Normative decision asymmetries with symmetric priors but asymmetric evidence

Tahra Eissa[1], Joshua I Gold[2,*], Krešimir Josić[3,4,*], and Zachary P Kilpatrick[1,5,*]

[1]Department of Applied Mathematics, University of Colorado Boulder, Boulder, CO, USA
[2]Department of Neuroscience, University of Pennsylvania, Philadelphia, PA, USA
[3]Department of Mathematics, University of Houston, Houston, TX, USA
[4]Department of Biology and Biochemistry, University of Houston, Houston, TX, USA
[5]Institute of Cognitive Science, University of Colorado Boulder, Boulder, CO, USA
[*]equal contributions

Abstract

Decisions based on rare events are challenging because rare events alone can be both informative and unreliable as evidence. How humans should and do overcome this challenge is not well understood. Here we present results from a preregistered study of 200 on-line participants performing a simple inference task in which the evidence was rare and asymmetric but the priors were symmetric. Consistent with a Bayesian ideal observer, most participants exhibited choice asymmetries that reflected a tendency to rationally interpret a rare event as evidence for the alternative likely to produce slightly more events, even when the two alternatives were equally likely *a priori*. A subset of participants exhibited additional biases based on an under-weighing of rare events. The results provide new quantitative and theoretically grounded insights into rare-event inference, which is relevant to both real-world problems like predicting stock-market crashes and common laboratory tasks like predicting changes in reward contingencies.

1 Introduction

Decision biases are common [1–3]. These biases can result from both non-normative factors, including incomplete or improper knowledge of the structure of the world or the task [4], and normative factors, including asymmetric priors or rewards [5]. Another potential source of decision bias is asymmetries in the quality or availability of evidence associated with the alternatives. Under these conditions, a decision-maker might be more likely to choose the alternative associated with the stronger evidence, even if the alternatives are all equally likely *a priori*. Despite the general nature of this problem, little is known about how evidence asymmetries are addressed by human decision-makers in relation to normative theory.

This problem is particularly acute when the evidence is both asymmetric and rare (e.g., diagnosing a patient based on sporadic, rare symptoms). Consider the following gambling problem. A one-armed bandit has a 50/50 chance of being one of two types, which differ in their probability of yielding a reward, say on 1% versus 2% of pulls. Using observations of its outcomes, the observer must decide whether it is the low-probability bandit, and therefore go elsewhere, or the high-probability bandit, and therefore stay to use it. Here the evidence is asymmetric because a non-rewarded pull provides a slight amount of evidence in favor of the low-probability bandit (e.g., computing the log-likelihood ratio as $\log(P_{low}/P_{high}) = \log(0.99/0.98) = 0.0102$), whereas a rewarded pull provides much stronger evidence in favor of the high-probability bandit($\log(0.2/0.1) = 0.6931$). As a consequence of this asymmetry, before collecting a lot of evidence, a rewarded pull is highly likely to tilt the weight of evidence strongly in favor of the high-reward bandit, thus potentially driving a decision bias.

This kind of rare-event bias has been identified under a variety of conditions and is thought to reflect a number of possible causes. In particular, people often underweight [6–8] or overweight [9, 10] the probability of rare events, based on the evidence provided. Theories for these differential weightings have been framed in terms of a "description-experience gap". When the task structure and associated event probabilities are described explicitly, people tend to overweight event probabilities [9, 11]. This overweighting from descriptions is thought to follow from prospect theory [9]: people tend to overestimate the probability of events described as rare, especially those representing negative outcomes [12]. In contrast, when the task structure is learned from experience, people tend to underweight event probabilities [13, 14].

This underweighting from experience is thought to arise primarily from errors in estimating event probabilities [6, 13], which can be attributed to sampling biases inherent to small samples (i.e., rare events occur less often than their objective probability dictates) [12, 15–17]. However, this delineation comes with caveats: subjects respond to feedback by adopting more objective probability weighting [18], modify estimates based on a regression to mean [17, 19], display recency effects [20, 21], and exhibit a diversity of other biases [22].

Despite the wealth of prior studies on this description-experience gap, there have been few attempts to understand the resulting decision biases, including their large individual variability, in the context of normative decision theory [14]. Here we derive and validate probabilistic principles associated with these evidence-driven decision asymmetries. We focus on two main questions. First, what choice asymmetries do normative and near-normative decision models exhibit, given limited, asymmetric evidence weights? Second, how well do these models account for human decision-making behavior, both on average and with respect to individual variability, under these conditions? To answer these questions, we developed a variant of a classic ball drawing task that requires the subject to infer which of two equally likely jars filled with red and blue balls is being drawn from, with replacement. We compared the decision statistics of an optimal observer that maximizes correct choice probability, plus related models, to those of human subjects. The results show that humans tend to present choice asymmetries during rare-event decision-making that reflect normative principles but with individual variability in additional suboptimalities that can amplify these biases and further degrade performance.

2 Results

We used a task that was based on a classic inference problem (Fig. 1A). On each trial, the subject observed a sample of 2,5, or 10 red and/or blue balls drawn with replacement from one of two jars. The details about the content of each jar were visible at all times, and balls drawn in the sample were shown all at once. The subject had to determine which of the two jars was the source of the sample (Fig. S1). The subject was told that each jar was *a priori* equally likely to be the source. Changes in the ideal evidence weight of each observation were imposed by varying the proportions of red and blue balls in each jar. Under "symmetric" conditions, the ratio of the two ball colors in each jar were reciprocal. In contrast, under "asymmetric" conditions, the ratio was non-reciprocal. We included asymmetric conditions in which one ball color could be rare in both jars (either blue or red, selected randomly for each subject), such that "rare" was defined as having a probability < 0.5. Observing a rare ball thus provided much more evidence about the source of the draws than a common ball.

2.1 Ideal Bayesian Observer

We first derived the strategy of an ideal Bayesian observer that knows the task structure and accounts for all known parameters to make the best guess about which jar was the source of the observed ball sample. We defined the ratios of balls in the jars as h_+ and h_-, where h refers to the fraction of rare colored balls. Below we use the term 'rare' colored balls to refer to the ball color that met the 'rare' criteria for asymmetric conditions above and 'common' colored balls for the other ball color. These labels are used consistently, even in symmetric conditions. The h_+ jar (termed the 'high' jar) included more rare balls, whereas the h_- jar (termed the 'low' jar) included fewer rare balls. Because the two jars were always visible, we assumed the fraction of rare balls, h_+ and h_-, in the low and high jars is known to the ideal observer.

The ideal observer sees a sample of draws, $\xi_{1:n}$, where $\xi_i = 1 (\xi_i = -1)$ if a rare (common) ball is drawn, and computes the log odds ratio (*belief*), $z_n = \log \frac{P(h_+|\xi_{1:n})}{P(h_-|\xi_{1:n})}$, between the probabilities that the sample of draws came from either jar. Balls are seen all at once, but it is revealing to write the recursive update equation

$$\begin{aligned} z_n &= \log \frac{P(\xi_n|h_+)}{P(\xi_n|h_-)} + z_{n-1} \\ &= \Psi(\xi_n) + z_{n-1}, \end{aligned}$$

where the belief increment resulting from the nth observation is

$$\Psi(\xi_n) = \log \begin{cases} h_+/h_-, & \xi_n = +1 \\ (1-h_+)/(1-h_-), & \xi_n = -1. \end{cases}$$

The most likely choice given n observations is given by the sign of z_n ($z_n > 0 \mapsto$ choose the high jar; $z_n < 0 \mapsto$ choose the low jar).

In *symmetric* environments, where the ratio of ball colors were reciprocals, $h_+ = 1 - h_-$, so the magnitude of the belief increment is the same for either observation: $|\Psi(+1)| = \log \frac{h_+}{1-h_+} = -\log \frac{1-h_+}{h_+} = |\Psi(-1)|$. In *asymmetric* environments, we have $h_+ \neq 1 - h_-$, and when one color is rare: $0 < h_- < h_+ < 0.5$. Therefore, observing a rare ball

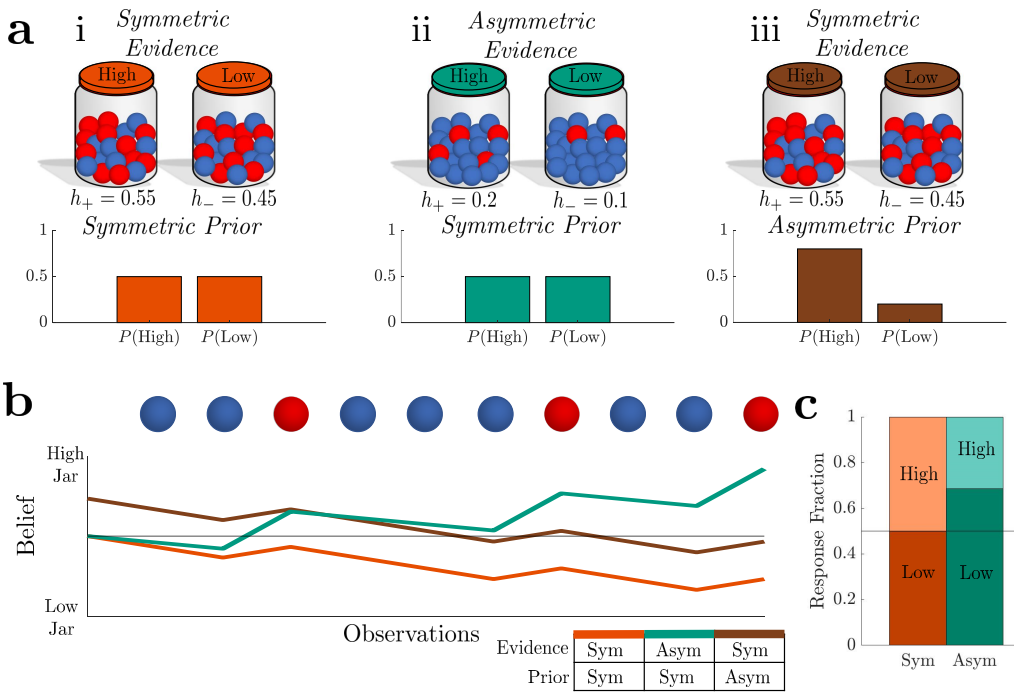

Figure 1: **The basis of decision asymmetries** **a.** Schematic of the Jar Discrimination Task. Balls are drawn with replacement from one of two equally probable jars with different ratios of red to blue balls. Here $h_\pm$ denotes the probability that a red ball is drawn from the "high" probability (h_+) and "low" probability (h_-) jar. Asymmetries can be included in two ways: 1) asymmetric evidence, when $h_- \neq 1 - h_+$ (ii) or 2) asymmetric priors, $P(\text{High} \neq P(\text{Low})$ (iii). Though subjects saw balls in the sample all at once, in models we represented the balls as being presented sequentially for illustration. Likewise, the 'rare' ball (as defined in the case of asymmetric evidence) could be either red or blue. Therefore, we use the term 'rare' and 'common' to refer to specific ball color frequencies, even in symmetric evidence jars. **b.** Example belief state of a Bayesian observer as a function of observations (balls drawn) using the task structures shown in **a**. In the asymmetric evidence environment, rare (here red) balls are weighed more heavily, causing larger increments in belief. With asymmetric priors, an initial belief bias still results in belief updates of equal magnitude. **c.** Ratio of low to high jar responses for an ideal observer across a block of simulated trials with < 10 observations per trial for the symmetric evidence- symmetric prior (orange) vs asymmetric evidence- symmetric prior (green) environments shown in **a**.

leads to a larger belief update than observing a common ball ($|\Psi(+1)| > |\Psi(-1)|$; Fig. 1B, green trace). In contrast, asymmetric priors only shifts the observer's initial belief (Fig. 1B, brown trace).

Because of these weight differences between the symmetric and asymmetric conditions and the limited evidence available, the ideal observer responds differently to the same set of observations. If the observer knows the two jars are equally likely to be the source of the observed sample *a priori*, we might expect the observer to choose both jars in equal proportions over the course of many trials. This is the case for symmetric jars. However, given rare asymmetric ball ratios, the ideal observer chooses the jar containing fewer rare balls (fraction h_-, low jar) more often (Fig. 1C). The reason for this response asymmetry is clearest in the case of small samples of 2 or 5 draws, for which rare balls are often not observed at all. If no rare balls are observed in a sample, an ideal observer will always choose the-lower frequency jar, resulting in a response asymmetry in favor of that jar. More generally, in all asymmetric cases with limited evidence, the observer is less likely to receive sufficient evidence in favor of the high jar (even when it is correct) than to be guided to an accurate low jar response. Thus, the optimal choice strategy produces a response asymmetry, even when the prior is symmetric.

2.2 Human Responses are Similar to the Ideal Observer On Average

To determine whether and to what extent human observers exhibit these response asymmetries, we recruited 201 subjects using the crowdsourcing platform Amazon Mechanical Turk (MTurk) to perform the Jar Discrimination task shown in Fig. 1Ai and ii. Each subject performed $T = 42$ trials ($T = 60$ in the control (CT) block) under four conditions that varied in difficulty and evidence asymmetry: hard asymmetric (HA), hard symmetric (HS), easy asymmetric (EA), and easy symmetric (ES) (Fig. 2A).

Like for the ideal observer, the subjects tended, on average, to exhibit choice asymmetries in the asymmetric, but

not symmetric, conditions that were slightly larger when the task was difficult (the HA versus the EA block; Fig. 2B) and asymmetries were reduced as trial length (balls observed) increased (Fig. 2C). These trends also tended to match the predictions of the ideal observer in terms of the impact of observing a rare ball under these conditions. Specifically, we identified all trials in which at least one rare ball was observed (rare-ball trials). There was an inverse relationship between the percentage of these rare-ball trials and the low-jar response frequency within a given block, such that both human and ideal observers tended to be more/less biased towards the low jar when rare-ball trials were less/more common (Fig. S6) Moreover, accuracy (percent correct) tended to be higher for the high jar (light curves in Fig. 2F) and lower for the low jar (dark curves in Fig. 2F) on rare-ball trials versus all trials. These effects were strongest for short trials and asymmetric blocks (green versus orange curves in Fig. 2F), which mimicked the ideal observer.

2.3 When do Human Subjects' Responses Differ from the Ideal Observer?

Despite general similarities in the performance of human and ideal observers on the asymmetric tasks, they also showed quantitative differences in terms of the magnitude of response asymmetries and their dependence on trial length, especially in the hard-asymmetric task (Fig. 2B,C). To better understand these differences and how they varied across individual subjects, we considered the impact of non-normative repetition biases that are common in human decision-making [23–25]. Specifically, we measured the frequencies of four possible two-trial response sequences [25]: 1) high-jar repetition; 2) low-jar repetition; 3) high-low jar alternation; and 4) low-high jar alternation. Because the correct jar was drawn independently on each trial, we could identify each participant's stereotypical response patterns by comparing the relative frequencies of these response sequences during each block (Fig. 3A). This procedure allowed us to distinguish biased responses that reflected a preference for a jar (i.e., low versus high) versus a response sequence (i.e., alternations versus repetitions) for each block and subject and then compare those preferences to those predicted by an ideal observer presented with the exact same trial sequences.

