

BRIEF REPORT

Empirical evidence for resource-rational anchoring and adjustment

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

© Psychonomic Society, Inc. 2017

Abstract People's estimates of numerical quantities are systematically biased towards their initial guess. This anchoring bias is usually interpreted as sign of human irrationality, but it has recently been suggested that the anchoring bias instead results from people's rational use of their finite time and limited cognitive resources. If this were true, then adjustment should decrease with the relative cost of time. To test this hypothesis, we designed a new numerical estimation paradigm that controls people's knowledge and varies the cost of time and error independently while allowing people to invest as much or as little time and effort into refining their estimate as they wish. Two experiments confirmed the prediction that adjustment decreases with time cost but increases with error cost regardless of whether the anchor was self-generated or provided. These results support the hypothesis that people rationally adapt their number

of adjustments to achieve a near-optimal speed-accuracy tradeoff. This suggests that the anchoring bias might be a signature of the rational use of finite time and limited cognitive resources rather than a sign of human irrationality.

Keywords Bounded rationality · Heuristics · Cognitive biases · Probabilistic reasoning · Anchoring-and-adjustment

Tversky and Kahneman's demonstrations that human judgment systematically violates the normative principles of probability theory fundamentally challenged the belief that people are rational (Tversky & Kahneman, 1974). In one of these demonstrations, Tversky and Kahneman (1974) spun a wheel of fortune in front of their participants. Participants were asked to judge whether the number of African countries in the United Nations was larger or smaller than the number the wheel had stopped on. Afterwards, participants estimated exactly how many African countries there were in the United Nations. Curiously, participants' estimates of this quantity were significantly biased towards the irrelevant number they had just compared it to. According to Tversky and Kahneman (1974), this violation occurs because people use a two-stage process called *anchoring-and-adjustment*. In the first stage, people generate an initial guess called their *anchor*. In the second stage, they adjust their estimate away from this anchor to incorporate additional information, but the adjustment is usually insufficient. In the experiment described above, people appear to have anchored on the random number provided by the experimenter and adjusted their estimate insufficiently. Consequently, when the anchor was low, people's judgments were too low; and when the anchor was high, then their judgments were too high.

✉ Falk Lieder
falk.lieder@berkeley.edu

¹ Helen Wills Neuroscience Institute, University of California, Berkeley USA

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

³ Department of Psychology, Stanford University, Stanford, USA

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

⁵ Department of Psychology, University of California, Berkeley, USA

At first sight, anchoring appears to be irrational, because it deviates from the standards of logic and probability which are typically used to assess rationality. But it could also arise from an optimal tradeoff between the cost of error in judgment and the cost of the time it takes to reduce this error. This hypothesis has been formalized by a computational model that chooses the number of adjustments that minimizes the expected sum of time cost and error cost (Lieder, Griffiths, & Goodman, 2012; Lieder, Griffiths, Huys, & Goodman, 2017a, under review). This model predicts that adjustment should decrease with time cost but increase with error cost regardless of whether the anchor was provided or self-generated. Here, we experimentally test these predictions by varying both the cost of error and the cost of time without imposing a deadline. This allows us to evaluate how much time and effort people choose to invest into adjusting their estimate away from the anchor against the normative prescriptions of resource-rational anchoring and adjustment (Lieder, et al., 2012, 2017a, under review).

After presenting the predictions of the resource-rational anchoring and adjustment model, we first report an experiment in which we evaluated the effect of time cost and error cost on adjustment from self-generated anchors, and then investigate the same effects in an experiment with provided anchors. We close by discussing the implications of our findings for theories of anchoring and the debate about human rationality (Stanovich, 2009).

Empirical predictions of resource-rational anchoring and adjustment

Resource-rational anchoring and adjustment postulates that the number of adjustments people perform achieves a near-optimal tradeoff between the cost of error and the cost of time (Lieder et al., 2012, 2013, 2017a, under review). In brief, optimal resource allocation implies that relative adjustment decreases with the relative cost of time. Therefore, the slope of the anchoring bias as a function of the distance from the anchor to the correct value should be highest when time cost is high and error cost is low. Conversely, the slope of the anchoring bias should be the shallowest when error cost is high and time cost is low. Lastly, when time cost and error cost are both high or both low, then the slope should be intermediate. Figure 1 illustrates these predictions.

