Testing models of anchoring and adjustment

Falk Lieder^{1,2,6}, Thomas L. Griffiths^{1,5}, Quentin J. M. Huys^{2,4} and
Noah D. Goodman³

¹Helen Wills Neuroscience Institute, University of California, Berkeley

²Translational Neuromodeling Unit, Institute for Biomedical
Engineering, University of Zürich and Swiss Federal Institute of
Technology (ETH) Zürich

³Department of Psychology, Stanford University

⁴Department of Psychiatry, Psychotherapy and Psychosomatics,
Hospital of Psychiatry, University of Zürich

⁵Department of Psychology, University of California, Berkeley

⁶Correspondence should be addressed to falk.lieder@berkeley.edu.

Abstract

This technical report compares alternative computational models of numerical estimation using Bayesian model selection. We find that people's estimates are best explained by a resource-rational model of anchoring and adjustment according to which the number of adjustments increases with error cost but decreases with time cost so as to achieve an optimal speed-accuracy tradeoff.

Keywords: bounded rationality; heuristics; cognitive biases; probabilistic reasoning; anchoring-and-adjustment; rational process models

Introduction

To make judgments under uncertainty people often rely on heuristics such as anchoring and adjustment (Tversky & Kahneman, 1974). Anchoring and adjustment is a two-stage process. In the first stage, people generate a preliminary judgment called their *anchor*. In the second stage, they adjust that judgment to incorporate additional information, but the adjustment is usually insufficient. Consequently, the resulting estimates tend to be biased towards the anchor. This phenomenon is known as the anchoring bias. Tversky and Kahneman (1974) demonstrated this cognitive bias by first asking people whether the percentage of African countries in the United Nations was smaller or larger than a randomly generated

number and then letting them estimate that unknown quantity. Strikingly, people’s median estimate was significantly larger when the random number was high than when it was low.

The anchoring bias has traditionally been interpreted as evidence against human rationality. By contrast, recent theoretical work (Lieder, Griffiths, & Goodman, 2012; Lieder, Griffiths, Huys, & Goodman, under review) and novel experiments (Lieder, Griffiths, Huys, & Goodman, under review; Lieder, Goodman, & Griffiths, 2013) suggested that people’s insufficient adjustments might be consistent with the rational use of their finite time and bounded computational resources (Griffiths, Lieder, & Goodman, 2015). Although this resource-rational theory can explain a wide range of anchoring phenomena (Lieder, Griffiths, Huys, & Goodman, under review), there are many alternative models of numerical estimation and the anchoring bias (Epley & Gilovich, 2006; Simmons, LeBoeuf, & Nelson, 2010; Strack & Mussweiler, 1997; Turner & Schley, 2016).

Here, we evaluate the resource-rational anchoring and adjustment model (Lieder et al., 2012; Lieder, Griffiths, Huys, & Goodman, submitted) against some of these alternative theories using Bayesian model selection (Stephan, Penny, Daunizeau, Moran, & Friston, 2009). To do so, we perform Bayesian model selection on the data from two anchoring by Lieder et al. (under review). Participants in these experiments predicted the departure time of a bus from the observation that it had not arrived by a given time and the distribution of bus departure times. One of the experiments anchored participants by asking them whether the bus’s delay would be shorter or longer than a certain number. In the other experiment participants generated their own anchors. Both experiments manipulated the cost of time and the cost of error and controlled participants’ prior knowledge. Consistent with the resource-rational anchoring and adjustment model, people’s anchoring biases in both experiments increased with time cost and decreased with error cost.

Computational models of anchoring and adjustment

We formalized four theories using eight probabilistic models of numerical estimation. The theories range from unbounded Bayesian rationality (theory 1) to random guessing (theory 4) with theories 2 and 3 formalizing intermediate levels of rationality: the sampling hypothesis (theory 3; Vul et al., 2014) and four models of the anchoring and adjustment heuristic that range from resource-rational anchoring and adjustment to less rational anchoring heuristics like the ones proposed by Epley and Gilovich (2006) and Simmons et al. (2010). To titrate exactly how rational our participants’ estimation strategies are, we determine which of these models best explains people’s adjustments from self-generated and provided anchors respectively. We start by presenting the computational models and then evaluate them on the self-generated and provided anchoring experiments by Lieder, Griffiths, Huys, and Goodman (submitted) using Bayesian model selection.

