Title
Computation-Limited Bayesian Updating

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

Authors
Zhu, Jian-Qiao
Sanborn, Adam
Chater, Nick
et al.

Publication Date
2023

Peer reviewed

Computation-Limited Bayesian Updating

Jian-Qiao Zhu (jz5204@princeton.edu)
Department of Computer Science
Princeton University

Adam N. Sanborn (a.n.sanborn@warwick.ac.uk)
Department of Psychology
University of Warwick

Nick Chater (nick.chater@wbs.ac.uk)
Warwick Business School
University of Warwick

Thomas L. Griffiths (tomg@princeton.edu)
Department of Psychology and Computer Science
Princeton University

Abstract

Effectively updating one's beliefs requires sufficient empirical evidence (i.e., data) and the computational capacity to process it. Yet both data and computational resources are limited for human minds. Here, we study the problem of belief updating under limited data and limited computation. Using information theory to characterize constraints on computation, we find that the solution to the resulting optimization problem links the data and computational limitations together: when computational resources are tight, agents may not be able to integrate new empirical evidence. The resource-rational belief updating rule we identify offers a novel interpretation of conservative Bayesian updating.

Keywords: Belief Updating; Bayes rule; Resource Rationality; Information Theory

Introduction

Intelligent agents need to be able to update their beliefs about the world based on empirical evidence. There are two fundamental challenges involved in this process: the empirical evidence is often based upon *limited data*, and it needs to be perceived and processed using *limited computational resources*. Each of these limitations has significant implications for belief updating.

Limited data means that even an ideal agent will have beliefs that incorporate uncertainty. Under reasonable assumptions, the optimal method of belief updating is Bayes' rule (Jaynes, 2003), which indicates how a probability distribution over hypotheses (the "prior" distribution) should be updated based on new data (becoming the "posterior" distribution). If the observed data are consistent with multiple hypotheses that have non-zero prior probability, these hypotheses will also have non-zero posterior probability: the posterior will reflect *data uncertainty*, which is irreducible even with more computation. For example, even after tossing three heads in a row, a coin is still likely to be fair; it feels wrong to have strong beliefs about the bias of a coin based on a few observations.

Limited computation can magnify this uncertainty. The exact belief-updating solution provided by Bayes' rule can require significant amounts of computation, becoming infeasible in realistic settings (e.g., Koller & Friedman, 2009). It has thus been proposed that people instead use an affordable amount of computation to obtain an approximation to the posterior (Sanborn & Chater, 2016; Griffiths, Vul, & Sanborn, 2012). This possibility has been explored using sampling- or optimization-based approximation algorithms, which have been shown to correspond well with human behaviors in a variety of cognitive domains including probability judgments (J.-Q. Zhu, Sanborn, & Chater, 2020; Dasgupta, Schulz, & Gershman, 2017), belief updating (Dasgupta, Schulz, Tenenbaum, & Gershman, 2020; Prat-Carrabin, Wilson, Cohen, & Azeredo da Silveira, 2021), risky choices (Vul, Goodman, Griffiths, & Tenenbaum, 2014), and causal inference (Bramley, Dayan, Griffiths, & Lagnado, 2017). These approximate solutions to Bayesian inference introduce errors that result in *computational uncertainty*, adding to the uncertainty in the resulting beliefs.

In this paper, we consider how a rational agent should manage these sources of uncertainty. In the spirit of resource-rationality (Griffiths et al., 2012; Lieder & Griffiths, 2020), we define an optimization problem in which we quantify the cost of computation and the benefits of following the Bayesian updating rule using information theory. The optimal solution, called *computation-limited Bayesian updating*, is then derived. We compare the predictions from the optimal updating rule with the behavioral data curated in the meta-analysis of human belief updating conducted by Benjamin (2019). Manipulating the computational cost also leads to updating behaviors that are more Bayesian with more efficient compute, suggesting a computational account of cognitive development in children (L. Zhu & Gigerenzer, 2006; Girotto & Gonzalez, 2007).

Resource-rational updating rules

We consider an agent that updates her beliefs by incorporating multiple sources of knowledge (e.g., by combining prior beliefs with data, by combining historical and present data, or by combining personal and social knowledge). Our focus is on the computational challenges that arise in this process, rather than on the specific sources of knowledge involved. However, we believe that the proposed solution can be applied to a variety of belief-updating situations in which knowledge comes from different sources.

To illustrate (see Figure 1), we consider a simplified scenario that has received extensive experimental investigation: an agent has to revise her prior beliefs (denoted q_{t-1}) in light of the knowledge from newly available data (denoted d_t) where t denotes the temporal order of the beliefs and the data. We adopt a probabilistic interpretation of knowledge. In this regard, $q_{t-1}(\theta)$ represents the subjective beliefs

In M. Goldwater, F. K. Anggoro, B. K. Hayes, & D. C. Ong (Eds.), *Proceedings of the 45th Annual Conference of the Cognitive Science Society*.

