Title
Making Predictions Without Data: How an Instance-Based Learning Model Predicts Sequential Decisions in the Balloon Analog Risk Task

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors
Bugbee, Erin H.
Gonzalez, Cleotilde

Publication Date
2022

Peer reviewed

Making Predictions Without Data: How an Instance-Based Learning Model Predicts Sequential Decisions in the Balloon Analog Risk Task

Erin H. Bugbee (ebugbee@cmu.edu)
Department of Social and Decision Sciences, Carnegie Mellon University
Pittsburgh, PA 15213 USA

Cleotilde Gonzalez (coty@cmu.edu)
Department of Social and Decision Sciences, Carnegie Mellon University
Pittsburgh, PA 15213 USA

Abstract

Many models in Cognitive Science require data to calibrate parameters. Some modelers calibrate their models' parameters for each individual in a data set, and others work at the aggregate level. Generally, the accuracy of a model is judged by the degree to which human data are replicated, and the model parameters are interpreted accordingly. It is not too surprising that models that are developed for a particular task and fit to each individual's data in such a task replicate the human data well. The question is, however, whether those models can make predictions in the absence of human data. In this paper, we present a theory-driven model of a well-known sequential decision task (the Balloon Analog Risk Task, BART) which is able to make predictions in the absence of human data. The cognitive model is grounded on the processes and mechanisms of Instance-Based Learning (IBL) Theory of experiential choice. We demonstrate the simulation predictions from an IBL model and those of a well-known model of the BART, which depends on the fits to human data. We further show that when making predictions without data, the IBL model provides predictions that are both theoretically founded and accurate, while the Two-Parameter model performs much worse than when fit to data. We conclude with a discussion of the benefits of making theory-based predictions in the absence of human data for our community.

Keywords: sequential decision making, instance-based learning, balloon analog risk task

Introduction

With the exponential increase in data sources, the availability of large volumes of data, and the possibility to collect data on the Web from a large number of individuals, statistical models have increased their potential to make accurate predictions of human choices. From movie recommendations to self-driving cars, models have become part of our daily decisions; with more data, models can make more accurate predictions of our preferences and choices. In particular, computational cognitive science has taken advantage of the availability of larger amounts of behavioral data to advance the "explanation" of cognitive processes involved in various types of tasks, notably decision making using Bayesian approaches (Griffiths, 2015).

However, beyond fitting models, computational cognitive science has been largely interested in a deeper theoretical understanding of human learning processes and the predictions made from the learning histories of populations or particular individuals. Some approaches are based on probability theory and their statistical analyses have been used for many decades to study many aspects of human cognition, including language and memory (Chater, Oaksford, et al., 2008) and decision making (Lee, 2006; Guan, Lee, & Vandekerckhove, 2015; Thomas, Coon, Westfall, & Lee, 2021). Although most of these models require large human data sets, they can also provide significant insights into human mental decision processes and inform the development of more powerful computational tools. However, it is unclear whether statistical models and Bayesian cognitive models of decision making in particular, have predictive and explanatory power (Hofman et al., 2021; Shmueli, 2010).

For the advancement of cognitive science, it is generally important not simply to make accurate predictions, but also to provide an explanation and understanding of how and why people behave the way they do. The development of computational cognitive models based on cognitive theories is expected to provide prediction power without a strong dependence on data (Hofman et al., 2021). A cognitive theory is a general postulation of mechanisms and processes that are globally applicable to families of tasks and types of activity rather than being dependent on a particular task; while cognitive models are very specific representations of parts or all aspects of a cognitive theory that apply to a particular task or activity (Gonzalez, 2017). Cognitive models of sequential decisions can be used to simulate the interaction of theoretical cognitive processes with a particular task environment. These models can make predictions of how human choices are made in such tasks that can be compared to data collected from human participants in the same tasks, without fitting to human data (Roberts & Pashler, 2000, 2002). Thus, an advantage of theoretically driven cognitive models is that no data is required to fit the models.

Using a theory-driven approach to build computational models is especially necessary in sequential decision making tasks, such as the Optimal Stopping and the Balloon Analog (Analogue) Risk Task (BART) (Lejuez et al., 2002), because in these tasks, each decision is contingent on the specific sequence of past decisions (Guan, Stokes, Vandekerckhove, & Lee, 2020; Thomas et al., 2021). Researchers of sequential decision making tasks have developed particular models that are applicable only to a particular sequential decision task (Guan et al., 2020; Guan, 2019). The combinatorial complexity of particular sequences of decisions would make these data-driven approaches difficult to predict the performance "out of sample" in novel data sets involving new individu-

In J. Culbertson, A. Perfors, H. Rabagliati & V. Ramenzoni (Eds.), *Proceedings of the 44th Annual Conference of the Cognitive Science Society*.

als and problems. Meanwhile, theory-driven models are expected to have predictive and explanatory power (Shmueli, 2010), with the ability to make a priori predictions and generalize in situations (Busemeyer & Wang, 2000).