To titrate how rational people’s adjustments are, we consider four theories that range from fully rational to purely heuristic. According to the first theory, people draw Bayes-optimal inferences and the observed biases merely reflect a regression towards their prior expectation. We formalized this explanation in terms of Bayesian decision theory (m_{BDT} ; Equations 1-4). To connect the deterministic predictions of Bayesian decision theory to people’s variable responses, measurement and response errors are included in the model. According to the second theory, people approximate optimal inference by drawing a single sample from the posterior distribution (posterior probability matching, cf. Vul, Goodman,

Griffiths, & Tenenbaum, 2014, m_{PPM} , Equations 5-7). However, generating even a single perfect sample can require an intractable number of computations. Therefore, according to the third theory, the mind approximates sampling from the posterior by anchoring-and-adjustment (Lieder, Griffiths, & Goodman, 2012). We modeled adjustment using the probabilistic mechanisms illustrated in Figure 1. We modified the stopping criterion to model several variants of anchoring-and-adjustment. Existing theories of anchoring-and-adjustments commonly assume that people adjust their estimate until it is sufficiently plausible (Epley & Gilovich, 2006; Mussweiler & Strack, 1999; Simmons et al., 2010). Our first anchoring-and-adjustment model formalizes this assumption by terminating adjustment as soon as the estimate’s posterior probability exceeds a certain plausibility threshold ($m_{\text{A\&A}_s}$, Equations 8-16). The plausibility threshold and the average size of the adjustment are free parameters. According to the second anchoring-and-adjustment model, people make a fixed number of adjustments to their initial guess and report the result as their estimate ($m_{\text{A\&A}}$, Equations 17-24). Here the number of adjustments replaces the plausibility-threshold as the model’s second parameter. According to the third anchoring-and-adjustment model, people adapt the number of adjustments and the adjustment step size to optimize their speed-accuracy tradeoff ($m_{\text{aA\&A}}$, Equations 25-36; Lieder, Griffiths, & Goodman, 2013). The optimal speed-accuracy tradeoff depends on the unknown time $\tau_{\text{adjustment}}$ it takes to perform an adjustment, so this time constant is a free-parameter. The fourth anchoring-and-adjustment model extends the third one by assuming that there is an intrinsic error cost in addition to the extrinsic error cost imposed by the experimenter, and this intrinsic cost is an additional model parameter (m_{aAAi} , Equations 37-38). All anchoring models assumed that the anchor in Experiment 1 was the estimate reported in the previous section, that is 7.5 minutes. Finally, we also included a fourth theory. According to this “null hypothesis”, our participants chose randomly among all possible responses (m_{random} , Equation 39).

Except for the null model, the response distributions predicted by our models are a mixture of two components: the distribution of responses expected if people perform the task and the distribution of responses expected when they do not. The relative contributions of these two components are determined by an additional model parameter: the percentage of trials p_{cost} in which participants fail to perform the task. Not performing the task is modeled as random choice according to the null model. Performing the task is modeled according to the assumed estimation strategies described above. Each model consists of two parts: the hypothesized mechanism and an error distribution.

Bayes-optimal estimation

The first model (m_{BDT}) formalizes the hypothesis that people’s estimates are Bayes optimal. According to Bayesian decision theory, the optimal estimate of a quantity X given observation y is

$$\hat{x} = \arg \min_{\hat{x}} \mathbb{E}[\text{cost}(X, \hat{x}) | y]. \quad (1)$$

The error distribution accounts for both errors in reporting the intended estimate as well as trials in which people do not comply with the task and guess randomly. The model

combines these two types of errors with the Bayes-optimal estimate as follows:

$$R = \begin{cases} \hat{x} + \varepsilon, & \hat{x} = \arg \min_{\hat{x}} \mathbb{E} [\text{cost}(x, \hat{x})|y], \varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon), & \text{with prob. } 1 - p_{\text{cost}} \\ R \sim \text{Uniform}(\mathcal{H}), & & \text{with prob. } p_{\text{cost}} \end{cases}, \quad (2)$$

where R denotes people's responses based on y , p_{cost} is the probability that people guess randomly, $\mathcal{H}$ is their hypothesis space, and ε is people's error in reporting their intended estimate. This model has two free parameters: the probability p_{cost} that people guess randomly on a given trial and the standard deviation of the response error σ_ε . The model's prior distributions on these parameters are

$$p(\sigma_\varepsilon) = \mathcal{U}([0, \max_{h_i, h_j \in \mathcal{H}} |h_i - h_j|]) \quad (3)$$

$$p_{\text{cost}} \sim \text{Uniform}([0, 1]). \quad (4)$$

Posterior probability matching

Posterior probability matching (m_{PPM}) assumes that people approximate Bayes-optimal estimation by drawing one sample from the posterior distribution $P(X|y)$:

$$\hat{X}_y \sim P(X|y). \quad (5)$$

The error model assumes that with probability p_{cost} people guess at random on given trial:

$$P(R = x) = (1 - p_{\text{cost}}) \cdot P(X = x|y) + p_{\text{cost}} \cdot \frac{1}{|\mathcal{H}|} \quad (6)$$

