

Bayesian signatures of confidence and central tendency in perceptual judgment

Yang Xiang,¹ Thomas Graeber,² Benjamin Enke,³ Samuel J. Gershman^{4,5}

¹Sloan School of Management, MIT

²Harvard Business School

³Department of Economics, Harvard University

⁴Department of Psychology and Center for Brain Science, Harvard University

⁵Center for Brains, Minds and Machines

*Corresponding author: yyx@mit.edu

October 12, 2020

Abstract

Despite a surge of interest in theoretical models of Bayesian cognition, direct empirical evidence for their validity remains limited. This paper studies Bayesian signatures in perceptual judgment, focusing on the *central tendency effect*: perceptual judgments are consistently biased towards the center of the stimulus distribution. Based on a formal Bayesian framework, we show that lower subjective confidence as a measure of posterior uncertainty about a judgment predicts (i) a lower sensitivity of magnitude estimates to objective stimuli; (ii) a higher sensitivity to the mean of the stimulus distribution; (iii) a stronger central tendency effect at higher stimulus magnitudes; and (iv) higher response variability. To test these predictions, we collect a tailored large-scale experimental data set and additionally re-analyze perceptual judgment data from several previous experiments. Across data sets, subjective confidence is strongly predictive of the central tendency effect and response variability, both correlationally and when we exogenously manipulate the magnitude of sensory noise. Our results lend support to Bayesian explanations of both confidence and the central tendency effect.

Introduction

Processing of sensory information is unreliable, corrupted by various sources of noise. Bayesian theories of perception hold that our interpretation of sensory information compensates for noise by relying on prior beliefs (Knill and Richards, 1996). Despite biasing perception towards the prior, this compensation minimizes perceptual error on average by suppressing unsystematic errors due to noise.

Recently, Petzschner et al. (2015) proposed that a Bayesian framework may account for a large number of stylized facts from perceptual judgment tasks that are related to the well-known central tendency (or regression) effect: across perceptual domains, participants' estimates of stimulus magnitude are consistently biased towards the center of the magnitude distribution (Hollingworth, 1910, Stevens and Greenbaum, 1966). While there is indirect evidence consistent with a link between the observed judgment patterns and Bayesian models, there is no contribution that comprehensively tests the predictions of Bayesian models in these settings.

This paper aims to provide such evidence by using measures of subjective confidence as a proxy for the subjective posterior probability that the magnitude estimate is correct (the Bayesian confidence hypothesis, Lebreton et al. (2015), Meyniel et al. (2015)). We begin by laying out a formal framework for analyzing magnitude estimation that directly links subjective confidence to the central tendency effect. This model makes several main predictions, elaborated in the next section: (i) confidence should decrease with sensory noise; (ii) the sensitivity of magnitude estimates to actual magnitudes should increase with subjective confidence and decrease with sensory noise; (iii) the central tendency effect should decrease with confidence and increase with sensory noise; (iv) both confidence and sensitivity should decrease with stimulus magnitude, assuming magnitude-dependent sensory noise; and (v) response variability should increase with sensory noise and decrease with confidence when prior uncertainty is high relative to sensory noise.

We test these model predictions in two ways. First, we implement a new large-scale magnitude estimation experiment. We elicit both magnitude estimates and subjective confidence. A key feature of our experiment is that we exogenously vary the objective stimulus, the mean of the stimulus distribution, and the magnitude of sensory noise. We collect these data because, as explained below, existing data sets on magnitude estimation and subjective confidence lack sufficient variation in some of the key elements of Bayesian models. Second, moving beyond our new experiment, we re-analyze data from several earlier studies that measured both continuous reports of stimulus magnitude and subjective confidence. Although these earlier studies are limited in several ways (which motivated our new experiment), they provide converging evidence for some of our predictions¹.

Theoretical Framework

To motivate our empirical predictions, we will first lay out a theoretical framework based on a simple Bayesian estimation problem, which mirrors the experimental tasks given to subjects.

Bayesian estimation

We model a task in which subjects are asked to estimate the magnitude of a stimulus x from a noise-corrupted signal $s = x + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$ is Gaussian-distributed sensory noise. If the subject has a Gaussian prior over the magnitude, $x \sim \mathcal{N}(\mu_x, \sigma_x^2)$, then the posterior is also Gaussian:

$$P(x|s) = \mathcal{N}(x; \lambda s + (1 - \lambda)\mu_x, \sigma_x^2), \quad (1)$$

where the posterior variance is

$$\sigma_x^2 = (1 - \lambda)\sigma_x^2, \quad (2)$$

and the *sensitivity* is

$$\lambda = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_\epsilon^2}, \quad (3)$$

which takes values between 0 and 1. Intuitively, the magnitude estimate will be more sensitive (higher λ) to the signal when the sensory noise variance (σ_ϵ^2) is small relative to the prior uncertainty (σ_x^2).

¹The data and code are available at https://github.com/yyyxiang/confidence_central_tendency

We follow most of the literature in assuming that the prior approximates either the experienced magnitude distribution or the instructed distribution, depending on the paradigm. For analytical tractability, we assume a Gaussian prior even when the experimentally controlled stimulus distribution is non-Gaussian (e.g., uniform). In these cases, the mean is assumed to be equal to the center of the uniform distribution (e.g., Acerbi et al., 2012, Roach et al., 2017). Our experimental predictions do not depend strongly on the choice of prior variance, provided that it is a constant function of the magnitude distribution range.

