour, B., Daw, N. D., Roiser, J. P., Dayan, P., & Dolan, R. (2012). Serotonin selectively modulates reward value in human decision-making. *Journal of Neuroscience*, *32*, 5833–5842.

Sharp, P. B., & Eldar, E. (2019). Computational models of anxiety: Nascent efforts and future directions. *Current Directions in Psychological Science*, *28*, 170–176.

Stafford, T. (2009). What use are computational models of cognitive processes? In J. Mayor, N. Ruh, & K. Plunkett (Eds.), *Connectionist models of behaviour and cognition ii: Proceedings of the 11th neural computation and psychology workshop.* World Scientific.

Stan Development Team. (2018). RStan: The R interface to Stan. Retrieved from http://mc-stan.org/

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction.* MIT press.

Thagard, P. (2018). Computational models in science and philosophy. In *Introduction to formal philosophy* (pp. 457–467). Springer.

Treede, R.-D., Rief, W., Barke, A., Aziz, Q., Bennett, M. I., Benoliel, R., … others. (2015). A classification of chronic pain for icd-11. *Pain*, *156*, 1003–1007.

Van Rossum, G., & Drake, F. L. (2011). *The python language reference manual.* Network Theory Ltd.

Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, *27*, 1413–1432.

Vervliet, B., & Indekeu, E. (2015). Low-cost avoidance behaviors are resistant to fear extinction in humans. *Frontiers in Behavioral Neuroscience*, *9*, 351.

Vlaeyen, J. W., Crombez, G., & Linton, S. J. (2016). The fear-avoidance model of pain. *Pain*, *157*, 1588–1589.

Wagenmakers, E.-J., Beek, T., Dijkhoff, L., Gronau, Q. F., Acosta, A., Adams Jr, R., ... others. (2016). Registered replication report: Strack, martin, and stepper (1988). *Perspectives on Psychological Science*, *11*, 917–928.

Wang, X.-J., & Krystal, J. H. (2014). Computational psychiatry. *Neuron*, *84*, 638–654.

Wilson, R. C., & Collins, A. G. (2019). Ten simple rules for the computational modeling of behavioral data. *eLife*, *8*, e49547.

Worthy, D. A., Pang, B., & Byrne, K. A. (2013). Decomposing the roles of perseveration and expected value representation in models of the Iowa gambling task. *Frontiers in Psychology*, *4*, 640.

Table 1

Model selection table. The Model names, first column, refers to the models as those referred in the hBayesDM R package (Ahn, et al. 2017). Please see Ahn, et al. (2017) for a detailed description of the parameters of each model. The LOOIC columns includes the leave-out-One information criterion for each model. Lastly, the Reference column includes the reference to where the relevant models are described.

Model	LOOIC	Reference
Hierarchical Bayesian Modeling of the Iowa Gambling Task using Outcome-Representation Learning Model	1945.88	Haines et al., (2018)
Hierarchical Bayesian Modeling of the Iowa Gambling Task using Value-Plus-Perseverance	1983.07	Worthy et al. (2013)
Hierarchical Bayesian Modeling of the Probabilistic Reversal Learning Task using Reward-Punishment Model	2159.89	Ouden et al., (2013)
Hierarchical Bayesian Modeling of the Iowa Gambling Task using Prospect Valence Learning Delta	2637.82	Ahn et al. (2008)
Hierarchical Bayesian Modeling of the Iowa Gambling Task using Prospect Valence Learning	2637.84	Ahn, et al., (2014)
Hierarchical Bayesian Modeling of the 3-Armed Bandit Task using 4 Parameter Model, without 'Choice perseveration' but with 'noise'	2638.78	Aylward et al., (2019)

Hierarchical Bayesian Modeling of the 3-Armed Bandit Task using 5 Parameter Model, without ’Choice perseveration’ but with ’noise’	2638.84	Seymour et al. (2012)
Hierarchical Bayesian Modeling of the 3-Armed Bandit Task using 3 Parameter Model, without ’choice perseveration’, ’reward sensitivity’, and ’punishment sensitivity’, but with ’noise’	2638.84	Aylward et al. (2019)
Hierarchical Bayesian Modeling of the 3-Armed Bandit Task using 5 Parameter Model, without ’Choice perseveration’ but with ’noise’	2640.33	Aylward et al. (2019)

Table 2

Mean, standard deviations, and description for each parameter for the winning model. Detailed explanation of the parameters is included in Haines et al. (2018).