This model has only one free parameter: the error probability p_{cost} . The prior on this parameter is the standard uniform distribution:

$$p_{\text{cost}} \sim U([0, 1]). \quad (7)$$

Anchoring-and-Adjustment with a simple stopping rule

The anchoring-and-adjustment model with a simple stopping rule (m_{AAs}) starts from an anchor a and adjusts the estimate until its plausibility (i.e. posterior probability) reaches a threshold ψ . We model adjustment as a Markov chain that converges the posterior distribution $P(X|y)$. Consequently, the estimate $\hat{X}_n$ becomes a random variable whose distribution changes $Q(\hat{X}_n)$ depends on the number of adjustments n . The initial distribution assigns all of its probability mass to the anchor a : $Q_0(x) = \delta(x - a)$. The probability $P(\hat{X}_n = h_l | \hat{X}_{n-1} = h_k)$ of adjusting estimate $\hat{X}_n = h_k$ to estimate $\hat{X}_{n+1} = h_l$ is defined as the probability that this adjustment is proposed ($P(X_n^{\text{prop}} | \hat{X}_{n-1})$) times the probability that it will be accepted according to the Metropolis-Hastings algorithm (Hastings, 1970):

$$P(\hat{X}_n = h_l | \hat{X}_{n-1} = h_k) = P(X_n^{\text{prop}} = h_l | \hat{X}_{n-1} = h_k) \cdot \min \left\{ 1, \frac{p(X = h_l | y)}{p(X = h_k | y)} \right\} \quad (8)$$

$$P(X_n^{\text{prop}} = h_k | \hat{X}_{n-1} = h_l) \propto \text{Poisson}(|k - l|; \mu_{\text{prop}}), \quad (9)$$

where μ_{prop} is the expected step-size of a proposed adjustment. If the current estimate's plausibility is above the threshold ψ then adjustment terminates. The set of states in which adjustment would terminate is

$$\mathcal{S} = \{h \in \mathcal{H} : P(X = h|y) > \psi\}. \quad (10)$$

If the current estimate is not in this set, then adjustment continues. Consequently, the number of adjustments is a random variable and we have to sum over its realizations to compute the distribution of the estimate $\hat{X}$:

$$Q_{AAs}(\hat{X} = h) = \sum_n Q_{AAs}(\hat{X}_n \in \mathcal{S} \wedge \forall m < n : \hat{X}_m \notin \mathcal{S}) \cdot Q_{AAs}(\hat{X}_n = h | \hat{X}_n \in \mathcal{S}) \quad (11)$$

$$Q_{AAs}(\hat{X}_n = x) = \sum_{k=1}^{|\mathcal{H}|} Q_{AAs}(\hat{X}_{n-1} = h_k | \hat{X}_{n-1} \notin \mathcal{S}) \cdot P(\hat{X}_n = x | \hat{X}_{n-1} = h_k). \quad (12)$$

As in the posterior probability matching model the response distribution combines takes into account that people guess randomly on some of the trials:

$$P(R = x) = (1 - p_{\text{cost}}) \cdot Q_{AAs}(\hat{X} = x) + p_{\text{cost}} \cdot \frac{1}{|\mathcal{H}|} \quad (13)$$

