

MIT Open Access Articles

Resource-rational decision making

The MIT Faculty has made this article openly available. ***Please share*** how this access benefits you. Your story matters.

Citation: Bhui, Rahul, Lai, Lucy and Gershman, Samuel J. 2021. "Resource-rational decision making." *Current Opinion in Behavioral Sciences*, 41.

As Published: 10.1016/J.COBEHA.2021.02.015

Publisher: Elsevier BV

Persistent URL: <https://hdl.handle.net/1721.1/144282>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-NonCommercial-NoDerivs License



Resource-rational decision making

Rahul Bhui¹, Lucy Lai², and Samuel J. Gershman^{3,4,*}

¹MIT Sloan School of Management

²Program in Neuroscience, Harvard University

³Department of Psychology and Center for Brain Science, Harvard University

⁴Center for Brains, Minds and Machines, MIT

*Correspondence: gershman@fas.harvard.edu

Abstract

Across many domains of decision making, people seem both rational and irrational. We review recent work that aims to reconcile these apparently contradictory views by modeling human decisions as optimal under a set of cognitive resource constraints. This “resource-rational” analysis connects psychology and neuroscience to ideas from engineering, economics, and machine learning. Here, we focus on an information-theoretic formalization of cognitive resources, highlighting its implications for understanding three important and widespread phenomena: reference dependence, stochastic choice, and perseveration. While these phenomena have traditionally been viewed as irrational biases or errors, we suggest that they may arise from a rational solution to the problem of resource-limited decision making.

Introduction

A paradox resides at the heart of behavioral science. Much work focuses on the complex tasks that people successfully solve in domains ranging from sensory perception to social communication [1, 2, 3, 4, 5, 6, 7]; yet, an equally important body of research lists our seemingly countless deviations from optimality, often at the same time [8, 9, 10]. How might one explain the “apparent contradictions between the great triumphs and the dramatic failures of the human mind” [11]?


Here, we review recent work that seeks to resolve this paradox by formalizing the idea that people are doing as best as they can subject to constraints on cognitive resources. This *resource-rational* (or *computationally rational*) perspective [12, 13, 14, 15] lets us rigorously capture both the method and the madness of human judgment in the same thrust—that is, the adaptive logic behind decision making as well as its systematic limitations. The world is vast, but we have only limited cognitive capacities with which to grasp it. We are thus forced to mentally represent information about the world as economically as possible. We take shortcuts by compressing or distorting or flat-out ignoring parts of the world—at least, when we can get away with it. Models of resource rationality clarify which corners are worth cutting and how we should cut them.

Resource rationality refines traditional views of bounded rationality by conceiving of judgment as optimal under cognitive costs or constraints, rather than merely feasible. By casting mental

processes in the mathematics of optimization, this approach enables precise descriptions of how we deploy limited cognitive resources where they are most valuable, and lets us trace out how decision making is shaped as a result. To achieve this blend of tractability and realism, resource-rational analysis links psychology and neuroscience with concepts from engineering, economics, information theory, statistics, and machine learning. In this article, we focus on analyses using information theory, which offers a versatile mathematical language for defining efficiency. Below, we outline some examples of important behavioral phenomena that this approach has helped to characterize.

Box I. Formalizing efficiency

Figure I. A view of cognition as a communication channel.



Perceptual and cognitive processes can be described as the transmission of information through a communication channel with limited capacity, illustrated in Figure I. A stimulus (x with prior distribution $p(x)$) must be encoded in an internal representation (r given by the encoder $p(r|x)$) which can only imperfectly represent the stimulus. A judgment is formed, which could reflect a decoded reconstruction of the stimulus (\hat{x} following $p(\hat{x}|r)$) or an action (a according to policy $p(a|\hat{x})$), yielding state-dependent reward value $v(a, x)$).

