When do you buy?
Predicting an individual's decision in optimal stopping problems

Manuel Guth[1] **and Marco Ragni**[1,2]
[1] Cognitive Computation Lab, Technical Faculty, Albert-Ludwigs-University Freiburg
[2] Danish Institute of Advanced Studies, Denmark
guthm@tf.uni-freiburg.de, ragni@sdu.dk

Abstract

Prices, e.g., for flight tickets can change almost daily. To minimize the costs, we have to decide when to take an action, i.e., when to buy. Suchs decision tasks are called *optimally stopping problems*. This paper reconsiders the strongest cognitive models that are able to predict the average decision maker, adapts them and investigate their predictive accuracy *on the individual level*, i.e., how good are models in predicting *when* a participant decides for an action. To perform this analyses, several steps are necessary: (i) Identify data sets that provide raw data for an individual, (ii) develop an individual testing framework to assess the models, (iii) implement and adapt existing models for the individual, and (iv) consider baseline models to assess the goodness-of-fit of the models for the individual. The best and second-best models achieved an overall prediction accuracy of 84.9% and 84.1% respectively. Five of the ten examined models managed to beat a strong baseline, proving that they did in fact managed to model the individual decision process.

Introduction

The *Optimal Stopping Problem* is implicitly present in many aspects of everyday life. When searching for the partner to spend life with, buying airplane tickets for the next holiday trip, or deciding when and with whom to fill an open job position. All tasks demand to decide whether to keep the current option (partner/ticket price/applicant) or to keep on searching for a better option. When declining an option, it is not known if a better option will eventually present itself. It is also often not possible to go back to one of the previous options, as possible partners might not be available anymore, ticket prices change from day to day, and a once rejected applicant might have started to work for another company.

Formally, the *Optimal Stopping Problem* is the task of finding in a sequence of timepoints $1 \leq i \leq n$ with a possibly unknown n for associated values (or options) $(y_i)_i$ the time i when to perform an action to maximize the desired potential outcome, i.e., increase profits or minimize costs. The options can change randomly and the quality of the future options cannot be estimated. In some cases, the number of options can be limited, e.g., one has to buy a plane ticket eventually if the vacation is planned on a certain date. If the last option is reached, it has to be chosen. The difficulty lies in evaluating if either the currently presented option is worth keeping, given the knowledge about the previously seen options and some domain knowledge (e.g., average plane ticket prices), or if a better option will occur in the future.

Most current models are assessed by a "fitting"-analyses of the response distribution and not on assessing the predictive accuracy of the next decision before an individual makes it (Guan & Lee, 2018; Lee & Wagenmakers, 2014; Seale & Rapoport, 2000; von Helversen & Mata, 2012; Zwick, Rapoport, Lo, & Muthukrishnan, 2003). The advantage of the latter method is that it allows even to falsify models, i.e., if they do not predict the right decisions and it can identify the underlying decision processes. Hence, we propose to assess models in the predictive setting using *CCOBRA* (Brand, Riesterer, & Ragni, 2020), a cognitive reasoning framework that allows predicting and adapting to an individual reasoner while evaluating a models performance. The data used stems from Baumann, Singmann, Gershman, and von Helversen (2020), which presents participants in an experiment with the task to buy an item for the cheapest possible price. The data includes artificially generated prices for fictional items with a normal, left- and right-skewed price distribution, as well as real prices for real items that can be used to evaluate the real-life performance of the predictive models. As part of this work, four models for human reasoning in *Optimal Stopping Problems* were implemented and adapted to individual human reasoning. The original models are presented in Baumann et al. (2020) and are based on a threshold heuristic. For comparison of the performance of the adapted models, two baseline models were implemented, one random one and one that will follow the optimal strategy (Gilbert & Mosteller, 2006) to find the best option. The models were fine-tuned on one of the available data sets (= training data) and then evaluated on the other data sets. The results were then compared to each other and the similarities and differences between them were examined. The findings were also compared to the findings of previous studies, which showed similar results. An insight into the prediction performance of each model on each individual reasoner is also given as the final step to analyze the models.