Instance-Based Learning Theory (IBLT) is a general cognitive theory of experience-based decisions that emerged from the need to explain the sequential decision process of dynamic decision making, where interdependent decisions are made in a sequence and over time (Gonzalez, Lerch, & Lebiere, 2003). IBLT provides a single general algorithm and mathematical formulations of memory retrieval that rely on the ACT-R cognitive architecture (Anderson & Lebiere, 2014). An IBL model uses IBLT in a particular task, where the memory "instances" represent the state, action, and utilities in the context of the task being modeled. Many models based on IBLT have been developed in a wide diversity of contexts and domains, from highly specialized and complex tasks such as cyber defense (Aggarwal et al., 2020; E. A. Cranford et al., 2020) and anti-phishing detection (E. Cranford, Lebiere, Rajivan, Aggarwal, & Gonzalez, 2019); where tasks need multidimensional representations and real-time interactivity (Nguyen, Phan, & Gonzalez, 2021); to simple and abstract binary choice dynamics (Gonzalez & Dutt, 2011; Lejarraga, Dutt, & Gonzalez, 2012; Konstantinidis, Harman, & Gonzalez, 2020). Importantly, all IBL models use the process and mechanisms of IBLT.

In this paper, we present an IBL model that makes theoretically-based predictions in a sequential decision task (BART) in the absence of data. We compare the IBL predictions to human data and to the predictions of the Two-Parameter model, which was built for the BART task and fits its parameters to each individual in the same human data set by using Bayesian inference. The Two-Parameter model data fits are an excellent representation of the data; and the IBL model captures the behavioral variability in the data set by making predictions in the absence of data. We then compare the models for their predictive power by running both models with standard parameter distributions, not fit to human data. The results show that the Two-Parameter model cannot generalize as well as the theory-driven, IBL model.

Balloon Analog Risk Task

The Balloon Analog Risk Task (BART) is a well-known task used to measure human risk-taking behavior. The task was introduced by Lejuez et al. (2002) to measure risk-taking and compare with other measures of risky behavior. Since its development, BART has been used to study risky decision making in many domains and contexts, such as adolescent decision making (Lejuez, Aklin, Zvolensky, & Pedulla, 2003; Aklin, Lejuez, Zvolensky, Kahler, & Gwadz, 2005), decision making of smokers versus non-smokers (Lejuez, Aklin, Jones, et al., 2003), impulsivity (Reynolds, Ortengren, Richards, & de Wit, 2006), and decision making under stress (Lighthall et al., 2011), to name a few.

BART is a sequential decision making task in which a decision maker inflates a balloon, and the level of inflation corresponds to a monetary value. At each time point, the decision maker decides whether to pump the balloon and increase its value or bank the current monetary amount. However, with each pump of the balloon, there is a probability that the balloon bursts and the decision maker receives a reward of 0. This leads to the need to balance exploration through pumping with exploitation through banking, with the goal of maximizing total reward. Each balloon has a predefined probability of bursting, although participants are not informed about this probability (Guan et al., 2020). Because BART is a sequential risk-taking task, the riskiness of the decision increases with each pump decision, and the number of pumps an individual makes represents the level of risk the participant is willing to take. A risk-seeking individual would likely pump the balloon more times than a risk-averse individual.

To analyze and predict risk propensities in BART, many researchers have proposed computational models that represent risky behaviors. The *Two-Parameter BART* model is perhaps the most well-known model of the BART task and has been used successfully in a large number of research studies (Pleskac, Wallsten, Wang, & Lejuez, 2008; van Ravenzwaaij, Dutilh, & Wagenmakers, 2011; Guan et al., 2020).

Two-Parameter BART Model

In essence, the model depends on an individual propensity to take risks (γ_{ik}) and the belief in the probability of bursting of the balloon (p_k), for the individual i and the condition k. The p_k is often assumed to be the true probability of a balloon burst. The p_k and γ_{ik} determine the target number of pumps to make (ω_{ik}). This model assumes that participants have a unique target number of pumps and that they apply the same target number of pumps to all problems. That is, the model assumes no learning[1]; instead, the model assumes that the probability of pumping in a particular trial of a problem depends directly on ω_{ik} and an additional parameter representing the behavioral consistency of the individual (β_{ik}) which is determined by a logistic function. These parameters are inferred using Bayesian inference. The model's equations have been presented and explained in many past studies (e.g. Thomas et al. (2021)), which we repeat here for completeness:

The Two-Parameter model assumes that each participant has a target number of pumps, denoted ω_{ik} for the participant i and the condition k. The target number of pumps is calculated according to Equation 1, where γ_{ik}^{+} is participant i's propensity for risk-taking in condition k and $\gamma_{ik}^{+} \geq 0$.