Efficiency obtains when some measure of the representation's fidelity (or consequent reward) is optimized given a cost or constraint on information transmission. Models are typically constructed by making assumptions about some aspects of this process and characterizing the rest based on what is optimal according to the objective function. Three classes of objective functions are commonly used in information-theoretic analyses:

1. Maximization of mutual information, $I(x, r) \equiv H(r) - H(r|x)$ where $H(r) \equiv E[-\log p(r)]$ is the entropy [16] that expresses uncertainty about variable r , and $H(r|x)$ is the conditional entropy that expresses uncertainty about r conditional on x . This criterion can be maximized exactly or via an approximation like Fisher information [17]. It can be interpreted as the reduction in uncertainty about r after observing x , or vice versa, because it is a symmetric measure.
2. Minimization of squared loss, $E_{\hat{x}, x}[(\hat{x} - x)^2]$. This criterion can be minimized subject to a cost or constraint on information capacity, often defined in terms of mutual information. It is often used in rate-distortion theory and economic analyses of capacity-limited decision making [18, 19].

3. Maximization of reward, $E_{a,x}[v(a, x)]$. This criterion can be maximized subject to a cost or constraint on information capacity often defined in terms of mutual information. It explicitly captures the consequences of actions taken based on the representation.

These objective functions formalize constraints that are internal to an agent, but another impetus for efficiency comes from external constraints, namely the limited data available to the agent. Results in statistical learning theory show that in order to generalize, an agent must be able to compress its data [20, 21]. This principle has motivated recent algorithms in machine learning, such as the variational autoencoder [22], that extract low complexity structure by pushing data through an information bottleneck, typically a low-dimensional layer in a deep neural network.

Reference dependence and efficient coding

Abundant research [23] and personal experience tell us that our feelings about an outcome depend on how we compare it to our other experiences. However, exactly what kinds of comparisons do we make? And for that matter, why do we even make comparisons at all? Answers lie in the *principle of efficient coding* [24, 25], which asserts that brains represent information in a cost-effective manner by adapting our neural representations to the local statistics of our environment. While an idealized agent would be able to perfectly assign values to all possible outcomes and have completely rational preferences, real humans cannot be sensitive to the entire spectrum of conceivable values at the same time. To operate within this limitation, the brain must preferentially dedicate its resources to representing the most likely stimulus values [26].

To formalize a version of this idea, consider the problem of finding a mapping $p(r|x)$ from a real-valued stimulus to an internal representation so as to maximize information transmission (Box I), defined as the mutual information between the two, $I(x, r)$. For simplicity, suppose the encoding process is deterministic, but the representation only comprises a limited number of internal states; we can think of it as a step function with output normalized to the unit interval. Since the set of internal states is coarser than the set of possible values, compression is required. When the number of internal states grows large, the optimal solution in this setting approaches the cumulative distribution function (CDF) of the stimulus distribution [27]:

$$r^* = F(x) \equiv \int_{-\infty}^x p(x') dx'.$$

The same solution also results from minimizing the probability of making errors when ranking two values drawn from the prior distribution [28].

Observe that this representation is not fixed, but instead mirrors the context as defined by the environmental distribution of stimulus values. It is most sensitive to variation within the most likely ranges of values, as the CDF is steep where the probability density function is high. For example, if values follow a bell-shaped prior (as illustrated in Figure 1), the representation is an S-shaped function, sensitive near the median but insensitive at the extremes. It shifts when the




Figure 1: Context-dependent value function resulting from efficient coding. The optimal representation with a fixed output range is given by the CDF of the prior stimulus distribution. It is sensitive to gradations near the peak of the prior, where values are most frequently encountered, and insensitive at the tails, which are rarely experienced. As the prior changes, such as from p_1 to p_2 , so does the efficient representation, from r_1 to r_2 . Adapted from [29].

prior shifts and flattens when the prior flattens. This property reveals how preferences should be shaped by the entire contextual distribution, rather than only a single reference point as in standard models of reference dependence. Consistent with the efficient coding hypothesis, experiments have demonstrated that sensitivity to value differences between options tends to decrease (a) when values are far away from the reference point (the prior mean), reflecting a form of diminishing sensitivity similar to Weber’s law for magnitude discrimination [30, 31]; and (b) when the variance of values across options is large [32]. Furthermore, the same value is judged as more or less attractive to the degree that it is higher or lower (respectively) than the reference point [33, 34, 35].

