A STUDY ON INFORMATION-EFFICIENT INFERENCE IN HUMAN DECISION-MAKING

By

Kristen Li

An Undergraduate Thesis submitted as part of the

WHARTON RESEARCH SCHOLARS

Faculty Advisors:

Jacob Parker, Ph.D. Candidate in Neuroscience

Dr. Joseph Kable, Professor of Psychology & Marketing

Dr. Joshua Gold, Professor of Neuroscience

THE WHARTON SCHOOL, UNIVERSITY OF PENNSYLVANIA

JANUARY 2025

ABSTRACT

Performance on inference tasks varies considerably across individuals and task conditions. One possible explanation to this variation in performance is the use of strategies that differ in the amount of cognitive resources expended. In this paper, we apply the information bottleneck (IB) framework to study how performance (inference predictiveness) relates to the amount of information used (inference capacity). The IB framework computes an upper bound on inference predictiveness as a function of inference capacity, and distinguishes between information-efficient inference on the bound and information-inefficient inference off the bound. We ran four versions of a classic inference task online to examine both between-subject and within-subject variation in inference capacity and predictiveness under different task conditions. We found that most participants remained information-efficient across task conditions and that classic manipulations of decision-making performance can reflect shifts along the IB bound. These results suggest that people are sensitive to the cost of using information during inference, and that individual variability in decision-making performance can be attributed to the efficient use of limited information.

Keywords: decision-making, inference, bounded rationality, information bottleneck

CONTENTS

1. INTRODUCTION

Many approaches to studying human decision-making assume that the underlying process is inherently rational, but this assumption has been challenged through work that shows consistent deviations from optimality (Kahneman & Tversky, 1979; Tversky & Kahneman, 1974; Wason, 1968). These suboptimalities are attributed to a limited reservoir of cognitive resources in an idea that is commonly known as bounded rationality (Simon, 1955). It has therefore been proposed that researchers should instead examine how humans might adapt to optimize decisions under these limitations (Lieder & Griffiths, 2020). This study of bounded rationality raises questions related to the nature of limited cognitive resources, the mechanisms through which the brain flexibly allocates these resources under different conditions, and the effects of these limitations on performance in various decision-making tasks.

Along this line of inquiry, we propose that in human decision-making, a relevant limitation is the amount of information considered. The more information considered to make a decision, the more cognitive resources that one would presumably expend in the process. Past studies have shown that decision-makers pursue different strategies that may differ in the amount of information used (Eissa et al., 2022; Filipowicz, Glaze, et al., 2020; Glaze et al., 2018; Tavoni et al., 2022). However, considering more information could, in principle, lead to more accurate decisions, and vice versa. Therefore, we posit that people face a tradeoff between minimizing the amount of information considered, which we refer to as inference capacity, and maximizing the accuracy of their decisions, which we refer to as inference predictiveness.

This tradeoff between inference predictiveness and inference capacity has been formalized under the information bottleneck (IB) framework (Tishby et al., 2000). This framework

examines inference predictiveness in relation to inference capacity and allows us to compute an information bottleneck (IB) bound, which is an upper bound on inference predictiveness as a function of inference capacity. We can therefore use this framework to distinguish between decision-makers on the IB bound who are information-efficient and achieve maximal or near maximal level of inference predictiveness given their inference capacity, and decision-makers off the IB bound who are information-inefficient and do not. As such, the IB framework delineates a measure of optimality under limitations in resource capacity.

The IB framework has been applied in several domains to assess information efficiency. For example, it has been used to show that the retina efficiently extracts features of a stimulus (i.e., position and velocity) to maximize the amount of predictive information encoded in sensory processing (Palmer et al., 2015). It has similarly been used to show that artificial neural networks extract the maximum amount of information about the signal it has been trained to predict given how much they compress the input (Shwartz-Ziv & Tishby, 2017).

We have been applying the same IB framework to examine the information efficiency of human decision-making via inference tasks, where people infer some latent state based on available evidence. Notably, in a task that involved predicting the color of beads drawn probabilistically from one of two hidden source jars that switched randomly, we observed variations in inference predictiveness between individuals that we can attribute to variations in inference capacity. In other words, while most participants failed to maximize absolute accuracy on the task, the majority of participants fell near the IB bound, indicating optimal performance given the amount of information used (Filipowicz, Forkin, et al., 2020). This result suggested that people may be sensitive to the cost of using information and that suboptimal decision-making performance can reflect optimal use of limited information.

In this paper, we aim to build on the findings of the above beads task to: 1) determine whether the relationship between inference capacity and inference predictiveness observed generalizes to other kinds of inference tasks; 2) explore whether there are conditions that lead

to shifts in inference capacity within individuals along with a corresponding shift in inference predictiveness; and 3) identify conditions under which individuals would systematically fail to make efficient use of a given amount of information and fall off the IB bound.

2. METHODS

2.1. The Information Bottleneck (IB)[1]

In the original paper on the Information Bottleneck (IB), Tishby et al. (2000) describes a signal $x \in X$ that contains information about another signal $y \in Y$. We can compress the signal X while minimizing the amount of information lost about Y by finding a lossy representation $r \in R$ of X that contains the maximum information about Y. Assuming that it is costly to increase the amount of information R retains about X, we define the following optimization problem:

$$\min_{p(r|x)} I(X;R) - \beta I(R;Y), \tag{2.1}$$

where $I(X;R)$ is the mutual information between X and R, and $I(R;Y)$ is the mutual information between R and Y. Equation (2.1) therefore captures a trade-off between minimizing $I(X;R)$, the cost of retaining information from X, and $I(R;Y)$, the amount of information that R retains about Y, in lossy compression.

We can apply the IB to study human inference by defining X to be some observation probabilistically associated with some latent state Y. The human decision-maker makes an inference R about the latent state Y that encodes some information about the observation X. We can consider $I(X;R)$ as a proxy for inference capacity, or how much information the decision maker considers, and $I(R;Y)$ as a proxy for inference predictiveness, or how accurate the decision-maker is about the latent state. We can therefore interpret equation (2.1) to represent the trade-off between minimizing the cost associated with inference and maximizing the accuracy of inference. Note that one advantage of the IB is that it is model-agnostic, meaning that it does not make assumptions about the strategies that a given decision-maker uses.

[1]This section summarizes work done by Jacob Parker.

We observe that the trade-off in equation (2.1) is governed by the parameter β, which controls the relative importance of the two objectives in the optimization problem. Larger values of β place greater emphasis on preserving the predictive information about Y, enabling more accurate inferences that require greater cognitive effort or resources. Conversely, smaller values of β prioritize compression over maintaining the relevance of R to Y, reducing the mental expenditure required at the expense of less accurate inferences about the underlying latent state.

Under this setup, the choice behavior of the decision-maker on an inference task is described by $p(r|x)$, which reflects the probability of making choice $r \in R$ given the observation $x \in X$. For an information-efficient decision-maker (a decision-maker that satisfies equation (2.1) for some value of β), the choice behavior follows the conditional probability distribution $p^*(r|x)$[2]

$$p^*(r|x) = \frac{e^{-\beta D_{KL}[p(y|x)||p^*(y|r)] + \log p^*(r)}}{\sum_r e^{-\beta D_{KL}[p(y|x)||p^*(y|r)] + \log p^*(r)}}, \tag{2.2}$$

where β is the same β from equation (2.1), $D_{KL}[p(y|x)||p^*(y|r)]$ is the KL divergence between $p(y|x)$ and $p(y|r)$, and $p^*(y|r)$ and $p^*(r)$ are probability distributions derived from $p^*(r|x)$ in the following manner:

$$p^*(y|r) = \sum_x p(y|x)p^*(x|r) \tag{2.3}$$

$$p^*(r) = \sum_x p(r|x)p^*(x) \tag{2.4}$$

Assuming the Markov chain condition $Y \leftarrow X \leftarrow R$ and given joint distribution $p(x, y)$, we can find $p^*(r|x)$ for a given value of β by sequentially iterating through equations (2.2) - (2.4). We can compute the empirical IB bound (fig. 2.1) by doing this over a range of $\beta \in [0, \infty)$ with the EMBO package described in section 2.6.

When $p(x, y)$ contains certain symmetries common within decision-making tasks, we hypothesize that information-efficient inference along the IB bound corresponds to optimal

[2]See Tishby et al. (2000) for full derivation.

Bayesian inference corrupted by different degrees of softmax noise (Parker, 2023). In other words, the choice behavior of the information-efficient decision-maker $p^*(r|x)$ can be modeled as a softmax transformation of the posterior probability distribution of Y given observation X:

$$p^*(r|x) = \frac{e^{\beta\alpha^* p(y=r|x)}}{\sum_r e^{\beta\alpha^* p(y=r|x)}}, \tag{2.5}$$

where the same β parameter controls the degree of softmax noise, and $\alpha^* = \log \frac{p^*(y=r)}{p^*(y\neq r)}$ measures the contrast between how accurate versus how inaccurate a decision-maker can expect to be. Low values of β correspond to noisier strategies that fall on the bottom left portion of the IB curve, whereas high values of β correspond to less noisy strategies that fall on the upper right portion of the IB curve. As β approaches infinity, we arrive at the fully optimal solution with no noise at the upper end of the IB curve (see Figure 2.1).