$$\omega_{ik} = \frac{-\gamma_{ik}^{+}}{\ln(1-p_k)} \tag{1}$$

The target number of pumps and a behavioral consistency

[1] Recently, Zhou, Myung, and Pitt (2021) introduced the Scaled Target Learning (STL) model as a learning model for BART. The STL model describes learning as the adjustment of the number of pumps in response to previous outcomes, the adjustments reflecting the sensitivity to wins and losses.

parameter β_{ik} determine the probability that the participant i in problem j in condition k will pump the balloon in the trial t, which is indicated by θ_{ijkt} and calculated according to Equation 2. Higher values of β correspond to greater consistency in decision making and $\beta_{ik} \geq 0$.

$$\theta_{ijkt} = \frac{1}{1+\exp(\beta_{ik}(t-\omega_{ik}))} \tag{2}$$

The Two-Parameter model makes a number of assumptions regarding information and decision elements for human participants in BART studies. As explained above, the true burst probability is assumed in the calculation of a target number of pumps. This target number is assumed to be determined by each individual and applied to every problem; some function of behavioral consistency is assumed; and importantly, it is assumed that there is no learning, such that the number of pumps is never adjusted as a result of experience. Actually, learning by human participants is often prevented in the design of the human studies used to demonstrate the Two-Parameter BART model. For example, the data set obtained from Guan et al. (2020) and used in Thomas et al. (2021) was obtained in a study in which individuals face problems with different actual bursting trials in every problem, and where problems are presented in random order across participants. This is essentially done because the Two-Parameter BART model is not a learning model, and it would be unable to capture patterns of learning over time.

We argue that at least some of these assumptions are behaviorally unreasonable. Participants in the BART task are not told the probability of the balloon bursting (Guan et al., 2020; Thomas et al., 2021), there is no evidence that they have a number of pumps in mind, but there is evidence indicating that the number of pumps is different in different problems (Guan et al., 2020). Learning across problems is prevented by experimental design (Guan et al., 2020). The actual bursting time is pre-determined based on the probability, it is set to be different in each problem, and problems are randomized. However, it seems unreasonable to assume that participants would not learn from their experience and adjust the number of pumps.

In addition, the Two-Parameter BART model needs human data to establish individual-level parameters that represent each individual's risk propensity. Although researchers have used the calibrated BART model to make inferences "out of sample" (e.g., Thomas et al. (2021)), the model depends on human data to determine the parameter values.

We offer a cognitive model for the BART task that makes none of these assumptions. The model makes predictions based on the theoretical principles of Instance-Based Learning Theory (IBLT) (Gonzalez et al., 2003), and in the absence of human data.

Instance-Based Learning Theory

IBLT's process and mechanisms have been published in many past publications (see a recent publication from Nguyen, Phan, and Gonzalez (2021)); we repeat the theory here for completeness.

In IBLT, an "instance" is a memory unit that results from the potential alternatives evaluated. These memory representations consist of three elements that are constructed over time: a situation state s which is composed of a set of features f; a decision or action a taken corresponding to an alternative in state s; and an expected utility or experienced outcome x of the action taken in a state. Concretely, for an IBL agent, an option $k = (s,a)$ is defined by action a in state s. At time t, assume that there are n_{kt} different considered instances (k_i, x_{ik_it}) for $i = 1,...,n_{kt}$, associated with k. Each instance i in memory has an *Activation* value, which represents how readily available this information is in memory and is expressed as follows (Anderson & Lebiere, 2014):

$$\Lambda_{ik_it} = \ln\left(\sum_{t' \in T_{ik_it}} (t-t')^{-d}\right) + \alpha \sum_j Sim_j(f_j^k, f_j^{k_i}) + \sigma \ln \frac{1-\xi_{ik_it}}{\xi_{ik_it}} \tag{3}$$

where d, α, and σ are the decay, mismatch penalty, and noise parameters, respectively, and $T_{ik_it} \subset \{0,...,t-1\}$ is the set of previous timestamps in which the instance i was observed, f_j^k is the j-th attribute of the state s, and Sim_j is a similarity function associated with the j-th attribute. The second term is a partial matching process that reflects the similarity between the current state s and the state of the option k_i. The rightmost term represents noise to capture individual variation in activation, and ξ_{ik_it} is a random number drawn from a uniform distribution $U(0,1)$ at each timestep and for each instance and option.

Activation of an instance i is used to determine the probability of retrieval of an instance from memory. The probability of an instance i is defined by a soft-max function:

$$P_{ik_it} = \frac{e^{\Lambda_{ik_it}/\tau}}{\sum_{j=1}^{n_{kt}} e^{\Lambda_{jk_jt}/\tau}} \tag{4}$$

where τ is the Boltzmann constant (i.e., the "temperature") in the Boltzmann distribution. For simplicity, τ is defined as a function of the σ used in the activation equation $\tau = \sigma\sqrt{2}$.

