Title

The Information-Computation Gap: Computational Complexity of Belief Updating Predicts Quality of Human Beliefs

Permalink

https://escholarship.org/uc/item/9s99v0vk

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 48(0)

Authors

Xu, Zhongyu Andy
Murawski, Carsten
Garagnani, Michele
et al.

Publication Date

2026

The Information-Computation Gap: Computational Complexity of Belief Updating Predicts Quality of Human Beliefs

Zhongyu Andy Xu[1], **Michele Garagnani**[1], **Shinsuke Suzuki**[1,2,3], **Carsten Murawski**[1]

[1]Centre for Brain, Mind and Markets, Faculty of Business and Economics, The University of Melbourne, Australia
[2]Faculty of Social Data Science, Hitotsubashi University, Japan
[3]Brain Research Center, Hitotsubashi Institute for Advanced Study, Hitotsubashi University, Japan

Abstract

Rational updating of beliefs requires the processing of new information, which, in turn, requires computational resources, which are limited. Limited computational resources may give rise to an information-computation gap: the quality of an agent's beliefs may not reach the information-theoretic limit due to limited computational resources. Here, we quantify the computational resources required for belief updating in an objective, ex-ante, and task-independent way and document, using a laboratory experiment, that indeed the quality of human beliefs declines in proportion to the computational demands of updating. Further, we validated this metric by showing that it makes unique predictions regarding belief quality. These findings advance a prescriptive theory of beliefs updating that accounts for systematic deviations from rational behaviour. Our results highlight the importance of incorporating computational constraints when evaluating decision-making performance.

Keywords: Belief updating; computational complexity; decision-making; cognition

Introduction

Belief updating, defined as the process of revising beliefs in response to new information, shapes our understanding of the world, guides decision-making and influences economic and social behaviour (Benjamin, 2019). It manifests across many domains of life, including in financial markets, where participants update the distribution of returns whenever news arrive; in medical diagnosis, where physicians revise the probability of illness as new test results emerge; and in public policy, where governments adjust forecasts and strategies in response to evolving health crises.

In updating their beliefs, agents face two fundamental challenges: Limited evidence imposes an information-theoretic limit, that is, a fundamental bound on the accuracy (or error) achievable by any inference procedure. It reflects the intrinsic hardness of an inference task given the amount of evidence and the underlying statistical structure. Limited computational resources, on the other hand, impose a computational limit, that is, a fundamental bound on the accuracy achievable by a computational system, regardless of the algorithm used. It reflects the intrinsic hardness of a task given the computational resource limit and the structure of the problem. Given that belief updating requires computational resources, these constraints lead to an information-computation gap: the information-theoretic limit might not be obtainable due to a lack of sufficient computational resources (Zdeborová & Krzakala, 2016).

Numerous studies have documented systematic deviations of human beliefs from the information-theoretic limit (Bayes optimum). These deviations are typically attributed to cognitive biases, heuristics and approximation strategies. For example, recency bias leads individuals to overweight the most recent evidence (Edenborough, 1975; Pitz & Reinhold, 1968), base-rate neglect causes them to ignore informative priors (Bar-Hillel, 1980; Kahneman, 1972), and the representativeness heuristic drives probability judgements based on stereotypes rather than base rates (Grether, 1992). Other strategies include reinforcement learning, inertia, non-updating, and reliance on sample proportions (Alós-Ferrer & Garagnani, 2023; Beach et al., 1970). Several task-specific predictors of the accuracy of belief updating have been proposed, such as the size of the state space (Banovetz & Oprea, 2023; Camara, 2022; Oprea, 2020), and the number of distinct outcomes (Bernheim & Sprenger, 2020; Puri, 2025). However, most of those studies are descriptive in nature. Only recently has there been an interest in mechanisms explaining these deviations (Zénon et al., 2019).

