The mean-variance signature of Bayesian probability judgment

Joakim Sundh¹, Jian-Qiao Zhu¹, Nick Chater², & Adam Sanborn¹

¹Department of Psychology, University of Warwick

²Warwick Business School, University of Warwick

Author note:

Sundh, Zhu, and Sanborn were supported by a European Research Council consolidator grant [817492-SAMPLING]. Chater was supported by the ESRC Network for Integrated Behavioural Science [ES/P008976/1]. Correspondence should be sent to: Joakim Sundh, Department of Psychology, University of Warwick, Coventry, CV4 7AL, United Kingdom. Email: joakim.j.sundh@warwick.ac.uk

Abstract

Human probability judgments are variable and subject to systematic biases. Sampling-based accounts of probability judgment have successfully explained such idiosyncrasies by assuming that people remember or simulate instances of events and base their judgments on sampled frequencies. Biases have been explained either by noise corrupting sample accumulation (the Probability Theory + Noise account), or as a Bayesian adjustment to the uncertainty implicit in small samples (the Bayesian sampler). While these two accounts closely mimic one another, here we show that they can be distinguished by a novel linear regression method that relates the variance of repeated judgments to their means. First, the efficacy of the method is confirmed by model recovery, and it more accurately recovers parameters than computationally complex methods. Second, the method is applied to both existing and new probability judgment data, which confirm that judgments are based on a small number of samples that are adjusted by a prior, as predicted by the Bayesian sampler.

Keywords: sampling, probability, biases, Bayes, noise

Introduction

During the last few decades, sampling-based models have emerged as one of the most promising accounts of human probability judgment. This perspective is based on the assumption that, when judging the probability of an event, the brain samples a number of instances from some internal representation, such as drawing them from long-term memory or performing mental simulation, and then bases the judgment on the frequencies in this sample. For example, when looking out the window and judging the probability of rain, we might sample a number of similar days and base our judgments on how many of those days it actually did rain. Sampling-based models have been very successful modeling a range of human behaviors and naturally account for the stochasticity of human judgment and decision making (Bonawitz, Denison, Griffiths, & Gopnik, 2014; Griffiths, Vul, & Sanborn, 2012; Juslin, Winman, & Hansson, 2007; Sanborn, & Chater, 2016).

Apart from varying stochastically, human probability judgment is subject to a number of biases, and these cannot be explained by sampling alone, which is of course unbiased. The most fundamental of these biases is conservatism: that people's probability judgments are usually less extreme than what would be expected (Costello & Watts, 2014; Erev, Wallsten, & Budescu, 1994; Hilbert, 2012; Kaufman, Lord, Reese, & Volkmann, 1949).¹ Conservatism can, in turn, partly explain other biases, such as the conjunction and disjunction fallacies. A conjunction fallacy occurs when a conjunction between events is judged as more likely than either of the marginal events (such as judging the probability of a person being a bank teller and a feminist as more likely than being a bank teller) and, conversely, a disjunction fallacy occurs when a disjunction between events is judged as less likely than either of the marginal events (such as judging the probability of a person being a bank teller or a feminist as less

¹ Note that we focus specifically on conservatism in probability judgments, rather than conservatism in updates on probability estimates in light of new evidence (e.g., Peterson & Beach, 1967).

likely than being a feminist; Bar-Hillel & Neter, 1993; Costello, 2009; Moro, 2009; Tversky & Kahneman, 1983). If one presumes that the conservatism bias is stronger for conjunctions and disjunctions than for single events, then conjunction and disjunction fallacies can occur (Costello & Watts, 2016; Zhu, Sanborn, & Chater, 2020).

From the point of view of sampling accounts, where does conservatism come from? Two recent accounts make almost indistinguishable predictions but embody very different viewpoints on conservatism. One influential approach, the Probability Theory plus Noise model (PT+N; Costello & Watts, 2014; 2016), suggests that conservatism arises because the process of retrieving memories is noisy. So even if a person recalls, say, five days all of which were rainy, there is a good chance they will erroneously recall one (or more) as dry, pulling their probability away from 1; and similarly, noise will tend to push up probabilities of low probability events. This account is similar to the noisy memory channels suggested by Hilbert (2012), but while Hilbert's theory includes various types of noise associated with various types of memory, PT+N focuses specifically on noisy retrieval of binary outcomes. More recently, Zhu et al. (2020) proposed a model in which conservatism arises not through noise, but as a result of Bayesian inference with small samples. Remarkably, the mean predictions of these models turn out to be identical. However, we will show that the two approaches differ regarding the relationship between the means and variances in probability judgments. Indeed, Bayesian models have a distinctive "signature" relationship between mean and variance, which is, we will see, empirically observed.

Sampling-based models of human probability judgment

We first outline the two classes of theory in more detail. The PT+N model is based on the idea that for each sampled instance there is a certain probability (denoted d) that the outcome will be misread, so that an occurrence is read as a non-occurrence or vice versa. We will refer to this error as *retrieval noise*, to distinguish it from the *sampling noise* inherent in

all sampling-based models with finite sample size. This is an important distinction because retrieval noise causes biases while sampling noise does not. For an event A , the probability of reading an outcome as A according to the PT+N model is the weighted mean of the probabilities $P(A)$ and $P(not-A)$ weighted by d . Because the judgment of the PT+N model is based directly on the proportion of outcomes read as A , the average estimate is

$$E(\hat{P}(A)) = (1 - 2d)P(A) + d. \quad 1$$

This means that whenever $d > 0$ the judgment will be regressed towards the middle of the scale, and also that the value of the d parameter will be directly related to the amount of bias that people show in their judgments (see Figure 1).