We posit that the subjective estimate \hat{x} is chosen to minimize expected loss $\mathcal{L}(\hat{x}, x)$ conditional on the signal:

$$\hat{x} = \operatorname{argmin}_{\hat{x}} \mathbb{E}[\mathcal{L}(\hat{x}, x)|s]. \quad (4)$$

The optimal subjective estimate under the quadratic loss $\mathcal{L}(\hat{x}, x) = (x - \hat{x})^2$ equals the posterior mean (see, e.g., Berger, 1985):

$$\hat{x} = \lambda s + (1 - \lambda)\mu_x. \quad (5)$$

For the Gaussian estimation problem described here, this prediction holds true for several other loss functions. For example, the Bayes-optimal estimator for the absolute loss $\mathcal{L}(\hat{x}, x) = |x - \hat{x}|$ is the posterior median, and the Bayes-optimal estimator for the relaxed 0-1 loss ($\mathcal{L}(\hat{x}, x) = 1$ if $|x - \hat{x}| \geq \kappa$, 0 otherwise) is the posterior mode. Since the mean, median, and mode for a Gaussian distribution have the same value, our predictions are invariant across these choices of loss function.

Subjective confidence

Magnitude estimation tasks require subjects to report a single point estimate of the magnitude, but Bayesian models hypothesize that subjects are representing an entire distribution over magnitudes. Subjective confidence judgments can potentially provide a probe of this distributional representation. According to the Bayesian confidence hypothesis (Aitchison et al., 2015, Fleming and Daw, 2017, Meyniel et al., 2015, Pouget et al., 2016, Rahnev et al., 2015, Sanders et al., 2016), subjective confidence corresponds to the posterior probability that an action is optimal (in this context, the probability that the subjective estimate equals the objective stimulus magnitude). Note that for continuous magnitudes, the probability that the point estimate \hat{x} equals the objective magnitude x is 0. However, we can evaluate the posterior probability that the estimate falls within $\pm\kappa$ of the objective magnitude; this is 1 minus the relaxed 0-1 loss, which we denote by $\mathcal{L}_\kappa(\hat{x}, x)$. Taking the limit as κ goes to 0 gives us the probability of the objective magnitude falling within an infinitesimally small region around the stimulus estimate. For the Gaussian estimation problem, this is:

$$1 - \lim_{\kappa \rightarrow 0} \mathcal{L}_\kappa(\hat{x}, x) = \mathcal{N}(x; \hat{x}, \sigma_{\hat{x}}^2), \quad (6)$$

where we have assumed that \hat{x} is the posterior mean (or equivalently the median or mode for Gaussian posteriors). Since x is unobserved, we take the expectation conditional on the signal s :

$$\mathbb{E}[1 - \lim_{\kappa \rightarrow 0} \mathcal{L}_\kappa(\hat{x}, x)|s] = \frac{\sqrt{2/\pi}}{\sigma_{\hat{x}}} \equiv c. \quad (7)$$

Thus, we see that Bayesian confidence (c) for the Gaussian estimation problem is inversely proportional to the posterior standard deviation. This result motivates the use of the standard deviation as a measure of “cognitive uncertainty” (Enke and Graeber, 2020).² Most importantly, note that under a constant prior variance σ_x^2 , Bayesian confidence c decreases in the variance of sensory noise, σ_ϵ^2 . Because our experimental design controls the prior variance as argued above, we will in the following interpret Bayesian confidence as an approximate measure of sensory noise.

Predictions

Using this framework, we can now lay out a set of theoretical predictions, which we will empirically test in subsequent sections.

Prediction 1 *Sensitivity (λ) monotonically increases with confidence (c).*

This prediction follows from the expression relating sensitivity to confidence:

$$\lambda = 1 - \frac{2}{\pi c^2 \sigma_x^2}. \quad (8)$$

Although λ is not directly observable, we can estimate it empirically using linear regression. The next prediction concerns the relationship between confidence and the central tendency effect, a straightforward corollary of Prediction 1 because $\frac{\partial \hat{x}}{\partial \mu_x} = 1 - \lambda$. As discussed above, we follow previous work in assuming that the prior is determined by the experienced or instructed magnitude distribution, and hence μ_x is the mean of the magnitude distribution.

Prediction 2 *The central tendency effect ($\frac{\partial \hat{x}}{\partial \mu_x}$) monotonically decreases with confidence.*

The aforementioned predictions are correlational in nature and rest on naturally-occurring variation in confidence. We now state the causal analogs of these predictions.

Prediction 3 *An exogenous increase in sensory noise (σ_ϵ^2), which reduces subjective confidence, decreases sensitivity.*

Prediction 4 *An exogenous increase in sensory noise increases the central tendency effect.*

Next, we theoretically explore how the strength of the central tendency effect depends on the stimulus magnitude. A well-known phenomenon is that response variability increases with stimulus magnitude according to Weber’s law. A common interpretation is that the signal-to-noise ratio decreases with stimulus magnitude, due either to a non-linear transformation of magnitude (e.g., Fechner, 1860, Nieder and Miller, 2003, Petzschner and Glasauer, 2011, Roach et al., 2017, Stevens, 1961) or to magnitude-dependent scaling of sensory noise (e.g., Gibbon, 1977, Gibbon and Church, 1981, Treisman, 1964). For concreteness, we examine the implications of the latter assumption:

$$\frac{\partial \sigma_\epsilon^2}{\partial x} > 0 \Rightarrow \frac{\partial c}{\partial x} < 0 \quad \text{and} \quad \frac{\partial \lambda}{\partial x} < 0 \quad (9)$$

²Note that the measure of Bayesian confidence described here does not directly map onto our experimental instructions, which simply asked subjects to report their subjective confidence using a Likert scale. However, this formulation has the advantage of being directly related to earlier formulations for discrete decisions.