Here, we investigate the relation between computational resource requirements (computational complexity) of belief updating and the quality of human beliefs. In a pre-registered laboratory study, participants were presented with a set of instances of an balls-and-urns belief-updating task. We assume that belief updating is based on computation. We quantify computational demands of individual instances of belief updating using a novel framework that allows us to quantify the minimum number of arithmetic operations required to update beliefs in accordance with Bayes' theorem. Critically, the instances participants were asked to solve differed in computational demands. We find that the quality of beliefs deteriorates systematically as computational demands of belief updating increases.

Method

Unless noted otherwise, all experimental methods and results described below were pre-registered at OSF (available here).

Participants

We recruited 104 participants (67 female, mean age 26 with SD = 5.41), predominantly university students, for an in-person, laboratory experiment. The experiment lasted 90 minutes on

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

average. Participants were paid a $A\$10$ show-up fee plus a performance bonus based on their responses in four randomly selected trials. The bonus was equal to $A\$\max(0, 10 - d)$, where d is the percentage point deviation from the Bayesian optimum. The average total payment was $A\$30$ per participant ($sd = A\$9$).

Belief updating task

In the main task, participants completed the same 12 predetermined instances of a modified balls-and-urns belief-updating paradigm (Alós-Ferrer & Garagnani, 2023; Alós-Ferrer et al., 2022; Barron, 2021; Grether, 1992; Kraemer & Weber, 2004). Instances were predetermined to maximise information gain by systematically varying computational complexity while holding other characteristics constant, and were presented to participants in random order. In each instance, the computer randomly and secretly selected one of k urns with equal probability, where $k \in \{2, 3, 4\}$ defined the *state space*. Each urn contained balls of c colours ($c \in \{2, 3, 4\}$), defining the *signal space*. Three balls were then drawn *with replacement* from the selected urn and revealed sequentially. After each draw, participants estimated, on a 0–100% scale, the posterior probability that each urn had been selected and the probability of each colour being drawn next. These two questions assessed participants' belief quality. The task screen is shown in Figure 1.

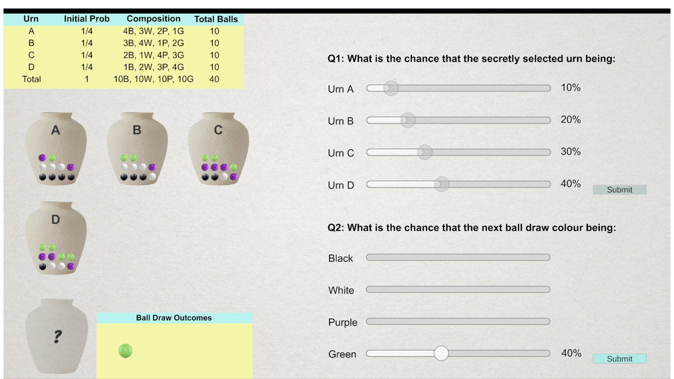

Figure 1. Task screen: In each trial, k urns containing 10 balls each are presented. Participants report posterior probabilities for each of the urns and for the colour of the next ball draw.

No external aids such as pen, paper, or calculators were permitted. Participants were informed of the full composition and selection probabilities of all urns. The colour distribution of each urn was displayed both graphically and in tabular form, which also summarised priors and the number of balls per urn. No feedback was provided in the main experiment. During the second and third draws, participants were reminded of their previous responses. Reported probabilities were constrained to sum to one.

We will refer to the task to compute the posterior over the set of urns given a sequence of ball draws as *hidden state inference (HSI)* task. We also consider the probability of the next ball drawn being colour C, given a sequence of n ball draws. We will refer to the task of computing the probability of the next ball drawn of being a given colour as *future forecasting (FF)* task. Below, we only report results related to the HSI task due to word limits.

Computational complexity

We aim to quantify the minimum computational resources needed to update beliefs optimally (i.e., via Bayesian updating) for each individual instance, drawing on computational complexity theory (Arora & Barak, 2009; Bossaerts & Murawski, 2017; Cook, 2007; Hartmanis & Stearns, 1965).