The expected utility of option k is calculated based on *Blending* as specified in choice tasks (Lejarraga et al., 2012; Gonzalez & Dutt, 2011):

$$V_{kt} = \sum_{i=1}^{n_{kt}} P_{ik_it} x_{ik_it} \tag{5}$$

The choice rule is to select the option that corresponds to the maximum blended value. In particular, at the l-th step of an episode, the agent selects the option (s_l, a_l) with

$$a_l = \arg\max_{a \in A} V_{(s_l,a),t} \tag{6}$$

When the agent receives delayed results, the agent updates expected utilities using a credit assignment mechanism (Nguyen, McDonald, & Gonzalez, 2021).

IBL BART Model

We built an IBL model of the BART using PyIBL, a Python implementation of IBLT (Morrison & Gonzalez, 2021). In this task, the state s is the number of pumps up to the current trial, the action a is the action to pump or bank, and the utility x is the amount of hypothetical money obtained from the problem.

An instance is stored for each decision: for each pump that does not burst the balloon, the utility is the reward that would be obtained from banking on the next trial; for each bank decision, the reward is the monetary amount collected; and for each pump that results in the balloon bursting, the utility is \$0 since no money is collected. In this model, we use linear similarity between the current instance and past instances and the decay and noise parameters are set to the ACT-R default values of $d = 0.5$ and $\sigma = 0.25$ respectively.

Data Set

We retrieved and analyzed the data from Guan et al. (2020), in which 56 human participants completed the BART. The participants were presented with balloons that had a fixed probability of bursting with each pump (either $p = 0.1$ or $p = 0.2$).[2] The design was within-subjects, so each participant completed 50 problems with each probability, and the order of the problems and conditions was randomized across participants. Each problem started with a balloon with a value of \$1. For each decision, the participant had the option to pump the balloon ("Pump") and increase its monetary value by \$1, or stop ("Bank") and collect the current monetary value. However, each pumping action risks bursting the balloon, resulting in collecting \$0. The participant continued making decisions until either the balloon burst or the participant chose the Bank action and collected the money. The goal was to maximize the total reward in all problems.

Results

Explaining decisions through simulations

We simulated 56 IBL model agents making decisions in the BART using default parameter values for decay and noise. The stimuli faced by each agent correspond to that of a human participant. That is, each agent experienced the problems in the same order as the corresponding human participant, so that the balloon bursts on the same number of pumps.

Following Guan et al. (2020) and Thomas et al. (2021) and by modifying their provided scripts, we used the Two-Parameter BART model developed by Pleskac et al. (2008). We first fit the Two-Parameter model to both the $p = 0.1$ and $p = 0.2$ conditions, and then simulated choices using the inferred parameters for each participant for that condition. The simulated problems are the same problems that the corresponding individual experienced in each condition. Here, we compare the model predictions to the choices of the human participants to determine the accuracy of the predictions relative to the observed human choices.

The distributions of the **number of pumps** from the IBL and Two-Parameter models are displayed in Figure 1 compared to human data. A greater number of pumps tends to correspond to a greater risk propensity.

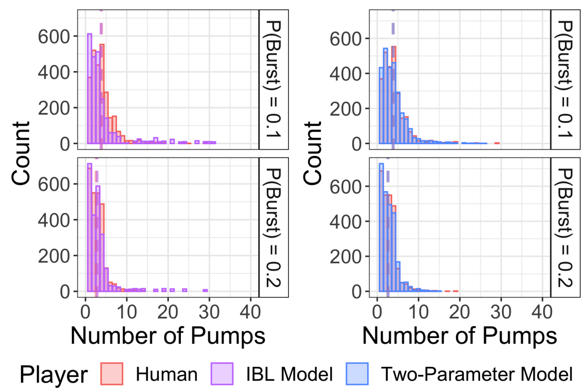

Figure 1: The distributions of the number of pumps of the balloon for all participants and problems, for the humans, IBL model, and Two-Parameter model. The dashed lines indicate the mean number of pumps.

T-tests show that for $p = 0.1$, the mean number of pumps is not significantly different when comparing the human ($M = 3.86, SD = 3.07$) and IBL model ($M = 3.77, SD = 5.64$), $t(5598) = 0.72, p = 0.469; d = 0.019$, and when comparing the human and Two-Parameter model ($M = 3.85, SD = 3.00$), $t(5598) = 0.15, p = 0.878; d = 0.004$. For $p = 0.2$, the mean for humans ($M = 2.58, SD = 1.95$) is significantly lower than the IBL model ($M = 2.83, SD = 3.62$), $t(5598) = -3.24, p = 0.001; d = 0.087$. There is no difference between the means of humans and the Two-Parameter model ($M = 2.61, SD = 1.97$), $t(5598) = -0.65, p = 0.517; d = 0.017$.

