Toward a precise theory of imprecise representations

Arthur Prat-Carrabin[1,⋆] and Samuel J. Gershman[1]

[1]Department of Psychology and Center for Brain Science, Harvard University, Cambridge, MA

⋆*email: arthurpc@fas.harvard.edu*

April 9, 2026

Abstract

Models in cognitive science typically explain behavioral patterns by positing the kinds of computations, often presumed suboptimal, that the brain might implement. Implicitly, these computations are assumed to operate on representations that faithfully reflect the relevant information. In the brain, however, representations are imprecise, with broad repercussions for many cognitive processes. Here we develop the basis for a theory of imprecise representations, and we draw out some of its implications. From a single principle of efficient coding, we derive four quantitative laws of psychophysics, which we show are verified in data. These results, supported by neural evidence, suggest that imprecise representations are strongly context-dependent, and help explain many behavioral patterns across perception and decision-making.

A venerable theoretical tradition in cognitive science treats the mind as an information processing system that performs computations over internal representations. Cognitive scientists thus seek to infer people's computational strategies in various tasks, such as: how they search through problem space in problem-solving tasks; how they update expected values, in reinforcement-learning scenarios; how they weigh gains and losses and distort probabilities in decisions under uncertainty; and how they implement belief updates in inference tasks. Many theories propose explanations regarding how people manipulate the information presented to them, implicitly assuming that this information is perfectly represented. Consider, for example, risky choices between lotteries. Here the information to be processed is minimal: just a few unambiguous numbers (e.g., '\$1 with probability 0.5'). In traditional models, these numbers undergo in the mind of the decision-maker a series of transformations (e.g., utility function, probability-weighting function, and comparison) that determine the decision. Behavioral patterns are then explained as resulting from the properties of these computations.

We propose instead that the information presented in most tasks is not precisely represented in the brain, with far-reaching consequences for decision-making. Numbers,

Highlights

Human cognition and decision-making rely on internal representations that are inherently imprecise.

Recent theoretical, behavioral, and neurophysiological results provide a coherent quantitative account of such imprecise representations in the case of magnitudes.

A single endogenous efficient-coding principle yields four quantitative psychophysical laws, linking prior, task demands, imprecision, and bias, and supported by evidence across multiple paradigms.

Representations are context-dependent, adapting to both stimulus statistics and the observer's objective.

Neural evidence supports flexible range adaptation, with coding resources dynamically reallocating across stimuli.

Imprecise representations explain a broad range of behavioral patterns.

to continue our example, are mapped to internal representations that are imprecise (in ways we detail below), such that the numbers cannot be unambiguously recovered from them. Yet it is on these 'fuzzy' representations that decisions are made. In risky choices, behavioral and neurophysiological results demonstrate that cognitive imprecision accounts for choice behavior[1–7], in departure from traditional theories of risk attitudes. This suggests that before studying complex human decisions, it is useful to first understand how the brain represents the elements that inform them.

A theory of representations would provide sound foundations for more complex cognitive models of decision, inference, learning, memory, etc. We argue that such a theory is within reach, at least for magnitudes, such as dollar amounts and probabilities, as in the example above, but also other kinds such as durations, distances, temperatures, ratings, and so on, and potentially more abstract concepts that lend themselves to comparisons (e.g., "better than", "tastier than", etc.). Such scalar representations are fundamental to many perceptual and decision tasks, including many economic decisions; accordingly, behavioral economists have shown growing interest in various kinds of imperfect information processing[8–18]. We present the theoretical efforts undertaken in this direction, the experimental results supporting them, and the predictions that remain to be tested.

The theory we present, in its current state, is a careful formalization of the idea that the brain forms optimal representations, subject to the constraints under which it operates. This perspective is related to **efficient coding** (see Glossary) in neuroscience[19–30], computational or resource rationality in cognitive science[31–34], and rational inattention in economics[8–11]. Each of these frameworks proposes that the brain optimizes some objective under some constraint; but none seem to agree on the objective nor on the constraint[35]. Yet these elements determine representations in these models, and should therefore be clearly identified. Recent theoretical and experimental work has made progress in the simple case of magnitudes, with important implications for psychophysics and decision-making.