Intuitively, updating beliefs after observing a single signal (ball draw) may require only a small number of computations, whereas incorporating multiple signals requires repeatedly combining and renormalising probabilities across possible states. For example, inferring which urn was selected after observing one ball colour involves evaluating a small set of likelihoods, while updating beliefs after two or three draws requires recomputing these likelihoods multiple times and aggregating intermediate results. We define the computational demands of an individual updating step as the minimum number of arithmetic operations required to perform Bayesian updating.

We distinguish two key forms of computational complexity: *step-wise computational complexity (CC)* and *cumulative computational complexity (CCC)*. The former refers to the incremental number of computations required to update beliefs based one additional piece of information (one additional ball draw), whereas *cumulative computational complexity* captures the total computational resources required to update beliefs a number of times given n ball draws. Then these measures will be related to step-wise and overall belief quality, respectively (see sec Information-computation gap). The distinction between step-wise and cumulative measures is motivated by the need to preclude the carry-over of accumulated errors from prior updating steps, thereby isolating the computational demands specific to each incremental inference.

In the main analysis, we assign a weight of 2 to multiplications and divisions, and a weight of 1 to additions and subtractions. We assign a greater weight to multiplication and division than to addition and subtraction, reflecting two converging considerations. First, in standard computer architectures multiplication and division are more expensive than addition and subtraction (Hennessy & Patterson, 2011). Second, in human arithmetic task, multi-digit multiplication and division are associated with longer response times, less use of retrieval strategies and greater use of procedural strategies compared to multi-digit addition or subtraction (Campbell & Xue, 2001; Tronsky, 2005). Crucially, our results remain qualitatively unchanged when employing equal weighting across arithmetic operations and when restricting the complexity measure to multiplication and division exclusively.

We first consider the number of computations required for updating in the HSI task. Let $\mathbf{r}_n = (r_1, \ldots, r_c)$ denote the vector of colour counts after n ball draws, where r_j is the number of balls of colour j observed and $\sum_{j=1}^{c} r_j = n$ with c is the number of possible colours. Let $j_n \in \{1, ..., c\}$ denote the colour of the ball drawn at step n. Applying Bayes' rule sequentially, the posterior probability that the secretly selected urn is U_i given $\mathbf{r}_n$ is

$$P(U_i|\mathbf{r}_n) = \frac{P(C = j_n \mid U_i)P(U_i|\mathbf{r}_{n-1})}{\sum_{l=1}^{k} P(C = j_n \mid U_l)P(U_l|\mathbf{r}_{n-1})}, \quad (1)$$

where $P(C = j_n \mid U_i)$ is the probability of drawing a ball of colour j_n in Urn i, $P(U_i|\mathbf{r}_{n-1})$ is the posterior for Urn i after preceding $n-1$ ball draws (with $P(U_i|\mathbf{r}_0) = 1/k$) and k is the number of urns.

We provide an example of how step-wise computational complexity is calculated below. After the first ball drawn, computing the posterior for a single hidden state requires one multiplication and one division under Bayes' rule (as the evidence term is directly observable from the experimental interface), yielding a complexity of four per state. Computing all k posteriors thus results in a total complexity of $4k$. Alternatively, an individual may compute only $k-1$ posteriors and derive the final one via the law of total probability, which requires an additional $k-1$ operations (additions or subtractions), yielding a total of $5k-5$. Assuming the most efficient strategy, the step-wise computational complexity after the first ball draw is therefore $CC^1_{HSI} = \min(4k, 5k-5)$ in HSI task. Similarly, the step-wise computational complexity for all subsequent draws is given by $CC^2_{HSI} = CC^3_{HSI} = \min(8k-6, 7k-1)$ in HSI task, where $k \in \{2, 3, 4\}$ denotes the number of urns (size of the state space). This arises because, for all subsequent ball draws, the evidence term must be computed rather than directly observed. The step-wise computational complexity of Bayesian updating is therefore always lower for the updating step after the first ball draw than for later draws in HSI task. This property motivates our preregistered hypothesis that trials with a single observed signal should yield higher belief quality than trials with two or three signals, reflecting their lower computational demands. Furthermore, belief quality should not differ between two- and three-signal trials, consistent with the fact that these conditions impose equivalent computational requirements.