Furthermore, resource-rational anchoring and adjustment also predicts that (an upper bound on) the anchoring bias increases linearly with the distance of the anchor from the correct value. This prediction leads to a linear regression model that allows us to estimate people's anchor




Fig. 1 According to resource-rational anchoring-and-adjustment the negative anchoring bias increases linearly with the distance from the anchor to the true value. Importantly, this model predicts that the slope of this function increase with time cost (TC) and decreases with error cost (EC)

a , their relative adjustments ($\frac{\mathbb{E}[\hat{X}|x] - a}{x - a}$), and the resulting anchoring bias $\text{Bias}_t(x, a)$ by regressing their estimates \hat{X} on the correct value x :

$$\hat{X} = \alpha + \beta \cdot x + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2) \quad (1)$$

$$\frac{\mathbb{E}[\hat{X}|x] - a}{x - a} = \beta, \quad a = \frac{\alpha}{1 - \beta} \quad (2)$$

$$\text{Bias}_t(x, a) = \alpha - (1 - \beta) \cdot x. \quad (3)$$

Contrary to Epley and Gilovich (2006), resource-rational anchoring and adjustment assumes that people adjust not only from self-generated anchors but also from provided anchors. If this assumption is correct, then error cost should increase adjustment and time cost should decrease adjustment regardless of whether the anchor was self-generated (Experiment 1) or provided (Experiment 2).

Experiment 1: Self-generated anchors

In most previous anchoring experiments the biases in people's judgments resulted not only from anchoring but also from the discrepancy between the truth and what people actually know (Lieder et al., 2017a). To avoid this confound, we designed a prediction task that controls people's knowledge about the quantity to be estimated. To test if people adapt their number of adjustments rationally, we manipulated both the cost of time and the cost of error within subjects. In this first experiment, no anchor was provided to the participant. Instead, we estimated our participants' self-generated anchor from how the bias of their responses changes with the correct value (1–2).

Method

Participants We recruited 30 participants (14 male, 15 female, 1 unreported) on Amazon Mechanical Turk. Our participants were between 19 and 65 years old, and their level of education ranged from high school to graduate degrees. Participants were paid \$1.05 for their time and could earn a performance-dependent bonus of up to \$0.80. Six participants were excluded because they incorrectly answered questions designed to test their understanding of the task (see Procedure).

Materials The experiment was presented as a website programmed in HTML and JavaScript. Participants predicted when a person would get on a bus given when he had arrived at the bus stop based on the bus's timetable and examples of previous departure times. Figure 2 shows a screenshot from one of the trials. The timeline at the top of the screen was used to present the relevant information and record predictions. At the beginning of each trial the bus's timetable (orange bars) and the person's arrival at the bus stop (blue bars) were highlighted on the timeline. Participants indicated their prediction by clicking on the corresponding point on the timeline. When participants were incentivized to respond quickly, a falling red bar indicated the passage of time and its cost, and the costs of time and error were conveyed in the bottom left corner; see Fig. 2. Feedback was provided by a pop-up window informing participants about how many points they had earned and a green bar highlighting the actual departure time on the number line.

Procedure After completing the consent form, each person participated in four scenarios corresponding to the four conditions of a 2×2 within-subject design with the independent variables time cost (0 vs. 30 points/sec) and error cost (0 vs. 10 points/unit error). The order of the four conditions was randomized between subjects. At the end of the experiment participants received a bonus payment proportional to the number of points they had earned in the experiment. The conversion rate was 1 cent per 100 points, and participants could earn up to 100 points per trial.

Each scenario comprised a cover story, instructions, 10 examples, 5 practice trials, 5 attention check questions, 20 prediction trials, 3 test questions, and one demographic question. Each cover story was about a person repeatedly taking the same bus route in the morning, for example "Jacob commutes to work with bus #22. On average, the first bus departs at 8:01 AM, and the second bus departs at 8:21 AM but departure times vary. On some days Jacob misses the first bus and takes the second bus." In each scenario both the person and the bus route were different. The task instructions informed participants about the cost of

time and error and encouraged them to attentively study the examples and practice trials so that they would learn to make accurate predictions. After the cover story, participants were shown when the bus had arrived on the ten workdays of the two preceding weeks (10 examples); see Fig. 3.

Next, participants made 5 practice predictions with feedback. The ensuing attention check questions verified the participants' understanding of the time line and the costs of time and error. Participants were allowed to go back and look up this information if necessary. Participants who made at least one error were required to retake this test until they got all questions correct. Once they had answered all questions correctly, participants proceeded to 20 prediction trials with feedback. In both the practice trials and the prediction trials the feedback comprised the correct departure time, the incurred error cost, the incurred time cost, and the resulting number of points for the trial. The times at which the fictitious person arrived at the bus stop were chosen such that the probability that he had missed the first bus approximately covered the full range from 0 to 1 in equal increments. In the 1st, 3rd, . . . , 2nd-last prediction trial the person arrived early and the bus was on time. The purpose of these odd-numbered trials was to set the anchor on the even-numbered trials to a low value. After each scenario's prediction trials we tested participants' understanding of the number line, the cost of time, and the cost of error once again. We excluded six participants, because their answers to these questions revealed that they had misunderstood the number line, the cost of time, or the cost of error in at least one condition. After this they reported one piece of demographic information: age, gender, level of education, and employment status respectively. On the last page of each block, participants were informed about the bonus they had earned in the scenario.