The CDF transformation is approximated by the psychological process model of decision by sampling [36, 37, 38] which has been proposed to account for malleability in representations of economic attributes (like values, probabilities, and delays) beyond the reach of traditional models. The key idea is that when judging an attribute, a decision maker constructs a relativistic representation by sampling attribute magnitudes stored in memory and comparing these samples ordinally to the target attribute magnitude. If memory is assumed to reflect the ecological distribution of magnitudes, then we can understand the sampling and comparison process as a Monte Carlo approximation of the CDF. Importantly, this means that the encoding process is noisy and the optimal internal representation needs to compensate for this noise, which can be accomplished by kernel smoothing the samples [38]. One consequence of this modification is that judgments will be sensitive not only to the rank of attribute magnitudes (via the CDF) but also to their range. A fixed number of samples spread across a wider range of magnitudes means that the attribute space is more sparsely sampled, and therefore needs a greater smoothing bandwidth to compensate for noise. This results in a form of range normalization: discriminability of magnitudes decreases with range. Many

studies of judgment and decision making have provided evidence for range normalization in choice [39, 40, 41]. More specifically, the balance between rank and range effects may be predicted based on the optimal degree of smoothing for a given level of noise in the encoding process [38].

Variations on this theme emerge when making alternative assumptions about the elements depicted in Box I. For example, the stimulus might consist of multiple items or attributes; in this case, efficiency requires encoding to be based on combinations of these dimensions that increase the statistical independence of the representational components. This function may be implemented by divisive normalization, which could account for violations of classic axioms of rational choice [42, 43, 44, 45, 46]. Other objectives or constraints may also be considered [47, 48, 49]; for instance, reward, rather than information capacity, might be taken as the criterion to maximize [50, 51, 52]. While evidence for reference dependence is abundant in both brain [26, 33] and behavior [23, 53, 54], these lines of theorizing could help pin down its nature more precisely by clarifying the conditions under which we should expect its different guises. Recent empirical evidence indicates that decision by sampling provides a better quantitative account of incentivized decision making compared to other models of efficient coding which maximize accuracy or reward, suggesting that it may be a particularly robust and pervasive cognitive process [51].

The context dependence arising in these types of models recapitulates the influential idea of ecological rationality, according to which rationality emerges from the match between the cognitive strategy of the agent and the structure of its environment [55, 56]. The same concept is plainly seen in the domain of sensory perception, where the principle of efficient coding originated; just as our eyes acclimate to light or dark and our ears get accustomed to noise or silence, our sense of value adapts to the level of rewards in our local environment. In fact, much of our understanding of context dependence in preferences stems from analogous research on perception [57], from the current era (e.g., [58, 59, 60, 61, 17, 62]) back to the very inception of behavioral economics [30, 63]. These deep interdisciplinary links underscore the unifying power of resource rationality.

Stochastic choice and perseverance

Even when faced with the same decision, our choices may vary in a random fashion [64]. This inconsistency can occur because the quality of judgment is limited by how precisely our brains represent values [65] (or states of the world which indicate values). While randomness in choice has often been treated as exogenously determined, an alternative perspective emerges from the idea that agents can control the richness of their representations. Higher-fidelity representations let us cleanly distinguish between different states and make it clear which action is best, but are costlier to implement [66]. Thus, there is a tradeoff between the complexity of the representation and the reward that the corresponding action policy achieves.

