



## OPEN ACCESS

## EDITED BY

Jochen Ditterich,  
University of California, Davis,  
United States

## REVIEWED BY

Varun Dutt,  
Indian Institute of Technology Mandi,  
India  
Ronald Van Den Berg,  
Stockholm University, Sweden

## \*CORRESPONDENCE

Cecilia Lindig-León  
cecilia.lindig-leon@uni-ulm.de

†These authors have contributed  
equally to this work

## SPECIALTY SECTION

This article was submitted to  
Decision Neuroscience,  
a section of the journal  
Frontiers in Neuroscience

RECEIVED 28 March 2022

ACCEPTED 09 September 2022

PUBLISHED 29 September 2022

## CITATION

Lindig-León C, Kaur N and Braun DA  
(2022) From Bayes-optimal to heuristic  
decision-making in a two-alternative  
forced choice task with an  
information-theoretic bounded  
rationality model.  
*Front. Neurosci.* 16:906198.  
doi: 10.3389/fnins.2022.906198

## COPYRIGHT

© 2022 Lindig-León, Kaur and Braun.  
This is an open-access article  
distributed under the terms of the  
[Creative Commons Attribution License  
\(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or  
reproduction in other forums is  
permitted, provided the original  
author(s) and the copyright owner(s)  
are credited and that the original  
publication in this journal is cited, in  
accordance with accepted academic  
practice. No use, distribution or  
reproduction is permitted which does  
not comply with these terms.

# From Bayes-optimal to heuristic decision-making in a two-alternative forced choice task with an information-theoretic bounded rationality model

Cecilia Lindig-León<sup>\*†</sup>, Nehchal Kaur<sup>†</sup> and Daniel A. Braun

Faculty of Engineering, Computer Science and Psychology, Institute of Neural Information Processing, Ulm University, Ulm, Germany

Bayes optimal and heuristic decision-making schemes are often considered fundamentally opposed to each other as a framework for studying human choice behavior, although recently it has been proposed that bounded rationality may provide a natural bridge between the two when varying information-processing resources. Here, we investigate a two-alternative forced choice task with varying time constraints, where subjects have to assign multi-component symbolic patterns to one of two stimulus classes. As expected, we find that subjects' response behavior becomes more imprecise with more time pressure. However, we also see that their response behavior changes qualitatively. By regressing subjects' decision weights, we find that decisions allowing for plenty of decision time rely on weighing multiple stimulus features, whereas decisions under high time pressure are made mostly based on a single feature. While the first response pattern is in line with a Bayes-optimal decision strategy, the latter could be considered as an instantiation of heuristic decision-making with cue discounting. When fitting a bounded rational decision model with multiple feature channels and varying information-processing capacity to subjects' responses, we find that the model is able to capture subjects' behavioral change. The model successfully reflects the simplicity of heuristics as well as the efficiency of optimal decision making, thus acting as a bridge between the two approaches.

## KEYWORDS

Bayes-optimality, heuristics, bounded rationality, decision-making, information theory

## 1. Introduction

In the cognitive and neural sciences, Bayes-optimal decision-making models have advanced over the last few decades to one of the dominant formal approaches to understand reasoning and acting in humans, animals, and machines (Doya et al., 2007). These models have strong theoretical appeal due to their axiomatic simplicity and elegance that allow for a normative interpretation of how decisions should be taken

optimally by rational actors (von Neumann and Morgenstern, 1944; Savage, 1954). Bayes-optimal models typically consist of two distinct components: first a model of beliefs that represent the states of unobserved variables taking into account the subjective uncertainty which is updated with new incoming data, and second an optimal decision-making scheme that selects actions in a way that maximizes the expected utility of the decision-maker given the current beliefs. Although it has been noted early on in the economic sciences that human decision-makers can substantially deviate from Bayes-optimality (Allais, 1953; Ellsberg, 1961), there is an extensive body of literature that has applied Bayesian modeling successfully not only to describe economic decision-making, but also to capture basic sensorimotor behavior like cue integration, perceptual categorization, visuomotor adaptation, and movement control in both humans and animals (van Beers et al., 1999; Ernst and Banks, 2002; Knill and Pouget, 2004; Körding and Wolpert, 2004; Körding and Wolpert, 2006; Braun et al., 2009).

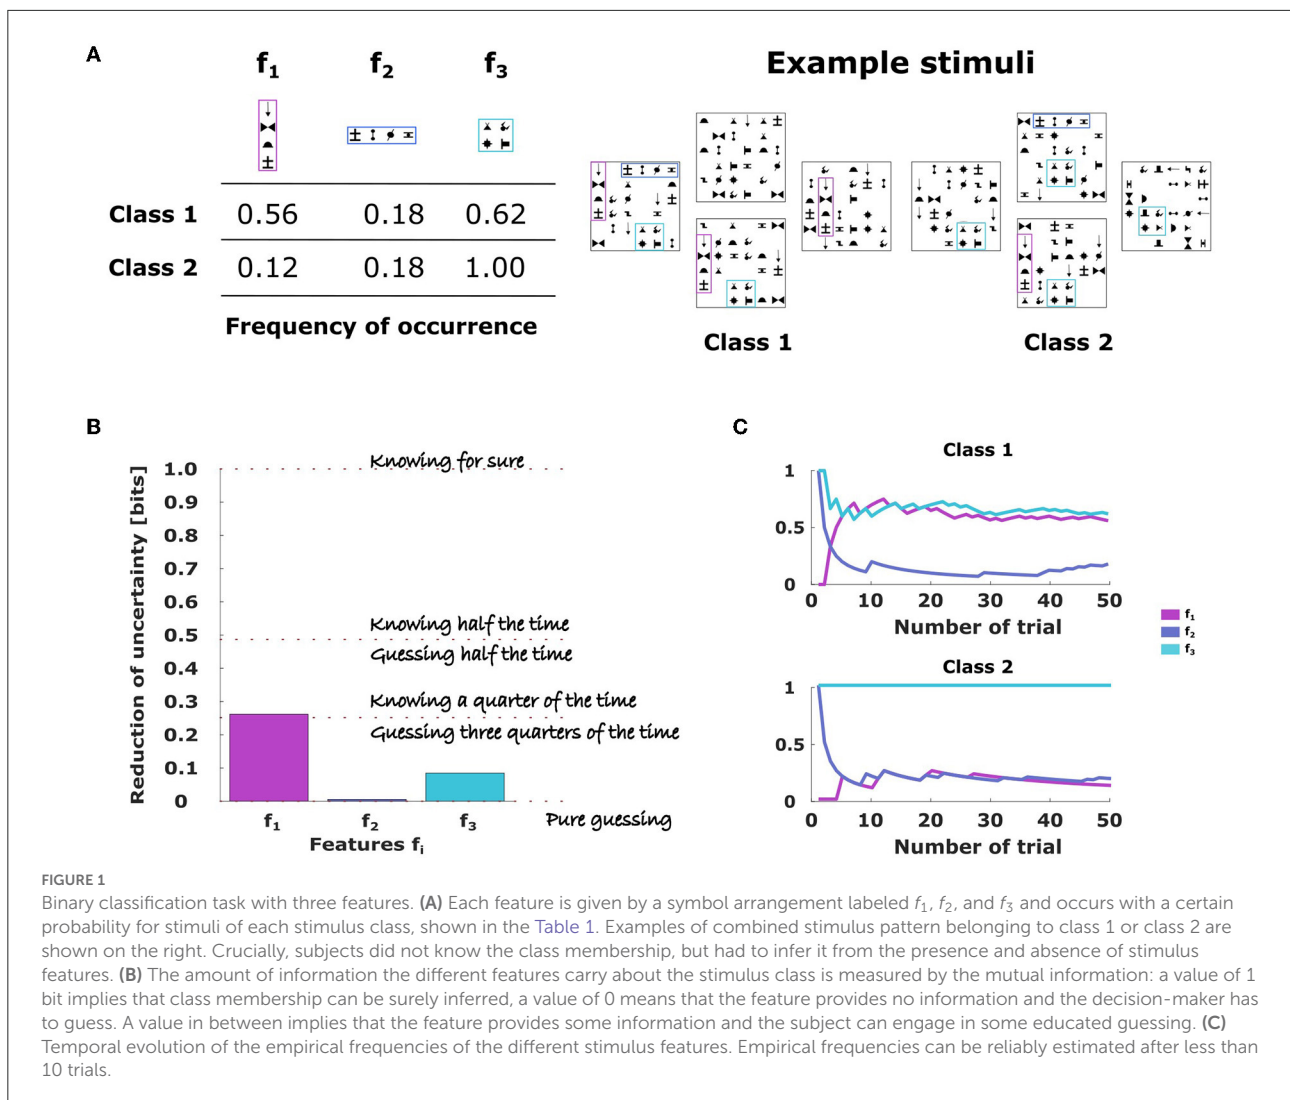
Nevertheless, Bayesian models have been criticized on several grounds, foremost for their computational intractability, in particular in complex environments, their assumption of perfect rationality, their disregard for model mis-specification as well as their explanatory flexibility when allowing for arbitrary prior beliefs, likelihood models, or cost functions (Jones and Love, 2011). An alternative view to Bayes-optimal decision-making that does not assume an overarching optimization process or principle is sometimes called the Adaptive Toolbox approach to the mind based on many different specialized decision heuristics that are adapted to various environmental contexts (Gigerenzer and Gaissmaier, 2011; Gigerenzer et al., 2011). A key feature of heuristic strategies is that they are fast to apply and that they ignore a large part of the information at hand. Well-known heuristic strategies include the recognition heuristic (Goldstein and Gigerenzer, 2002)—assuming a higher value for an option whose name is recognized—, the take-the-best-heuristic (Gigerenzer and Goldstein, 1996a)—compare two options by their most important criterion and ignore other criteria—and elimination by aspects (Tversky, 1972) that gets rid of alternatives that do not meet the criteria of a certain aspect. Over the years a considerable body of evidence has built up in the literature showing how human decision-makers rely on heuristic decision-making in many real-world scenarios, for example in the financial and medical professions (Julian and Gerd, 2012; Forbes et al., 2015).

Usually, Bayes-optimal decision strategies that integrate all the available information into an optimal decision are regarded as incompatible with the idea of having multiple heuristics that ignore large amounts of information and that are not optimal in any particular sense, even though such heuristic strategies have been shown to work well in practical scenarios where exact models are unavailable or misspecified. Recently, however, there have been several proposals that suggest that heuristics can be regarded as Bayes-optimal strategies under certain constraints.

Parpart et al. (2018), for example, have proposed that ignoring certain aspects of data in a heuristic decision-making scheme could be enforced within a Bayesian linear regression framework by putting a high weight on the corresponding prior over regression parameters. In another study (Belousov et al., 2016), Belousov and colleagues have suggested that different reactive and predictive strategies of catching a ball that were previously suggested as heuristic strategies can arise from an optimal stochastic controller when varying the level of observation noise and reaction time. While such adaptive changes in strategy selection depending on task constraints are well-known in the psychological literature (Svenson and Edland, 1987; Payne et al., 1988; Rieskamp and Hoffrage, 2008; Pachur and Bröder, 2013), a general mathematical formalization in terms of a single framework that explains both heuristic and Bayes-optimal decision-making poses an ongoing research problem. Such a framework might not only elucidate the mathematical principles underlying heuristic decision-making, but also shed further light on the question of how heuristics may be learned.