Related Work

Lee (2006) proposes a hierarchical *Bayesian Model* to predict human decisions. Participants in the experiment were confronted with the problem of choosing the maximum out of a sequence of few numbers. The participants knew the generation principle of the numbers. The results supported a threshold-based model to explain the decisions of the participants, choosing the first maximal number that exceeds a certain threshold for each index. Since the overall performance did not increase during the experiment, participants did not learn from the previous problems.

An optimal strategy (that has the highest expected value) to select the best (in this case highest) option out of a finite list of options is described in Gilbert and Mosteller (2006). The idea is to start with a high threshold and decrease it over time based on the distribution the options were sampled from.

Guan, Lee, and Vandekerckhove (2015) considers the *Optimal Stopping Problem* for a sequence length of both, 5 and 10. In both cases, the reasoner had to find the highest option of the sequence. The authors propose a threshold model for human reasoning that takes into account how far the individual reasoner deviates from the optimal threshold at the current step. The deviation is controlled by two parameters β and γ. β determines how far above or below the threshold is from the optimal one and γ controls how fast the bias increases/decreases as the sequence progresses. Their results show that reasoners that set their initial threshold higher than the optimal, tend to decrease it faster than optimal, and reasoners that set their initial threshold too low, decrease it slower than optimal. Furthermore, β and γ remains stable for participants in the sequence of length 5 and 10. That allows to transfer the observed values in one tasks to this individual for other tasks.

The data used for this research and the models that set the foundation for the adapted models are presented in Baumann et al. (2020). The authors describe four models for predicting human reasoning in an *Optimal Stopping Problem* where the goal is to find (and buy) the cheapest price for an item in a sequence of 10 prices. The four models are the *Independent Threshold Model*, the *Linear Threshold Model*, the *Biased Optimal Model*, and the *Cut-off Model*. All of the models are based on the calculation of an acceptance probability θ_i that implements a sigmoid choice function with sensitivity parameter β and the current item i with price x_i and the position-dependent threshold t_i:

$$\theta_i = \frac{1}{1 + exp\{\beta(x_i - t_i)\}}$$

The goal of each model is to provide the threshold t_i which changes depending on the task and the position in the sequence.

The ***Independent Threshold Model (ITM)*** assumes no dependence between the thresholds, it samples N independent random thresholds (from a uniform distribution) $t_1, ..., t_N$ where at position $N+1$ the price must be accepted.

In the ***Linear Threshold Model (LTM)***, the thresholds are constrained by a linear relation to each other and are defined by the initial threshold t_0 and the linear scaling factor δ:

$$t_{i+1} = t_i + \delta \cdot i$$

The ***Biased Optimal Model (BOM)*** is based on the model presented in Guan et al. (2015). It uses the optimal threshold t_i^*, a systematic bias parameter γ that reflects the divergence from the optimal threshold, and the parameter α which describes how much the threshold decreases or increases as the sequence progresses.

$$t_i = t_i^* + \gamma + \alpha \cdot i$$

The ***Cut-off Model (CoM)*** assumes that the reasoner has a fixed cutoff value k that determines how long the sequence is explored before the first value that is lower than the already seen minimum is accepted.

Benchmark Data

The data used in this project stems from (Baumann et al., 2020). It recreates a scenario in which the decision-maker is planning a vacation and wants to buy the flight tickets online. The prices vary randomly from day to day and the customer wants to find the cheapest ticket. Each day the decision-maker checks the price and can either buy the ticket or reject the offer and wait for the next day. Since the vacation will start in ten days, the decision maker has to accept the tenth offer, no matter the price. Once a price is accepted the search is also finished. A total of three experiments with different price distributions are reported.

For the *first experiment*, 129 participants were set to answer the described ticket-shopping task. The prices were sampled from a normal distribution with a mean of 180 and a standard deviation of 20. Each subject finished 200 trials of the ticket-shopping task. In each trial, the participants searched through a sequence of ten prices. The subjects were aware that they could see up to ten prices and were always informed about the number of remaining prices. However, they could only see the price of the current product. It was not possible to go back to an already rejected price. If the subjects reached the tenth price they were forced to accept it.