To formalize core elements of this *reward-complexity tradeoff*, consider the problem of finding a policy $p(a|x)$ which maps states of the world to actions in order to maximize expected reward (Box I). Critically, this mapping incurs a cost proportional to the complexity of the state representation, as measured by the mutual information between states and actions, $I(a, x)$. The agent thus selects a policy to maximize $\beta E_{a,x}[v(a, x)] - I(a, x)$, where β reflects the added value of a 1-bit increase in the fidelity of the state representation. The optimal action policy that solves this trade-off is given




Figure 2: Complexity-dependent binary choice probabilities. The focal option here tends to be better in more states of the world. When the benefit of a complex representation is high (or equivalently, the cost of information transmission is low), the agent almost deterministically selects whichever option is best in each situation (yellow line). When information is relatively costly, choice is stochastic with a bias toward the action which tends to be chosen more often across states (red line). As capacity becomes severely constrained, the bias dominates and the agent almost always selects the marginally high frequency action, ignoring the true state (blue line).

by [67, 68, 69, 70, 71, 72]

$$p^*(a|x) = \frac{\exp[\beta v(a, x) + \log p^*(a)]}{\sum_{a'} \exp[\beta v(a', x) + \log p^*(a')]}$$

where $p^*(a) = \sum_x p^*(a|x)p(x)$ is the marginal action distribution.

This policy has three key features, which can be seen in Figure 2. First, it takes the familiar form of the softmax function (also known as the multinomial logit model or the Boltzmann distribution). Note that this form results from the entropy-based cost function rather than from convenient assumptions about the shape of the noise distribution. Second, the inverse temperature parameter, which regulates the degree of choice stochasticity, is given by the β parameter which quantifies the benefit of a complex representation. When β is high, the policy will be concentrated on the best action in each state. Third, the bias term is interpretable as the marginal action distribution, $p^*(a)$, meaning that choices are biased toward actions that are frequently chosen across states. When it is not very beneficial to discern the best action (β is low), choices become less sensitive to payoffs of actions and more dependent on the marginal distribution.

This line of thinking can be extended to various dynamic settings. The above expression could summarize the choice probabilities resulting from a noisy process of deliberation taking place in real time [73, 74]. Such findings help normatively ground and expand upon classic sequential sampling models that jointly capture choice stochasticity and response time. Or, if learning happens over

multiple trials, a kind of perseveration occurs as choice probabilities appear to recapitulate one's prior action history [75, 76]. Inertia can be adaptive when states are correlated across time, because learning about state transitions is costly, so the agent tends to rely on their prior experience and repeat their previous actions [77]. In complex multi-step tasks such as navigation, policy compression can yield a noisier version of the ideal trajectory or even an alternative path that is more tolerant of noise [68, 78]. Moreover, the coarsened representation reflects efficient state abstraction (e.g., clustering together locations in the same room) and subgoals naturally emerge at the transition from one state cluster (e.g., room) to the next [79, 80].

In sum, we may endure randomness when it's not worth eliminating, and use strategies that compensate for its consequences. These strategies can manifest as biases or imperfections despite having adaptive value, like taking a longer, but simpler, route between two points in a city to avoid the risk of getting lost. Errors and biases can thus be two sides of the same coin. These kinds of mechanisms may shed light on properties of noisy and biased perceptions of value [48, 51] and probability [81, 82, 83], risky [84] and intertemporal choice [85], and a wide variety of further economic outcomes [19].

Conclusions

We have dipped our toes into the ocean of resource rationality [15], focusing on its implications for decision making. The idea that cognitive resources are spent when and where they are most needed can be flexibly specified and applied to a vast array of decision problems, illuminating a number of mental shortcuts. Different disciplines vary in their willingness to trade off psychological realism for formal tractability [86]. Resource-rational analysis targets the middle ground and aims to expand the frontier of possibilities—to achieve a greater mix of realism and tractability than traditionally attainable.

