

SCIENTIFIC REPORTS



OPEN

A nonlinear updating algorithm captures suboptimal inference in the presence of signal-dependent noise

Seth W. Egger ^{1,2} & Mehrdad Jazayeri ^{1,2}

Bayesian models have advanced the idea that humans combine prior beliefs and sensory observations to optimize behavior. How the brain implements Bayes-optimal inference, however, remains poorly understood. Simple behavioral tasks suggest that the brain can flexibly represent probability distributions. An alternative view is that the brain relies on simple algorithms that can implement Bayes-optimal behavior only when the computational demands are low. To distinguish between these alternatives, we devised a task in which Bayes-optimal performance could not be matched by simple algorithms. We asked subjects to estimate and reproduce a time interval by combining prior information with one or two sequential measurements. In the domain of time, measurement noise increases with duration. This property takes the integration of multiple measurements beyond the reach of simple algorithms. We found that subjects were able to update their estimates using the second measurement but their performance was suboptimal, suggesting that they were unable to update full probability distributions. Instead, subjects' behavior was consistent with an algorithm that predicts upcoming sensory signals, and applies a nonlinear function to errors in prediction to update estimates. These results indicate that the inference strategies employed by humans may deviate from Bayes-optimal integration when the computational demands are high.

Sensorimotor control depends on accurate estimation of internal state variables^{1–5}. Numerous experiments have used Bayesian estimation theory to demonstrate that humans estimate internal states by integrating multiple sources of information including prior beliefs and sensory cues from various modalities^{6–16}. Bayesian estimation is typically formulated in terms of three components: prior distributions representing *a priori* beliefs about state variables, likelihood functions derived from noisy sensory measurements, and cost functions that characterize reward contingencies¹⁷. In this formulation, the likelihood function and prior distribution are combined to compute a posterior distribution and the cost function is used to extract an estimate that maximizes expected reward. This formulation is the basis of most psychophysical studies of Bayesian integration^{9–15,18–20}.

Implicit in this formulation is the assumption that the brain has access to priors, likelihoods, and cost functions. Access to these quantities is appealing as it could support rapid and optimal state estimation without the need to learn new policies for novel behavioral contexts^{21,22}. However, in most experiments, Bayes-optimal behavior can also be achieved by simpler algorithms that do not depend on direct access to likelihoods, priors and cost functions^{21–23}. For example, optimal cue combination in the presence of Gaussian noise may be implemented by a weighted sum of measurements⁶. Similarly, integration of noisy evidence with prior beliefs may be implemented by a suitable functional mapping between measurements and estimates^{24,25}. Finally, online estimation of a variable from sequential measurements that are subject to Gaussian noise can be achieved by a *Kalman filter* that only keeps track of the mean and variance²⁶ without representing and updating the full posterior distribution.

In contrast to simple laboratory tasks, optimal inference in natural settings is often intractable and involves approximations that may deviate from optimality²⁷. Therefore, it is critical to go beyond statements of optimality and suboptimality, and assess the inference algorithms humans use during sensorimotor and cognitive tasks^{28,29}.

¹McGovern Institute for Brain Research, Massachusetts Institute of Technology, Cambridge, MA, USA. ²Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA, USA. Correspondence and requests for materials should be addressed to M.J. (email: mjaz@mit.edu)

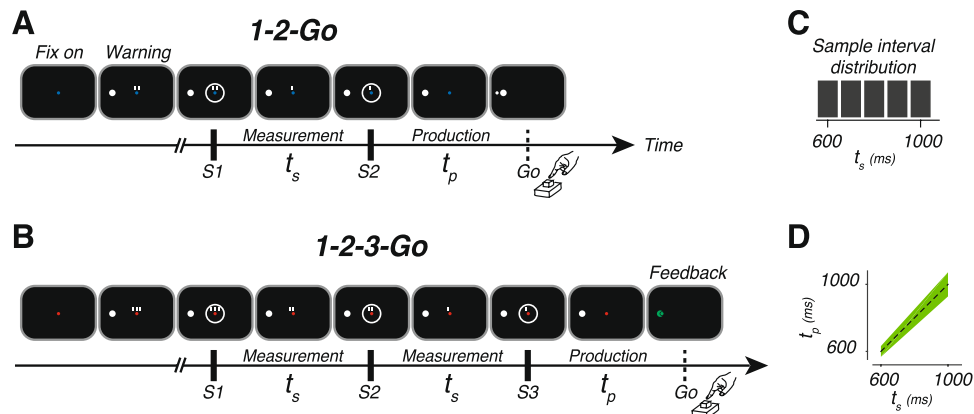


Figure 1. The 1-2-Go and 1-2-3-Go interval reproduction task. **(A,B)** Task design. Each trial began with the appearance of a fixation spot (Fix on). The color of the fixation spot informed the subject of the trial type: blue for 1-2-Go, and red for 1-2-3-Go. After a random delay, a warning stimulus (large white circle) appeared. Additionally, two or three small white rectangles were presented above the fixation spot. The number of rectangles was associated with the number of upcoming flashes. After another random delay, two (S1 and S2 for 1-2-Go) or three (S1, S2 and S3 for 1-2-3-Go) white annulae were flashed for 100 ms in a sequence around the fixation spot. Consecutive flashes were separated by the duration of the sample interval (t_s). With the disappearance of each flash, one of the small rectangles also disappeared (rightmost first and leftmost last). The white rectangles were provided to help subjects keep track of events during the trial. Subjects had to press a button after the last flash to produce an interval (t_p) that matched t_s . Immediately after button press, subjects received feedback. The feedback was a small circle that was presented to the left or right of the warning stimulus depending on whether t_p was larger or smaller than t_s , respectively. The distance of the feedback circle to the center of the warning stimulus was proportional to the magnitude of the error ($t_p - t_s$). **(C)** Experimental distribution of sample intervals. **(D)** Feedback. Subjects received positive feedback if production times fell within the green region. The width of the positive feedback window was scaled with t_s .

Further, as articulated by Marr³⁰, characterization of the underlying algorithms could establish a bridge between behaviorally relevant computations and neurobiological mechanisms.

We devised an experiment in which the computational demands for optimal Bayesian estimation were incompatible with simple algorithmic solutions. Subjects had to reproduce an interval by integrating their prior belief with one or two measurements of the interval. Several previous experiments have reported a decrease in perceptual or motor variability when subjects are given multiple intervals to measure^{31–40}, but the underlying algorithms are not characterized. An important constraint for developing a suitable algorithmic model is that noise associated with measurement and production of time intervals scales with duration⁴¹. A consequence of this so-called scalar property of noise is that simple algorithms that only update certain parameters of the posterior (e.g., mean and/or variance) cannot emulate Bayes-optimal behavior. Therefore, optimal behavior in this paradigm would provide strong evidence that the underlying inference algorithm involves updating probability distributions. Conversely, suboptimal behavior would suggest that subjects rely on a simpler algorithm. We found that when subjects made two measurements their performance was suboptimal. Furthermore, comparison of behavior with various models indicated that subjects relied on an inference algorithm that used measurements to update point estimates using point nonlinearities.

Results

Subjects integrate interval measurements with prior knowledge. Subjects performed an interval reproduction task consisting of two randomly interleaved trial types (Fig. 1A,B). In “1-2-Go” trials, two flashes (S1 followed by S2) demarcated a sample interval (t_s). Subjects had to reproduce t_s immediately after S2. The interval between the onset of S2 and when the keyboard was pressed was designated as the production interval (t_p). In “1-2-3-Go” trials, t_s was presented twice, demarcated once by S1 and S2 and once by S2 and S3, providing the opportunity to make two measurements (Fig. 1B). Similar to 1-2-Go, subjects had to match t_p (the interval between S3 and keyboard press) to t_s . Across trials, t_s was drawn from a discrete uniform distribution ranging between 600 and 1000 ms (Fig. 1C). Subjects received two forms of trial-by-trial feedback based on the magnitude and sign of the error. First, a feedback stimulus was presented whose location relative to the warning stimulus reflected the magnitude and sign of the error (Fig. 1A,B; see Methods). Second, if the error exceeded a threshold window (Fig. 1D), stimuli remained white and a tone denoting incorrect response was presented. Otherwise, the stimuli turned green and a tone denoting correct was presented. The threshold window for correct performance was proportionally larger for longer t_s to accommodate the scalar variability of timing due to signal-dependent noise^{34,42–46}. The threshold was adjusted adaptively and on a trial-by-trial basis according to performance (see Methods).

Subjects’ timing behavior exhibited three characteristic features (Fig. 2). First, t_p increased monotonically with t_s . Second, t_p was systematically biased toward the mean of the prior, as evident from the tendency of responses to deviate from t_s (diagonal) and gravitate toward the mean t_s . As proposed previously^{24,47–49}, this so-called




Figure 2. Performance in the interval reproduction task. (A) Production interval (t_p) as a function of sample interval (t_s) for a low sensitivity subject (GB). Filled circles and error bars show the mean and standard deviation of t_p for each t_s in the 1-2-Go (blue) and 1-2-3-Go (red) conditions. The dotted unity line represents perfect performance and the colored lines show the expected t_p from a Bayes Least-Squares (BLS) model fit to the data. Inset: root-mean-square error (RMSE) in the 1-2-Go (blue) and 1-2-3-Go (red) conditions differed significantly (asterisk, p -value < 0.01; permutation test). (B) Same as (A) for a high sensitivity subject (LB). (C) The histogram of changes in RMSE across conditions for all subjects. See also Supplementary Table 1.

regression to the mean indicated that subjects relied on their knowledge of the prior distribution of t_s . Third, performance was better in 1-2-3-Go condition in which subjects made two measurements, as evidenced by a lower root-mean-square error (RMSE) in 1-2-3-Go compared to 1-2-Go condition (Fig. 2C; permutation test; p -value < 0.01 for all subjects; see Supplementary Table 1 for a summary of RMSE data by subject and condition). This observation indicates that subjects combined the two measurements to improve their estimates by decreasing variability and systematic biases, corroborating reports from other behavioral paradigms^{31–40,50,51}. Combined with the systematic bias toward the mean of t_s , these results indicated that subjects integrated prior information with one or two measurements to improve their performance.