Here, we test an information-theoretic bounded rationality model (Braun et al., 2011; Ortega and Braun, 2011, 2013) for its ability to capture both Bayes-optimal and heuristic decision strategies in human subjects when varying the available information-processing resources in a binary classification task. In this model, limited information-processing capabilities are formalized abstractly by information constraints that quantify how much decision-makers can deviate from a given prior strategy when they face a decision task that is represented abstractly by a utility function. Decision-makers with more resources will be able to deviate more from their priors and adapt to the utility function, whereas decision-makers with limited resources will have to stick more to their prior strategy. The basic rationale is that any kind of resource constraint (e.g., time, memory, money, effort,...) will ultimately translate into a reduced ability for uncertainty reduction (Gottwald and Braun, 2019), which can be considered as an application of the rate-distortion principle of information theory (Cover and Thomas, 2006; Tishby and Polani, 2011). Previously, we have used such information-theoretic bounded rationality models to study effects of resource limitations in motor planning (Schach et al., 2018), absolute identification (Lindig-León et al., 2019), and sensorimotor interactions (Lindig-León et al., 2021).

In our study, we expose human subjects to a binary classification task by two-alternative forced choice, where a stimulus (here: a panel of abstract symbols) has to be assigned to one of two classes based on three different stimulus features (here: certain symbol arrangements)—compare Figures 1A,B. The features have different probabilities for the two stimulus classes and are independent, that means given the class membership the presence of a feature has no information about any other feature. The optimal decision-maker in such a scenario is the Naive Bayes classifier that provides the performance baseline irrespective of the available



information-processing resources. This is in contrast to bounded rational decision-makers that optimize performance under given constraints on resources. In particular, we investigate the hypothesis that subjects behave like bounded rational decision-makers that focus on the most informative features under time pressure, thus, ignoring other stimulus information, which corresponds to some kind of Take-the-Best decision heuristic (Gigerenzer and Goldstein, 1996b). To this end, we measure the amount of information subjects extract from the different features and compare that to the information reduction expected by different bounded rationality models—compare. We find that subjects’ behavior is best captured by an information-theoretic bounded rationality model that optimizes performance assuming capacity constraints on three separate information channels representing the three features, accounting for the fact that feature information is processed in separate chunks.

## 2. Methods

### Participants

There were a total of 16 participants. They belonged to the age group 20–50, 11 of them conducted the task using remote access. The study was approved by the Ethics Committee of Ulm University. Subjects were informed about the purpose and method of the experiment as well as the ensuing data handling, and that they could withdraw from the study any time. During the initial familiarization phase of the experiment, subjects were explicitly instructed that stimuli were composed of three different features that could occur in isolation or in combination with each other along with distractor symbols. They were shown the three features which were each made up of four abstract symbols inspired by a previous study (Orbán et al., 2008). Subjects were told that they have to classify the stimulus

patterns that they will observe into one of two classes and that they can learn this classification from feedback after each trial.

## Experimental design

**Trial design.** In each trial subjects were shown a  $6 \times 6$  square panel (ca.  $15 \times 15$  cm) made of 36 tiles full of blanks and symbols and they were required to press the keys 1 and 2 on the keyboard to indicate whether they believed that the presented stimulus belonged either to class 1 or 2, respectively. The key press had to occur within a given time window depending on the condition (fast condition 1 s, medium condition 3 s, slow condition 5 s, otherwise the trial was counted as a missed trial). After each trial a dialogue box at the center of the screen would indicate to subjects whether the trial was correct, incorrect, or missed. Subjects could have breaks after every block of 50 consecutive trials with the same time constraint (fast, medium, or slow). In total, there were six blocks, that is 300 trials altogether per session, and two sessions per subject. The first three blocks of 50 trials in the first session were executed in the order slow, medium, fast, in the ensuing blocks the order was randomized. Each session lasted approximately half an hour. All evaluations in the paper are based on data from the second session.

**Stimulus classes.** There were two equiprobable classes of stimuli, where each class is defined like a bag-of-words model based on three independent features. In a bag-of-words model such features typically are given by characteristic words whose frequency of occurrence distinguishes, for example, spam from non-spam email. Similarly, in our task, the class-specific occurrence probabilities of specified visual features, each of which could either be present or absent, allows for the classification of stimuli into two separate classes. Based on the three features, we can distinguish eight different stimulus types that can be denoted by a binary vector  $\mathbf{x}$  and that are associated with the two classes with different probabilities. In our experiment we used the following feature probabilities to generate 50 stimuli for each class and condition as a random template:

Classes	F1	F2	F3
Class 1	0.616	0.089	0.680
Class 2	0.198	0.227	0.985

When subsequently counting the occurrences of features in this template, the empirical frequencies with which features occur did not exactly match the aforementioned probabilities due to the finite amount of stimulus samples. Instead, the empirical frequencies were given by:

Classes	F1	F2	F3
Class 1	0.560	0.180	0.620
Class 2	0.120	0.180	1.000

We used the same random trial sequence template for all subjects, so we could ensure that each subject experienced exactly the same empirical frequencies. It should also be noted that the empirical frequencies of the trial sequence template could not be known immediately by subjects, but had to be learned. However, as can be seen from Figure 1C, the empirical estimates of the feature frequencies stabilizes already after 10 trials around their final values.

**Stimulus design.** The stimulus design shown in Figure 1 was inspired by a previous study (Orbán et al., 2008). Each stimulus was a  $6 \times 6$  square panel, where each square could contain a symbol or a blank. To balance stimuli in terms of luminosity, each stimulus contained exactly 25 symbols and 11 blanks. The stimuli were characterized by three independent visual features where each feature corresponds to a spatially order chunk of four symbols. The chunks appeared in a fixed part of the panel, but their exact position could vary by two squares at most. Depending on the number of chunks present, more or less randomly placed distractor symbols were used to arrive at the total number of 25 symbols. There was an inventory of 20 different symbols out of which 10 were used to define the three chunks and all of them could be used as distractor symbols.

## Models

The eight possible stimulus pattern can be described by a binary input vector  $\mathbf{x} \in \{0, 1\}^3$ , for example for the first stimulus pattern with all cues present it would be  $\mathbf{x} = (111)$ . The distribution  $P(\mathbf{x})$  over stimulus pattern is given by  $P(\mathbf{x}) = \sum_{y \in \{1, 2\}} P(\mathbf{x}|y)\pi(y)$ , where  $\pi(y) = \frac{1}{2}$  is the equal prior probability of presenting a stimulus from either class, and  $P(\mathbf{x}|y)$  is the generative model of the stimulus pattern in each class  $y$ . For our task, the generative model is given by

$$P(\mathbf{x}|y = 1) = \prod_{j=1}^3 \mu_j^{x_j} (1 - \mu_j)^{1-x_j}$$

$$P(\mathbf{x}|y = 2) = \prod_{j=1}^3 v_j^{x_j} (1 - v_j)^{1-x_j}$$

where  $\mu_j$  and  $v_j$  are the probabilities of the features occurring in class 1 and class 2, respectively, as determined from the experiment. The decision-maker's choice can be described by a binary output variable  $a \in \{1, 2\}$  corresponding to a selection of class 1 and class 2, respectively. A choice strategy then corresponds to a conditional distribution  $P(a|\mathbf{x})$  detailing the probability of selection class  $a$  given the stimulus pattern  $\mathbf{x}$ . For example, if  $\mathbf{x} = (111)$ , the choice strategy is to assign  $\mathbf{x}$  to class 1 with a probability  $p$  and to class 2 with probability  $1 - p$ .

**Naive bayes model.** The Bayes-posterior for our task is given by the sigmoid  $P(y = 1|\mathbf{x}) = \sigma(\mathbf{w}^T \mathbf{x} + w_0)$

and  $P(y = 2|\mathbf{x}) = 1 - P(y = 1|\mathbf{x})$  with the weight parameters

$$w_j = \log \frac{\mu_j}{1 - \mu_j} - \log \frac{v_j}{1 - v_j}$$

$$w_0 = \sum_{j=1}^3 \log \frac{1 - \mu_j}{1 - v_j}$$

Assuming a 0/1-utility of  $U(y, a) = \delta_{y,a}$  with the Kronecker delta  $\delta_{y,a} = 1$  if  $y = a$  and  $\delta_{y,a} = 0$  otherwise, we can then define the optimal choice strategy  $P(a|\mathbf{x}) = \delta_{a,a^*(\mathbf{x})}$  that maximizes the expected utility according to

$$a^*(\mathbf{x}) = \arg \max_a \sum_y P(y|\mathbf{x})U(y, a)$$

**Bounded rational model with a single information channel for action.** Assuming the same expected utility as above given by  $EU(\mathbf{x}, a) = \sum_{y \in \{1,2\}} P(y|\mathbf{x})U(y, a)$ , the bounded rational decision-maker with a single information channel for action seeks to find the optimal strategy  $P^*(a|\mathbf{x})$  that maximizes the expected utility

$$P^*(a|\mathbf{x}) = \arg \max_{P(a|\mathbf{x})} \sum_{\mathbf{x}} \sum_a P(\mathbf{x})P(a|\mathbf{x})EU(\mathbf{x}, a)$$

subject to the information constraint

$$I(\mathbf{X}; A) = \sum_{\mathbf{x}} \sum_a P(\mathbf{x})P(a|\mathbf{x}) \log \frac{P(a|\mathbf{x})}{P(a)} \leq K$$

where the prior is given by  $P(a) = \sum_{\mathbf{x}} P(\mathbf{x})P(a|\mathbf{x})$  and  $K$  is a positive real-valued bound. The optimization problem is equivalent to the unconstrained optimization problem

$$P^*(a|\mathbf{x}) = \arg \max_{P(a|\mathbf{x})} \sum_{\mathbf{x}} \sum_a P(\mathbf{x})P(a|\mathbf{x})EU(\mathbf{x}, a) - \frac{1}{\beta} I(\mathbf{X}; A)$$

where the precision parameter  $\beta$  is chosen to match the information bound  $K$ . The optimal strategy is given by

$$P^*(a|\mathbf{x}) = \frac{1}{Z(\mathbf{x})} P^*(a) \exp[\beta EU(\mathbf{x}, a)]$$

with the optimal prior  $P^*(a) = \sum_{\mathbf{x}} P(\mathbf{x})P^*(a|\mathbf{x})$  and the normalization constant  $Z(\mathbf{x}) = \sum_a P^*(a) \exp(\beta EU(\mathbf{x}, a))$ .