For the *second experiment*, 172 participants were in the ticket-shopping task but with changing distributions from which the prices were sampled. Three different sample methods were used: Exp 2a – prices were sampled from a left-skewed PERT distribution PERT(40, 195, 200) with a mean of 170. Exp 2b – a normal PERT distribution PERT(90, 140, 190) with a mean of 140 was used. Exp 2c – a right-skewed PERT distribution (PERT(120, 125, 400)) with a mean of 170 was used. Each participant was assigned to only one experiment and had to select the lowest price out of a sequence of 10 prices for 200 trials.

The *third experiment* simulates an online shopping experience where the goal is to buy a certain product for the lowest price possible with the prices being presented sequentially. A total of 60 commodity products were selected and the prices collected from an online shop. Only products with an approximately normal distributed price range were chosen. For the experiment, the prices were sampled from a normal distribution with the mean and standard deviation estimated from the real prices. All 100 participants performed 120 trials divided into two blocks containing the same sixty products. The subjects were always aware of the number of remaining prices and were also informed about the mean price of the product. Once a price was rejected it could not be chosen again and the tenth price had to be chosen if no previous buy was performed.

All models were evaluated using *CCOBRA* (Brand et al., 2020). CCOBRA is a cognitive reasoning framework that sequentially presents per person data to the model that is currently evaluated. In each step of the evaluation sequence, the model is presented with the current task, in this case the task would be one price of the sequence of the ten prices out of a problem the participant had to face. With the presented task, the models have to predict the answer the current reasoner gave for this task, which is then used to evaluate the models performance. After predicting, CCOBRA provides an adaption function in which not only the task had to be predicted just now, but also the given answer is presented to the model. This information can than be used to adapt the model to the current reasoner. For the evaluation in this paper, the data was prepared in a way that the available information for the task is the price for the current ticket/item and the sequence of the current task (how often a price was rejected in this iteration). The reaction time (how long the individual took to make the decision) and the mean price (180 for experiments 1, 170 for the left- and right-skewed, 140 for the normal task in experiment 2, and variable for the third experiment) are also given as further information.

Adapting Models to Predict the Individual

The previously presented models from Baumann et al. (2020) were adapted to work with variable price means by scaling the individual parameters with the mean of the task. The models were also able to adapt to the individual reasoner by updating the parameters during the prediction process which will be presented in the following sections. A genetic algorithm was used to search for the optimal parameters for every 30 questions asked. The current parameters were then updated with new optimal values by setting the new parameters to be 70% *current parameter* + 30% *searched parameter*. For all models, the previously presented β parameter was initialized to 0.21 and the genetic algorithm searched in the interval of $[0,2]$.

Random Model (RM)

The *Random Model* represents the most simple decision maker by randomly selecting one of the options. A model that can't beat the random baseline would probably be better off just guessing the answer.

Independent Threshold Model (ITM)

The *Independent Threshold Model* samples its ten independent thresholds from a *uniform distribution* between 60% and 120% of the mean value for the current task. During the adaption to the individual reasoner, only β is searched with the genetic algorithm.

Linear Threshold Model (LTM)

The *Linear Threshold Model* starts with an initial threshold t_0, which in this case is a percentage of the mean value of the current task. It is defined as: $t_0 = meanvalue \cdot t_{\%}$. The linear increase δ is also represented as a percentage of the mean value of the current task. It is defined as: $\delta = meanvalue \cdot \delta_{\%}$. The threshold calculation is then done via: $t_{i+1} = t_i + meanvalue \cdot \delta_{\%} \cdot i$. For the basic version and the adaption, the parameters are initialized with $\delta_{\%} = 0.005$ and $t_{\%} = 0.7$. During the search for a better fitting value in the adaption phase, $\delta_{\%}$ was limited to the interval of $[0,0.1]$ and $t_{\%}$ searched between 0% and 100%.

Optimal Threshold Model (OTM)

The *Optimal Threshold Model* uses the mathematical optimal threshold to determine whether to buy for the current price or to wait for the next opportunity. A way to calculate the optimal thresholds to find the highest number in a sequence is described in Gilbert and Mosteller (2006). In order to calculate the thresholds for the lowest number in the sequence, the threshold generation process was inversed. This results in a list of optimal thresholds (percentage of the mean price):

Table 1: Optimal thresholds for each value in the sequence. At 10 there is a 'must buy'. Any option below the threshold is accepted.