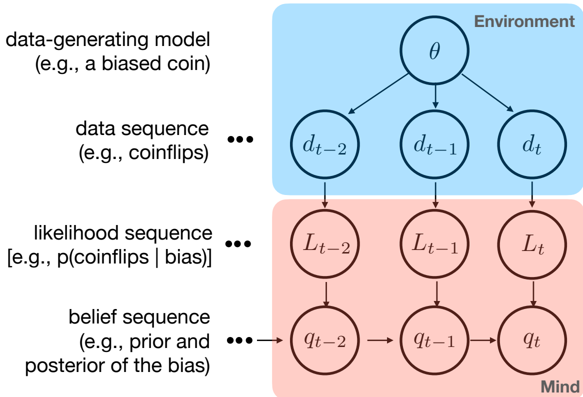

Figure 1: A schematic illustration of the human belief updating process. People are asked to update their beliefs (q) about the unobservable state (θ). Information about the true value of θ can only be accessed through the data (d) that is generated by the data-generating model. t denotes the temporal order of the dataset, likelihoods, and beliefs where q_t is calculated based on the likelihood (L_t) and the prior (q_{t-1}).

about θ (e.g., θ could represent the bias of a coin or any other objects that people can assign credence to), expressed as a probability distribution over θ. The "subjective" knowledge from the new data, d_t, is described via the likelihood function $L_t(\theta) = p(d_t|\theta)$. A belief-updating rule specifies a way to calculate the new beliefs $q_t(\theta)$ based on $q_{t-1}(\theta)$ and $L_t(\theta)$.

Deriving optimal updating rules

The standard belief-updating rule is Bayes' rule, which indicates that the posterior distribution (denoted $p_t(\theta|d_t)$) should combine prior and likelihood as follows:

$$p_t(\theta|d_t) = \frac{p(d_t|\theta)q_{t-1}(\theta)}{p(d_t)} \tag{1}$$

$$\propto L_t(\theta)q_{t-1}(\theta). \tag{2}$$

Our agent could thus simply take $q_t(\theta) = p_t(\theta)$, satisfying the various optimality criteria that justify Bayesian inference. However, the computational challenges involved in applying Bayes' rule can make it infeasible for agents with limited computational resources.

Various methods for approximating Bayesian inference have been proposed by theorists as an alternative for understanding human behaviors (Griffiths et al., 2012; Vul et al., 2014; J.-Q. Zhu, León-Villagrá, Chater, & Sanborn, 2022; Dasgupta et al., 2020; Sanborn & Chater, 2016), but there is yet to be a theoretical investigation on the optimal updating rule under computational constraints. In this section, we formally develop the idea of a computation-limited agent by performing a resource-rational analysis of belief updating where people are assumed to best approximate the Bayes-optimal posterior under a limited computational budget (c.f. Lieder & Griffiths, 2020).

To derive an optimal updating rule, we need to define the agent's goal in revising old beliefs. We assume that the agent tries to revise her belief towards the Bayes-optimal posterior as close as possible, subject to the constraint that too drastic changes in beliefs are too computationally costly.

Consider the Bayes-optimal posterior, p_t, that the agent attempts to approximate. We can model the approximation process as the process of minimizing the information-theoretic distance between the belief state and the posterior:

$$\min_{q_t} D_{KL}(q_t||p_t) \tag{3}$$

where q_t is the new belief state, p_t is the Bayes-optimal posterior (i.e., calculated as in Equation 2 with q_{t-1} as the prior), and $D_{KL}(a||b)$ denotes the Kullback–Leibler divergence between two probability distributions a and b, which is a type of statistical distance that measures how a is different from b. The KL-divergence can be intuitively interpreted as the expected surprise or dissimilarity from using b as a model when the actual distribution should be a.

Next, we define the computational cost (C) of changing belief states from q_{t-1} to q_t as follows:

$$C(q_t) = \lambda D_{KL}(q_t||q_{t-1}) \tag{4}$$

where λ is the conversion factor between one unit of information and one unit of compute for approximation algorithms. In other words, to change an agent's belief state by one unit (as measured in the information-theoretic distance), it will cost λ. The parameter λ is also conceptually similar to the conversion factor between one unit of information and one unit of energy in thermodynamically inspired formalization of bounded rationality (Ortega & Braun, 2013). Similar perspectives can also be found in the reinforcement learning literature (Todorov, 2006; Ho, Abel, Cohen, Littman, & Griffiths, 2020) and the rational inattention model in economics (Sims, 2003).