Figure 2 divides the problems into blocks of 10 problems in the order in which they were presented to the participants. This shows that the IBL model becomes increasingly better at predicting human behavior with more experience. The Two-Parameter model fits the human data well, as expected given the fitting to each individual's choices, but it does not learn.

The **burst rate** is the proportion of problems for which the participant bursts the balloon. Participants with higher risk propensities will likely tend to burst the balloon on a higher proportion of problems.

Figure 3 shows the distributions of the burst rates. In the $p = 0.1$ condition, the t-tests show that the difference in the mean burst rates for humans ($M = 0.40, SD = 0.16$) and the IBL model ($M = 0.37, SD = 0.20$) is not significant $t(110) = 0.65, p = 0.516; d = 0.123$. The same holds for the humans and Two-Parameter model ($M = 0.39, SD = 0.16$) comparison, $t(110) = 0.29, p = 0.772, d = 0.055$. In the $p = 0.2$ condition, the t-tests show that the difference for humans ($M = 0.46, SD = 0.17$) and the IBL model ($M = 0.48, SD = 0.20$)

[2]The constant probability of bursting is not the standard BART design. Typically, the probability of bursting increases with each pump of the balloon (Lejuez et al., 2002).

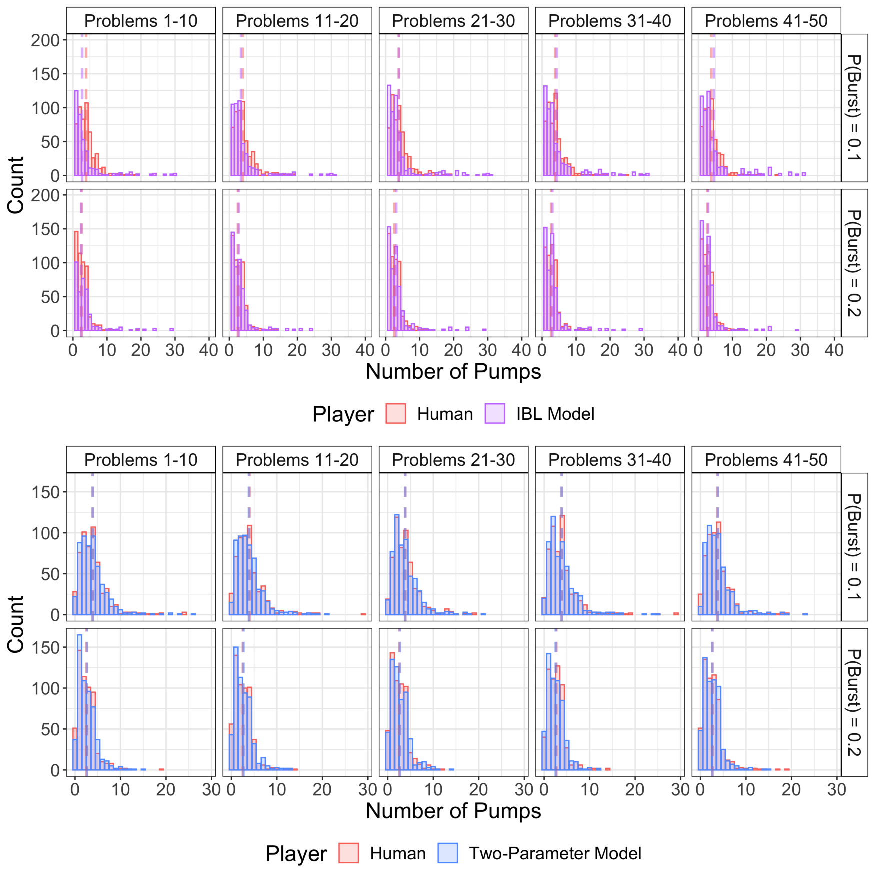

Figure 2: The distributions of the number of pumps of the balloon for all participants by block of 10 problems, for the humans, IBL model, and Two-Parameter model. The IBL model uses default parameter values and the Two-Parameter model parameters are fit for each individual and condition. The dashed lines indicate the mean number of pumps.

is not significant, $t(110) = -0.61, p = 0.545; d = 0.115$; the same is true for the humans and the Two-Parameter model ($M = 0.46, SD = 0.16$), $t(110) = -0.09, p = 0.927; d = 0.017$. Together, these results show that both the IBL and Two-Parameter model predict human burst rates accurately; however, the IBL model can do this without fitting to human data and relying only on IBLT.