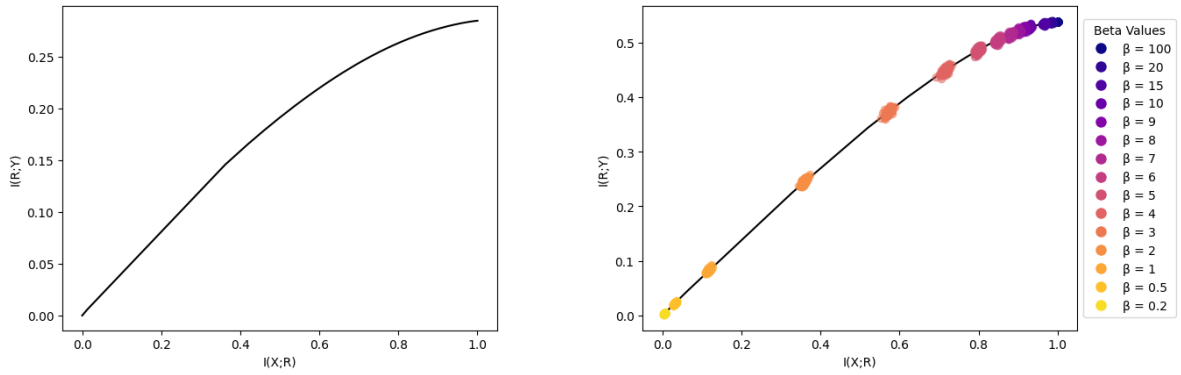

Figure 2.1: empirical IB bound constructed from the EMBO package (left) and with simulated choice behavior of the optimal decision-maker corrupted by different degrees of softmax noise over 10,000 trials of an inference task (right).

The idea that information compression leads to softmax-distributed choice variability has been previously proposed in the theory of rational inattention, where a similar correspondence is found in decision-makers who maximize utility with limited information (Denti et al., 2020; Maćkowiak et al., 2023; Matějka & McKay, 2015). This correspondence is illustrated by simulations depicted in fig. 2.1, where the choice behavior of an optimal agent characterized by different values of β across a large number of trials falls almost exactly along the IB bound.

2.2. The Weather Prediction Task

The weather prediction task is a classic inference task. In the original design, patients with amnesia and control subjects were shown a series of four shapes probabilistically associated with one of two outcomes (represented as a "sunny" and a "rainy" state, hence the name "weather prediction"); by comparing the performance between the two groups, the researchers showed that learning could occur in the absence of declarative knowledge (Knowlton et al., 1994, 1996). A number of subsequent studies used a similar experimental design to explore the inference strategies that people adopted, and to study procedural memory in patients with different neurodevelopmental and mental disorders (Ashby & Maddox, 2005; Bochud-Fragnière et al., 2022; Bochud-Fragnière et al., 2023; Gabay et al., 2015; Gluck et al., 2002; Kelmendi et al., 2016).

A variation of this weather prediction task has also been used in monkeys to study the neural basis of decision-making. In particular, recordings from neurons in the lateral intraparietal cortex (LIP) demonstrated the integration of probabilistic information from a sequence of shapes that favor one of two latent states (Gold & Shadlen, 2007; Kira et al., 2015; Yang & Shadlen, 2007). The design of these weather prediction studies lends itself to our questions of interest because unlike the beads task, evidence is presented within each trial rather than across multiple trials, providing us with more freedom to manipulate the information content of the evidence presented and interleave trials with different task conditions.

2.3. Hypotheses

Using the weather prediction task, we first conjecture that there exists variation in inference capacity between individuals, and relate this variation in inference capacity to any variation in inference predictiveness. We seek to replicate the findings in the beads task that most individuals fall along the IB bound, achieving maximal or near maximal levels of

inference predictiveness for a given level of inference capacity, and are therefore information-efficient.

Likewise, we hypothesize that information-efficient strategies (i.e., those that lead to performance along or close to the IB bound) are equivalent to optimal inference corrupted by softmax noise, as proposed in section 2.1. Furthermore, we theorize that under certain task conditions, the use of information-inefficient heuristics causes decision-makers to perform biased inference and therefore fall off the IB bound.

In addition to between-subject variation, we are also interested in exploring within-subject variation in choice behavior. In particular, we explore whether participants adjust the amount of information used (i.e., inference capacity) in response to changing task conditions. We hypothesize that participants show an increase in inference capacity that corresponds to an increase in inference predictiveness while maintaining a consistent level of information efficiency when 1) previously exposed to the task, 2) when viewing times are increased, and 3) when monetary incentives are awarded.

To study choice behavior across individuals and under different task conditions, we designed four versions of the weather prediction task: 1) a learning task that involves repeated exposure to the same task, 2) a speed-accuracy task that imposes different time constraints, 3) a heuristics task that distinguishes different inference strategies, and 4) a reward task that offers different incentive levels for accurate inferences.

2.4. Participants

60 participants were recruited via Prolific for each of the tasks based on a power analysis, given our hypothesis that inference capacity and predictiveness would increase from the baseline to the manipulation condition in the learning task, speed-accuracy task, and the reward task. Assuming a true within-subject change of 0.05 bits in inference capacity, we simulated choice behavior of N participants with inference capacity uniformly distributed

on the interval $[0, 0.9]$ in the baseline condition. We found that $N = 60$ provided $> 80\%$ power to detect a significant change in both inference capacity and inference predictiveness at $p = 0.05$ level using a paired-sample Wilcoxon signed-rank test described in section 2.6.

We screened for participants who self-reported to be between 18–65 years old, were fluent in English, and of US nationality. Participants must also have a Prolific task approval rating of 99–100%, have completed at least 20 tasks on Prolific, and have not completed any of the other tasks or pilots of these tasks. During data collection, we excluded participants who did not complete the entire experiment or provided an invalid completion code. We also excluded participants whose inference capacity fell below 0.05 bits (equivalent to approximately 58% accuracy) in all conditions during data analysis. Participants were paid a base amount of $9 for an approximate mean completion time of 45 minutes, and were able to earn a bonus amount in each of the tasks as specified below.

2.5. Task Structure

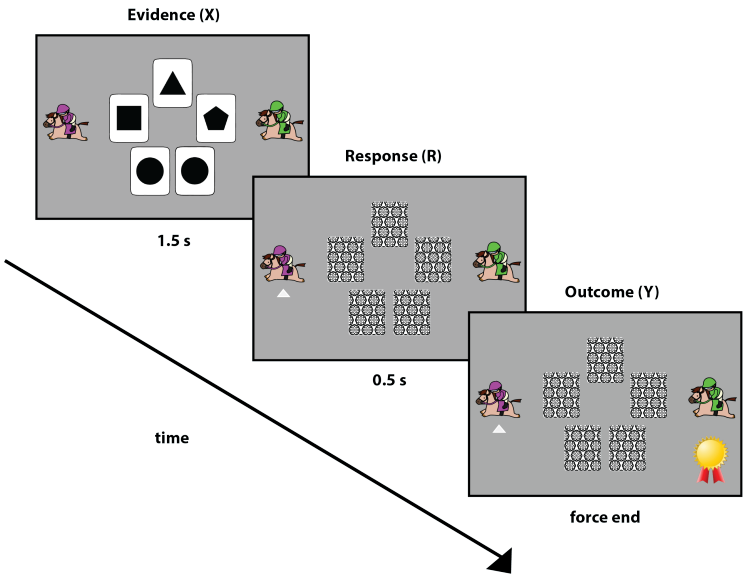

Figure 2.2: Trial structure with a baseline viewing time of 1.5 seconds

For all of the tasks, the participant was simultaneously presented with five shapes, each probabilistically associated with one of two equally likely latent states, over a series of trials. The participant was asked to indicate which of the two latent states the presented shapes

were generated from, via a keypress that must occur within a limited viewing time period before the trial elapsed. The two latent states were represented in the experiment as two race horses of different colors on either side of the screen, analogous to the two weather states in the original design proposed by Knowlton et al. (1994). The position (left/right) and the color (green/purple) of the two latent states were randomized for each participant and remained the same throughout the experiment.

Each participant first completed 100 training trials before proceeding to 100 pretesting trials and 600 testing trials. The training and the pretesting trials were designed to familiarize participants with the task, minimizing learning effects and stabilizing their strategies before the start of testing trials, from which we performed procedures outlined in section 2.6. While the cards remained on the screen until the participants responded in the training trials, the cards were turned over after a prespecified period of viewing time in the pretesting trials as in the testing trials. Upon the cards turning over, the participant must indicate a decision within 0.5 seconds. The actual outcome of the trial was only shown if a response was entered within this time period, and any missed trials were repeated up to three more times.