	Mean	SD	Interpretation
Reward learning rate	0.05	0.04	Learning rate after receiving a reward (i.e., not receiving a painful stimulus) Values close to 1: More weight is given to recent outcomes. Values close to 0: More weight is placed on older outcomes.
Punishment learning rate	0.26	0.07	Learning rate after receiving a punishment (i.e., receiving a painful stimulus) Values close to 1: More weight is given to recent outcomes. Values close to 0: More weight is placed on older outcomes.
Perseverance decay	0.81	0.62	Negative values: tendency to switch the stimulus chosen on the previous trials. Positive values: tendency to persevere the stimulus chosen on the previous trials.
Outcome frequency weight	-1.72	0.37	Negative values: Preference for stimuli with low win frequency (i.e., high resistance) Positive values: Preference for stimuli with high win frequency (i.e., low resistance)
Perseverance weight	-0.10	0.55	Negative values: Switch from recently selected stimuli Positive values: Non-switch from recently selected stimuli

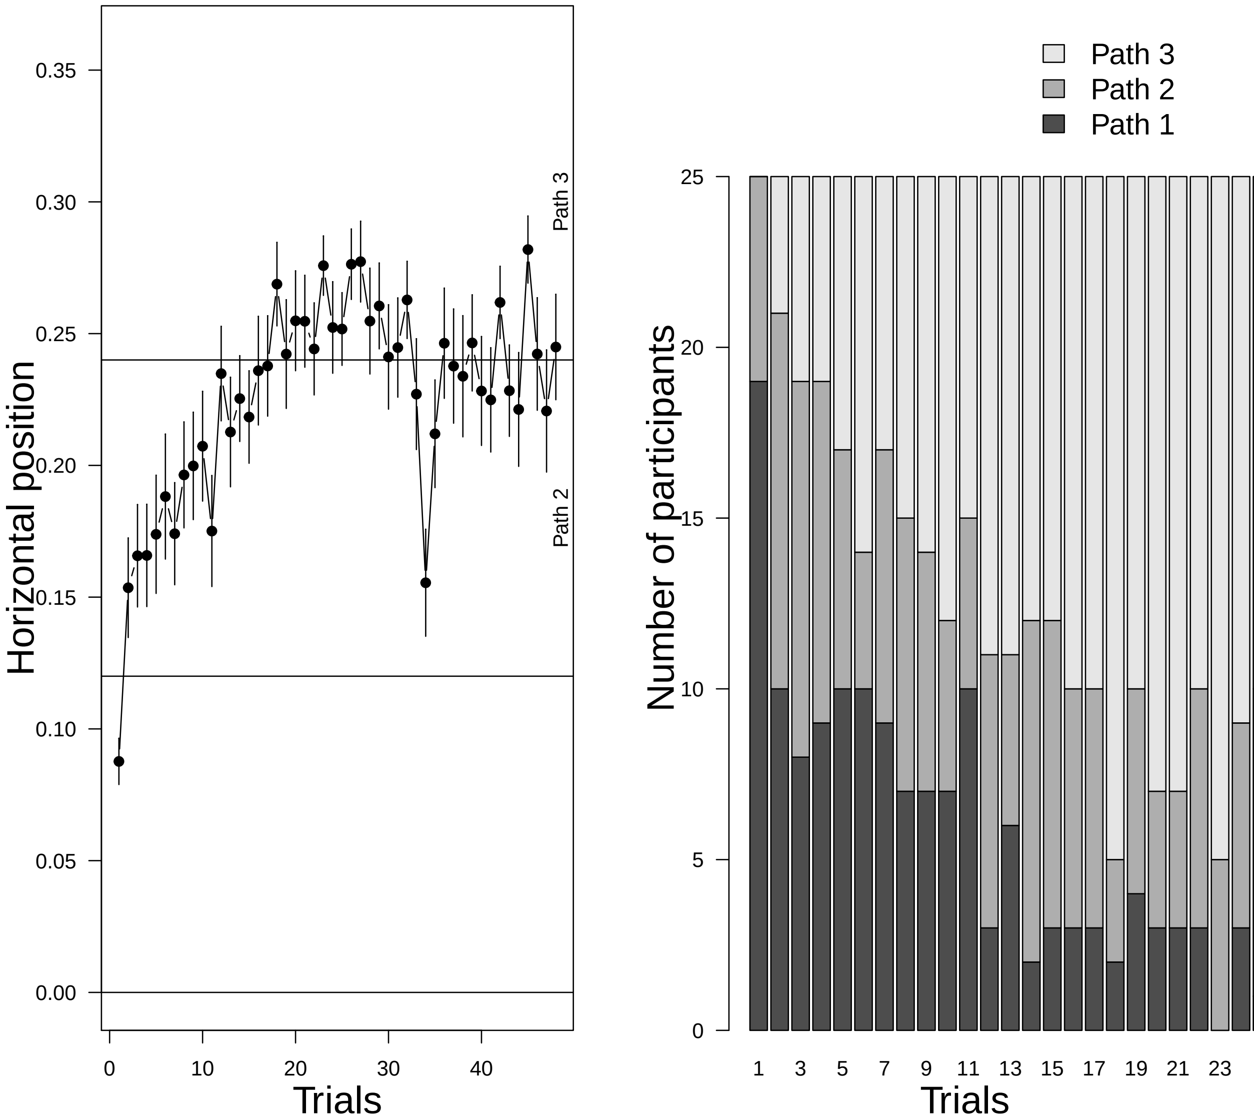

Figure 1. Left panel: Mean horizontal position that the haptic arm was moved during each extinction trial. With horizontal lines we visualise the limits for each path. Right panel: Frequencies of each one of the selected paths per trial (x-axis) and per participant (y-axis) of Meulders et al. (2016).