Cumulative computational complexity quantifies the total number of operations required to compute Bayesian posteriors efficiently up to the current ball draw. As sequential updating conditional on previous beliefs is the most efficient inference strategy in our task, cumulative complexity equals the sum of all prior step-wise complexities. Formally, $CCC^1_{HSI} = CC^1_{HSI}$, $CCC^2_{HSI} = CCC^1_{HSI} + CC^2_{HSI}$ and $CCC^3_{HSI} = CCC^2_{HSI} + CC^3_{HSI}$.

Information-computation gap

We operationalise the information-computation gap as the difference in information contained in two posterior distributions, the Bayes optimal posterior distribution and posterior distributed reported by a participant. We quantify this difference using the Jensen-Shannon distance (JSD). JSD is used instead of KL divergence since it is symmetrical and bounded. For two distributions P and Q, the JSD difference is given by

$$JSD(P \parallel Q) = \sqrt{\frac{1}{2}KL(P \parallel M_2) + \frac{1}{2}KL(Q \parallel M_2)}, \quad (2)$$

where $M_2 = \frac{1}{2}(P + Q)$, $KL(P \parallel M_2)$ is the Kullback-Leibler divergence of P from M_2 and $KL(Q \parallel M_2)$ is the Kullback-Leibler divergence of Q from M_2. The Kullback-Leibler divergence between P and Q is defined as $KL(P \parallel Q) = \sum_i P(i) \log_2 \frac{P(i)}{Q(i)}$, where $P(i)$ and $Q(i)$ are the probabilities of outcome i in distributions P and Q, respectively (Lin, 1991). JSD quantifies how different Q is from P in information terms. Applied to our setting, we hypothesise that the JSD between participant reported posterior and the Bayesian posterior increases as the computational resource requirements of computing the posterior increases.

In human belief updating processes, two sources of error can be distinguished: errors originating from prior beliefs and errors arising from the updating process itself and. We quantify the latter using step-wise belief quality, and the combined effect of both sources using overall belief quality. Step-wise belief quality captures participants' performance incrementally processing new information. For example, it evaluates whether a participant appropriately updates their beliefs after observing the second draw, conditional on their beliefs after the first. This measure allows us to isolate local updating errors at each step and distinguish them from cumulative errors driven by priors. In contrast, overall belief quality captures participants' performance given all information available up to a given point within an instance. For example, after three signal draws, it reflects how accurately participants integrate the full sequence of observations. We denote the posterior reported by a participant as P_R, the Bayesian posterior as P_B, and P_S as the posterior a participant would obtain by applying Bayes' rule after each draw using their own prior (i.e., their reported posterior from the previous step within the same instance[1]). Step-wise belief quality is defined as $BQ_s = 1 - \text{JSD}(P_R \parallel P_S)$, and overall belief quality as $BQ_o = 1 - \text{JSD}(P_R \parallel P_B)$. Both measures are bounded between 0 and 1[2].

Data analysis

All statistical analyses were conducted in a Bayesian framework. We pre-registered both intercept-only random-effects

[1]In the first trial of each instance, own prior equals to Bayesian prior

[2]Values of 0 and 1 are transformed to 10^{-10} and $1 - 10^{-10}$, respectively, ensuring that all observations lie within the support of the Beta distribution used for estimation.

models and intercept-plus-random-slope models for each Bayesian hierarchical regression. When the more complex model achieved satisfactory convergence, we conducted model selection based on the Bayes factor (BF) with base 10. Specifically, placing the intercept-plus-slope model in the numerator, the more complex model was preferred if $BF > 3$; otherwise, the simpler intercept-only model was selected. Convergence was assessed using the $\hat{R}$ diagnostic, with all winning models achieving $\hat{R} < 1.1$ (Brooks & Gelman, 1998). Model stability was further evaluated using the Effective Sample Size (ESS), and all winning models with at least one independent variable achieved $ESS > 3{,}000$ (Gong & Flegal, 2016).