Prediction 5 *As the stimulus magnitude increases, confidence and sensitivity decrease. The latter effect implies a stronger central tendency effect for larger magnitudes.*

Finally, we turn to response variability. Theoretically, variability is affected both by the effect of sensory noise on estimates (which increases variability) and the counteracting central tendency effect (which decreases variability; see Enke and Graeber, 2020). The expression for response variability is given by:

$$\text{Var}(\hat{x}) = \lambda^2 \sigma_\epsilon^2 = \left(\frac{\sigma_x^2}{\sigma_x^2 + \sigma_\epsilon^2} \right)^2 \sigma_\epsilon^2. \quad (10)$$

The relationship between response variability and other quantities depends on the degree of prior uncertainty relative to sensory noise.

Prediction 6 *When sensory noise is small relative to prior uncertainty ($\sigma_\epsilon^2 < \sigma_x^2$), response variability increases in sensory noise variance and decreases in confidence. When sensory noise is large relative to prior uncertainty ($\sigma_\epsilon^2 > \sigma_x^2$), response variability decreases in sensory noise variance and increases in confidence.*

Experiment

Although a considerable number of studies have implemented magnitude estimation tasks while also measuring subjective confidence, none of these studies explicitly vary all of the variables that are of interest in light of our theoretical predictions above. Specifically, we require a study setup that features (i) variation in stimuli; (ii) a meaningful degree of exogenous variation in the mean of the stimulus distribution; (iii) exogenous variation in sensory noise; and (iv) a large sample size to allow for sufficiently powered statistical analyses of the relationship between confidence and the precise mechanics of central tendency. Earlier studies typically had sample sizes of fewer than 50 participants (see Table 2).

Materials and Methods

Participants

We recruited 300 participants from Amazon Mechanical Turk (MTurk). All participants gave informed consent prior to testing. To ensure that participants fully understood the experiment, they completed a comprehension check immediately after the instructions. Participants who failed the comprehension check were asked to leave the study and were compensated (\$0.5) for their time. Participants who passed the comprehension check proceeded with the experiment and were paid (\$4.5) for their participation. Participants also had the opportunity to earn a performance-based bonus payment (\$1) if their estimate on a randomly chosen trial was within 2 of the stimulus magnitude. The experiment was approved by the Harvard Institutional Review Board.

Stimuli

The stimuli were arrays of black dots on a white background. The number of dots ranged between 15 and 65. Each participant completed six blocks with a total of 240 arrays. For each participant,

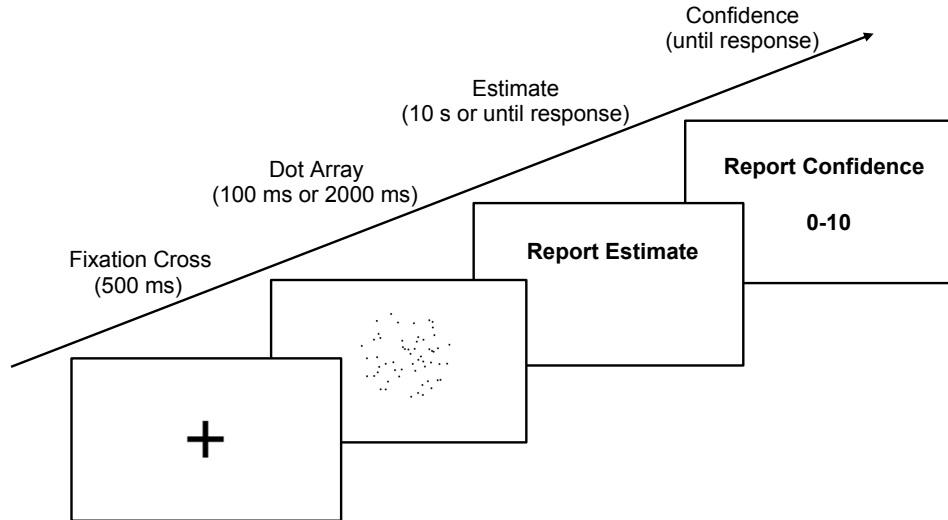


Figure 1: Illustration of the experiment. Each trial begins with a fixation cross, followed by a dot array. Participants then report their numerosity estimate and confidence.

the average stimulus magnitude within each block was randomly drawn from a discrete uniform distribution on the integers from 30 to 50 (sampled without replacement). For each block, the within-block distribution was randomly drawn from a uniform distribution centered at the average stimulus magnitude, ranging between ± 15 (sampled with replacement). This procedure ensured that our trials feature substantial variation both in actual stimuli within-block and in average stimuli across blocks. We conceptualize the latter as the perceptual prior and use the variation to test Prediction 4. To make the prior salient to subjects, they were informed about the average stimulus value within a block at the beginning of each block, before making any estimates.

The stimulus duration on each trial was either 100 ms or 2000 ms, intermixed randomly within each block; each of the two duration conditions appeared on half of the trials on each block. This feature of our experiment allows us to leverage exogenous variation in sensory noise and confidence to test Prediction 3. This is motivated by the common assumption that sensory evidence is accumulated across time, such that the signal-to-noise ratio is higher for longer durations (Cheyette and Piantadosi, 2019, 2020, Inglis and Gilmore, 2013).

Procedure

As illustrated in Figure 1, each trial began with the presentation of a fixation cross at the center of the screen for 500 ms. Then a random dot array was presented for either 100 ms or 2000 ms. Next, participants had 10 s to type in their numerosity estimate using the number pad on the keyboard. Participants were then prompted to provide a confidence rating by clicking on a discrete slider from 0-10.