A Bayesian model of behavior. Building on previous work^{11,18,22,52}, we asked whether subjects' behavior could be accounted for by a Bayesian observer model based on the Bayes-Least Squares (BLS) estimator. For the 1-2-Go trials, the observer model (1) makes a noisy measurement of t_s , which we denote by t_{m_1} , (2) combines the likelihood function associated with t_{m_1} , $p(t_{m_1}|t_s)$, with the prior distribution of t_s , $p(t_s)$, to compute the posterior, $p(t_s|t_{m_1})$, and (3) uses the mean of the posterior as the optimal estimate, t_{e_1} . We modeled $p(t_{m_1}|t_s)$ as a Gaussian distribution centered at t_s with standard deviation, σ_m , proportional to t_s with constant of proportionality, w_m ; i.e., $\sigma_m = w_m t_s$ (Fig. 3A, left box). We assumed that the production process was also perturbed by noise and modeled t_p as a sample from a Gaussian distribution centered at t_{e_1} with standard deviation, σ_p , proportional to t_{e_1} with constant of proportionality, w_p ; i.e., $\sigma_p = w_p t_{e_1}$ (Fig. 3A, right box). Note that the entire operation of the BLS estimator can be described in terms of a deterministic mapping of t_{m_1} to t_{e_1} using a nonlinear function, which we denote as $f_{BLS_1}(t_{m_1})$ (Fig. 3B)²⁴.

For the 1-2-3-Go trials, the observer model (1) makes two measurements, t_{m_1} and t_{m_2} , (2) combines the likelihood, $p(t_{m_1}, t_{m_2}|t_s)$, with the prior, $p(t_s)$, to compute the posterior, $p(t_s|t_{m_1}, t_{m_2})$, and (3) uses the mean of the posterior to derive an optimal estimate, t_{e_2} . When the measurements are conditionally independent, the posterior is proportional to $p(t_{m_2}|t_s)p(t_{m_1}|t_s)p(t_s)$ which can be rewritten as $p(t_{m_2}|t_s)p(t_{m_1}|t_s)$. This revised formulation can be interpreted in terms of an updating strategy in which the observer uses the posterior after one measurement, $p(t_s|t_{m_1})$, as the prior for the second measurement (Fig. 3C; see Methods). In these trials, the mapping from t_{m_1} and t_{m_2} to t_{e_2} can be described in terms of a two-dimensional nonlinear function, denoted by $f_{BLS_2}(t_{m_1}, t_{m_2})$ (Fig. 3D). Note that the iso-estimate contours of $f_{BLS_2}(t_{m_1}, t_{m_2})$ are nonlinear and convex (Fig. 3D, red). The nonlinearity indicates that the effect of t_{m_1} and t_{m_2} on t_{e_2} is non-separable, and the convexity indicates that t_{e_2} is more strongly influenced by the larger of the two measurements. These features are direct consequences of scalar noise and are not present when measurements are perturbed by Gaussian noise (see Appendix). Note that the BLS model for the 1-2-3-Go task reduces RMSE by reducing both prior-induced biases and variability, as previous studies have reported^{31–40,50,51}.

We fit the model to each subject's data assuming that responses in both 1-2-Go and 1-2-3-Go conditions were associated with the same w_m and w_p (Methods). The model was augmented in two ways to ensure that estimates of w_m and w_p were accurate. First, we included an offset parameter to absorb interval-independent biases (e.g., consistently pressing the button too early or too late). Second, trials in which t_p grossly deviated from t_s were designated as "lapse" trials (see Methods).

Model fits captured subjects' behavior for both conditions as shown by a few representative subjects (Figs 2A,B, 4A; see Supplementary Fig. 1 for fits to all the subjects). Following previous work²⁴, we evaluated model fits using two statistics, an overall bias, BIAS, and an overall variability, $\sqrt{\text{VAR}}$ (see Methods). As shown in Fig. 4B, the model broadly captured the bias and variance for all subjects in both 1-2-Go and 1-2-3-Go conditions. We did not

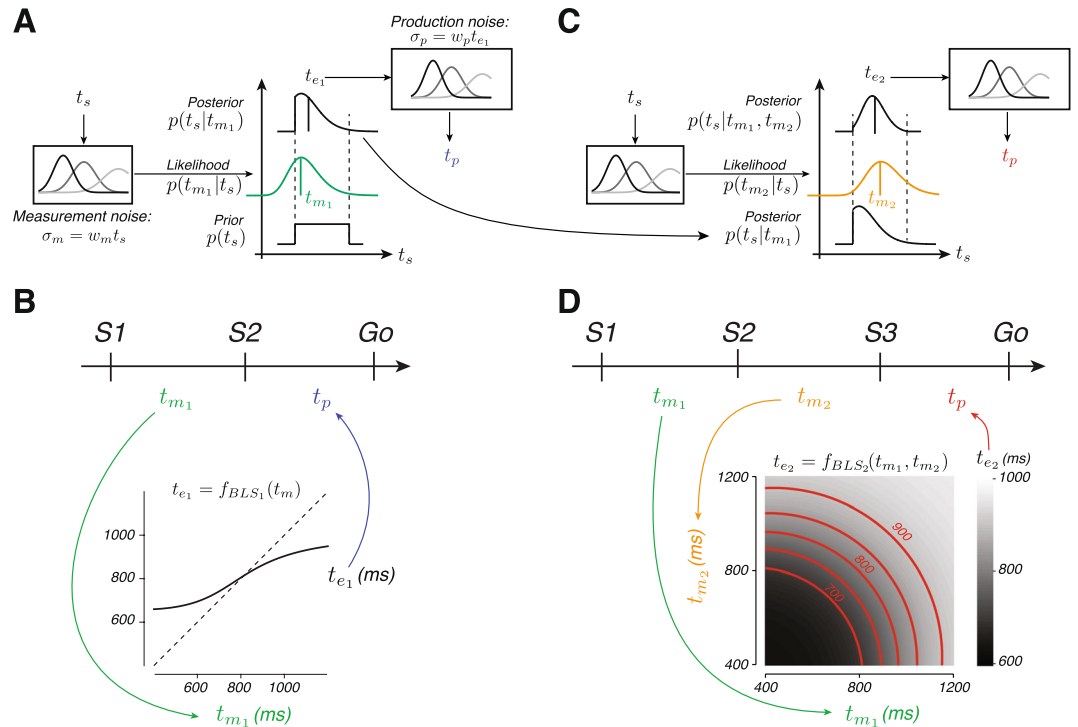


Figure 3. BLS model of interval integration. **(A)** BLS model for 1-2-Go trials. The left panel illustrates the measurement process. The measured interval, t_{m_1} , is perturbed by zero-mean Gaussian noise whose standard deviation is proportional to the sample interval, t_s , with constant of proportionality w_m ($\sigma_m = w_m t_s$). The middle panel illustrates the estimation process. The model multiplies the likelihood function associated with t_{m_1} (middle panel, green) with the prior (bottom), and uses the mean of the posterior (top) to derive an interval estimate (t_{e_1} , black vertical line on the posterior). The right panel illustrates the production process. The produced interval, t_p , is perturbed by zero-mean Gaussian noise with standard deviation proportional to t_{e_1} , with constant of proportionality w_p ($\sigma_p = w_p t_{e_1}$). **(B)** The effective mapping function (f_{BLS_1} , black curve) from the first measurement, t_{m_1} , to the optimal estimate, t_{e_1} . The dashed line indicates unity. **(C)** BLS model for 1-2-3-Go trials. The model uses the posterior after the first measurement, $p(t_s|t_{m_1})$, as the prior and combines it with the likelihood of the second measurement (t_{m_2} , orange) to compute an updated posterior, $p(t_s|t_{m_1}, t_{m_2})$. The mean of the updated posterior is taken as the interval estimate (t_{e_2}). **(D)** The effective mapping function (f_{BLS_2} , grayscale) from each combination of measurements, t_{m_1} and t_{m_2} , to the optimal estimate, t_{e_2} . Red lines indicate combinations of measurements that lead to identical estimates (shown for $t_{e_2} = 700, 750, 800, 850$, and 900 ms).

find any systematic difference in $\sqrt{\text{VAR}}$ between the model and data (see Supplementary Fig. 2). In contrast, the observed BIAS was significantly larger than predicted by the model fits in the 1-2-3-Go condition (Fig. 4B, inset; two tailed t-test, $t(8) = 4.6982$, $p\text{-value} = 0.0015$), but not in the 1-2-Go condition (two tailed t-test, $t(8) = -0.3236$, $p\text{-value} = 0.7546$).

We quantified this observation across subjects by normalizing each subject's RMSE in the 1-2-Go and 1-2-3-Go conditions to the RMSE expected from the BLS model in the 1-2-3-Go condition (Fig. 4C). We found that the observed RMSE in the 1-2-3-Go condition was significantly larger than expected (two tailed t-test, $t(8) = 3.5484$, $p\text{-value} = 0.007$). Further, the drop in observed RMSE in the 1-2-3-Go was significantly less than expected by the BLS model (see Supplementary Fig. 3). These analyses indicate that subjects were able to integrate the two measurements but failed to optimally update the posterior by the likelihood information associated with the second measurement.

Our original BLS model assumed that the noise statistics for t_{m_1} and t_{m_2} were identical. However, for two reasons, the noise statistics between the two measurements may differ. First, the internal representation of the first measurement may be subject to additional noise since it has to be held longer in working memory. Second, after the first measurement is made, subjects may be able to benefit from anticipatory and attentional mechanisms to make a more accurate second measurement. Both possibilities can be straightforwardly modeled by a modified BLS model that accommodates different levels of noise for the two measurements. Therefore, we also formulated an optimal estimator, BLS_{mem} , in which the two measurements were associated with different noise statistics (see Methods, Supplementary Fig. 4). Despite having additional parameters, BLS_{mem} failed to capture the BIAS observed from data in 1-2-3-Go trials (two-tailed t-test, $t(8) = 4.9690$, $p\text{-value} = 0.0011$; Supplementary Fig. 4).