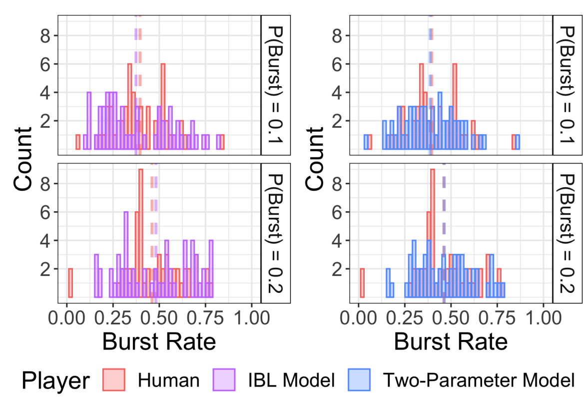

Figure 3: The distributions of the burst rates for the humans, IBL model, and Two-Parameter model. The dashed lines indicate the mean burst rate.

Evaluation of predictions through simulations

To evaluate the IBL and Two-Parameter models on equivalent grounds, we generated predictions from the Two-Parameter model in this task, sampling values for β and γ for each individual from a truncated normal distribution with a lower bound of 0, centered at the mean fitted value of each parameter with a standard deviation of the standard deviation of that parameter's fitted values. We used these sampled values to simulate choices in the task; the parameters are not directly fit to the choices of the human participants, yet they are a reasonable representation of the parameter distribution. For the IBL model, we sampled values for the decay and noise parameters from truncated normal distributions centered on their default value with a lower bound of 0. We used the standard deviation of the fitted β parameter from the Two-Parameter model for each condition, to give similar variability in parameters for both models. These simulations allow us to evaluate the Two-Parameter model's ability to generalize and make predictions without parameter fitting.

The distributions of the **number of pumps** are displayed in Figure 4. T-tests show that for $p = 0.1$, the mean number of pumps is significantly lower for the IBL model ($M = 3.32, SD = 5.00$) relative to humans ($M = 3.86, SD = 3.07$), $t(5598) = 4.91, p < 0.001; d = 0.131$. The mean number of pumps is significantly higher for the Two-Parameter model ($M = 5.96, SD = 5.33$) relative to humans, $t(5598) = -18.05, p < 0.001; d = 0.482$. For $p = 0.2$, the mean for humans ($M = 2.58, SD = 1.95$) is not significantly different than for the IBL model ($M = 2.62, SD = 3.40$), $t(5598) = -0.58, p = 0.562; d = 0.015$. In contrast, the mean for the Two-Parameter model ($M = 3.70, SD = 3.02$) is much higher than for humans, $t(5598) = -16.49, p < 0.001; d = 0.441$.

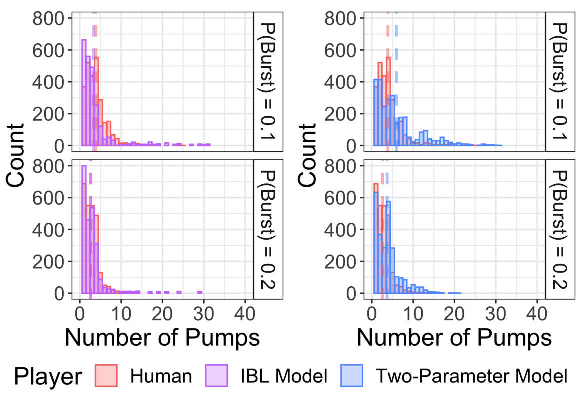

Figure 4: The distributions of the number of pumps of the balloon for all participants and problems, for the humans and the IBL model and Two-Parameter model with sampled parameters.

Figure 5 shows the distributions of **burst rates**. For the $p = 0.1$ condition, t-tests show that the difference in mean burst rates for humans ($M = 0.40, SD = 0.16$) and the IBL model ($M = 0.34, SD = 0.19$) is not significant, $t(110) = 1.62, p = 0.108; d = 0.306$. The difference is significant for humans and the comparison of the Two-Parameter model ($M = 0.55, SD = 0.24$), $t(110) = -4.12, p < 0.001; d = 0.779$, with the Two-Parameter model having a significantly higher burst rate compared to human participants, indicating that the Two-Parameter model predicts a higher risk propensity compared to that actually observed in human data. In the $p = 0.2$ condition, the t-tests show that the difference for the humans ($M = 0.46, SD = 0.17$) and the IBL model ($M = 0.44, SD = 0.19$) is not significant, $t(110) = 0.60, p = 0.552; d = 0.113$. The difference is significant for the comparison of human and Two-Parameter model ($M = 0.65, SD = 0.22$), $t(110) = -5.20, p < 0.001; d = 0.982$. The Two-Parameter model predicts a significantly higher burst rate compared to the human participants. Without fitting to participant choices, the Two-Parameter model becomes less accurate.

Discussion

To be able to provide explanations and accurate "out-of-sample" predictions of human decisions, particularly in sequential decision making tasks, it is important to rely on theory-driven models. Models that rely on human data to fit their parameters to particular human data sets can provide interpretable parameters (e.g., regarding risk propensity). However, the same model immediately becomes less accurate and significantly different from the human data (e.g., burst rate) when making predictions, even on the same data set. The IBL model, in contrast, makes theory-based predictions that are statistically indistinguishable from human data, in the complete absence of human data.