**Bounded rational model with multiple information channels for action and perception.** Assuming an internal state  $\mathbf{s} = (s_1 s_2 s_3)$  for perceptual processing and feature-specific information channels  $P(s_j|x_j)$  with  $j = 1, 2, 3$  that learn to specialize in processing the three different symbol arrangements serving as features, the bounded rational decision problem with multiple information channels for both perception and action can be stated as a joint optimization problem, that does not only seek to optimize the choice strategy  $p(a|\mathbf{s})$  given the perceptual

information  $\mathbf{s}$ , but also learns to co-optimize the perceptual filters  $P(s_j|x_j)$ , yielding:

$$\arg \max_{\substack{P(s_1|x_1) \\ P(s_2|x_2) \\ P(s_3|x_3) \\ P(a|\mathbf{s})}} E_{\mathbf{x}, \mathbf{s}, a} [EU(\mathbf{x}, a)] - \frac{1}{\beta} \sum_{j=1}^3 I(X_j; S_j) - \frac{1}{\beta_U} I(A; \mathbf{S})$$

where

$$I(X_j; S_j) = \sum_{x_j} \sum_{s_j} P(x_j)P(s_j|x_j) \log \frac{P(s_j|x_j)}{P(s_j)}$$

$$I(A; \mathbf{S}) = E_{\mathbf{x}, \mathbf{s}, a} \left[ \log \frac{P(a|\mathbf{s})}{P(a)} \right]$$

with the expectation

$$E_{\mathbf{x}, \mathbf{s}, a} [f(\cdot)] = \sum_{x_1 x_2 x_3} \sum_{s_1 s_2 s_3} \sum_a P(\mathbf{x})P(s_1|x_1)P(s_2|x_2)P(s_3|x_3)P(a|\mathbf{s})f(\cdot)$$

The optimal information channels for perceptual processing of the features are given by

$$P^*(s_1|x_1) = \frac{1}{Z(x_1)} P^*(s_1) \exp \left( \beta \sum_{x_2 x_3} \sum_{s_2 s_3} P(s_2|x_2)P(s_3|x_3) \times P(x_2 x_3|x_1)F(\mathbf{x}, \mathbf{s}) \right)$$

$$P^*(s_2|x_2) = \frac{1}{Z(x_2)} P^*(s_2) \exp \left( \beta \sum_{x_1 x_3} \sum_{s_1 s_3} P(s_1|x_1)P(s_3|x_3) \times P(x_1 x_3|x_2)F(\mathbf{x}, \mathbf{s}) \right)$$

$$P^*(s_3|x_3) = \frac{1}{Z(x_3)} P^*(s_3) \exp \left( \beta \sum_{x_1 x_2} \sum_{s_1 s_2} P(s_1|x_1)P(s_2|x_2) \times P(x_1 x_2|x_3)F(\mathbf{x}, \mathbf{s}) \right)$$

with

$$F(\mathbf{x}, \mathbf{s}) = \sum_a P^*(a|\mathbf{s})EU(\mathbf{x}, a) - \frac{1}{\beta_U} \sum_a P^*(a|\mathbf{s}) \log \frac{P^*(a|\mathbf{s})}{P^*(a)}$$

and the optimal action channel

$$P^*(a|\mathbf{s}) = \frac{1}{Z(\mathbf{s})} P^*(a) \exp \left( \beta_U \sum_{\mathbf{x}} P(\mathbf{x}|\mathbf{s})EU(\mathbf{x}, a) \right)$$

where

$$P(\mathbf{x}|\mathbf{s}) = \frac{P(\mathbf{x})P^*(s_1|x_1)P^*(s_2|x_2)P^*(s_3|x_3)}{\sum_{x_1 x_2 x_3} P(\mathbf{x})P^*(s_1|x_1)P^*(s_2|x_2)P^*(s_3|x_3)}$$

The overall strategy of the decision-maker is given by

$$P(a|\mathbf{x}) = \sum_{s_1 s_2 s_3} P^*(s_1|x_1)P^*(s_2|x_2)P^*(s_3|x_3)P^*(a|\mathbf{s})$$

**Bounded rational model by Parpart et al. (2018).** As an alternative model we consider the bounded rational model by Parpart et al. (2018), that is designed for binary comparisons, where in our case the input  $\mathbf{x} \in \{-1, 0, +1\}^3$  indicates whether each cue  $x_j$  with  $j \in \{1, 2, 3\}$  is in favor of class 1 ( $x_j = +1$ ), in favor of class 2 ( $x_j = -1$ ) or neither ( $x_j = 0$ ) and where the output  $y \in \{-1, +1\}$  indicates whether the stimulus  $\mathbf{x}$  should be assigned to class 1 ( $y = +1$ ) or class 2 ( $y = -1$ ). Due to the feature frequencies in our stimulus set, the presence of the first feature favors class 1, the presence of the second feature is irrelevant, and the presence of the third feature favors class 2. The first stimulus pattern with all cues present would then be expressed as  $\mathbf{x} = (+1, 0, -1)$ .

For each feature dimension  $j$ , the model assumes a different linear mapping between feature vector and outputs

$$y = \mathbf{W}^j \mathbf{x} + W_0^j$$

where each of the three different weight vectors  $\mathbf{W}^j$  represents a different linear mapping and  $W_0^j$  is an offset that we added to the model to improve classification performance. The weight vectors are determined by subjects' responses  $y_i$  toward stimulus  $\mathbf{x}_i$  in trial  $i$  through ridge regression

$$\hat{\mathbf{W}}^j = (\Lambda^j + X^T X)^{-1} X^T \mathbf{y}$$

$$W_0^j = \mu_y - \hat{\mathbf{W}}^j \mu_{\mathbf{x}}$$

where  $\mathbf{y} = (y_1 - \mu_y, \dots, y_N - \mu_y)$  is the response vector over all trials  $N$  and  $X^T = (\mathbf{x}_1 - \mu_{\mathbf{x}}, \dots, \mathbf{x}_N - \mu_{\mathbf{x}})$  is a matrix comprising all stimulus vectors with  $\mu_{\mathbf{x}} = \frac{1}{N} \sum_i \mathbf{x}_i$  and  $\mu_y = \frac{1}{N} \sum_i y_i$ , and

$$\Lambda^1 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \frac{\sigma^2}{\eta^2} & 0 \\ 0 & 0 & \frac{\sigma^2}{\eta^2} \end{pmatrix} \Lambda^2 = \begin{pmatrix} \frac{\sigma^2}{\eta^2} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \frac{\sigma^2}{\eta^2} \end{pmatrix} \Lambda^3 = \begin{pmatrix} \frac{\sigma^2}{\eta^2} & 0 & 0 \\ 0 & \frac{\sigma^2}{\eta^2} & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

express the weight that is put on the prior. For  $\eta \rightarrow \infty$  the prior is ignored and all three regression weights are identical. For  $\eta \rightarrow 0$  each regressor ignores all cue dimensions but one. Each regressor can be seen as a decision hyperplane where  $\text{sign}(\hat{\mathbf{W}}^j \mathbf{x} + W_0^j)$  determines the class membership  $\pm 1$ . To bring the three regressors together for a single action choice, the decision hyperplane is determined which induces the maximum distance

$$j^* = \arg \max_j |\hat{\mathbf{W}}^j \mathbf{x} + W_0^j|$$

from the stimulus  $\mathbf{x}$ . The class selection is then determined by

$$a = \text{sign}(\hat{\mathbf{W}}^{j^*} \mathbf{x} + W_0^{j^*})$$

For  $\eta \rightarrow \infty$  this corresponds to a standard regression of the decision hyperplane. For  $\eta \rightarrow 0$  each stimulus is classified by

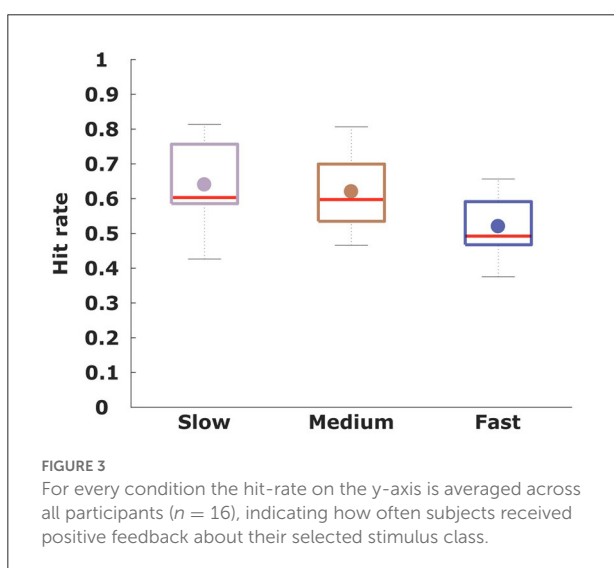
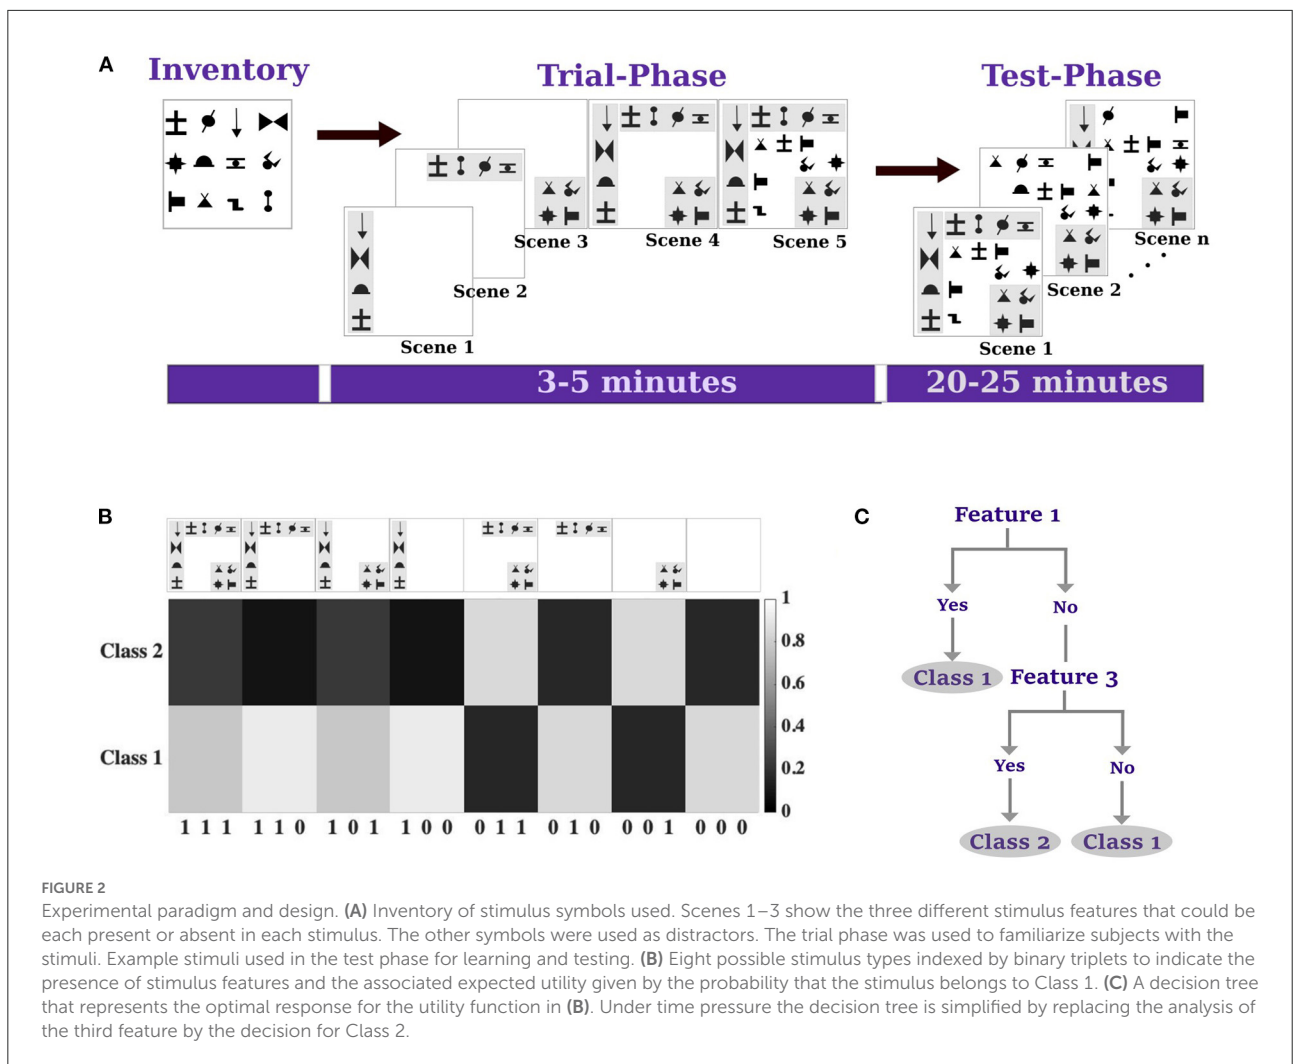
the regressor with maximum distance to one of the hyperplanes that ignore all cue dimensions but one, reminiscent of the Take-the-best heuristic.