Pos.	1	2	3	4	5	6	7	8	9
Opt.	72	78	84	91	99	109	121	137	160

At each option that is to be predicted, the *Optimal Threshold Model* checks whether the current price is below the optimal threshold and predicts the *buy* option. If the current price is above the optimal threshold, it decides not to buy. Since there are no free parameters that can be optimized for this model, there is also no adaption variant for it.

Biased Optimal Model (BOM)

The *Biased Optimal Model* takes the optimal threshold for the current position in the sequence and adds two parameters to it (γ and α). Since those are also dependent on the magnitude of the current prices, they are also represented by a percentage of the optimal threshold t^*. The calculation for the *Biased Optimal Model* is therefore done like the following:

$$t_i = t_i^* + \gamma \cdot t_i^* + \alpha \cdot i \cdot t_i^*$$

Both parameters γ and α are both initialized to 0, which means that without any adaption, the *BOM* is equal to the *OTM*. During the adaption to the individual reasoners, the genetic algorithm searches for the optimal α value in the range of $[-0.2,0.1]$ and for the optimal γ value in the range $[-0.5,0.3]$.

Cut-off Model (CoM)

The *Cut-off Model* explores the sequence a fixed number of steps (k) and then takes the next option that is lower than the previously seen lowest value. During the initialization of the *CoM* k is set to 5. The genetic algorithm for the adaption part of the model searches for k in the range of $[1, 10]$.

Results

All models were tested and tuned on the data of Exp 1. These results were used to improve the models by adapting the parameters for better performance. The data of all other experiments was not used in the training process, and only used in the final evaluation, to avoid overfitting on the data. Experiment 2 and its variant for skewnesses demonstrates the models' power to adapt, given different price distributions. A model that performs well on a left-skewed distribution (more cheap prices in the sequence) might in term perform worse on a right-skewed distribution (more costly prices). The dataset of Exp 3 (real prices) gives insight into the models' ability to adapt to real-life situations. During the evaluation, each model was run five times on each dataset, to account for the randomness of the genetic algorithm (Ritter, Schoelles, Quigley, & Klein, 2011; Byrne, 2013). All later presented results are therefore the mean of five evaluation runs. Overall the results show a good performance for most models on the datasets of the first and second experiment. The best performing models were the *Biased Optimal Model* and the *Linear Threshold Model* both with adaption. The best performing model was the *LTM* with adaption, reaching an accuracy of 88.74%. On the third experiment, however, all models had a significant drop in performance, with the *Cut-off Model* being the best performing one. Overall the *Independent Threshold Model* showed the worst performance.

Experiment 1

The evaluation results for this experiments can be found in Table 2. The clear prediction performance winner is the *Biased Optimal Model* with adaption. It achieved an 86.7% mean accuracy on the prediction. Next up is the *Linear Threshold Model* with adaption which scored an 85.0% accuracy. The *Cut-off Model* with and without adaption as well as the *Linear Threshold Model* without adaption scored at around 81% accuracy. Due to the nature of how the data is presented in the datasets, simply predicting that the current reasoner will not buy for the current price will lead to a high prediction performance (in this case 79.1%). This is because once a reasoner accepted a certain price, the remaining prices for this sequence were skipped. This leads to an over-representation of *don't buy* answers in the dataset. This prediction performance can therefore be seen as the barrier that shows if the model truly learned the reasoning process. As to be expected, the *Biased Optimal Model* without adaption and the *Optimal Threshold Model*, with 73.5% prediction accuracy, share the same performance since the *BOM* without adaption represents the *OTM*. The *Independent Threshold Model* with and without adaption achieved around 68% accuracy with the adaption model even performing slightly worse. The random model scored around 50% accuracy as it is to be expected in an two possible outcomes random choice.