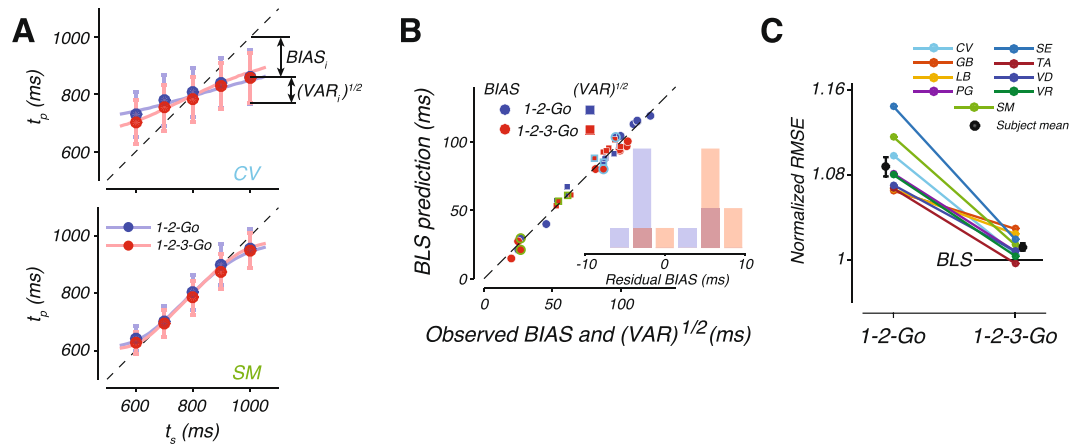


Figure 4. BLS model fits to data. (A) Behavior of two subjects and the corresponding BLS model fits with the same format as in Fig. 2A,B. (B) BIAS (circles) and $\sqrt{\text{VAR}}$ (squares) of each subject (abscissa) and the corresponding values computed from simulations of the fitted BLS model (ordinate). Red and blue points correspond to 1-2-Go and 1-2-3-Go, respectively. The dotted line plots unity. Data points corresponding to subjects SM and CV are marked by light green and light blue, respectively. Inset: difference between the BIAS observed from data and that expected by the BLS model fit for 1-2-Go (blue) and 1-2-3-Go (red) conditions. (C) Comparison of behavioral performance to model predictions. Each line connects the RMSE for 1-2-Go (left) and 1-2-3-Go (right) conditions for one subject. To facilitate comparison across subjects, RMSE values for each subject were normalized by the RMSE of the BLS model in the 1-2-3-Go condition. The black circles and error bars correspond to the mean and standard error of the normalized RMSE across subjects. See also Supplementary Figs 1, 2, and 3.

An algorithmic view of Bayesian integration. The success of the BLS model in capturing behavior in the 1-2-Go condition^{24,47,49} and its failure in the 1-2-3-Go condition suggests that subjects were unable to update the posterior by the second measurement. We examined a number of simple inference algorithms that could account for this limitation. One of the simplest algorithms proposed for integrating sequential measurements is the Kalman filter. The Kalman filter only updates the mean and variance of the posterior²⁶. This strategy is optimal when measurement noise is Gaussian because a Gaussian distribution is fully determined by its mean and variance. More generally, when integrating the likelihood function leaves the parametric form of the posterior distribution unchanged, a simple inference algorithm that updates those parameters can implement optimal integration.

First, we asked whether there exists a similarly simple and optimal updating algorithm when the noise is signal-dependent (i.e., scalar noise). For the posterior to have the same parametric form after one and two measurements, it is necessary that the product of two likelihood functions have the same parametric form as a single likelihood function. We tested this property analytically and verified that the parametric form of the likelihood function associated with scalar noise was not invariant under multiplication (see Appendix). As a result, the updating algorithm requires adjustment within each trial depending on t_{m_1} (Supplementary Fig. 5). In other words, any inference algorithm that only updates certain statistics of the posterior (e.g., mean and variance) is expected to behave suboptimally when multiple time intervals have to be integrated. Therefore, we hypothesized that subjects might have used a simple updating algorithm analogous to the Kalman filter to integrate multiple measurements.

A linear-nonlinear estimator (LNE) model for approximate Bayesian inference. The first algorithm we tested was one in which the observer combines the last estimate \bar{t}_{n-1} with the current measurement, t_{m_n} , using a linear updating strategy. If we denote the corresponding weights by $1-k_n$ and k_n and set $k_n = 1/n$, this algorithm tracks the running average of the measurements, \bar{t}_n (k_1 and k_2 are 1 and 0.5, respectively). This and similar models with linear updating schemes^{32,53–60} would certainly fail to account for the observed nonlinearities in subjects' behavior (Supplementary Fig. 6). Therefore, we constructed a linear-nonlinear estimator (LNE) that augmented the linear updating by a point nonlinearity that could account for the observed prior dependent biases in t_p (Fig. 5A). The nonlinear function, $f_{BLS_1}(\bar{t}_n)$, was chosen to match the BLS estimator for a single measurement ($n = 1$), which is determined by w_m .

Simulation of LNE verified that it could indeed integrate multiple measurements and exhibit prior-dependent biases (see Supplementary Fig. 7). However, the behavior of LNE was qualitatively different from BLS. The contrast between the two models was evident from a comparison of the relationship between measurements and estimates. Unlike BLS (Fig. 3D), estimates derived from LNE are linear with respect to t_{m_1} and t_{m_2} , a feature that can be visualized by the linear iso-estimate contours of the LNE model (Fig. 5B).

We fitted LNE to each subject independently and asked how well it accounted for the observed statistics. The LNE model broadly captured the observed regression to the mean (Fig. 5C,D; see Supplementary Fig. 8 for fits to all subjects), but had a qualitative failure: subject's behavior exhibited significantly more BIAS in 1-2-Go condition (Fig. 5D, inset; two tailed t-test, $t(8) = 4.9304$, $p\text{-value} = 0.001$) and significantly less BIAS in 1-2-3-Go condition

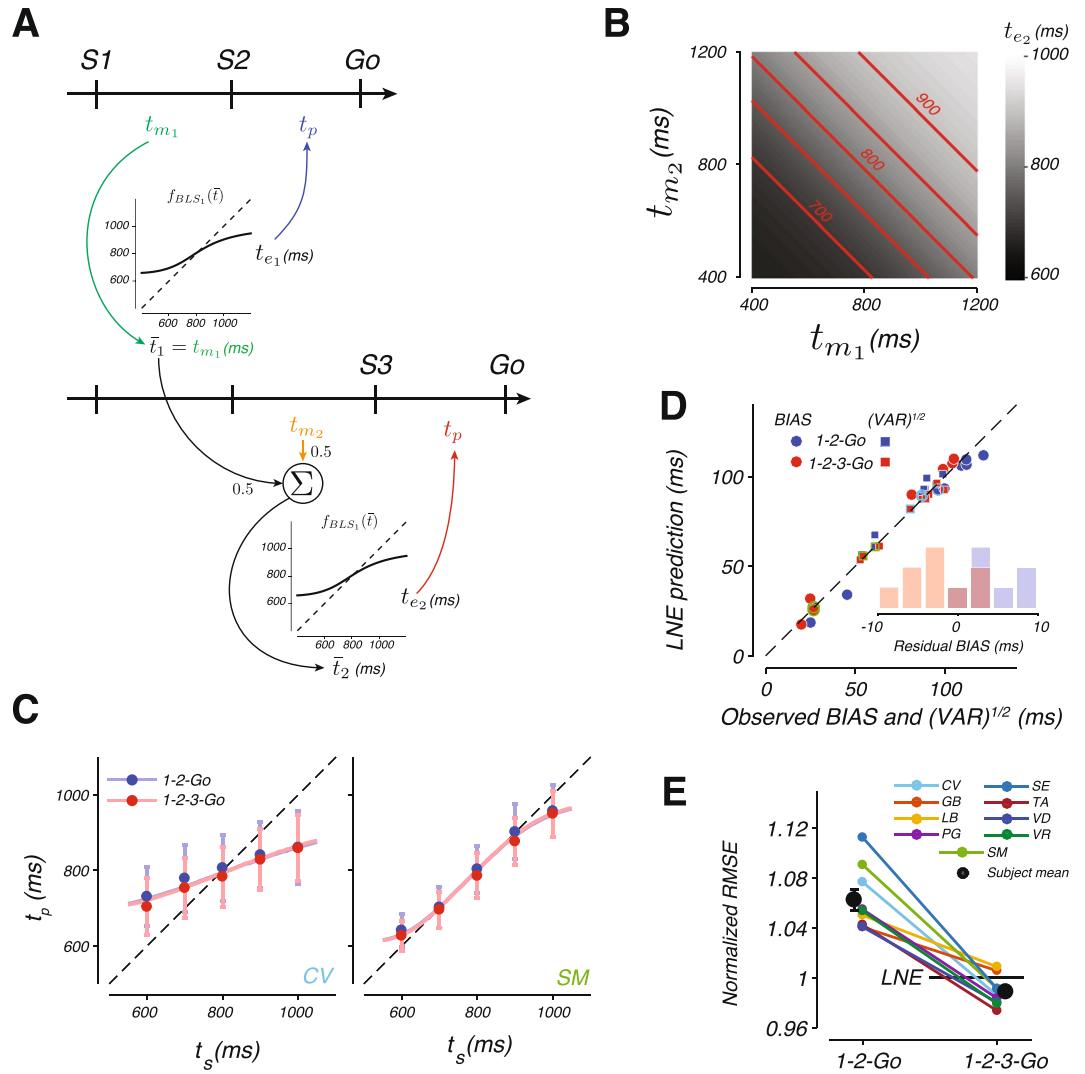


Figure 5. A linear-nonlinear estimator (LNE) model and its fits to the data. **(A)** LNE algorithm. LNE derives an estimate by applying a nonlinear function, f_{BLS_1} , to the average of the measurements. In the 1-2-Go trials (top), the average, \bar{t}_1 , is the same as the first measurement, t_{m_1} , and the estimate, t_{e_1} , is $f_{BLS_1}(\bar{t}_1)$. In 1-2-3-Go trials (bottom), the average, \bar{t}_2 , is updated by the second measurement, t_{m_2} ($\bar{t}_2 = 0.5[\bar{t}_1 + t_{m_2}]$), and the estimate, t_{e_2} , is $f_{BLS_1}(\bar{t}_2)$. In both conditions, the produced interval, t_p , is perturbed by zero-mean Gaussian noise with standard deviation proportional to the final estimate (t_{e_1} for 1-2-Go and t_{e_2} for 1-2-3-Go) with the constant of proportionality w_p , as in the BLS model. **(B)** The mapping from measurements to estimates (grayscale) for the LNE estimator in the 1-2-3-Go trials. Red lines indicate combinations of measurements that lead to identical estimates (shown for $t_{e_2} = 700, 750, 800, 850,$ and 900 ms). **(C)** Mean and standard deviation of t_p as a function of t_s for two example subjects (circles and error bars) along with the corresponding fits of the LNE model (lines). **(D)** BIAS (circles) and $\sqrt{\text{VAR}}$ (squares) of each subject (abscissa) and the corresponding values computed from simulations of the fitted LNE model (ordinate). Conventions match Fig. 4B. **(E)** The RMSE in the 1-2-Go and 1-2-3-Go conditions relative to the corresponding predictions from the LNE model (conventions as in Fig. 4C). See also Supplementary Figs 2, 3, 7, and 8.

(Fig. 5D, inset; two tailed t-test, $t(8) = -2.3782$, $p\text{-value} = 0.045$) than the biases predicted by the model. This failure can be readily explained in terms of how LNE functions. Since the static nonlinearity in LNE is the same for one and two measurements, the bias LNE generates is the same for the 1-2-Go and 1-2-3-Go conditions and improvements in estimation are achieved mainly through a reduction in VAR. Therefore, when we fitted LNE to data from both conditions, the model consistently underestimated BIAS for the 1-2-Go condition, and overestimated BIAS for the 1-2-3-Go condition (Fig. 5C, red and blue lines nearly overlap). Further, the LNE model made systematic errors predicting subject VAR in 1-2-Go trials (Supplementary Fig. 2).