### 3. Results

In the familiarization phase of the experiment, human subjects were first acquainted with the possible symbol inventory that was used to compose stimuli (Orbán et al., 2008). In particular, subjects were shown three different stimulus combinations—shown in Scene 1, Scene 2, and Scene 3 in Figure 2A—that would serve as three independent binary features that could each be present or absent in a stimulus buried among a set of meaningless distractor symbols, as shown exemplarily in Scene 5 of Figure 2A. Thus, based on all combinations of binary features, there were eight different kinds of stimuli (disregarding distractors), which we encode in the following by a triplet like 110, meaning feature 1 and 2 are present in the stimulus, while the third feature is absent. A stimulus class is defined by a distribution over the three features. Setting the three features to be independent, this distribution can be defined by three values, e.g., (0.616, 0.089, 0.680) meaning that for stimuli in class 1 the first feature is present in 61.6% of stimuli, the second feature with 8.9%, and so on—compare Figure S1. This way, we can define two different stimulus classes with two different sets of probabilities regarding the presence or absence of the three features and we can determine for each possible kind of stimulus the probability of belonging to class 1 or class 2, respectively—(see Figure 2B). Subjects were then trained on these probabilities by providing them with binary feedback when letting them categorize randomly composed stimuli into the two classes the optimal categorization can be represented by a decision tree - (see Figure 2C). Finally, subjects' performance was tested with the same paradigm after training with three different reaction time conditions ranging from fast, to medium, to slow.

### Behavioral analysis

To assess subjects' performance under the three time conditions, we investigated the overall hit rate in each condition averaged across all subjects, reflecting how often subjects experienced positive feedback for their choices. As the stimulus features are only stochastically associated with each class, a perfect hit rate is impossible with a theoretical maximum expectation of 0.80. As shown in Figure 3, the hit rate increased when subjects were given more deliberation time. The performance as measured by the hit rate was  $0.64 \pm 0.09$  in the slow condition,  $0.62 \pm 0.08$  in the medium condition and  $0.52 \pm 0.07$  in the fast condition. Paired  $t$ -tests revealed


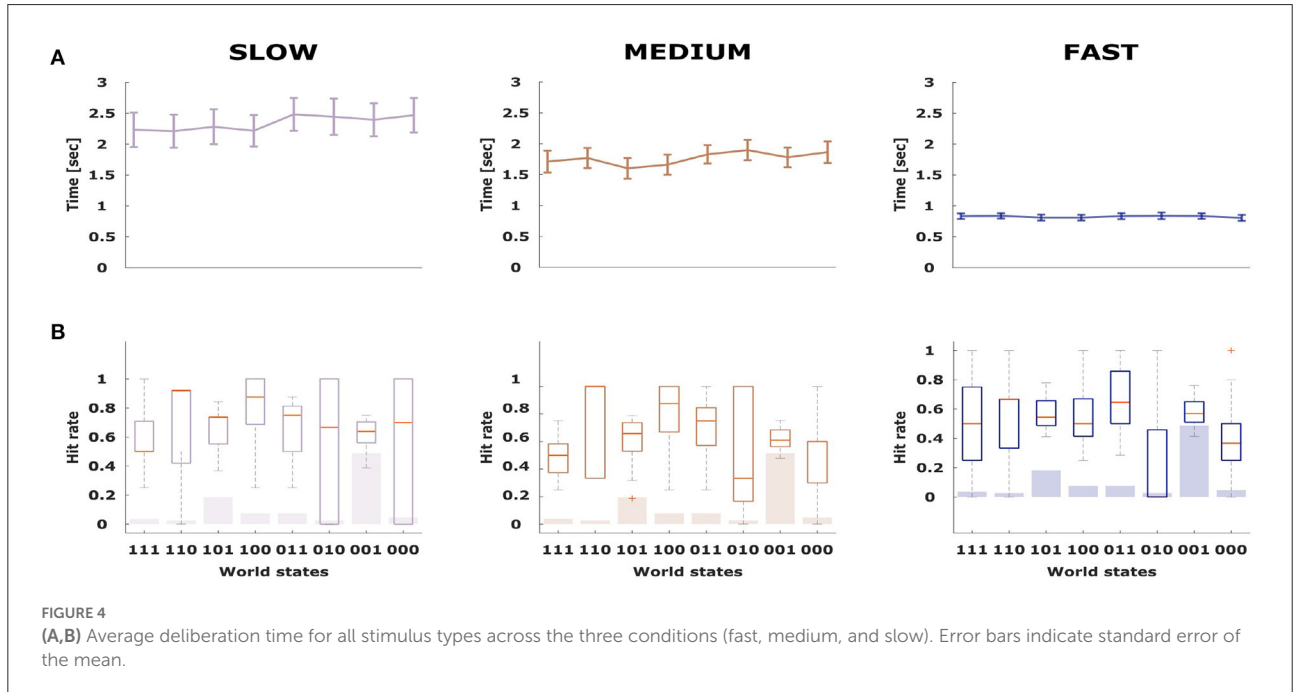


significant differences between the mean hit rate of the fast condition compared to the slow ( $p < 0.001$ ,  $t = 4.08$ ,  $df = 15$ ) as well as the medium condition ( $p < 0.001$ ,  $t = 4.08$ ,  $df = 15$ ). The difference in hit rate between the medium and slow condition was not significant ( $p = 0.2$ ,  $t = 1.26$ ,  $df = 15$ ). To investigate whether the hit-rate was particularly increased or decreased for specific stimulus types, we display the stimulus-specific hit rates in Figure 4 overlaid onto the frequency of the different stimuli. The top of the figure shows the corresponding average deliberation time taken by subjects when making their decisions. In the slow condition the deliberation time saturated at around 2.3 s. In the fast and medium condition, the average deliberation time is reduced to around 1.8 and 0.8 s, respectively.

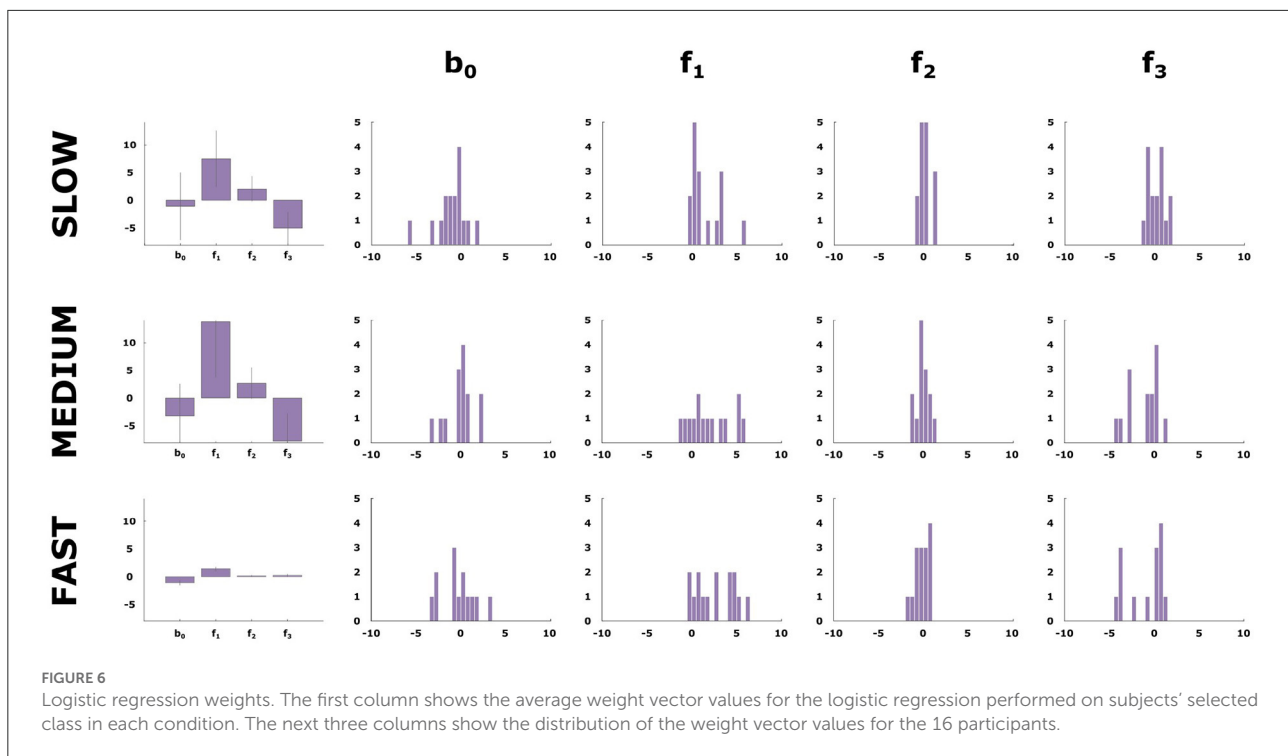
To compare subjects' performance to ideal behavior that always selects the most probable stimulus class in line with Figure 2B, we determined subjects' selection accuracy in Figure 5 for each possible kind of stimulus separately. For an ideal actor

it is therefore possible to achieve 100% accuracy (associated with selecting the most probably correct answer all the time), even though the hit rate is only 80% (because of randomness of

features). With the exception of the (000)- and (010)-stimulus, all other world states show performance well above the chance line. The median value for the selection accuracy is highest