Assuming larger changes in belief states will require greater cost, we can now combine the cost function with the approximation process, resulting in the following optimization problem:

$$\min_{q_t} D_{KL}(q_t||p_t) + C(q_t) \tag{5}$$

Theoretically, the optimization problem stated in Equation 5 is one in which an agent is trying to find the best balance between achieving a high level of accuracy in their approximation while also minimizing the cost of making those approximations. Many approximation algorithms can be characterized in this way. For example, sampling-based approximation algorithms such as Markov Chain Monte Carlo generate samples to approximate the Bayes-optimal posterior, but each sample is generated at a cost. Similarly, optimization-based approximation algorithms such as variational inference may improve the accuracy of parametric models through a costly process of gradient descent. In practice, the efficiency of the approximation algorithm can impact the conversion factor λ and thus the computational cost C, with more efficient algorithms having lower values. For more information on how the cost associated with sampling-based and optimization-based approximation algorithms, refer to Appendix A.

Solving the optimization problem in Equation 5, we show that the optimal update rule is

$$q_t^*(\theta) = \frac{1}{\mathcal{Z}} p_t(\theta)^{1/(1+\lambda^*)} q_{t-1}(\theta)^{\lambda^*/(1+\lambda^*)} \quad (6)$$

where $\mathcal{Z} = \int_\theta p_t(\theta)^{1/(1+\lambda^*)} q_{t-1}(\theta)^{\lambda^*/(1+\lambda^*)} d\theta$ is the normalizing constant. Detailed derivations for this rule can be found in the Appendix B. Moreover, because p_t is the Bayes-optimal posterior, which is proportional to the product of prior q_{t-1} and likelhood L_t, Equation 6 can be simplified to

$$q_t^*(\theta) \propto \left[q_{t-1}(\theta) L_t(\theta)\right]^{1/(1+\lambda^*)} q_{t-1}(\theta)^{\lambda^*/(1+\lambda^*)} \quad (7)$$

$$= L_t(\theta)^{1/(1+\lambda^*)} q_{t-1}(\theta) \quad (8)$$

which precisely localizes the impact of computational limitations to the processing of the data, as reflected in the exponent of the likelihood term. Since $\lambda \geq 0$, this exponent is less than or equal to 1. Consequently, computational limitations should be expected to decrease the extent to which agents are sensitive to empirical evidence.

Summary

Our computation-limited Bayesian updating model is motivated by an optimization problem where we use information theory to characterize the costs and benefits of accurate updating. When computational capacity is unbounded, standard Bayesian updating becomes the optimal information processing rule under our framework. Relaxing the psychologically implausible assumption of infinite computational resources, the optimal belief-updating rule can deviate from Bayes' rule and make more conservative use of data. Therefore, we should expect deviations from the Bayes' rule in the direction of decreased sensitivity to empirical evidence, and for these deviations to be exaggerated in settings where agents face greater computational constraints.

Behavioral signatures of computation-limited Bayesian updating

We applied the theory developed in the previous section to a recent meta-analysis of belief updating conducted by Benjamin (2019), which reviewed the vast empirical literature of belief updating using the classical *bookbag-and-poker-chip* experiments (e.g., Grether, 1992; Holt & Smith, 2009; Barron, 2021; Buser, Gerhards, & Van der Weele, 2016; Coutts, 2019; Gotthard-Real, 2017; Möbius, Niederle, Niehaus, & Rosenblat, 2022; Charness & Dave, 2017; Beach, Wise, & Barclay, 1970; Dave & Wolfe, 2003; Kraemer & Weber, 2004; Sasaki & Kawagoe, 2007). A typical experiment involves drawing poker chips (or balls) of different colors out of bookbags (or urns) in front of people. Those people are then asked to report their beliefs about which bookbag (or urn) was selected to generate the chips (or balls).

Our analyses center on the sequential updating problem within the context of bookbag-and-poker-chip experiments, wherein participants are required to report their posterior estimates sequentially as new data becomes accessible. Within this framework, the self-reported posterior from the preceding trial can be regarded as an appropriate candidate for the prior of the current trial. Therefore, sequential updating problems offer the most unambiguous definitions of the prior and posterior probabilities, facilitating the robust evaluation of our theoretical model.

More formally, these experiments are binomial updating problems where people have two discrete hypotheses about the binomial parameter θ. For example, the two hypotheses might be $A = \{\text{red-dominated bag}\}$ and $B = \{\text{blue-dominated bag}\}$. Using θ to denote the probability of sampling a red poker chip, under hypothesis A θ takes a value θ_A greater than 0.5. Under hypothesis B θ takes the value $\theta_B = 1 - \theta_A$ in a symmetric updating problem. The prior probabilities of the hypotheses A and B are typically $p(A) = 1 - p(B) = 0.5$. When red or blue chips are randomly drawn from the chosen bag with replacement, the likelihoods can be calculated using the binomial distribution with the appropriate parameter θ.