The Bayesian sampler, by contrast, does not assume that sampling is perturbed by retrieval noise, but rather that each judgment is adjusted by a prior after sampling is complete and frequencies are tallied. The function of this prior is to compensate for the inherent uncertainty in sampling, by weighting the sampled proportion according to a distribution that is (presumably) representative of one's previous experiences and general knowledge. In the Bayesian sampler the prior is represented by a symmetric Beta distribution, which is commonly used as a prior over probabilities. The symmetric Beta distribution has a single parameter, b , that determines its shape: for $b > 1$ it has a single peak at 0.5 and the probability is lowest at 0 and 1, for $b = 1$ it is a uniform distribution across the entire range, while for $b < 1$ it is a u-shape that peaks at 0 and 1 and the probability is lowest at 0.5.² If only a small number of sampled instances is used then such an adjustment will decrease average judgment error. For example, if one draws a red ball from an urn with an unknown proportion of red and blue balls, then it would be foolish to assume that the urn contained only red balls on that evidence alone; given no information on the proportion of red and blue

² Note that in many cases, including the article by Zhu et al. (2020), the parameter of the Beta distribution is represented by the letter β (or α and β) rather than b . We chose to change the notation to b in order to avoid confusion with the regression weight parameter, which is also traditionally denoted β .

balls (i.e., a uniform prior with $b = 1$) the estimate that will minimize average squared error is .67.

Notably, the Bayesian adjustment is a simple linear function of the sample size N and the prior parameter b . It also can be expressed in terms of the same parameter d as in the PT+N model and using the same equation (Costello & Watts, 2019; Zhu et al., 2020). In this case the d parameter does not represent retrieval noise but rather the influence of the symmetric prior as expressed by a Beta distribution, mitigated by the sample size. The average estimate of the Bayesian sampler is therefore identical to the average estimate of the PT+N model (see Equation 1), given the bridge conditions $d = \frac{b}{N+2b}$ and $(1 - 2d) = \frac{N}{N+2b}$. This implies that the bias expressed by the Bayesian sampler is adaptive in the sense that d will decrease as the number of sampled instances increases and approach zero as sample size approaches infinity (see Figure 1).

These rather different interpretations of the source of bias in human probability judgment are difficult to disentangle. First, the predictions of mean probability judgments from the two models are identical for simple probabilities (e.g., the probability that a day will be rainy) and for conjunctions of two events (e.g., the probability that a day will be rainy and cold). The only judgments for which mean estimates appear different for the two models are for conditional probabilities (e.g., the probability that it will rain on a cold day). For these judgments, Zhu et al. (2020) found that the Bayesian sampler model outperformed the current version of the PT+N for conditional probabilities (Costello & Watts, 2016). But, as Zhu et al. (2020) note, the predictions of a slightly modified variant of the PT+N model can precisely match the predictions of the Bayesian sampler for conditional, as well as unconditional, probabilities. Thus, mean judgments alone do not provide the power to distinguish between retrieval noise and Bayesian adjustment due to using a prior over probabilities. In addition,

using mean judgments alone only allows the overall bias d to be estimated, while the sample size N and prior parameter b are not identifiable.

A natural response to this state of affairs is to look at the individual judgments, because retrieval noise and the Bayesian sampler, as we discuss below, predict different distributions of responses. This is commonly done in a likelihood framework by calculating a model selection measure such as Bayes factors, AIC, BIC, etc. based on the likelihood of the model producing the observed data. Unfortunately for these two models this was not wise: each generally predicts that people will *never* use portions of the response scale. People do, at least on occasion, use any of the response values, and so the identity of the winning model will depend very heavily on the auxiliary assumptions about participant “response noise,” in particular the exact nature of the keypress and mouse movement errors that they make, rather than the theoretically interesting assumptions of each model (Acerbi, Ma, & Vijayakumar, 2014).

In order to avoid these issues, Zhu et al. (2020) evaluated the fit of the models to individual judgments using a Wasserstein distance measure instead of likelihood. Wasserstein distance, often called earth mover’s distance, is the minimum amount of “effort” needed to transform one discrete distribution into another. This fit measure allowed for an evaluation of the goodness of fit of the discrete distributions predicted by the response-noise-free PT+N and Bayesian sampler models to the discrete distributions of responses made by participants. It also allowed the sample size N and prior parameter b to be estimated for the Bayesian sampler. While this measure has the potential to distinguish the models in theory, Zhu et al. (2020) found that it did not prove diagnostic in practice: the fits were somewhat ambiguous, and the parameter estimates appeared implausibly biased towards very large sample sizes for both models.

In this paper, we take a different approach. We show that it is possible to distinguish generally between retrieval noise and adjustment based on a prior by modeling the relationship between the mean and the variance of the judgments using linear regression. Recall that according to the PT+N model retrieval noise is an inherent part of the sampling process. By contrast, the adjustment described by the Bayesian sampler takes place after sampling is completed and proportions are tallied. It turns out that this distinction has different implications for the relationship between the mean and the variance of the probability judgments. And, in addition, linear regression can be used to extract parameter estimates of the sample size N and, for the Bayesian sampler, the parameter b of the symmetric prior distribution $\text{Beta}(b, b)$.

The mean/variance relationship as a linear regression

In sample-based models of human probability judgment, samples are distributed according to a binomial distribution, meaning that the variance of the estimate $\hat{P}(A)$ is dependent on the true probability³ $P(A)$ as well as the sample size N , according to the equation

$$V(\hat{P}(A)) = \frac{1}{N} P(A)(1 - P(A)). \quad 2$$

This formula illustrates that variance of repeated judgments will be lower the larger the sample size N and be higher in the middle of the probability scale than at the edges of the scale. Although we typically do not have access to the true probabilities for subjective probability estimates, for repeated judgments we can generally use the mean of the estimates as a proxy. Thus, if we plot the variance of the estimates against the mean of the same estimates, we will observe an inverse U-shaped curve, the height of which will vary according to sample size (see Figure 1).

³ Note that “true probability” here refers to distributions in peoples’ minds, which might or might not correspond to the statistical properties of their environment.

Crucially, when we plot the variance of the estimates against the mean of the estimates, the retrieval noise of the PT+N model will not affect the shape of this curve, while the adjustment of the Bayesian sampler implies that the endpoints of the curve will move away from the extremes and towards the middle of the probability scale (see Figure 1). The reason for this might not be immediately obvious, but it follows logically from the fact that retrieval noise will only affect the variance in so far as it affects the probability of sampling an outcome; in other words, retrieval noise will change the value of $V(\hat{P}(A))$ only because it changes the value of $P(A)$, meaning that the relationship described in Equation 2 is retained. A Bayesian adjustment made after sampling, on the other hand, will restrict the width of possible estimates, in effect adding an additional (negative) term to the equation.