Data analysis

We analyzed the data using linear mixed-effects models. Model 1 regresses subjective estimates on the true stimulus value (Stimulus) and the average stimulus magnitude in a block (AveStim). In the absence of a central tendency effect, the coefficient of the true stimulus should be one, while the coefficient of average stimulus should be zero. The central tendency effect is indicated by a stimulus coefficient of less than one and an average stimulus coefficient of greater than zero.

To explore the role of sensory noise and confidence for the central tendency effect, we ran Model 2. Here, we added as additional regressors subjective confidence, the interaction between confidence and the stimulus, and the interaction between confidence and average stimulus. As derived above, our hypothesis is that higher confidence (as a consequence of lower sensory noise) is associated with a higher responsiveness to the stimulus and a lower responsiveness to the average stimulus, meaning that the first interaction effect should be positive and the second one negative.

In addition, Model 2 also accounts for our exogenous variation in stimulus duration, which we conceptualize as exogenous variation in sensory noise that translates into confidence. We include as regressors a binary Condition indicator (0 for the 100 ms condition, 1 for the 2000 ms condition), the interaction between Stimulus and Condition, and the interaction between average stimulus magnitude and Condition.

To see how the proportional regression to the prior changes in stimulus magnitude, we specified Model 3. Here, we regressed an empirical estimate of λ on stimulus magnitude and hypothesize a negative coefficient.³

Turning to the analysis of response variability, in Model 4, we regressed Variability (response standard deviation⁴ across multiple repetitions of the same magnitude) on Confidence (averaged across stimulus repetitions). In Model 5, we regressed Variability on Condition. For all five models, we included random effects for the intercept and each regressor. For better interpretability, we standardized the Confidence coefficients.

Results

Preliminaries

Our study rests on two prerequisites: (i) The existence of a central tendency effect; and (ii) variation in subjective confidence as a function of stimulus duration.

The results of Model 1 (column 1 of Table 1) confirm the existence of a central tendency effect in our data: the coefficient of the true stimulus is substantially smaller than one [$F(1, 69599) = 1577.20, p < 0.0001$], and the coefficient of the average stimulus is considerably larger than zero [$F(1, 69599) = 117.34, p < 0.0001$]. To visualize the central tendency effect, we plotted subjective estimates as a function of stimulus magnitude. As shown in Figure 2A, the slope of the estimation

³Given estimate \hat{x} , the empirical analogue of λ is constructed by noting that:

$$\tilde{\lambda} := \mathbb{E}[\lambda|x] = \frac{\hat{x} - \mu_x}{x - \mu_x} \quad (11)$$

That is, the empirical analogue of the model quantity λ is the distance between mean perceptual judgments and the prior mean, $(\hat{x} - \mu_x)$, as a fraction of the distance between the true stimulus and the prior, $(x - \mu_x)$. In this regression, we excluded $\tilde{\lambda}$ values of \pm infinity and those where $\tilde{\lambda}$ was undefined due to a denominator of zero. We excluded 3% of the sample this way.

⁴We used standard deviation instead of variance because we found it to be slightly better behaved. The predictions are qualitatively the same for both standard deviation and variance.

functions has a slope considerably smaller than 1, and the level of the estimation function increases in the mean of the stimulus distribution. These are the two canonical features of the central tendency effect.

Our regression analysis revealed that confidence was higher for longer durations [$t(71405) = 16.20, p < 0.0001$]⁵, consistent with the hypothesis that lower sensory noise registers as higher confidence (Figure 2C).

Sensory noise and the central tendency effect: Correlational evidence

We begin our main analysis by testing Predictions 1 and 2. The results of Model 2 (column 2 of Table 1) show that the central tendency effect is strongly mediated by subjective confidence. The positive interaction effect of confidence and stimulus implies that, for every standard deviation of confidence, the responsiveness of subjective estimates to the stimulus value increases by 5 percentage points [$t(71399) = 6.71, p < 0.0001$]. This result confirms Prediction 1. Similarly, the negative interaction effect of confidence and average stimulus shows that confident subjects place substantially lower weight on the mean of the stimulus distribution [$t(71399) = -3.18, p < 0.01$], supporting Prediction 2 that subjective estimates are pulled towards the prior mean to a greater extent when confidence is low. These two interaction effects correspond to the central predictions of the model about the role of sensory noise for the central tendency effect. Figure 2B visualizes the results by revealing stronger central tendency effects when confidence is low, i.e., the implied response curve is substantially flatter.

Sensory noise and the central tendency effect: Causal evidence

To confirm our correlational findings, we leverage the exogenous variation in confidence that is induced by stimulus duration. As shown in column 2 of Table 1, longer stimulus durations increased the effect of objective stimulus magnitude [$t(71399) = 20.91, p < 0.0001$] and decreased the effect of average stimulus magnitude [$t(71399) = -12.24, p < 0.0001$], consistent with a weaker central tendency effect when sensory noise is lower. In combination with the strong effect of duration on confidence, these results suggest that sensory noise causally influences both confidence and central tendency effects, confirming Prediction 3 and Prediction 4. Figure 2D visualizes these patterns by showing that the central tendency effect is considerably more pronounced under shorter stimulus duration, i.e., the implied response curve is flatter.

Sensory noise and the magnitude-dependent central tendency effect

We find strong support for Prediction 5. First recall from Figure 2C that confidence decreases in the stimulus magnitude, indicating larger sensory noise. Our model therefore predicts proportionally larger compression (lower λ) at higher magnitudes. The results of Model 3 show that the empirical analog of λ indeed decreases in stimulus magnitude [$t(69087) = -17.17, p < 0.0001$], meaning that participants display proportionally stronger central tendency effects at higher stimulus magnitudes. We can also intuitively gauge this effect in Figure 2A, where estimates are further away from the 45-degree line at higher magnitudes.