We further evaluated LNE by asking how it accounted for the observed performance improvement in the 1-2-3-Go condition compared to the 1-2-Go condition. We normalized each subject's RMSE from the 1-2-Go and 1-2-3-Go conditions to the RMSE expected from the behavior of the fitted LNE model in the 1-2-3-Go condition (Fig. 5E). Most subjects surpassed the predictions of the LNE model (horizontal line) for the 1-2-3-Go




Figure 6. An extended Kalman filter (EKF) model and its fits to the data. **(A)** EKF algorithm. EKF is a real-time inference algorithm that uses each measurement to update the estimate. After the first flash, EKF uses the mean of the prior as its initial estimate, t_{e_0} . The second flash furnishes the first measurement, t_{m_1} . EKF computes a new estimate, t_{e_1} , using the following procedure: (1) it measures the difference between t_{m_1} and t_{e_0} to compute an error, x_1 , (2) it applies a nonlinear function, $f^*(x)$, to x_1 , (3) it scales $f^*(x_1)$ by a gain factor, k_1 , whose magnitude depends on the relative reliability of t_{m_1} and t_{e_0} , and (4) it adds $k_1 f^*(x_1)$ to t_{e_0} to compute t_{e_1} . In the 1-2-Go condition (top), t_{e_1} is the final estimate used for the production of t_p . In the 1-2-3-Go condition (bottom), the updating procedure is repeated to compute a new estimate t_{e_2} by adding t_{e_1} to $k_2 f^*(x_2)$ where x_2 is the difference between the second measurement, t_{m_2} , and t_{e_1} , and k_2 is the scale factor determined by the relative reliability of t_{e_1} and t_{m_2} . t_{e_2} is then used as the final estimate for the production of t_p . We assumed that the produced interval, t_p , is perturbed by zero-mean Gaussian noise with standard deviation proportional to the final estimate (t_{e_1} for 1-2-Go and t_{e_2} for 1-2-3-Go) with the constant of proportionality w_p , as in the BLS model. **(B)** The mapping from measurements to estimates (grayscale) for the EKF estimator in the 1-2-3-Go condition. Red lines indicate combinations of measurements that lead to identical estimates (shown for $t_{e_2} = 700, 750, 800, 850,$ and 900 ms). **(C)** Mean and standard deviation of t_p as a function of t_s for two example subjects (circles and error bars) along with the corresponding fits of the EKF model (lines). **(D)** BIAS (circles) and $\sqrt{\text{VAR}}$ (squares) of each subject (abscissa) and the corresponding values computed from simulations of the fitted EKF model (ordinate). Conventions match Fig. 4B. **(E)** The RMSE in the 1-2-Go and 1-2-3-Go conditions relative to the corresponding predictions from the EKF model (conventions as in Fig. 4C). See also Supplementary Figs 2, 3, 9, and 10.

condition, and average RMSE reached values that were significantly lower than expected (0.990; two tailed t-test, $t(8) = -2.463$, $p\text{-value} = 0.039$). Based on these results, we concluded that LNE fails to capture subjects' behavior both qualitatively and quantitatively.

An extended Kalman filter (EKF) model for approximate Bayesian inference. We considered a moderately more sophisticated algorithm inspired by the extended Kalman filter (EKF)⁶¹. This algorithm is

Model	Subject								
	CV	GB	LB	PG	SM	SE	TA	VR	VD
LNE	1.1678 (0.5012)	-1.8267 (0.2976)	-0.3094 (0.2020)	0.1516 (0.3407)	0.4896 (0.3427)	-1.7151 (0.3493)	-0.0365 (0.2806)	0.3456 (0.3593)	0.3793 (0.2175)
EKF	1.1900 (0.1998)	0.0380 (0.1326)	0.7168 (0.2234)	0.3828 (0.1428)	-1.4262 (0.2755)	1.0909 (0.1543)	0.2358 (0.1250)	0.3623 (0.1581)	0.6537 (0.1002)
BLSmem	0.4621 (0.1832)	-0.0821 (0.0267)	-0.2445 (0.0743)	0.0086 (0.0518)	-0.0637 (0.2714)	0.0440 (0.0876)	-0.0191 (0.0195)	-0.0359 (0.0249)	-0.0421 (0.0268)

Table 1. Predictive log likelihood ratio for each model and subject. For each subject, the model with the highest predictive log likelihood is shown in bold. All log likelihood ratios were all smaller than 2 and consistent with what is expected based on simulations of each model (Supplementary Figs 7, 9 and 11).

shown in Fig. 6A. Upon each new measurement, EKF uses the error between the previous estimate and the current measurement to generate a new estimate. The difference between EKF and the Kalman filter is that the error is subjected to a nonlinear function before being used to update the previous estimate. This nonlinearity is necessary for the algorithm to be able to account for the nonlinear prior-dependent biases observed in behavior.

In our experiment, immediately after the first flash, the only information about the sample interval, t_s , comes from the prior distribution. Accordingly, we set the initial estimate, t_{e_0} , to the mean of the prior distribution. After the first measurement, EKF computes an “innovation” term by applying a static nonlinearity, $f^*(x)$ to the error, x_1 , between t_{m_1} and t_{e_0} . This innovation is multiplied by a gain, k_1 , and added to t_{e_0} to compute the new estimate, t_{e_1} . In the 1-2-Go condition in which only one measurement is available, t_{e_1} serves as the final estimate that the model aims to reproduce.

For the 1-2-3-Go condition, EKF repeats the updating procedure after the second measurement, t_{m_2} . It computes the difference between t_{m_2} and t_{e_1} to derive a prediction error, x_2 , which is subjected to the same nonlinear function, $f^*(x)$, to yield a second innovation. This innovation is then scaled by an appropriate gain, k_2 , and added to t_{e_1} to generate an updated estimate, t_{e_2} , which the model aims to reproduce.

The two important elements that determine the overall behavior of EKF are the nonlinear function $f^*(x)$ and the gain factor(s) applied to the innovation(s) (k_1 and k_2) to update the estimate(s). We set the form of the nonlinear function $f^*(x)$ such that biases in t_{e_1} after one measurement are the same between EKF and BLS models. This ensures that EKF and BLS behave identically in the 1-2-Go condition. Note that our implementation of EKF assumes that the same nonlinear function is applied after every measurement. If one allows this nonlinear function to be optimized separately for each measurement, EKF would be able to replicate the behavior of BLS exactly (Supplementary Fig. 5).

For the gain factors, we reasoned that the most rational choice is to set the weight of each innovation based on the expected reliability of the corresponding estimate, $t_{e_{n-1}}$, relative to the new measurement, t_{m_n} , as in the Kalman filter (see Methods). This causes the gain factor to decrease with the number of measurements, and ensures that the influence of each new measurement is appropriately titrated. With these assumptions, EKF remains suboptimal for the 1-2-3-Go condition. However, it captures certain aspects of the nonlinearities associated with the optimal BLS estimator as shown by Fig. 6B (compare to Fig. 3D).

The algorithm implemented by EKF is appealing as it uses a simple updating strategy that can be straightforwardly extended to multiple sequential measurements and is a nonlinear version of error correcting mechanisms proposed for related synchronization tasks^{54,56,58}. Furthermore, EKF captures important features of human behavior. First, integration of each new measurement causes a reduction in RMSE, as seen in 1-2-3-Go compared to 1-2-Go condition. Second, the nonlinear function applied to innovations allows EKF to incorporate prior information and capture prior-dependent biases. Third, since the nonlinearity is applied to each innovation (as opposed to the final estimate), EKF, unlike LNE, is able to capture the reduction in BIAS in 1-2-3-Go compared to 1-2-Go condition.

After simulating the model to ensure it integrates measurements and exhibits prior-dependent biases (Supplementary Fig. 9), we fitted EKF to each subject independently and asked how well it accounted for the observed statistics. Similar to BLS and LNE, EKF broadly captured the observed regression to the mean in the 1-2-Go trials (Fig. 6C,D, blue). This is not surprising since the EKF algorithm is identical to BLS when the prior is integrated with a single measurement. EKF was also able to capture the mean t_p as a function of t_s in the 1-2-3-Go trials (Fig. 6C,D red). Across subjects, EKF provided a better match to the data when compared to BLS and LNE although it modestly underestimated the BIAS in 1-2-3-Go condition (Fig. 6D, inset; two tailed t-test, $t(8) = 4.6055$, $p\text{-value} = 0.02639$). See Supplementary Fig. 10 for fits of the EKF model to all subjects.

We also asked if EKF could account for the observed RMSEs. To do so, we performed the same analysis we used to evaluate the BLS and LNE models. We normalized each subject’s RMSE from the 1-2-Go and 1-2-3-Go conditions to the RMSE expected from the EKF model for 1-2-3-Go (Fig. 6E). We found no significant difference between observed and predicted RMSEs for the 1-2-3-Go condition (two-tailed t-test, $t(8) = 1.5506$, $p\text{-value} = 0.160$), and no significant difference between the observed and predicted change in RMSE from the 1-2-Go to the 1-2-3-Go condition (Supplementary Fig. 3). These results indicate that subjects’ suboptimal behavior is consistent with the approximate Bayesian integration implemented by the EKF algorithm.

To further validate the superiority of the EKF model, we directly compared various models to BLS using log likelihood ratio. Specifically, we computed the ratio of the log likelihood of the data given each model and maximum likelihood parameters ($\log \mathcal{L}(t_s, t_p | \mathcal{M}_i, \Theta^{\text{ML}})$ see Methods) to the log likelihood of the BLS model, $\log \mathcal{L}(t_s, t_p | \mathcal{M}_{\text{BLS}}, \Theta^{\text{ML}})$, for each subject. We found that EKF provided the best fit for 8 out of 9 subjects (Table 1). For one of the subjects, the fits were poor for all models but LNE provided the best fit.

Discussion

The neural systems implementing sensorimotor transformations must rapidly compute state estimates to effectively implement online control of behavior. Behavioral studies indicate that, at a computational level, state estimation may be described in terms of Bayesian integration^{6–15}. However, describing behavior with a Bayesian model does not necessarily indicate that the brain implements these computations by representing probability distributions^{21,22,62}. Here, we focused on integration of multiple time intervals and found evidence that the brain relies on simpler algorithms that approximate optimal Bayesian inference.