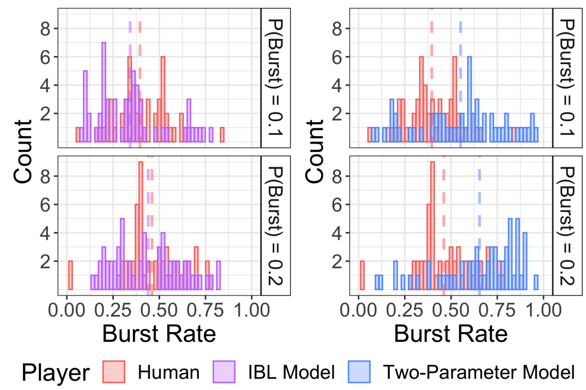

Figure 5: The distributions of the burst rates for the humans and the IBL model and Two-Parameter model with sampled parameters.

We presented and demonstrated how a theory-driven cognitive model of the BART task can make accurate predictions in the absence of human data. The IBL model uses the theoretical principles of IBLT, which have been used to model many other decision making tasks across contexts. The predictions emerge from the simulation of the IBL algorithm using default generic parameters. The model stores only information that is available to human participants and makes no unreasonable assumptions, such as knowledge of the bursting probability or assuming that participants have a predefined and consistent number of pumps to make in the task.

By evaluating the performance of the models "out-of-sample," we aim to point out that often a model may work well for a fixed data set to which the parameters are fitted, but the same model may fail to generalize to settings that involve even the same tasks and individuals. In fact, different models have been constructed for related sequential decision tasks, suggesting that the models are only applicable to a particular sequential decision task and can generalize poorly (Guan et al., 2020). In future work, we seek to investigate how the same IBL model of sequential decisions can be generalized across other tasks. This is a feasible avenue for future research, given the currently available data sets in which the same individuals completed four different sequential decision tasks (Guan et al., 2020), and given the theory-driven approach to building computational models of sequential decision making tasks that we are pursuing.

Acknowledgments

An Open Science Framework project is available at https://osf.io/cty4q with data, code, and analysis files. This research was supported by AFRL Award FA8650-20-F-6212 sub-award number 1990692. We are grateful to our colleagues, Christian Lebiere, Leslie Blaha, Chase McDonald, Erin McCormick, and Alexander Hough, for exchanging conversations related to this project. We thank Michael Lee for assistance with the BART data set from Guan et al. (2020) and for helpful feedback.

References

Aggarwal, P., Thakoor, O., Mate, A., Tambe, M., Cranford, E. A., Lebiere, C., & Gonzalez, C. (2020). An exploratory study of a masking strategy of cyberdeception using cybervan. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *64*(1), 446-450. doi: 10.1177/1071181320641100

Aklin, W. M., Lejuez, C., Zvolensky, M. J., Kahler, C. W., & Gwadz, M. (2005). Evaluation of behavioral measures of risk taking propensity with inner city adolescents. *Behaviour research and therapy*, *43*(2), 215–228.

Anderson, J. R., & Lebiere, C. J. (2014). *The atomic components of thought*. Psychology Press.

Busemeyer, J. R., & Wang, Y.-M. (2000). Model comparisons and model selections based on generalization criterion methodology. *Journal of Mathematical Psychology*, *44*(1), 171–189. doi: 10.1006/jmps.1999.1282

Chater, N., Oaksford, M., et al. (2008). *The probabilistic mind: Prospects for bayesian cognitive science*. Oxford University Press, USA.

Cranford, E., Lebiere, C., Rajivan, P., Aggarwal, P., & Gonzalez, C. (2019). Modeling cognitive dynamics in end-user response to phishing emails. *Proceedings of the 17th International Conference on Cognitive Modeling*, 35-40.

Cranford, E. A., Gonzalez, C., Aggarwal, P., Cooney, S., Tambe, M., & Lebiere, C. (2020). Toward personalized deceptive signaling for cyber defense using cognitive models. *Topics in Cognitive Science*, *12*(3), 992–1011.

Gonzalez, C. (2017). Decision making: a cognitive science perspective. *The Oxford handbook of cognitive science*, *1*, 1–27.

Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating decisions from experience in sampling and repeated choice paradigms. *Psychological Review*, *118*(4), 523–51.

Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, *27*. doi: 10.1016/S0364-0213(03)00031-4

Griffiths, T. L. (2015). Manifesto for a new (computational) cognitive revolution. *Cognition*, *135*, 21–23.