Indeed this theoretical framework predicts laws of psychophysics, including longestablished laws (e.g., **Weber's law**), but also new laws that we show are verified in data, and new laws that have yet to be tested. In particular, it predicts non-trivial scaling laws for the variability of percepts and for biases. We emphasize that all the behavioral effects we present here are *derived* from first principles (in fact, from a single equation). The goal is to unearth these fundamentals, rather than to exhibit behavioral curiosities. We also point to neural evidence that directly supports the theory.

The most important lesson that emerges from this body of work is that representations are exquisitely sensitive to the context in which they are needed. For instance, the representation of a magnitude is different depending on whether it is used for estimating the magnitude, or for comparing it with another magnitude. In other words, representations depend on what they are used for. In addition, they depend on the statistics of the represented information (i.e., the **prior**). The argument is not only that context influences *decisions* (as Bayesian decision theory predicts), but that it determines *representations*, on which reasoning and decision-making operate. This sensitivity has important implications for the study of human cognition. Careful attention should be paid to the precise conditions under which behavior is elicited; these conditions shape representations, and therefore they shape cognition, and behavior.

A theory of magnitude representations: endogenous efficient coding

Given an observer presented with a magnitude x, our objective is to formulate some description of the representation r elicited in the observer's brain. The theory we present instantiates the principle of efficient coding, i.e., the notion that the brain chooses optimal representations, under constraints. This is formalized as an optimization problem. We then show that from this single optimization problem follow four laws of psychophysics that are widely supported in data.

Characterizing representations To define the optimization problem, we must specify what is optimized over. Here this is the 'encoding', i.e., the distribution $p(r|x)$ of the representation r conditional on the magnitude x. In economics, this would be called the 'signal structure'. In neuroscience, this means for instance choosing neuronal **tuning curves**, i.e., specifying how the (stochastic) neural activity encodes the magnitude, a complex and often ill-defined problem. Instead we characterize the encoding by its **Fisher information**, $I(x)$, which quantifies its precision (Box 1). Different neural implementations may induce the same Fisher information, and are therefore equivalent for our purposes: working with $I(x)$ allows us to abstract away from idiosyncratic features of the code (e.g., the precise parametrization of tuning curves) and to formulate general principles of neural coding.

This provides a compact description of representations while retaining direct links to behavior: in the small-noise regime, the standard deviation of estimates and the discrimination threshold, or just-noticeable difference (JND), are both approximated by $\sigma(x) = 1/\sqrt{I(x)}$, and it also defines an important component of the **bias** (see Box 1 and Eq. 19). We make one structural assumption on the encoding: we assume that it results from the accumulation of n independent, identically-distributed signals ('samples'), each with Fisher information $I_1(x)$. The resulting total Fisher information is the sum $I(x) = nI_1(x)$. Increasing n improves the overall precision of the encoding, while by choosing the function $I_1(x)$, under some constraint, detailed below, one can choose to be more precise about some stimuli than about others.

Box 1: Fisher information: key properties and implications

The neuroscience of perception suggests that representations materialize in the activity of sensory neurons, whose stochastic firing patterns depend on stimuli, which are thereby thought to be *encoded*. The precision of the representation is therefore the precision of this neural code. Two stimuli will easily be distinguished if the neural patterns they elicit are different. The Fisher information, $I(x)$, formalizes this idea. It is a function of the derivative of the probability distribution characterizing the neural activity, $p(r|x)$, with respect to the encoded stimulus, x. Thus it measures the sensitivity of the encoding to x. Formally,

$$I(x) = \mathbb{E}\left[\left(\tfrac{\mathrm{d}}{\mathrm{d}x}\log p(r|x)\right)^2 \;\middle|\; x\right]. \tag{1}$$