Meta-analyzing the use of likelihoods and priors

In short, the computation-limited Bayesian updating rule obtained in Equation 8 suggests that an agent should revise her beliefs by raising the likelihood to a power. If we consider the following generic updating rule:

$$q_t^* \propto L_t^\alpha q_{t-1}^\beta, \quad (9)$$

the computation-limited Bayesian updating then requires that the power of likelihoods, $\alpha = 1/(1+\lambda^*)$, should be in the range between 0 and 1, and that of priors should be 1, $\beta = 1$. Note that the standard Bayesian updating rule instead requires that $\alpha = \beta = 1$.

Experimenters analyzing the results of bookbag-and-poker-chip experiments often use linear models to quantify the contributions of likelihoods and priors in arriving at subjective posterior estimates. By taking natural logarithms on both sides of the generic updating rule (Equation 9), the multiplicative relationship between likelihood and prior becomes additive. When we have two hypotheses, we can rewrite this updating rule in terms of the log-odds, with

$$\ln \frac{q_t(A)}{q_t(B)} = \alpha \ln \frac{p(d_t|A)}{p(d_t|B)} + \beta \ln \frac{q_{t-1}(A)}{q_{t-1}(B)} + \varepsilon \quad (10)$$

where d_t is the new data made available in the period between $t-1$ and t, and ε is the residual of the model, which captures response error and is assumed to be distributed as a Gaussian with mean 0. We also assume that the posterior-odds that the participants reported at previous timestep $t-1$ become their

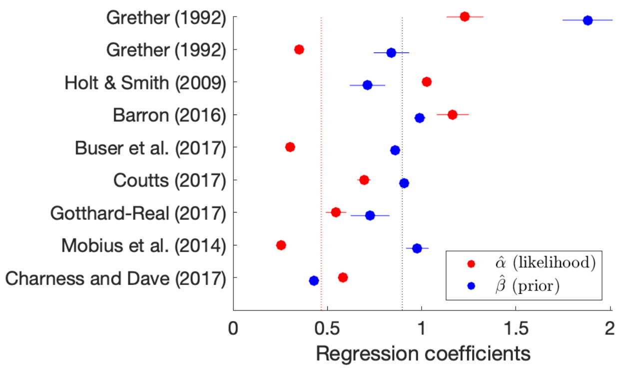

Figure 2: Meta-analyzing regression coefficients for likelihoods and priors (i.e., $\hat{\alpha}$ and $\hat{\beta}$). Regression results are drawn from 8 papers based on the regressions performed within individual paper. Error bars denote standard errors for regression coefficients. The data was curated by Benjamin (2019). Inverse-variance-weighted means for likelihoods and priors are 0.47 (red dashed line; 95%CI, $[0.46, 0.48]$) and 0.90 (blue dashed line; 95%CI, $[0.88, 0.91]$) respectively. Estimates from Charness and Dave (2017) did not count towards the meta-analyzed mean due to absence of standard errors.

prior-odds for the next update at t. In this linear model, α and β have clear behavioral interpretations in terms of the weight given to likelihoods and priors.

All papers in Figure 2 ran linear regressions similar to the linear model presented in Equation 10 except that some included an intercept term in their linear regressions (Charness & Dave, 2017; Grether, 1992). The detailed specifications of the linear regressions and their methods of estimation were summarized in the Online Appendix of Benjamin (2019). We reproduced their meta analysis in Figure 2. Using the inverse-variance-weighted mean calculation, the meta-analysed estimates of regression coefficients are that the $\hat{\alpha}$ estimate of likelihoods is 0.47 and the $\hat{\beta}$ estimate of priors is 0.90.[1] This result is consistent with a large body of prior findings that demonstrate conservatism in belief updating, downweighting likelihoods (e.g., Edwards, 1968).

Reanalyzing sequential updating problems

In addition to meta-analyzing past regression results whose linear models and estimation methods vary across studies, we also conducted our own linear regression (applying Equation 10) directly on behavioral data collected from people solving the sequential updating problems. As shown in Figure 3, the meta-analysis sample that reports subjective posterior-odds includes 127 observations drawn from four papers (Benjamin, 2019). All regressions were estimated using the ordinary least squares methods. The estimated coefficients $(\hat{\alpha}, \hat{\beta})$ and their standard errors are displayed in Figure 3, with meta-analyzed mean estimates equal to 0.33 and 0.95 for likelihoods and priors respectively.