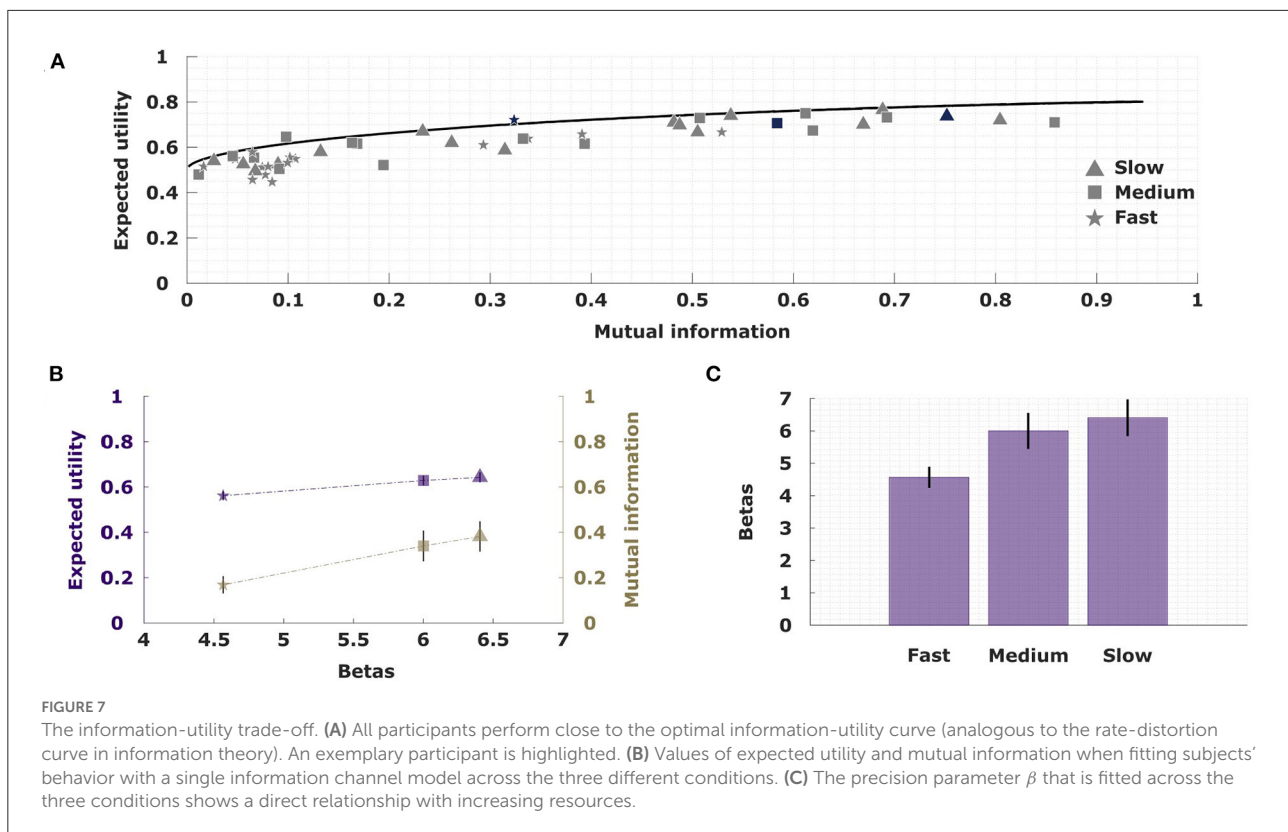
for the (100)-stimulus ( $0.75 \pm 0.05$ ), the (111)- and the (001)-stimulus with values of  $0.74 \pm 0.05$  and  $0.73 \pm 0.04$ , respectively. As evident in Figure 5B, with more time there is generally an increase in the mean value of selection accuracy between the three conditions. When performing a repeated measures ANOVA, we find that both factors stimulus ( $DF = 7$ ,  $F = 5.5$ ,  $p < 0.001$ ) and time condition ( $DF = 2$ ,  $F = 12.02$ ,  $p < 0.001$ ) are significant for subjects' selection accuracy. The only exception to this pattern is the (111)-stimulus, despite being associated with a high selection accuracy.

Finally, we assess the relative weight subjects give to the three different stimulus features when making their decisions under the three different time conditions. To this end, we carried out a logistic regression relating subjects' choices to the three different feature dimensions. In Figure 6, it can be clearly seen that there is a considerable dependence on the first feature across all conditions. In contrast, the value of the weight coefficient of the third feature changes significantly from the fast to the slow condition. The presence of this feature corresponds to a higher likelihood for class 2 and it is only with more time resources that the subjects successfully can take this stimulus feature into account. The histogram reveals the distribution of weights across subjects. This shows a consistent strategy change for most subjects with a decision strategy that concentrates on a single feature in the fast condition, something similar to a take-the-best heuristic, and the integration of multiple features in the slow condition.

## Bounded rational modeling

Similar to the rate-distortion curve in information theory (Cover and Thomas, 2006), bounded rational decision-making models with a single information constraint allow us to quantify the maximally achievable utility by a decision-maker that has a certain amount of information at their disposal. Here, information measures the extra amount of bits available to encode a posterior class selection strategy that takes into account a particular stimulus pattern compared to a prior class selection strategy that does not know about the stimulus (for example a completely uninformed uniform choice probability, or an optimal prior that is tuned to the statistics of stimulus occurrence). Naturally, the more information-processing capability is available, the higher the maximally achievable utility. This optimal information-utility trade-off is illustrated for our task in Figure 7 assuming a single information channel with an optimally tuned prior. In order to achieve maximum utility with perfection requires a little bit less than one bit of information for binary selection, as the optimal prior strategy is non-uniform due to the non-uniform occurrence probabilities of the stimulus pattern. The non-uniformity of the prior is also evident in the other extreme of zero information, where the expected utility is slightly above the chance level of 0.5.

We can now assess subjects' behavior by measuring the average utility they achieved for each condition and by determining their mutual information between stimulus pattern



and class selection based on their empirical choice frequencies. As can be seen in Figure 7A, subjects' performance data lie below the efficiency frontier as expected, but generally close to it. Due to statistical fluctuations a few data points are slightly above. Subjects' data for the fast condition clusters mainly in the low information part of the diagram, whereas performance data of the medium and slow condition are stretched along the efficiency frontier for higher information values. The efficiency frontier is parameterized by a scalar precision parameter, so that we can fit an average precision parameter across all subjects for each condition in a way that minimizes the distance between the data points and the curve. In Figures 7B,C it can be seen that the fitted precision parameter increases across conditions from fast to slow, and that the average utility achieved by subjects increases together with the average precision parameter.

Depending on the available information-processing resources, the bounded rationality model predicts that the decision-maker's choice strategy should become sharper and more concentrated the higher the available precision is. The choice strategy of the bounded optimal decision-maker more and more optimizes the utility function shown in Figure 2B with increasing time resources. Figure 8C shows the average theoretical posterior of the bounded rational model across subjects fitted by the information constraint determined by the maximum likelihood precision parameter  $\beta$  for each condition and subject. The color code indicates the probability

of choosing class 1 or class 2 for each stimulus pattern, and therefore represents a conditional distribution over the class selection variable. Figure 8A shows the average empirical choice strategy formed by subjects across the three time conditions for comparison. The corresponding best responses by the Bayes-optimal model are depicted in Figure 8B and are deterministic and invariant across conditions since they ignore time constraints. We also compare to the average response probabilities of a logistic regression of subjects' responses in Figure 8E where we decrease the number of free parameters across conditions: four free parameters in the slow condition, three free parameters in the medium condition, and two free parameters in the fast condition. Finally, we compare to an alternative regression-based bounded rationality model (COR, Covariance Orthogonalizing Regularization) by Parpart et al. (2018) in Figure 8F. The parameter fits for all models were obtained by maximum likelihood. Table 1 shows the grand average across subjects for each model, the corresponding values for each subject and model for the different time conditions can be found in the Supplementary material in Tables S1–S3.

In the slow condition, the predicted choice probabilities of all models match subjects' data quite well. Compared to the deterministic Bayes-optimal model and the COR model, the bounded rationality model also captures the stochasticity in the choice, including the increase in stochasticity in the medium condition. In the fast condition, the predicted choice

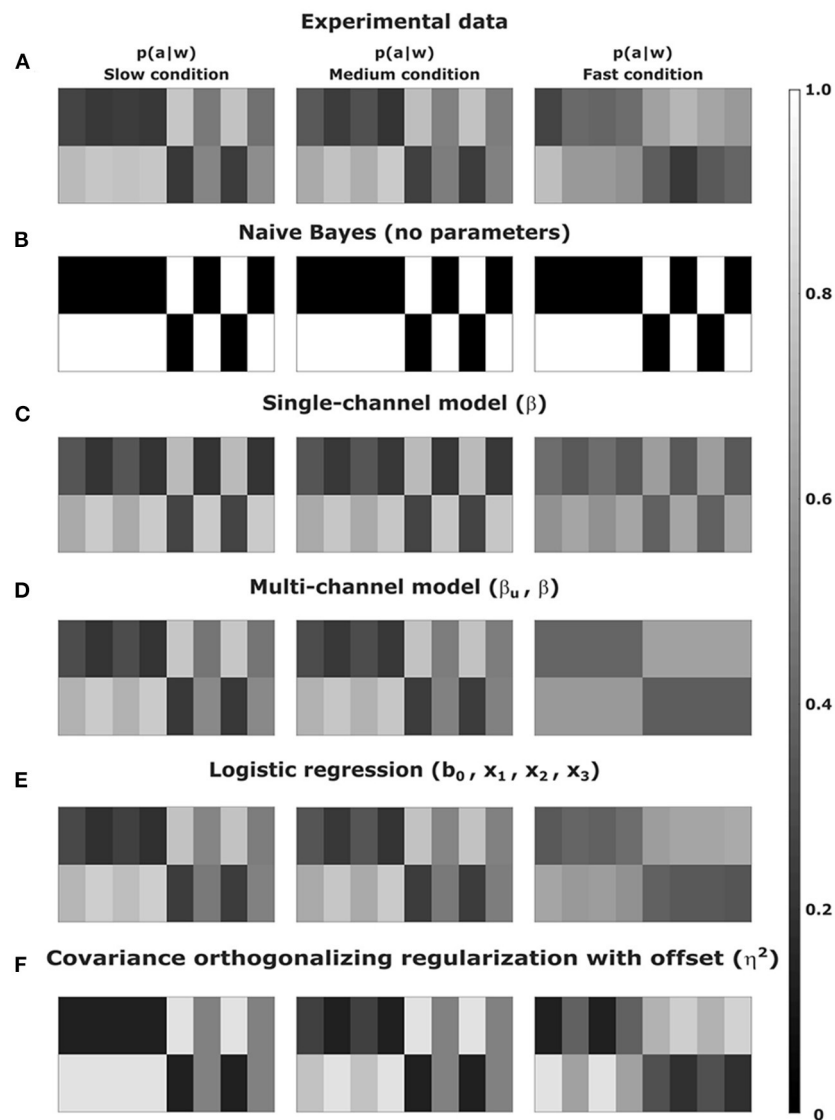


FIGURE 8

Averaged selection strategies recorded experimentally (A) and compared to averaged theoretical selection strategies predicted by a naive Bayes-optimal model (B), a single information channel model (C) and a bounded rational model with multiple feature-specific information channels (D), a logistic regression (E), and the bounded rational regression model (F) by Parpart et al. (2018).

probabilities do not match subjects' data so well anymore. As can be seen in the rightmost panel of Figure 8A, subjects' responses tend to fall into two blocks, where one half of the stimulus space is classified predominantly as class 1, and the other half of the stimuli is classified as class 2, mostly. This can be interpreted as a footprint of a decision heuristic that coincides with the simplified decision tree in Figure 2C. However, this bipartite response profile is not reproduced by the naive Bayes model, the COR model or the bounded rational model with a single information constraint, although it is clearly reflected in the logistic regression fit depicted in the rightmost panel of Figure 8E. As suggested by the weights of the logistic regression

shown in Figure 6, this change in response profile is due to a change in feature selection.

To compare how much weight the different models put on the different feature dimensions of the stimulus with different levels of available information processing capability, we determine in Figure 9 the mutual information between each stimulus feature and the choices subjects made. In line with the logistic regression analysis, we see in Figure 9A that subjects put most weight on the first stimulus feature, largely ignore the second feature, and modulate the importance of the third feature down to the level of the second feature in case of the fast condition with little information resources.

TABLE 1 Mean model parameters averaged across subjects for the three conditions slow, medium, and fast.