Each of the five shapes shown on a trial was drawn with replacement from a set of 4 unique shapes with a probability distribution governed by the latent state on that trial. Arranged in order, the log-likelihood ratios of these shapes were $[-0.5, -0.2, 0.2, 0.5]$ for one latent state (with the signs flipped for the other latent state) for the baseline distribution used in the learning task, the speed-accuracy task, and the reward task, so that each of the two latent states was associated with one strong-evidence shape and one weak-evidence shape. Note that negative values of log-likelihood ratios provide evidence against a given state and positive values of log-likelihood ratios provide evidence for a given state. A different distribution was used for each version of the heuristics task, as described below.

Each probability distribution can be uniquely characterized by a set of log-likelihood ratios and a value p_1, the conditional probability of one shape given a latent state. We

conducted a parameter search to select a probability distribution that: 1) maximized distinguishability between information-efficient and -inefficient strategies (see "equal weights" and "ignore weak" heuristics described in the heuristics task), 2) provided stable measures of inference capacity and predictiveness, and 3) corresponded to a maximum attainable accuracy between approximately 75% to 90%. See Appendix A for more details on distribution generation.

The probability distribution was used to generate the trial sequence matched with the relative frequencies of occurrence of each shape combination under each of the two equally likely latent states. For a sequence of 300 testing trials, there were 150 trials associated with one latent state and 150 trials associated with the other latent state randomly interleaved together. Out of the 150 trials associated with each latent state, the number of trials with each shape combination is determined based on their probability of occurrence. For example, if a given shape combination had an 8% probability of being generated for a latent state, there would be 12 trials with that shape combination.

2.5.1. Learning Task

The learning task was designed to detect any change in choice behavior with repeated exposure to the same trial sequence. In the task, 60 participants completed 100 pretesting trials and 600 testing trials. The 600 testing trials were broken up into 2 sequential blocks of 300 identical trials that were randomized in order and in the positions of the shapes shown on the screen. All trials were generated according to the baseline probability distribution with log-likelihood ratios of [-0.5, -0.2, 0.2, 0.5] and $p_1 = 0.08$. The shapes were shown for 1.5 seconds before being covered. Across all pretesting and testing trials, participants earned \$0.0012 for each correct choice and lost \$0.0012 for each incorrect choice. The total accumulated bonus within each block was not allowed to drop below \$0 at any point.

2.5.2. Speed-Accuracy Task

The speed-accuracy task was designed to detect any change in choice behavior with increased viewing time of the evidence, which consisted of the 5 shapes shown on each trial. The task consisted of a fast condition where the shapes were shown for 0.75 seconds before being covered, and a slow condition where the shapes were shown for 1.75 seconds before being covered. 60 participants first completed 100 pretesting trials and then 600 testing trials (both comprising an equal number of fast and slow trials), which were presented as one contiguous block. The fast and the slow trials were interleaved throughout the block, and were randomized in order and in the positions of the shapes shown on the screen. All trials were generated according to the baseline probability distribution with log-likelihood ratios of [-0.5, -0.2, 0.2, 0.5] and $p_1 = 0.08$. Across all pretesting and testing trials, participants earned \$0.0012 for each correct choice and lost \$0.0012 for each incorrect choice. The total accumulated bonus within each block was not allowed to drop below \$0 at any point.

2.5.3. Reward Task

The reward task was designed to detect any change in choice behavior with increased monetary incentive to accurate inferences, or inferences that matched the underlying latent state on a given trial. In the reward task, 60 participants first completed a block of 100 pretesting trials at a baseline reward rate, where they earned \$0.0012 for each correct choice and lost \$0.0012 for each incorrect choice. Then, 30 participants completed a block of 300 testing trials with no reward followed by a block of 300 testing trials with high reward, while the other 30 participants completed a block of 300 testing trials with high reward followed by a block of 300 testing trials with no reward. The stakes in the high reward condition were double the stakes of the baseline reward condition, where participants earned \$0.0024 for each correct choice and lost \$0.0024 for each incorrect choice.

The two blocks of 300 identical testing trials for the high reward and the no reward

conditions were randomized in order and in the positions of the shapes shown on the screen. All trials were generated according to the baseline probability distribution with log-likelihood ratios of $[-0.5, -0.2, 0.2, 0.5]$ and $p_1 = 0.08$. The total accumulated bonus within each block was not allowed to drop below \$0 at any point. The shapes were shown for 1.5 seconds before being covered in both the high reward and no reward conditions.

2.5.4. Heuristics Task

The heuristics task was designed to distinguish between information-efficient and inefficient inference strategies. In the task, 60 participants completed 100 pretesting trials followed by 600 testing trials, which were presented as one contiguous block. Shapes were shown for 1.5 seconds before being covered. Across all pretesting and testing trials, participants earned \$0.0012 for each correct choice and lost \$0.0012 for each incorrect choice. The total accumulated bonus within each block was not allowed to drop below \$0 at any point.

30 of the 60 participants completed the task with trials generated from a distribution that favored an "equal weights" heuristic with log-likelihood ratios of $[-0.585, -0.45, 0.45, 0.585]$ and $p_1 = 0.06$. An "equal weights" heuristic involved weighing the strong-evidence and the weak-evidence shapes the same for each of the two latent states. Under this generative distribution, the IB curve computed for the "equal weights" heuristic was almost indistinguishable from the optimal IB bound, but was distinguishable from the "ignore weak" heuristic described below.

The other 30 participants completed the task with trials generated from a distribution that favored an "ignore weak" heuristic with log-likelihood ratios of $[-1.134, -0.18, 0.18, 1.134]$ and $p_1 = 0.02$. An "ignore weak" heuristic involved using only the strong-evidence shapes when they appeared and considering the weak-evidence shapes only when there was no strong-evidence shape on a given trial. Under this generative distribution, the IB bound computed from the "ignore weak" heuristic was almost indistinguishable from the optimal

IB bound, but was distinguishable from the "equal weights" heuristic described above.

2.6. Data Analyses

We used the EMBO (Empirical Bottleneck) Python package to analyze the choice behavior of participants in the weather prediction task (Piasini et al., 2021; Strouse & Schwab, 2016; Tishby et al., 2000).[3] First, we defined the evidence X to be the sequence of 5 shapes shown on a given trial. Given that shapes were drawn with replacement and the order of shapes did not matter in the inference process, there were a total of 56 unique shape combinations. The response R corresponded to the binary choice the participants made about the latent state. The state Y corresponded to the actual latent state on each trial.

After X, R, and Y were assigned, we empirically computed the IB bound with the EMBO package for each version of the weather prediction task, where the x-axis is the mutual information between X and R, or inference capacity as noted above. The y-axis is the mutual information between R and Y, or inference predictiveness as noted above. Furthermore, the upper bound was also computed using the EMBO package, which calculated the stochastic mapping R of X that encoded the most information about Y, subject to a constraint on the information that R was allowed to retain about X for different degrees of choice noise.

For all versions of the weather prediction task, we computed inference capacity and predictiveness separately for each participant based on their responses across the testing trials, with correction techniques to account for the finite nature of the data. We then plotted these measures to examine where each participant fell relative to the optimal IB bound. Across all the tasks, we broadly interpreted being on or near the bound as an indication of using an information-efficient strategy, akin to performing optimal inference with varying degrees of choice noise. On the other hand, we broadly interpreted being far off the bound as an indication of using a biased strategy that involved systematic errors in

[3]See the package documentation: https://pypi.org/project/embo/

the inference process (in addition to any noise).

For the learning task, the speed-accuracy task, and the reward task, we computed each participant's inference capacity, inference predictiveness, and distance from the IB bound in each condition (block 1 vs block 2, fast vs slow viewing time, and high vs no reward). The distance from IB bound was computed as the ratio between the observed inference predictiveness and the inference predictiveness of the IB bound interpolated from the observed inference capacity. We removed participants whose inference capacity was negative in any condition and replaced negative values of inference predictiveness with zeroes. We performed a paired-sample Wilcoxon signed-rank test to detect any significant difference in these measures between the conditions at $p = 0.05$ and concluded whether the manipulation in question led participants to change inference capacity and predictiveness while remaining information-efficient.

We also concatenated the two blocks in the learning task to construct individual bootstrap intervals for inference predictiveness. If an interval overlapped with the optimal IB bound, we would consider the participant to be information-efficient, and vice versa. Taking into account all 600 trials as opposed to 300 trials in each block produced shorter confidence intervals that allowed us to better distinguish between optimal and suboptimal choice behavior on the task. In addition, we also compared the individual confidence intervals for inference capacity to that of the group based on the combined 600 trials of the learning task as a measure of between-subject variation.