Guan, M. (2019). *A cognitive modeling analysis of risk in sequential choice tasks*. Unpublished doctoral dissertation, UC Irvine. Retrieved from https://escholarship.org/uc/item/802684nb

Guan, M., Lee, M. D., & Vandekerckhove, J. (2015). A hierarchical cognitive threshold model of human decision making on different length optimal stopping problems. In *Proceedings of the 37th annual meeting of the cognitive science society.*

Guan, M., Stokes, R., Vandekerckhove, J., & Lee, M. D. (2020). A cognitive modeling analysis of risk in sequential choice tasks. *Judgment and Decision Making*, *15*, 823-850.

Hofman, J. M., Watts, D. J., Athey, S., Garip, F., Griffiths, T. L., Kleinberg, J., ... others (2021). Integrating explanation and prediction in computational social science. *Nature*, *595*(7866), 181–188.

Konstantinidis, E., Harman, J. L., & Gonzalez, C. (2020). Memory patterns for choice adaptation in dynamic environments.

Lee, M. D. (2006). A hierarchical bayesian model of human decision-making on an optimal stopping problem. *Cognitive Science*, *30*, 1-26. doi: 10.1207/s15516709cog0000$_6$9

Lejarraga, T., Dutt, V., & Gonzalez, C. (2012). Instance-based learning: A general model of repeated binary choice. *Journal of Behavioral Decision Making*, *25*(2), 143–153.

Lejuez, C., Aklin, W. M., Jones, H. A., Richards, J. B., Strong, D. R., Kahler, C. W., & Read, J. P. (2003). The balloon analogue risk task (bart) differentiates smokers and nonsmokers. *Experimental and clinical psychopharmacology*, *11*(1), 26.

Lejuez, C., Aklin, W. M., Zvolensky, M. J., & Pedulla, C. M. (2003). Evaluation of the balloon analogue risk task (bart) as a predictor of adolescent real-world risk-taking behaviours. *Journal of Adolescence*, *26*(4), 475-479. doi: https://doi.org/10.1016/S0140-1971(03)00036-8

Lejuez, C., Read, J., Kahler, C., Richards, J., Ramsey, S., Stuart, G., ... Brown, R. (2002). Evaluation of a behavioral measure of risk taking: The balloon analogue risk task (bart). *Journal of experimental psychology. Applied*, *8*, 75-84. doi: 10.1037/1076-898X.8.2.75

Lighthall, N. R., Sakaki, M., Vasunilashorn, S., Nga, L., Somayajula, S., Chen, E. Y., ... Mather, M. (2011). Gender differences in reward-related decision processing under stress. *Social Cognitive and Affective Neuroscience*, *7*(4), 476-484.

Morrison, D., & Gonzalez, C. (2021). *Pyibl version 4.0.* Retrieved 2021-03-30, from http://pyibl.ddmlab.com/

Nguyen, T. N., McDonald, C., & Gonzalez, C. (2021). *Credit assignment: Challenges and opportunities in developing human-like ai agents* (Tech. Rep.). Carnegie Mellon University.

Nguyen, T. N., Phan, D. N., & Gonzalez, C. (2021). Speedyibl: A solution to the curse of exponential growth in instance-based learning models of decisions from experience. *arXiv preprint arXiv:2111.10268*.

Pleskac, T. J., Wallsten, T. S., Wang, P., & Lejuez, C. (2008). Development of an automatic response mode to improve the clinical utility of sequential risk-taking tasks. *Experi-*

mental and clinical psychopharmacology, *16*(6), 555.

Reynolds, B., Ortengren, A., Richards, J. B., & de Wit, H. (2006). Dimensions of impulsive behavior: Personality and behavioral measures. *Personality and Individual Differences*, *40*(2), 305-315. doi: https://doi.org/10.1016/j.paid.2005.03.024

Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? a comment on theory testing. *Psychological Review*, *107*(2), 358–367. doi: 10.1037/0033-295X.107.2.358

Roberts, S., & Pashler, H. (2002). Reply to rodgers and rowe (2002). *Psychological Review*, *109*(3), 605–607. doi: 10.1037/0033-295X.109.3.605

Shmueli, G. (2010). To explain or to predict? *Statistical Science*, *25*(3), 289–310.

Thomas, B., Coon, J., Westfall, H. A., & Lee, M. D. (2021). Model-based wisdom of the crowd for sequential decision-making tasks. *Cognitive Science*, *45*. doi: 10.1111/cogs.13011

van Ravenzwaaij, D., Dutilh, G., & Wagenmakers, E.-J. (2011). Cognitive model decomposition of the bart: Assessment and application. *Journal of Mathematical Psychology*, *55*(1), 94-105. (Special Issue on Hierarchical Bayesian Models) doi: https://doi.org/10.1016/j.jmp.2010.08.010

Zhou, R., Myung, J. I., & Pitt, M. A. (2021). The scaled target learning model: Revisiting learning in the balloon analogue risk task. *Cognitive Psychology*, *128*, 101407.