Experiment 2a: left skewed prices

The results for the second experiment's first condition, with a left-skewed distribution (more cheap prices), are presented in Table 2. Interestingly, all models managed to improve their performance in comparison to the first experiment. This time, the *Linear Threshold Model* with adaption with 88.7% accuracy performed slightly better than the *Biased Optimal Model* with adaption that reached 88.7% prediction performance. With 87.2% accuracy, the *LTM* without adaption managed to improve its performance drastically compared to the first experiment.

Experiment 2b: normally distributed prices

In Table 2 the results for the second condition of the second experiment (normal-distributed prices) are shown. Once again, the *BOM* and *LTM*, both with adaption, are the best performing models. This time, like in the first experiment where the prices were also normal-distributed, the *Biased Optimal Model* with adaption performed slightly better, with 88.0% accuracy, than the *Linear Threshold Model* with adaption which reached 86.6% accuracy.

Experiment 2c: right skewed prices

The results for the second experiment's third condition (left-skewed distribution, more expensive prices) are presented in Table 2. The *Biased Optimal Model* with adaption continued with the trend of being one of the strongest models and showed the best performance of all models with 87.3%. However, compared to the other two conditions of the second experiment, the *BOM* with adaption showed the worst performance in this condition. With an accuracy of 86.1%, the *Linear Threshold Model* with adaption also showed slightly worse performance than under the previous two conditions of the second experiment. Nevertheless, the *LTM* with adaption proved to be a solid predictor of the decision-makers in the second experiment. Both, the *Cut-off Model* with and without prediction, showed constantly good results in the entirety of the second experiment and always managed to beat the *never buy* threshold. In this case, they achieved an accuracy of 81.2% and 80.5% respectively. The *BOM* without adaption and the *Optimal Threshold Model* consistently achieved a performance of around 74% accuracy. The trend of both *Independent Threshold Models* being the worst in the portfolio also continued under this condition and with an accuracy of 68.8% both reached a performance as low as ever.

Experiment 3: real prices

In the third experiment, the performance of all models dropped significantly in comparison to the first and second experiment (cp. Table 2). None of the models managed to beat the *don't buy* threshold. Most notably the *Cut-off Model* with and without adaption is now leading the scoreboard with an accuracy of 77.8% and 76.1% respectively. This represents a loss of roughly three percentage points compared to the first and second experiments. The *Linear Threshold Model* with adaption lost roughly thirteen percentage points compared to the previous experiments and dropped

Table 2: The median predictive accuracy of the cognitive models for each experiment and all experiments, and the random and don't buy baseline models. In bold is the highest predictive performance. The median is calculated from five evaluation runs.

	Random	dontbuy	ITM_A	ITM	OTM	BOM	LTM	CoM	CoM_A	LTM_A	BOM_A
Exp. 1	49.8	79.1	67.9	68.1	73.5	73.5	81.0	81.1	81.4	85.0	**86.7**
Exp. 2 (left skew.)	50.0	81.1	69.7	69.8	73.9	73.9	87.2	81.5	82.9	**88.7**	88.5
Exp. 2 (normal)	50.1	79.2	70.0	70.0	74.7	74.7	82.1	81.2	81.7	86.6	**88.0**
Exp. 2 (right skew.)	50.0	78.7	68.8	68.8	74.1	74.1	83.7	80.5	81.2	86.1	**87.2**
Exp. 3	50.0	78.2	64.2	64.1	71.2	71.2	73.1	76.1	**77.8**	74.0	73.8
All	50.0	80.5	68.1	68.2	73.5	73.5	81.4	80.1	81.0	84.1	**84.9**

to an accuracy of 73.9%. Surprisingly, the *LTM* without adaption scored only slightly lower with 73.1% accuracy, showing that in this case, the adaption did not bring any meaningful advantage. Overall the *LTM* without adaption dropped about ten percentage points in performance compared to the previous experiments. A similar observation can be done with the *Biased Optimal Model*. The adaption-based model scored 73.8% accuracy while the non-adapting one scored only slightly lower with 71.2% prediction accuracy. The adaption-based *BOM* loses roughly fourteen percentage points of prediction performance compared to experiments one and two, while the non-adaption model loses only around three percentage points, which is of course the same development as the *Optimal Threshold Model*. Once again, the *Independent Threshold Model*, with and without adaption, showed the worse performance of all models with only around 64% accuracy. Compared to the other experiments this represents a loss of around four to five percentage points.