Perhaps the deepest implication of the resource-rational approach is to help fundamentally change the way we treat apparently irrational behavior. Phenomena such as those discussed in this paper are often viewed as strange defects of the human mind, challenging our ability to systematically explain them. Resource-rational analysis provides us with a path forward that respects both the intelligent nature of our judgment and the constraints we face as real flesh and blood creatures. It is a unified lens through which to view the entirety of the mind, including key facets of judgment and decision making alongside core cognitive functions like perception, attention, and memory. And rather than treating biases and errors as bizarre quirks, oddities, or eccentricities of the mind, sometimes we find that they reflect sensible adaptations to cognitive limits, as sure as our eyes adapt to light and dark every day of our lives.

References

- [1] David C Knill and Whitman Richards. *Perception as Bayesian Inference*. Cambridge University Press, September 1996.

- [2] Emanuel Todorov. Optimality principles in sensorimotor control. *Nature Neuroscience*, 7(9):907–915, September 2004.
- [3] Alan Yuille and Daniel Kersten. Vision as bayesian inference: analysis by synthesis? *Trends in Cognitive Science*, 10(7):301–308, July 2006.
- [4] Thomas L Griffiths and Joshua B Tenenbaum. Optimal predictions in everyday cognition. *Psychological Science*, 17(9):767–773, September 2006.
- [5] Mike Oaksford and Nick Chater. Bayesian rationality the probabilistic approach to human reasoning. *Oxford cognitive science series.*, 330, 2007.
- [6] Rafal Bogacz. Optimal decision-making theories: linking neurobiology with behaviour. *Trends in Cognitive Science*, 11(3):118–125, March 2007.
- [7] Michael C Frank and Noah D Goodman. Predicting pragmatic reasoning in language games. *Science*, 336(6084):998, May 2012.
- [8] Dan Ariely. Predictably irrational: The hidden forces that shape our decisions. 280, 2008.
- [9] Gary Marcus. *Kluge: The Haphazard Evolution of the Human Mind*. Houghton Mifflin Harcourt, April 2009.
- [10] Daniel Kahneman. *Thinking, Fast and Slow*. Macmillan, October 2011.
- [11] Lee Ross and Richard E Nisbett. The person and the situation: Perspectives of social psychology. *McGraw-Hill series in social psychology.*, 286, 1991.
- [12] Richard L Lewis, Andrew Howes, and Satinder Singh. Computational rationality: linking mechanism and behavior through bounded utility maximization. *Topics in Cognitive Science*, 6(2):279–311, April 2014.
- [13] Samuel J Gershman, Eric J Horvitz, and Joshua B Tenenbaum. Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, 349(6245):273–278, July 2015.
- [14] Thomas L Griffiths, Falk Lieder, and Noah D Goodman. Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science*, 7(2):217–229, 2015.
- [15] Falk Lieder and Thomas L Griffiths. Resource-rational analysis: understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, 43, 2020.
- [16] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2006.
- [17] Xue-Xin Wei and Alan A Stocker. Mutual information, fisher information, and efficient coding. *Neural Computation*, 28(2):305–326, February 2016.
- [18] Christopher A Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690, April 2003.