Many subjects exhibited consistent low-jar preferences in asymmetric blocks but no clear preferences in symmetric blocks (Fig. 3B; percentage of individual subjects with low-jar preferences=20–22% in symmetric blocks, 67% in HA blocks, 58% in EA blocks). The low-jar preference on asymmetric trials remained for some subjects even after normalizing by the preference expected by the ideal observer, implying an excess, non-normative bias (Fig. 3C; percentage of individual subjects with this excess low-jar preference=0.5–2% in symmetric blocks, 22% in HA blocks, 8% in EA

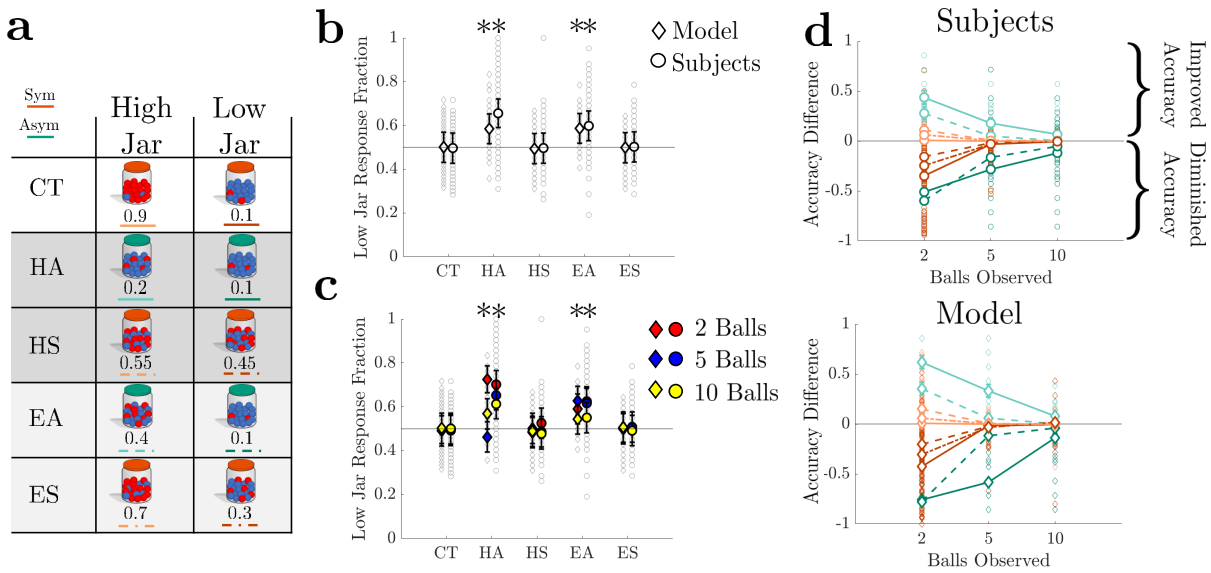

Figure 2: **Ideal observers and human participants show similar behaviors, on average. a.** Schematic showing the different variants of the task and parameters used with human subjects (n=201): CT = Control; HA = Hard Asymmetric; HS = Hard Symmetric; EA = Easy Asymmetric; ES = Easy Symmetric. Colors and line type shown under each jar correspond with plots in **d**. **b.** Low jar response fractions for each participant (grey circles) and sample-matched ideal observer responses (grey diamonds) for each block. Population bootstrapped means (1000 iterations) and % confidence intervals shown in black. * denote multiple comparison corrected population p-vals <0.05. **c.** Low jar response fraction as in **b** for each trial length: 2, 5, or 10 balls observed. (HA population p-vals <0.05: model- 2 ball, subject- 2,5 ball; EA: model- 2, 5 ball, subject- 2, 5 ball). **d.** Difference in performance on trials with at least one rare ball compared to performance from all trials for the given condition for subject (top) and ideal observer (bottom) responses, plotted as a function of the number of balls observed (trial length) for each condition. Changes in accuracy that corresponded with response asymmetries tend to be larger on asymmetric blocks with fewer observations in both human and ideal observers.

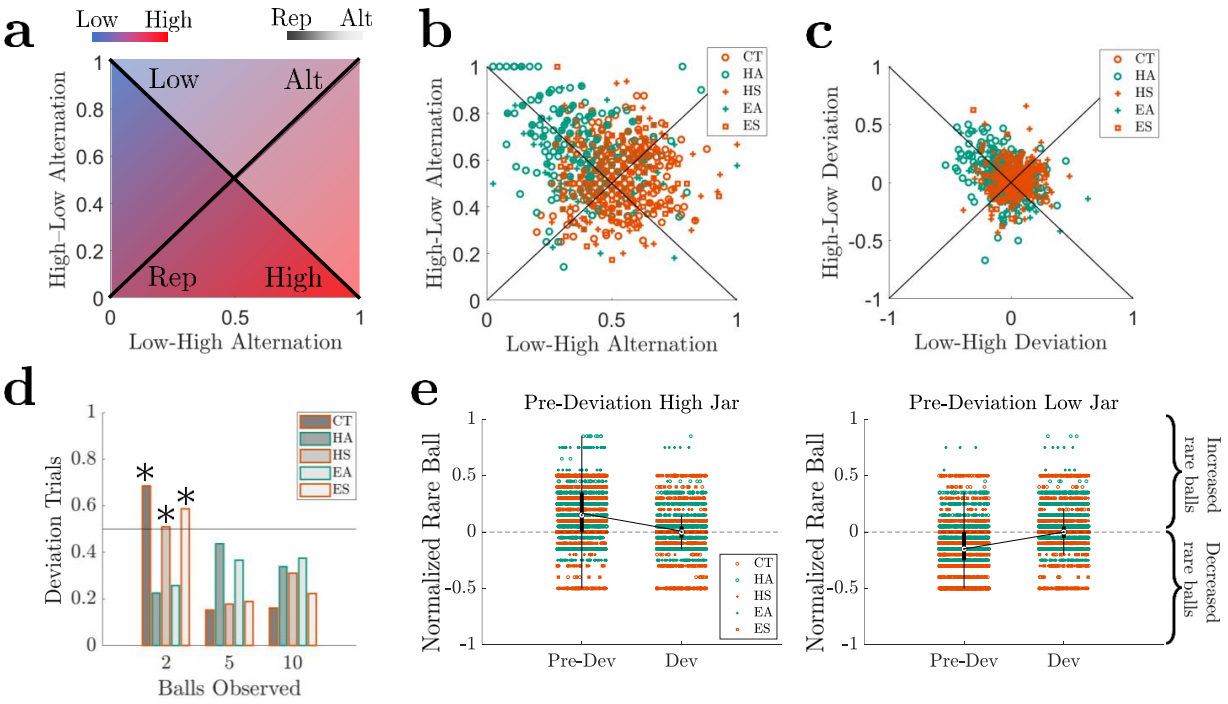

Figure 3: **Human subject deviations from the ideal observer.** **a.** Schematic depicting the preferences that can be inferred by comparing the response-alternation probabilities (low-high or high-low). The increasing diagonal corresponds to response sequence preference. The decreasing diagonal corresponds to jar preference. The intersection point represents unbiased responses. **b.** Human subjects' alternation plots. There is a relatively strong low-jar preference in asymmetric blocks (percentage of individual subjects 67% in HA blocks, 58% in EA blocks) but not repetition preferences (11–18% of the subjects in a given block.) **c.** Alternation plots normalized by the matched ideal observer. There was an excess low-jar preference for some subjects in asymmetric blocks. **d.** Fraction of trials where subjects' responses do not match the ideal observer (deviations) for each trial length. Asterisks denote blocks in which > 0.5 of deviation trials included only two ball draws. **e.** Normalized rare ball observations (percentage of rare balls in a trial sample compared to the average percentage) on the trial prior to (pre-dev) and on trials subjects deviated from the ideal observer (dev). Positive (negative) values represent an enhanced (decreased) number of rare balls compared to the average expected number of rare balls. Left plot represents sequences in which the pre-deviation trial was a high jar trial, right plot is equivalent for pre-deviation low jar trials. Raw data from individual subjects are shown with box and whisker overlays.

blocks). Some subjects showed an additional repetition bias, but these biases were not common or consistent across blocks (11–18% of the subjects exhibited repetition biases in at least one block).

To understand what factors might drive the excessive, non-normative preference for the low jar for some subjects, we analyzed performance from all individual trials on which the subject gave a response that differed from the sample-matched ideal observer (deviation trials). For all blocks, the deviations tended to occur with relatively constant frequency throughout a block (Fig. S8) and thus did not appear to be a learned compensatory mechanism to a realization that asymmetric choice distributions did not match the expected symmetric prior probabilities. Moreover, the deviations were not associated with particular response sequences (Fig. S9). However, these deviation trials were most common for short (2 ball) trials in symmetric blocks, implying a straightforward sampling effect (Fig. 3D), but trial length did not account for these deviations in asymmetric blocks.

Instead, the deviations in asymmetric blocks were related to the number of rare balls observed on the previous trial, regardless of trial length. Specifically, when a deviation trial was preceded by a high-jar trial, the preceding trial tended to have more rare balls than usual. When a deviation trial was preceded by a low-jar trial, the preceding trial tended to have fewer rare balls than usual (Fig. 3E). When an average number of rare balls were observed on the deviation trial, a deviation often occurred, regardless of whether it was a low- or high-jar trial. Thus, deviations from predictions of the ideal observer on asymmetric trials tended to reflect an excessive low-jar preference and follow trials with strong evidence, suggesting that subjects were not necessarily swayed by the sequence of previous responses but by the number of rare balls previously observed and possibly a misestimation of their weight.

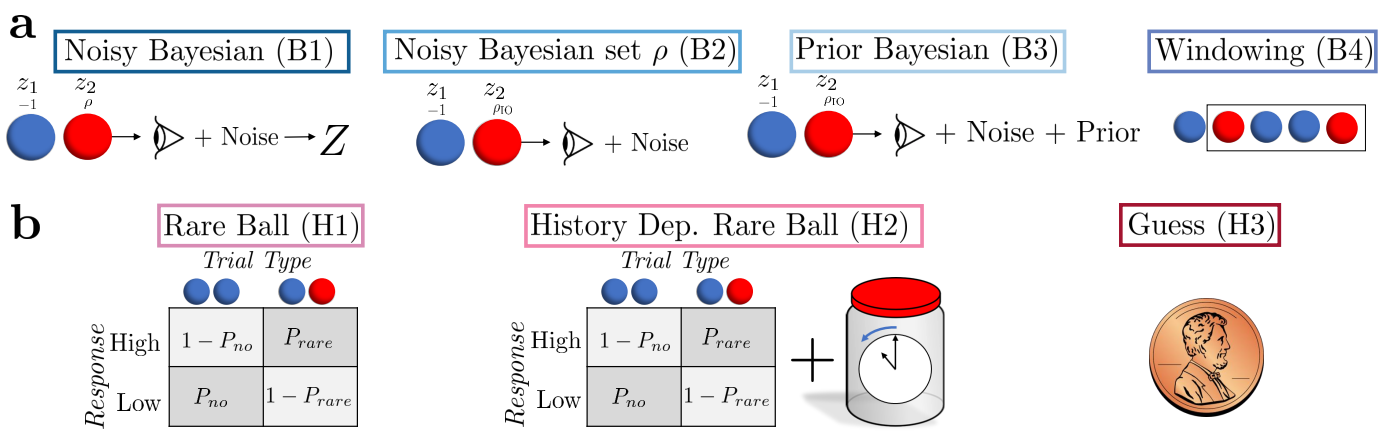

Figure 4: **Models for comparison a.** Bayesian models. B1: Noisy Bayesian model, which includes noise and a variable rare ball weight ρ. B2: Noisy Bayesian with a set ρ, which sets ρ to the rare ball weight of the ideal observer (ρ_{IO}). B3: Prior Bayesian, which includes a jar bias (prior) but ρ matching the ideal observer. B4: Windowing, which applies the noisy Bayesian model to a selected window of data. Here,we begin the window at the rightmost ball in the sample (the final ball, assuming subjects view the sample moving from left to right). **b.** Heuristic models. H1: Rare Ball, in which an observer has a preference for each response based on the presence of a rare ball (rare or no rare). H2: History-Dependent Rare Ball, which extends the Rare Ball model to include the previous response. H3: Guess, in which the observer guesses the high jar with some probability.

2.4 Formal Model Comparisons

Given that the subjects' decision strategies showed both similarities and differences with the ideal observer, we sought to more precisely quantify their strategies in terms of evidence-accumulation models. We compared their responses to a Bayesian observer and the following suboptimal models. The Noisy Bayesian model was defined by assuming that the model observer computes a noisy version of the log likelihood,

$$z_n = \left\{ \begin{array}{cc} \log h_+/h_-, & \xi_n = +1 \\ \log(1-h_+)/(1-h_-), & \xi_n = -1 \end{array} \right\} + z_{n-1} + w_n,$$

where w_n that is a normally distributed random variable with zero mean and variance a. To generalize to a detuned (suboptimal) version of the Bayesian observer and because only the sign of z_n matters for decisions, we considered a reduced form of the model

$$z_n = \left\{ \begin{array}{cc} \rho, & \xi_n = +1 \\ -1, & \xi_n = -1 \end{array} \right\} + z_{n-1} + w_n$$

where ρ is a free parameter representing the observer's rare ball weight (corresponding to when a rare ball is observed $\xi_n = 1$) (Fig. 4A, model B1). Changes in ρ away from 1 reflect an asymmetric weighting of the 'rare' ($\xi_n = 1$) as opposed to common ball ($\xi_n = -1$) (Fig. S10). Also, setting $\rho \equiv [\log h_+/h_-] \,/\, [\log(1-h_+)/(1-h_-)]$ recovers a rescaled noisy version of the ideal observer (Noisy Bayesian set ρ, model B2).

We considered two other extensions of the Noisy Bayesian model (Fig. 4A), including: 1) a noisy Bayesian observer with a free parameter for the (potentially incorrect) prior, but ρ set as in model B2 (Prior Bayesian, model B3); and 2) a windowing model (Windowing, model B4) in which the number of observations used by the noisy Bayesian model is a free parameter. The number of observations could be smaller or larger than that drawn on a single trial, and thus could include draws from past trials . We assume that the window begins at the rightmost ball in the sample (the final ball, assuming subjects view the sample moving from left to right).

We also considered several heuristic models (Fig. 4B). The first was a rare-ball model (model H1), in which the observer's probability of choosing the high jar is swayed by the observation (or lack thereof) of a rare ball, resulting in a model with two parameters P_{rare} and P_{no}. If a rare ball was observed ($\max(\xi_{1:N}) = 1$), the observer responds in favor of the high jar with probability P_{rare} and low jar with probability $1 - P_{\text{rare}}$. In contrast, if no rare balls were seen ($\max(\xi_{1:N}) = -1$), the observer responds in favor of the low jar with probability P_{no} and high jar with probability $1 - P_{\text{no}}$. The second (model H2) was an extension of the rare-ball model that accounted for the previous trial's response (high or low jar) by now defining four probabilities of response: $P_{\text{rare, high}}$, $P_{\text{rare, low}}$, $P_{\text{no, high}}$, and $P_{\text{no, low}}$. The third was a simple guessing model (model H3) in which the observer selects the high jar with some probability P_{guess}, regardless of the observations (Fig. 4B).

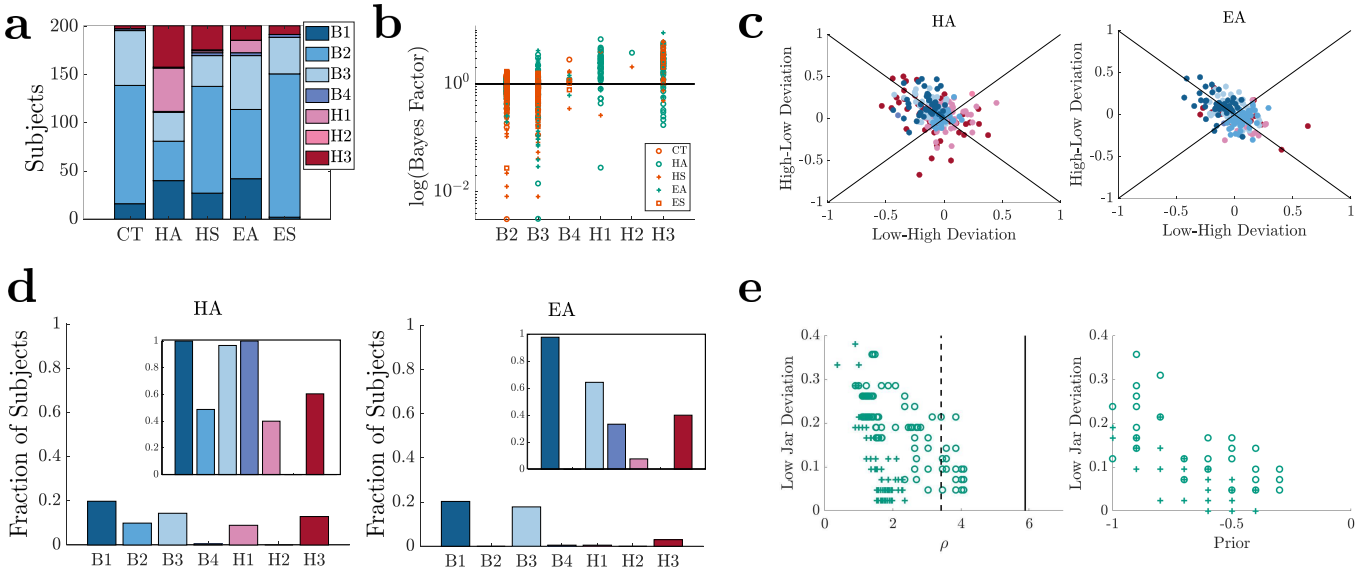

Figure 5: **Subjects predominantly apply Bayesian, but not necessarily optimal, approaches a.** Subjects categorized by most likely model, colors and labels as in Fig. 4. **b.** Bayes factors computed between the noisy Bayesian (B1) and each alternative model. Given that model B1 was used as our comparison, subjects whose most likely model is the noisy Bayesian (Bayes factor values of zero) are not shown. Bayes factors $\log > 1$ (black line) are considered to be strong evidence [26]. **c.** Asymmetric block alternation plots from Fig. 3**c**, color coded by subjects' most likely model. **d.** Fraction of subjects that showed excess low jar responses for asymmetric blocks based on most likely model. Subplot: Percentage of subjects with a specific model match that had excess low jar responses. Note that most subjects that matched with the Noisy Bayesian or Prior Bayesian models show this increase. **e.** Left: Maximum likelihood estimation of the rare ball weight ρ versus increase in low jar responses (shown as the deviation from the ideal observer's response asymmetry) for subjects best matched to the Noisy Bayesian model in the asymmetric block (marker legend as in **b**). Black traces note the ideal observer's corresponding rare ball weight (HA: solid line, EA: dashed line). Right: Maximum likelihood estimation of response bias (Prior) compared to increased low jar response asymmetry for subjects best matched to the Prior Bayesian model in the asymmetric block (marker legend as in **b**). Negative values correspond to a bias in favor of the low jar.