Furthermore, to establish the effect of reward on participants' choice behavior in the reward task, we fit a linear mixed effects model for inference capacity and for inference predictiveness across all 60 participants. Of these 60 participants, 30 completed the high reward condition block first and 30 completed the no reward condition block first, so this approach was necessary to disambiguate a reward effect from a block order effect. Both linear mixed-effects models included main effects of block order and reward magnitude, the

interaction between the two, and a random effect of participants (modeled as random intercepts). We assessed the statistical significance of reward magnitude separately for inference capacity and predictiveness using Satterthwaite's degrees of freedom method as implemented in lmerTest.[4]

For the heuristics task, we computed the inference capacity and inference predictiveness of each participant across the 600 testing trials. To determine whether a given participant was on the optimal IB bound or the inefficient heuristic curve, we constructed a 95% confidence interval for their inference predictiveness via bootstrapping. The confidence interval quantified how much the measure would be expected to vary for a given participant if they were to complete the same experiment many times (with the variability coming from the fact that we were observing probabilistic choices from a participant over a finite number of trials). If a participant's bootstrap interval overlapped with the optimal IB bound or the inefficient heuristic curve, then we concluded that their strategy on the task was indistinguishable from that of the corresponding curve. Given the greater number of trials and the separation between the optimal and the inefficient heuristic bound in this task, we could effectively distinguish between optimal behavior and behavior that was consistent with the inefficient heuristic.

Lastly, we compared the model-agnostic IB-based method of distinguishing between information-efficient and -inefficient strategies in inference with a more traditional, model-based approach. In particular, we fit the choice behavior of each participant in the heuristics task to three distinct strategies: the optimal strategy, the "equal weights" strategy, and the "ignore weak" strategy (see section 2.5.4). For each of these strategies, we fixed the log-likelihood ratio and varied β as a free parameter. Maximum likelihood estimation (MLE) was used to determine the best fitting model and the corresponding β for each participant. We then compared results of this model-fitting to measures based on the IB in section 3.4

[4]See package documentation: https://cran.r-project.org/web/packages/lmerTest/lmerTest.pdf

to establish a correspondence between capacity-limited information-efficient strategies and noisy optimal inference. See Appendix B for more details.

3. RESULTS

3.1. Between-Subject Variation in Inference Capacity Is Larger than Expected Within-Subject Variation

First, we observed variability in inference capacity between individuals based on choice behavior of participants across both blocks of the learning task. We computed 95% bootstrap interval for inference capacity across 60 participants and found that all of them fell within the 95% confidence interval of the group. We observed similar results in the speed-accuracy task, the reward task, and the heuristics task. This indicated that measurement variability for inference capacity across participants was larger than what can be explained by the variability in any given participant, suggesting that individuals adopted strategies that vary in the level of inference capacity on the weather prediction task.

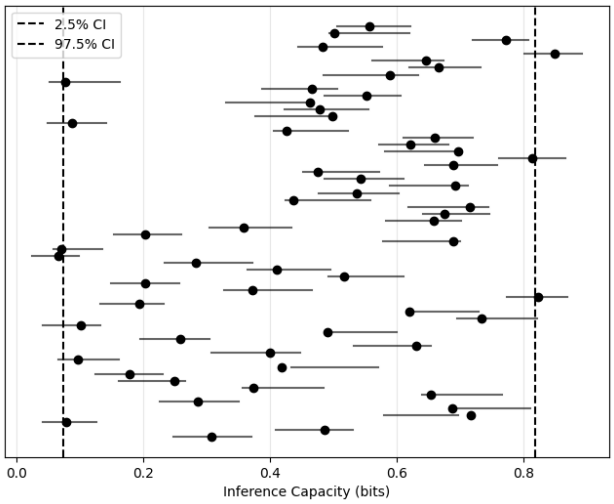

Figure 3.1: 95% bootstrap confidence interval of individual inference capacity across both conditions of the learning task relative to the 95% confidence interval of the group

3.2. Between-Subject Variation in Inference Predictiveness Reflects an Efficient Information Bottleneck in Human Inference

Furthermore, we found that between-subject variation in inference predictiveness can largely be explained by between-subject variation in inference capacity described in section 3.1, based on choice behavior of 60 participants in the learning task and 30 participants in the "equal weights" version of the heuristics task. This was evident in the overlap between the individual bootstrap interval for inference predictiveness and the optimal IB bound for most participants in those task conditions, which allowed us to classify these participants to be on or near the optimal IB bound and therefore information-efficient.

In particular, across the two blocks of the learning task, we observed that the 95% confidence interval of 53 out of 60 participants overlapped with the optimal IB bound. Likewise, across the 30 participants in the "equal weights" version of the heuristics task, the 95% confidence interval of 29 participants overlapped with the optimal IB bound (if we removed the 2 participants whose confidence interval also overlapped with the heuristic bound, there were still 27 participants considered to be information-efficient). These results demonstrated that the variability in performance can be attributed to optimal use of varying amounts of information, illustrating an efficient information bottleneck in human inference.

3.3. Distinguishing Between Information -Efficient and -Inefficient Strategies

We hypothesized that information-efficient strategies along the IB bound correspond to Bayesian-optimal inference corrupted by various degrees of choice noise (see section 2.1). We further hypothesized that noisy biased strategies would be characterized as information-inefficient strategies off the IB bound. We found evidence of such a correspondence in the "ignore-weak" version of the heuristics task, where a considerable number of participants fell off the bound, in contrast to the "equal-weights" version discussed in section 3.2 where most remained on the bound.

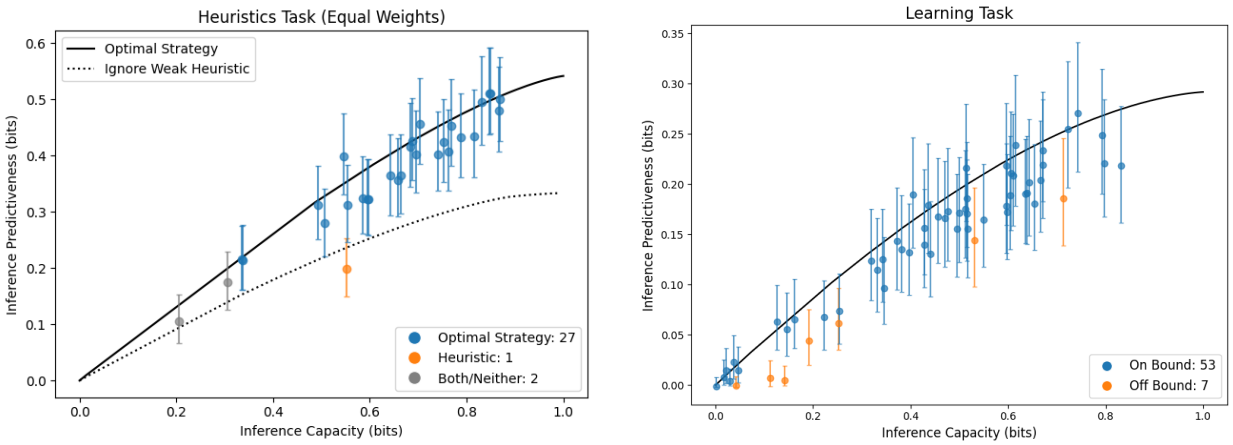

Figure 3.2: 95% bootstrap confidence interval of inference predictiveness for 30 participants on the "equal weights" heuristics task (left) and 95% bootstrap confidence interval of individual inference predictiveness across both conditions of the learning task (right)

In particular, of the 30 participants in the "ignore-weak" version of the heuristics task, the 95% bootstrap confidence interval of 15 participants overlapped with the optimal IB bound, which was almost indistinguishable from the "ignore weak" heuristic bound, while the 95% bootstrap confidence interval of 11 participants overlapped with the "equal weights" heuristic bound. The 95% bootstrap confidence interval of the remaining participants either overlapped with both bounds (1 participant) or with neither bound (3 participants). We refer to this method of distinguishing between information-efficient and -inefficient strategies with the abbreviation IB (table 3.1).

We identified the strategy that best fit the choice behavior of each of the 30 participants on the task using maximum likelihood estimation (MLE) described in section 2.6, which we considered to be a model-based approach (abbreviated MB in table 3.1). Since the task was designed such that the optimal IB bound was almost indistinguishable from the "ignore weak" heuristic bound, we considered participants best fit by the optimal and the "ignore weak" model to be using the "optimal" strategy and those best fit by the "equal weights" model to be using the "heuristic" strategy. Out of the 26 participants whose bootstrap confidence interval overlapped with just the optimal or the heuristic bound, 23 were best fit