To get some intuition of why this is, imagine an event Q with probability $P(Q) = 0$. When the PT+N model estimates $P(Q)$ it will potentially, depending on the value of d , mistakenly read a sample as Q rather than *not-Q*. This in turn implies that the probability of sampling Q is above zero, and therefore the variance of the judgments will also be above zero. In other words, the variance of the judgments has increased, but so has the mean, and it has moved along the same curve. The Bayesian sampler, by contrast, will sample exactly zero occurrences of Q when $P(Q) = 0$, but will adjust the judgment to a value above zero, because it will compensate for small samples. In this case, because the probability of sampling Q is still zero, the variance of the judgment is also equal to zero, so the mean changes while the variance does not. In effect, the endpoints of the curve when using the Bayesian sampler will be adjusted in towards the middle (see Figure 1).

Figure 1. Upper panels show the underlying subjective probability (x-axis) and the corresponding predicted model estimation (y-axis) for different levels of bias. Lower panels show the mean model estimate (x-axis) and the corresponding variance for the same levels of bias. As demonstrated, the same levels of bias are associated with different relationships between mean estimate and variance for the PT+N model (left side panels) and the Bayesian sampler (right side panels).

Putting the above considerations together, we can let $X = E(\hat{P}(A))(1 - E(\hat{P}(A)))$, β

$= \frac{1}{N}$, and $\alpha = \frac{1}{N}d(d - 1)$ (see Appendix A for the mathematical details), which means that

we can express the variance of the estimates of the Bayesian sampler in terms of linear regression

$$V(\hat{P}(A)) = \alpha + \beta X. \quad 3$$

This equation can be fitted to data with repeated judgments of different events, regardless of the relationship between events and even in ignorance of the true probabilities underlying the sampling process. This method has some notable advantages over a related method presented

in Howe and Costello (2020) that fit a relationship between the true objective probabilities and the variance of estimates. Firstly, the mean-variance curve we describe is invariant to retrieval noise while the relationship between true probabilities and variance is not, so our method provides a better test of the use of a prior. Secondly, as our method does not require knowing participants' subjective probabilities it is applicable to a much wider range of experimental designs. Thirdly, because we fit parameters with linear regression, we can take advantage of its methodological advances (e.g., pooling data across participants using random effects) as well as its fast and easy-to-use implementations (e.g., we illustrate how the parameters can be fit using Microsoft Excel in the supplemental information).

The parameter estimates of the intercept α and the regression weight β can be used to in turn extract the parameters N , d , and b , by using the following equations:

$$N = \frac{1}{\beta} \quad 4$$

$$d = \frac{1 - \sqrt{4N\alpha + 1}}{2} \quad 5$$

$$b = \frac{Nd}{1 - 2d} \quad 6$$

Because this linear regression is invariant to retrieval noise (as explained above), the d parameter only pertains to the bias introduced by the adjustment of the Bayesian sampler, and therefore the intercept can be used to measure the contribution of this adjustment. In simple terms, the presence of a negative intercept indicates the presence of a Bayesian adjustment,

and the value of the intercept can be used to evaluate the size of the bias d and, in turn, the shape of the prior $\text{Beta}(b, b)$.⁴

Importantly, other types of error or adjustments, such as retrieval error, typing errors, or misunderstanding the question, typically either increase total variance or do not affect total variance at all. We have already demonstrated that the linear regression described is invariant to retrieval error and it is easy to see that this holds for any type of noise that directly affect the probability of sampling an outcome, since changes in the value of variable $P(A)$ will not affect the relationship described in Equation 2. Typing errors or misunderstandings, on the other hand, both imply additional variability since they each represent a probabilistic change to the judgment; if we conceptualize typing errors and misunderstandings as independent random variables in their own right then they will each add to the total variance, since the variance of a sum of two independent random variables equals the sum of the variance. By contrast, neither the PT+N model or the Bayesian sampler implies any additional random variables; the retrieval error of the PT+N model is subsumed by the sampling error and the Bayesian sampler does not include any random components beyond the sampling error. Therefore, a negative intercept is a strong indication of a Bayesian adjustment and (to our knowledge) cannot be explained by other factors.

Model Recovery

As described above, the linear regression can be used to estimate the sample size N , as well as identify the presence of post-sampling adjustment and estimate the prior $\text{Beta}(b, b)$ according to which judgments are adjusted. In previous research Zhu et al. (2020) estimated N and $\text{Beta}(b, b)$ parameters using Wasserstein distance, but because overall fit for all models was better the larger the sample size, there is reason to suspect that this method results in a

⁴ It is important to note that only a strictly negative intercept implies a Bayesian adjustment according to a prior, because when $N > 0$ and $0 > d > .5$ then $\alpha = \frac{1}{N}d(d - 1) < 0$, meaning that Equations 5 and 6 are no longer applicable for a positive intercept.

positive bias. To compare the accuracy of parameter estimates from the linear regression with estimates from Wasserstein distance we performed model recovery using data with the same design as in Experiment 1 in Zhu et al. (2020). We created a total of 100 simulated participants, each with a set of true probabilities drawn from the uniform distribution $U(0, 1)$, a sample size randomly drawn from the values [1 2 3 4 5 6 7 8 9 10 15 20 50 100 250], and a Bayesian adjustment according to a prior that was randomly drawn from the values [0.2 0.25 0.33 0.5 1 2 3 4 5] (see Appendix B for details). Results indicated that estimates of N and b from using the linear regression indeed had both less total error and less bias than estimates from using Wasserstein distance, but that the estimate of d was comparable (see Table 1).⁵ We can conclude that Wasserstein distance is suitable for estimating the total amount of bias, but less so for estimates of sample size and prior.