- [19] Bartosz Mackowiak, Filip Matejka, and Mirko Wiederholt. Rational inattention: A review. Technical Report 15408, October 2020.
- [20] Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K Warmuth. Occam’s razor. *Information Processing Letters*, 24:377–380, 1987.
- [21] Avrim Blum and John Langford. PAC-MDL bounds. In *Learning Theory and Kernel Machines*, pages 344–357. Springer, 2003.
- [22] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *The 2nd International Conference on Learning Representations*, 2013.
- [23] Ted O’Donoghue and Charles Sprenger. Reference-dependent preferences. In *Handbook of Behavioral Economics: Applications and Foundations 1*, volume 1, pages 1–77. Elsevier, 2018.
- [24] F Attneave. Some informational aspects of visual perception. *Psychological Review*, 61(3):183–193, May 1954.
- [25] Horace B Barlow. Possible principles underlying the transformation of sensory messages. *Sensory communication*, 1:217–234, 1961.
- [26] Antonio Rangel and John A Clithero. Value normalization in decision making: theory and evidence. *Current Opinion in Neurobiology*, 22(6):970–981, December 2012.
- [27] Simon Laughlin. A simple coding procedure enhances a neuron’s information capacity. *Zeitschrift für Naturforschung c*, 36(9-10):910–912, 1981.
- [28] Arthur J Robson. The biological basis of economic behavior. *Journal of Economic Literature*, 39(1):11–33, 2001.
- [29] Barry Wark, Brian Nils Lundstrom, and Adrienne Fairhall. Sensory adaptation. *Current Opinion in Neurobiology*, 17(4):423–429, 2007.
- [30] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 1979.
- [31] Ido Erev, Eyal Ert, and Eldad Yechiam. Loss aversion, diminishing sensitivity, and the effect of experience on repeated decisions. *Journal of Behavioral Decision Making*, 21:575–597, 2008.
- [32] Francesco Rigoli, Karl J Friston, Cristina Martinelli, Mirjana Selaković, Sukhwinder S Shergill, and Raymond J Dolan. A Bayesian model of context-sensitive value attribution. *ELife*, 5:e16127, 2016.
- [33] Francesco Rigoli, Karl J Friston, and Raymond J Dolan. Neural processes mediating contextual influences on human choice behaviour. *Nature Communications*, 7:1–11, 2016.
- [34] Uri Simonsohn and George Loewenstein. Mistake #37: The effect of previously encountered prices on current housing demand. *Economic Journal*, 116:175–199, 2006.
- [35] Christoph Ungemach, Neil Stewart, and Stian Reimers. How incidental values from the environment affect decisions about money, risk, and delay. *Psychological Science*, 22:253–260, 2011.

- [36] Neil Stewart, Nick Chater, and Gordon D A Brown. Decision by sampling. *Cognitive Psychology*, 53(1):1–26, August 2006.
- [37] Timothy L Mullett and Richard J Tunney. Value representations by rank order in a distributed network of varying context dependency. *Brain and Cognition*, 82(1):76–83, June 2013.
- [38] Rahul Bhui and Samuel J Gershman. Decision by sampling implements efficient coding of psychoeconomic functions. *Psychological Review*, 125(6):985–1001, November 2018.
- [39] Allen Parducci. Happiness, pleasure, and judgment: The contextual theory and its applications. 1995.
- [40] Alireza Soltani, Benedetto De Martino, and Colin Camerer. A range-normalization model of context-dependent choice: a new model and evidence. *PLoS Computational Biology*, 8:e1002607, 2012.
- [41] Aldo Rustichini, Katherine E Conen, Xinying Cai, and Camillo Padoa-Schioppa. Optimal coding and neuronal adaptation in economic decisions. *Nature Communications*, 8:1–14, 2017.
- [42] Kenway Louie, Lauren E Grattan, and Paul W Glimcher. Reward value-based gain control: divisive normalization in parietal cortex. *Journal of Neuroscience*, 31(29):10627–10639, 2011.
- [43] Kenway Louie, Mel W Khaw, and Paul W Glimcher. Normalization is a general neural mechanism for context-dependent decision making. *Proceedings of the National Academy of Sciences*, 110(15):6139–6144, April 2013.
- [44] Mel W Khaw, Paul W Glimcher, and Kenway Louie. Normalized value coding explains dynamic adaptation in the human valuation process. *Proceedings of the National Academy of Sciences*, 114(48):12696–12701, 2017.
- [45] Kai Steverson, Adam Brandenburger, and Paul Glimcher. Choice-theoretic foundations of the divisive normalization model. *Journal of Economic Behavior and Organization*, 164:148–165, August 2019.
- [46] Ryan Webb, Paul W Glimcher, and Kenway Louie. The normalization of consumer valuations: Context-dependent preferences from neurobiological constraints. *Management Science*, 2020.
- [47] Luis Rayo and Gary S Becker. Evolutionary efficiency and happiness. *Journal of Political Economy*, 115(2):302–337, April 2007.
- [48] Michael Woodford. Prospect theory as efficient perceptual distortion. *American Economic Review*, 102(3):41–46, 2012.
- [49] Rafael Polanía, Michael Woodford, and Christian C Ruff. Efficient coding of subjective value. *Nature Neuroscience*, 22(1):134–142, January 2019.
- [50] Nick Netzer. Evolution of time preferences and attitudes toward risk. *American Economic Review*, 99(3):937–955, June 2009.
- [51] Joseph A Heng, Michael Woodford, and Rafael Polania. Efficient sampling and noisy decisions. *Elife*, 9, September 2020.