To pose a different prediction problem on every trial of each block despite the limited number of meaningfully different arrival times, we varied the distribution of the bus's delays between blocks. There were four delay distributions in total. All of them were Pearson distributions that differed only in their variance. Their mean, skewness, and kurtosis were based on the bus lateness statistics from Great Britain.¹ The order of the delay distributions was randomized between participants independently of the incentives. The 10 examples of bus departure times were chosen such that their mean, variance, and skewness reflected the block's delay distribution as accurately as possible. For each trial, a "correct" departure time x was sampled from the conditional distribution of departure times given that the fictitious person departs after his arrival at the bus stop. The

¹Bus Punctuality Statistics GB 2007 report; <http://estebanmoro.org/2009/01/waiting-for-the-bus/>

Day 11

Today Will arrived at the bus stop at **7:26:40** AM. I predict that he will board a bus at





Fig. 2 Screenshot of a prediction trial from Experiment 1 with time cost and error cost. The number line on the top conveys the bus schedule and when the person arrived at the bus stop. The cost of error and time are shown in the bottom left corner, and the red bar in the bottom right corner shows the passage of time and the cost associated with it. The complete experiment can be inspected online at <http://cocosci.berkeley.edu/mturk/falk/PredictionExperiment1/experiment.html>

condition's cost of time c_t and cost of error c_e determined the number of points a participant would receive according to

$$\text{points} = \max\{0, 100 - c_e \cdot \text{PE} - c_t \cdot \text{RT}\}, \quad (4)$$

$$\text{PE} = |\hat{x} - x|, \quad (5)$$

where PE is the absolute prediction error between the estimate \hat{x} and the true value x , and RT is the response time. The bottom part of Fig. 2 shows how time cost and error cost were conveyed to the participants during the trials. The red bar on the right moved downward and its position indicates how much time has passed and how many points have consequently been lost.

Results and discussion

To assess whether our participants' predictions were systematically biased, we inspected their average prediction for a range of true bus delays. The true bus delays were sampled from a distribution, of which subjects had seen 10 samples. We binned participants' average predictions when the true bus delay was 0.5 ± 2.5 min, 5.5 ± 2.5 min, ..., or 35.5 ± 2.5 min. Participants showed a systematic bias, overestimating the delay when its true value was less than 3 minutes ($t(815) = 16.0$, $p < 10^{-15}$), but underestimating

it when its true value was larger than 7 minutes (all $p \leq 0.0011$; see Fig. 4). Visual inspection suggested that the bias was approximately proportional to the correct value. Applying the linear regression model of anchoring (1–3) confirmed that the linear correlation between correct value and bias was significantly different from zero ($P(\text{slope} \in [-0.6148, -0.5596]) = 0.95$). As shown in Fig. 4, the bias was positive when the delay was greater than 7.5min and negative for greater delays. Our participants thus appeared to anchor around 7.5min and adjust their initial estimate by about 41.3% of the total distance to the true value (95%-CI: [38.52%, 44.04%]).

Since the data showed standard anchoring effects, we can now proceed to testing its novel predictions. Consistent with our theory's prediction, we found that time cost made participants faster and less accurate whereas error cost made them slower and more accurate (Fig. 5). To determine whether the anchoring bias increased with time cost and decreased with error cost we performed a repeated-measures ANOVA of participants' relative adjustments as a function of time cost and error cost. To be precise, we first estimated each participant's relative adjustment separately for each of the four conditions using our linear regression model of anchoring and adjustment (1). We then performed an ANOVA on the estimated relative

Week 1

If you study the following examples closely, you will earn more points by guessing more accurately. Each example is presented as points on a timeline: the **scheduled departure time** (i.e. when the bus is supposed to arrived), and **when the first bus actually arrived**.