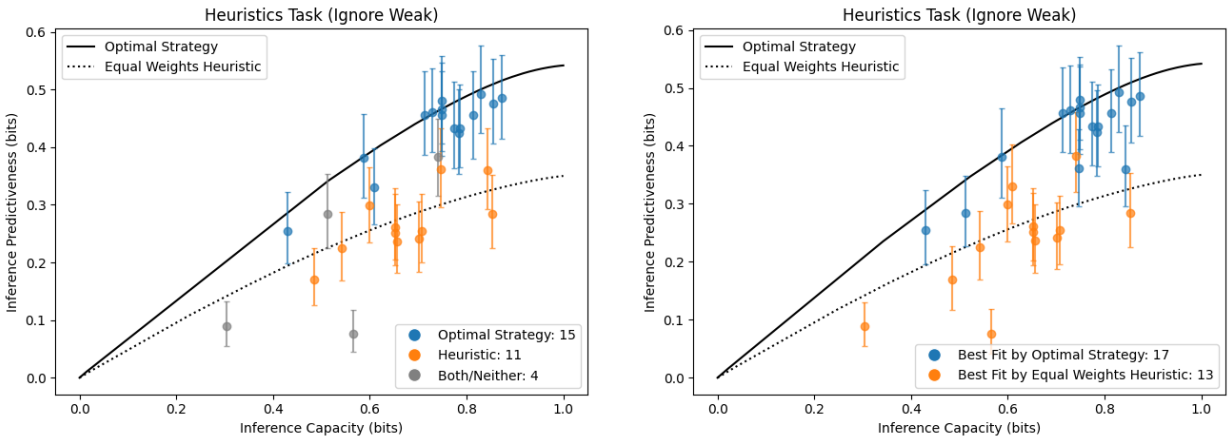

Figure 3.3: 95% bootstrap confidence interval of inference predictiveness for 30 participants on the "ignore weak" heuristics task colored based on overlap with optimal and/or heuristic bound (left) colored based on the model that best fit to choice behavior (right)

by the corresponding model based on MLE (see table 3.1)

	MB: Optimal	MB: Heuristic	Total
IB: Optimal	14	1	15
IB: Heuristic	2	9	11
Total	16	10	26

Table 3.1

The "ignore weak" version of heuristics task not only highlighted a task condition where a substantial number of participants (approximately a third) fell off the bound, but also provided the opportunity for us to compare the model-agnostic IB method of distinguishing information-efficient and -inefficient strategies with more traditional, model-based approaches (see section 2.6). We found that choice behavior that overlapped with the optimal IB bound tend to be best fit by the Bayesian optimal model, and those that overlapped with the inefficient heuristic bound tend to be best fit by the heuristic model, which in this case is based on the "equal weights" strategy.

These findings supported the hypothesis that capacity-limited information-efficient inference along the IB bound is equivalent to performing optimal Bayesian inference and

adding different degrees of softmax noise, as denoted by the β parameter. Conversely, capacity-limited information-inefficient inference off the IB bound can be equated to performing noisy, biased inference. In fact, we observed a significant positive linear relationship between inference capacity and the β parameter of the model that best fit choice behavior across all 30 participants in the "ignore weak" version of heuristics task as shown in fig. 3.4, regardless of whether they were classified as using an "optimal" or "heuristic" strategy under the IB-based or model-based method.

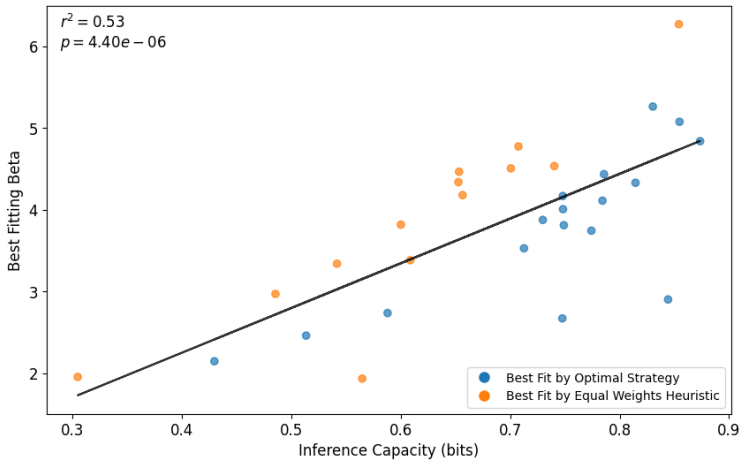

Figure 3.4: Plot of best fitting β parameter against inference capacity for 30 participants in the "ignore weak" heuristic task, color based on the best fitting model

3.4. Classic Manipulations in Decision-Making Performance Reflect Shifts along IB Bound

We explored how choice behavior in IB space shifted from repeated exposure to the same task, increased viewing times, and greater reward incentives using the learning task, the speed-accuracy task, and the reward task.

Across 60 participants in the learning task, the paired-sample Wilcoxon signed-rank test indicated no significant difference in inference capacity ($W = 665, p = 0.0657$) or predictiveness ($W = 848, p = 0.622$) between the first and the second blocks. We also observed no significant difference in distance from the IB bound between the two blocks ($W = 576, p = 0.0701$). These results suggested that while there was no systematic shift in

either inference capacity or predictiveness across repeated exposure of the task, the two were coupled in a way that showed consistent movement along the bound; increases or decreases in inference capacity were accompanied by similarly scaled increases or decreases in inference predictiveness (see fig. 3.5).

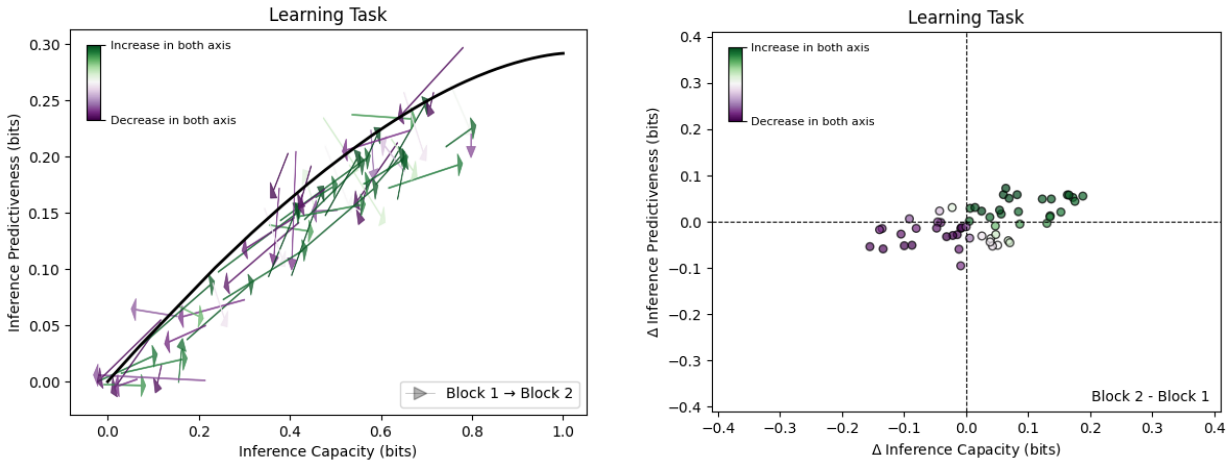

Figure 3.5: Shift in choice behavior from block 1 to block 2 (left) and change in inference predictiveness against change in inference capacity between block 1 and block 2 (right)

Across the 60 participants in the speed-accuracy task, inference capacity and predictiveness were higher in the slow condition compared to the fast condition as we had expected. The paired-sample Wilcoxon signed-rank test indicated a significant difference in inference capacity ($W = 20, p < 0.001$) and a significant difference in inference predictiveness ($W = 194, p < 0.001$) between the fast and the slow conditions. We observed no significant difference in distance from the IB bound between the fast and the slow conditions ($W = 576, p = 0.0701$). These results suggested that longer viewing times led to a systematic increase in inference capacity that was accompanied by a corresponding increase in inference predictiveness, or a shift up and along the IB bound (see fig. 3.6).

Across the 60 participants in the reward task, inference capacity and predictiveness tended to be higher in the high reward condition compared to the no reward condition as we had expected; the paired-sample Wilcoxon signed-rank test indicated a significant

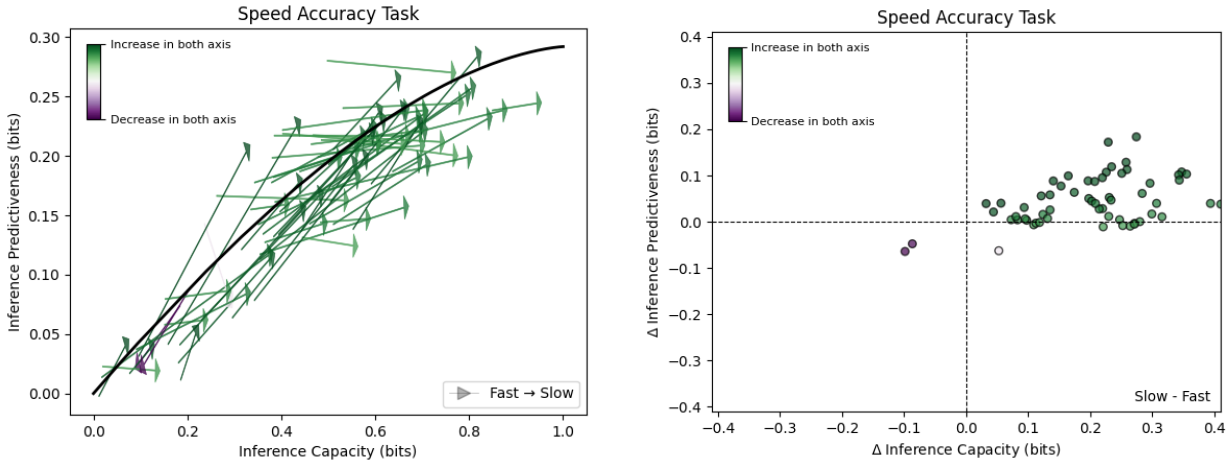

Figure 3.6: Shift in choice behavior from fast to slow condition (left) and change in inference predictiveness against change in inference capacity between the fast and slow conditions (right)

difference in inference capacity ($W = 448, p < 0.001$) and a significant difference in inference predictiveness ($W = 487, p < 0.001$) between the high reward and the no reward conditions. We observed no significant difference in distance from the IB bound between the high reward and the no reward conditions ($W = 718, p = 0.287$). These results suggested that higher reward incentives similarly led to a shift up and along the IB bound (see fig. 3.7)