⁵This result came from an additional regression, where we regressed Confidence on Stimulus and Condition.

Sensory noise and response variability

Figures 2E and 2F visualize the results on the relationship between sensory noise and response variability. Note that Prediction 6 implies a hump-shaped relationship between response variability and confidence: response variability should increase at low confidence levels (sensory noise variance high relative to prior uncertainty) but decrease at high confidence levels (sensory noise variance relatively low). Figure 2E strongly supports this distinctive prediction of the Bayesian model. A plausible assumption is that prior uncertainty usually exceeds sensory noise variance. Correspondingly, we find that response variability decreases across most of confidence range in Figure 2E and results from our regression Model 4 show a negative effect of confidence [$t(11979) = -6.93, p < 0.0001$]. Turning to the causal manipulation, we see from Model 5 results that variability is higher under shorter stimulus duration [$t(18181) = -5.36, p < 0.0001$; also see Figure 2F]. Thus the data clearly support the model predictions about the relationship between variability, sensory noise and subjective confidence.

Discussion

Taken together, the results establish support for all six predictions spelled out above. In particular, (i) lower sensory noise (high confidence) is associated with a *higher* sensitivity of estimates to the true stimulus; (ii) lower sensory noise (higher confidence) is associated with a *lower* sensitivity of estimates to the average stimulus; (iii) these results hold both in correlational analyses and when we exogenously manipulate confidence; (iv) central tendency increases with the stimulus magnitude; and (v) response variability decreases with confidence and increases under shorter stimulus duration.

Re-analysis of Earlier Studies

While our experiment has the advantage of being specifically tailored to investigate the predictions associated with the core elements of a Bayesian model, we sought to test the validity of our hypothesis more generally. Owing to the recent publication of the Confidence Database (Rahnev et al., 2020), we were able to additionally address the relationship between confidence and the central tendency effect by re-analyzing data from several earlier studies.

Materials and Methods

Data sets

Out of a total of 145 studies contained in the Confidence Database (Rahnev et al., 2020), we identified six studies that elicited *continuous* reports of both stimulus magnitude and subjective confidence (Table 2). Two of these studies use circular stimuli (motion direction and grating orientation), and hence are not relevant for analyses of magnitude-dependent noise. Nonetheless, these studies are still informative about the interplay between objective magnitude and confidence in determining subjective judgments.

Data analysis

To measure the central tendency effect, we specified the average stimulus magnitude as the cumulative mean of all across-block stimulus magnitudes that preceded the current trial. We excluded