The prior distributions on the models' free parameters are given below:

$$p(\psi) = \exp(-\psi) \quad (14)$$

$$p(\mu_{\text{prop}}) = \mathcal{U}([\min_{h_i, h_j \in \mathcal{H}} |i - j|, \max_{h_i, h_j \in \mathcal{H}} |i - j|]) \quad (15)$$

$$p_{\text{cost}} \sim \text{Uniform}([0, 1]) \quad (16)$$

Anchoring-and-adjustment with a fixed number of adjustments

The anchoring-and-adjustment model with a fixed number of adjustments (m_{AA}) differs from the previous model in that adjustment stops after a fixed, but unknown, number of adjustments (N) regardless of the plausibility of the current estimate:

$$Q_{\text{AA}}(\hat{X}) = Q_{\text{AA}}(\hat{X}_n) \quad (17)$$

$$Q_{\text{AA}}(\hat{X}_0 = x) = \delta(x - a) \quad (18)$$

$$Q_{\text{AA}}(\hat{X}_n = h_l | \hat{X}_{n-1} = h_k) = P(X_n^{\text{prop}} = h_l | \hat{X}_{n-1} = h_k) \cdot \min \left\{ 1, \frac{P(X = h_l | y)}{P(X = h_k | y)} \right\} \quad (19)$$

$$P(X_n^{\text{prop}} = h_l | \hat{X}_{n-1} = h_k) \propto \text{Poisson}(|l - k|; \mu_{\text{prop}}) \quad (20)$$

The error model is the same as before:

$$P(R = x) = (1 - p_{\text{cost}}) \cdot Q_{\text{AA}}(\hat{X} = x) + p_{\text{cost}} \cdot \frac{1}{|\mathcal{H}|} \quad (21)$$

The prior distributions on the model parameters are given below:

$$P(N) = \mathcal{U}(\{0, 100\}) \quad (22)$$

$$p(\mu_{\text{prop}}) = \mathcal{U}([\min_{h_i, h_j \in \mathcal{H}} |i - j|, \max_{h_i, h_j \in \mathcal{H}} |i - j|]) \quad (23)$$

$$p_{\text{cost}} \sim \text{Uniform}([0, 1]) \quad (24)$$

Figure 1. The figure illustrates the resource-rational anchoring-and-adjustment. The three jagged lines are examples of the stochastic sequences of estimates the adjustment process might generate starting from a low, medium, and high anchor respectively. In each iteration a potential adjustment is sampled from a proposal distribution p_{prop} illustrated by the bell curves. Each proposed adjustment is stochastically accepted or rejected such that over time the relative frequency with which different estimates are considered $q(\hat{x}_t)$ becomes the target distribution $p(x|k)$. The top of the figure compares the empirical distribution of the samples collected over the second half of the adjustments with the target distribution $p(x|k)$. Importantly, this distribution is the same for each of the three sequences. In fact, it is independent of the anchor, because the influence of the anchor vanishes as the number of adjustments increases. Yet, when the number of adjustments (iterations) is low (e.g., 25), the estimates are still biased towards their initial values. The optimal number of iterations i^* is very low as illustrated by the dotted line. Consequently, the resulting estimates indicated by the red, yellow, and red cross are still biased towards their respective anchors.

Adaptive Anchoring-and-Adjustment

According to the adaptive anchoring-and-adjustment model (m_{aAA}), the mind adapts the expected step-size of its adjustments μ_{prop} and the number of adjustments n . Concretely, the model chooses the optimal combination $(n^*, \mu_{\text{prop}}^*)$ of adjustments and the step-size such as to minimize the expected sum of time cost and error cost given the relative time cost per adjustment γ and the posterior variance σ :

$$Q_{\text{aAA}}(\hat{X} = x) = Q_{\text{aAA}}(\hat{X}_{n^*} = x) \quad (25)$$

$$(n^*, \mu_{\text{prop}}^*) = \arg \min_{n, \mu_{\text{prop}}} \mathbb{E}_{P(\mu), P(\sigma)} \left[\mathbb{E}_{\mathcal{N}(\tilde{X}; \mu, \sigma)} \left[\mathbb{E}_{\tilde{Q}(\hat{X}_n; \mu, \sigma)} \left[\text{cost}(\tilde{X}, \hat{X}) \right] \right] \right] + \gamma \cdot n, \quad (26)$$

where $\tilde{Q}(\hat{X}_n|\hat{X}_{i-1})$ is the probability to transition from one estimate to the next, if the posterior distribution is a normal distribution with mean μ and standard deviation σ :

$$\tilde{Q}(\hat{X}_n = h_l|\hat{X}_{i-1} = h_k; \mu, \sigma) = P(X_n^{\text{prop}} = h_l|\hat{X}_{i-1} = h_k) \cdot \min \left\{ 1, \frac{\mathcal{N}(h_l; \mu, \sigma)}{\mathcal{N}(h_k; \mu, \sigma)} \right\} \quad (27)$$

$$P(\mu) = P(X), P(\sigma) = \mathcal{U} \left(\sigma; \min_y \sqrt{\text{Var}(X|y)}, \max_y \sqrt{\text{Var}(X|y)} \right). \quad (28)$$