Using Bayes factors, we selected the most likely model for each subject for each block (after confirming model identifiability; see Figs. S11, S12, S13 and Methods for details). For symmetric blocks, the behavior of most subjects (86-98%) was most consistent with one of the Bayesian models (Figs. 5A and B). For asymmetric blocks, over half the subjects (HA: 56%, EA: 86%) also showed responses consistent with Bayesian approaches. This result suggests that subjects use a variety of strategies when confronted with asymmetric evidence, often but not always, consistent with a Bayesian observer.

Bayesian strategies also tended to account best for the behavior of subjects who had a strong preference for the low jar in asymmetric blocks. Specifically, the behavior of those subjects tended to be most consistent with either the Noisy Bayesian (B1) or Prior Bayesian (B3) models (Fig. 5C and D). Likewise, subjects that used one of those two strategies tended to have excess response asymmetries (Fig. 5D, subplots). From these fits, we identified two factors that could account for excessive preferences for the low jar by our subjects. First, subjects with an increasingly excessive low-jar preference tended to increasingly underweight the evidence from a rare-ball observation, ρ, compared to the ideal observer, based on parameter fits of the Noisy Bayesian model (Fig. 5E; left). Second, subjects with an increasingly excessive low-jar preference and matched with the Prior Bayesian model also tended to have a prior that was increasingly biased towards the low jar (Fig. 5E; Right). In summary, subjects who showed an increased low jar preference in asymmetric blocks were most often matched to a Noisy Bayesian model, in which the evidence associated with a rare ball was under-weighed but the prior was unbiased, or a Prior Bayesian model, in which the evidence associated with a rare ball was weighed appropriately but the prior was biased in favor of the low jar.

2.5 Subjects display a trade off between information use and accuracy

To better understand the decision strategies used by the participants without relying as strongly on specific models, we also used a complementary, model-free analysis. We computed the mutual information (MI) between the observations

and responses from each participant separately across trials in each block:

$$I(f_B^i, r_B^i) = \sum_{f_B^i \in F} \sum_{r_B^i \in R} P(f_B^i, r_B^i) \log_2 \frac{P(f_B^i, r_B^i)}{P(f_B^i)P(r_B^i)},$$

where f_B^i is a vector of observed 'features', i is the individual subject, B is the block, and r_B^i is the vector of responses. Based on rate-distortion theory [27], we compared these values to the subject's accuracy and the maximum possible mean accuracy based on strategies with a fixed level of mutual information (MI) between the features and the response (the accuracy "bound"). We computed a subject's MI using two type of features, either the balls observed on a given trial $\boldsymbol{\xi}_B^i$ ($I_{O,B}^i$) or the balls observed $\boldsymbol{\xi}_B^i$ and the previous trial's response, R_{t-1} ($I_{OR,B}^i$) (Fig. 6A). This approach allowed us to identify three important aspects of a subject's strategy: 1) its complexity, such that higher values of I correspond to more complex strategies; 2) the relevant features, such that higher MI with the inclusion of the previous response suggests attention to irrelevant task features; and 3) how effectively that level of complexity was used to make decisions, such that accuracy closer to the bound implies more effective use of the given strategy, for example by focusing only on relevant information.

In general, subjects who used more complex strategies also tended to use those strategies more effectively (i.e., they were close to the optimal bound and did not show any shift in MI with the addition of previous responses to the feature space $I_{OR,B}^i \approx I_{O,B}^i$), for all five blocks (Fig. 6B and C; note that points exceeding the bound resulted from the limited number of trials for each subject). In contrast, subjects with lower complexity strategies were suboptimal and showed signs of considering irrelevant information in their strategy ($I_{OR,B}^i > I_{O,B}^i$; Fig. S15).

To relate this finding to our model-based analyses, we simulated responses from observers using each of the seven models. Based on these responses, we computed each observer's MI and accuracy and identified where most of a given model's simulated observers fell in MI-accuracy space. We then estimated which regimes of MI-Accuracy space different classes of models were most likely to occupy (Figs. 6D and S16). Bayesian models (Noisy Bayesian, Noisy Bayesian set ρ, Prior Bayesian) clustered together strongly at high MI values near the accuracy bound. The Windowing model tended to show intermediate accuracy and MI. Heuristic models occupied a broad range low MI space, including some regions that were close and others that were far from the accuracy bound. Creating similar plots using the subjects' MI results and their most likely model match, we found a high level of correlation in MI-accuracy space between the simulated and human model maps (Figs. S17 and S18).

Given that heuristic models occupied the same regime as subjects leveraging low complexity strategies (low MI). We considered whether these subjects' suboptimality may be a result of constraints due to the model type by computing rate-distortion bounds for each alternative model. Even when compared to the bound for a subject's most likely model, the performance drop off persists (Fig. S19), suggesting that subjects' suboptimality is not a result of their strategy directly but of some other cause.

Finally, we readdressed the strategy complexity similarities between subjects and ideal observers by comparing subjects' performance to that of an ideal observer at different MI levels (Fig. 6E). These results showed a decrease in subject performance relative to the ideal observer at lower values of both MI and accuracy, as well as an increase in the distance of the subject from the ideal-observer accuracy bound, confirming the results of our model group maps.

Thus, individual differences in task performance could be characterized using data-driven metrics that yielded insights into the model-fitting results. Subjects who used more complex strategies (higher MI) performed closer to the optimal bound on accuracy and better matched Bayesian models, specifically the ideal observer. Subjects who used less complex strategies (lower MI) were more likely to match heuristic models and deviate from the optimal bound, reflecting the relative ineffectiveness of those strategies.

3 Discussion

We examined how asymmetries in the available evidence impact decision-making. In agreement with the prediction of a normative model, we found that a majority of human subjects in our study exhibited choice asymmetries when presented with limited, asymmetric evidence but a symmetric objective prior. In particular, their choices were most consistent with those of noisy, but well-tuned, Bayesian observers, albeit with some additional biases in evidence weights and/or subjective priors. Variability in response accuracy was well explained by subjects' effective use of task-relevant information (i.e., balls drawn). For example, subjects who performed poorly also displayed low mutual information between evidence and response, but improved mutual information values when prior, irrelevant choices were considered. These results suggest that human observers integrate asymmetric evidence in Bayesian, but not necessarily optimal ways. This is in line with prior findings showing people underweight the relevance of rare events when making decisions [6,13].

A number of previous studies focused on the effects of asymmetric priors [5, 28, 29] or rewards [30]. In contrast, we asked how asymmetries in evidence contribute to choice asymmetries. Our experimental setup mimics classical [31,

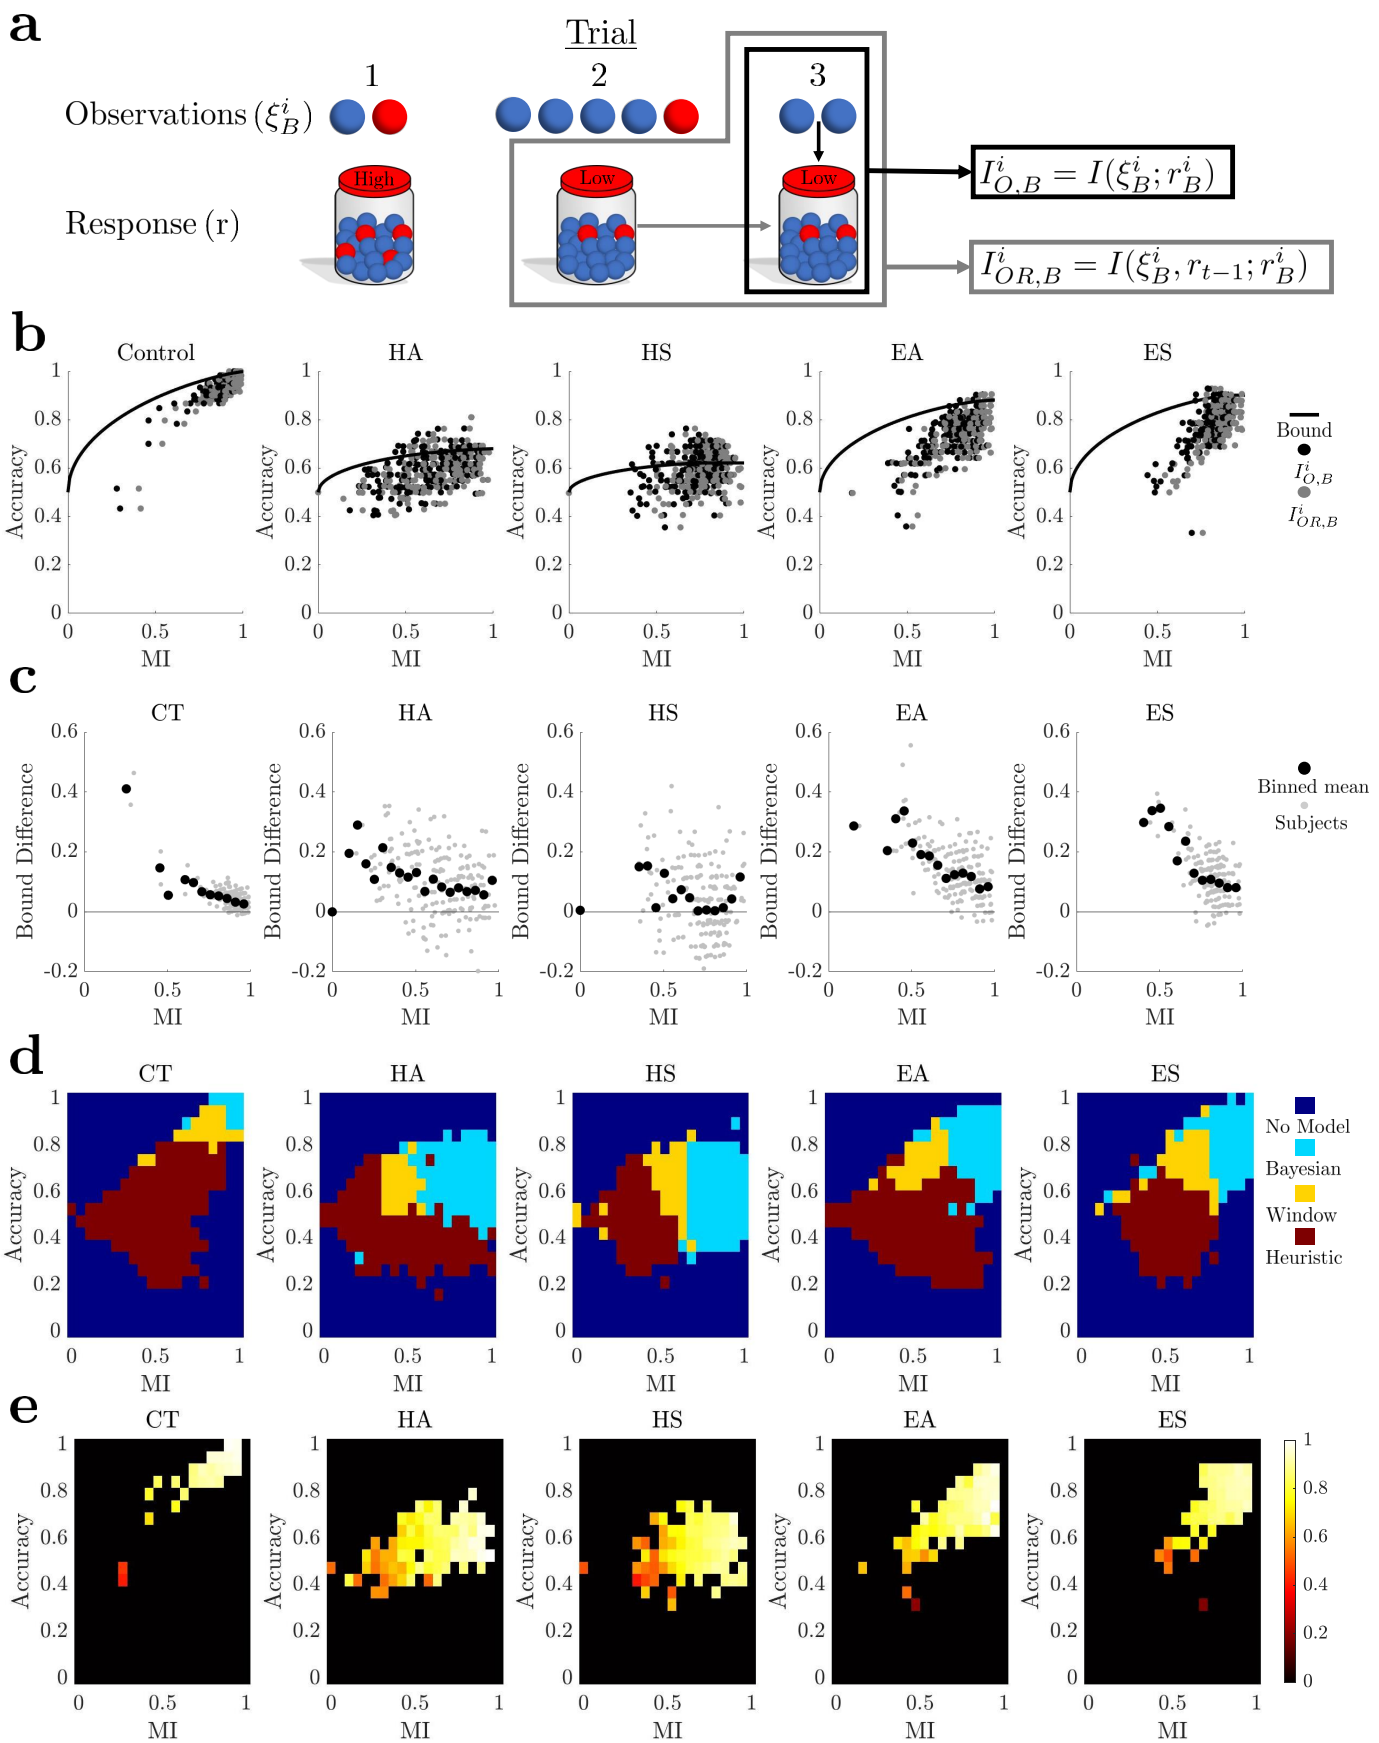

Figure 6: **Human strategy can be described by a tradeoff in the effective use of task information. a.** Schematic of the mutual information between task details and responses. $I^i_{O,B}$ includes only the relevant observations in the feature set, while $I^i_{OR,B}$ includes both the current observations and the irrelevant previous response. **b.** Rate distortion plots for each block portraying the bootstrapped means (1000 iterations) for $I^i_{O,B}$ (black dots) and $I^i_{OR,B}$ (grey dots) for each subject. Bound shows the maximum predicted accuracy for each MI value. Note that points exceeding the bound resulted from the limited number of trials for each subject **c.** Normalized distance from the rate distortion bound. **d.** Model likelihood map. Models have been grouped as: Bayesian (Noisy Bayesian, Noisy Bayesian set ρ, Prior Bayesian), Windowing model, or heuristic (Guess, Rare Ball, History-dependent Rare Ball). **e.** Subject optimality heatmap representing percentage of matched responses between bootstrapped subject data and ideal observer.

32] and recent task designs [33] in which the animal (or subject) must discriminate between two similar stimuli (or environments): lights or frequencies in the case of animals, and jars for our subjects. Unlike these previous studies, our work highlights the effect that asymmetric evidence weight has on the decision making strategy.

Our approach can be extended to answer other questions about how humans make decisions: For example, modeling and experiments have helped identify how ideal observers and humans adaptively integrate evidence in dynamic environments [34–37]. Learning the probabilistic change rate (hazard rate) is crucial to optimizing evidence discounting arriving from dynamic sources [38, 39]. An observer learns this by counting infrequent change points when change rates are low, and thus relies on asymmetric evidence. Similar evidence asymmetries arise when an observer must decide between two choices supported by rare evidence (*e.g.* deciding which of two areas is more prone to earthquakes [40], or which of two states is more susceptible to a severe terrorist event [41]). Our current work thus informs future studies by identifying the types of choice asymmetries exhibited by ideal observers, and the biases exhibited by human subjects.

A primary question of our work was to compare how well human subjects performed given limited, asymmetric evidence but symmetric priors to the predictions expected by an ideal observer. Limited evidence can also produce choice asymmetries in rational decision makers in alternate contexts, such as judging the process generating a binary sequence [42]. Short random sequences tend to provide very weak evidence of their randomness, so there are many competing hypotheses that would endow the underlying processes that could be concluded by an ideal observer with an expansive enough prior. This is consistent with the observation that people's choices can be strongly altered by cues predicted so seeing structure in randomness may follow near-Bayesian models [42, 43].

Serial biases typically play a crucial role in suboptimal inference in subjects performing sequences of decision-making trials [23, 24, 44, 45], especially in the face of limited data or rare observations [11, 20, 21, 46–48]. Such effect are consistent with accumulator models that incorporate an initial belief shift in the direction of the previous response [49,50]. The effects of recency biases were fairly weak or not present in our subjects' decision strategies, but has a far less significant impact on subject decisions that the presence of a rare ball.