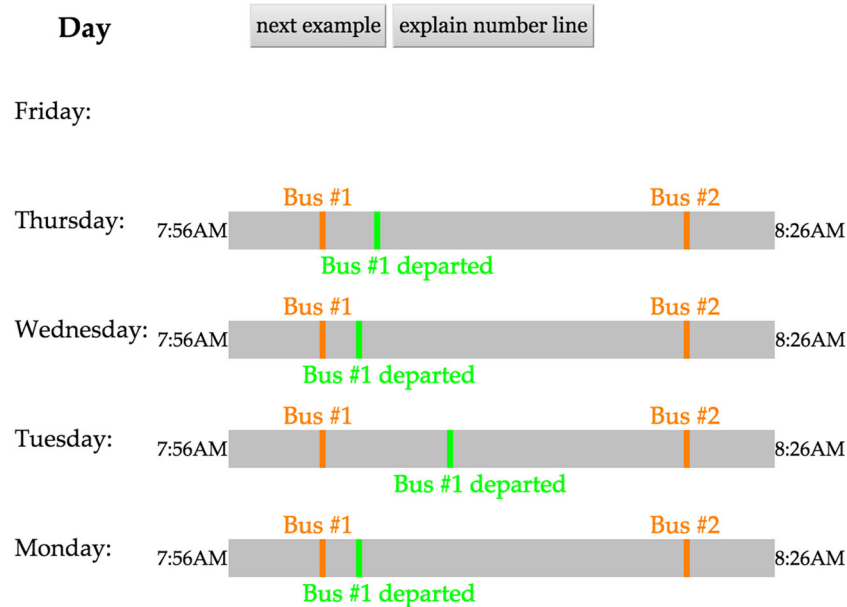


Fig. 3 Screenshot of the first examples screen of Experiment 1

adjustments with the factors time cost and error cost (fixed-effects) as well as participant number (random effect) and the interaction effect of time cost and error cost. We found that time cost significantly reduced relative adjustment from 50.7% to 31.0% ($F(1, 69) = 21.86, p < 0.0001$) whereas error cost significantly increased it from 31.6% to 50.1% ($F(1, 69) = 19.49, p < 0.0001$) and the interaction was non-significant. The mean relative adjustments of each condition are shown in Table 1. Consequently, as predicted by our theory (Fig. 1), the anchoring bias increased more rapidly with the true delay when time cost was high or error cost was low (Fig. 6). This is consistent with the hypothesis that people rationally adapt the number of adjustments to the relative cost of time.²

Quantitative comparisons of our resource-rational model against alternative theories, including the anchoring and adjustment heuristic proposed by Epley and Gilovich (2006), also provided strong evidence for rational adjustment, and we provide a detailed summary of these results in a technical report (Lieder et al., 2017b). While the results presented here demonstrate that people adaptively trade off being biased for being fast, our analysis had to postulate and estimate people's self-generated anchors. Therefore, we

cannot be sure whether people really self-generated anchors and adjusted from them, or whether their responses merely look as if they did so. If people really use anchoring and adjustment in this task, then we should be able to shift the biases shown in Fig. 4 by providing different anchors; we tested this prediction in Experiment 2.




Fig. 4 In Experiment 1 the magnitude of the anchoring bias grew linearly with the correct value. The error bars indicate 95% confidence intervals on the average bias, that is ± 1.96 standard errors of the mean

²Estimating relative adjustment under the assumption that people anchor on their previous estimate led to the same conclusions.




Fig. 5 Mean absolute errors and reaction times as a function of time cost and error cost indicate an adaptive speed-accuracy tradeoff

Experiment 2: Provided Anchors

To test whether the biases observed in Experiment 1 resulted from anchoring and to evaluate whether the effects of time cost and error cost also hold for provided anchors, we ran a second experiment in which anchors were provided by asking participants to compare the to-be-predicted delay to a low versus a high number before every prediction.

Method

The materials, procedures, models, and data analysis tools used in Experiment 2 were identical to those used in Experiment 1 unless stated otherwise.

Participants We recruited 60 participants (31 male, 29 female) on Amazon Mechanical Turk. They were between 18 and 60 years old, and their level of education ranged from high school diploma to PhD. Participants were paid \$1.25 for participation and could earn a bonus of up to \$2.20 for the points they earned in the experiment.

Materials Experiment 2 was presented as a website programmed in HTML and JavaScript. Experiment 2 was mostly identical to Experiment 1. The relevant changes are summarized below. The complete experiment can be inspected online at <http://cocosci.berkeley.edu/mturk/falk/PredictionExperiment2/experiment.html>.