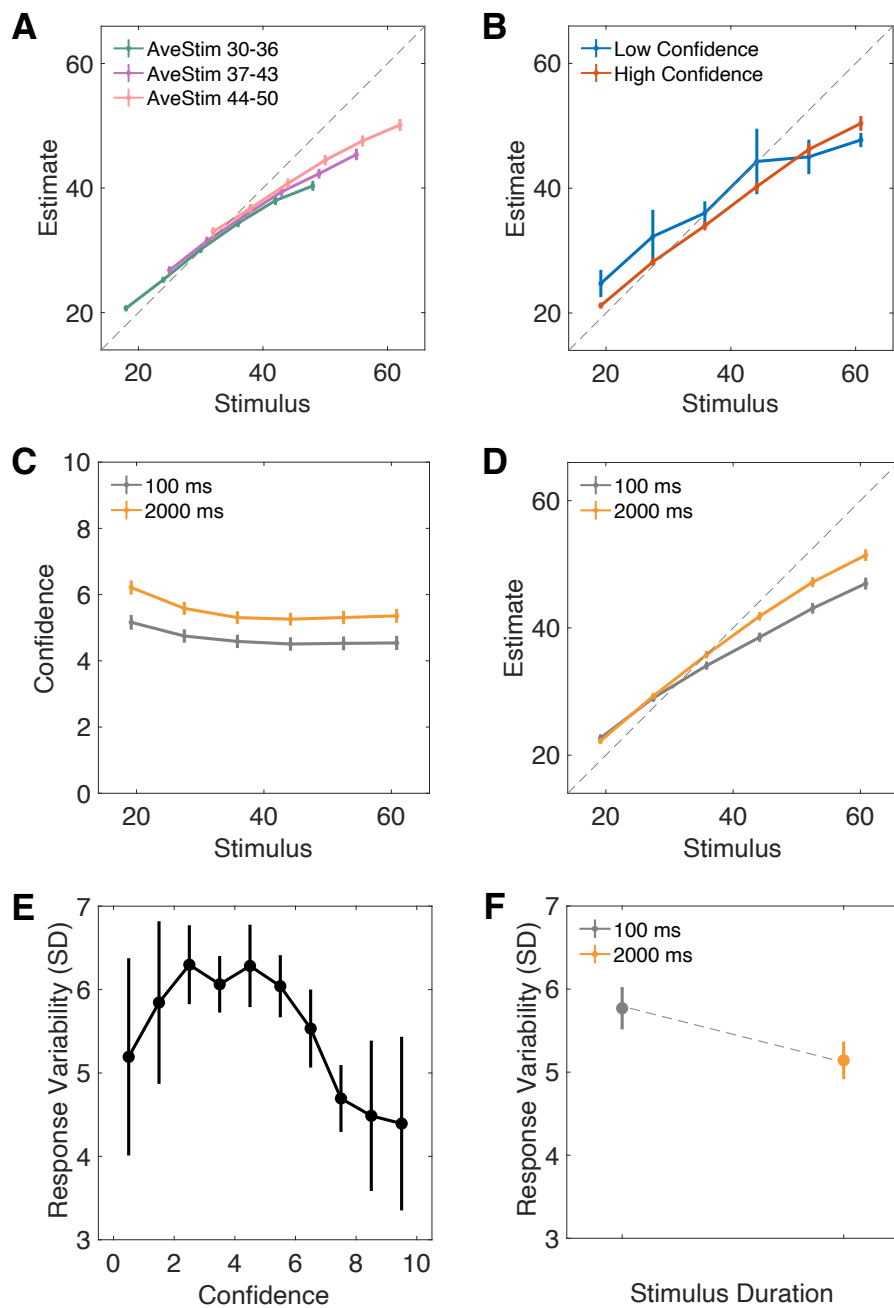


Figure 2: (A) Subjective estimates as a function of objective stimulus magnitude, shown separately for three within-block average stimulus ranges. (B) Subjective estimates as a function of objective stimulus magnitude, shown separately for low confidence (confidence levels 0-3) and high confidence (confidence levels 7-10). (C) Confidence as a function of objective stimulus magnitude, shown separately for different stimulus duration conditions. (D) Subjective estimates as a function of objective stimulus magnitude, shown separately for different stimulus durations. (E) Response standard deviation as a function of confidence. (F) Response standard deviation under short and long stimulus duration. Error bars indicate 95% confidence intervals.

Table 1: Regression coefficients and standard errors for stimulus estimates in the new data set (Models 1 and 2).

(Intercept)	7.435**** (0.464)	6.666**** (0.560)
Stimulus	0.637**** (0.010)	0.537**** (0.011)
AveStim	0.110**** (0.013)	0.210**** (0.017)
Confidence		-1.483**** (0.370)
Stimulus:Confidence		0.049**** (0.007)
AveStim:Confidence		-0.036** (0.011)
Condition		2.481**** (0.496)
Stimulus:Condition		0.178**** (0.009)
AveStim:Condition		-0.176**** (0.014)
Observations	71408	71408

Note: Condition refers to the 2000 ms stimulus duration condition.
The Confidence coefficients are standardized.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, **** $p < 0.0001$

each participant’s first trial in regressions involving average stimulus, because first trials do not have preceding trials. In our analysis, we excluded estimates that were clearly out of range and trials on which participants did not comply with the experiment rules.⁶

To assess the relationship between confidence and central tendency, we constructed linear mixed-effects models akin to the ones used to analyze our own experiment, except that here we do not have treatment conditions that exogenously vary sensory noise and confidence through stimulus duration.

Results and discussion

Subjective estimates correlate more strongly with objective stimulus magnitude under high confidence

Consistent with the model and our own experimental results, we document significant positive interactions between stimulus and confidence in all data sets (Table 3): AB17 [$t(15254) = 4.66, p < 0.0001$], DB18 [$t(14432) = 3.78, p < 0.001$], DB19 [$t(8563) = 3.58, p < 0.001$], DB20 [$t(9651) = 6.16, p < 0.0001$], RZ14 [$t(8974) = 22.84, p < 0.0001$], and SP17 [$t(14059) = 11.57, p < 0.0001$]. In

⁶These are trials with estimates 999 or NaN in DB18, trials with estimates 9999 or 999 in DB19, and trials on which participants did not respond.

Table 2: Descriptive information about the included studies.

Data set	Authors	Year	Stimuli	# of Participants
AB17	Akdogan & Balçı	2017	Duration of Time	76
DB18	Duyan & Balçı	2018	Number of Beeps	44
DB19	Duyan & Balçı	2019	Number of Dots	18
DB20	Duyan & Balçı	2020	Length of Lines	41
RZ14	Rausch & Zehetleitner	2014	Direction of Coherent Motion	20
SP17	Samaha & Postle	2017	Orientation of Grating	40

Note: Experiment 4 in data set AB17 was excluded because each participant was presented with a single stimulus magnitude. Experiment 3 in data set SP17 was excluded because it did not elicit continuous report of stimulus magnitude.

other words, subjective estimates correlated more strongly with the objective stimulus magnitude when confidence was high, as predicted by Bayesian models of perception.

In contrast to the consistent interaction between objective stimulus magnitude and confidence across studies, the interaction between average stimulus magnitude and confidence was not consistently observed across studies (Table 3). We found a significant effect in only one of the studies [AB17, $t(15254) = -3.88, p < 0.001$], which had the predicted sign. We conjecture that this is due to insufficient variation in the average stimulus magnitude in prior experiments. In our new study reported above, we remedied this limitation by exploring a wider range of average stimulus magnitudes to be sufficiently statistically powered.

The low variation in stimulus magnitude in the earlier studies also makes our analysis of Prediction 5 (Model 3) difficult. We hypothesized that the magnitude of regression decreases in stimulus magnitude. The results of Model 3 are mixed, with a negative effect for only two studies, and neither of them is significant [AB17, $t(12454) = -0.77, p = 0.44$; DB19, $t(8564) = -0.31, p = 0.75$]. Again, this null result is to be expected given insufficient variation in (average) stimulus magnitude. Our new experiment deliberately remedies this shortcoming.

Response variability decreases with confidence

Finally, we again find strong evidence for confidence-dependent variability. Five studies show a significantly negative coefficient: AB17 [$t(150) = -2.07, p < 0.05$], DB18 [$t(86) = -3.96, p < 0.001$], DB19 [$t(466) = -5.62, p < 0.0001$], RZ14 [$t(2538) = -10.58, p < 0.0001$], and SP17 [$t(4169) = -4.60, p < 0.0001$]. These results are very similar to the ones observed in our own experimental data.