We had also hypothesized that subjects would be uncomfortable with the discrepancy between the symmetric prior and the asymmetric response distribution supported by the asymmetric evidence. As a result, we thought they might adjust their evidence weighting or within-trial priors, so their response fractions over a block more closely matched the true prior. In contrast to this prediction, most subjects continued to produce asymmetric responses and in some cases, excessively favored the low jar. Strategies of subjects who showed excess asymmetry were in line with one of two Bayesian but suboptimal approaches. One strategy involved underweighting evidence provided by the rare ball in comparison to the ideal observer. We speculate that subjects may undervalue rare observations [6] or may be attempting to symmetrize the evidence weight of either balls, so the infrequency of rare balls further exacerbates response asymmetries. Another strategy subjects use is to adopt a prior in favor of the low jar, which may be innate [51] or built up over the course of a block in response to seeing many early choices for the low jar. Since we provided subjects an explicit description of both the prior and choice statistics to mitigate effects of pre-existing priors, we suspect the latter explanation is plausible, in line with work on base-rate neglect [1, 52–55].

The description-experience gap is an oft cited theory use to distinguish the tendency to overweight the importance of rare events when their frequency is described and underweight when they learn rare event frequency through experience [6, 11, 15, 56]. Unlike our task, rare events in these tasks were typically losses in reward [6, 13, 57], so subjects either make a decision with a known risk (description) or learn a payout distribution with some uncertainty (experience). Small sample sizes thus play an important role in subjects misestimating and underappreciating rare losses in the experience condition [11]. Even without risk or uncertainty about the rare event probability, subjects in our task underweight rare events, though the event probabilities are known to them. Future work could consider approaches that combine described and experienced rare event probabilities could help further disentangle how these two avenues for learning probabilities contribute to how rare events are weighted. Such work would be relevant to better understanding people's decision making concerning financial investments [58–61] and climate change [62]. Subjects provided both descriptive and experiential information about event probabilities can perform closer to optimality [63].

Resource rational theories of decision making suggest people may simultaneously seek to maximize reward while constraining strategy complexity [64]. We offered no reward to motivate this trade off, but subjects using high complexity strategies hugged the accuracy bound. On the other hand, subjects using low complexity strategies deviated strongly from the accuracy bound as in [65, 66], which is not explained by the use of heuristics alone. These subjects may have been more distracted, devoted less attention, or fewer cognitive resources. It is possible this may also have resulted in their taking longer to learn the optimal strategy, but given that subjects trained and were tested prior to the main task, this latter explanation seems less relevant.

Complementary to our findings, recent work analyzing the response trends of subjects across multiple trials with correlated evidence show subjects with more complex strategies tend to fall away from an ideal observer's bound on the ability to predict future events [39, 67]. In such tasks, humans with less complex strategies tend to lie closer to the perfomance bound. The difference in task design as well as the different measure being used to assess performance

(predictivity versus accuracy) may account for these contrasting results. Moreover, in our task, evidence from past trials is irrelevant, and so the pool of evidence that is useful for future prediction is far more limited. Thus, in tasks with limited evidence, it may be easier for subjects to attain performance bounds, while those that underutilize evidence are more likely to do so inefficiently.

Our results show that limited and asymmetrically weighted evidence shift the response preferences of both humans and optimal observers. Some people became even more entrenched in these decision asymmetries than ideal observers, exhibiting strongly biased response statistics. How humans respond to asymmetric and limited evidence is an important consideration for future work on decision making. Since people often make decisions, in everyday life, on short notice or with lopsided observations, it is important to understand if they respond to these limitations by using available evidence as efficiently as possible, or by becoming more deeply rooted in the belief and decision towards which these asymmetries push them. Generating computationally grounded theories of bias in people can help us better understand how to keep people well informed, and to determine how individuals are susceptible to receiving biased or limited information.

4 Methods

4.1 Task

We recruited 201 consenting subjects to perform the Jar Discrimination Task on the Amazon Mechanical Turk crowd-sourcing platform (95 female, 105 male, 1 non-disclosed). Subjects were compensated $4.50 for completing the task and were only recruited if they had a 95% or better approval rating and had performed at least 100 previous approved tasks. Location restrictions were limited to the United States. Human subject protocols were approved by the University of Pennsylvania Internal Review Board. The task was preregistered at osf.io prior to data acquisition.

For the task, subjects were shown a string of 2,5, or 10 red/blue balls randomly drawn with replacement and asked to determine which of two equally probable jars (displayed on the screen) was used to draw the balls (Fig. 1A and Fig. S1). The ratio of colored in each jar were varied to create five blocks of trials and can be described by the fraction of balls that are one color, termed the ”rare ball” color. Each subject was randomly assigned a rare ball color of either red or blue, which was consistent throughout all blocks. Blocks were defined by the following rare ball fractions for the high (more rare balls) and low (fewer rare balls) jars: Control (0.9/0.1), Hard Asymmetric (0.2/0.1), Hard Symmetric (0.55/0.45), Easy Asymmetric (0.4/0.1), Easy Symmetric (0.7/0.3).

Before beginning the full task, subjects were shown a training slideshow and performed 24 trials using the control block. Subjects were required to achieve at least 80% accuracy on these trials to proceed to the full task. Full sessions included pseudo-randomized block orders for the remaining 4 test blocks interspersed with 12 control trials between test blocks. Each test block had 42 trials with the jar used for ball draws and trial length pseudo-randomized within the block. All trial lengths, block lengths, and jar ratios were validated using synthetic data prior to data acquisition to ensure model and parameter identifiability (See model fitting below).

4.2 Models

4.2.1 Bayesian Models

Ideal Observer. In the simplest case without noise, an ideal Bayesian observer sees a sample of n ball draws, $\xi_{1:n}$, where $\xi_i = 1(\xi_i = -1)$ if a rare (common) ball color is drawn, uses this information to appropriately weigh evidence from each draw. To do so, the ideal Bayesian observer updates its log likelihood ratio (*belief*), $z_n = \log \frac{P(h_+|\xi_{1:n})}{P(h_-|\xi_{1:n})}$, between the probabilities the sample came from the jar with the high or low rare ball fraction. For illustration, we write the update iteratively:

$$\begin{aligned} z_n &= \log \frac{P(\xi_n|h_+)}{P(\xi_n|h_-)} + z_{n-1} \\ &= \log \left\{ \begin{array}{cc} h_+/h_-, & \xi_n = +1 \\ (1-h_+)/(1-h_-), & \xi_n = -1 \end{array} \right\} + z_{n-1}. \end{aligned}$$

The most likely choice after the nth observation is given by the sign of z_n ($z_n > 0 \mapsto$ choose the high jar; $z_n < 0 \mapsto$ choose the low jar). Given a flat prior, $z_0 = \log \frac{P(h_+)}{P(h_-)} = 0$, and a symmetric environment, $h_+ = 1 - h_-$,

$$z_n = \xi_n \cdot \log \frac{h_+}{1-h_+} + z_{n-1},$$

so the belief increment ($\left|\log \frac{P(\xi_n|h_+)}{P(\xi_n|h_-)}\right|$) is the same for either observations. When the environment is asymmetric, $h_- < 1 - h_+$, each ball is given a different evidence weight.

For n ball draws and an arbitrary trial, we can compute the response probability for either jar

$$P(r = h_\pm) = P(r = h_\pm|h_+)P(h_+) + P(r = h_\pm|h_-)P(h_-)$$

using binomial distributions.

$$P(r = h_\pm) = 0.5 \left[\sum_{j=1}^{n} \left(\begin{array}{c} n \\ j \end{array} \right) h_\pm^j h_\mp^{n-j} \right].$$

Conditioning on trial type, we can extend this analysis to identify the minimum number rare balls B that must be observed to produce a high jar response, given n balls total. This number is dependent on the values of h_+ and h_-, and underlies the "sawtooth" accuracy trend as a function of ball count n. When the jars are symmetric ($h_+ = 1 - h_-$), $B = n/2$. In asymmetric cases, $B < n/2$ if $h_+ + h_- < 1$ ($B > n/2$ if $h_+ + h_- > 1$, but we do not study this regime in subjects). Thus,

$$P(r = h_+|h_\pm) = \sum_{k>B}^{n} \left(\begin{array}{c} n \\ k \end{array} \right) h_\pm^k (1 - h_\pm)^{n-k}$$

and

$$P(r = h_-|h_\pm) = \sum_{k=1}^{\lfloor B \rfloor} \left(\begin{array}{c} n \\ k \end{array} \right) h_\pm^k (1 - h_\pm)^{n-k}.$$

Noisy Bayesian. We extended the ideal observer model to include noise $w_n \sim \mathcal{N}(0, a)$ that is normally distributed with zero mean and variance a and deviations in the weight given to rare ball observations (since only the sign of z_n matters for the decision):

$$z_n = \left\{ \begin{array}{cc} \rho, & \xi_n = +1 \\ -1, & \xi_n = -1 \end{array} \right\} + z_{n-1} + w_n$$

where ρ is a free parameter representing the observer's rare ball weight (corresponding to when $\xi_n = 1$).

Noisy Bayesian set ρ Model. If we take a special case of the noisy Bayesian model for which

$$\rho \equiv \frac{\log h_+/h_-}{\log(1 - h_+)/(1 - h_-)},$$

the resulting model has a single free parameter a the variance of the noise and is equivalent to the ideal observer model with noise added to each observation.

Prior Bayesian Model. We can modify the noisy Bayesian set ρ model further by assuming the observer has a biased prior, which deviates from the underlying symmetric probability of drawing a jar on either trial, so that

$$z_0 \neq 0 \neq \log \frac{\tilde{P}(h_+)}{\tilde{P}(h_-)},$$

where $\tilde{P}$ represents the observer's assumed probability which is different than the true probability. A positive value of z_0 denotes a preference for the high jar and a negative value is in favor of the low jar.

Windowing Model. The noisy Bayesian model can be extended to include a free parameter y that fixes the number of ball observations used on any given trial. If the current trial has more than y observations, the last y of the sample are used, beginning from the final/ rightmost displayed ball. If the current trial has fewer than y balls, the observations includes the last $y - n$ balls from the previous trial.

4.2.2 Heuristic Models

Rare Ball Model. The rare ball model utilizes probabilities of response based on the observation (or lack there of) of a rare ball on any of the draws

$$P_{response} = \left\{ \begin{array}{c} P(r = h_+|\text{rare}) = P_{rare} \\ P(r = h_-|\text{rare}) = 1 - P_{rare} \\ P(r = h_-|\text{no rare}) = P_{no} \\ P(r = h_+|\text{no rare}) = 1 - P_{no} \end{array} \right\},$$

where rare $\equiv (\max(\xi_{1:N}) = 1)$ and no rare $\equiv (\max(\xi_{1:N}) = -1)$.

History Dependent Rare Ball Model. An extension of the rare ball model, we consider that the probability of repeating a response is influenced by the previous trial's response. Thus, we now have four free parameters:

$$P_{response|high} = \left\{ \begin{array}{c} P(r_N = h_+|\text{rare}, r_{N-1} = h_+) = P_{rareH} \\ P(r_N = h_-|\text{rare}, r_{N-1} = h_+) = 1 - P_{rareH} \\ P(r_N = h_-|\text{no rare}, r_{N-1} = h_+) = P_{noH} \\ P(r_N = h_+|\text{no rare}, r_{N-1} = h_+) = 1 - P_{noH} \end{array} \right\},$$

$$P_{response|low} = \left\{ \begin{array}{c} P(r_{=}N = h_+|\text{rare}, r_{N-1} = h_-) = P_{rareL} \\ P(r_N = h_-|\text{rare}, r_{N-1} = h_-) = 1 - P_{rareL} \\ P(r_N = h_-|\text{no rare}, r_{N-1} = h_-) = P_{noL} \\ P(r_N = h_+|\text{no rare}, r_{N-1} = h_-) = 1 - P_{noL} \end{array} \right\}.$$

Guess Model. The guess model includes one free parameter that determines the probability of responding with the high jar.

$$P(r = h_+) = P_{guess}, \qquad P(r = h_-) = 1 - P_{guess}.$$

4.3 Model Fitting and Comparison

Prior to data acquisition, pilot and synthetic data from the proposed models was used to confirm that models were identifiable under the task conditions (Fig. S11).

Model parameter fits were performed using Bayesian maximum likelihood estimation, where the parameter posteriors were generated by considering the response data $r_{1:42}$ and observation vectors $\boldsymbol{\xi}_{1:42}$ across all 42 trials of a block. For instance, to infer the noise variance a and rare ball weight ρ of the noisy Bayesian model, we could apply Bayes rule and then compute the probability of a response r_N in a given trial conditioned on observations $\boldsymbol{\xi}_N$ as

$$p(a, \rho|r_N, \boldsymbol{\xi}_N) = \frac{p(r_N|a, \rho, \xi_N)p(a, \rho)}{p(r_N|\xi_N)}.$$

Since the denominator only provides a normalization of the probability densities of a and ρ, the primary contributions are the probability of a response r_N given the parameters and observations times the parametric priors. Below, we explain how we set the parametric priors. Since all models were either simple binary random variables or thresholded Gaussians, we could evaluate the associated likelihood functions analytically. These were used as functions in MATLAB then to compute parametric posteriors across blocks

$$p(a, \rho|r_{1:42}, \boldsymbol{\xi}_N) = \frac{p(a, \rho)}{p(r_{1:42}|\boldsymbol{\xi}_{1:42})} \prod_{j=1}^{42} p(r_j|a, \rho, \boldsymbol{\xi}_j).$$

The maximum of this function is the maximum likelihood estimate, and the portion of the posterior within 5% of the maximum is the credible interval we used.

Model comparisons were computed using log Bayes factors comparing the likelihood that a particular data set came from one of two models by computing the log likelihood ratio of the marginal likelihoods for two possible models.

$$\log \text{BF} = \log \frac{P(D|M_1)}{P(D|M_2)} = \log \frac{P(M_1|D)P(M_2)}{P(M_2|D)P(M_1)},$$

where D is the given data for a block ($r_{1:42}$ and $\boldsymbol{\xi}_{1:42}$), and M_1 and M_2 are two possible models. For example, to compare the noisy Bayesian model to the prior ideal observer model for a given block, we would integrate the probability of responses conditioned on observations and parameters against the parametric priors:

$$\log \mathrm{BF} = \log \frac{\int_{-\infty}^{\infty}\int_{0}^{\infty} p(r_{1:42}|a, z_0, \boldsymbol{\xi}_{1:42})p(a, z_0)dadz_0}{\int_{0}^{\infty}\int_{0}^{\infty} p(r_{1:42}|a, \rho, \boldsymbol{\xi}_{1:42})p(a, \rho)dad\rho}.$$

For all comparisons, we assumed a null hypothesis of the noisy Bayesian model. Prior to performing our model comparisons on data, we had run them on synthetic data, sampling models from parametric priors. We found that two candidate models (one which assumed an asymmetric repetition bias and one that included a biased prior and free parameter for rare ball weight) were not identifiable, and so excluded it from our model comparison analysis of the data (Fig.S12).

Based on analyses with synthetic and pilot data, we found that the assuming a flat prior over parameter space led to inconclusive parameter identification in the noisy Bayesian model due to dependencies in noise variance and rare ball weight parameters (Fig. S20A). To mitigate this effect in our human data analysis, we used pilot data from 20 subjects to create an informative prior based on subjects' posteriors that were weakened and smoothed (Fig. S20B and C).

Testing these informative priors on synthetic data, we found that credible intervals (parameter sets with the top 5% of all likelihoods in a posterior) for our best fit parameters usually contained true parameters (Fig. S21). While applying informative priors did limit identifiability at high noise variance and rare ball weight values, we ensured that our models could correctly identify parameters of interest (namely the rare ball weight ρ) near or below the values predicted by the ideal observer.

Subjects' best fit models were determined by computing the $\log$ Bayes factor between the noisy Bayesian model and each alternative model. Positive values were evidence in favor of a particular alternative model over the noisy Bayesian model, with larger values being stronger evidence. The best fit model was selected based on the maximal Bayes factor value across all alternative models. If no values were greater than 0, the noisy Bayesian model was selected.

4.4 Information Theory

We used rate distortion theory to compare subjects accuracy to a maximal accuracy bound with respect to the information their strategy used. The mutual information (MI) between an observer's response and underlying task features (observations and possibly prior responses) is

$$I(f_B^i, r_B^i) = \sum_{f_B^i \in F}\sum_{r_B^i \in R} P(f_B^i, r_B^i)\log_2 \frac{P(f_B^i, r_B^i)}{P(f_B^i)P(r_B^i)},$$

where f describes the task features, i is a subject or model, B is the block, and r_B^i describes the responses. Two sets of task features were defined: the balls observed on the response trial $I_{O,B}^i$ and the observations combined with the observer's response on the previous trial $I_{OR,B}^i$. For example, when computing MI using just the observations, we summed over all the viewed samples, disregarding ball order, ($n+1$ possible counts for trials with $n = 2, 5, 10$ ball draws) on the 42 trials of a block

$$I(f_B^i, r_B^i) = \sum_{f_B^i \in F} P(f) \sum_{r_B^i \in R} \pi(r_B^i|f_B^i)\log_2 \frac{\pi(r_B^i|f_B^i)}{p(r_B^i)}$$

where $\pi(r_B^i|f_B^i)$ is the policy used to generate responses across the block.