Table 1. *Mean Error and Mean Absolute Error for Recovered Parameter Estimates.*

Parameter		Linear regression	Wasserstein distance
N	Mean error	.794	20.2
	Mean abs. error	7.11	20.3
b	Mean error	-.115	.709
	Mean abs. error	.729	.825
d	Mean error	-.007	-.011
	Mean abs. error	.017	.015

Because one of the most important aspects of the linear regression is its ability to distinguish Bayesian adjustment by the presence or absence of a negative intercept, we also

⁵ When recovering parameters, bounds are imposed so that $\alpha \leq 0$ and $\beta \geq 0$, in order to conform to the same conditions as for the Wasserstein modelling in Zhu et al. (2020). Without these bounds the results are broadly the same but there are occasional additional outliers (see Appendix B). These bounds are not used when recovering the number of statistically significant negative intercepts.

created 100 simulated datasets each with 80 simulated participants (the average number of participants in our empirical datasets) generated in the same manner as above and applied a mixed-effects model to each dataset with random slopes and intercepts for each participant. For each simulated dataset, we found a statistically significant (negative) intercept at $p < .005$. This procedure was then repeated but with simulated participants without a Bayesian adjustment, which resulted in no statistically significant intercepts for any of the simulated datasets. While for individual participants negative intercepts close to zero may be difficult to detect (see Appendix B), we can conclude that, on an aggregate level, the linear regression will reliably detect a Bayesian adjustment if it is present.

Empirical Results

Having validated our method with simulated data, we now use it to distinguish between the Bayesian sampler and PT+N models on empirical data. Specifically, we applied the linear regression described above to data from two previous experiments (Zhu et al., 2020)⁶ and two new datasets with the same experimental design. In Experiment 1-3 participants made probability judgments on the format: “What is the probability that the weather is [X] on a random day in England?” Various weather events were used, and the queries included both marginal events, conditional events, conjunctions, and disjunctions. In Experiment 4 participants instead made probability judgments on future events, such as: “What is the probability that there will be an early UK general election?” Again, the queries included both marginal events, conditional events, conjunctions, and disjunctions, using the same structure as in Experiment 1-3 (see Appendix C for a description of the experimental design). Crucially, for all experiments, each individual query was repeated three times, which provided a mean and a variance for each participant’s judgment of each query.

⁶ Note that in Zhu et al. (2020), the event pair *{warm, snowy}* was excluded from Exp. 2, due to the fact that such highly dependent events can induce schematic reasoning if the participant consider, for example, the conjunction *warm and snowy* to be a logical contradiction. We found the results of the linear regression to be very similar in both cases and therefore chose to include also *{warm, snowy}* in the current analysis.

The linear regression described above was applied to each dataset as a mixed-effects model with participant as a random effect, in order to explore overall tendencies while also allowing for between-subjects variability.⁷ The best fitting parameter values are summarized in Table 2. The results have three important implications. Firstly, we estimate that only a small number of instances were sampled for each judgment, as few as approximately three in most cases. This is in contrast to the notably higher medians of N (median $N_{Exp1} = 10$, median $N_{Exp2} = 48$) found by Zhu et al. (2020) using Wasserstein distance and, following the model recovery, we consider the results of the linear regression to be more reliable. Smaller samples are arguably more psychologically plausible as well; previous research has indicated that people only generate a small number of thoughts when coming to a decision (e.g., 3.6 on average; Weber, Johnson, Milch, Chang, Brodscholl, & Goldstein, 2007) and that only very small samples are optimal when sampling takes time or effort (Vul, Goodman, Griffiths, & Tenenbaum, 2014).

Secondly, the linear regression confirms that judgments were indeed adjusted after the sampling process, as predicted by the Bayesian sampler. All four experiments showed a negative intercept ($p_{Exp1-3} < .001$, $p_{Exp4} = .005$). While the results in Zhu et al. (2020) did favor the Bayesian sampler overall, the fits were sometimes ambiguous, and the distinction relied primarily on the conditional probabilities. Conversely, by focusing the relationship between mean and variance, the linear regression clearly indicates that, on aggregate, a Bayesian adjustment did take place. This does not necessarily preclude other sources of bias, such as retrieval noise, but it does strongly indicate that the Bayesian sampler is indeed a good description of the process that takes place.

⁷ The random effect was applied to both the intercept and the weight, to account for individual variation in both sample size and prior.

Thirdly, the adjustments suggest a U-shaped prior, with higher priors associated with extreme probabilities. Again, this contrast with the findings by Zhu et al. (2020) and, again, we consider the results of the linear regression more reliable, since model recovery indicates that the estimates of the Wasserstein distance are positively biased. Interestingly, the estimated values are remarkably similar to the empirical prior that Zhu et al. (2020) calculated from the statistics of probability words in natural language observed by Stewart, Chater, and Brown (2006). In the latter study, the frequencies of when a range of probability related words and phrases were used to describe a probability were collected, and participants were then asked to judge the probability they associated with each phrase. When fitting a symmetric Beta distribution to the resulting distribution of values, the best fitting distribution was Beta(.27, .27), which is very similar to the results of the linear regression.

Table 2. *Best Mixed-Effects Model Parameter Estimates for Exp. 1-4.*

Experiment	Parameter		
	<i>N</i>	<i>b</i>	<i>d</i>
Exp. 1*	2.68	.239	.076
Exp. 2*	6.66	.347	.047
Exp. 3**	3.09	.293	.080
Exp. 4**	6.43	.290	.041

*From Zhu, Sanborn & Chater, 2020; **New dataset

Examining the random effects of the mixed-effects model also indicates very strong negative correlations between intercepts and regression weights, $r_{Exp1} = -.811$, 95% CI [-.884, -.700], $r_{Exp2} = -1$, 95% CI [-1, -1], $r_{Exp3} = -.966$, 95% CI [-.979, -.946], $r_{Exp4} = -1$, 95% CI [-1, -1] which in turn implies that bias is indeed adaptive (i.e., we observe more bias for smaller

samples). Because intercepts and the regression weights are statistically related however,⁸ it is more prudent to examine the extracted parameters.

Figure 2. Histograms of the distributions of fitted sizes of mental samples in all experiments.

To explore parameter correlations, we applied the linear regression to each participant individually, in order to avoid the random effects shrinking the individual estimates towards the mean. Results confirm that sample sizes are centered around relatively small values (median $N_{Exp1} = 2.65$, median $N_{Exp2} = 8.25$, median $N_{Exp3} = 3.00$; median $N_{Exp4} = 5.91$ see

⁸ In a linear regression of the form $y = \alpha + \beta x$, intercept and slope are negatively correlated as long as $x > 0$. Simulations indicate that randomly generated $N = \text{unif}(2, 10)$ and $b = \text{unif}(.25, 1.25)$ will create a correlation of only $r = -.164$ however, so the observed correlation is still notably higher than what would be expected for chance alone.