Properties The Fisher information has several important and useful properties. Consider a stimulus x, encoded in n stochastic signals collectively denoted by r, with total Fisher information $I(x)$. For the Bayesian-mean estimator, $\hat{x}(r) = \mathbb{E}[x|r]$, (i) the variance of the estimator, $\mathrm{Var}[\hat{x}|x]$, (ii) its mean squared-error, $\mathbb{E}[(\hat{x}-x)^2|x]$, and (iii) the variance of the Bayesian posterior, $\mathrm{Var}[x|r]$, are all approximately (with large n and under regularity conditions[29,36])

$$\sigma^2(x) = \frac{1}{I(x)}. \tag{2}$$

We clarify that here $\sigma(x)$ is not the standard deviation of x: rather, it is the imprecision of the estimator when the stimulus is x, where 'imprecision' can be understood in any of the senses (i-iii) above. (We focus on the first two, but experimental measures of the third one, i.e., of observers' uncertainty, indeed correlate with the first two[37].) This property is useful to characterize the loss in an estimation task (e.g., quadratic loss; see Box 2) as well as the behavioral data. Indeed $\sigma(x)$ approximates the standard deviation of estimates, in estimation tasks, and it can be shown that it is proportional to the discrimination threshold (or JND), in discrimination tasks.

Another important property is that the variation of the Fisher information across stimuli affect the bias of the estimator, by pulling it towards the *less-precisely* encoded stimuli ('encoding effect'; see Fig. I in this box, and Eq. 19).

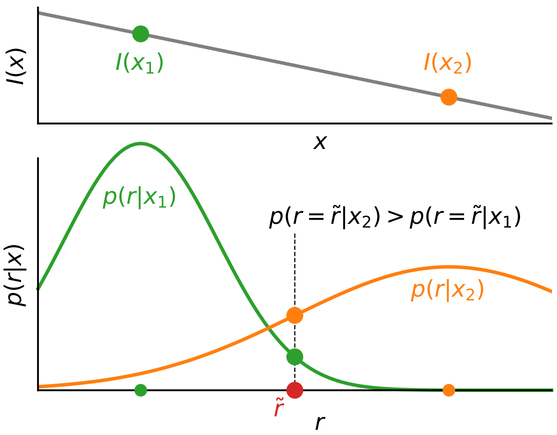

Fig. I: Encoding effect: a bias toward less-precisely encoded stimuli. In this illustration, the Fisher information for x_1 is greater than for x_2 (top panel). Thus the representation r is more precise for x_1 than for x_2 (narrower distribution $p(r|x)$; bottom panel). Given a realization $\tilde{r}$ of the representation that could result from either x_1 or x_2, the likelihood of x_2 is greater than that of x_1, (i.e., $p(r = \tilde{r}|x_2) > p(r = \tilde{r}|x_1)$), due to the stronger specificity of the encoding for x_1. This greater likelihood induces a bias toward x_2 over x_1.

Another property of the Fisher information that we use is that for independent signals it simply sums, i.e., n signals with Fisher information $I_1(x)$ together have total Fisher information $I(x) = nI_1(x)$.

Capacity constraint Finally, given a differentiable transformation $z = \varphi(x)$, the Fisher information with respect to z, $I_z(z)$, is related to the Fisher information with respect to x, $I_x(x)$, as

$$I_x(x) = I_z(\varphi(x))(\varphi'(x))^2. \tag{3}$$

Informally, the sensitivity to x is equal to the sensitivity to $z = \varphi(x)$ times the sensitivity of z to x (chain rule). An important implication arises if the encoding must occur through a fixed, unchanging 'channel', $p_z(r|z)$, such that $p(r|x) = p_z(r|z = \varphi(x))$, with φ invertible, and where z describes for instance some internal, psychological scale. From this assumption follows the inequality[29]:

$$\int_{D_x} \sqrt{I_x(x)}\mathrm{d}x = \int_{\varphi(D_x)} \sqrt{I_z(z)}\mathrm{d}z \leq \int_{D_z} \sqrt{I_z(z)}\mathrm{d}z = \text{constant}, \tag{4}$$

where D_x is the domain of x and D_z that of z, and where the last integral is assumed finite. Thus even if one is free to choose the 'transducer' function, φ, the fixed-channel assumption puts a capacity constraint on feasible encoding schemes. Conceptually, the capacity constraint is that the total Fisher information across the stimulus space is bounded. This is this constraint assumed in Eq. 11.