We evaluated our hypotheses using Bayes factor criteria, where $BF > 3$ indicates substantial evidence, $BF > 20$ strong evidence, and $BF > 100$ decisive evidence in favour of a hypothesis compared to its alternative (Kass & Raftery, 1995). For each hypothesis test, we report the Bayes factor defined as the ratio of the marginal likelihood of model under alternative hypothesis to that of model under null hypothesis, denoted BF; median of the posterior distribution for the interested variable (β_{var}); and its corresponding 95% credible interval ($ETI_{0.95}$).

Results

Computational complexity and hidden-state inference

Overall belief quality (BQ_o) was regressed on cumulative computational complexity (CCC), the cumulative computational resources required to update beliefs using all available information, and compared it with an intercept-only control model. As predicted, cumulative computational complexity was strongly negatively associated with overall belief quality, with decisive evidence favouring the model including CCC over the control ($BF > 100$, $\beta_{CCC} = -0.014$, $ETI_{0.95} = [-0.015, -0.013]$; Figure 2). This indicates that greater cumulative computational demands systematically impair belief quality.

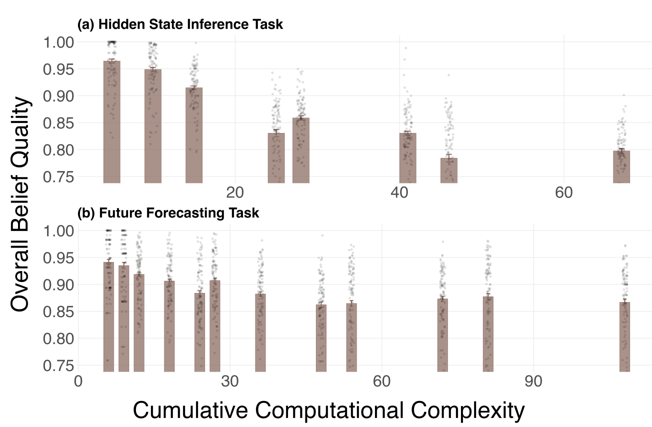

Figure 2. Overall Belief Quality Split by Cumulative Computational Complexity. Error bars represent standard error, dots represent mean overall belief quality of each participant.

Our second hypothesis predicted a negative relationship between step-wise belief quality (BQ_s), a measure of belief updating quality at a single step based on the new information received at this step, and step-wise computational complexity (CC), the incremental number of computations required to update beliefs with one additional piece of information. CC was regressed on BQ_s and compared with an intercept-only control model. Consistent with this hypothesis, higher step-wise computational complexity was associated with lower step-wise belief quality, with decisive evidence favouring the model including step-wise complexity over the control model ($BF > 100$, $\beta_{CC} = -0.037$, $ETI_{0.95} = [-0.041, -0.033]$; Figure 3).

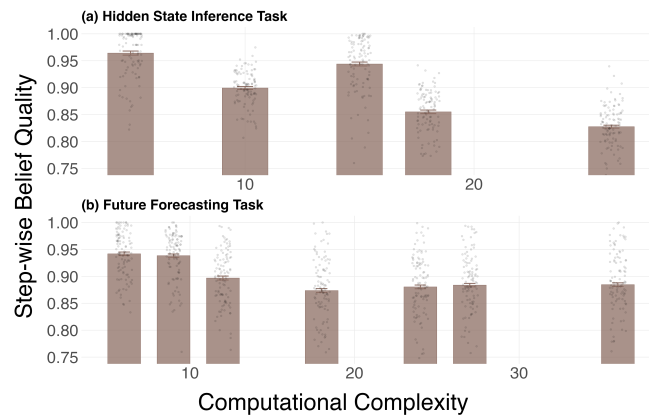

Figure 3. Step-wise Belief Quality Split by Step-wise Computational Complexity. Error bars represent standard error, dots represent mean step-wise belief quality of each participant.