Figure 2). If bias, as measured by the parameter d , is adaptive, then we should find a negative correlation between N and d , and a positive correlation between d and b , because d is defined as $d = \frac{b}{N+2b}$. Conversely, we do not expect a correlation between N and b , because we have no a priori reason to expect sample size and the prior to be related. These predictions were all confirmed (see Table 3). While only a limited proportion of individual participants in the datasets had statistically significant intercepts (Exp. 1: 20/59, Exp. 2: 4/84, Exp. 3: 26/69 Exp 4: 24/111), this is to be expected with the current experimental design and parameter values (as demonstrated by model recovery in Appendix B).

Table 3. *Correlation Coefficients Between Individual Parameters Estimates for Exp. 1-4.*

Experiment	Correlation		
	N vs. d	N vs. b	d vs. b
Exp. 1*	$r = -.440$ [-.625, -.207]	$r = -.198$ [-.432, .061]	$r = .783$ [.659, .866]
Exp. 2*	$r = -.535$ [-.673, -.363]	$r = -.177$ [-.377, .040]	$r = .742$ [.628, .825]
Exp. 3**	$r = -.396$ [-.579, -.176]	$r = -.014$ [-.250, .223]	$r = .794$ [.686, .868]
Exp. 4**	$r = -.391$ [-.538, -.221]	$r = -.003$ [-.189, .184]	$r = .611$ [.480, .716]

*From Zhu, Sanborn & Chater, 2020; **New dataset

Discussion

In this article we demonstrate how to use the relationship between the mean and variance of repeated judgments in a linear regression to determine whether the conservatism bias in human probability judgment is due to a Bayesian adjustment after sampling or arises from a Probability Theory + Noise account. We confirm through model recovery that this method will recover the relevant parameters. We apply the method to data from two

published experiments and two new experiments, where we find that parameters are consistent with small samples and an adjustment after sampling, as predicted by the Bayesian sampler. It is worth noting that, although previous research has demonstrated that the Bayesian sampler fits data well (Zhu et al., 2020), this is the first direct demonstration of the presence of such an adjustment.

The parameters extracted from all three experiments are consistent with previous empirical data, in two important aspects: Firstly, the estimated number of samples are generally small. This is consistent with research on the limitations of working memory (Cowan, 2001; Miller, 1956; Weber et al., 2007) and the computational restraints of the brain (Gershman, Horvitz, & Tenenbaum, 2015; Griffiths, Lieder, & Goodman, 2015), as well as results indicating that only a very small number of samples is optimal when sampling takes time or effort (Vul et al. 2014). Additionally, it has been shown that some cognitive biases can be explained as a direct consequence of a small number of samples (Juslin et al., 2007).

Secondly, the parameters are consistent with a U-shaped prior, implying that most events are either very likely or very unlikely and that intermediate probabilities are more rarely encountered. This matches the distribution of probability words in natural language (Stewart et al., 2006), which tentatively suggests that the prior is based on the frequencies of probability concepts in day-to-day life. If this is indeed representative of general tendencies in people, then this might have interesting implications regarding how our usage of probability in everyday language affects the way we approach uncertainty, and vice versa.

Successful models will naturally explain data well, meaning that distinguishing two successful models only by their relative fit to data will be difficult. This is particularly the case with the PT+N model and the Bayesian sampler, since they make the exact same average predictions in many situations. The linear regression presented here is an alternative way of distinguishing between the models, by taking a conceptual difference between the two (i.e.,

the source of the conservatism effect) and exploring its effect on the variance of the data. Results indicate that the Bayesian sampler is consistent with data, insofar that we can confirm that estimates does indeed appear to be adjusted after sampling. Such an adjustment is necessary, as well as intuitively compelling, whenever sample sizes are small, meaning that the resulting bias is “adaptive” in the sense that it generally increases the accuracy of probability judgments at the cost of introducing inconsistencies such as conjunction fallacies.

Although the linear regression validates the Bayesian sampler, responses could potentially be affected by other factors as well, such as reading errors. In fact, the mechanisms of the PT+N model are not explicitly contradicted by the Bayesian sampler, and it is possible that both processes occur simultaneously, in which case a hybrid model might be called for. Because the mean-variance relationship we base our method on is invariant to retrieval noise, our results do not exclude this possibility. It should be noted however, that previous work on the PT+N model generally assumed a relatively low value of approximately $d = 0.1$ (Howe & Costello 2020; Costello & Watts, 2017), which might allow little room for retrieval noise after the Bayesian adjustment is accounted for. This also raises further questions regarding the source and function of the proposed retrieval noise and its relationship to the Bayesian interpretation. From the perspective of the PT+N account, although it is reasonable to assume that recall from long term memory is less than perfect, it appears odd that the same limitations should apply to mental simulation, as has been previously suggested in order to account for unique events (e.g., Costello & Watts, 2018; Ludwin-Peery, Bramley, Davis, & Gureckis, 2020). From the perspective of the Bayesian sampler, on the other hand, it seems reasonable to assume that, if the Bayesian sampler compensates for small samples, then it should also adjust for errors in sampling, such as retrieval noise. A hybrid model of human probability judgment would need to take these points into account.

Although we here focus on the source of conservatism in human probability judgment and the distinction between two successful models in the area, we believe that the method we used has promise for wider application, particularly for estimating sample sizes for sampling-based accounts of judgments. Sampling-based models have generally been relatively non-committal regarding the number of sampled instances for each judgment, but there are indications that the human mind is limited to processing only a small number of samples (Cowan, 2001; Miller, 1956; Weber et al., 2007). A wider application of the linear regression presented here, or variants based on the same general idea, could go a long way to confirm or reject such claims.