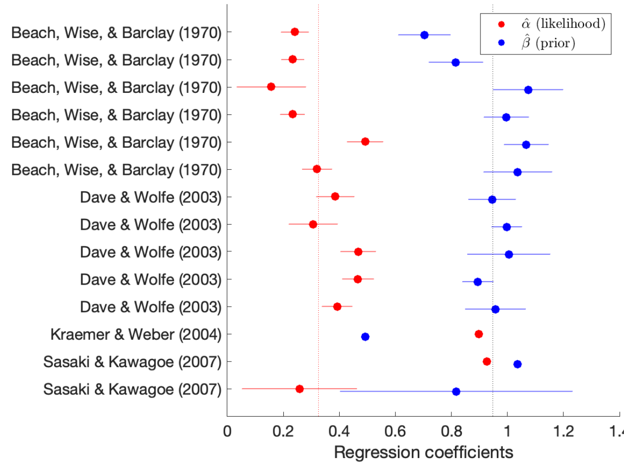

Figure 3: Reanalyzing sequential belief updating problems studied in four papers. Regression coefficients for likelihoods and priors (i.e., $\hat{\alpha}$ and $\hat{\beta}$) were computed using participants' mean posterior odds. Error bars denote standard errors for regression coefficients. The data was curated by Benjamin (2019). Inverse-variance-weighted means for likelihoods and priors are 0.33 (red dashed line; 95%CI, $[0.29, 0.36]$) and 0.95 (blue dashed line; 95%CI, $[0.90, 1.00]$) respectively. Estimates from Kraemer and Weber (2004) and one experiment of Sasaki and Kawagoe (2007) did not count towards the meta-analyzed mean due to missing standard errors.

This analysis provides further evidence that is consistent with conservative updating. The majority of α estimates are below 1, suggesting that people revised their beliefs less than requested by the likelihood. The β estimates, however, are closer to 1, indicating that priors or previous subjective posteriors are used in an almost unbiased manner to calculate the updated posterior. Our reanalysis and recent work (e.g., Powell, 2022) both point to the robust phenomenon of conservatism as the main form of deviations from Bayesian rationality in belief updating.

Conservatism as computation-limited updating

As discussed above, conservatism is a well-established property of people's belief updating. A range of explanations have been proposed in the literature for conservatism, including misperception of likelihoods, difficulty of aggregating different sources of information, and/or response bias (Phillips & Edwards, 1966; Slovic & Lichtenstein, 1971; Edwards, 1968). In more recent literature, conservatism has been rationalized in various ways. Under circumstances of higher-order uncertainty, such as when the correctness of the likelihood function is uncertain, maintaining a degree of conservatism may yield more precise posterior probabilities

[1] The inverse-variance-weighted mean for regression coefficient is $\beta_{combined} = (\sum_{i=1}^{M} \frac{\beta_i}{\sigma_i^2})/(\sum_{i=1}^{M} \frac{1}{\sigma_i^2})$, where M is the total number of regressions.

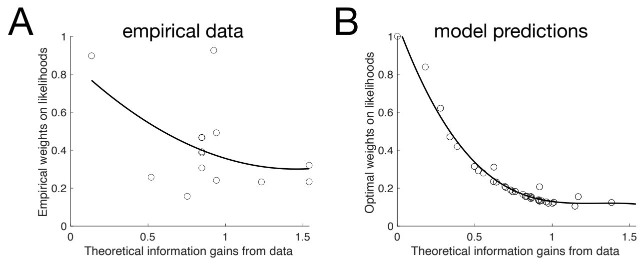

Figure 4: **(A)** Empirical and **(B)** predicted relationships between the amount of information gain from data and the weights on likelihoods. The empirical weights on likelihoods were adapted from Figure 3 and the information gains were calculated as mean absolute log-likelihood-ratios. The predicted optimal weight on likelihood were calculated with a fixed computational cost, $C = 0.1$. The solid lines represent the best-fitting polynomial curves of the highest degree 3 by minimizing least squares errors.

(Bramley, Lagnado, & Speekenbrink, 2015). Furthermore, in cases where the causal representation is not comprehensive (or where the world model of the agent becomes increasingly complex), deliberately moderating the likelihood, as exhibited in conservatism, can facilitate the relaxation of an otherwise inefficient optimization problem (Ramscar, Hendrix, Shaoul, Milin, & Baayen, 2014; Bramley et al., 2017).

Here, our computation-limited Bayesian updating model offers a resource-rational interpretation of this behavior that does not require appealing to any of these other factors. As shown in Equation 8, computational constraints naturally lead to a weaker-than-rational belief updating.

There are other regularities in human belief updating that are potentially consistent with our computation-limited model. For example, the impact of evidence (i.e., the value of α) becomes smaller when data diagnosticity is stronger (Phillips & Edwards, 1966) and when more data are presented (Peterson, Schneider, & Miller, 1965). Both stronger signal per datapoint and larger dataset result in greater information gains from data. Indeed, the empirical weights on likelihoods decrease as the information gains from the data increases (see Figure 4A). In these experiments, the data presented to participants was administered by randomly drawing poker chips (or balls) from a selected bookbag (or urn). Assuming a fixed computational cost, the computation-limited Bayesian updating rule also reproduces this negative relationship between the biased use of likelihood and the amount of information provided to participants (see Figure 4B).