Individual Prediction Performance

Fig. 1 shows the individual participant performance for all models on the third experiment. The *COM* model and its adaption show no meaningfull difference in the overall distribution of the individual prediction performance other than that the adaption manages to score slightly higher on almost all participants showing a better fit to the individual. The adaption of the *LTM* clearly reduces the overall spread of the individuals, while also the outliers both at the bottom and at the top are clearly reduced by the adaption.

For the *Biased Optimal Models*, the non-adapting model shows a wider spread of the individual performances. However, there are no real outliers to the bottom of the performance so the basic model manages to fit all participants at least to a certain amount. The model with adaption clearly shows how the performance is improving by shrinking the spread between the individuals and bringing them all to higher accuracy. Nevertheless, there are some participants the adaption model does not quite manage to fit and therefore there are also more outliers to the bottom. The *ITM* and its adaption did not show any improvements. It is therefore most likely, that the β value does not yield any performance increase in a setting with random thresholds.

Discussion

The overall bad performance of the *Independent Threshold Model* with and without adaption reinforces the previous findings (Baumann et al., 2020; Guan & Lee, 2018; Guan et al., 2015) that human reasoners change their thresholds following a certain strategy (e.g. linear). A random (bounded) threshold for each step in the sequence would mean that a decision-maker could refuse a cheap price in the current step only to accept a way to high price in the next step. This behavior is not typical for a human reasoner, which is shown by the bad performance of the models. The below-average performance of the *Optimal Threshold Model* and the *Biased Optimal Model* without adaption, which are the same, confirms the findings of Guan et al. (2015) that human reasoners in *Optimal Stopping Problems* tend to set their initial threshold too low or too high, and then either decrease them to slow or increase them to fast. The *Cut-off Model* showed a solid performance and managed to beat the *don't buy threshold* consistently. This shows that the reasoners tend to at least somewhat try to explore the sequence rather than making a hasty decision. This effect has already been mentioned in Baumann et al. (2020). The outstanding performance of the *Linear Threshold Model* with and without adaption once again strengthens the assumption that human reasoners tend to use linear threshold in *Optimal Stopping Problems*. The good performance of the *LTM* shows that most reasoners, at least to a certain extend, choose a linear threshold to guide their decision. However, some outliers might apply a different technique and an in-depth evaluation of those outliers could help to further increase the ability to predict the individual reasoner. The *Biased Optimal Model* with adaption shows a similar prediction accuracy to the *LTM* with adaption. This behavior is to some extent understandable since the optimal thresholds for the first ten options in the sequence are roughly linear as well. Additionally, the parameters for γ and α found by the genetic algorithm during the adaption were mostly negative meaning they also counteracted the nonlinear effect of the optimal threshold.

Both the *BOM* and *LTM* show a strong performance over all experiments (cp. Table 2). Only in the third experiment, both did not outperform the *don't buy* baseline. There might

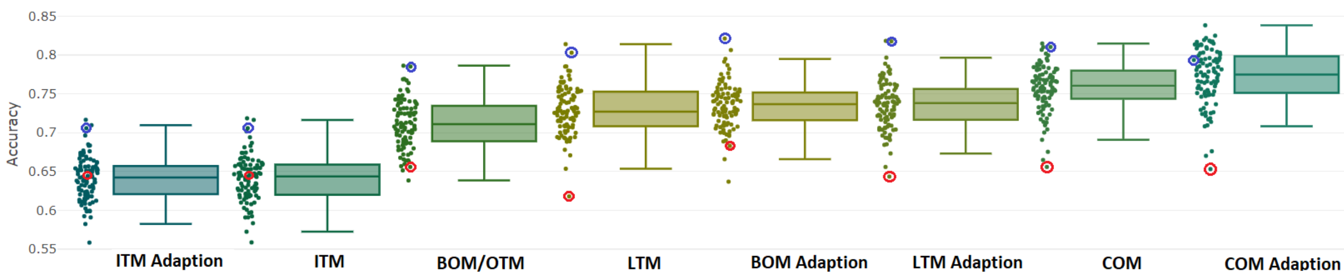

Figure 1: Prediction results of the models on each participant in the third experiment. Random model omitted for space reasons, BOM/OTM are the same with no adaption. Red and blue circle mark the same participants respectively.