References

- Acerbi, L., Ma, W. J., & Vijayakumar, S. (2014, January). A Framework for Testing Identifiability of Bayesian Models of Perception. In *NIPS* (pp. 1026-1034).
- Bar-Hillel, M., & Neter, E. (1993). How alike is it versus how likely is it: A disjunction fallacy in probability judgments. *Journal of Personality and Social Psychology*, 65(6), 1119.
- Bonawitz, E., Denison, S., Griffiths, T. L., & Gopnik, A. (2014). Probabilistic models, learning algorithms, and response variability: sampling in cognitive development. *Trends in Cognitive Sciences*, 18(10), 497-500.
- Costello, F. (2009). Fallacies in probability judgments for conjunctions and disjunctions of everyday events. *Journal of Behavioral Decision Making*, 22(3), 235-251.
- Costello, F., & Watts, P. (2014). Surprisingly rational: probability theory plus noise explains biases in judgment. *Psychological Review*, 121(3), 463.
- Costello, F., & Watts, P. (2016). People's conditional probability judgments follow probability theory (plus noise). *Cognitive Psychology*, 89, 106-133.
- Costello, F., & Watts, P. (2017). Explaining high conjunction fallacy rates: The probability theory plus noise account. *Journal of Behavioral Decision Making*, 30(2), 304-321.
- Costello, F., & Watts, P. (2018). Probability theory plus noise: Descriptive estimation and inferential judgment. *Topics in Cognitive Science*, 10(1), 192-208.
- Costello, F., & Watts, P. (2019). The rationality of illusory correlation. *Psychological Review*, 126(3), 437.
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24(1), 87-114.
- Erev, I., Wallsten, T. S., & Budescu, D. V. (1994). Simultaneous over-and underconfidence: The role of error in judgment processes. *Psychological Review*, 101(3), 519.
- Gershman, S. J., Horvitz, E. J., & Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, 349(6245), 273-278.
- 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.
- Griffiths, T. L., Vul, E., & Sanborn, A. N. (2012). Bridging levels of analysis for probabilistic models of cognition. *Current Directions in Psychological Science*, 21(4), 263-268.
- Hilbert, M. (2012). Toward a synthesis of cognitive biases: how noisy information processing can bias human decision making. *Psychological Bulletin*, 138(2), 211.

- Howe, R., & Costello, F. (2020). Random variation and systematic biases in probability estimation. *Cognitive Psychology*, 123, 101306.
- Juslin, P., Winman, A., & Hansson, P. (2007). The naïve intuitive statistician: A naïve sampling model of intuitive confidence intervals. *Psychological Review*, 114(3), 678.
- Kaufman, E. L., Lord, M. W., Reese, T. W., & Volkmann, J. (1949). The discrimination of visual number. *The American Journal of Psychology*, 62(4), 498-525.
- Ludwin-Peery, E., Bramley, N. R., Davis, E., & Gureckis, T. M. (2020). Broken physics: A conjunction-fallacy effect in intuitive physical reasoning. *Psychological Science*, 31(12), 1602-1611.
- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63(2), 81.
- Moro, R. (2009). On the nature of the conjunction fallacy. *Synthese*, 171(1), 1-24.
- Peterson, C. R., & Beach, L. R. (1967). Man as an intuitive statistician. *Psychological Bulletin*, 68(1), 29.
- Sanborn, A. N., & Chater, N. (2016). Bayesian brains without probabilities. *Trends in Cognitive Sciences*, 20(12), 883-893.
- Stewart, N., Chater, N., & Brown, G. D. (2006). Decision by sampling. *Cognitive Psychology*, 53(1), 1-26.
- Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90(4), 293.
- Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? Optimal decisions from very few samples. *Cognitive Science*, 38(4), 599-637.
- Weber, E. U., Johnson, E. J., Milch, K. F., Chang, H., Brodscholl, J. C., & Goldstein, D. G. (2007). Asymmetric discounting in intertemporal choice: A query-theory account. *Psychological Science*, 18(6), 516-523.
- Zhu, J. Q., Sanborn, A. N., & Chater, N. (2020). The Bayesian sampler: Generic Bayesian inference causes incoherence in human probability judgments. *Psychological Review*.

Appendix A – Mathematical details

The variance of a probability estimate $\hat{P}(A)$, based on the proportion of occurrences S in N mental samples, is

$$\begin{aligned}
 V(\hat{P}(A)) &= V\left(\frac{S}{N}\right) \\
 &= \frac{V(S)}{N^2} \\
 &= \frac{NP(A)(1 - P(A))}{N^2} \\
 &= \frac{P(A)(1 - P(A))}{N}.
 \end{aligned}
 \tag{A1}$$

This can also be written as

$$V(\hat{P}(A)) = \frac{1}{N}P(A)(1 - P(A)), \tag{A2}$$

which, if we let $\beta = \frac{1}{N}$ and $X = P(A)(1 - P(A))$, can be expressed as a linear model,

$$V(\hat{P}(A)) = \beta X. \tag{A3}$$

We generally do not have access to the values of $P(A)$, but both the PT+N model and the Bayesian sampler can be rewritten to use the expected value of the model predictions instead. In this case, although we do not know the “true” value of $P(A)$, we can substitute the mean response for each individual query. For both models $E(\hat{P}(A)) = (1 - 2d)P(A) + d$. In the PT+N model, d represents the probability of misreading a sampled instance, so that a positive outcome is read as a negative outcome or vice versa. In the Bayesian sampler however, d represents an adjustment according to a prior so that $d = \frac{b}{N+2b}$, where the prior distribution is defined as Beta(b, b). Because the PT+N model adjust the probabilities of sampling A before outcomes are tallied, and the Bayesian sampler adjusts the estimate of $P(A)$ after outcomes are tallied, the variance is expressed in different ways.

For the PT+N model,

$$\begin{aligned}
V(\hat{P}(A)) &= V\left(\frac{S}{N}\right) \\
&= \frac{V(S)}{N^2} \\
&= \frac{NP(\text{read as } A)(1 - P(\text{read as } A))}{N^2} \\
&= \frac{P(\text{read as } A)(1 - P(\text{read as } A))}{N} \\
&= \frac{((1 - 2d)P(A) + d)(1 - ((1 - 2d)P(A) + d))}{N}.
\end{aligned}
\tag{A4}$$

By contrast, for the Bayesian sampler,

$$\begin{aligned}
V(\hat{P}(A)) &= V\left((1 - 2d)\frac{S}{N} + d\right) \\
&= V(S)\left(\frac{1 - 2d}{N}\right)^2 \\
&= \frac{NP(A)(1 - P(A))(1 - 2d)^2}{N^2} \\
&= \frac{P(A)(1 - P(A))(1 - 2d)^2}{N}.
\end{aligned}
\tag{A5}$$