Loss function The choice of representation is guided by a loss function. We posit that representations are chosen to optimize what the observer is currently trying to do. This may seem like an obvious choice, but it has profound implications, and in many models it is not the chosen objective. Many efficient-coding models assume that the brain maximizes the **mutual information** (MI) between stimulus and representation. This is not an objective that corresponds in any obvious way to a goal one might have in everyday life; but it may be a sensible general-purpose objective, appropriate in early sensory systems[26]. However, a fixed objective such as MI implies that representations do not change when the observer's goal changes; but experimental evidence, presented below, suggests that they do.

A more flexible hypothesis is that representations are adapted to the task at hand. Myriad tasks are conceivable; here we model a subset of them (Box 2). Each task induces a loss, which can be expressed as a functional of the Fisher information, $L_{task}[I]$. These tasks are naturally organized using two indices. The first index, $m \in \{1, 2\}$, is the number of relevant stimuli in the task: $m = 1$ in **estimation tasks**, where the goal is to estimate the stimulus magnitude x; $m = 2$ in **choice tasks**, where the goal is to choose the stimulus $x \in \{x_1, x_2\}$ with greater expected utility (often greater magnitude, in psychophysical experiments). The second index, $p > 0$, determines how errors are penalized; in estimation tasks, it indicates the p-norm error (see Box 2; below, unless otherwise specified, we consider squared errors, i.e., $p = 2$). The loss also depends on the stimulus statistics, i.e., the prior. We denote by $L_{f,m,p}[I]$ the loss function associated with a task characterized by m and p, in the context of a prior f.

Box 2: A taxonomy of tasks, and corresponding loss functions

Different tasks imply different notions of error, and therefore different optimal representations. In the small-noise regime, these task-dependent losses can be expressed as functionals of the Fisher information, $I(x)$.

Estimation tasks ($m = 1$, $p = 2$) A magnitude x is presented and the observer must report it as accurately as possible. For the sake of simplicity we use the term *estimation task* to refer specifically to the case of squared-error loss. In the small-noise regime, the mean squared error is approximately the inverse of the Fisher information (see Box 1, Eq. 2). The expected loss is then

$$L_{est}[I] = \frac{1}{2} \int \frac{f(x)}{I(x)} \, \mathrm{d}x. \tag{5}$$

Other p-norm error metrics lead to the same structure but with $1/I(x)^{p/2}$.

Discrimination tasks ($m = 2$) These tasks involve comparing two magnitudes. Because performance depends on pairs of stimuli, the relevant prior is the joint probability of drawing two nearby magnitudes ($x \approx y$), which is $f(x)f(y) \approx f(x)^2$.

Choice tasks ($p = 2$). Two magnitudes are presented; the observer chooses one; and the reward equals the chosen magnitude (or is proportional to it). Errors are therefore more costly when the difference is large, and less so when the two options are similar. In the small-noise regime, this yields[29]

$$L_{choice}[I] = \frac{1}{2} \int \frac{f(x)^2}{I(x)} \, \mathrm{d}x. \tag{6}$$

Classification tasks ($p = 1$). If instead a fixed reward is given for correctly identifying the greater of two magnitudes, all mistakes are penalized equally. In that case,[29]

$$L_{class}[I] = \int \frac{f(x)^2}{\sqrt{I(x)}} \, \mathrm{d}x. \tag{7}$$

We emphasize that unlike choice tasks, confusing two nearly equal magnitudes is just as costly as any other error.