General Discussion

Using data from earlier studies and a new data set, we have established a relationship between confidence and central tendency that conforms with the generic predictions of Bayesian models. First, we showed that the central tendency effect is lower on high confidence trials. Second, we showed that this relationship is causal: when sensory noise was exogenously increased via a stimulus duration manipulation, the central tendency effect increased and confidence decreased. Third, we showed a stronger central tendency effect at higher magnitudes, which is in line with subjective

Table 3: Regression coefficients and standard errors for the re-analysis of stimulus estimates (Model 2).

Dataset	AB17	DB18	DB19	DB20	RZ14	SP17
(Intercept)	0.512**** (0.116)	0.363 (0.929)	6.367**** (1.253)	93.025**** (14.287)	72.216**** (16.608)	123.56**** (8.076)
Stimulus	0.849**** (0.017)	0.900**** (0.042)	0.631**** (0.040)	0.671**** (0.021)	0.352**** (0.031)	0.537**** (0.019)
AveStim	0.027 (0.036)	0.153 (0.078)	-0.103 (0.082)	-0.023 (0.035)	0.244** (0.083)	0.072 (0.094)
Confidence	-0.030 (0.037)	-0.357 (0.243)	-3.527** (1.237)	-2.558 (8.303)	-49.055*** (14.176)	-26.668*** (7.255)
Stimulus:Confidence	0.035**** (0.007)	0.058*** (0.015)	0.076*** (0.021)	0.056**** (0.009)	0.244**** (0.011)	0.190**** (0.016)
AveStim:Confidence	-0.054*** (0.014)	-0.027 (0.022)	0.086 (0.080)	-0.035 (0.022)	0.036 (0.077)	0.108 (0.085)
Observations	15260	14438	8569	9657	8980	14065

Note: The Confidence coefficients are standardized.

$p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, **** $p < 0.0001$

confidence decreasing in stimulus magnitude. Fourth, we showed that across-trial variability in responses decreased in subjective confidence and increased in sensory noise whenever prior uncertainty is relatively large.

Our findings bridge several disparate lines of research on confidence and central tendency effects. One line of research has presented evidence that confidence judgments in perceptual decision making tasks reflect the posterior probability of being correct—the *Bayesian confidence hypothesis* (Aitchison et al., 2015, Fleming and Daw, 2017, Meyniel et al., 2015, Pouget et al., 2016, Rahnev et al., 2015, Sanders et al., 2016). While most work on the Bayesian confidence hypothesis has focused on discrete choice tasks, here we analyzed continuous report tasks, which allowed us to relate confidence judgments to the central tendency effect. Our finding that this relationship held across several different stimulus domains (time, visual numerosity, auditory numerosity, line length, motion direction, and grating orientation) lends support to the generality of our conclusions.

While our findings are specifically consistent with the Bayesian confidence hypothesis, they might also be compatible with alternative models. For example, Adler and Ma (2018) developed several models that map probability representations of uncertainty onto confidence in a non-Bayesian way, and presented experimental evidence that some of these models outperformed the Bayesian model in predicting confidence judgments. Li and Ma (2020) developed a different non-Bayesian model, which determined confidence based on the difference in probability between the top two hypotheses. For our purposes, all of these non-Bayesian models share the key property that confidence is lower when uncertainty (due to sensory noise) is greater.

A second line of research bridged by our theory is on the central tendency effect. Recent research has shown that cognitive load (which ostensibly increases sensory noise) strengthens the central tendency effect, broadly consistent with Bayesian models of perception. For example, Allred et al. (2016) found that asking participants to memorize six-digit numbers (high load condition) increased the central tendency effect in the estimation of line length, compared to a low load

condition in which participants memorized two-digit numbers. Similarly, Olkkonen et al. (2014) found that increasing chromatic noise or the delay between stimulus presentation and estimation increased the central tendency effect in a color estimation task. Relatedly, there is evidence that the central tendency effect is stronger when sensory information is less reliable or when the magnitude distribution is more concentrated around the center (Allred et al., 2016, Ashourian and Loewenstein, 2011, Huttenlocher et al., 2000, Olkkonen et al., 2014), again consistent with Bayesian models. Our study took this line of research one step further, showing that confidence judgments respond to an exogenous manipulation of sensory noise, while explaining significant additional response variance not explained by noise alone.

Our study is limited in several ways that can be addressed by future research. One is that we did not provide a detailed information processing model of behavior. This was a deliberate choice: our aim was to test generic principles rather than the predictions of specific Bayesian models, which would inevitably require a number of *ad hoc* assumptions. A second limitation is that we did not explore the effects of magnitude distribution shape on the central tendency effect.

The elementary psychophysical regularities studied here may have broader implications. Economists have begun to consider the significance of noisy Bayesian cognition for a range of phenomena, including sensitivity to risk and ambiguity, belief updating, and survey responses (Enke and Graeber, 2020, Frydman and Jin, 2019, Gabaix, 2019, Payzan-LeNestour and Woodford, 2020, Woodford, 2019). Models of noisy Bayesian cognition may therefore hold promise in unifying empirical evidence from disparate fields, explaining not only lower-level perceptual processes but also higher-level cognition.

Acknowledgments

We are grateful to Lucy Lai for helpful guidance. This research was supported by the Office of Naval Research (N00014-17-1-2984), the Center for Brains, Minds and Machines (funded by NSF STC award CCF-1231216), and a research fellowship from the Alfred P. Sloan Foundation.