First, let us consider the PT+N model. Because d only affects the probability of sampling A (or rather, the probability of reading a sampled outcome as A , but these are equivalent) and because the estimate of the model is equal to the sampled proportion of A , we can assume that $P(\text{read as } A) = E(\hat{P}(A))$. Therefore, we can ignore the last line of Eq. 4 and write,

$$\begin{aligned}
V(\hat{P}(A)) &= \frac{E(\hat{P}(A))(1 - E(\hat{P}(A)))}{N} \\
&= \frac{1}{N}E(\hat{P}(A))(1 - E(\hat{P}(A))),
\end{aligned}
\tag{A6}$$

which gives us the linear model

$$V(\hat{P}(A)) = \beta X, \quad \text{A7}$$

where $X = E(\hat{P}(A)) \left(1 - E(\hat{P}(A))\right)$ and $\beta = \frac{1}{N}$. The parameter d is not part of the equation, because the degree of sampling noise in the PT+N model will affect the probability of sampling A but not the functional relationship of $P(A)$ to $V(\hat{P}(A))$. Therefore, regardless of the value of d , the variance $V(\hat{P}(A))$ will follow the same curve, defined by the parameter β .

The Bayesian sampler, on the other hand, assumes that the estimate is adjusted according to a prior *after* the sampling process is completed and the proportion of occurrences has been tallied. In the case of the PT+N model, the endpoints of the curve are located at $P(A) = 0$ and $P(A) = 1$ respectively, but this is no longer the case for the Bayesian sampler. Instead, the variance of the estimates is

$$\begin{aligned} V(\hat{P}(A)) &= \frac{P(A)(1 - P(A))(1 - 2d)^2}{N} \\ &= \frac{1}{N} P(A)(1 - P(A))(1 - 2d)^2. \end{aligned} \quad \text{A9}$$

We can rewrite the equation in terms of the expected value of the responses, because the expected value of the responses is

$$E(\hat{P}(A)) = (1 - 2d)P(A) + d, \quad \text{A10}$$

which we can rewrite as

$$P(A) = \frac{d - E(\hat{P}(A))}{2d - 1}. \quad \text{A11}$$

Inserting this into the above equation gives us

$$V(\hat{P}(A)) = \frac{1}{N} \left(\frac{d - E(\hat{P}(A))}{2d - 1} \right) \left(1 - \left(\frac{d - E(\hat{P}(A))}{2d - 1} \right) \right) (1 - 2d)^2, \quad \text{A12}$$

which can be simplified to

$$V(\hat{P}(A)) = \frac{1}{N} E(\hat{P}(A)) \left(1 - E(\hat{P}(A))\right) + \frac{1}{N} d(d-1). \quad A13$$

This equation effectively gives us a linear model similar to the one described above, with $X = E(\hat{P}(A)) \left(1 - E(\hat{P}(A))\right)$ and $\beta = \frac{1}{N}$ being exactly the same, plus the addition of the intercept $\alpha = \frac{1}{N} d(d-1)$. This does imply that the value of the intercept α is not wholly independent from the value of β , but because d is also dependent on the prior, α is still free to vary within certain bounds. This means that we can express the variance of the estimates of the Bayesian sampler as the linear model

$$V(\hat{P}(A)) = \alpha + \beta X. \quad A14$$

These two linear models can then be compared to evaluate whether the pattern of variances is more consistent with the PT+N model or the Bayesian sampler.

Appendix B – Model recovery

Simulated data

The simulated data for the model recovery was created to match the empirical data as closely as possible. Therefore, each simulated participant was based on two sets of 20 queries, which was each based on two event pairs. For each event pair, queries included judgments of the marginal events and their negations as well as all possible conjunctions, disjunctions, and conditionals, resulting in 20 queries for each event pair. Each simulated participant was assigned a sample size N randomly drawn from the values [1 2 3 4 5 6 7 8 9 10 15 20 50 100 250]. These values were chosen to represent a range of both small and large samples sizes from previous literature, while also including steps representative of differences in scale (e.g., the step between $N = 5$ and $N = 6$ is of more relative importance than the step between $N = 105$ and $N = 106$). Each simulated participant was also assigned a Bayesian adjustment according to a prior parameter b that was randomly drawn from the values [0.2 0.25 0.33 0.5 1 2 3 4 5], and a set of “true” probabilities for each marginal probability in each event pair drawn from the uniform distribution $U(0, 1)$. The true probabilities for each query were extrapolated from the marginal probabilities; for simplicity we assume independent events, so that $P(A|B) = P(A)$. To create the simulated judgment data, three binomial samples with sample size N was drawn for each query and adjusted according to the Bayesian sampler with a prior $\text{Beta}(b, b)$ (Zhu, et al. 2020), creating a dataset with a total of 120 judgments for each participant.

Parameter error

Mean error and mean absolute error of parameter estimates indicates that the linear regression gives parameter estimates that are both less biased and more accurate than with Wasserstein distance for the N and b parameters, while the estimates of the d parameter are

similarly unbiased and accurate for both methods (see Figure B1 and Table 1 in the main text).

Figure B1. Scatterplots of the true parameter values (x-axis) and the model estimates (y-axis) for the linear regression (upper panels) and the Wasserstein distance (lower panels). Note that, for readability, the plots of the N parameter (leftmost panels) are shown in log scale.

Model recovery without constraining parameter values

As demonstrated in Figure B2 and Table B2, running the linear regression with no bounds will result in a number of additional outliers, including several values that do not have a meaningful interpretation ($N \leq 0$, $b \leq 0$, $d < 0$). These outliers will contribute to larger bias and lower accuracy when compared to Wasserstein distance. Bounded parameters are therefore important in order to ensure reliable estimates of individual parameters.

Figure B2. Scatterplots of the true parameter values (x-axis) and the model estimates (y-axis) for the linear regression (upper panels) and the Wasserstein distance (lower panels). Note that, for readability, the plots of the N parameter (leftmost panels) are shown in log scale, and one extreme outlier with linear regression estimates $N = 412$, $b = -107$, and $d = -.542$ is excluded from all panels.