We demonstrated that humans integrate prior knowledge with one or two measurements to improve their performance. A key observation was that the integration was nearly optimal for one measurement but not for two measurements. In particular, when two measurements were provided, subjects systematically exhibited more BIAS toward the mean of the prior than expected from an optimal Bayesian model. This observation motivated us to investigate various algorithms that could lead to similar patterns of behavior.

Analytical and numerical analyses suggested that simple inference algorithms that update certain parameters of the posterior instead of the full distribution can not integrate multiple measurements optimally when the noise is signal-dependent. We then systematically explored simple inference algorithms that could perform sequential updating and account for the behavioral observations. One of the simplest updating algorithms is the Kalman filter⁶³. However, this algorithm updates estimates linearly and thus cannot account for the nonlinearities in subjects' behavior, even for a single measurement (Supplementary Fig. 6). The LNE model augmented the Kalman filter such that the final estimate was subjected to a point-nonlinearity. This allowed LNE to generate nonlinear biases but since the nonlinearity was applied to the final estimate, LNE failed to capture the decrease in bias observed in the 1-2-3-Go compared to 1-2-Go condition. Finally, we adapted the EKF, which is a more sophisticated variant of the Kalman filter that applies a static nonlinearity to the errors in estimation at every stage of updating. This algorithm accounted for optimal behavior in the 1-2-Go condition and exhibited the same patterns of suboptimality observed in humans in the 1-2-3-Go condition. Therefore, EKF provides a good characterization of the algorithm the brain uses when there is need to integrate multiple pieces of information presented sequentially. This finding implies that subjects may only rely on the first few moments of a distribution and use nonlinear updating strategy to track those instead of updating the entire posterior. This strategy is simple and in many scenarios could lead to optimal behavior with little computational cost. Moreover, the recursive nature of EKF's updating strategy allows it to readily generalize to scenarios when it is necessary to update estimates in real time, even when the number of available samples is not known *a priori*, which could be tested by extensions of our experiment. Finally, the error-correcting nature of the EKF updates are consistent with the correlations between response intervals observed in synchronization and continuation tasks^{53,57,64–66}, corrections to timing perturbations during synchronization^{53,58,67}, and the influence of recent temporal inputs on the perception of interval duration^{68–70}. Therefore, EKF may provide an algorithmic understanding across a range of timing tasks including interval estimation, synchronization, discrimination, and reproduction.

Maintaining and updating probability distributions is computationally expensive. Moreover, it is not currently known how neural networks might implement such operations²². In contrast, EKF is relatively simple to implement. The only requirement is to use the current estimate to predict the next sample, and use a nonlinear function of the error in prediction to update estimates sequentially. Predictive mechanisms that EKF relies on are thought to be an integral part of how brain circuits support perception and sensorimotor function^{1,5,8,71–74}. As such, the relative success of EKF may be in part due to its compatibility with predictive mechanisms that the brain uses to perform sequential updating. This observation makes the following intriguing prediction: when performing 1-2-3-Go task, subjects do not make two measurements; instead, they use the prior to predict the time of the second flash, use the prediction error to update their estimate, and the new estimate to predict the third flash. Two lines of evidence from recent physiological experiments support this prediction. First, individual neurons in the primate medial frontal cortex encode interval duration similarly for single and multiple interval tasks⁷⁵, providing a basis for maintaining the current interval estimate during a given task. Second, neural signals in several regions of the brain encode intervals prospectively^{76–80}, providing a basis for predicting the timing of upcoming sensory inputs. Future modeling efforts and electrophysiological experiments will be required to link neural signals to the implementation of the EKF algorithm.

While EKF provides a better account of the observed data in our experiments, it may be that our specific formulation of the Bayesian model did not capture the underlying process. Our BLS model was based on three assumptions: (1) that likelihood function is characterized by signal-dependent noise, (2) that the subjective prior matches the experimentally imposed uniform prior distribution, and (3) that the final estimates are derived from the mean of the posterior, which implicitly assumes that subject rely on a quadratic cost function, as was previously demonstrated²⁴. Our formulation of the likelihood function is particularly important, as it is the key factor that prohibits simple algorithms such as EKF to optimally integrate multiple measurements. The inherent signal-dependent noise in timing causes the likelihood function to be skewed toward longer intervals (see Appendix). This characteristic feature was particularly important for explaining human behavior in a task requiring interval estimation following several measurements⁴⁰. Moreover, it has been shown that subjects exhibit larger biases for longer intervals within the domain of the prior indicating that the brain has an internal model of this signal-dependent noise^{24,49,81,82}. These results support our formulation of the likelihood function. However, one aspect of our formulation that deserves further scrutiny is the assumption that noise perturbing the two measurements was independent. This may need revision given the long-range positive autocorrelations in behavioral variability^{83–85}, and because S2 is shared between the two measurements, which may lead to correlations between t_{m_1} and t_{m_2} .

Our formulation of the prior and cost function should also be further evaluated. For example, humans may not be able to correctly encode a uniform prior probability distribution for interval estimation^{47,49}. Similarly,

the cost function may not be quadratic⁸⁶. However, since priors and the cost functions impact both 1-2-Go and 1-2-3-Go conditions, moderate inaccuracies in modeling these components may not be able to explain optimal behavior in 1-2-Go and suboptimal behavior in the 1-2-3-Go condition simultaneously. Finally, recent results suggest the performance may be limited by imperfect integration^{87,88} and imperfect memory^{38,89}, which future models of sequential updating should incorporate.

Methods

Subjects and apparatus. All experiments were performed in accordance with relevant regulations and guidelines for the ethical treatment of subjects, as approved by the Committee on the Use of Humans as Experimental Subjects at MIT, after receiving informed consent. Eleven human subjects (6 male and 5 female) between 18 and 33 years of age participated in the interval reproduction experiment. Of the 11 subjects, 10 were naive to the purpose of the study.

Subjects sat in a dark, quiet room at a distance of approximately 50 cm from a display monitor. The display monitor had a refresh rate of 60 Hz, a resolution of 1920 by 1200, and was controlled by a custom software (MWorks; <http://mworks-project.org/>) on an Apple Macintosh platform.

Interval reproduction task. Experiment consisted of several 1 hour sessions in which subjects performed an interval reproduction task (Fig. 1). The task consisted of two randomly interleaved trial types referred to as “1-2-Go” and “1-2-3-Go”. On 1-2-Go trials, two flashes (S1 followed by S2) demarcated a sample interval (t_s) that subjects had to measure²⁴. On 1-2-3-Go trials, t_s was presented twice, once demarcated by S1 and S2 flashes, and once by S2 and S3 flashes. For both trial types, subjects had to reproduce t_s immediately after the last flash (S2 for 1-2-Go and S3 for 1-2-3-Go) by pressing a button on a standard Apple keyboard. On all trials, subjects had to initiate their response proactively and without any additional cue (no explicit Go cue was presented). Subjects received graded feedback on their accuracy.

Each trial began with the presentation of a 0.5 deg circular fixation point at the center of a monitor display. The color of the fixation was blue or red for the 1-2-Go and 1-2-3-Go trials, respectively. Subjects were asked to shift their gaze to the fixation point and maintain fixation throughout the trial. Eye movements were not monitored. After a random delay with a uniform hazard (100 ms minimum plus an interval drawn from an exponential distribution with a mean of 300 ms), a warning stimulus and a trial cue were presented. The warning stimulus was a white circle that subtended 1.5 deg and was presented 10 deg to the left of the fixation point. The trial cue consisted of 2 or 3 small rectangles 0.6 deg above the fixation point (subtending 0.2×0.4 deg, 0.5 deg apart) for the 1-2-Go and 1-2-3-Go trials, respectively. After a random delay with a uniform hazard (250 ms minimum plus an interval drawn from an exponential distribution with mean of 500 ms), flashes demarcating t_s were presented. Each flash (S1 and S2 for 1-2-Go and S1, S2, and S3 for 1-2-3-Go) lasted for 6 frames (100 ms) and was presented as an annulus around the fixation point with an inside and outside diameter of 2.5 and 3 deg, respectively (Fig. 1A,B). The time between consecutive flashes, which determined t_s , was sampled from a discrete uniform distribution ranging between 600 and 1000 ms with a 5 samples (Fig. 1C). To help subjects track the progression of events throughout the trial, after each flash, one rectangle from the trial cue would disappear (starting from the rightmost).

Produced interval (t_p) was measured as the interval between the time of the last flash and the time when the subject pressed a designated key on the keyboard (Fig. 1A,B). Subjects received trial-by-trial visual feedback based on the magnitude and sign of the relative error, $(t_p - t_s)/t_s$. A 0.5 deg circle (“analog feedback”) was presented to the right (for error < 0) or left (error > 0) of the warning stimulus at a distance that scaled with the magnitude of the error. Additionally, when the error was smaller than a threshold, both the warning stimulus and the analog feedback turned green and a tone denoting “correct” was presented. If the production error was larger than the threshold, the warning stimulus and analog feedback remained white and a tone denoting “incorrect” was presented. The threshold was constant as a function of the relative error and therefore scaled with the sample interval (Fig. 1D). This accommodated the scalar variability of timing that leads to more variable production intervals for longer sample intervals. The scaling factor was initialized at 0.15 at the start of every session and adjusted adaptively using a one-up, one-down scheme that added or subtracted 0.001 to the scaling factor for incorrect or correct responses, respectively. These manipulations ensured that the performance across conditions, subjects, and trials remained approximately at a steady state of 50% correct trials.

To ensure subjects understood the task design, the first session included a number of training blocks. Training blocks were conducted with the supervision of an experimenter. Training trials were arranged in 25 trial blocks. In the first block, the subjects performed the 1-2-Go condition with the sample interval fixed at 600 ms. In the second block, we fixed the interval to be 1000 ms. In the third block, the subject performed the 1-2-3-Go task with the interval fixed at 1000 ms. In the fourth block, the subject continued to perform the 1-2-3-Go task, but with the intervals chosen at random from the experimental distribution. In the final training block, the task condition and sample intervals were fully randomized, as in the main experiment. The subject then performed 400 trials of the main experiment. Subjects completed 10 sessions total, performing 800 trials in each of the remaining 9 experimental sessions. To ensure subjects were adapted to the statistics of the prior²⁴, we discarded the first 99 trials of each session. We also discarded any trial when the subject responded before S2 (for 1-2-Go) or S3 (for 1-2-3-Go) or 1000 ms after the veridical t_s . Supplementary Table 2 summarizes the number of completed trials for each subject. Data from two subjects were not included in the analyses because they were not sensitive to the range of sample intervals we tested and their production interval distributions were not significantly different for the longest and shortest sample intervals.