Procedure Experiment 2 proceeded like Experiment 1 except for three changes: First, each prediction was preceded by the question “Do you think he will depart before or after X am?”, where X is the anchor. This question was presented between the sentence reporting the time the person reached the bus stop and the number line. Participants were

Table 1 Relative size of our participants’ adjustments of their initial guesses towards the correct answer by incentive condition with 95% confidence intervals

	No error cost	High error cost
No time cost	43.6 ± 11.2%	57.8 ± 4.8%
High time cost	19.6 ± 9.0%	42.5 ± 9.8%

required to answer this question by selecting “before” or “after”. This is the standard procedure for providing anchors (Jacowitz & Kahneman, 1995; Russo & Schoemaker, 1989; Tversky & Kahneman, 1974). In the two conditions with time cost, participants were given 3 seconds to answer this question before the timer started. Participants were not allowed to make a prediction until they had answered. We incentivized them to take this question serious by awarding +10 points for correct answers and −100 points for incorrect ones. For each participant the anchor was high in half of the trials of each condition and low in the other half. The low anchor was 3 minutes past the scheduled departure of the first bus, and the high anchor was 3 minutes past the scheduled departure of the second bus. The list of anchors was shuffled separately for each block and participant. Second, the 1st, 3rd, 5th, . . . , 2nd-last trial were no longer needed, because they merely served to set the anchor on the even numbered trials of Experiment 1 to a small value. We therefore replaced those trials by 10 trials whose query times tighten the grid of those in the even-numbered trials. Thus for each participant, each block includes ten prediction trials with low anchors and ten prediction trials with high anchors. Third, the conversion of points into bonuses remained linear but was scaled up accordingly. The instructions were updated to reflect the changes.

We excluded one participant due to incomplete data, and 16 participants whose answers to our test questions indicated they had misunderstood the time line used to present information and record predictions, or the cost of time or error in at least one condition.³

Results and Discussion

Our participants answered the anchoring questions correctly in 74.8% of the trials. As in Experiment 1, people’s predictions were systematically biased: Our participants significantly overestimated delays smaller than 8 min (all $p < 10^{-11}$) and significantly underestimated delays larger than 13 min (all $p < 10^{-4}$); see Fig. 7. Furthermore, the biases were shifted upwards when the anchor was

³This exclusion rate would be high in a laboratory experiment, but it is not unusual for long online experiments run on Amazon Mechanical Turk.

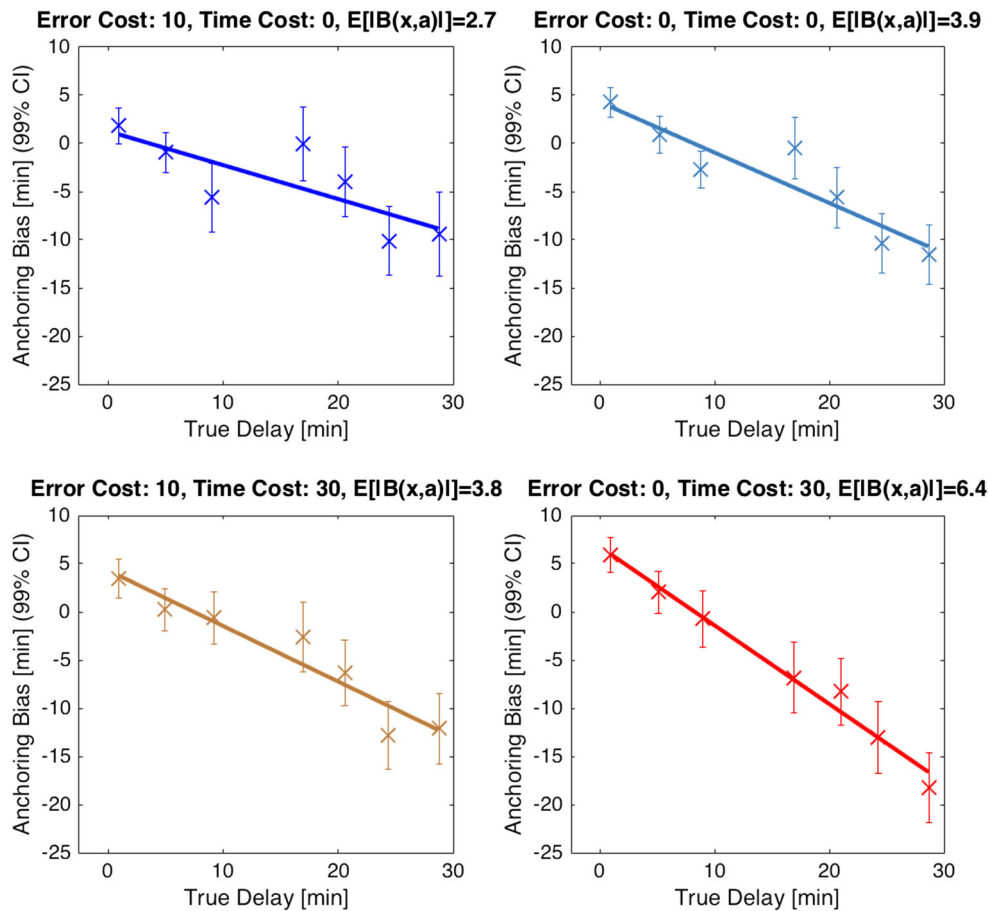


Fig. 6 Anchoring bias in Experiment 1 by time cost and error cost confirms our theoretical prediction; compare Fig. 1