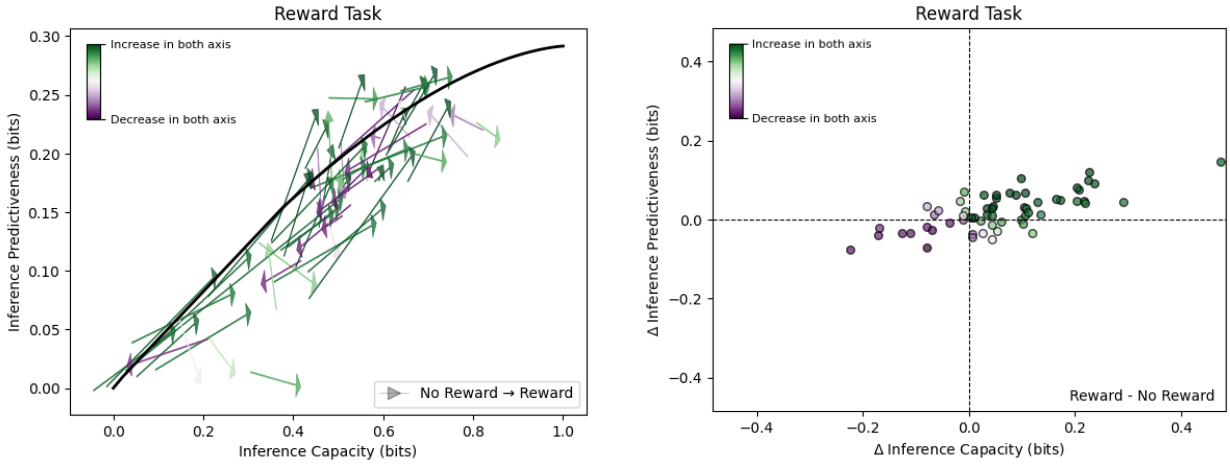

Figure 3.7: Shift in choice behavior from no reward to reward condition (left) and change in inference predictiveness against change in inference capacity between the no reward and reward conditions (right)

Furthermore, we fit a linear mixed effects model to choice behavior of all 60 participants

in the reward task to isolate the effect of reward from potential effects of block order. In the linear mixed effects model for inference capacity, we observed a significant effect of reward magnitude ($t = -2.517, p < 0.05$), but no significant effect of block order ($t = 0.147, p = 0.884$) or interaction between the two terms ($t = -0.091, p = 0.928$). On the other hand, in the linear mixed-effects model for inference predictiveness, the effect of reward magnitude was not significant at the 0.05 level ($t = -1.982, p = 0.0522$), nor was the effect of block order ($t = 0.224, p = 0.824$) or interaction between the two terms ($t = -0.71, p = 0.480$).

While the Wilcoxon signed-rank tests showed significant changes in both inference capacity and predictiveness, the main effect of reward was significant for inference capacity but not for predictiveness in the linear mixed effects model at the $p = 0.05$ level. This suggested that the difference in inference predictiveness between high reward and no reward conditions could be partly explained by the effect of block order and its interaction with reward magnitude. However, given the lack of a significant learning effect in the learning task and a marginal $p = 0.0522$, we further investigated whether the additional complexity of block order and the interaction term was warranted.

In particular, we compared the full linear mixed effects model described in section 2.5.3 (with the block order effect and the interaction term) with a simpler model that included just the main effect of reward magnitude along with random effect of participant. A likelihood ratio test (LRT) indicated that the addition of the block order effect and the interaction term did not significantly improve the fit of the model ($\chi^2 = 0.5197, p = 0.7712$). We therefore reasonably resorted to the simpler model, for which the effect of reward magnitude on inference predictiveness was also significant at the 0.05 level ($t = -3.529, p < 0.001$), concluding that an increase in reward incentives did lead to an increase in inference predictiveness in addition to inference capacity.

In sum, we found that classic manipulations of decision making performance, including changes in viewing times and reward incentives, reflected a shift along the IB curve. As

hypothesized, we observed that longer viewing times and higher incentives led to an increase in inference capacity with a corresponding increase in inference predictiveness while maintaining comparable levels of information efficiency. While participants likewise exhibited consistent levels of information efficiency across repeated exposure to the task, we did not observe any systematic change in inference capacity or predictiveness.

4. DISCUSSION

In this paper, we studied choice behavior using data collected online from a classic inference task called the weather prediction task. Using the IB framework, we assessed whether variability in inference capacity can explain variability observed in inference predictiveness across individuals. We also examined how classic manipulations of decision-making performance—repeated exposure to the same task, increased viewing times, and greater reward incentives—impacted measures of inference capacity and predictiveness. Finally, we studied whether individuals were information-efficient or not under different task conditions using both an IB-based and model-based approach.

First, we found that there was substantial variability in inference capacity across all participants, and that this variability could largely explain differences in inference predictiveness between participants, consistent with the findings of the beads task. Across the learning task and the "equal weights" version the heuristics task, we considered most participants to be performing information-efficient inference, based on the overlap between bootstrap estimates of individual inference predictiveness and the optimal IB bound.

In addition to variability in choice behavior between individuals, we studied variability within individuals between different task conditions using the learning task, the speed-accuracy task, and the reward task. In the learning task, we found no significant change in inference capacity or inference predictiveness, as well as no significant change in distance from IB bound between the first and the second block. This suggested that while participants maintained a consistent level of information efficiency, there was no significant learning that led to a shift along the bound.

Learning seemed to have plateaued before the start of the testing blocks based on the high levels of training ($\sim 70\%$) and pretesting ($\sim 67\%$) performance compared to the testing

performance across both blocks ($\sim 70\%$). While the manipulation in the task (i.e., repeated exposure to the trials) did not produce the intended effect of shifting participants along the IB bound, the results of the task did provide evidence that participants' inference strategies had stabilized by the start of the testing trials and that any observed change in choice behavior between the conditions would not be attributed to learning effects in the other tasks (i.e., speed-accuracy task, reward task) discussed below.

In the speed-accuracy task and the reward task), on the other hand, we did observe a significant increase in inference capacity and inference predictiveness from the fast to the slow condition (in the speed-accuracy task) and from no to high reward (in the reward task, with caveats discussed in section 3.4) while maintaining comparable levels of information efficiency. In other words, the expected increase in task performance from longer viewing times and task rewards can be attributed to an efficient use of more information. We therefore showed that these classic manipulations of decision-making performance resulted in noisy adjustments along the bound, as opposed to deviations from the bound that would arise from any systematic biases.

Last but not least, while the majority of participants performed optimal noisy inference that positioned them along the IB bound, there were a considerable number of participants who used the inefficient heuristic and fell along the heuristics bound in the "ignore weak" version of the heuristics task, according to measures of overlap with the respective bounds and model-based maximum likelihood estimation. In particular, more than a third of the participants fell along the "equal weights" heuristic bound, a heuristic that led to biased inference in a task where the "ignore weak" heuristic was near optimal. This task showed that the use of such biased strategy caused decision-makers to fall off the IB bound and demonstrated that there existed conditions under which decision-makers were not information-efficient.

Overall, results from the weather prediction tasks illustrated that the amount of information used to make inferences (i.e., inference capacity) was an important source of resource

limitation that accounted for the variation in performance observed across individuals, and that can be adjusted under different task conditions where performance was more valuable or when constrained by greater time pressure. Across these different tasks, we also observed that most participants remained information-efficient, the equivalent of performing optimal inference with different degrees of choice noise. We therefore demonstrated how the IB framework can be used to characterize a specific form of bounded rationality, and to show that most people performed optimally in inference tasks under this form of bounded rationality.

Future works in this direction include exploring the neural basis of optimal or suboptimal inference, such as identifying neural correlates of adjustments in inference capacity. Furthermore, while we have shown that some participants adopted information-inefficient strategies in a version of the heuristics task, it might be worthwhile to explore other task conditions that would cause most decision-makers to systematically fall off the bound in addition to shifting along the bound. Last but not least, future work might also entail exploring other forms of resource limitation that decision-makers might balance performance with. For example, the heuristics task suggested that the complexity of inference strategy can also constitute another form of bounded rationality that would account for variations in task performance.