Models. We considered several models for the interval estimation: (1) an optimal Bayes Least-Squares model (BLS), (2) an optimal Bayes Least-Squares model that allowed different noise levels for the two measurements in 1-2-3-Go trials, (3) an extended Kalman filter model (EKF), and (4) a linear-nonlinear estimation model (LNE).

All models were designed to be identical for the 1-2-Go task where only one measurement was available but differed in their prediction for the 1-2-3-Go trials.

BLS model. We used the Bayesian integration model that was previously shown to capture behavior in the 1-2-Go task²⁴. This model assumes that subjects combine the measurements and the prior distribution probabilistically according to Bayes' rule:

$$p(t_s|t_m) = \frac{\lambda(t_m|t_s)p(t_s)}{p(t_m)}, \quad (1)$$

$$\lambda(t_m|t_s) = \frac{1}{\sqrt{2\pi w_m^2 t_s^2}} e^{-\frac{(t_m-t_s)^2}{2w_m^2 t_s^2}}, \quad (2)$$

where $p(t_s)$ represents the prior distribution of the sample intervals and $p(t_m)$ the probability distribution of the measurements. The likelihood function, $\lambda(t_m|t_s)$, was formulated based on the assumption that measurement noise was Gaussian and had zero mean. To incorporate scalar variability into our model, we further assumed that the standard deviation of noise scales with t_s with constant of proportionality w_m representing the Weber fraction for measurement.

Following previous work²⁴, we further assumed that subjects' behavior can be described by a BLS estimator that minimizes the expected squared error, and uses the expected value of the posterior distribution as the optimal estimate:

$$t_{e_1} = f_{BLS_1}(t_{m_1}) = E[t_s|t_{m_1}], \quad (3)$$

where f_{BLS_1} denotes the BLS function that maps the measurement (t_{m_1}) to the Bayesian estimate after one measurement (t_{e_1}). The subscript 1 is added to clarify that this equation corresponds to the condition with a single measurement (i.e., 1-2-Go). The notation $E[\bullet]$ denotes expected value. Given a uniform prior distribution with a range from t_s^{min} to t_s^{max} , the BLS estimator can be written as:

$$f_{BLS_1}(t_{m_1}) = \frac{\int_{t_s^{min}}^{t_s^{max}} t_s \lambda(t_{m_1}|t_s) dt_s}{\int_{t_s^{min}}^{t_s^{max}} \lambda(t_{m_1}|t_s) dt_s}, \quad (4)$$

We assume that t_s^{min} and t_s^{max} match the minimum and maximum of the experimentally imposed sample interval distribution. We extended this model for the 1-2-3-Go task to two measurements. To do so, we incorporated two likelihood functions in the derivation of the posterior. Assuming that the two measurements are conditionally independent, the posterior can be written as:

$$p(t_s|t_{m_1}, t_{m_2}) = \frac{\lambda(t_{m_1}|t_s)\lambda(t_{m_2}|t_s)p(t_s)}{\int \lambda(t_{m_1}|t_s)\lambda(t_{m_2}|t_s)p(t_s) dt_s}, \quad (5)$$

where t_{m_1} and t_{m_2} denote the first and second measurements, respectively, and the likelihood function, λ , is from Eq (2). Because measurements are taken in a sequence, it is intuitive to rewrite Eq (5) in a recursive form:

$$p(t_s|t_{m_1}, t_{m_2}) = \frac{1}{Z} \lambda(t_{m_2}|t_s) p(t_s|t_{m_1}), \quad (6)$$

where $p(t_s|t_{m_1})$ is the posterior as specified in Eq (1) and Z is a normalization factor which ensures the integral over the density sums to one. Note that Z includes terms related to $p(t_{m_1})$, allowing it to appropriately normalize the density after propagating the posterior related to the first measurement forward. Therefore, although the posterior for Eqs (5) and (6) are identical, specifying the posterior in this way allows for the algorithm to be updated sequentially.

The corresponding BLS estimator can again be written as the expected value of the posterior:

$$t_{e_2} = f_{BLS_2}(t_{m_1}, t_{m_2}) = E[t_s|t_{m_1}, t_{m_2}], \quad (7)$$

$$f_{BLS_2}(t_{m_1}, t_{m_2}) = \frac{\int_{t_s^{min}}^{t_s^{max}} t_s \lambda(t_{m_1}|t_s)\lambda(t_{m_2}|t_s) dt_s}{\int_{t_s^{min}}^{t_s^{max}} \lambda(t_{m_1}|t_s)\lambda(t_{m_2}|t_s) dt_s}, \quad (8)$$

where f_{BLS_2} denotes the BLS function that uses two measurements (t_{m_1} and t_{m_2}) to compute t_{e_2} . The subscript 2 indicates the mapping function is for two measurements (i.e., 1-2-3-Go). We performed the integrations for the BLS model numerically using Simpson's quadrature.

BLS_{mem} model. We also considered the possibility that the brain may not be able to hold representations of the first measurement or the associated posterior perfectly over time until the time for integration. To model this we assumed two Weber fractions – w_m as formulated in the BLS model and w_{mem} which adjusts the Weber fraction of the first measurement in 1-2-3-Go trials to account for noisy memory or inference processes. In 1-2-Go trials, the posterior was set according to Eq (1) with w_m controlling the signal dependent noise. In 1-2-3-Go trials, the posterior was set according to:

$$p(t_s | t_{m_1}, t_{m_2}) = \frac{\lambda_{mem}(t_{m_1} | t_s) \lambda(t_{m_2} | t_s) p(t_s)}{\int \lambda_{mem}(t_{m_1} | t_s) \lambda(t_{m_2} | t_s) p(t_s) dt_s}, \quad (9)$$

With the likelihood function associated with the first measurement defined as:

$$\lambda_{mem}(t_m | t_s) = \frac{1}{\sqrt{2\pi w_{mem}^2 t_s^2}} e^{-\frac{(t_m - t_s)^2}{2w_{mem}^2 t_s^2}}, \quad (10)$$

This formulation allows the measurement noise to be different for the two measurements. The optimal estimator was then calculated as:

$$f_{BLS_2}(t_{m_1}, t_{m_2}) = \frac{\int_{t_s^{min}}^{t_s^{max}} t_s \lambda_{mem}(t_{m_1} | t_s) \lambda(t_{m_2} | t_s) dt_s}{\int_{t_s^{min}}^{t_s^{max}} \lambda_{mem}(t_{m_1} | t_s) \lambda(t_{m_2} | t_s) dt_s}, \quad (11)$$

EKF model. EKF implements an updating algorithm in which, after each flash, the observer updates the estimate, t_{e_n} , based on the previous estimate, $t_{e_{n-1}}$, and the current measurement, t_{m_n} . The updating rule changes $t_{e_{n-1}}$ by a nonlinear function of the error between $t_{e_{n-1}}$ and t_{m_n} , which we denote by x_n .

$$t_{e_n} = t_{e_{n-1}} + k_n f^*(x_n), \quad (12)$$

$$x_n = t_{m_n} - t_{e_{n-1}}, \quad (13)$$

f^* is a nonlinear function based on the BLS estimator, f_{BLS_1} :

$$f^*(x_n) = f_{BLS_1}(x_n + t_{e_0}) - t_{e_0}, \quad (14)$$

k_n is a gain factor that controls the magnitude of the update and is set by the relative reliability of $t_{e_{n-1}}$ and t_{m_n} , which were formulated in terms of two Weber fractions, w_{n-1} and w_m , respectively:

$$k_n = \frac{w_{n-1}^2}{w_{n-1}^2 + w_m^2}, \quad (15)$$

To track the reliability of t_{e_n} we used a formulation based on optimal cue combination under Gaussian noise. For Gaussian likelihoods, the reliability of the estimate is related to the inverse of the variance of the posterior. Similarly, the reliability of the interval estimate is related to the inverse of the Weber fraction. Therefore, we used the following algorithm to track the Weber fraction of the estimate, w_n :

$$w_n = \frac{w_{n-1} w_m}{\sqrt{w_{n-1}^2 + w_m^2}}, \quad (16)$$

As in the case of Gaussians, this algorithm ensures that w_n decreases with each additional measurement, reflecting the increased reliability of the estimate relative to the measurement. This ensures that the weight of the innovation respects information already integrated into the estimate by previous iterations of the EKF algorithm.

At S1, no measurements are available. Therefore, we set the initial estimate, t_{e_0} to the mean of the prior, and its reliability, w_0 , to ∞ . After S2 (one measurement), the EKF estimate is identical to the BLS model. For two measurements, the process is repeated to compute t_{e_2} , but the estimate is suboptimal. This formulation can be readily extended to more than two measurements.

LNE model. LNE uses a linear updating strategy similar to a Kalman filter to update estimates by measurements as follows:

$$\tilde{t}_n = (1 - k_n) \tilde{t}_{n-1} + k_n t_{m_n}, \quad (17)$$

The algorithm is initialized such that $\tilde{t}_1 = t_{m_1}$ and we chose the weighting to be $k_n = 1/n$. This choice minimizes the squared errors in 1-2-3-Go trials. Note that any other choice for k_n would deteriorate LNE's performance. Following this sequential and linear updating scheme, LNE passes the final estimate through a nonlinear transfer function specified by the BLS model for one measurement (f_{BLS_1}):

$$f_{LNE_n}(t_{m_1}, \dots, t_{m_n}) = f_{BLS_1}(\bar{t}_n), \quad (18)$$

where f_{LNE_n} denotes the linear-nonlinear estimator after n measurements. This formulation ensures that LNE is identical to the BLS in 1-2-Go trials.

Interval production model. In all models, the final estimate is used for the production phase. Following previous work²⁴, we assumed that the production of an interval is perturbed by Gaussian noise whose standard deviation scales with the estimated interval. The model was additionally augmented by an offset term to account for stimulus-independent biases observed in responses:

$$p(t_p|t_e) = \frac{1}{\sqrt{2\pi w_p^2 t_e^2}} e^{-\frac{(t_p - t_e - b)^2}{2w_p^2 t_e^2}}, \quad (19)$$

where w_p is the Weber fraction for production, b is the offset term, and t_e can refer to the estimate for either 1-2-Go and 1-2-3-Go trials.

All models accommodated “lapse trials” in which the produced interval was outside the mass of the production interval distribution. The lapse trials were modeled as trials in which the production interval was sampled from a fixed uniform distribution, $p(t_p|\text{lapse})$, independent of t_s . With this modification, the production interval distribution can be written as:

$$p(t_p|t_e, \gamma) = (1 - \gamma)p(t_p|t_e) + \gamma p(t_p|\text{lapse}), \quad (20)$$

where γ represents the lapse rate. With this formulation, we could identify lapse trials as those for which the likelihood of lapse exceeded the likelihood of a nonlapse. To limit cases of falsely identified lapse trials, we set the width of this uniform distribution conservatively to the range of possible production intervals (between 0 and 2000 ms).