Table B2. *Mean Error and Mean Absolute Error for Recovered Parameter Estimates with No Bounds.*

Parameter		Linear regression	Wasserstein distance
N	Mean error	6.52	20.2
	Mean abs. error	12.8	20.3
b	Mean error	-1.73	.709
	Mean abs. error	2.35	.825
d	Mean error	-.015	-.011
	Mean abs. error	.025	.015

Intercepts for individual participants

In order to explore the proportion of statistically significant intercepts when analyzing individual participants, we created 100 simulated participants for each of $N = [1\ 2\ 3\ 4\ 5\ 10\ 15\ 20\ 50\ 100\ 250]$ and each of $b = [0.33\ 1\ 3]$ as well as without adjustment according to prior, then calculated the mean value of the intercept and the number of simulated participants with statistically significant intercept on the $p < .05$ level (see Table B3).

Table B3. *Mean Value of Intercepts and Number of Statistically Significant Intercepts Out of 100 Simulated Participants for Different Parameter Values*

N	No adjustment		$b = 0.33$		$b = 1$		$b = 3$	
	mean	stat. sig.	mean	stat. sig.	mean	stat. sig.	mean	stat. sig.
1	-5.29×10^{-17}	91	-0.24	100	-0.332	100	-0.368	100
2	-9.4×10^{-04}	0	-0.063	89	-0.112	99	-0.137	100
3	-8.29×10^{-04}	0	-0.031	50	-0.059	87	-0.082	97
4	-2.28×10^{-04}	0	-0.015	19	-0.039	79	-0.059	89
5	8.22×10^{-04}	0	-0.014	11	-0.024	54	-0.046	92
10	2.64×10^{-04}	0	-0.002	1	-0.008	16	-0.017	58
15	-5.93×10^{-04}	1	-0.001	0	-0.003	10	-0.007	37
20	-7.67×10^{-05}	0	-8.46×10^{-04}	0	-0.003	6	-0.005	25
50	1.62×10^{-05}	0	-3.44×10^{-04}	0	-6.6×10^{-04}	0	-0.001	3
100	-6.15×10^{-05}	0	1.22×10^{-05}	1	-2.08×10^{-04}	0	-2.36×10^{-04}	0
250	7.14×10^{-06}	0	-1.93×10^{-06}	2	-8.37×10^{-06}	0	-5.05×10^{-05}	1

We can see that the linear regression consistently outputs a negative intercept when an adjustment is present, but that for larger samples ($N > 10$) it is generally small enough that it will be difficult to detect with the current sizes of datasets. This is, to some extent, dependent

on the nature of the prior, since a stronger prior implies larger adjustment and a larger (negative) intercept. It should be noted, however, that these results are based on an experimental paradigm mimicking that in Experiment 1 in Zhu et al. (2020), and that a higher-powered experiment would no doubt detect more statistically significant intercepts, though the mean values would presumably be the same.

It is also worth noting that the method detects a surprising number of statistically significant intercepts when no adjustment is made and the sample size is $N = 1$, although in this case the mean value is astronomically small. The reason for this is presumably that data can only take on two different values when the sample size is one, which can create illusory systematicity due to the lack of variance. This appears unlikely to happen in real data, but even if it did, the value of the intercept clearly demonstrate that no adjustment is actually taking place. Additionally, it would be very easy to detect if any participant was actually using only two different values in their responses.

Appendix C – Experimental design

All experiments were based on the basic probability judgment task introduced by Costello and Watts (2014, 2016), where participants were instructed to estimate the probability of a number of events. Responses were typed on a scale from 0 to 100.

Each set of queries was based on a number of event pairs (e.g., {icy, frosty}, see Table C1 for a comprehensive list). For each event pair, queries included the marginal events and their negations (e.g., *icy*, *not icy*, *frosty*, and *not frosty*) as well as all possible conjunctions, disjunctions, and conditionals between the marginal events and their negations (e.g., *icy and frosty*, *icy or frosty*, *icy given frosty*, *icy and not frosty* etc.). For each event pair, this made for a total of 20 unique queries. The total set of unique queries formed a block within which the presentation order was randomized for each participant, and each experiment consisted of three blocks, so that all participants responded to each unique query three times.

In Exp. 1-3, queries concerned the probability of various weather events occurring on a random day in England. Queries on marginal events, conjunctions, and disjunctions were expressed on the format “*What is the probability that the weather will be [some event] on a random day in England?*” and queries on conditional events were expressed on the format “*If the weather in England is [some marginal event] on a random day, what is the probability that weather will also be [another marginal event] on that same day?*”

In Exp. 4, queries concerned the probability of future events (e.g., the probability that Joe Biden will win a second term). Queries on marginal events, conjunctions, and disjunctions were expressed on the format “*What is the probability that [some event]?*” and queries on conditional events were expressed on the format “*If [some marginal event occurs], what is the probability that [another marginal event also occurs]?*” In order to avoid misinterpretation while keeping the wording of the actual queries relatively brief, the marginal events were all fully explained as part of the instructions (e.g., the instructions

specified the marginal event as “that Joe Biden wins a second term as US president in the 2024 US election” while the query was expressed as “that Joe Biden will win a second term”). For conditional events, participants were specifically instructed that it does not matter which event occurs first or if one causes the other or not, instead they should assume that they (somehow) were given certain information about one event but not about the other.

All experiments took approximately 30-40 minutes to complete. Participants for Exp. 1, 2, and 4 were recruited through the University of Warwick Student Research Experience Subject Panel and completed the experiment in exchange for course credit. Participants for Exp. 3 were recruited through the University of Warwick Research Participation System and were compensated with £3. Exp. 1-3 were administered in person at the psychology lab at the Department of Psychology, University of Warwick. Exp. 4 was administered online between 26 April and 7 May 2021, meaning that the events in the queries were still well in the future. See Table C1 for details.

Table C1. *Number of Participants Tested and Event Pairs for All Experiments*

	No. participants	Event pairs
Exp. 1	59	{icy, frosty} {normal, typical}
Exp. 2	84	{cold, rainy} {windy, cloudy} {warm, snowy}
Exp. 3	69	{snowy, stormy} {thundery, humid} {mild, foggy}
Exp. 4	111	{Joe Biden winning a second term, the world reaching the 2050 climate goals} {early UK election, the UK economy recovering within current year}