A common family These examples can be unified in the family[29]

$$L_{f,m,p}[I] = \frac{1}{p} \int \frac{f(x)^m}{I(x)^{p/2}} \, \mathrm{d}x. \tag{8}$$

The index m reflects how many stimuli jointly matter for the objective ($m = 1$ for estimation, $m = 2$ for discrimination), while p reflects how strongly errors are penalized.

Mutual-information objective A commonly used alternative in efficient-coding models is to maximize the mutual information between stimulus and representation. If the representation consists of n signals, one has in the large-n limit[38]

$$MI(X, R) = \frac{1}{2} \int f(x) \ln I(x) \, \mathrm{d}x + H(X) - \frac{1}{2} \ln(2\pi e) + o(1), \tag{9}$$

where $H(X)$ is the entropy. This corresponds to the limiting case

$$L_{MI}[I] = \lim_{p \to 0} \frac{1}{p} \int \frac{f(x)}{I(x)^{p/2}} \, \mathrm{d}x, \tag{10}$$

i.e., $m = 1$ and $p \to 0$ within the same family.

These different objectives imply different optimal allocations of representational resources across stimuli. As a result, the same stimulus may be encoded with different precision depending on the task, leading to distinct patterns of variability and bias.

Capacity constraint We have assumed that the encoding results from the accumulation of n samples. Each sample brings information about the stimulus, x, and helps minimize the loss; but we posit a limit on the information $I_1(x)$ provided by a sample. Specifically, we assume the constraint

$$\int \sqrt{I_1(x)}\mathrm{d}x \leq \sqrt{K}. \tag{11}$$

This constraint is used elsewhere[24,25,27–29,39]; we explain in Box 1 why it is a natural one. Informally, it means that the encoding cannot be precise about all stimuli. It can be seen as a constraint on a 'precision budget' (notwithstanding the possibility of sampling several times, to which we turn below). For instance, if the encoding is carried by a single Poisson neuron with a monotonic tuning curve, this constraint is equivalent to assuming an upper bound on its firing rate.

Sampling cost We posit that each sample comes with a cost. This could represent, for instance, the energy required to fire action potentials[40–42]; it can also be construed more theoretically as an opportunity cost corresponding to the loss incurred in all other tasks by employing resources in this particular one[34]. Here we posit that the cost of one sample is λ, and thus the cost of n samples is λn.

Optimization problem and solution Putting together these ingredients yields an efficient-coding problem in which the observer chooses how to allocate encoding precision across stimuli, and how many samples to collect, so as to minimize the task loss under a sampling cost and a constraint on feasible encoding schemes (Box 3). The resulting optimal Fisher information, for task losses of the form $L_{f,m,p}[I]$, is[43]

$$I(x) = \frac{f(x)^{\gamma_{task}}}{C_{f,m,p}}, \tag{12}$$

where

$$\gamma_{task} = \frac{2m}{p+1}, \tag{13}$$

and $C_{f,m,p}$ is a constant (see Box 3).

This shows, first, that the encoding should be more precise for more frequent stimuli: encoding resources 'spread out' across the stimulus space according to the relative frequencies of stimuli—a typical prediction of efficient coding. Second, the Fisher information depends not only on the prior but also on m and p: i.e., the representation depends on the task. We now present laws of psychophysics that follow from Eq. 12.

Box 3: Endogenous efficient coding

Combining the task loss (Box 2), the capacity constraint (Eq. 11), and the cost of sampling yields the optimization problem[43]

$$\min_{n,I_1} L_{f,m,p}[nI_1] + \lambda n \text{ subject to } \int \sqrt{I_1(x)}\mathrm{d}x \leq \sqrt{K}. \tag{14}$$

The optimal encoding takes the form

$$I(x) = \frac{f(x)^{\gamma_{task}}}{C_{f,m,p}}, \quad \text{where} \quad \gamma_{task} = \frac{2m}{p+1}, \tag{15}$$

and where $C_{f,m,p}$ does not depend on x (see Supplemental Information).