To further test whether step-wise computational complexity is a task-independent predictor of belief updating, we evaluated additional preregistered hypotheses that exploit the computational structure of the HSI task. In this task, computational requirements after the second and third ball draws are identical and both exceed those after the first draw (see Computational complexity). We therefore predicted that step-wise belief quality would be lower after the second and third draws than after the first, but indistinguishable between the second and third draws. We tested this prediction by comparing Bayesian regression models differing in whether they allowed distinct mean levels of BQ_s across draw numbers. Relative to one-draw trials, step-wise belief quality was substantially lower after both two draws ($\beta_{DV(sDraws=2)} = -0.979$, $ETI_{0.95} = [-1.041, -0.916]$) and three draws ($\beta_{DV(sDraws=3)} = -0.978$, $ETI_{0.95} = [-1.041, -0.917]$) with decisive evidence favouring models that included these mean shifts ($BF > 100$).

By contrast, when two-draw trials were used as the reference condition, allowing an additional mean shift for three-draw trials did not improve model performance ($BF > 100$, $\beta_{DV(sDraws=3)} = 0.000$, $ETI_{0.95} = [-0.055, 0.055]$). Together, these results indicate that step-wise belief quality is highest during initial belief updating and uniformly lower during subsequent updates, consistent with predictions derived from step-wise computational complexity. This pattern is robust against signal diagnosticity, observed outcomes, and state- and signal-space size, providing convergent evidence that computational complexity systematically predicts human belief updating behaviour.

Discussion

How people update beliefs and perform probabilistic inference have been investigated widely in the literature. Several studies have examined the effects of complexity on the quality of beliefs (Bernheim & Sprenger, 2020; Enke & Graeber, 2023; Enke & Shubatt, 2023; Enke et al., 2023; Kurzban et al., 2013; Oprea, 2020, 2022; Sasaki & Kawagoe, 2007; Westbrook & Braver, 2015; Zhu et al., 2023). However, they typically defined complexity in terms of individual task attributes and are often descriptive in nature. The aim of our study was to investigate the effects of computational complexity on the quality of beliefs, using a mechanistic, domain-general framework based on computation. To this end, we conducted a pre-registered laboratory experiment in which participants had to update their beliefs in different instances of a modified balls-and-urns belief-updating task that varied in computational resource requirements. We observe a robust negative association between the quality of beliefs and computational complexity both when considering individual updating steps and the a series of updating steps jointly. In other words, computational complexity of belief updating systematically predicts the quality of human beliefs.

Our study has several limitations. First, the metric of computational complexity we use is an approximation of the actual computational complexity people might be facing. While the metric counts the number of operations required for inference, it does not account for the varying difficulty of each calculation; for example, computing 0.5×0.5 is cognitively easier than 0.14567×0.24567. Future research could refine this approach by incorporating the inherent difficulty of computations developing a more nuanced complexity metric better capturing aspects of human computation. Second, our metric does not account for space complexity (Arora & Barak, 2009). Participants were required to retain intermediate results in memory, which may have imposed additional cognitive load. Future research could explore the effect of cognitive offloading tools—such as providing pen and paper—on reducing space complexity and improving belief quality. Third, the current research is not able to disentangle belief updating size from computational complexity in order to determine whether both measures independently predict belief updating quality, or whether one subsumes the other, this is left to future studies.

Despite these limitations, our findings consistently support the central claim that the computational resources required to rationally solve belief updating problems systematically predict human belief quality. When considered alongside prior research in human decision making and computational complexity (Bossaerts & Murawski, 2017; Franco & Murawski, 2023; Franco et al., 2021, 2022, 2024), these results strengthen the case that computational resource demands are a fundamental driver of human cognitive performance. We would argue that many existing behavioural patterns can be reinterpreted as manifestations of the underlying computational processes and constraints. Accordingly, computational complexity should be regarded as a key explanatory factor in understanding human cognition. Future research should incorporate our task-independent definition of computational complexity as a complementary lens for explaining a wide range of human behaviours.