Variation in computational cost

Our cognitive abilities develop as we grow older. While it is evident that adults can integrate different sources of information in probabilistic reasoning (Griffiths & Tenenbaum, 2006), the question of how we acquire such an ability during childhood remains unanswered. Here, we entertain the hypothesis that people approximate Bayesian updating in the way suggested by our computation-limited account, but that their computational cost gradually decreases as they age, resulting in a more efficient use of compute. As a consequence, we should expect a decrease in conservatism as people become more able to engage in the computations required by belief updating.

To test this prediction, we reanalyzed behavioral data collected by Girotto and Gonzalez (2007) and L. Zhu and Gigerenzer (2006), who studied changes in belief updating among 4- to 7-year-olds children and adults (see Figure 5A). Because limited information was provided to children in Girotto and Gonzalez (2007) whereas all relevant information to reach a Bayes-optimal posterior was provided in L. Zhu and Gigerenzer (2006), we simulated two different amounts of information gain from data (0.8 and 1.6 bits) while evenly varying the computational cost from 1 (preschoolers) to 0.1 (adults) among the 8 age groups (see Figure 5B). Our model qualitatively reproduces the developmental trajectories of Bayesian reasoning in children, which are the result of a decreasing computational cost.

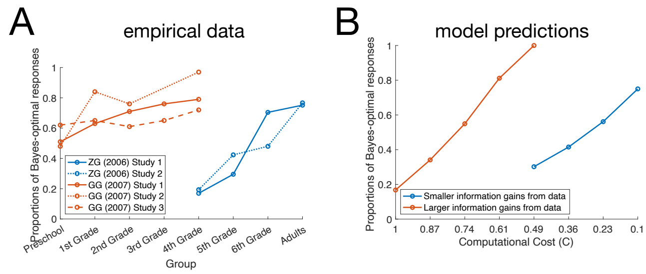

Figure 5: **(A)** Different age groups exhibit different belief-updating behaviours. Human data show an increase in the proportion of responses consistent with Bayesian updating with age (L. Zhu & Gigerenzer, 2006; Girotto & Gonzalez, 2007). **(B)** Decreasing computational cost with age predicts progressively better adherence to the Bayesian updating rule.

Discussion

Updating beliefs can be computationally costly. We introduced a new framework for understanding how these costs might be expected to affect human belief updating. By construing computational constraints as constraints on the degree to which beliefs can be changed by any new piece of data, the theory predicts that people should underweight evidence. Specifically, when compared to optimal Bayesian inference, people should discount the likelihood but make accurate use of the prior. We confirmed these predictions by reanalyzing experiments that asked people to sequentially update their beliefs. Our model also predicts that we should see greater discounting of evidence under stronger computational constraints, a prediction that is borne out by examining developmental data: the progressive increase in the correspondence of beliefs with Bayesian inference with age supports the idea that computational cost plays a central role in belief updating.

We also note a number of limitations in our theory. First, we made some convenient assumptions about the belief updating process, which may weaken the generalizability of the theory. In practice, priors and likelihoods may not necessarily match the ones we have assumed. Second, the KL-divergence may not be the correct quantity for measuring the disparity between two probability distributions. There are other more generalized information measures such as the Renyi entropy, which may better describe the information processed by the mind (Sajid et al., 2022).

The biggest advantage of our approach is that our theory subsumes a large variety of algorithmic-level models that engage with the computational-level problem of belief updating (c.f. Griffiths et al., 2012). These approximation algorithms include, but are not limited to, sampling-based methods such as Markov chain Monte Carlo and importance sampling, optimization-based methods such as variational Bayes, and other approximation techniques such as neural networks. As long as the degree of belief updating increases with the amount of computation, these approximation schemes as well as heuristics are expected to behave similarly to computation-limited Bayesian updating. To gain a deeper understanding of how humans update their beliefs in a way that unifies a variety of model predictions and behaviors, we must consider not only the cost of limited data but also the cost of limited computational resources.

Appendix A. Evaluating Approximation Algorithms

Evaluating approximation algorithms can involve considering both the rate at which the algorithm approaches the target distribution and the cost of each step of the algorithm. Here, we focus on the former metric, the convergence rate. The speed of convergence is often measured in total variation distance between the target distribution and its approximation, which is upper bounded by the KL-divergence. An Δ-accuracy means that the total variation distance is less than $\Delta \in [0,1]$ (Bishop, 2006). Therefore, a small value of Δ results in a more accurate approximation of the target distribution. Moreover, a faster reduction in Δ typically indicates a larger convergence rate for the approximation algorithm, thus a smaller conversion factor λ and a smaller cost of compute C in Equation 5.