The relative iteration cost γ is determined by the time cost c_t , the error cost c_e , and the time $\tau_{\text{adjustment}}$ it takes to perform one adjustment

$$\gamma = \frac{\tau_{\text{adjustment}} \cdot c_t}{c_e}. \quad (29)$$

Note that the choice of the number of iterations and the step-size of the proposal distribution is not informed by the distance from the anchor to the posterior mean since this would presume that the answer was already known. Instead, the model minimizes the expected value of the cost under the assumption that the posterior mean will be drawn from the prior distribution. The model also does not presume the shape of the posterior distribution was known a priori; instead it makes a Gaussian approximation with matching mean and variance. Given the number of adjustment and the step-size of the proposal distribution, the adjustment process and response generation work as in the previous model:

$$P(R = x|y) = (1 - p_{\text{cost}}) \cdot Q_{\text{aAA}}(\hat{X}_{n^*} = x) + p_{\text{cost}} \cdot \frac{1}{|\mathcal{H}|} \quad (30)$$

$$Q_{\text{aAA}}(\hat{X}_0 = x) = \delta(x - a) \quad (31)$$

$$Q_{\text{aAA}}(\hat{X}_n = h_l|\hat{X}_{i-1} = h_k) = P(X_n^{\text{prop}} = h_l|\hat{X}_{i-1} = h_k) \cdot \min \left\{ 1, \frac{P(X = h_l|y)}{P(X = h_k|y)} \right\} \quad (32)$$

$$P(X_n^{\text{prop}} = h_l|\hat{X}_{i-1} = h_k) \propto \text{Poisson}(|l - k|; \mu_{\text{prop}}^*) \quad (33)$$

The prior distributions on the model's parameters are given below:

$$p(\tau_{\text{adjustment}}) = \text{Exp}(\tau_{\text{adjustment}}; \mu = 50\text{ms}) \quad (34)$$

$$p(\sigma_\varepsilon) = \mathcal{U}([0, \max_{h_i, h_j \in \mathcal{H}} |h_i - h_j|]) \quad (35)$$

$$p_{\text{cost}} \sim \text{Uniform}([0, 1]) \quad (36)$$

Adaptive Anchoring-and-Adjustment with intrinsic error cost

The adaptive anchoring-and-adjustment model with intrinsic error cost (m_{aAAi}) extends the adaptive model m_{aAA} by one parameter: a constant $c_{\text{intrinsic}}$ that is added to the error cost:

$$\gamma = \frac{\tau_{\text{adjustment}} \cdot c_t}{c_e + c_{\text{intrinsic}}} \quad (37)$$

The prior over $c_{\text{intrinsic}}$ was

$$p(c_{\text{intrinsic}}) = \text{Uniform}([0, 100]) \quad (38)$$

Random Choice

According to the random choice model, people’s responses are independent of the task and uniformly distributed over the range of all possible responses:

$$R \sim \text{Uniform}(\mathcal{H}) \quad (39)$$

Which model best explains adjustments from self-generated anchors?

To determine which model best explains adjustment from self-generated anchors we perform Bayesian model selection on the data from Experiment 1 by Lieder et al. (submitted). In this experiment, participants estimated bus departure times based on prior knowledge about the distribution of departure times and the observation that the bus had not departed by a given time. Each participant completed this task under the four incentive conditions of a 2×2 factorial design with the independent variables time cost (high vs. none) and error cost (high versus none). In this experiment, people appeared to anchor their estimate of the bus’s delay on about 7.5 minutes and their adjustments increased with error cost and decreased with time cost (Lieder et al., submitted).

To formally test the four theories—anchoring and adjustment, posterior probability matching, Bayesian decision theory, and random choice—and which of the seven models that instantiate them against each other, we performed random-effects Bayesian model selection at the group level (Stephan et al., 2009) and family-level Bayesian model selection (Penny et al., 2010) as implemented in SPM8. For each model we separately approximated the log-probability of each participant’s predictions using the Laplace approximation (Tierney & Kadane, 1986) when applicable, that is when the likelihood function is differentiable with respect to the parameters, and numerical integration of the joint density otherwise. Numerical integration was necessary for discrete-valued parameters such as the number of adjustments. Numerical integration was also necessary for continuous parameters that affect the resource-rational number of adjustments. This is because the likelihood function changes abruptly by a non-differential step when the resource-rational number of adjustments jumps from one number to another. Numerical integration with respect to continuous parameters was performed using the functions *integral* and *integral2* available in Matlab 2013b.