Using simulations, we verified that our model was able to detect lapses for the range of w_p , w_p , and γ values inferred from the behavior of our subject pool. Most subjects had a small probability of a lapse trial that was consistent with previous reports²⁴. Two subjects had relatively unreliable performance with a larger number of lapse trials. However, our conclusions do not depend on the inclusion of these two subjects.

Analysis and model fitting. All analyses were performed using MATLAB R2014b or MATLAB R2017a, The MathWorks, Inc., Natick, Massachusetts, United States. We used a predictive maximum likelihood procedure to fit each model to the data. The likelihood of t_p given t_s and a set of parameters Θ (specific to each model) was defined as:

$$p(t_p|t_s, \Theta) = \int p(t_p, t_m|t_s, \Theta) dt_m, \quad (21)$$

For 1-2-Go trials and

$$p(t_p|t_s, \Theta) = \int \int p(t_p, t_{m_1}, t_{m_2}|t_s, \Theta) dt_{m_1} dt_{m_2}, \quad (22)$$

For 1-2-3-Go trials. The integrand for Equation 22 is

$$p(t_p, t_{m_1}, t_{m_2}|t_s, \Theta) = \lambda(t_{m_1}|t_s, \Theta) \lambda(t_{m_2}|t_s, \Theta) p(t_p|f_X[t_{m_1}, t_{m_2}, \Theta], \Theta), \quad (23)$$

With the output of $f_X[t_{m_1}, t_{m_2}, \Theta]$ corresponding to the equivalent mapping function for f_{BLS} , f_{LNE} , f_{EKF} , or $f_{BLS_{mem}}$. $\lambda(t_{m_i}|t_s, \Theta)$ corresponds to the likelihood function as defined above, with its dependence on the parameters Θ made explicit.

Assuming that production intervals are conditionally independent across trials, the log likelihood of model parameters can be formulated as:

$$\log_{\text{el}} p(t_p^1, t_p^2, \dots, t_p^N|t_s, \Theta) = \sum_{i=1}^N \log_{\text{el}} p(t_p^i|t_s, \Theta), \quad (24)$$

where the superscripts denote trial number. Maximum likelihood fits were derived from N-100 trials and cross validated on the remaining 100 trials (leave N out cross validation, LNOCV). This process was performed iteratively until all the data was fit. The final model parameters were taken as the average of parameter values across all the fits to the data. Fits were robust to changes in the amount of left out data. See Supplementary Figs 1, 8, and 10 for a summary of maximum likelihood parameters and predictions of each model fit to our subjects. We also computed the maximum likelihood parameters using the full data set for each subject. Parameters found using either the full data set or LNOCV were nearly identical (see Supplementary Table 3).

We evaluated model fits by generating simulated data from that model and comparing various summary statistics (BIAS, $\sqrt{\text{VAR}}$, and RMSE) observed for each subject to those generated by model simulations. For the observed data, summary statistics were computed for non-lapse trials and after removing the offset (b). Model

simulations were performed without the lapse term and after setting the offset to zero. The summary statistics were computed as follows:

$$\text{BIAS}^2 = \frac{1}{N} \sum_{i=1}^N (\bar{t}_{p_i} - t_{s_i})^2, \quad (25)$$

$$\text{VAR} = \frac{1}{N} \sum_{i=1}^N \sigma_i^2, \quad (26)$$

$$\text{RMSE} = \sqrt{\text{BIAS}^2 + \text{VAR}}, \quad (27)$$

BIAS² and VAR represent the average squared bias and average variance over the N distinct t_s 's of the prior distribution. The terms $\bar{t}_{p_i}, \sigma_i^2$ represent the mean and variance of production intervals for the i -th sample interval (t_s). The overall RMSE was computed as the square root of the sum of BIAS² and VAR. To find the BIAS² and VAR of each model we took the mean value of each after 1000 simulations of the model with the trial number matched to each subject. This ensured an accurate estimate of these quantities that includes the systematic deviations from the true model behavior due to a finite number of trials.

To perform model comparison, we measured the likelihood of the data, given a model and the maximum likelihood model parameters fit to training data, $\mathcal{L}(t_s, t_p | \mathcal{M}_i, \Theta^{\text{ML}})$. We then computed the ratio $\mathcal{L}(t_s, t_p | \mathcal{M}_i, \Theta^{\text{ML}})$ and $\mathcal{L}(t_s, t_p | \mathcal{M}_{\text{BLS}}, \Theta^{\text{ML}})$, the likelihood of the BLS model, and computed the logarithm of that value to measure the log likelihood ratio. To generate confidence intervals, we evaluated the likelihoods using 100 trials of test data that were left out of model fitting. We iterated this process until all the data was used as training data, allowing us to measure the variability of the log likelihood ratio for each subject.

Finally, we further validated our fitting procedure, analyses, and model comparison on data simulated using a generative process that emulates each model. Using this as ground truth data, we confirmed that our analyses identify the correct model and parameters using similar numbers of trials and subjects (Supplementary Figs 7, 9, and 11).

References

1. Wolpert, D. M., Ghahramani, Z. & Jordan, M. I. An internal model for sensorimotor integration. *Sci.* **269**, 1880–1882 (1995).
2. Ariff, G., Donchin, O., Nanayakkara, T. & Shadmehr, R. A real-time state predictor in motor control: study of saccadic eye movements during unseen reaching movements. *J. Neurosci.* **22**, 7721–7729 (2002).
3. Shadmehr, R., Smith, M. A. & Krakauer, J. W. Error correction, sensory prediction, and adaptation in motor control. *Annu. Rev. Neurosci.* **33**, 89–108 (2010).
4. Franklin, D. W. & Wolpert, D. M. Computational mechanisms of sensorimotor control. *Neuron* **72**, 425–442 (2011).
5. Todorov, E. & Jordan, M. I. Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.* **5**, 1226–1235 (2002).
6. Landy, M. S., Maloney, L. T., Johnston, E. B. & Young, M. Measurement and modeling of depth cue combination - in defense of weak fusion. *Vis. Res.* **35**, 389–412 (1995).
7. Mamassian, P. & Landy, M. S. Observer biases in the 3D interpretation of line drawings. *Vis. Res.* **38**, 2817–2832 (1998).
8. van Beers, R. J., Sittig, A. C. & Gon, J. J. Integration of proprioceptive and visual position-information: An experimentally supported model. *J. Neurophysiol.* **81**, 1355–1364 (1999).
9. Ernst, M. O. & Banks, M. S. Humans integrate visual and haptic information in a statistically optimal fashion. *Nat.* **415**, 429–433 (2002).
10. Battaglia, P. W., Jacobs, R. A. & Aslin, R. N. Bayesian integration of visual and auditory signals for spatial localization. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **20**, 1391–1397 (2003).
11. Körding, K. P. & Wolpert, D. M. Bayesian integration in sensorimotor learning. *Nat.* **427**, 244–247 (2004).
12. Adams, W. J., Graf, E. W. & Ernst, M. O. Experience can change the 'light-from-above' prior. *Nat. Neurosci.* **7**, 1057–1058 (2004).
13. Alais, D. & Burr, D. The ventriloquist effect results from near-optimal bimodal integration. *Curr. Biol.* **14**, 257–262 (2004).
14. Oruç, I., Maloney, L. T. & Landy, M. S. Weighted linear cue combination with possibly correlated error. *Vis. Res.* **43**, 2451–2468 (2003).
15. Fetsch, C. R., Pouget, A., DeAngelis, G. C. & Angelaki, D. E. Neural correlates of reliability-based cue weighting during multisensory integration. *Nat. Neurosci.* **15**, 146–154 (2012).
16. Knill, D. C. & Richards, W. Perception as Bayesian Inference (Cambridge University Press, 1996).
17. Blackwell, D. A. & Girshick, M. A. Theory of Games and Statistical Decisions (Courier Corporation, 1979).
18. Stocker, A. A. & Simoncelli, E. P. Noise characteristics and prior expectations in human visual speed perception. *Nat. Neurosci.* **9**, 578–585 (2006).
19. Narain, D., van Beers, R. J., Smeets, J. B. J. & Brenner, E. Sensorimotor priors in nonstationary environments. *J. Neurophysiol.* **109**, 1259–1267 (2013).
20. Brainard, D. H. *et al.* Bayesian model of human color constancy. *J. Vis.* **6**, 1267–1281 (2006).
21. Maloney, L. T. & Mamassian, P. Bayesian decision theory as a model of human visual perception: testing bayesian transfer. *Vis. Neurosci.* **26**, 147–155 (2009).
22. Ma, W. J. & Jazayeri, M. Neural coding of uncertainty and probability. *Annu. Rev. Neurosci.* **37**, 205–220 (2014).
23. Raphan, M., Simoncelli, E. P., Scholkopf, B., Platt, J. & Hoffman, T. Learning to be bayesian without supervision. In *Neural Information Processing Systems*, 1145–1152 (2006).
24. Jazayeri, M. & Shadlen, M. N. Temporal context calibrates interval timing. *Nat. Neurosci.* **13**, 1020–1026 (2010).
25. Simoncelli, E. P. Optimal estimation in sensory systems. *The Cogn. Neurosci.* **IV**, 525–535 (2009).
26. Kalman, R. E. A new approach to linear filtering and prediction problems. *J. Basic Eng* **82**, 35–45 (1960).
27. Simon, H. A. A behavioral model of rational choice. *Q. J. Econ.* **69**, 99–118 (1955).
28. Griffiths, T. L., Vul, E. & Sanborn, A. N. Bridging levels of analysis for probabilistic models of cognition. *Curr. Dir. Psychol. Sci.* **21**, 263–268 (2012).
29. Griffiths, T. L., Lieder, F. & Goodman, N. D. Rational use of cognitive resources: levels of analysis between the computational and the algorithmic. *Top. Cogn. Sci.* **7**, 217–229 (2015).
30. Marr, D. Vision: A computational approach (1982).