Model		Slow	Medium	Fast
sCh	$\beta$	1.8823 $\pm$ 1.3872	1.7623 $\pm$ 1.3214	0.7840 $\pm$ 0.7357
	$\beta_u$	3.3158 $\pm$ 1.5922	3.3026 $\pm$ 1.5403	2.7500 $\pm$ 1.5655
mCh	$\beta$	45.4135 $\pm$ 40.2344	45.0724 $\pm$ 36.6945	19.2679 $\pm$ 16.3503
	$b_0$	-1.3514 $\pm$ 22.3324	-2.9647 $\pm$ 21.5730	-1.0700 $\pm$ 1.7163
logReg	$x_0$	7.3524 $\pm$ 19.0123	13.3082 $\pm$ 37.3361	1.4091 $\pm$ 1.7384
	$x_1$	1.9988 $\pm$ 8.5445	2.6053 $\pm$ 10.6771	0.1400 $\pm$ 0.6096
	$x_2$	-4.5859 $\pm$ 10.2636	-7.4739 $\pm$ 18.4158	0.2378 $\pm$ 0.9275
COR	$\eta^2$	0.0017 $\pm$ 0.0062	0.0046 $\pm$ 0.0085	0.0019 $\pm$ 0.0060

Qualitatively, none of the models is able to reproduce this modulation. The naive Bayes model in Figure 9B gives the highest weight to the first feature and a medium weight to the third feature, but the absolute value of the mutual information is too high due to the lack of stochasticity in the response, and more importantly, this profile does not change across conditions, because the model is blind to resource constraints. The COR model in Figure 9F shows a similar response profile that also does not modulate across conditions, because in our experiment there are no correlations between features that the model could downweight. The bounded rational model with information constraint in Figure 9C downregulates the mutual information between actions and features when the task becomes more demanding, however, it does so for all features equally, without showing the observed concentration on the first stimulus features.

The finding that the simple bounded rationality model in Figure 9C fails to capture the transition in the fast condition that is dominated by a single feature may not be too surprising, as the model only allows for a single information channel that cannot distinguish between the different feature dimensions, and therefore lacks internal structure (Luce, 2003). Hence, we also study an extension of the bounded rationality model where we introduce three independent perceptual channels corresponding to the three feature dimensions that are then integrated by the information channel that corresponds to the decision-making process—see section Methods for details. The assumption of a structured internal sensory state with its own precision values allows the model to prioritize between different perceptual states based on their feature grouping and their impact on choice utility, even though *a priori* all feature channels have the same precision. In Figure 8D it can be seen, that, unlike the other models, this extended bounded rational model with feature groupings is able to capture the transition to a bipartite response profile. Moreover, Figure 9D shows that the extended bounded rational model achieves this response profile by downregulating its mutual information between actions and stimulus features the same way that

subjects do, namely by focusing almost exclusively on the first stimulus feature.

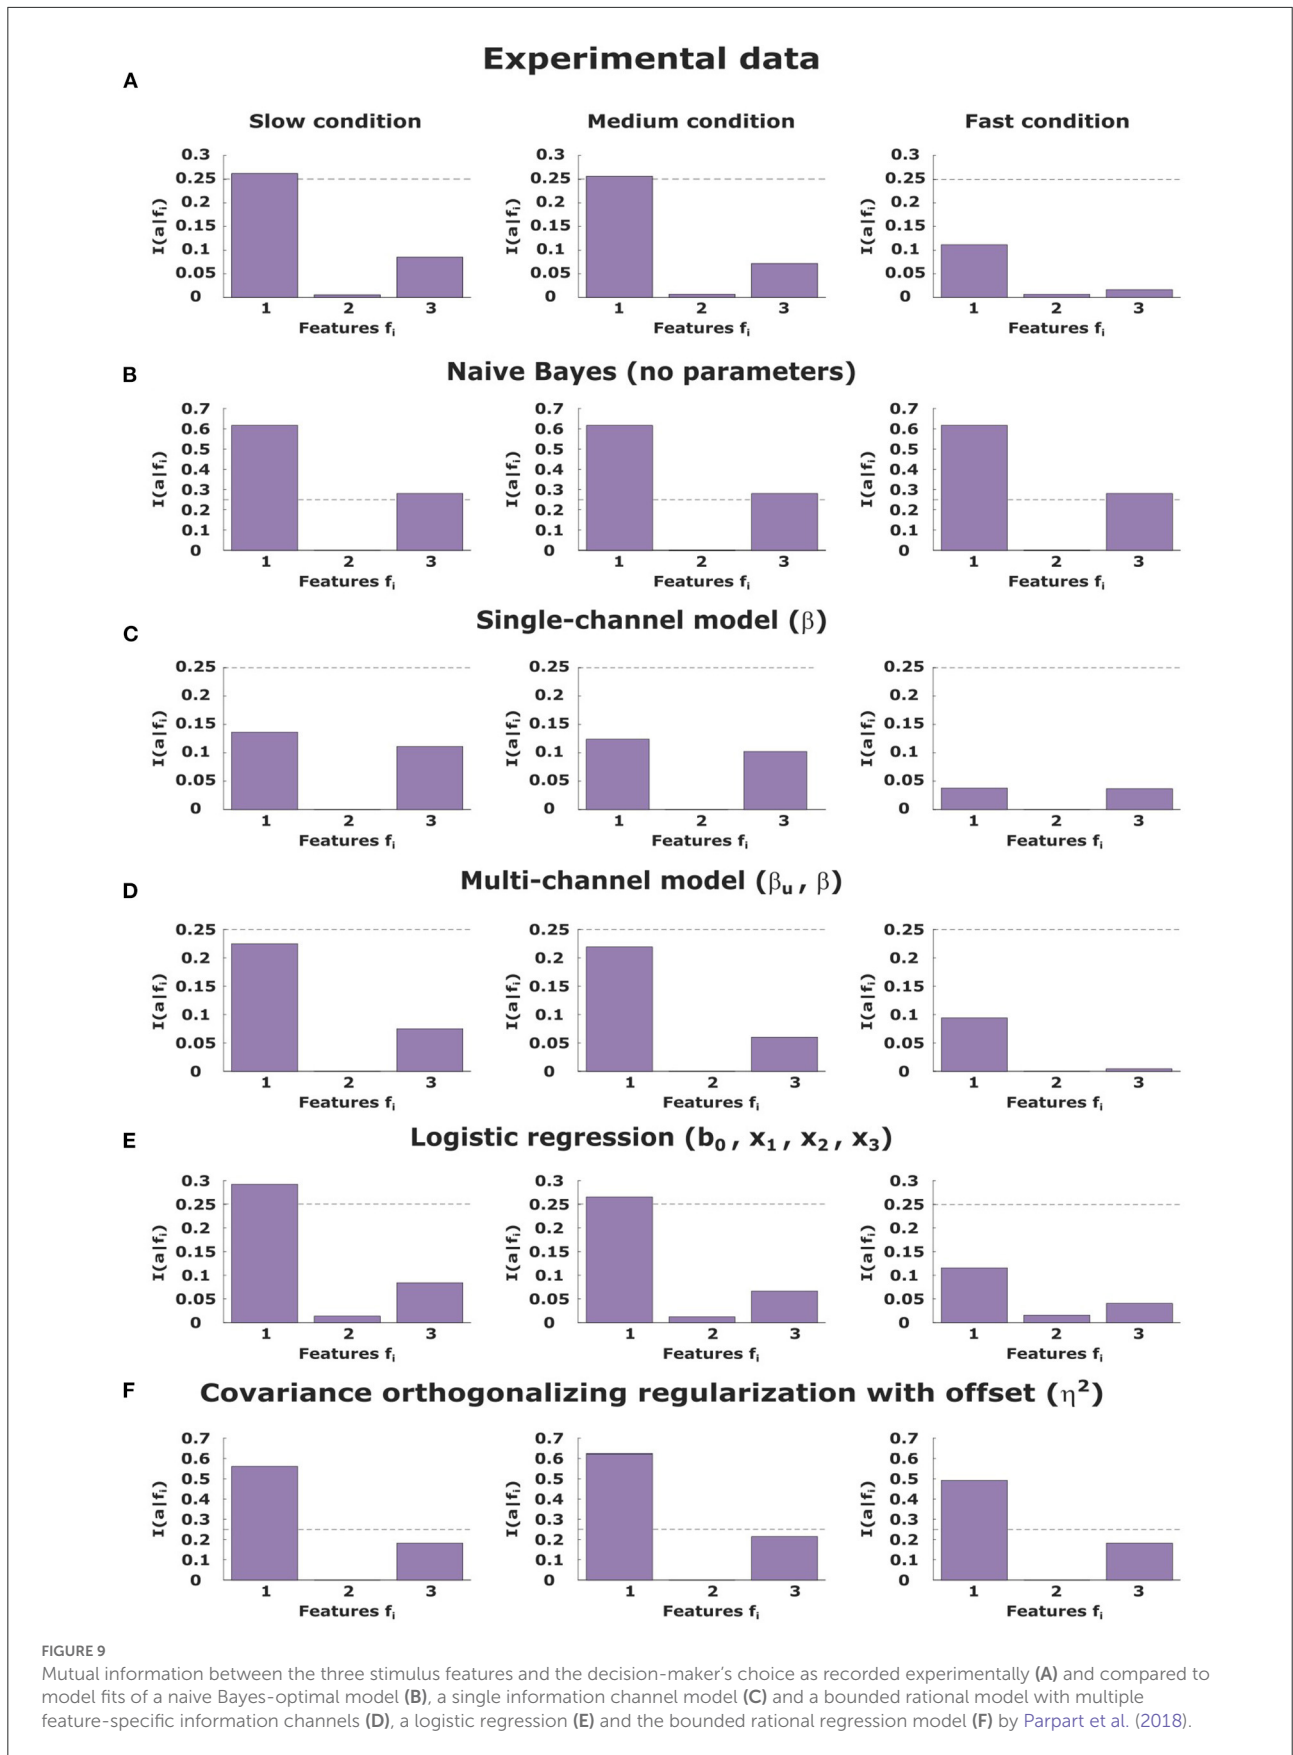
## Model comparison

To compare the different model fits quantitatively, we take a three-pronged approach: (i) we compare the likelihood of subjects' choices under the probability distributions induced by the various models, (ii) we compare the similarity of the model posteriors to the experimental posterior, (iii) we compare the similarity of the models' mutual information profiles to the experimental profile. The comparison of the mutual information profile is particularly relevant to reveal the models' ability in adapting the feature selectivity across conditions. As a similarity measure we use the Euclidean norm.

To compare the likelihood of subjects' behavior under the different models, we perform a 1,000-fold cross validation, where in each fold we use a random partition corresponding to half the data for fitting and the other half for evaluation of the likelihood from the predicted probabilities of subjects' choices. The cross-validated likelihoods for each subject and condition are shown in Tables S4–S6. The predicted likelihood of subjects' behavior under the multi-channel model achieves the highest average likelihood of all models. When performing a pairwise comparison, the multi-channel model predicts subjects' choices significantly better than the naive Bayes model (paired *t*-test,  $p < 0.001$ ), the COR model (paired *t*-test,  $p < 0.001$ ), the logistic regression model (paired *t*-test,  $p < 0.001$ ), and the single-channel model (paired *t*-test,  $p = 0.028$ ).

To compare the similarity between the models' posterior and subjects' empirical distribution, we compute the Euclidean norm between the corresponding conditional probability tables and the normalized histograms—compare Tables S7–S9. We find that the multi-channel model has the lowest norm on average. When performing a pairwise comparison, we find that the multi-channel model predicts subjects' conditional choice distribution significantly better than the naive Bayes model (paired *t*-test,  $p < 0.001$ ), the COR model (paired *t*-test,  $p < 0.001$ ), the logistic regression model (paired *t*-test,  $p < 0.001$ ), and the single-channel model (paired *t*-test,  $p = 0.0207$ ).