According to Bayesian model selection, adaptive anchoring-and-adjustment with intrinsic error cost (m_{aAAi}) explained the data better than any of the alternative models: we can be 99.99% confident that the adaptive anchoring-and-adjustment with intrinsic error is the best model for a larger percentage of people (64.4%) than any of the alternative models; see Figure 2, top panel. In addition to this random-effects analysis we also performed a Bayesian fixed effects analysis by computing the group Bayes factor for each pair of models. Reassuringly, this analysis led to the same conclusion: according to the posterior odds ratios, the adaptive anchoring-and-adjustment with intrinsic error cost was at least $\exp(220)$ times as likely as any of the other models we considered. Next, we applied family level inference to determine which theory best explains our data; see Figure 2, bottom left panel. According to this method, we can be 99.99% confident that anchoring-and-adjustment is the most probable explanation for a significantly larger proportion of participants (78.2%) than either posterior probability matching (11.0%), Bayesian decision theory (7.2%), or random choice (3.6%). Finally, we compared adaptive to non-adaptive models; see Figure 2, bottom right

panel. According to the result, we can be 99.86% confident that for the majority of people (79.2%) our adaptive models' predictions are more accurate than the predictions of their non-adaptive counterparts.

To validate that people perform more adjustments when errors are costly and fewer adjustments when time is costly, as assumed by the adaptive resource-rational model, we computed the maximum-a-posteriori estimates of the parameters of the second anchoring-and-adjustment model (m_{AA}) separately for each of the four incentive conditions. Figure 3 shows the estimated number of adjustments as a function of the incentives for speed and accuracy. For five of the six pairs of conditions, we can be more than 96.9% confident that the number adjustments differ in the indicated direction, and for the sixth pair we can be more than 92% confident that this is the case. Therefore, this analysis supports the conclusion that our participants adapted the number of adjustments to the cost of time and error. To determine whether this pattern is consistent with choosing the number of adjustments adaptively we fit the parameters determining the rational number of adjustments to these estimates. We found that rational resource allocation predicts a qualitatively similar pattern of adjustments for reasonable parameter values (convergence rate: 0.71, time per adjustment: 27ms, assumed initial bias: 6.25min).

Which model best explains adjustments from provided anchors?

To determine which model best explains adjustment from provided anchors, we performed Bayesian model selection on the data from Experiment 2 by Lieder et al. (submitted). In this experiment, participants were first asked to compare a bus's expected departure time to a low versus high anchor and then estimated when the bus would depart based on prior knowledge about the distribution of departure times and the observation that the bus had not departed by a given time. Again, each participant completed this task under four different incentive conditions (Time Cost vs. No Time Cost $\times$ Error Cost vs. No Error Cost). Participants' judgments were biased in the direction of the provided anchors, and the magnitude of this bias increased with time cost but decreased with error cost (Lieder et al., submitted).

Consistent with the biases and the effects of time cost and error cost reported by Lieder et al. (submitted), we found that the two adaptive anchoring-and-adjustment models explained their participants' predictions significantly better than any of the alternative models; see Figure 4, top panel. Concretely, the first adaptive anchoring-and-adjustment model (m_{aAA}) was the best explanation for 36.9% of the participants, and the adaptive anchoring-and-adjustment model with an additional intrinsic error cost parameter (m_{aAAi}) was the best explanation for another 24.8% of them. Thus for the majority of participants, responses were best explained by adaptive anchoring-and-adjustment. Furthermore, we can be 85.9% confident that the first adaptive anchoring-and-adjustment is the best model for a larger percentage of people than any of the alternative models. In addition to this random effects analysis, we also ran a Bayesian fixed-effects analysis by computing the group Bayes factors. This analysis confirmed that the two adaptive anchoring-and-adjustment models explain the data substantially better than any of the alternatives, but among these two models it strongly favored the more complex model with intrinsic error cost: According to the posterior-odds ratios this model is at least 10^{30} times as likely as any other model we considered. In conclusion, we found that most participants performed adaptive anchoring-