31. Keele, S. W., Roberto, N., Ivry, R. I. & Pokorny, R. A. Mechanisms of perceptual timing: Beat-based or interval-based judgements? *Psychol. Res.* **50**, 251–256 (1989).
32. Schulze, H. H. The perception of temporal deviations in isochronic patterns. *Percept. Psychophys.* **45**, 291–296 (1989).
33. Drake, C. & Botte, M. C. Tempo sensitivity in auditory sequences: evidence for a multiple-look model. *Percept. Psychophys.* **54**, 277–286 (1993).
34. Ivry, R. B. & Hazeltine, R. E. Perception and production of temporal intervals across a range of durations: evidence for a common timing mechanism. *J. Exp. Psychol. Hum. Percept. Perform.* **21**, 3–18 (1995).
35. Burr, D., Banks, M. S. & Morrone, M. C. Auditory dominance over vision in the perception of interval duration. *Exp. Brain Res.* **198**, 49–57 (2009).
36. Ogden, R. S. & Jones, L. A. More is still not better: testing the perturbation model of temporal reference memory across different modalities and tasks. *Q. J. Exp. Psychol.* **62**, 909–924 (2009).
37. Hartcher-O'Brien, J., Di Luca, M. & Ernst, M. O. The duration of uncertain times: audiovisual information about intervals is integrated in a statistically optimal fashion. *PLoS One* **9**, e89339 (2014).
38. Cai, M. B. & Eagleman, D. M. Duration estimates within a modality are integrated sub-optimally. *Front. Psychol.* **6** (2015).
39. Shi, Z., Ganzenmüller, S. & Müller, H. J. Reducing bias in auditory duration reproduction by integrating the reproduced signal. *PLoS One* **8**, e62065 (2013).
40. Di Luca, M. & Rhodes, D. Optimal perceived timing: Integrating sensory information with dynamically updated expectations. *Sci. Rep.* **6**, 28563 (2016).
41. Gallistel, C. R. Mental magnitudes. In *Space, Time and Number in the Brain*, 3–12 (Elsevier, 2011).
42. Gallistel, C. R. & Gibbon, J. Time, rate, and conditioning. *Psychol. Rev.* **107**, 289–344 (2000).
43. Merchant, H., Zarco, W. & Prado, L. Do we have a common mechanism for measuring time in the hundreds of millisecond range? evidence from multiple-interval timing tasks. *J. Neurophysiol.* **99**, 939–949 (2008).
44. Rakitin, B. C. *et al.* Scalar expectancy theory and peak-interval timing in humans. *J. Exp. Psychol. Anim. Behav. Process.* **24**, 15–33 (1998).
45. Gibbon, J. Scalar expectancy theory and weber's law in animal timing. *Psychol. Rev.* **84**, 279–325 (1977).
46. Getty, D. J. Discrimination of short temporal intervals: A comparison of two models. *Percept. Psychophys.* **18**, 1–8 (1975).
47. Acerbi, L., Wolpert, D. M. & Vijayakumar, S. Internal representations of temporal statistics and feedback calibrate motor-sensory interval timing. *PLoS Comput. Biol.* **8**, e1002771 (2012).
48. Miyazaki, M., Nozaki, D. & Nakajima, Y. Testing bayesian models of human coincidence timing. *J. Neurophysiol.* **94**, 395–399 (2005).
49. Cicchini, G. M., Arrighi, R., Cecchetti, L., Giusti, M. & Burr, D. C. Optimal encoding of interval timing in expert percussionists. *J. Neurosci.* **32**, 1056–1060 (2012).
50. Miller, N. S. & McAuley, J. D. Tempo sensitivity in isochronous tone sequences: the multiple-look model revisited. *Percept. Psychophys.* **67**, 1150–1160 (2005).
51. Elliott, M. T., Wing, A. M. & Welchman, A. E. Moving in time: Bayesian causal inference explains movement coordination to auditory beats. *Proc. Royal Soc. Lond. B: Biol. Sci.* **281**, 20140751 (2014).
52. Knill, D. C. & Pouget, A. The bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci.* **27**, 712–719 (2004).
53. Michon, J. A. *Timing in temporal tracking* (Institute for Perception RVO-TNO Soesterberg, The Netherlands, 1967).
54. Mates, J. A model of synchronization of motor acts to a stimulus sequence I. timing and error corrections. *Biol. Cybern.* **70**, 463–473 (1994).
55. Mates, J. A model of synchronization of motor acts to a stimulus sequence II. stability analysis, error estimation and simulations. *Biol. Cybern.* **70**, 475–484 (1994).
56. Pressing, J. Error correction processes in temporal pattern production. *J. Math. Psychol.* **42**, 63–101 (1998).
57. Semjen, A., Schulze, H. H. & Vorberg, D. Timing precision in continuation and synchronization tapping. *Psychol. Res.* **63**, 137–147 (2000).
58. Repp, B. H. Sensorimotor synchronization: a review of the tapping literature. *Psychon. Bull. Rev.* **12**, 969–992 (2005).
59. Schulze, H.-H. & Vorberg, D. Linear phase correction models for synchronization: Parameter identification and estimation of parameters. *Brain Cogn.* **48**, 80–97 (2002).
60. Vorberg, D. & Schulze, H.-H. Linear Phase-Correction in synchronization: Predictions, parameter estimation, and simulations. *J. Math. Psychol.* **46**, 56–87 (2002).
61. Stengel, R. F. *Optimal control and estimation* (Dover Publications Inc., 1994).
62. Bowers, J. S. & Davis, C. J. Bayesian just-so stories in psychology and neuroscience. *Psychol. Bull.* **138**, 389–414 (2012).
63. Shi, Z. & Burr, D. Predictive coding of multisensory timing. *Curr Opin Behav Sci* **8**, 200–206 (2016).
64. Wing, A. M. & Kristofferson, A. B. Response delays and the timing of discrete motor responses. *Percept. Psychophys.* **14**, 5–12 (1973).
65. Hary, D. & Moore, G. P. Synchronizing human movement with an external clock source. *Biol. Cybern.* **56**, 305–311 (1987).
66. Schulze, H.-H. The error correction model for the tracking of a random metronome: Statistical properties and an empirical test. In Macar, F., Pouthas, V. & Friedman, W. J. (eds) *Time, Action and Cognition: Towards Bridging the Gap*, 275–286 (Springer Netherlands, Dordrecht, 1992).
67. Repp, B. H. Phase correction following a perturbation in sensorimotor synchronization depends on sensory information. *J. Mot. Behav.* **34**, 291–298 (2002).
68. Barnes, R. & Jones, M. R. Expectancy, attention, and time. *Cogn. Psychol.* **41**, 254–311 (2000).
69. Taatgen, N. & van Rijn, H. Traces of times past: representations of temporal intervals in memory. *Mem. Cogn.* **39**, 1546–1560 (2011).
70. Burr, D., Della Rocca, E., Rocca, E. D. & Morrone, M. C. Contextual effects in interval-duration judgements in vision, audition and touch. *Exp. Brain Res.* **230**, 87–98 (2013).
71. Rao, R. P. N. & Ballard, D. H. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* **2**, 79–87 (1999).
72. Vaziri, S., Diedrichsen, J. & Shadmehr, R. Why does the brain predict sensory consequences of oculomotor commands? optimal integration of the predicted and the actual sensory feedback. *J. Neurosci.* **26**, 4188–4197 (2006).
73. Stevenson, I. H., Fernandes, H. L., Vilares, L., Wei, K. & Körding, K. P. Bayesian integration and non-linear feedback control in a full-body motor task. *PLoS Comput. Biol.* **5**, e1000629 (2009).
74. Friston, K. The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.* **11**, 127–138 (2010).
75. Merchant, H., Pérez, O., Zarco, W. & Gámez, J. Interval tuning in the primate medial premotor cortex as a general timing mechanism. *J. Neurosci.* **33**, 9082–9096 (2013).
76. Wang, J., Narain, D., Hosseini, E. A. & Jazayeri, M. Flexible timing by temporal scaling of cortical responses. *Nat. Neurosci.* **21**, 102–110 (2018).
77. Merchant, H. & Averbeck, B. B. The computational and neural basis of rhythmic timing in medial premotor cortex. *J. Neurosci.* **37**, 4552–4564 (2017).
78. Komura, Y. *et al.* Retrospective and prospective coding for predicted reward in the sensory thalamus. *Nat.* **412**, 546–549 (2001).
79. Mello, G. B. M., Soares, S. & Paton, J. J. A scalable population code for time in the striatum. *Curr. Biol.* **25**, 1113–1122 (2015).
80. Xu, M., Zhang, S.-Y., Dan, Y. & Poo, M.-M. Representation of interval timing by temporally scalable firing patterns in rat prefrontal cortex. *Proc. Natl. Acad. Sci.* **111**, 480–485 (2014).

81. Petzschner, F. H. & Glasauer, S. Iterative bayesian estimation as an explanation for range and regression effects: A study on human path integration. *J. Neurosci.* **31**, 17220–17229 (2011).
82. Hudson, T. E., Maloney, L. T. & Landy, M. S. Optimal compensation for temporal uncertainty in movement planning. *PLoS Comput. Biol.* **4**, e1000130 (2008).
83. Kwon, O.-S. & Knill, D. C. The brain uses adaptive internal models of scene statistics for sensorimotor estimation and planning. *Proc. Natl. Acad. Sci. USA* **110**, E1064–73 (2013).
84. Gilden, D. L., Thornton, T. & Mallon, M. W. 1/f noise in human cognition. *Sci.* **267**, 1837–1839 (1995).
85. Farrell, S., Wagenmakers, E.-J. & Ratcliff, R. 1/f noise in human cognition: is it ubiquitous, and what does it mean? *Psychon. Bull. Rev.* **13**, 737–741 (2006).
86. Körding, K. P. & Wolpert, D. M. The loss function of sensorimotor learning. *Proc. Natl. Acad. Sci. USA* **101**, 9839–9842 (2004).
87. Acerbi, L., Vijayakumar, S. & Wolpert, D. M. On the origins of suboptimality in human probabilistic inference. *PLoS Comput. Biol.* **10**, e1003661 (2014).
88. Drugowitsch, J., Wyart, V., Devauchelle, A.-D. & Koechlin, E. Computational precision of mental inference as critical source of human choice suboptimality. *Neuron* (2016).
89. Ma, W. J., Husain, M. & Bays, P. M. Changing concepts of working memory. *Nat. Neurosci.* **17**, 347–356 (2014).

Acknowledgements

We would like to thank Christopher Stawarz for his help in programing the psychophysical paradigm. We also thank Evan Remington, Jing Wang, Hansem Sohn, Devika Narain, and Eghbal Hosseini for their helpful discussions of experimental results. Finally, we would like to thank Joshua Tenenbaum and Evan Remington for helpful comments on an earlier versions of this manuscript, and Rossana Chung for her editorial assistance.

Author Contributions

S.W.E. and M.J. conceived of and designed the tasks. S.W.E. performed the experiments and analyzed the data. S.W.E. and M.J. interpreted the results and wrote the paper.

Additional Information

Supplementary information accompanies this paper at <https://doi.org/10.1038/s41598-018-30722-0>.

Competing Interests: The authors declare no competing interests.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2018