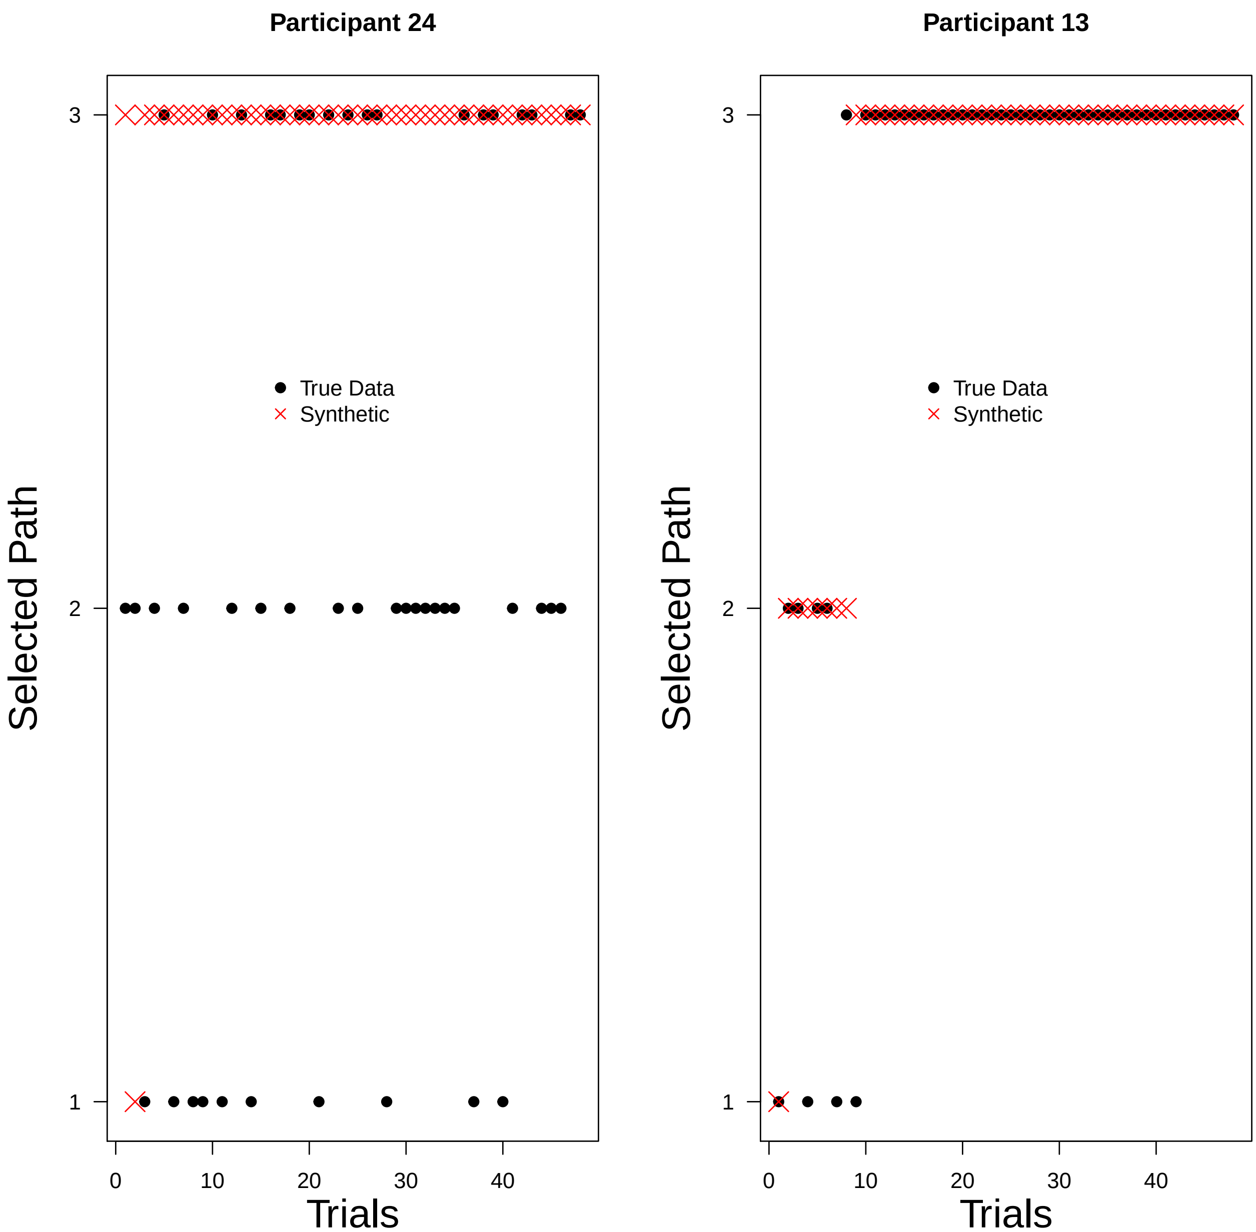

Figure 2. Post prediction check plots for Participant 1 (left panel) and 5 (right panel) of the used data set. The black dots refer to the choices made by each participant. The red crosses refer to the choices predicted by the model that fitted the data best (i.e., winning model). The model does a poor job for participant 1, as 14 out of the 24 trials were not predicted well by the model. On the other hand, the model does a good job for participant 5 where the model missed only 1 point.