high compared to when the anchor was low ($z = 7.26$, $p < 10^{-12}$; see Fig. 7). This effect was also evident in our participants' average predictions: when the anchor was




Fig. 7 Biases when the provided anchor was high versus low. Solid lines show the results of linear regression. Shaded areas are 95% confidence bands, the diamonds with error bars are the average biases within a five minute window and their 95% confidence intervals; that is ± 1.96 SEM

high, then participants predicted significantly later departures than when the anchor was low: 12.06 ± 0.29 min versus 10.03 ± 0.15 min ($t(3438) = 6.16$, $p < 10^{-15}$). To estimate our participants' anchors and quantify their adjustments, we applied the linear regression model described above (1). Overall, the estimated anchor was significantly higher in the high anchor condition (12.69 min) than in the low anchor condition (9.74 min, $p < 10^{-15}$). Adjustments tended to be small: on average, participants adjusted their estimate only 29.86% of the distance from the anchor to the correct value when the anchor was low (95% CI: [26.38%; 30.85%]) and 27.25% of this distance when the anchor was high (95% CI: [24.00%; 30.50%]). Thus, the relative adjustments were significantly smaller than in Experiment 1 (95% CI: [38.52%, 44.04%]) and they did not differ between the high and low anchor condition ($z = 1.16$; $p = 0.12$). Thus the linear relationship between the bias and the true delay and difference between the biases for the high versus the low anchor (Fig. 7) may result from insufficient adjustment away from different anchors. This also explains why the average predictions were higher in the high anchor condition than in the low anchor condition.

Table 2 ANOVA of relative adjustment as a function of time cost and error cost

Source	d.f.	Sum Sq.	Mean Sq.	F	p
error cost	1	1.0318	1.03178	15.5	0.0001
time cost	1	0.6912	0.69115	10.39	0.0014
subject	42	11.3544	0.27034	4.06	10^{-12}
anchor (high vs. low)	1	0.0774	0.07739	1.16	0.2817
error cost \times time cost	1	0.1066	0.10659	1.6	0.2066
Error	297	19.7643	0.06655		
Total	343	33.0256			

Next, we assessed whether the amount by which participants adjusted their initial estimate increased with error cost and decreased with time cost. To answer this question we performed a repeated-measures ANOVA of relative adjustment as a function of time cost and error cost. To be precise, we first estimated each participant's relative adjustment separately for each of the four conditions and the two anchors using our linear regression model of anchoring and adjustment (1). We then performed an ANOVA on the estimated relative adjustments with the factors time cost, error cost, and high vs. low anchor (fixed-effects) as well as participant number (random effect) and the interaction effect of time cost and error cost; see Table 2. We found that time cost significantly reduced relative adjustment from 37.2% to 28.2% ($F(1, 297) = 15.5$, $p = 0.0001$) whereas error cost significantly increased it from 31.2% to 34.2% ($F(1, 297) = 10.39$, $p = 0.0014$), and the interaction was non-significant. The mean relative adjustments of each condition are shown in Table 3. Figure 8 shows the effects of incentives for speed and accuracy on the anchoring bias in the provided anchors experiment; note that the slope of each line is 1 minus the relative size of the adjustments as estimated with the linear regression model of anchoring (1–3). As predicted by our theory (cf. Fig. 1) and observed for self-generated anchors (cf. Fig. 6), the slope of the anchoring bias was largest when time cost was high and errors were not penalized. Table 3 summarizes the relative adjustments sizes in the four incentive conditions. Figure 8 suggests that the effects of time cost and error cost were weaker in the high anchor condition than in the low anchor condition.

As for Experiment 1, quantitative model comparisons confirmed that Experiment 2 provides strong evidence for rational adjustment (Lieder et al., 2017b).