Figure 2. Results of Bayesian model selection given the data from Experiment 1 by Lieder, Griffiths, Huys, and Goodman (submitted). The top panel shows the posterior probabilities of individual models. The bottom left panel shows the posterior probabilities of the four theories (BDT: Bayesian decision theory, PPM: posterior probability matching, AA: anchoring-and-adjustment, random: predictions are chosen randomly). The bottom right panel shows the posterior probabilities of adaptive versus non-adaptive models.

and-adjustment (m_{aAA} and m_{aAAi}) and while the contribution of the intrinsic error cost is

Figure 3. Estimated (left panel) and predicted (right panel) number of adjustments.

negligible in many participants it is crucial in others. Next, we asked which theory best explains people’s anchoring biases; see Figure 4, bottom left panel. According to family-level Bayesian model selection, we can be 99.99% confident that anchoring-and-adjustment is the most probable explanation for a significantly larger proportion of people (76.9%) than either posterior probability matching (10.6%), Bayesian decision theory (10.4%), or random choice (2.1%). Furthermore, we can be 98.5% confident that for the majority of people (67.6%) our adaptive models’ predictions are more accurate than the predictions of their non-adaptive counterparts; see Figure 4, bottom right panel.

Conclusion

Our model-based analysis of the experiments reported by Lieder et al. (submitted) supported anchoring-and-adjustment over alternative models of numerical estimation. Furthermore, it strongly supported the resource-rational anchoring and adjustment model over non-adaptive anchoring-and-adjustment models with a fixed number of adjustments or a fixed threshold. These findings were robust to whether the anchor was provided or self-generated. These results replicate earlier findings based on a different experimental paradigm Lieder, Goodman, and Griffiths (2013).

These findings support the view that people rationally adapt their number of adjustments to the cost of time and error. Therefore, the resulting anchoring bias might be a window on resource-rational information processing rather than a sign of human irrationality.

Figure 4. Model selection results for Experiment 2. The top panel shows the posterior model probabilities. The bottom panel shows the results of Bayesian inference on the level of model families.

References

- Epley, N., & Gilovich, T. (2006). The anchoring-and-adjustment heuristic. *Psychological Science*, 17(4), 311–318.
- Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science*, 7(2), 217–229.
- Hastings, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1), 97–109.
- Lieder, F., Goodman, N. D., & Griffiths, T. L. (2013). *Reverse-engineering resource-efficient algorithms* [Paper presented at NIPS-2013 Workshop Resource-Efficient ML, Lake Tahoe, USA].
- Lieder, F., Griffiths, T. L., & Goodman, N. D. (2012). Burn-in, bias, and the rationality of anchoring. In P. Bartlett, F. C. N. Pereira, L. Bottou, C. J. C. Burges, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems 26*. Red Hook: Curran Associates, Inc.
- Lieder, F., Griffiths, T. L., Huys, Q. J. M., & Goodman, N. D. (n.d.-a). The anchoring bias reflects rational use of cognitive resources.

- Lieder, F., Griffiths, T. L., Huys, Q. J. M., & Goodman, N. D. (submitted). Empirical evidence for resource-rational anchoring-and-adjustment.
- Mussweiler, T., & Strack, F. (1999). Hypothesis-consistent testing and semantic priming in the anchoring paradigm: A selective accessibility model. *Journal of Experimental Social Psychology*, 35 (2), 136–164.
- Penny, W. D., Stephan, K. E., Daunizeau, J., Rosa, M. J., Friston, K. J., Schofield, T. M., & Leff, A. P. (2010). Comparing families of dynamic causal models. *PLoS Computational Biology*, 6(3), e1000709.
- Simmons, J. P., LeBoeuf, R. A., & Nelson, L. D. (2010). The effect of accuracy motivation on anchoring and adjustment: do people adjust from provided anchors? *Journal of Personality and Social Psychology*, 99(6), 917–932.
- Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *Neuroimage*, 46(4), 1004–1017.
- Strack, F., & Mussweiler, T. (1997). Explaining the enigmatic anchoring effect: Mechanisms of selective accessibility. *Journal of Personality and Social Psychology*, 73(3), 437.
- Tierney, L., & Kadane, J. B. (1986). Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association*, 81(393), 82–86.
- Turner, B. M., & Schley, D. R. (2016). The anchor integration model: A descriptive model of anchoring effects. *Cognitive Psychology*, 90, 1–47.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131.
- Vul, E., Goodman, N. D., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? Optimal decisions from very few samples. *Cognitive Science*, 38, 599–637.