Efficient coding The dependence $I(x) \propto f(x)^{\gamma_{task}}$ reflects the efficient allocation of precision across stimuli: more frequent stimuli are encoded more precisely. The exponent γ_{task} determines how concentrated this allocation is. Larger values of γ_{task} place relatively more precision on the most frequent stimuli, while smaller values spread precision more evenly across the stimulus space. Because γ_{task} depends on (m, p), the optimal allocation of precision is task-dependent (see Law #1).

Endogenous precision The number of samples n depends on the prior, and is therefore part of the optimal solution. For instance, with a uniform prior of width w, one obtains $n \propto w$ in estimation tasks, and $n \propto \sqrt{w}$ in choice tasks: broader ranges lead to more samples. Why should that be the case? As the prior widens and encoding resources are increasingly spread thin, errors get larger as well; to mitigate this, it is optimal to 'pay' more in encoding resources (i.e., to increase n). This endogenous adjustment of n underlies the sublinear scaling of imprecision (Law #1), and the quantile invariance of the bias (Law #4). We note that the resulting sublinear scaling means that representations are *relatively* more precise with wider priors. This implies, for instance, that the scale used to communicate a quantity may affect the precision with which it is perceived.

Laws of psychophysics

• Law #1: task-dependent sublinear scaling of imprecision.

Because encoding resources are costly and must be allocated across stimuli, wider priors yield greater imprecision. But how the imprecision scales with the prior width depends on the observer's objective, i.e., on the task. For uniform priors of width w, one obtains $I(x) \propto 1/w$ in estimation tasks and $I(x) \propto 1/w^{3/2}$ in choice tasks. Therefore the standard deviation of estimates in estimation tasks and the discrimination threshold in choice tasks verify, respectively, the following laws first derived by Prat-Carrabin & Woodford[43]:

$$\sigma_{est} \propto \sqrt{w}, \tag{16}$$

$$\text{and} \quad \sigma_{choice} \propto w^{3/4}. \tag{17}$$

These laws relate the scale of the internal imprecision to the scale of the prior. The scaling is sublinear, and task-dependent: if the prior width doubles, the imprecision is multiplied by $\sqrt{2} \approx 1.41$ in estimation tasks, and by $2^{3/4} \approx 1.68$ in choice tasks. These theoretical predictions are verified quantitatively in numerosity estimation, motion-direction estimation, numerosity discrimination, and risky choice[3,43–45] (Fig. 1a,b, Key Figure).

Several lessons can be drawn from these results. First, representations adjust to changing priors, suggesting that efficient coding is dynamic rather than fixed to a long-run prior. This is consistent with evidence of prior sensitivity in perceptual tasks involving orientations, colors, shapes, and numerosity[39,46–49]. Second, representations depend on the task. This is consistent with evidence that sensory encoding is sensitive to reward structure[50–53], and more generally that task demands shape representations[54–56]. Psychophysical measurements such as JNDs or biases are therefore not intrinsic, immutable characteristics of subjects, but can vary with the objective of the observer, and the JND measured in one task may not equal the JND measured in another task. Third, perhaps counter-intuitively, the scaling is sublinear (by contrast, 'range normalization' models[57–60] predict proportionality, $\sigma \propto w$). The reason is that the optimal number of signals, n, also grows with the prior width, thereby increasing the precision, partially compensating for the spreading of precision over a wider range (see Box 3). More generally, precision is a property of the brain's representations that is endogenously chosen to optimally adapt to context. A similar idea has been proposed in a model of intertemporal choice[61].

• Law #2: task-dependent generalized Weber's law.