APPENDIX A: DISTRIBUTION GENERATION

	$i=1$	$i=2$	$i=3$	$i=4$
$P(s_i\|A)$	0.08	0.258	0.409	0.253
$P(s_i\|B)$	0.253	0.409	0.258	0.08
$log_{10}\frac{P(s_i\|A)}{P(s_i\|B)}$	-0.5	-0.2	0.2	0.5

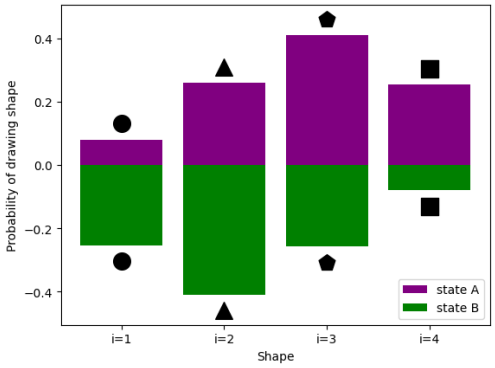

Figure 4.1: The baseline distribution characterized by parameters $p_1 = 0.08, b = 0.2, m = 2.5$ yields a maximum attainable accuracy of 83.3% across 300 testing trials

	$i=1$	$i=2$	$i=3$	$i=4$
$P(s_i\|A)$	0.06	0.186	0.524	0.231
$P(s_i\|B)$	0.231	0.524	0.186	0.06
$log_{10}\frac{P(s_i\|A)}{P(s_i\|B)}$	-0.585	-0.45	0.45	0.585

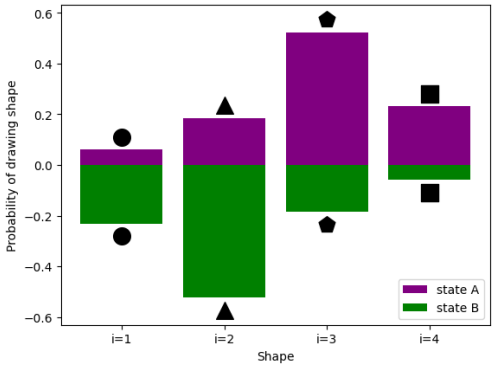

Figure 4.2: The "equal weights" distribution characterized by parameters $p_1 = 0.06, b = 0.45, m = 1.3$ yields a maximum attainable accuracy of 90.33% across 600 testing trials

	$i=1$	$i=2$	$i=3$	$i=4$
$P(s_i\|A)$	0.02	0.25	0.63	0.1
$P(s_i\|B)$	0.1	0.63	0.25	0.02
$log_{10}\frac{P(s_i\|A)}{P(s_i\|B)}$	-1.134	-0.18	0.18	1.134

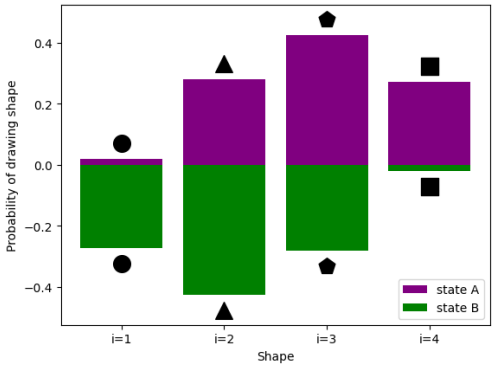

Figure 4.3: The "ignore weak" distribution characterized by parameters $p_1 = 0.02, b = 0.18, m = 6.3$ yields a maximum attainable accuracy of 90.33% across 600 testing trials

We arrived at the generative distribution of each weather prediction task based on a parameter search. Let $\{s_1, s_2, s_3, s_4\}$ be four different shapes, randomly assigned to be triangle, pentagon, circle, and square for each participant. Let $\{A, B\}$ be the two latent states, randomly assigned to be the green or the purple horse for each participant. As noted above, the two latent states are equally likely to occur on any given trial, so we have $P(A) = P(B) = 0.5$. For each trial with latent state $L \in \{A, B\}$, we independently draw 5 shapes with replacement from the set of 4 possible shapes, where the probability of drawing shape s_i is given by the conditional distribution $P(s_i \mid L)$.

The conditional distribution is uniquely characterized by $p_1 = P(s_1|A)$, probability of a shape given a latent state, and a set of log likelihood ratios $log_{10}\frac{P(s_i|A)}{P(s_i|B)}$ for $i \in 1, 2, 3, 4$. Limiting ourselves to symmetrical log-likelihood of the form $[-m * b, -b, b, m * b]$, we varied parameters multiplier m, base b, and p_1 to search for the distributions that 1) maximized distinguishability between information-efficient and information-inefficient strategies 2) provided stable measures of inference capacity and predictiveness and 3) produced a maximum attainable accuracy across testing trials within a reasonable range.

We assessed 1) distinguishability between information-efficient and -inefficient strategies by comparing the area under the IB curve computed using an optimal encoding of the evidence versus using a biased encoding of evidence. We consider two biased encoding schemes, corresponding to the "equal weights" and the "ignore weak" heuristics described in the heuristics task. Greater distinguishability between information-efficient strategies – the class of strategies we consider optimal – and information-inefficient heuristics would provide greater power to detect whether individuals remained information-efficient under different task conditions.

We selected the baseline distribution (see fig. 4.1) for the learning task, the speed-accuracy task, and the reward task based on distinguishability between the optimal strategy and *both* the "equal weights" and the "ignore weak" heuristics. On the other hand, we

selected the distribution favoring the "equal weights" strategy (see fig. 4.2) based on distinguishability between the optimal strategy and the "ignore weak" heuristic, whereas we selected the distribution favoring the "ignore weak" strategy (see fig. 4.3) based on distinguishability between the optimal strategy and the "equal weights" heuristic for the two versions of the heuristics task, respectively.

Given a conditional distribution, we can also assess the choice behavior of the Bayesian-optimal decision maker. On a trial with shape combination $S = (s_1, s_2, s_3, s_4, s_5)$, the fully optimal decision-maker selects the latent state $L^* = \underset{L \in \{A,B\}}{\arg\max}\, P(L|S)$, with the posterior probabilities given by

$$P(A|S) = \frac{P(S|A)}{P(S|A) + P(S|B)} \qquad P(B|S) = 1 - P(A|S) = \frac{P(S|B)}{P(S|A) + P(S|B)} \tag{4.1}$$

where $P(S|A) = \prod_{j=1}^{5} P(s_j|A)$ and $P(S|B) = \prod_{j=1}^{5} P(s_j|B)$, since the two latent states are equally likely to occur and the 5 shapes are drawn independently with replacement.

We can similarly reason about the choice behavior of the Bayesian-optimal decision maker by summing the log likelihood ratios of the shapes on a given trial. In fact, choosing the state with the greater posterior probability is equivalent to choosing state A when the the sum of the log likelihood ratios exceeds 0:

$$\text{LLR sum} = \sum_{j=1}^{5} \log \frac{P(s_j|A)}{P(s_j|B)} = \log \prod_{j=1}^{5} \frac{P(s_j|A)}{P(s_j|B)} = \log \frac{P(S|A)}{P(S|B)} = \log \frac{P(A|S)}{P(B|S)} \tag{4.2}$$

Since the prior probabilities of the two states are equal in this task, the sum of the log likelihood ratios is also equivalent to the log posterior odds:

$$\text{LLR sum} = \log \frac{P(A|S)}{P(B|S)} = \log \frac{P(S|A)}{P(S|B)} \tag{4.3}$$

The posterior probabilities can therefore be computed from the sum of the log likelihood

ratio as follows as well:

$$P(A|S) = \frac{10^{\text{LLR sum}}}{1 + 10^{\text{LLR sum}}} \qquad P(B|S) = 1 - P(A|S) = \frac{1}{1 + 10^{\text{LLR sum}}} \tag{4.4}$$

As hypothesized in section 2.1, information-efficient inference can be equated to optimal inference corrupted by different degrees of softmax noise (i.e., β). Given a trial with shape combination S, an information-efficient decision-maker with a trade-off between inference predictiveness and inference capacity defined by the β parameter chooses latent state L with choice probabilities given by

$$P_\beta(A|S) = \frac{e^{\beta\alpha^* P(A|S)}}{e^{\beta\alpha^* P(A|S)} + e^{\beta\alpha^* P(B|S)}} \qquad P_\beta(B|S) = \frac{e^{\beta\alpha^* P(B|S)}}{e^{\beta\alpha^* P(A|S)} + e^{\beta\alpha^* P(B|S)}}, \tag{4.5}$$

where $P(A|S)$ and $P(B|S)$ represent the respective posterior probabilities in eq. (4.4) and α^* represents the contrast between how accurate versus how inaccurate a decision-maker can expect to be. Note that as $\beta \to \infty$, the choice behavior of the information-efficient decision-maker is fully optimal.