References

- Acerbi, L., Wolpert, D. M., and Vijayakumar, S. (2012). Internal representations of temporal statistics and feedback calibrate motor-sensory interval timing. *PLoS Computational Biology*, 8:e1002771.
- Adler, W. T. and Ma, W. J. (2018). Comparing Bayesian and non-Bayesian accounts of human confidence reports. *PLoS Computational Biology*, 14:e1006572.
- Aitchison, L., Bang, D., Bahrami, B., and Latham, P. E. (2015). Doubly Bayesian analysis of confidence in perceptual decision-making. *PLoS Computational Biology*, 11:e1004519.
- Allred, S. R., Crawford, L. E., Duffy, S., and Smith, J. (2016). Working memory and spatial judgments: Cognitive load increases the central tendency bias. *Psychonomic Bulletin & Review*, 23:1825–1831.
- Ashourian, P. and Loewenstein, Y. (2011). Bayesian inference underlies the contraction bias in delayed comparison tasks. *PloS One*, 6:e19551.

- Berger, J. O. (1985). *Statistical Decision Theory and Bayesian Analysis*. Springer Science & Business Media.
- Cheyette, S. J. and Piantadosi, S. T. (2019). A primarily serial, foveal accumulator underlies approximate numerical estimation. *Proceedings of the National Academy of Sciences*, 116:17729–17734.
- Cheyette, S. J. and Piantadosi, S. T. (2020). A unified account of numerosity perception. *Nature Human Behaviour*, pages 1–8.
- Enke, B. and Graeber, T. (2020). Cognitive uncertainty. Technical report, National Bureau of Economic Research.
- Fechner, G. T. (1860). *Elemente der Psychophysik*. Breitkopf u. Härtel.
- Fleming, S. and Daw, N. (2017). Self-evaluation of decision-making: A general Bayesian framework for metacognitive computation. *Psychological Review*, 124:91–114.
- Frydman, C. and Jin, L. J. (2019). Efficient coding and risky choice. *Available at SSRN 3270773*.
- Gabaix, X. (2019). Behavioral inattention. In *Handbook of Behavioral Economics: Applications and Foundations 1*, volume 2, pages 261–343. Elsevier.
- Gibbon, J. (1977). Scalar expectancy theory and Weber’s law in animal timing. *Psychological Review*, 84:279–325.
- Gibbon, J. and Church, R. (1981). Time left: Linear versus logarithmic subjective time. *Journal of Experimental Psychology: Animal Behavior Processes*, 7:87–108.
- Hollingworth, H. L. (1910). The central tendency of judgment. *The Journal of Philosophy, Psychology and Scientific Methods*, 7:461–469.
- Huttenlocher, J., Hedges, L., and Vevea, J. (2000). Why do categories affect stimulus judgment? *Journal of Experimental Psychology: General*, 129:220–241.
- Inglis, M. and Gilmore, C. (2013). Sampling from the mental number line: How are approximate number system representations formed? *Cognition*, 129:63–69.
- Knill, D. C. and Richards, W. (1996). *Perception as Bayesian inference*. Cambridge University Press.
- Lebreton, M., Abitbol, R., Daunizeau, J., and Pessiglione, M. (2015). Automatic integration of confidence in the brain valuation signal. *Nature neuroscience*, 18(8):1159–1167.
- Li, H.-H. and Ma, W. J. (2020). Confidence reports in decision-making with multiple alternatives violate the Bayesian confidence hypothesis. *Nature Communications*, 11:1–11.
- Meyniel, F., Sigman, M., and Mainen, Z. F. (2015). Confidence as Bayesian probability: From neural origins to behavior. *Neuron*, 88:78–92.
- Nieder, A. and Miller, E. K. (2003). Coding of cognitive magnitude: Compressed scaling of numerical information in the primate prefrontal cortex. *Neuron*, 37:149–157.

- Olkkonen, M., McCarthy, P. F., and Allred, S. R. (2014). The central tendency bias in color perception: Effects of internal and external noise. *Journal of Vision*, 14:5–5.
- Payzan-LeNestour, E. and Woodford, M. (2020). ‘outlier blindness’: Efficient coding generates an inability to represent extreme values. *Available at SSRN 3152166*.
- Petzschner, F. H. and Glasauer, S. (2011). Iterative Bayesian estimation as an explanation for range and regression effects: a study on human path integration. *Journal of Neuroscience*, 31:17220–17229.
- Petzschner, F. H., Glasauer, S., and Stephan, K. E. (2015). A Bayesian perspective on magnitude estimation. *Trends in Cognitive Sciences*, 19:285–293.
- Pouget, A., Drugowitsch, J., and Kepecs, A. (2016). Confidence and certainty: distinct probabilistic quantities for different goals. *Nature Neuroscience*, 19:366.
- Rahnev, D., Desender, K., Lee, A. L., Adler, W. T., Aguilar-Lleyda, D., Akdoğan, B., Arbuzova, P., Atlas, L. Y., Balci, F., Bang, J. W., et al. (2020). The confidence database. *Nature Human Behaviour*, 4:317–325.
- Rahnev, D., Koizumi, A., McCurdy, L. Y., D’Esposito, M., and Lau, H. (2015). Confidence leak in perceptual decision making. *Psychological Science*, 26:1664–1680.
- Roach, N. W., McGraw, P. V., Whitaker, D. J., and Heron, J. (2017). Generalization of prior information for rapid Bayesian time estimation. *Proceedings of the National Academy of Sciences*, 114:412–417.
- Sanders, J. I., Hangya, B., and Kepecs, A. (2016). Signatures of a statistical computation in the human sense of confidence. *Neuron*, 90:499–506.
- Stevens, S. (1961). To honor Fechner and repeal his law. *Science*, 133:80–86.
- Stevens, S. and Greenbaum, H. B. (1966). Regression effect in psychophysical judgment. *Perception & Psychophysics*, 1:439–446.
- Treisman, M. (1964). Noise and Weber’s law: The discrimination of brightness and other dimensions. *Psychological Review*, 71:314–330.
- Woodford, M. (2019). Modeling imprecision in perception, valuation, and choice. *Annual Review of Economics*, 12.