- [52] Keno Juechems, Jan Balaguer, Bernhard Spitzer, and Christopher Summerfield. Optimal utility and probability functions for agents with finite computational precision. 2020.
- [53] Sophie Bavard, Maël Lebreton, Mehdi Khamassi, Giorgio Coricelli, and Stefano Palminteri. Reference-point centering and range-adaptation enhance human reinforcement learning at the cost of irrational preferences. *Nature Communications*, 9:1–12, 2018.
- [54] Francesco Rigoli. Reference effects on decision-making elicited by previous rewards. *Cognition*, 192:104034, 2019.
- [55] Herbert A Simon. Rational choice and the structure of the environment. *Psychological Review*, 63(2):129, 1956.
- [56] Peter M Todd and Gerd Gigerenzer. Bounding rationality to the world. *Journal of Economic Psychology*, 24(2):143–165, 2003.
- [57] Kenway Louie and Paul W Glimcher. Efficient coding and the neural representation of value. *Annals of the New York Academy of Sciences*, 1251(1):13–32, 2012.
- [58] Matteo Carandini and David J Heeger. Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, 13(1):51–62, January 2012.
- [59] Siwei Lyu and Eero P Simoncelli. Nonlinear extraction of independent components of natural images using radial gaussianization. *Neural Computation*, 21(6):1485–1519, June 2009.
- [60] Siwei Lyu. Dependency reduction with divisive normalization: Justification and effectiveness. *Neural Computation*, 23(11):2942–2973, 2011.
- [61] Xue-Xin Wei and Alan A Stocker. Efficient coding provides a direct link between prior and likelihood in perceptual bayesian inference. In F Pereira, C J C Burges, L Bottou, and K Q Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25, pages 1304–1312. Curran Associates, Inc., 2012.
- [62] Xue-Xin Wei and Alan A Stocker. Lawful relation between perceptual bias and discriminability. *Proceedings of the National Academy of Sciences*, 114(38):10244–10249, September 2017.
- [63] Daniel Kahneman. Maps of bounded rationality: Psychology for behavioral economics. *American Economic Review*, 93(5):1449–1475, 2003.
- [64] Michael Woodford. Modeling imprecision in perception, valuation, and choice. *Annual Review of Economics*, 12:579–601, 2020.
- [65] A Aldo Faisal, Luc PJ Selen, and Daniel M Wolpert. Noise in the nervous system. *Nature Reviews Neuroscience*, 9(4):292–303, 2008.
- [66] Alexandre Zenon, Oleg Solopchuk, and Giovanni Pezzulo. An information-theoretic perspective on the costs of cognition. *Neuropsychologia*, 123:5–18, 2019.
- [67] Naama Parush, Naftali Tishby, and Hagai Bergman. Dopaminergic balance between reward maximization and policy complexity. *Frontiers in Systems Neuroscience*, 5:22, May 2011.