References

Alós-Ferrer, C., Buckenmaier, J., Garagnani, M., & Ritschel, A. (2022). The reinforcement paradox: Monetary incentives and bayesian updating. *Economics Letters*, *211*, 110200.

Alós-Ferrer, C., & Garagnani, M. (2023). Part-Time Bayesians: Incentives and Behavioral Heterogeneity in Belief Updating. *Management Science*, *69*(9), 5523–5542. https://doi.org/10.1287/mnsc.2022.4584

Arora, S., & Barak, B. (2009). *Computational Complexity: A Modern Approach*. Cambridge University Press.

Banovetz, J., & Oprea, R. (2023). Complexity and Procedural Choice. *American Economic Journal: Microeconomics*, *15*(2), 384–413. https://doi.org/10.1257/mic.20210032

Bar-Hillel, M. (1980). The base-rate fallacy in probability judgments. *Acta Psychologica*, *44*(3), 211–233. https://doi.org/10.1016/0001-6918(80)90046-3

Barron, K. (2021). Belief updating: Does the 'good-news, bad-news' asymmetry extend to purely financial domains? *Experimental Economics*, *24*(1), 31–58. https://doi.org/10.1007/s10683-020-09653-z

Beach, L. R., Wise, J. A., & Barclay, S. (1970). Sample proportions and subjective probability revisions. *Organizational Behavior and Human Performance*, *5*(2), 183–190. https://doi.org/10.1016/0030-5073(70)90014-0

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

Bernheim, B. D., & Sprenger, C. (2020). On the Empirical Validity of Cumulative Prospect Theory: Experimental Evidence of Rank-Independent Probability Weighting. *Econometrica*, *88*(4), 1363–1409. https://doi.org/10.3982/ECTA16646

Bossaerts, P., & Murawski, C. (2017). Computational Complexity and Human Decision-Making. *Trends in Cognitive Sciences*, *21*(12), 917–929. https://doi.org/10.1016/j.tics.2017.09.005

Brooks, S. P., & Gelman, A. (1998). General Methods for Monitoring Convergence of Iterative Simulations. *Journal of Computational and Graphical Statistics*, *7*(4), 434–455. https://doi.org/10.1080/10618600.1998.10474787

Camara, M. K. (2022). Computationally Tractable Choice. *Proceedings of the 23rd ACM Conference on Economics and Computation*, 28–28. https://doi.org/10.1145/3490486.3538256

Campbell, J. I. D., & Xue, Q. (2001). Cognitive arithmetic across cultures. *Journal of Experimental Psychology: General*. https://doi.org/10.1037/0096-3445.130.2.299

Cook, S. A. (2007). An overview of computational complexity. In *ACM Turing award lectures* (p. 1982). ACM.

Edenborough, R. A. (1975). Order effects and display persistence in probabilistic opinion revision. *Bulletin of the Psychonomic Society*, *5*(1), 39–40. https://doi.org/10.3758/BF03336694

Enke, B., & Graeber, T. (2023). Cognitive Uncertainty*. *The Quarterly Journal of Economics*, *138*(4), 2021–2067. https://doi.org/10.1093/qje/qjad025

Enke, B., Graeber, T., & Oprea, R. (2023). Complexity and Hyperbolic Discounting. *Working Paper*. https://doi.org/10.3386/w31047

Enke, B., & Shubatt, C. (2023). Quantifying Lottery Choice Complexity. *Working Paper*. https://doi.org/10.3386/w31677

Franco, J. P., Bossaerts, P., & Murawski, C. (2024). The neural dynamics associated with computational complexity. *PLOS Computational Biology*, *20*(9), e1012447. https://doi.org/10.1371/journal.pcbi.1012447

Franco, J. P., Doroc, K., Yadav, N., Bossaerts, P., & Murawski, C. (2022). Task-independent metrics of computational hardness predict human cognitive performance. *Scientific Reports*, *12*(1), 12914. https://doi.org/10.1038/s41598-022-16565-w

Franco, J. P., & Murawski, C. (2023). Harnessing Computational Complexity Theory to Model Human Decision-making and Cognition. *Cognitive Science*, *47*(6), e13304. https://doi.org/10.1111/cogs.13304