In addition, f_B^i could also encode both the observations and the previous response. An increase in MI with the inclusion of the previous responses implies a subject is considering irrelevant past information. Each subject's trials within a block were bootstrapped 1000 times to estimate the MI and accuracy for the block.

Computing the Optimal Bound. The rate distortion bound can be computed according to a constrained optimization problem in which we identify the maximum possible accuracy for a given level of MI. In the ideal observer case, the policy applied to compute MI is

$$\pi(r = h_+|m) = \begin{cases} 1, & m > B \\ 0.5, & m = B, \\ 0, & m < B. \end{cases}$$

where m is the count of rare balls observed and B is the number required to trigger a high jar response. Additionally, we must compute the predictive accuracy using the value function

$$V(f, r) = \sum_f p(f) \sum_{r \in \{h_+, h_-\}} \pi(r|f)Q(f, r),$$

which sums over all possible combinations of observation samples ($n = 2, 5, 10$ in a trial and $n + 1$ possible unordered samples for each trial length) for which we can always compute the trial specific value function from the ideal observer $Q(f, r) = 1/(1 + \mathrm{e}^{-z_n})$, where z_n is the ideal observer's log likelihood ratio.

Combining the above details, we find

$$\begin{aligned} V(f,r) =& 0.5p(f = B) + 0.5 \sum_{m<B} \binom{n}{m} \frac{h_+^m(1-h_+)^{n-m} + h_-^m(1-h_-)^{n-m}}{1+\mathrm{e}^{z_n}} \\ &+ 0.5 \sum_{m>B} \binom{n}{m} \frac{h_+^m(1-h_+)^{n-m} + h_-^m(1-h_-)^{n-m}}{1+\mathrm{e}^{-z_n}}. \end{aligned}$$

Thus, to bound accuracy for a given MI $I(f, r) \leq C$, we maximize the value function

$$V^* = \max V(f, r) \quad \text{s.t.} \quad I(f, r) = C,$$

which generates the optimal predicted bounds.

Model Mapping To identify which regimes of MI-accuracy space each of our alternative models exist, we computed MI and accuracy values for 1000 synthetic datasets of each model and binned across MI- accuracy (Rate distortion) space. Each bin was assigned to the model with highest number datasets associated with that bin.

The same procedure was used for groups of models. Subject model and group maps were based on bootstrapped subject MI and accuracy values and subjects' best fit models.

Computing Bounds for Alternative Models

Unlike the ideal observer bound, which was computed explicitly, our alternative model bounds were computed using 1000 synthetic datasets with randomized parameters and block lengths matching the task structure. Initial bootstrapped MI-accuracy values were computed for all 1000 datasets and the top 10% of accuracy values in each MI bin were identified. The model parameters from these maximal values were then used to regenerate 1000 new synthetic data sets with increased block lengths of 1000 trials. The mean for the top 10% of these new binned values was then used to define the maximal bound for each alternative model.

Subject Deviation from Ideal Observer Based on subjects binning in MI-accuracy space, 1000 bootstrapped subject responses were selected and compared to ideal observer responses. The percentage of matching responses were reported.

4.5 Statistics

Population statistics were computed using bootstrapping to identify the mean and confidence intervals. Subject p-value statistics were computed using bootstrapping and applying the Benjamini-Hochberg correction for multiple comparisons.

References

[1] Tversky, A. & Kahneman, D. Judgment under uncertainty: Heuristics and biases. *Science* **185**, 1124–1131 (1974).

[2] Gigerenzer, G. & Brighton, H. Homo heuristicus: Why biased minds make better inferences. *Topics in cognitive science* **1**, 107–143 (2009).

[3] Hilbert, M. Toward a synthesis of cognitive biases: how noisy information processing can bias human decision making. *Psychological bulletin* **138**, 211 (2012).

[4] Yeshurun, Y., Carrasco, M. & Maloney, L. T. Bias and sensitivity in two-interval forced choice procedures: Tests of the difference model. *Vision Res.* **48**, 1837–51 (2008).

[5] Green, D. M., Swets, J. A. *et al. Signal detection theory and psychophysics*, vol. 1 (Wiley New York, 1966).

[6] Hertwig, R., Barron, G., Weber, E. U. & Erev, I. Decisions from experience and the effect of rare events in risky choice. *Psychological Science* **15** (2004).

[7] Barron, G. & Yechiam, E. The coexistence of overestimation and underweighting of rare events and the contingent recency effect. *Judgement and Decision Making* **4**, 447–460 (2009).

[8] Camilleri, A. R. & Newell, B. R. When and why rare events are underweighted: a direct comparison of the sampling, partial feedback, full feedback and description choice paradigm. *Psychon Bull Rev* (2011).

[9] Kahneman, D. & Tversky, A. Prospect theory: An analysis of decision under risk. *Econometrica* **47**, 263–292 (1979).

[10] Hertwig, R. & Erev, I. The description-experience gap in risky choice. *Trends in Cog. Sci.* **13**, 517–523 (2009).

[11] Wulff, D. U., Mergenthaler-Canseco, M. & Hertwig, R. A meta-analytic review of two modes of learning and the description-experience gap. *Psychological Bulletin* **144**, 140–176 (2018).

[12] Fox, C. R. & Hadar, L. " decisions from experience"= sampling error+ prospect theory: Reconsidering hertwig, barron, weber & erev (2004). *Judgment and Decision Making* **1**, 159 (2006).

[13] Barron, G. & Erev, I. Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making* **16** (2003).

[14] Lieder, F., Griffiths, T. L. & Hsu, M. Overrepresentation of extreme events in decision making reflects rational use of cognitive resources. *Psychological review* **125**, 1 (2018).

[15] Hau, R., Pleskac, T., Kiefer, J. & Hertwig, R. The description-experience gap in risky choice: the role of sample size and experienced probabilities. *J. Behav. Decis. Making* **21** (2008).

[16] Rakow, T., Demes, K. A. & Newell, B. R. Biased samples not mode of presentation: Re-examining the apparent underweighting of rare events in experience-based choice. *Organizational Behavior and Human Decision Processes* **106**, 168–179 (2008).

[17] Ungemach, C., Chater, N. & Stewart, N. Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)? *Psychological Science* **20** (2009).

[18] Jessup, R. K., Bishara, A. J. & Busemeyer, J. R. Feedback produces divergence from prospect theory in descriptive choice. *Psychological Science* **19**, 1015–1022 (2008).

[19] Glöckner, A., Hilbig, B., Henninger, F. & Fiedler, S. The reversed description-experience gap: Disentangling sourced of presentation format effects in risky choice. *J. Exp. Psych.* **145**, 486–508 (2016).

[20] Plonsky, O., Teodorescu, K. & Erev, I. Reliance on small samples, the wavy recency effect, and simlarity-based learning. *Psychological Review* **122**, 621–647 (2015).

[21] Szollosi, A., Liang, G., Konstantinidis, E., Donkin, C. & Newell, B. R. Simultaneous underweighting and overestimation of rare events: unpacking a paradox. *J. Exp. Psych* **148**, 2207–2217 (2019).

[22] Kellen, D., PAchur, T. & Hertwig, R. How (in)variant are subjective representations of descrived and experienced risk and rewards? *Cognition* **157**, 126–138 (2016).

[23] Urai, A. E., Braun, A. & Donner, T. H. Pupil-linked arousal is driven by decision uncertainty and alters serial choice bias. *Nature communications* **8**, 1–11 (2017).

[24] Abrahamyan, A., Silva, L. L., Dakin, S. C., Carandini, M. & Gardner, J. L. Adaptable history biases in human perceptual decisions. *Proceedings of the National Academy of Sciences* **113**, E3548–E3557 (2016).

[25] Meyniel, F., Maheu, M. & Dahaene, S. Human inferences about sequences: A minimal transition probability model. *Plos Comp. Biol.* **12** (2016).

[26] Kass, R. E. & Raftery, A. E. Bayes factors. *Journal of the american statistical association* **90**, 773–795 (1995).

[27] Berger, T. Rate-distortion theory. *Wiley Encyclopedia of Telecommunications* (2003).

[28] Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E. & Shadlen, M. N. Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *Journal of Neuroscience* **31**, 6339–6352 (2011).

[29] Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W. & Forstmann, B. U. Bias in the brain: a diffusion model analysis of prior probability and potential payoff. *Journal of Neuroscience* **32**, 2335–2343 (2012).

[30] Fan, Y., Gold, J. I. & Ding, L. Ongoing, rational calibration of reward-driven perceptual biases. *Elife* **7**, e36018 (2018).

[31] Mountcastle, V. B., Talbot, W. H., Sakata, H. & Hyvärinen, J. Cortical neuronal mechanisms in flutter-vibration studied in unanesthetized monkeys. neuronal periodicity and frequency discrimination. *Journal of neurophysiology* **32**, 452–484 (1969).

[32] Mountcastle, V., Steinmetz, M. & Romo, R. Frequency discrimination in the sense of flutter: psychophysical measurements correlated with postcentral events in behaving monkeys. *Journal of Neuroscience* **10**, 3032–3044 (1990).

[33] Odoemene, O., Pisupati, S., Nguyen, H. & Churchland, A. K. Visual evidence accumulation guides decision-making in unrestrained mice. *Journal of Neuroscience* **38**, 10143–10155 (2018).

[34] Glaze, C. M., Kable, J. W. & Gold, J. I. Normative evidence accumulation in unpredictable environments. *Elife* **4**, e08825 (2015).

[35] Veliz-Cuba, A., Kilpatrick, Z. P. & Josić, K. Stochastic models of evidence accumulation in changing environments. *SIAM Review* **58**, 264–289 (2016).

[36] Eissa, T. L., Barendregt, N. W., Gold, J. I., Josić, K. & Kilpatrick, Z. P. Hierarchical inference interactions in dynamic environments (Cosyne, 2019).

[37] Kilpatrick, Z. P., Holmes, W. R., Eissa, T. L. & Josić, K. Optimal models of decision-making in dynamic environments. *Current Opinion in Neurobiology* **58**, 54–60 (2019).

[38] Radillo, A. E., Veliz-Cuba, A., Josić, K. & Kilpatrick, Z. P. Evidence accumulation and change rate inference in dynamic environments. *Neural computation* **29**, 1561–1610 (2017).

[39] Glaze, C. M., Filipowicz, A. L., Kable, J. W., Balasubramanian, V. & Gold, J. I. A bias–variance trade-off governs individual differences in on-line learning in an unpredictable environment. *Nature Human Behaviour* **2**, 213–224 (2018).

[40] Livina, V. N., Havlin, S. & Bunde, A. Memory in the occurrence of earthquakes. *Physical review letters* **95**, 208501 (2005).

[41] Clauset, A., Young, M. & Gleditsch, K. S. On the frequency of severe terrorist events. *Journal of Conflict Resolution* **51**, 58–87 (2007).

[42] Williams, J. J. & Griffiths, T. L. Why are people bad at detecting randomness? a statistical argument. *J. Exp. Psychol.* **39**, 1473–1490 (2013).

[43] Cassey, P., Hawkins, G. E., Donkin, C. & Brown, S. D. Using alien coins to test whether simple inference is bayesian. *Journal of Experimental Psychology: Learning, Memory, and Cognition* (2015).

[44] Hwang, E. J., Dahlen, J. E., Mukundan, M. & Komiyama, T. History-based action selection bias in posterior parietal cortex. *Nat. Comm.* **8** (2017).

[45] Fründ, I., Wichmann, F. A. & Macke, J. H. Quantifying the effect of intertrial dependence on perceptual decisions. *Journal of Vision* **14** (2014).

[46] Miller, J. B. & Sanjuro, A. Surprised by the hot hand fallacy? a truth in the law of small numbers. *Econometrica* **86**, 2019–2047 (2018).

[47] Nevo, I. & Erev, I. On surprise, change, and the effect of recent outcomes. *Frontiers in Psychology* **3** (2012).

[48] Gilovich, T., Vallone, R. & Tversky, A. The hot hand in basketball: On the misperception of random sequences. *Cognitive Psychology* **17** (1985).

[49] T.D., K., M, K. & I., G. J. Coupled decision processes update and maintain saccadic priors in a dynamic environment. *J. Neurosci.* **37**, 3632–3645 (2017).

[50] Nguyen, K. P., Josić, K. & Kilpatrick, Z. Optimizing sequential decisions in the drift-diffusion model. *J. Math Psychol.* **88** (2019).

[51] Girshick, A. R., Landy, M. S. & Simoncelli, E. P. Cardinal rules: Visual orientation perception reflects knowledge of environmental statistics. *Nat. Neurosci.* **14**, 926–932 (2011).

[52] Goodie, A. S. & Fantino, E. Learning to commit or avoid the base-rate error. *Nature* **380**, 247–249 (1996).

[53] Goodie, A. S. & Fantino, E. What does and does not alleviate base-rate neglect under direct experience. *J. Behavioral Decision Making* **12**, 307–335 (1999).

[54] Welsh, M. B. & Navarro, D. J. Seeing is believing: Priors, trust, and base rate neglect. *Organizational Behavior and Human Decision Processes* **119**, 1–14 (2012).

[55] Bar-Hillel, M. The base-rate fallacy in probability judgments. *Acta Psychologica* **44**, 211–233 (1980).

[56] Hau, R., Pleskac, T. & Hertwig, R. Decisions from experience and statistical probabilities: Why they trigger different choices than a priori probabilities. *J. Behav. Decis. Making* **23** (2010).

[57] Weber, E. U., Shafir, S. & Blais, A. Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychol. Rev.* **111**, 430–445 (2004).

[58] Bradbury, M. A. S., Hens, T. & Zeisberger, S. Improving investment decisions with simulated experience. *Rev. Finance* **19** (2015).

[59] M.Hogarth, R. & Soyer, E. Communicating forecasts: The simplicity of simulated experience. *J. Business Research* **68** (2015).

[60] M.Hogarth, R. & Soyer, E. Providing information for decision making: Contrasting description and simulation. *J. Applied Research in Memory and Cognition* **4** (2015).

[61] Kaufmann, C., Weber, M. & Haisley, E. Public understanding of climate change in the united states. *Management Science* **59** (2012).

[62] Weber, E. U. & Stern, P. C. Public understanding of climate change in the united states. *American Psychologist* **66**, 315–328 (2011).

[63] FANTINO, E. & NAVARRO, A. Description–experience gaps: Assessments in other choice paradigms. *J. Behav. Dec. Making* **25** (2012).

[64] Lieder, F. & Griffiths, T. L. Resource-rational analysis: understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences* **43** (2020).

[65] Gershman, S. J. Origin of perseveration in the trade-off between reward and complexity. *bioRxiv* (2020).

[66] Gershman, S. J. & Lai, L. The reward-complexity trade-off in schizophrenia. *bioRxiv* (2020).

[67] Filipowicz, A. L., Glaze, C. M., Kable, J. W. & Gold, J. I. Pupil diameter encodes the idiosyncratic, cognitive complexity of belief updating. *Elife* **9**, e57872 (2020).

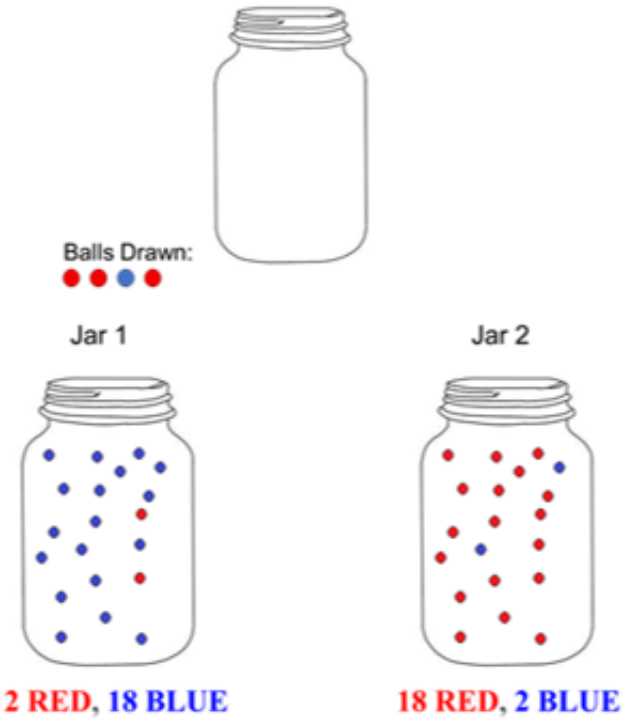

Figure S1: Example of the screen viewed by subjects on Amazon Mechanical Turk. Prompt at the bottom of the screen would request subjects to select the jar in use.