be several reasons for this drop in performance: One reason for the change in performance could lie in the participants themselves. The fact that they had to deal with real items that they know and have an understanding of the price developments, could have nudged a behavior change compared to the fictional price setting of the first and second experiment. Another reason might be that the reasoners just did not behave as predictable as in the other tasks. Since the prices varied quite a lot in their magnitude, a reasoner that previously accepted a costly product at a 20% discount could not accept the same 20% discount for a cheaper item since the absolute discount of the cheap item does not appear as high.

Consider two types of participants in Fig. 1 and 2, depicted by a red and blue circle accordingly. The blue participant is predicted adequately by almost all models in contrast to the low predictive accuracy for the red participant across models. This indicates that either the red-marked participant gave more random answers and is less predictable by any meaningful models or the reasoner developed a strategy that is beyond what the implemented models can cover. The well-predicted blue reasoner, however, seems to use a strategy that is covered by almost all models.

In conclusion, this paper adapted and evaluated core decision making models for *Optimal Stopping Problems* to predict decisions performed by individuals. Using genetic algorithms allowed the models to find the optimal parameters for each individual. The findings support previous research that showed that human reasoners tend to use a linear threshold in *Optimal Stopping Problems* to rate the current option. The analysis shows that thresholds are variable among decision makers and that adapting to the individual can bring a vast improvement in the prediction capabilities of the models.

Acknowledgements

Support to MR by the Danish Institute of Advanced Studies and the DFG (RA 1934/4-1, RA1934/8-1) is gratefully acknowledged.

References

Baumann, C., Singmann, H., Gershman, S. J., & von Helversen, B. (2020). A linear threshold model for optimal stopping behavior. *Proceedings of the National Academy of Sciences*, *117*(23), 12750–12755.

Brand, D., Riesterer, N., & Ragni, M. (2020). *Ccobra-framework.* `https://github.com/CognitiveComputationLab/ccobra`. GitHub.

Byrne, M. D. (2013). How many times should a stochastic model be run? an approach based on confidence intervals. In *Proceedings of the 12th International Conference on Cognitive Modeling, Ottawa.*

Gilbert, J. P., & Mosteller, F. (2006). Recognizing the maximum of a sequence. In *Selected Papers of Frederick Mosteller* (pp. 355–398). Springer.

Guan, M., & Lee, M. D. (2018). The effect of goals and environments on human performance in optimal stopping problems. *Decision*, *5*(4), 339.

Guan, M., Lee, M. D., & Vandekerckhove, J. (2015). A hierarchical cognitive threshold model of human decision making on different length optimal stopping problems. In *CogSci.*

Lee, M. D. (2006). A hierarchical bayesian model of human decision-making on an optimal stopping problem. *Cognitive Science*, *30*(3), 1–26.

Lee, M. D., & Wagenmakers, E.-J. (2014). *Bayesian cognitive modeling: A practical course*. Cambridge University Press.

Ritter, F. E., Schoelles, M. J., Quigley, K. S., & Klein, L. C. (2011). Determining the number of simulation runs: Treating simulations as theories by not sampling their behavior. In *Human-in-the-loop simulations* (pp. 97–116). Springer.

Seale, D. A., & Rapoport, A. (2000). Optimal stopping behavior with relative ranks: The secretary problem with unknown population size. *Journal of Behavioral Decision Making*, *13*(4), 391–411.

von Helversen, B., & Mata, R. (2012). Losing a dime with a satisfied mind: Positive affect predicts less search in sequential decision making. *Psychology and Aging*, *27*(4), 825.

Zwick, R., Rapoport, A., Lo, A. K. C., & Muthukrishnan, A. (2003). Consumer sequential search: Not enough or too much? *Marketing Science*, *22*(4), 503–519.