In summary, our participants' predictions were significantly biased towards the provided anchors. This bias increased linearly with the distance from the anchor to the correct value. Critically, as predicted by resource-rational anchoring and adjustment, the magnitude of this effect decreased with error cost and increased with time cost (compare Figs. 1 and 8). Thus the bias towards the provided anchors and the effects of time cost and error cost were qualitatively the same as for self-generated anchors (Fig. 6). Hence, consistent with Simmons et al. (2010) but contrary to Epley and Gilovich (2006), our results suggest that incentives increase adjustment from both provided and self-generated anchors. Incentives can thus be more effective at reducing the anchoring bias than initially assumed (Tversky & Kahneman, 1974; Epley & Gilovich, 2005), and anchoring and adjustment may be sufficient to explain the effects of both provided and self-generated anchors. For a detailed discussion of how the resource-rational anchoring-and-adjustment model evaluated in this paper goes beyond previous proposals, please see Lieder et al. (2017a).

Despite the qualitative commonalities between the results for self-generated versus provided anchors, there were quantitative differences: In three of the four conditions, the adjustments were significantly smaller for provided anchors than for self-generated anchors. One reason could be that in Experiment 1 the anchoring biases towards high versus low self-generated anchors cancelled each other out. A second reason could be that people treat a provided anchors as

Table 3 Relative size of our participants' adjustments of their initial guess towards the correct answer by incentive condition in the experiment with provided anchors with 95% confidence intervals

	No Error Cost	High error cost
No time cost	$30.0 \pm 7.4\%$	$44.4 \pm 8.4\%$
High time cost	$24.5 \pm 6.5\%$	$32.0 \pm 9.1\%$