- [68] Naftali Tishby and Daniel Polani. Information theory of decisions and actions, 2011.
- [69] Susanne Still and Doina Precup. An information-theoretic approach to curiosity-driven reinforcement learning. *Theory in Biosciences*, 131(3):139–148, September 2012.
- [70] Andrew Caplin and Mark Dean. Behavioral implications of rational inattention with shannon entropy. Technical report, National Bureau of Economic Research, 2013.
- [71] Pedro A Ortega and Daniel A Braun. 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, May 2013.
- [72] Filip Matějka and Alisdair McKay. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–298, January 2015.
- [73] Benjamin M Hébert and Michael Woodford. Rational inattention when decisions take time. Technical report, National Bureau of Economic Research, 2019.
- [74] Weijie Zhong. Optimal dynamic information acquisition. 2019.
- [75] Samuel J Gershman. Origin of perseveration in the trade-off between reward and complexity. *Cognition*, 204:104394, July 2020.
- [76] Samuel J Gershman and Lucy Lai. The reward-complexity trade-off in schizophrenia. *bioRxiv*, 2020.
- [77] Jakub Steiner, Colin Stewart, and Filip Matějka. Rational inattention dynamics: Inertia and delay in decision-making. *Econometrica*, 85(2):521–553, 2017.
- [78] Jonathan Rubin, Ohad Shamir, and Naftali Tishby. Trading value and information in mdps. In *Decision Making with Imperfect Decision Makers*, pages 57–74. Springer, 2012.
- [79] Sander G Van Dijk and Daniel Polani. Informational constraints-driven organization in goal-directed behavior. *Advances in Complex Systems*, 16(02n03):1350016, 2013.
- [80] Matthew Botvinick, Ari Weinstein, Alec Solway, and Andrew Barto. Reinforcement learning, efficient coding, and the statistics of natural tasks. *Current Opinion in Behavioral Sciences*, 5:71–77, 2015.
- [81] Jakub Steiner and Colin Stewart. Perceiving prospects properly. *American Economic Review*, 106(7):1601–1631, July 2016.
- [82] Olivier Gossner and Jakub Steiner. On the cost of misperception: General results and behavioral applications. *Journal of Economic Theory*, 177:816–847, September 2018.
- [83] Ishita Dasgupta, Eric Schulz, Joshua B Tenenbaum, and Samuel J Gershman. A theory of learning to infer. *Psychological Review*, 127(3):412, 2020.
- [84] Cary Frydman and Lawrence J Jin. Efficient coding and risky choice. August 2020.
- [85] Samuel J Gershman and Rahul Bhui. Rationally inattentive intertemporal choice. *Nature Communications*, 11(1):3365, July 2020.

- [86] Matthew Rabin. A perspective on psychology and economics. *European Economic Review*, 46(4):657–685, May 2002.

References of Interest

* of special interest

** of outstanding interest

[15] Lieder and Griffiths 2020** This modern review of resource rationality discusses a wide variety of applications, accompanied by commentary from many scholars.

[19] Mackowiak, Matejka, and Wiederhold 2020** This modern review provides a clear introduction to capacity-limited decision making, laying out its assumptions and implications in-depth and describing its applications to subfields of economics.

[38] Bhui and Gershman 2018* This theoretical article reveals how classic, influential models of context-sensitive economic judgment based on rank and range normalization can be derived from an efficient coding analysis.

[49] Polania, Ruff, Woodford 2019* This study finds that a theory of efficient coding and Bayesian decoding previously applied to perception can account for intricate patterns of bias in representations of value.

[51] Heng, Woodford, and Polania 2020* This study develops models of efficient coding that maximize accuracy and reward in a numerosity discrimination task, and compares their empirical performance to decision by sampling, which fits the data better.

[66] Zenon, Solopchuk, and Pezzulo 2019* This article elaborates on the connections between information costs and cognitive costs, translating between computational quantities and psychological content.

[78] Rubin, Shamir, and Tishby 2012* This piece characterizes a capacity-constrained version of reinforcement learning which the authors apply to navigation, demonstrating how efficient action policies are stochastic and may be altered to compensate for this stochasticity.