In the convex setting where global properties of the target distribution can be assessed through local information, optimization-based approximation algorithms typically converge faster than sampling-based approximation algorithms (Barber, 2012; Bishop, 2006). Approximation algorithms that use the gradient information are expected to converge with a rate of $1/k$ where k is the number of gradient descents.

However, recent theoretical analyses show that in a class of problems that are convex outside of a bounded region but non-convex inside of it, sampling-based algorithms converge to Δ-accuracy proportional to R/Δ iterations, whereas any optimization-based algorithm converges in proportional to $(1/\Delta)^R$ iterations (Ma, Chen, Jin, Flammarion, & Jordan, 2019). R denotes the radius of the non-convex region, with a larger value indicating a more challenging optimization problem. Therefore, as greater computing power allows for a larger number of iterations, the accuracy of approximation algorithms improves.

Appendix B. Solving the Optimization Problem

Here, we provide a detailed mathematical solution to the optimization problem stated in Equation 5, deriving the result stated in Equation 6. For simplicity we will use abbreviated notation for integrals over probability distributions. For example, instead of writing the Kullback-Leibler divergence between q_t and p_t as $D_{KL}(q_t||p_t) = \int q_t(\theta) \ln \frac{q_t(\theta)}{p_t(\theta)} d\theta$ we will write it as $\int q_t \ln \frac{q_t}{p_t}$, with all integrals being over θ and all functions taking the argument θ implicitly.

We begin rewriting the optimization problem using the following format:

$$L(q_t, \lambda) = \int q_t \ln \frac{q_t}{p_t} + C(q_t) \tag{11}$$

$$= \int q_t \ln \frac{q_t}{p_t} + \lambda \int q_t \ln \frac{q_t}{q_{t-1}} \tag{12}$$

$$= \int q_t \ln \frac{q_t^{1+\lambda}}{p_t q_{t-1}^{\lambda}} \tag{13}$$

$$= (1+\lambda) \int q_t \ln \left(\frac{q_t}{p_t^{1/(1+\lambda)} q_{t-1}^{\lambda/(1+\lambda)}} \right) \tag{14}$$

Note that the optimization problem reduces to a KL divergence:

$$(1+\lambda) D_{KL}\left(q_t \middle|\middle| p_t^{1/(1+\lambda)} q_{t-1}^{\lambda/(1+\lambda)} \right) \tag{15}$$

which is minimized when

$$q_t^* = \frac{1}{\mathcal{Z}} p_t^{1/(1+\lambda)} q_{t-1}^{\lambda/(1+\lambda)} \tag{16}$$

where $\mathcal{Z}$ is the normalizing constant $\int p_t^{1/(1+\lambda)} q_{t-1}^{\lambda/(1+\lambda)}$.

Substituting q_t^* back, we obtain the following optimal value for the conversion factor λ:

$$\lambda^* = \text{argmin}_{\lambda} (1+\lambda) \ln \int p_t^{1/(1+\lambda)} q_{t-1}^{\lambda/(1+\lambda)} \tag{17}$$

Acknowledgments. This research project and related results were made possible with the support of a European Research Council consolidator grant (81472-SAMPLING) and the NOMIS Foundation.

References

Barber, D. (2012). *Bayesian Reasoning and Machine Learning*. Cambridge University Press.

Barron, K. (2021). Belief updating: does the 'good-news, bad-news' asymmetry extend to purely financial domains? *Experimental Economics*, *24*(1), 31–58.

Beach, L. R., Wise, J. A., & Barclay, S. (1970). Sample proportions and subjective probability revisions. *Organizational Behavior and Human Performance*, *5*(2), 183–190.

Benjamin, D. J. (2019). Errors in probabilistic reasoning and judgment biases. *Handbook of Behavioral Economics: Applications and Foundations 1*, *2*, 69–186.

Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer.

Bramley, N. R., Dayan, P., Griffiths, T. L., & Lagnado, D. A. (2017). Formalizing Neurath's ship: Approximate algorithms for online causal learning. *Psychological Review*, *124*(3), 301–338.

Bramley, N. R., Lagnado, D. A., & Speekenbrink, M. (2015). Conservative forgetful scholars: How people learn causal structure through sequences of interventions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*(3), 708.

Buser, T., Gerhards, L., & Van der Weele, J. J. (2016). Measuring responsiveness to feedback as a personal trait. *Tinbergen Institute Discussion Paper 16-043/I*. (Available at SSRN: `https://ssrn.com/abstract=2789407`)

Charness, G., & Dave, C. (2017). Confirmation bias with motivated beliefs. *Games and Economic Behavior*, *104*, 1–23.

Coutts, A. (2019). Good news and bad news are still news: Experimental evidence on belief updating. *Experimental Economics*, *22*(2), 369–395.