Using simulations of choice behavior described above, we searched for probability distributions that produced 2) stable measures of inference capacity and inference predictiveness. We assessed the extent to which the same decision-maker with the same trade-off between inference capacity and predictiveness (i.e., same β parameter) would achieve approximately the same performance on repeated attempts of the inference task generated from a given probability distribution. In particular, we repeatedly simulated choice behavior for the optimal decision-maker using a range of β parameter on the same sequence of testing trials and visually inspected its variability when plotted in IB space.

Finally, we selected for probability distributions that 3) yielded a maximum attainable accuracy between approximately 75% to 90%. The maximum attainable accuracy was calculated based on the simulated choice behavior of the fully-optimal decision-maker over

300 testing trials of the baseline distribution (see fig. 4.1) used in the learning task, the speed-accuracy task, and the reward task and over 600 testing trials of the "equal weights" distribution (see fig. 4.2) and the "ignore weak" distribution (see fig. 4.3) in the heuristics task. Limiting maximum attainable accuracy to this range allowed for assessments of behavior under varying levels of uncertainty without making the task impossible.

APPENDIX B: MODEL FITTING

To provide a frame of reference for the IB-based method of distinguishing between information-efficient (i.e., on the bound) and -inefficient (i.e., off the bound) strategies of inference introduced in section 2.6, we also performed a model-fitting procedure for each of the 60 participants in two versions of the heuristics task. This procedure used maximum likelihood estimation (MLE) to determine whether the choice behavior of each participant is best fit by an optimal strategy, an "equal weights" strategy, or an "ignore weak" strategy (see section 2.5.4).

We defined a decision-maker who uses the optimal strategy as making choices based on the actual log likelihood ratios $[-m*b, -b, b, m*b]$ associated with the task. Likewise, we defined a decision-maker who uses the "equal weights" strategy as making choices based on log likelihood ratios $[-x, -x, x, x]$, where $x = log_{10}\frac{P(s_3|A)+P(s_4|A)}{P(s_1|A)+P(s_2|A)}$. Finally, we defined a decision-maker who uses the "ignore weak" strategy as making choices based on log likelihood ratios $[-m*b, 0, 0, m*b]$ when there are strong shapes and optimal actual log likelihood ratios when there are no strong shapes on a given trial. See Appendix A for the task parameters associated with each version of the heuristics task.

For each of the three strategies, we used the corresponding log likelihood ratios defined above to compute posterior probabilities on each given trial based on eq. (4.4) in Appendix A. Given a value of β, we computed the choice probabilities of an agent using that strategy with eq. (4.5) in Appendix A. These choice probabilities were then used to calculate the sum of negative log likelihood for each of the 60 participants across all three strategies, as an inverse measure of how well their choice behavior aligned with a particular strategy.

To determine the strategy that best fit each participant's choice behavior, we used SciPy's `optimize` package to minimize the negative log-likelihood for the three strategies

listed above, with β as a free parameter.[1] For each participant, we identified the strategy with the lowest minimum negative log-likelihood as the one that best fit their choice behavior. The corresponding β value from the best-fitting strategy was then used as an estimate of choice noise for that participant.

[1]See the package documentation: https://docs.scipy.org/doc/scipy/tutorial/optimize.html

BIBLIOGRAPHY

Ashby, F. G., & Maddox, W. T. (2005). Human category learning. *Annual Review of Psychology*, *56*(1), 149–178. https://doi.org/10.1146/annurev.psych.56.091103.070217

Bochud-Fragnière, E., Banta Lavenex, P., & Lavenex, P. (2022). What is the weather prediction task good for? a new analysis of learning strategies reveals how young adults solve the task. *Frontiers in Psychology*, *13*, 886339. https://doi.org/10.3389/fpsyg.2022.886339

Bochud-Fragnière, E., Lavenex, P., & Banta Lavenex, P. (2023). When and how do children solve the weather prediction task? *Developmental Psychobiology*, *65*(6), e22407. https://doi.org/10.1002/dev.22407

Denti, T., Marinacci, M., & Montrucchio, L. (2020). A note on rational inattention and rate distortion theory. *Decisions in Economics and Finance*, *43*(1), 75–89. https://doi.org/10.1007/s10203-019-00243-0

Eissa, T. L., Gold, J. I., Josić, K., & Kilpatrick, Z. P. (2022). Suboptimal human inference can invert the bias-variance trade-off for decisions with asymmetric evidence. *PLOS Computational Biology*, *18*(7), e1010323. https://doi.org/10.1371/journal.pcbi.1010323

Filipowicz, A. L., Forkin, G., Piasini, E., Krishnamurthy, K., Kable, J. W., Balasubramanian, V., & Gold, J. I. (2020). Preregistration: The optimality of human inference in uncertain environments.

Filipowicz, A. L., Glaze, C. M., Kable, J. W., & Gold, J. I. (2020). Pupil diameter encodes the idiosyncratic, cognitive complexity of belief updating. *eLife*, *9*, e57872. https://doi.org/10.7554/eLife.57872

Gabay, Y., Vakil, E., Schiff, R., & Holt, L. L. (2015). Probabilistic category learning in developmental dyslexia: Evidence from feedback and paired-associate weather prediction tasks. *Neuropsychology*, *29*(6), 844–854. https://doi.org/10.1037/neu0000194

Glaze, C. M., Filipowicz, A. L. S., Kable, J. W., Balasubramanian, V., & Gold, J. I. (2018). A bias–variance trade-off governs individual differences in on-line learning in an unpredictable environment. *Nature Human Behaviour*, *2*(3), 213–224. https://doi.org/10.1038/s41562-018-0297-4

Gluck, M. A., Shohamy, D., & Myers, C. (2002). How do people solve the "weather prediction" task?: Individual variability in strategies for probabilistic category learning. *Learning & Memory*, *9*(6), 408–418. https://doi.org/10.1101/lm.45202

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*(1), 535–574. https://doi.org/10.1146/annurev.neuro.29.051605.113038

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291. https://doi.org/10.2307/1914185

Kelmendi, B., Adams, T., Jakubovski, E., Hawkins, K. A., Coric, V., & Pittenger, C. (2016). Probing implicit learning in obsessive-compulsive disorder: Moderating role of medication on the weather prediction task. *Journal of Obsessive-Compulsive and Related Disorders*, *9*, 90–95. https://doi.org/10.1016/j.jocrd.2016.03.003

Kira, S., Yang, T., & Shadlen, M. N. (2015). A neural implementation of wald's sequential probability ratio test. *Neuron*, *85*(4), 861–873. https://doi.org/10.1016/j.neuron.2015.01.007

Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996). A neostriatal habit learning system in humans. *Science*, *273*(5280), 1399–1402. https://doi.org/10.1126/science.273.5280.1399

Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1994). Probabilistic classification learning in amnesia. *Learning & Memory*, *1*(2), 106–120. https://doi.org/10.1101/lm.1.2.106

Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, *43*, e1. https://doi.org/10.1017/S0140525X1900061X

Maćkowiak, B., Matějka, F., & Wiederholt, M. (2023). Rational inattention: A review. *Journal of Economic Literature*, *61*(1), 226–273. https://doi.org/10.1257/jel.20211524

Matějka, F., & McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, *105*(1), 272–298. https://doi.org/10.1257/aer.20130047

Palmer, S. E., Marre, O., Berry, M. J., & Bialek, W. (2015). Predictive information in a sensory population. *Proceedings of the National Academy of Sciences*, *112*(22), 6908–6913. https://doi.org/10.1073/pnas.1506855112

Parker, J. A. (2023). Theoretical, behavioral, and neural implications of information-limited inference [Unpublished doctoral preliminary proposal].

Piasini, E., Filipowicz, A. L. S., Levine, J., & Gold, J. I. (2021). Embo: A python package for empirical data analysis using the information bottleneck. *Journal of Open Research Software*, *9*(10). https://doi.org/10.5334/jors.322

Shwartz-Ziv, R., & Tishby, N. (2017). Opening the black box of deep neural networks via information. *arXiv*. https://doi.org/10.48550/arXiv.1703.00810

Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, *69*(1), 99. https://doi.org/10.2307/1884852

Strouse, D. J., & Schwab, D. J. (2016). The deterministic information bottleneck. *arXiv*. https://doi.org/10.48550/arXiv.1604.00268

Tavoni, G., Doi, T., Pizzica, C., Balasubramanian, V., & Gold, J. I. (2022). Human inference reflects a normative balance of complexity and accuracy. *Nature Human Behaviour*, *6*(8), 1153–1168. https://doi.org/10.1038/s41562-022-01357-z

Tishby, N., Pereira, F. C., & Bialek, W. (2000). The information bottleneck method. *arXiv*. http://arxiv.org/abs/physics/0004057

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*, 1124–1131. https://doi.org/10.1126/science.185.4157.1124

Wason, P. C. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology*, *20*(3), 273–281. https://doi.org/10.1080/14640746808400161

Yang, T., & Shadlen, M. N. (2007). Probabilistic reasoning by neurons. *Nature*, *447*(7148), 1075–1080. https://doi.org/10.1038/nature05852