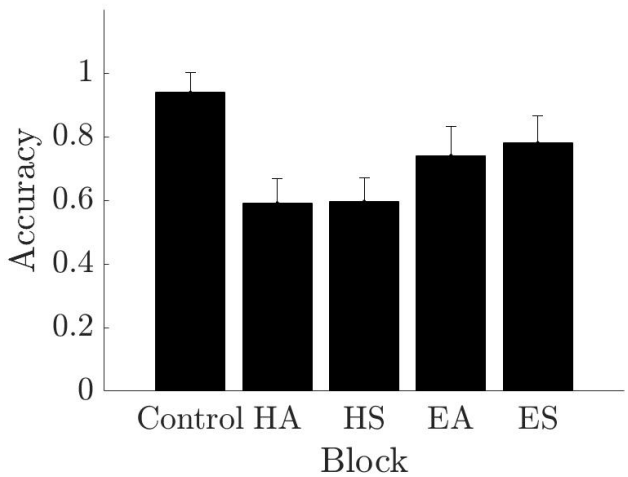

Figure S2: Mean and standard deviation of subjects overall accuracy for each block. Jar probabilities displays as h_-/h_+. CT (Control)-0.1/0.9; HA (Hard Asymmetric)-0.1/0.2; HS (Hard Symmetric)-0.45/0.55; EA (Easy Asymmetric)-0.1/0.4; ES (Easy Symmetric)-0.3/0.7. Overall accuracy depends on block difficulty rather than symmetry.

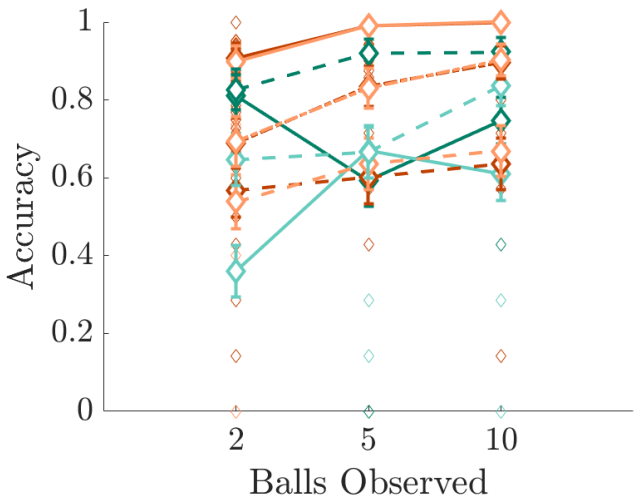

Figure S3: Plot as in Fig. 2D for the matched ideal observer. Dark traces represent trends in low jar accuracy, light traces note trends in high jar accuracy. Plot definitions as follows: CT-orange/full lines, HA-green/full lines, HS-orange/dashed lines, EA-green/dashed lines, ES-orange/dot dashed lines. Diamonds and errorbars represent bootstrapped population means and confidence intervals.

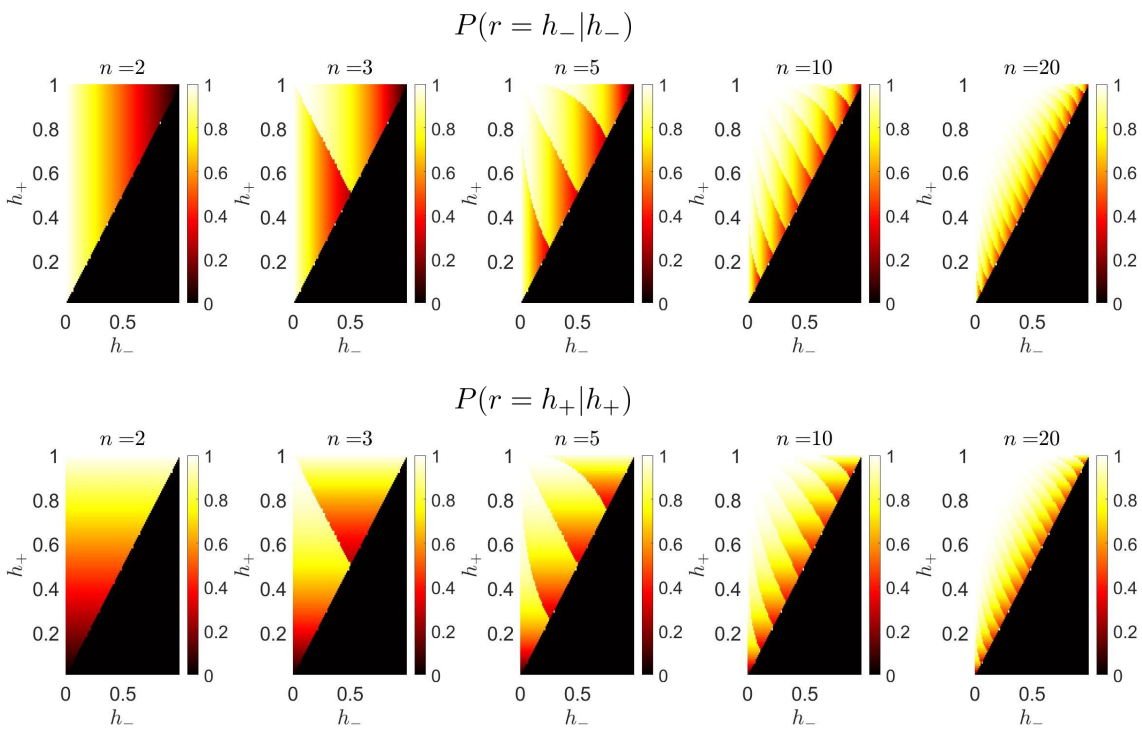

Figure S4: Accuracy heatmaps for the ideal observer while varying the number of observations (n). Black portion of plots correspond with jar probability pairs that do not meet our criteria ($h_+ > h_-$). Warmer colors correspond with higher accuracy. Top row: probability of a correct choice when the high (h_+) jar was selected. Bottom row: probability of a correct choice when the ow (h_-) jar was selected

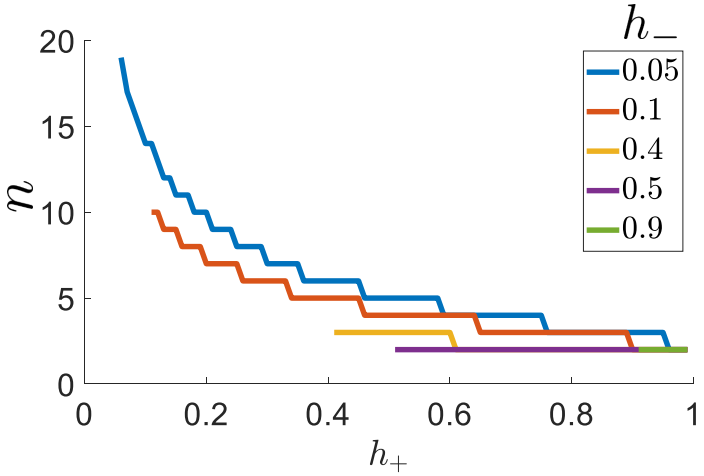

Figure S5: Estimated number of ball observations n until a rare ball is drawn with respect to various ball fraction h pairs. Rare asymmetric jars require a substantial number of observations to include a rare ball.

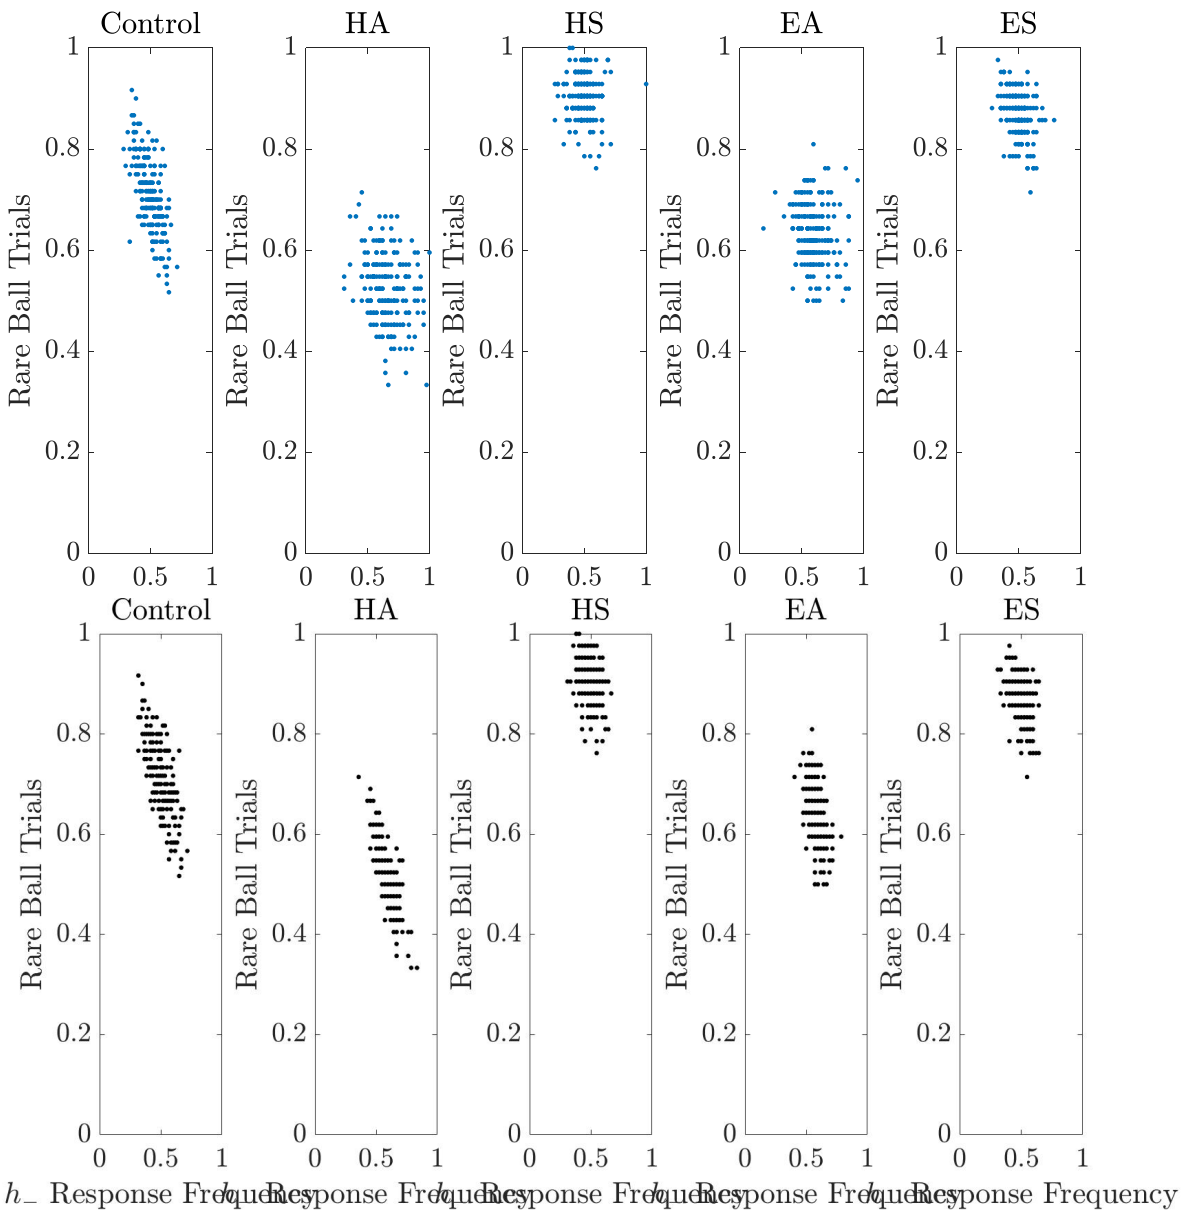

Figure S6: Fraction of rare ball trials compared to the low jar response frequency for individual subjects (Top, blue dots), and matched ideal observers (bottom, black dots).

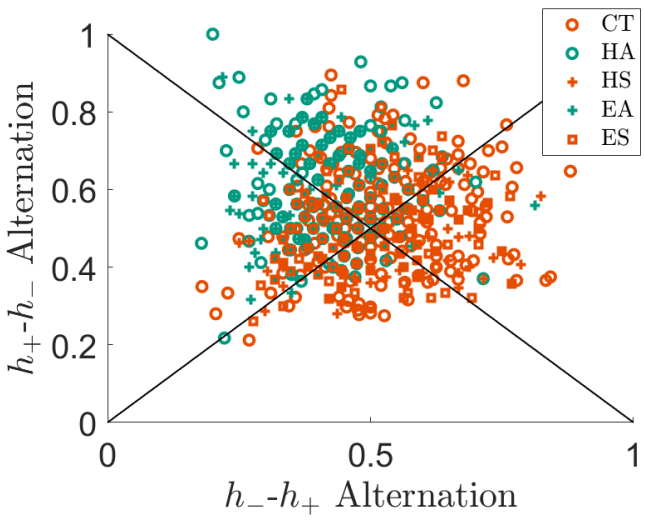

Figure S7: Alternation-bias plots for the matched ideal observer responses. Note the bias toward the low jar in asymmetric blocks (green).

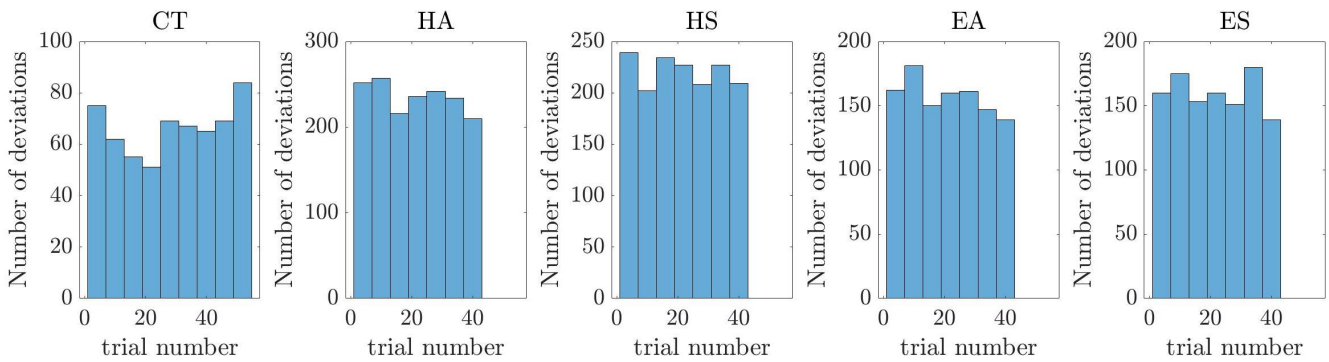

Figure S8: Histogram of frequency of subject deviations from the ideal observer based on trial number.

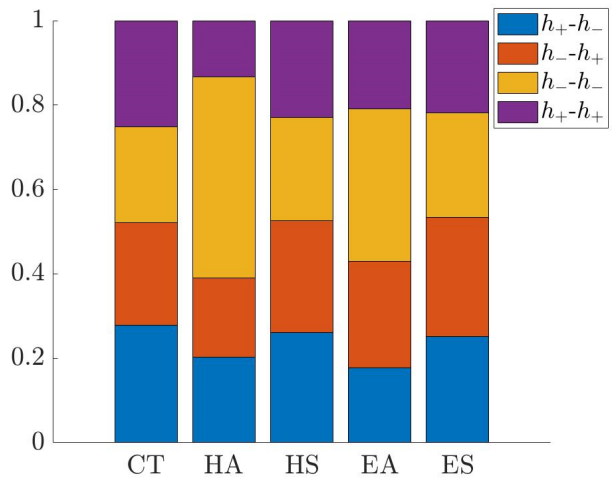

Figure S9: Fraction of deviation trials for each sequence (pre-deviation trial and deviation trial)

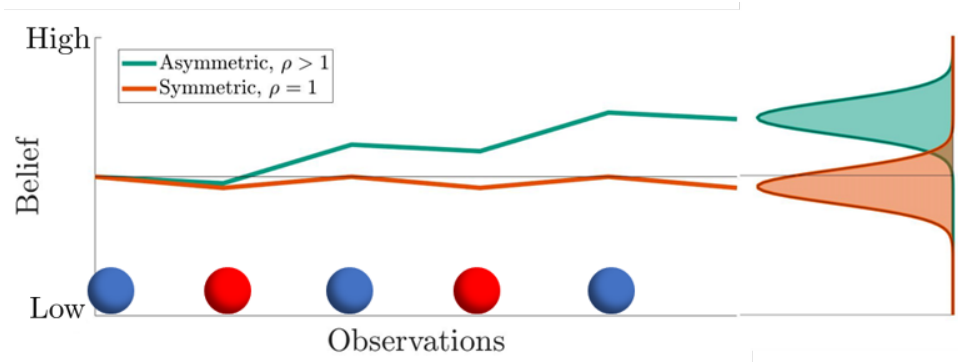

Figure S10: Schematic of the noisy Bayesian model whose mean evolves according to the solid lines in either case, but the beliefs are distributed according to the normal distributions shown.