Dasgupta, I., Schulz, E., & Gershman, S. J. (2017). Where do hypotheses come from? *Cognitive Psychology*, *96*, 1–25.

Dasgupta, I., Schulz, E., Tenenbaum, J. B., & Gershman, S. J. (2020). A theory of learning to infer. *Psychological Review*, *127*(3), 412–441.

Dave, C., & Wolfe, K. W. (2003). *On confirmation bias and deviations from Bayesian updating.* (Working Paper)

Edwards, W. (1968). Conservatism in human information processing. *Formal Representation of Human Judgment*.

Girotto, V., & Gonzalez, M. (2007). Children's understanding of posterior probability. *Cognition*, *106*(1), 325–344.

Gotthard-Real, A. (2017). Desirability and information processing: An experimental study. *Economics Letters*, *152*, 96–99.

Grether, D. M. (1992). Testing Bayes rule and the representativeness heuristic: Some experimental evidence. *Journal of Economic Behavior & Organization*, *17*(1), 31–57.

Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological Science*, *17*(9), 767–773.

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.

Ho, M. K., Abel, D., Cohen, J. D., Littman, M. L., & Griffiths, T. L. (2020). The efficiency of human cognition reflects planned information processing. *The 34th AAAI Conference on Artificial Intelligence*.

Holt, C. A., & Smith, A. M. (2009). An update on Bayesian updating. *Journal of Economic Behavior & Organization*, *69*(2), 125–134.

Jaynes, E. T. (2003). *Probability theory: The logic of science*. Cambridge University Press.

Koller, D., & Friedman, N. (2009). *Probabilistic graphical models: principles and techniques*. MIT press.

Kraemer, C., & Weber, M. (2004). How do people take into account weight, strength and quality of segregated vs. aggregated data? experimental evidence. *Journal of Risk and Uncertainty*, *29*(2), 113–142.

Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, *43*.

Ma, Y.-A., Chen, Y., Jin, C., Flammarion, N., & Jordan, M. I. (2019). Sampling can be faster than optimization. *Proceedings of the National Academy of Sciences*, *116*(42), 20881–20885.

Möbius, M. M., Niederle, M., Niehaus, P., & Rosenblat, T. S. (2022). Managing self-confidence: Theory and experimental evidence. *Management Science*.

Ortega, P. A., & Braun, D. A. (2013). Thermodynamics as a theory of decision-making with information-processing costs. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *469*(2153), 20120683.

Peterson, C. R., Schneider, R. J., & Miller, A. J. (1965). Sample size and the revision of subjective probabilities. *Journal of Experimental Psychology*, *69*(5), 522–527.

Phillips, L. D., & Edwards, W. (1966). Conservatism in a simple probability inference task. *Journal of Experimental Psychology*, *72*(3), 346–354.

Powell, D. (2022). A descriptive Bayesian account of optimism in belief revision. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 44).

Prat-Carrabin, A., Wilson, R. C., Cohen, J. D., & Azeredo da Silveira, R. (2021). Human inference in changing environments with temporal structure. *Psychological Review*, *128*(5), 879–912.

Ramscar, M., Hendrix, P., Shaoul, C., Milin, P., & Baayen, H. (2014). The myth of cognitive decline: Non-linear dynamics of lifelong learning. *Topics in Cognitive Science*, *6*(1), 5–42.

Sajid, N., Faccio, F., Da Costa, L., Parr, T., Schmidhuber, J., & Friston, K. (2022). Bayesian brains and the Rényi divergence. *Neural Computation*, *34*(4), 829–855.

Sanborn, A. N., & Chater, N. (2016). Bayesian brains without probabilities. *Trends in Cognitive Sciences*, *20*(12), 883–893.

Sasaki, S., & Kawagoe, T. (2007). *Belief updating in individual and social learning: A field experiment on the internet.* (ISER Discussion Paper)

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, *50*(3), 665–690.

Slovic, P., & Lichtenstein, S. (1971). Comparison of

Bayesian and regression approaches to the study of information processing in judgment. *Organizational Behavior and Human Performance*, *6*(6), 649–744.

Todorov, E. (2006). Linearly-solvable Markov decision problems. *Advances in Neural Information Processing Systems*, *19*.

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.

Zhu, J.-Q., León-Villagrá, P., Chater, N., & Sanborn, A. N. (2022). Understanding the structure of cognitive noise. *PLoS Computational Biology*, *18*(8), e1010312.

Zhu, J.-Q., Sanborn, A. N., & Chater, N. (2020). The Bayesian sampler: Generic Bayesian inference causes incoherence in human probability judgments. *Psychological Review*, *127*(5), 719–748.

Zhu, L., & Gigerenzer, G. (2006). Children can solve Bayesian problems: The role of representation in mental computation. *Cognition*, *98*(3), 287–308.