Above we considered how the encoding changes across priors of different sizes. We now turn to its variations across stimuli. We consider priors in the 'shifted power law' family, i.e., with density $f(x) \propto 1/(x+d)^{\alpha}$. For these priors we obtain (see Eqs. 2 and 12) the generalized Weber's law:

$$\sigma(x) \propto (x+d)^{\frac{1}{2}\alpha\gamma_{task}}. \tag{18}$$

In discrimination tasks, where σ is interpreted as the discrimination threshold, Weber's law corresponds to a proportionality relationship ($\sigma(x) \propto x$). Its generalization as in Eq. 18 accounts very well for discrimination data across different sensory modalities[72]. Here we thus provide an efficient-coding rationale for Weber's law (see also Ref. 73). Two further predictions follow: the scaling relationship is determined by the shape of the prior (via α and d), and by the task (via $\gamma_{task} = \frac{2m}{p+1}$). For several sensory modalities, studies have successfully modeled the prior as a generalized power law, and the encoding as a Fisher information proportional to a power of that prior[26,30], consistent with the model above over several orders of magnitude (Fig. 1c). But a careful examination of whether and how the scaling of the imprecision quantitatively varies with the shape of the prior remains to be conducted.

Similarly, to our knowledge a direct test of the sensitivity to the task has not been conducted. A prediction is that in choice tasks ($\gamma_{choice} = 4/3$) the imprecision should increase with the magnitude faster than in estimation tasks ($\gamma_{est} = 2/3$). This calls for experiments directly comparing the two tasks, with the same sensory modality (and the same prior). In the studies just mentioned, the operating hypothesis is that the encoding maximizes the MI, corresponding to $\gamma_{MI} = 2$ ($m = 1$ and $p \to 0$; the exponent $\gamma = 1/2$, also tested, poorly fits the data[26]). But the same exponent—thus the same allocation of precision, and thus the same perceptual behavior—is obtained in **classification tasks**