Franco, J. P., Yadav, N., Bossaerts, P., & Murawski, C. (2021). Generic properties of a computational task predict human effort and performance. *Journal of Mathematical Psychology*, *104*, 102592. https://doi.org/10.1016/j.jmp.2021.102592

Gong, L., & Flegal, J. M. (2016). A Practical Sequential Stopping Rule for High-Dimensional Markov Chain Monte Carlo. *Journal of Computational and Graphical Statistics*, *25*(3), 684–700. https://doi.org/10.1080/10618600.2015.1044092

Grether, D. M. (1992). Testing bayes rule and the representativeness heuristic: Some experimental evidence. *Journal of Economic Behavior & Organization*, *17*(1), 31–57. https://doi.org/10.1016/0167-2681(92)90078-P

Hartmanis, J., & Stearns, R. E. (1965). On the computational complexity of algorithms. *Transactions of the American Mathematical Society*, *117*(0), 285–306. https://doi.org/10.1090/S0002-9947-1965-0170805-7

Hennessy, J. L., & Patterson, D. A. (2011). *Computer Architecture: A Quantitative Approach*. Morgan Kaufmann.

Kahneman, D. (1972). On prediction and judgement. *ORI Research Monograph*, *1*(4).

Kass, R. E., & Raftery, A. E. (1995). Bayes Factors. *Journal of the American Statistical Association*, *90*(430), 773–795. https://doi.org/10.1080/01621459.1995.10476572

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. https://doi.org/10.1023/B:RISK.0000038940.62992.1b

Kurzban, R., Duckworth, A., Kable, J. W., & Myers, J. (2013). An opportunity cost model of subjective effort and task performance. *Behavioral and Brain Sciences*, *36*(6), 661–679. https://doi.org/10.1017/S0140525X12003196

Lin, J. (1991). Divergence measures based on the Shannon entropy. *IEEE Transactions on Information Theory*, *37*(1), 145–151. https://doi.org/10.1109/18.61115

Oprea, R. (2020). What Makes a Rule Complex? *American Economic Review*, *110*(12), 3913–3951. https://doi.org/10.1257/aer.20191717

Oprea, R. (2022). Simplicity Equivalents. *Working Paper*.

Pitz, G. F., & Reinhold, H. (1968). Payoff effects in sequential decision-making. *Journal of Experimental Psychology*, *77*(2), 249–257. https://doi.org/10.1037/h0025802

Puri, I. (2025). Simplicity and Risk. *The Journal of Finance*, *80*(2), 1029–1080. https://doi.org/10.1111/jofi.13417

Sasaki, S., & Kawagoe, T. (2007). Belief Updating in Individual and Social Learning: A Field Experiment on the Internet. *SSRN Scholarly Paper*. https://doi.org/10.2139/ssrn.989689

Tronsky, L. N. (2005). Strategy use, the development of automaticity, and working memory involvement in complex multiplication. *Memory & Cognition*. https://doi.org/10.3758/BF03193086

Westbrook, A., & Braver, T. S. (2015). Cognitive effort: A neuroeconomic approach. *Cognitive, Affective, & Behavioral Neuroscience*, *15*(2), 395–415. https://doi.org/10.3758/s13415-015-0334-y

Zdeborová, L., & Krzakala, F. (2016). Statistical physics of inference: Thresholds and algorithms. *Advances in Physics*, *65*(5), 453–552. https://doi.org/10.1080/00018732.2016.1211393

Zénon, A., Solopchuk, O., & Pezzulo, G. (2019). An information-theoretic perspective on the costs of cognition. *Neuropsychologia*, *123*, 5–18. https://doi.org/10.1016/j.neuropsychologia.2018.09.013

Zhu, J.-Q., Sanborn, A. N., Chater, N., & Griffiths, T. L. (2023). Computation-Limited Bayesian Updating. *PsyArXiv*. https://doi.org/10.31234/osf.io/znywg