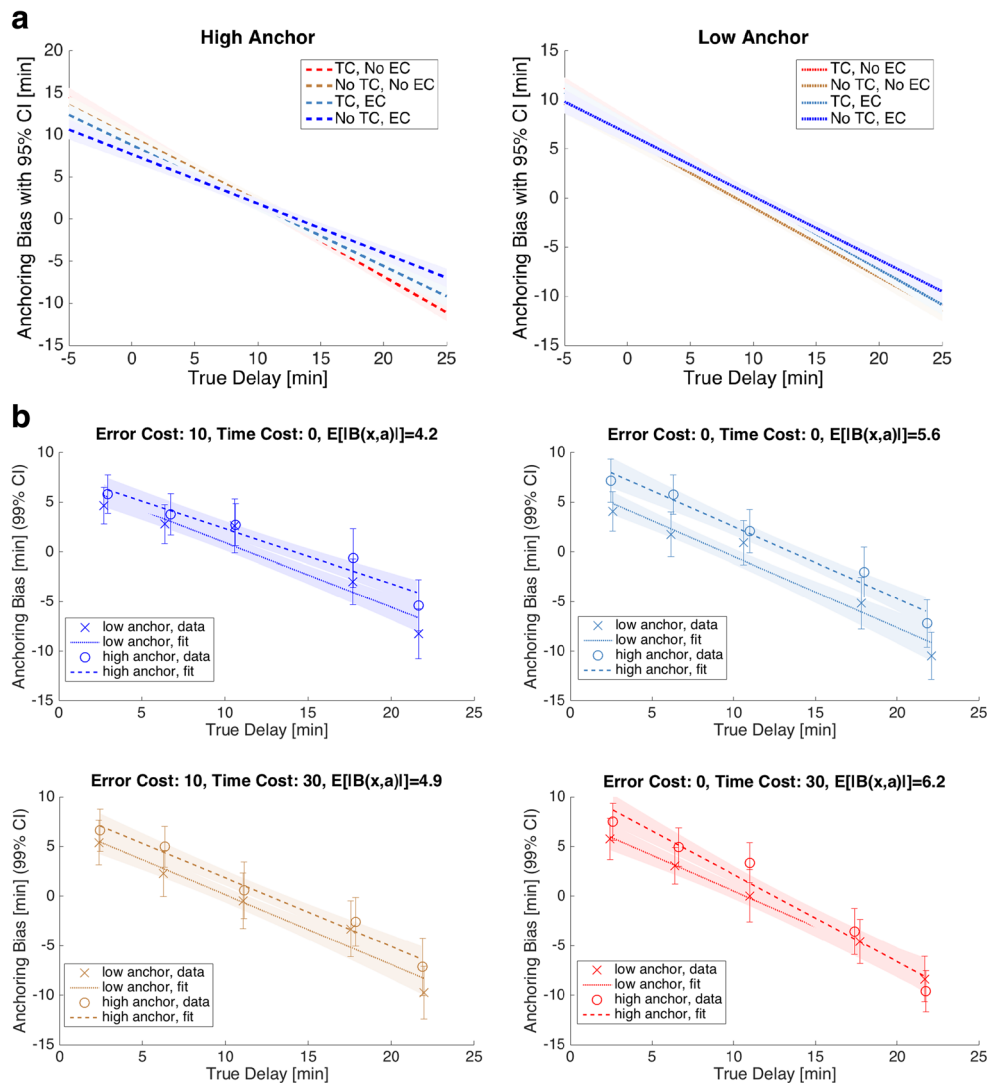


Fig. 8 Effect of incentives for speed and accuracy when a high anchor was provided confirm our theory's prediction; cf. Figure 1. **A**) Anchoring bias of our participants' judgments as a function of the true delay as estimated with the linear regression model of anchoring (1–3) for the low-anchor condition (*left*) and the high-anchor condition (*right*). **B**) Average anchoring bias of our participants' judgments binned by the true delay along with the fit of the linear regression model of anchoring (1–3) for each of the four incentive conditions

a conversational hint that the correct value is close to that value (Zhang and Schwarz, 2013).

Conclusion

Our experiments confirmed the predictions of resource-rational anchoring and adjustment. Most importantly, people appear to rationally adapt their number of adjustments to the cost of time and error. When errors were costly, people invested more time and were more accurate, their adjustments were larger, and their anchoring bias was smaller. By contrast, when time was costly then our participants were faster and less accurate, their adjustments appeared to be smaller, and their anchoring bias was larger. This

is consistent with the rational allocation of mental effort (Shenhav et al., 2017) and our hypothesis that the number of adjustments is chosen to achieve an optimal speed-accuracy tradeoff. However, since our experiment used only two levels of time cost and error cost, it remains to be investigated whether the number of adjustments changes gradually with its costs and benefits—as predicted by the resource-rational model—or whether people can only choose between a fast versus a slow mode of numerical estimation as postulated by dual-systems theories (Evans, 2008; Kahneman, 2011). Hence, while people's estimates are biased in the statistical sense of the word, this bias might be consistent with how people ought to reason. In this sense, the anchoring “bias” might not be a cognitive bias after all.

Acknowledgements This research was supported by grant number ONR MURI N00014-13-1-0341 from the Office of Naval Research (TLG and NDG), grant number FA-9550-10-1-0232 from the Air Force Office of Scientific Research (TLG), and a John S. McDonnell Scholar Award (NDG).

References

- Epley, N., & Gilovich, T. (2005). When effortful thinking influences judgmental anchoring: differential effects of forewarning and incentives on self-generated and externally provided anchors. *Journal of Behavioral Decision Making*, 18(3), 199–212.
- Epley, N., & Gilovich, T. (2006). The anchoring-and-adjustment heuristic. *Psychological Science*, 17(4), 311–318.
- Evans, J. S. B. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, 59, 255–278.
- Jacowitz, K. E., & Kahneman, D. (1995). Measures of anchoring in estimation tasks. *Personality and Social Psychology Bulletin*, 21(11), 1161–1166.
- Kahneman, D. (2011). *Thinking, fast and slow*. New York: Farrar, Strauss and Giroux.
- Lieder, F., Goodman, N. D., & Griffiths, T. L. (2013). *Reverse-engineering resource-efficient algorithms*. Lake Tahoe, USA: Paper presented at NIPS-2013 Workshop Resource-Efficient ML.
- Lieder, F., Griffiths, T., Huys, Q., & Goodman, N. (2017b). Testing models of anchoring and adjustment. PsyArXiv Preprint. Retrieved from <https://osf.io/preprints/psyarxiv/94yvy> (Version 3).
- Lieder, F., Griffiths, T. L., & Goodman, N. D. (2012). Burn-in, bias, and the rationality of anchoring. In Bartlett, P., Pereira, F.C.N., Bottou, L., Burges, C.J.C., & Weinberger, K.Q. (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. (2017a). The anchoring bias reflects rational use of cognitive resources. *Psychonomic Bulletin and Review*.
- Russo, J. E., & Schoemaker, P. J. H. (1989). Decision traps: Ten barriers to brilliant decision-making and how to overcome them. Simon and Schuster.
- Shenhav, A., Musslick, S., Lieder, F., Kool, W., Griffiths, T., Cohen, J., & Botvinick, M. (2017). Toward a rational and mechanistic account of mental effort. *Annual Review of Neuroscience*, 40.
- 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.
- Stanovich, K. E. (2009). Decision making and rationality in the modern world.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131.
- Zhang, Y. C., & Schwarz, N. (2013). The power of precise numbers: A conversational logic analysis. *Journal of Experimental Social Psychology*, 49(5), 944–946.