Finally, we compare how well the models can reproduce the mutual information profile in each condition exhibited by subjects in Figure 9A, by computing the Euclidean distance between the mutual information estimated from empirical choice frequencies and the mutual information profiles computed from model distributions—compare Tables S10–S12. We find that the multi-channel model predicts the best fitting information profile across all conditions compared to any other model (paired *t*-test, pairwise comparison to naive Bayes  $p < 0.001$ ), the COR model (paired *t*-test,  $p < 0.001$ ), the logistic regression model (paired *t*-test,  $p < 0.001$ ), and the single-channel model (paired *t*-test,  $p = 0.0308$ ). More



importantly, we find in a separate simulation that merely tries to reproduce the mutual information profile, that the single channel model is unable to generate a profile that focuses on the first stimulus feature in the fast condition, as both the first and the third feature always receive similar weights—compare [Figure S2](#). This suggests that the single-channel bounded rationality model lacks the necessary internal structure to reproduce the adaptive concentration on a single stimulus feature ([Luce, 2003](#)), and that a bounded rationality model where perceptual states are grouped according to the presented features provides overall the best explanation of our data.

## 4. Discussion

In this study, we set out to investigate in how far bounded rationality models based on information constraints may provide a bridge between heuristic and Bayes-optimal approaches to decision making in a binary classification task. We collected behavioral data in order to explore the change in strategy that human decision-makers exhibit when taking decisions under varying time pressure. Logistic regression performed on subjects' responses across trials depending on the presented stimulus features reveals a clear pattern for the weight allocated by subjects to each of the features. The dominant weight given to the first feature in the fast condition indicates a focus on a single criterion reminiscent of the take-the-best heuristic ([Brusovansky et al., 2018](#)). The broadening of the weights with respect to the other two features, as decision time increases in the medium and slow conditions, suggests that with the introduction of more resources, subjects start to integrate also less informative features to improve decision quality, in line with Bayes-optimal models of information integration ([Ernst and Banks, 2002](#); [Körding and Wolpert, 2004](#)). Decision-makers were, thus, not only able to extract and integrate different stimulus features, but were also able to do so in an adaptive manner across conditions ([Payne et al., 1988](#)).

The bounded rationality model that we investigate captures the change in strategy for different resource conditions with different information constraints that quantify how much a decision-maker can deviate from their prior choice strategy and how much sensory information they can process from the different stimulus features ([Ortega and Braun, 2013](#)). Depending on the information constraints, the model therefore provides a normative standard with which we can compare subjects' performances under different resource conditions ([Schach et al., 2018](#); [Lindig-León et al., 2019](#)), as shown in [Figure 7](#). When removing the information processing constraints, the model predictions converge to the deterministic Bayes-optimal decision rule maximizing the expected utility shown in [Figure 2B](#). As predicted by the bounded rational model, over the three time conditions from fast to slow the information divergence measures of subjects' strategy distributions with

respect to their prior strategies increase. Furthermore, subjects' average strategies resemble those exhibited by the model under different information constraints in that they become more deterministic in the slow condition, with higher information divergence from the prior, higher utility (lower error rate) and broader weight integration compared to the fast condition. Such a strategy change can of course not be captured by a Bayes-optimal model, like the Naive Bayes model that we evaluated, that does not take into account the available resources.

In the psychology literature, strategy changes under resource limitations have been studied extensively ([Payne et al., 1988, 1993](#)), in particular speed accuracy trade-offs that quantify the relationship between task-related error-rates and the time available to perform the task ([Heitz, 2014](#)). For example, [Rieskamp and Hoffrage \(2008\)](#) have found in a multi-attribute inference task that increasing time pressure led subjects to a strategy shift from linear additive attribute weighting to a lexicographic heuristic, irrespective of whether time pressure is induced by limiting response times or increasing opportunity costs. A recent study ([Oh et al., 2016](#)), very similar to ours, has investigated the strategic shift in a multi-cue classification task with abstract symbol features under varying conditions of time pressure. The task was to choose from a set of two stimuli, the stimulus with maximal winning probability under varying time constraints ranging between 2 s and 700 ms. The winning probability could be determined by integrating over the different symbol features, each of which was associated with a probability of being a winning stimulus. In conditions with little time pressure, the study found that subjects integrated information from multiple features, and that with more time pressure some of the features were ignored, suggesting some kind of cue discounting. The study also found that the observed behavior was not consistent with the take-the-best heuristic, but rather some kind of drop-the-worst heuristic. Such a discounting of less informative cues under time pressure has also been reported recently in a binary classification task modeled with reinforcement learning ([Wang and Rehder, 2017](#)) and a recent fMRI study ([Oh-Descher et al., 2017](#)) that investigated a shift from cortical to sub-cortical activity with increasing time pressure.

The gradual cue discounting observed under time pressure in previous studies is entirely compatible with our model—compare for example [Figure S3](#) where we apply the multi-channel bounded rationality model to the experiment by Oh et al. with three features ([Oh et al., 2016](#)). The model generally predicts that decision-makers should weight the cues depending on their predictive value and that the relative weight of the less predictive cues decreases faster under information constraints. This is consistent with the notion that less informative cues are discounted when time pressure is ramped up. By varying information constraints in a continuous fashion, the bounded rational model is able to describe both gradual transitions in information processing strategies and sharp phase transitions

depending on the defined task utilities. Therefore, the prediction of a behavior consistent with a take-the-best heuristic in our task is not intrinsic to the model, but a consequence of applying the model to the particular task. Also in our task, the model does not predict that the less informative features have weight exactly zero, only that the first feature is domineering in the low information regime.

There have been a number of previous studies relating heuristic decision-making rules to utility-based decision models in the context of making choices between gambles. [Olschewski and Rieskamp \(2021\)](#), for example, modeled human decision making under time pressure both with simple heuristics and a classical random utility model. They found that under time pressure decision-makers were more likely to choose risky gambles and to follow more non-compensatory decision strategies. However, their risk model is not applicable to the current task, which is a simple classification task and does not involve the comparison of gambles with different degrees of riskiness. Without risk, the random utility model with a logit link function can be seen as a special case of a single channel bounded rationality model with uniform prior, which does not capture subjects' behavior very well. In another study, [Pachur et al. \(2017\)](#) varied the parameters of a prospect theory model to emulate a number of heuristics for choosing between gambles, including minimax, maximax, and the priority heuristic. Again, our classification task does not involve the comparison of risky gambles, so probability and value distortions cannot reverse the choice between the two classes, only increase or decrease their utility difference, which is the same as a logit model with uniform prior and varying precision parameter. Nevertheless, in more complex decision scenarios the information-theoretic bounded rationality model can also implement minimax and maximax strategies ([Ortega and Braun, 2013](#)), while the relation to more elaborate heuristics remains to be explored. Finally, [Loewenstein et al. \(2015\)](#) has investigated the implementation of heuristic choice strategies with accumulator networks, which provides a possible mechanistic explanation for heuristic and bounded rational choices.

[Parpart et al. \(2018\)](#) have previously presented a framework that integrates classical Bayesian rational and boundedly rational decision-making by weighting prior and likelihood in a regression scenario. When the prior is ignored, the resulting logistic regression can take into account correlations between different cues. In the other extreme, each cue dimension is regressed separately and the single most predictive dimension is used for classification as in heuristic decision making. While their model can be applied to our classification task, all cue dimensions in our task are independent, and thus, ignoring cue correlations has no effect. Consequently, the model cannot explain the strategy shifts that we observed under varying time conditions. Previous applications of bounded rationality models with information constraints in tasks with time pressure include sensorimotor tasks like simple reaching movements with

varying time pressure for motor planning affecting endpoint precision ([Schach et al., 2018](#)), psychophysical tasks like absolute identification of geometrical shapes on different levels of granularity ([Lindig-León et al., 2019](#)), and cognitive puzzles like checking satisfiability of conjunctive normal forms with varying deliberation time ([Ortega and Stocker, 2016](#)). Compared to these previous studies with information constraints, our current study is the first to apply different information-theoretic bounded rationality models to a multi-attribute decision-making task.

Bounded rationality models come in many different flavors and have a long history going back at least to Herbert Simon's ideas ([Simon, 1979](#)) that human decision-makers cannot engage in full cost-benefit analyses when facing complex problems due to their limited cognitive abilities, and that accordingly humans have to exploit heuristics adapted to the structure of their environment in order to find satisfying solutions, that is solutions that are good enough, but not necessarily optimal in any sense. Particularly in psychology, this has led to some researchers completely abandoning the notion of optimality by focusing on process-models of decision-making, that are heuristics used in solving complex practical problems ([Gigerenzer and Selten, 2002](#)). While such rules-of-thumb may be outperformed in controlled environments, it has been repeatedly argued that heuristics can perform better than computationally expensive algorithms based on optimality principles, in particular when there is high model uncertainty ([Gonzalez, 2004](#); [Buckmann and Şimşek, 2017](#)). However, often the question remains how these heuristics are learned, where they come from and whether there are any underlying mathematical principles. The heuristics-and-biases approach to bounded rationality ([Tversky and Kahneman, 1974](#); [Taniguchi et al., 2016](#); [Itri and Patel, 2018](#)), for example, that has focused on the role of cognitive biases in human decision-making has made attempts to formalize biased decision-making procedures using the traditional concepts from economics, however, without normative cogency.

Bounded rationality based on mathematical principles often takes the shape of constrained optimization models in the economics and artificial intelligence literature ([Horvitz et al., 1988](#)), where the central idea is that bounded rational behavior can be still regarded as optimal when limitations are taken into account by the appropriate mathematical constraints. Often such models are portrayed as fundamentally irreconcilable with a heuristics approach to bounded rationality, as the remaining notion of optimality apparently introduces an infinite regress: decision-makers that are unable to solve the original utility optimization problem have to solve instead a more complicated optimization problem ([Zilberstein, 2008](#)). However, this is only a problem when the bounded rational model is interpreted as a process model, where the decision-maker is optimizing by calculating. When the bounded rationality model only serves as a description for an observer, the model simply allows to compare behavior to a normative baseline.

For example, for the bounded rationality model we use in this paper, we could imagine a simple satisficing agent that probabilistically considers random samples with respect to a certain aspiration level only dealing with utility to the best of its abilities, but for an observer it may appear that this decision-making system is trading off utility and information (Ortega and Braun, 2014). Accordingly, the bounded rational model in this paper is still somewhat descriptive and does not provide a detailed heuristic process model, rather it predicts a strategy distribution whose performance is compatible with such a heuristic in our simple task. How this generalizes to more complex tasks is an interesting area for further study.

## Data availability statement

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

## Ethics statement

The studies involving human participants were reviewed and approved by the Ethics Committee of Ulm University. The patients/participants provided their written informed consent to participate in this study.

## Author contributions

NK, CL-L, and DB conceived the experiment. NK conducted the experiment. NK and CL-L analyzed the data. DB developed

the model. NK and DB wrote the paper. All authors reviewed the manuscript.

## Funding

This study was funded by European Research Council (ERC-StG-2015-ERC Starting Grant, Project ID: 678082, BRISC: Bounded Rationality in Sensorimotor Coordination).

## Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

## Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fnins.2022.906198/full#supplementary-material>

## References

- Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine. *Econometrica* 21, 503–546. doi: 10.2307/1907921
- Belousov, B., Neumann, G., Rothkopf, C. A., and Peters, J. R. (2016). "Catching heuristics are optimal control policies," in *Advances in Neural Information Processing Systems*, Vol. 29, eds D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (Barcelona).
- Braun, D. A., Aertsen, A., Wolpert, D. M., and Mehring, C. (2009). Learning optimal adaptation strategies in unpredictable motor tasks. *J. Neurosci.* 29, 6472–6478. doi: 10.1523/JNEUROSCI.3075-08.2009
- Braun, D. A., Ortega, P. A., Theodorou, E., and Schaal, S. (2011). "Path integral control and bounded rationality," in *2011 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)* (Paris), 202–209. doi: 10.1109/ADPRL.2011.5967366
- Brusovansky, M., Glickman, M., and Usher, M. (2018). Fast and effective: intuitive processes in complex decisions. *Psychon. Bull. Rev.* 25, 1542–1548. doi: 10.3758/s13423-018-1474-1
- Buckmann, M., and Şimşek, Ö. (2017). Decision heuristics for comparison: how good are they? in *Imperfect Decision Makers: Admitting Real-World Rationality*, PMLR, Vol. 58 (Barcelona), 1–11.
- Cover, T. M. and Thomas, J. A. (2006). *Elements of Information Theory*. (Wiley Series in Telecommunications and Signal Processing). New York, NY: Wiley-Interscience.
- Doya, K., Ishii, S., Pouget, A., and Rao, R. P. N. (2007). *Bayesian Brain: Probabilistic Approaches to Neural Coding*. Cambridge, MA: The MIT Press. doi: 10.7551/mitpress/9780262042383.001.0001
- Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *Quart. J. Econ.* 75, 643–669. doi: 10.2307/1884324
- Ernst, M. O. and Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415, 429–33. doi: 10.1038/415429a
- Forbes, W., Hudson, R., Skerratt, L., and Soufian, M. (2015). Which heuristics can aid financial-decision-making? *Int. Rev. Finan. Anal.* 42, 199–210. doi: 10.1016/j.irfa.2015.07.002
- Gigerenzer, G., and Gaissmaier, W. (2011). Heuristic decision making. *Annu. Rev. Psychol.* 62, 451–482. doi: 10.1146/annurev-psych-120709-145346
- Gigerenzer, G., and Goldstein, D. (1996a). Reasoning the fast and frugal way: models of bounded rationality. *Psychol. Rev.* 62, 650–669. doi: 10.1093/acprofoso/9780199744282.003.0002



- Gigerenzer, G., and Goldstein, D. G. (1996b). Reasoning the fast and frugal way: models of bounded rationality. *Psychol. Rev.* 103, 650–669. doi: 10.1037/0033-295x.103.4.650
- Gigerenzer, G., Hertwig, R., and Pachur, T. (2011). *Heuristics*. Oxford: Oxford University Press.
- Gigerenzer, G., and Selten, R. (2002). *Bounded Rationality: The Adaptive Toolbox*. Cambridge, MA; London: MIT Press. doi: 10.7551/mitpress/1654.001.0001
- Goldstein, D., and Gigerenzer, G. (2002). Models of ecological rationality: the recognition heuristic. *Psychol. Rev.* 109, 75–90. doi: 10.1037/0033-295x.109.1.75
- Gonzalez, C. (2004). Learning to make decisions in dynamic environments: effects of time constraints and cognitive abilities. *Hum. Factors* 46, 449–460. doi: 10.1518/hfes.46.3.449.50395
- Gottwald, S., and Braun, D. (2019). Bounded rational decision-making from elementary computations that reduce uncertainty. *Entropy* 21:375. doi: 10.3390/e21040375
- Heitz, R. P. (2014). The speed-accuracy tradeoff: history, physiology, methodology, and behavior. *Front. Neurosci.* 8:150. doi: 10.3389/fnins.2014.00150
- Horvitz, E. J., Breese, J. S., and Henrion, M. (1988). Decision theory in expert systems and artificial intelligence. *Int. J. Approx. Reason.* 2, 247–302.
- Itri, J. N., and Patel, S. H. (2018). Heuristics and cognitive error in medical imaging. *Amer. J. Roentgenol.* 210, 1097–1105. doi: 10.2214/AJR.17.18907
- Jones, M., and Love, B. C. (2011). Bayesian fundamentalism or enlightenment? on the explanatory status and theoretical contributions of Bayesian models of cognition. *Behav. Brain Sci.* 34, 169–188. doi: 10.1017/S0140525X10003134
- Julian, M., and Gerd, G. (2012). Heuristic decision making in medicine. *Dialog. Clin. Neurosci.* 14, 77–89. doi: 10.31887/DCNS.2012.14.1/jmarewski
- Knill, D. C., and Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci.* 27, 712–719. doi: 10.1016/j.tins.2004.10.007
- Körding, K. P., and Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature* 427, 244–247. doi: 10.1038/nature02169
- Körding, K. P., and Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends Cogn. Sci.* 10, 319–326. doi: 10.1016/j.tics.2006.05.003
- Lindig-León, C., Gottwald, S., and Braun, D. A. (2019). Analyzing abstraction and hierarchical decision-making in absolute identification by information-theoretic bounded rationality. *Front. Neurosci.* 13:1230. doi: 10.3389/fnins.2019.01230
- Lindig-León, C., Schmid, G., and Braun, D. A. (2021). Bounded rational response equilibria in human sensorimotor interactions. *Proc. R. Soc. B* 288:20212094. doi: 10.1098/rspb.2021.2094
- Loewenstein, G., O'Donoghue, T., and Bhatia, S. (2015). Modeling the interplay between affect and deliberation. *Decision* 2, 55–81. doi: 10.1037/dec0000029
- Luce, R. D. (2003). Whatever happened to information theory in psychology? *Rev. Gen. Psychol.* 7, 183–188. doi: 10.1037/1089-2680.7.2.183
- Oh, H., Beck, J. M., Zhu, P., Sommer, M. A., Ferrari, S., and Egner, T. (2016). Satisficing in split-second decision making is characterized by strategic cue discounting. *J. Exp. Psychol. Learn. Mem. Cogn.* 42, 1937–1956. doi: 10.1037/xlm0000284
- Oh-Descher, H., Beck, J. M., Ferrari, S., Sommer, M. A., and Egner, T. (2017). Probabilistic inference under time pressure leads to a cortical-to-subcortical shift in decision evidence integration. *Neuroimage* 162, 138–150. doi: 10.1016/j.neuroimage.2017.08.069
- Olschewski, S., and Rieskamp, J. (2021). Distinguishing three effects of time pressure on risk taking: choice consistency, risk preference, and strategy selection. *J. Behav. Decis. Mak.* 34, 541–554. doi: 10.1002/bdm.2228
- Orbán, G., Fiser, J., Aslin, R. N., and Lengyel, M. (2008). Bayesian learning of visual chunks by human observers. *Proc. Natl. Acad. Sci. U.S.A.* 105, 2745–2750. doi: 10.1073/pnas.0708424105
- Ortega, P., and Braun, D. (2014). Generalized thompson sampling for sequential decision-making and causal inference. *Complex Adapt. Syst. Model.* 2, 1–23. doi: 10.1186/2194-3206-2-2
- Ortega, P. A., and Braun, D. A. (2011). “Information, utility and bounded rationality,” in *Artificial General Intelligence. AGI 2011. Lecture Notes in Computer Science*, Vol. 6830, eds J. Schmidhuber, K. R. Thórisson, and M. Looks (Berlin, Heidelberg: Springer). 269–274. doi: 10.1007/978-3-642-22887-2\_28
- Ortega, P. A., and Braun, D. A. (2013). Thermodynamics as a theory of decision-making with information-processing costs. *Proc. R. Soc.* 469, 1–18. doi: 10.1098/rspa.2012.0683
- Ortega, P. A., and Stocker, A. A. (2016). “Human decision-making under limited time,” in *30th Conference on Neural Information Processing Systems (NIPS 2016)* (Barcelona).
- Pachur, T., and Bröder, A. (2013). Judgment: a cognitive processing perspective: judgment: a cognitive processing perspective. *Wiley Interdiscip. Rev. Cogn. Sci.* 4, 665–681. doi: 10.1002/wcs.1259
- Pachur, T., Suter, R. S., and Hertwig, R. (2017). How the twin can meet: prospect theory and models of heuristics in risky choice. *Cognit. Psychol.* 93, 44–73. doi: 10.1016/j.cogpsych.2017.01.001
- Parpart, P., Jones, M., and Love, B. C. (2018). Heuristics as bayesian inference under extreme priors. *Cognit. Psychol.* 102, 127–144. doi: 10.1016/j.cogpsych.2017.11.006
- Payne, J. W., Bettman, J. R., and Johnson, E. J. (1988). Adaptive strategy selection in decision making. *J. Exp. Psychol. Learn. Mem. Cogn.* 14, 534–552. doi: 10.1037/0278-7393.14.3.534
- Payne, J. W., Bettman, J. R., and Johnson, E. J. (1993). *The Adaptive Decision Maker*. Cambridge: Cambridge University Press. doi: 10.1017/CBO9781139173933
- Rieskamp, J., and Hoffrage, U. (2008). Inferences under time pressure: how opportunity costs affect strategy selection. *Acta Psychol.* 127, 258–276. doi: 10.1016/j.actpsy.2007.05.004
- Savage, L. J. (1954). *The Foundations of Statistics*. New York, NY: Wiley Publications in Statistics.
- Schach, S., Gottwald, S., and Braun, D. A. (2018). Quantifying motor task performance by bounded rational decision theory. *Front. Neurosci.* 12:932. doi: 10.3389/fnins.2018.00932
- Simon, H. A. (1979). Rational decision making in business organizations. *Amer. Econ. Rev.* 69, 493–513.
- Svenson, O., and Edland, A. (1987). Change of preferences under time pressure. *Scand. J. Psychol.* 28, 322–330. doi: 10.1111/j.1467-9450.1987.tb00769.x
- Taniguchi, H., Shirakawa, T., and Takahashi, T. (2016). “Implementation of human cognitive bias on naïve bayes,” in *Proceedings of the 9th EAI International Conference on Bio-Inspired Information and Communications Technologies (Formerly BIONETICS)* (New York, NY: ACM), 483–489. doi: 10.4108/eai.3-12-2015.2262494
- Tishby, N., and Polani, D. (2011). “Information theory of decisions and actions,” in *Perception-Action Cycle*, eds V. Cutsuridis, A. Hussain, and J. G. Taylor (New York, NY: Springer New York), 601–636. doi: 10.1007/978-1-4419-1452-1\_19
- Tversky, A. (1972). Elimination by aspects: a theory of choice. *Psychol. Rev.* 79, 281–299. doi: 10.1037/h0032955
- Tversky, A., and Kahneman, D. (1974). Judgment under uncertainty: heuristics and biases. *Science* 185, 1124–1131. doi: 10.1126/science.185.4157.1124
- van Beers, R. J., Sittig, A. C., and Gon, J. J. (1999). Integration of proprioceptive and visual position-information: an experimentally supported model. *J. Neurophysiol.* 81, 1355–1364. doi: 10.1152/jn.1999.81.3.1355
- von Neumann, J., and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- Wang, S., and Rehder, B. (2017). Multi-attribute decision-making is best characterized by an attribute-wise reinforcement learning model. *bioRxiv* 2017:234732. doi: 10.1101/234732
- Zilberstein, S. (2008). *Metareasoning and Bounded Rationality. AAAI Workshop—Technical Report*.