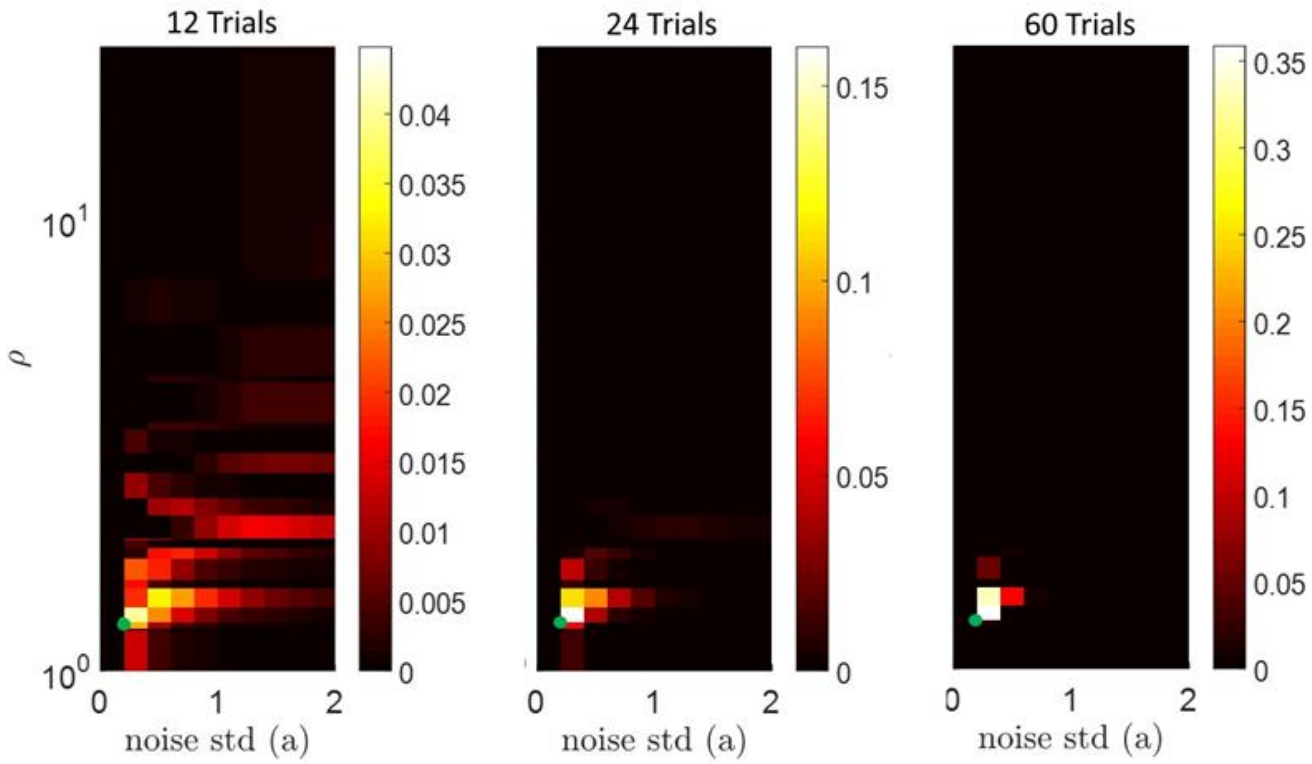

Figure S11: Example parameter fits when the block length is varied for the noisy Bayesian model. By 60 trials, the parameters are well identified.

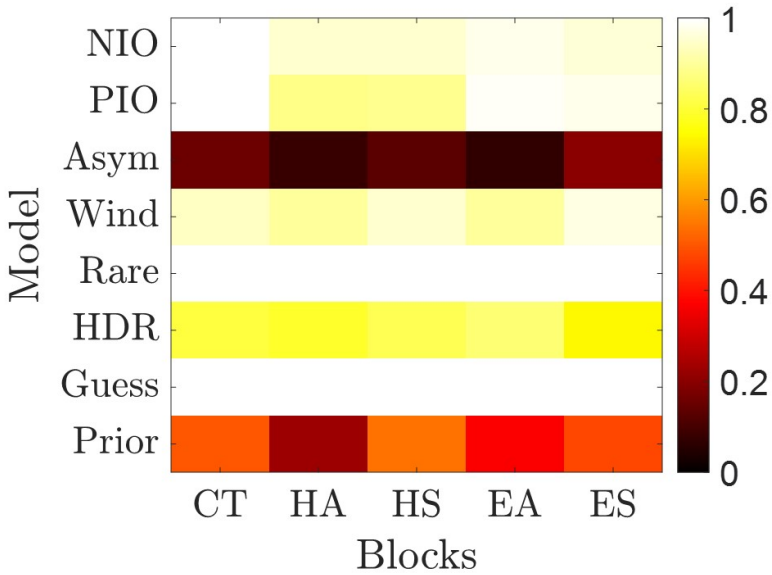

Figure S12: Fraction of times an alternative model was correctly identified as compared to the noisy Bayesian base model. Synthetic data produced for 100 subjects. Due to two of these models lacking identifiability (Asym and Prior), we eliminated them from the list of candidate models for subject analysis. NIO-assumes the ideal observer's rare ball weight ρ and only has a noise parameter. PIO-assumes the ideal observer's rare ball weight but a biased prior and noise. Asym-assumes an asymmetric repetition bias that occurs following an low jar response (not further used for analysis). Wind-assumed the subject applies a windowing procedure, that can include a particular number of observations (fewer or more than that given in the current trial) rather than information from a given trial exclusively. Rare-a rare ball model with some probability if a rare ball is observer. HDR-history dependent rare ball model. Guess-a guessing model. Prior-noisy Bayesian model with biased prior (not further used for analysis).

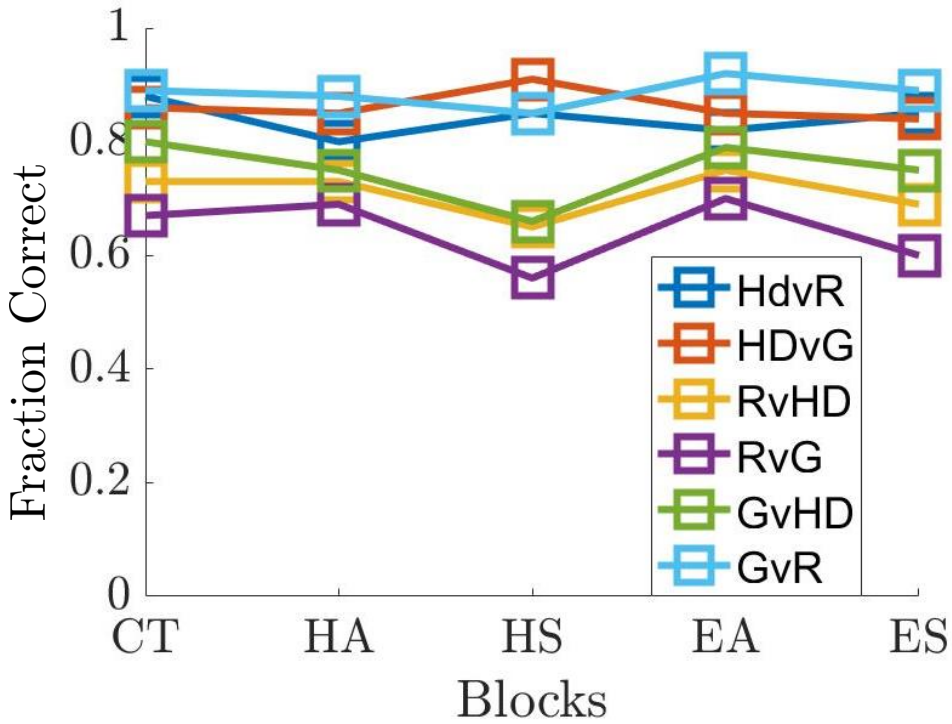

Figure S13: Fraction of times that heuristic models were correctly identified. Synthetic data produced for 100 subjects.

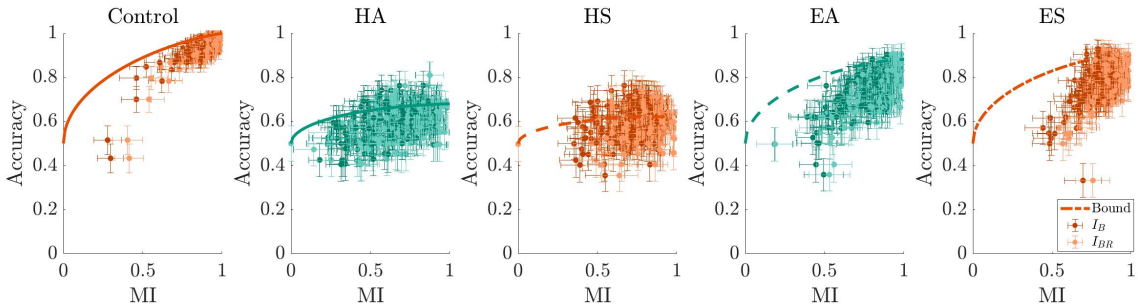

Figure S14: Subject rate distortion plots showing accuracy as a function of mutual information (MI) with mean and standard deviation of bootstrapped (1000 iterations) data.

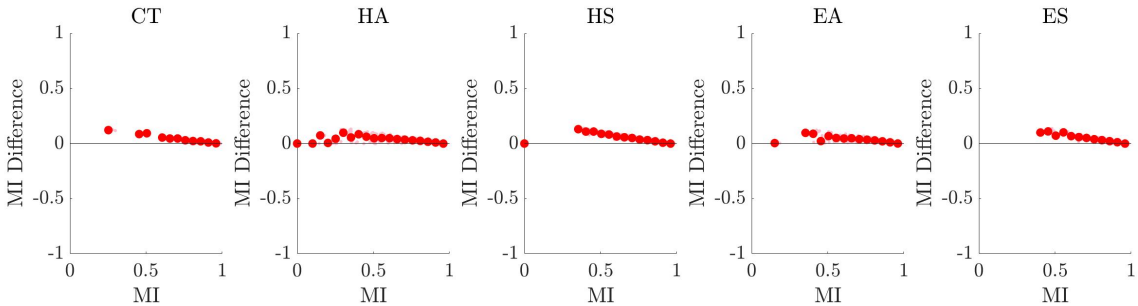

Figure S15: Mutual information difference for subjects comparing that involving prior responses (I_{BR}) and that involving only the balls in the current trial (I_B), $I_{BR} - I_B$, for individuals (pink points) and binned means (red points).

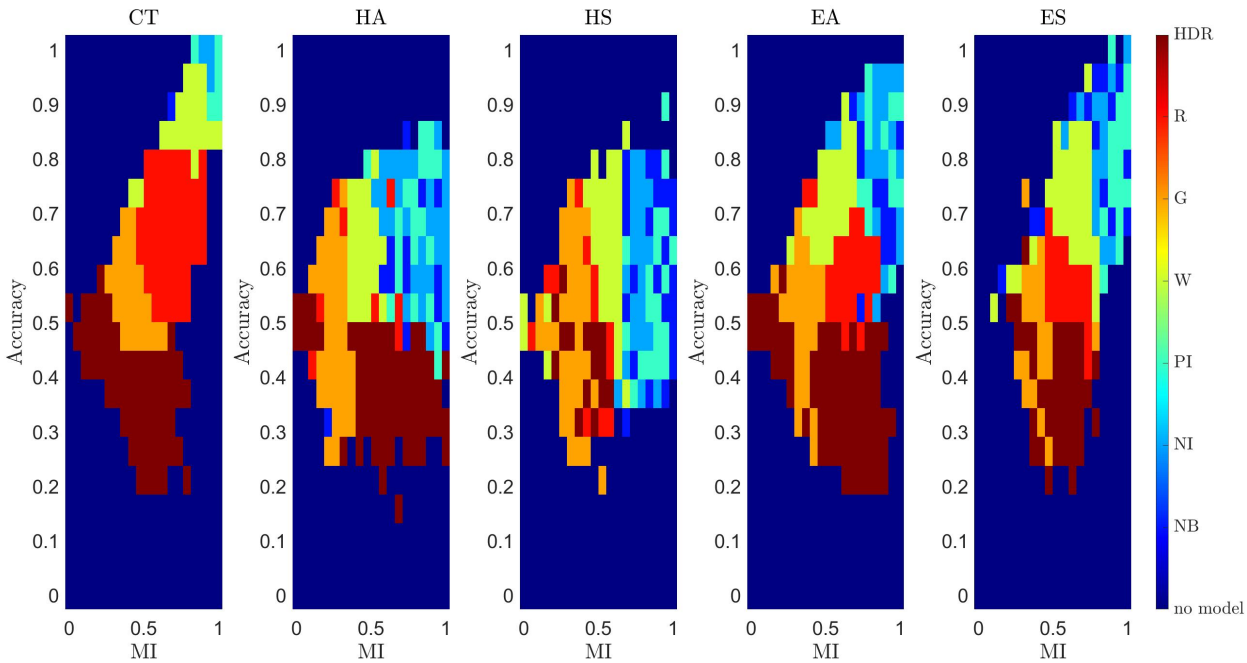

Figure S16: Map of most likely model in rate-distortion space based on synthetic data. All model labels as in Fig. 5.

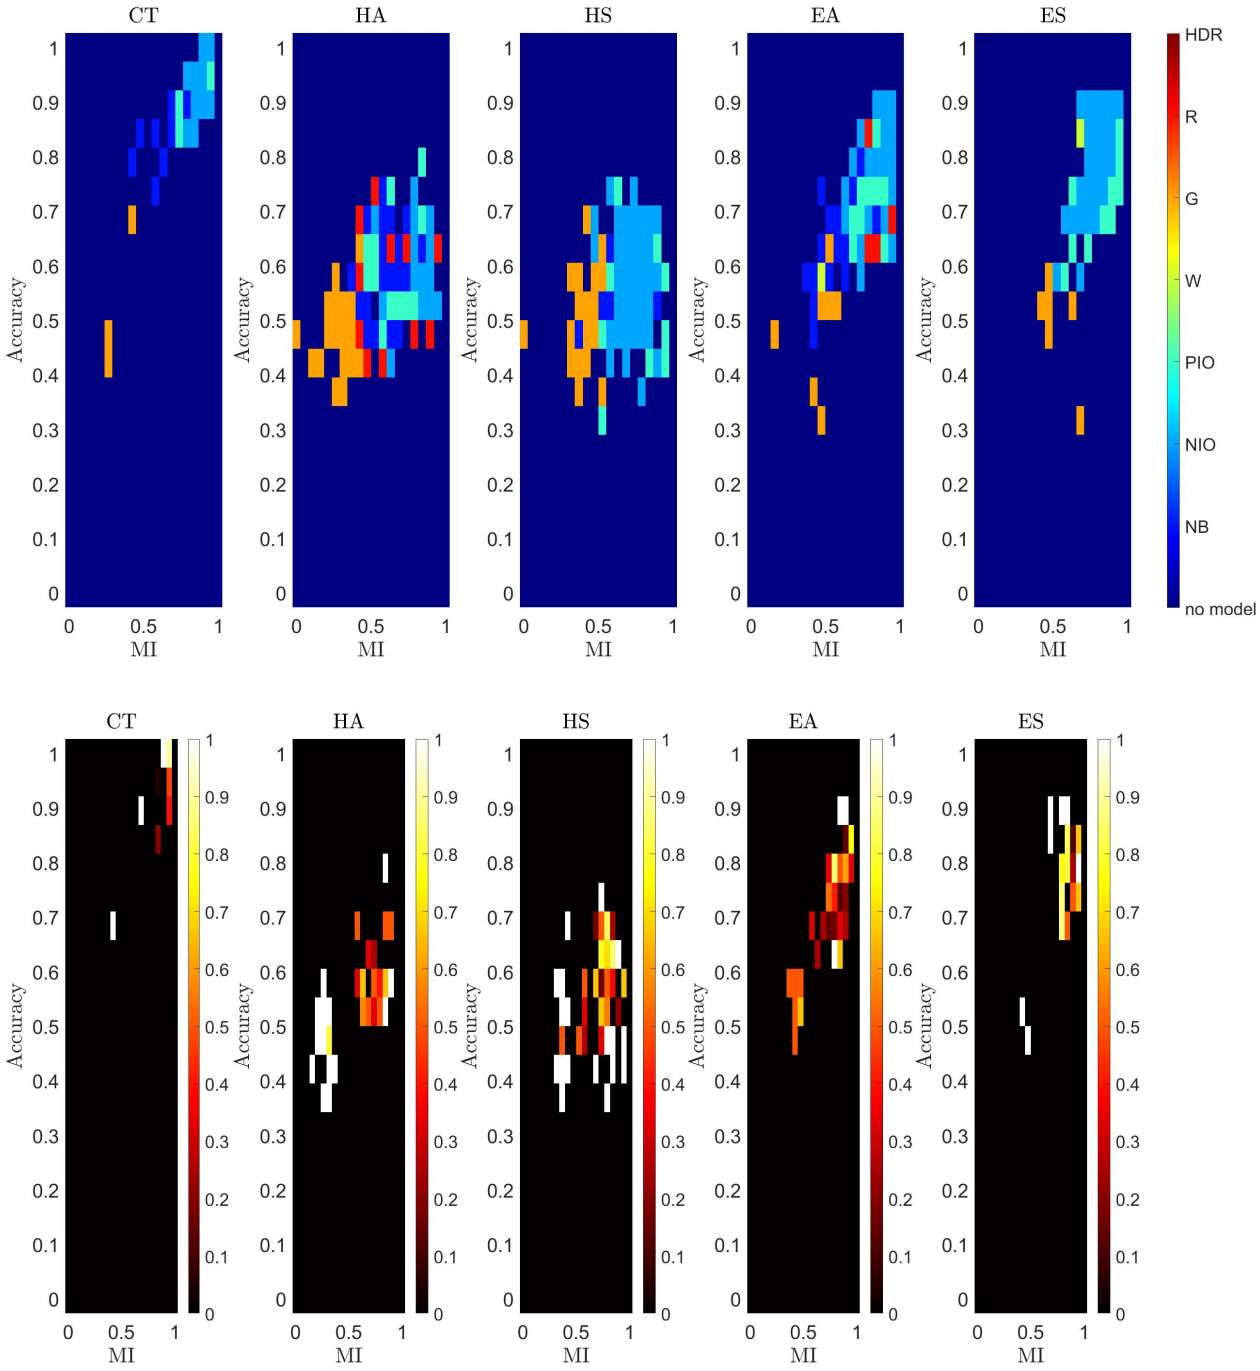

Figure S17: Top: Subject best fit models mapped to rate distortion space. Map shows most likely model at each point. Bottom: Fraction of subjects in a given bin whose most likely model matches the predicted model for that bin (based on the synthetic model maps from Fig. S16.

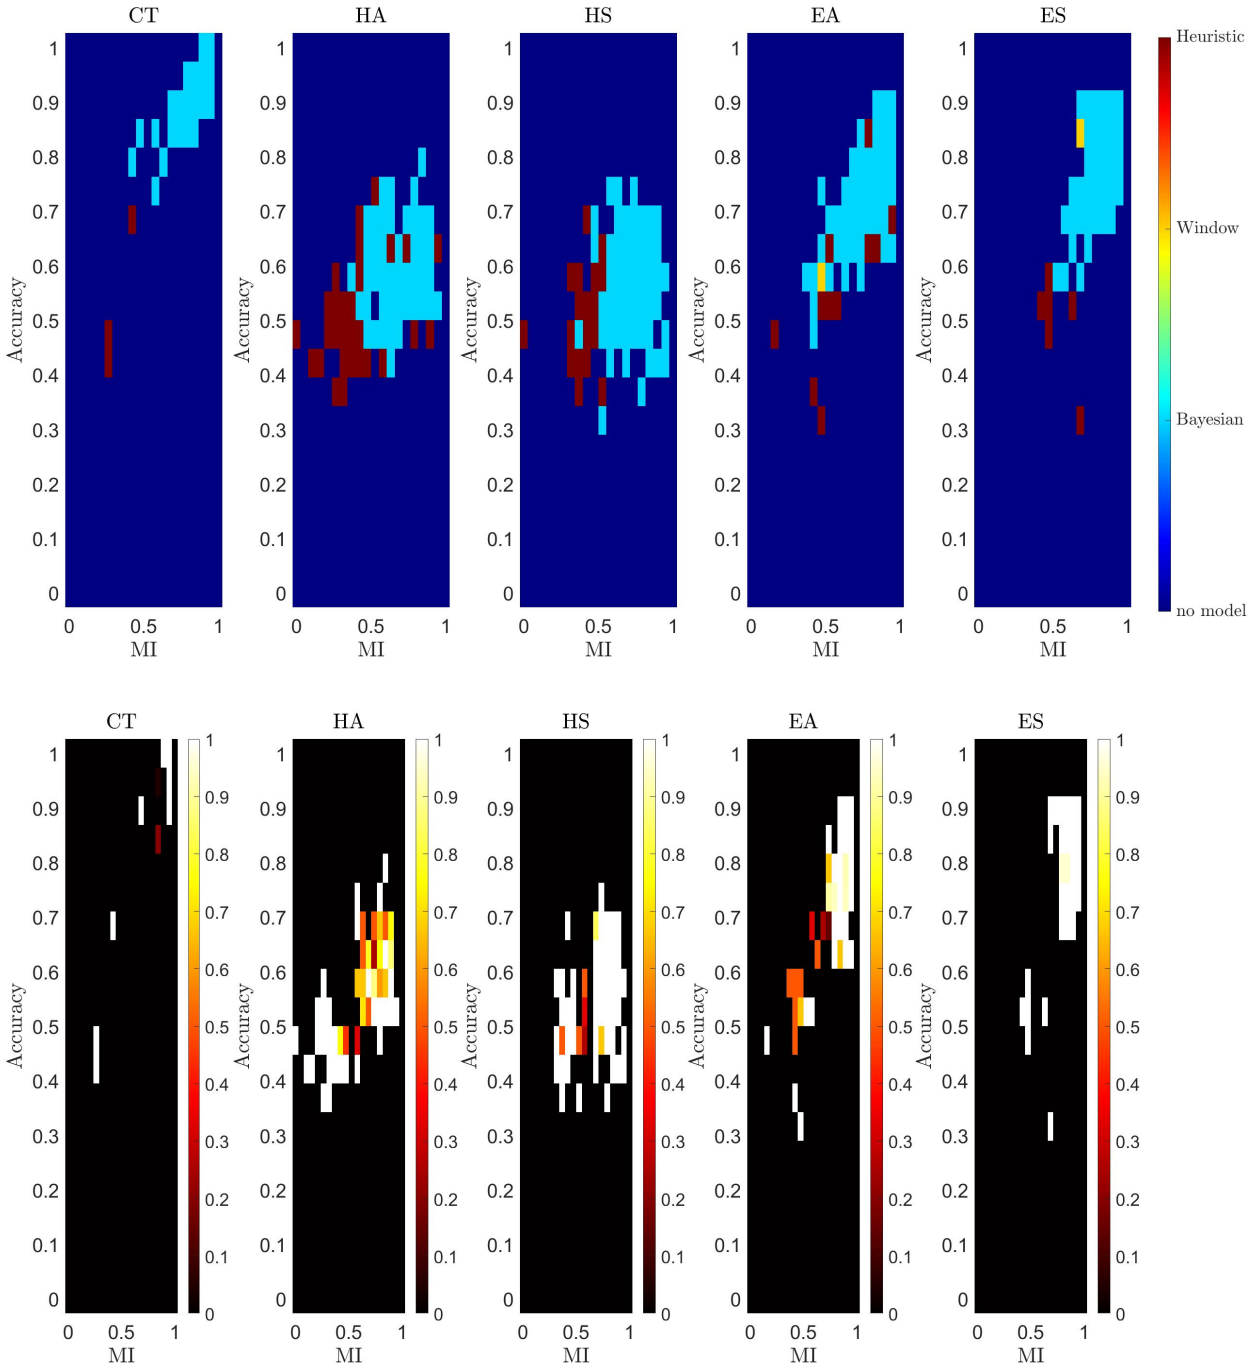

Figure S18: Top: Subject best fit models, by group (Bayesian, Window, Heuristic) mapped to rate distortion space. Map shows most likely model group at each point. Bottom: Fraction of subjects in a given bin whose most likely model group matches the predicted model group for that bin (based on the synthetic model maps from Fig. 6

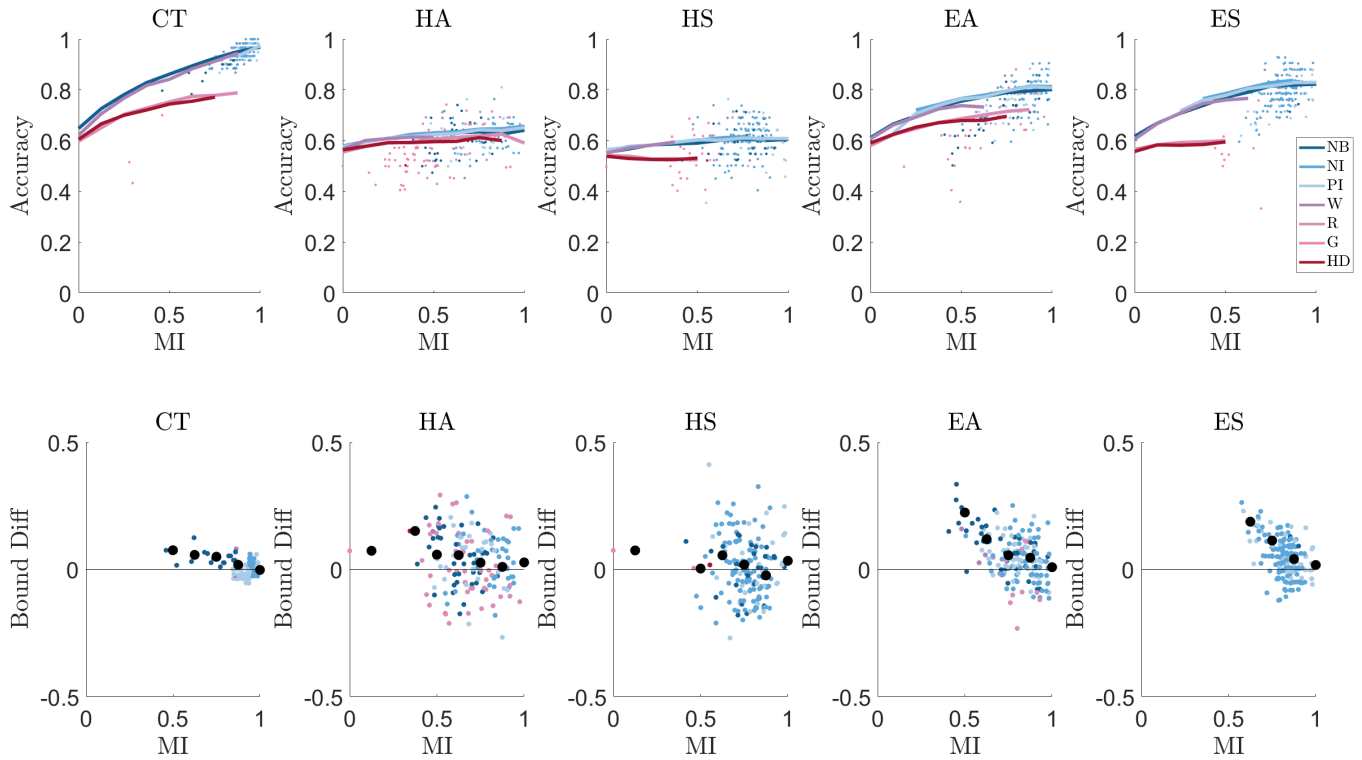

Figure S19: Top: Bounds for each model type computed using 1000 synthetic data sets with randomized parameters. Dots show subject data points color coded based on the most likely model. Bottom: Differences between the bound for a subject's most likely model at their MI value and subject accuracy. Black dots show binned average trend across models.

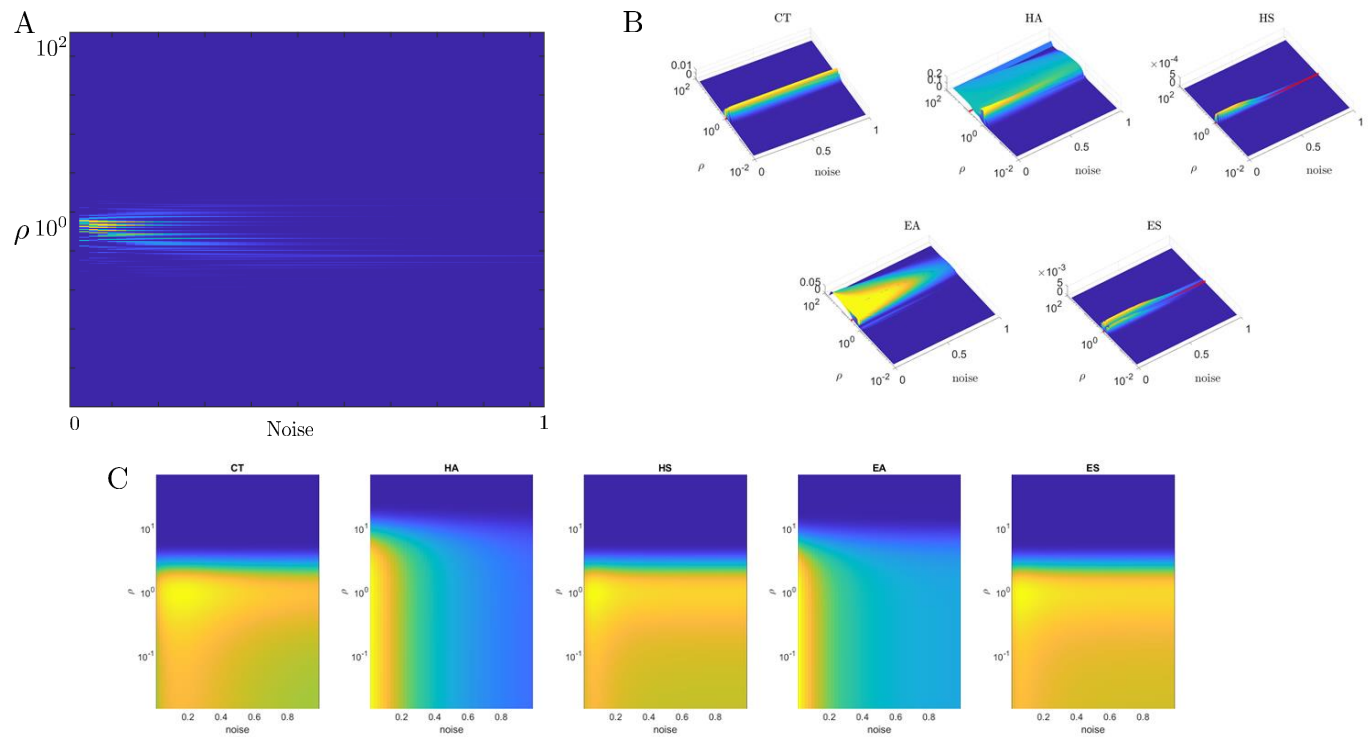

Figure S20: Examples for the weak informative prior. **A.** Example posterior for synthetic data from the noisy Bayesian model. While the true value of ρ is close to 1 with low noise, we see high likelihoods for alternative scenarios in which there is a tradeoff between noise and ρ. **B.** Averaged posteriors based on secondary human subject data set used to create informative priors. Red trace shows where the ideal observer's ρ value lies. **C.** Weakened informative priors based on **B**.

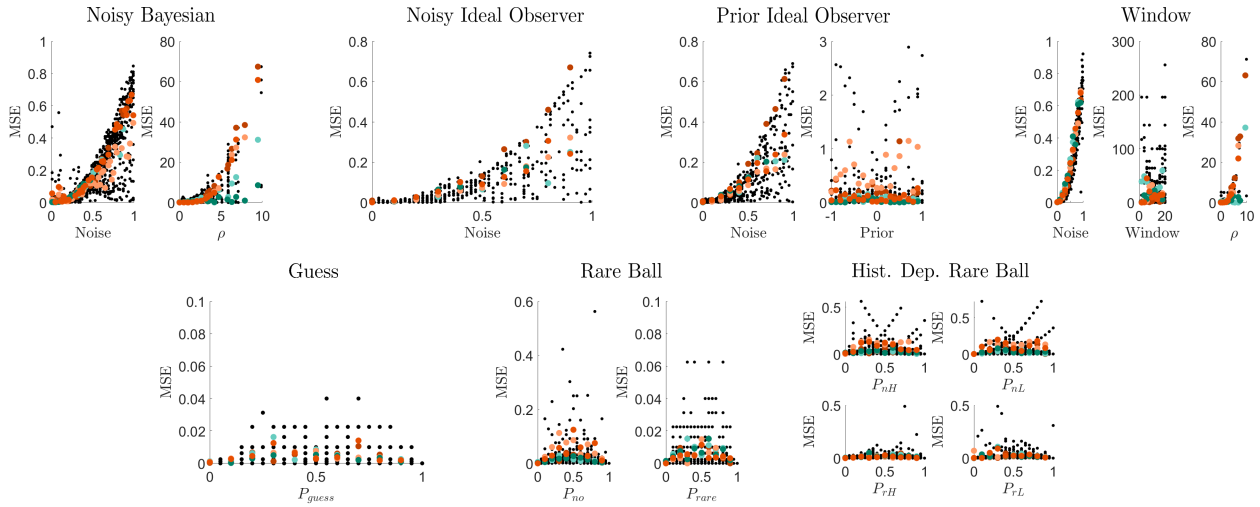

Figure S21: Parameter fits for each model based on synthetic data and applying the informative prior. Black traces show individual dataset fits. Colored dots show binned averages for each block (orange-symmetric blocks, green-asymmetric blocks). Top row: Bayesian models. Note that ρ is well fit for values near or below the ideal observer's rare ball weight. Thus, asymmetric averages (green dots) lay closer to zero at higher ρ values. Bottom row: Heuristic model parameter fits.