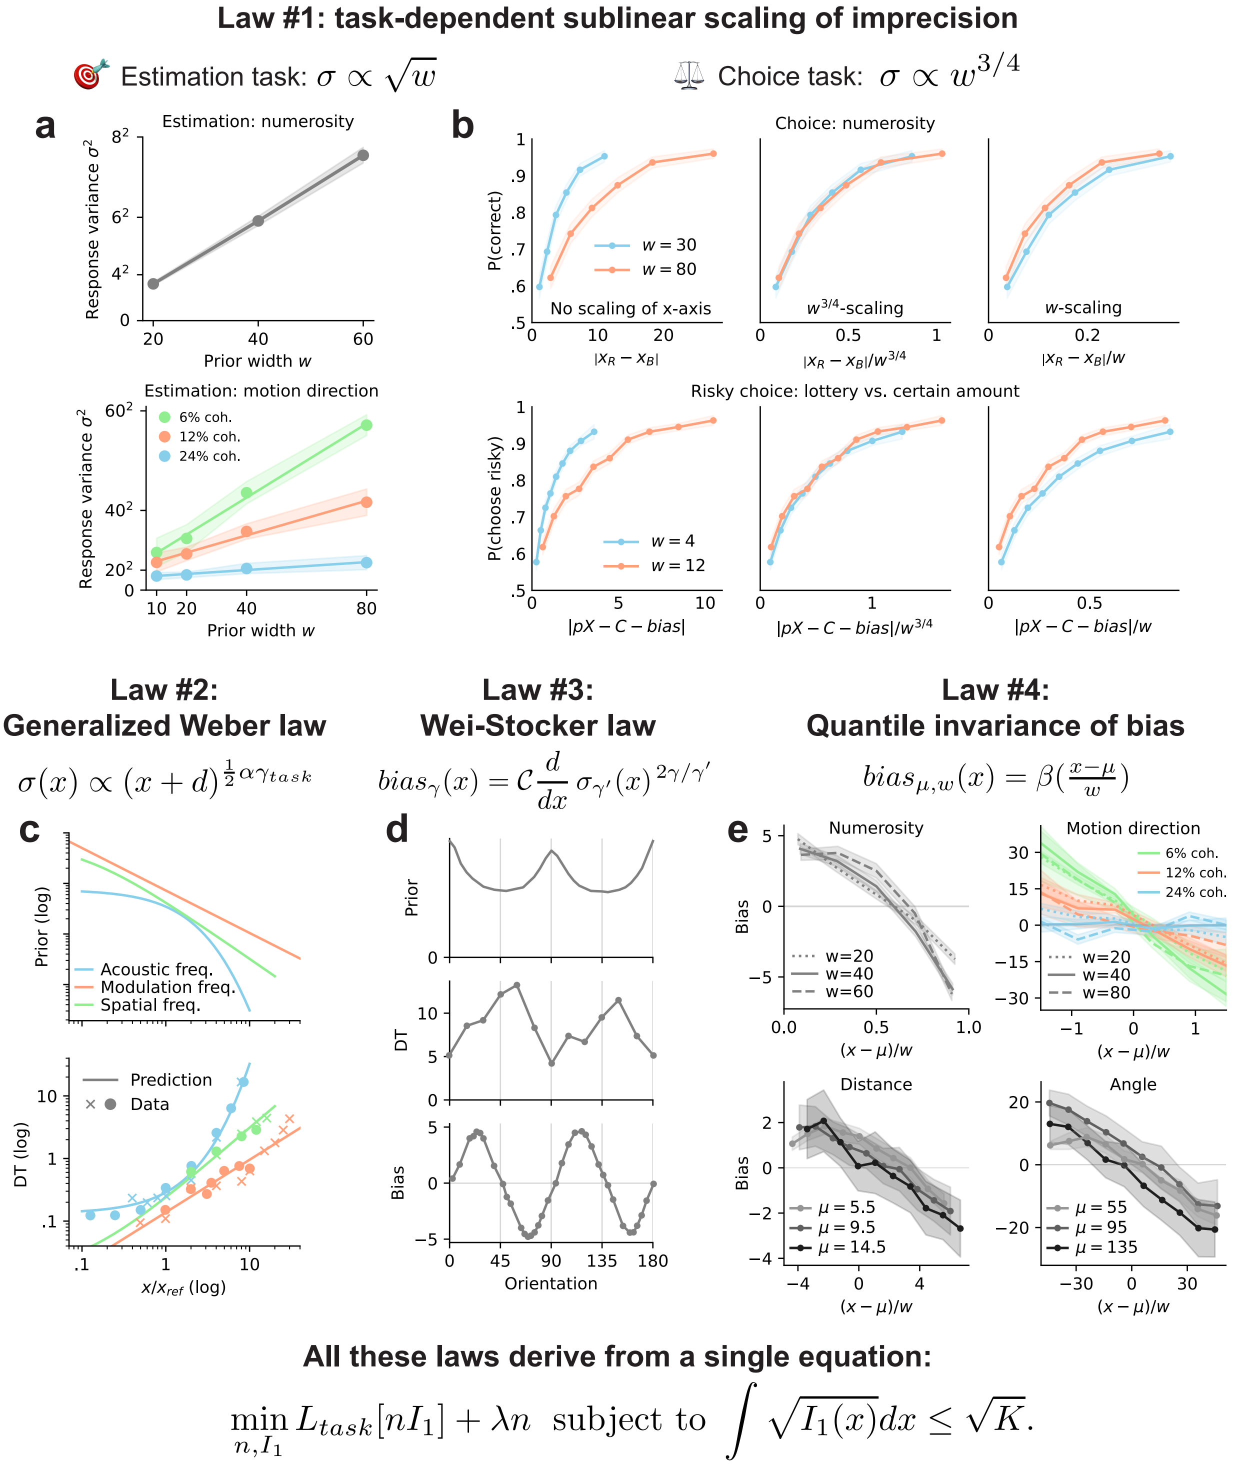

Fig. 1: Key Figure. Laws of psychophysics derived from endogenous efficient coding. **a.** Response variance increases linearly with the prior width in a numerosity estimation task[43] (top panel) and in a motion direction estimation task[45] (random dot kinematograms with different levels of coherence; bottom panel). (*caption continued on next page*)

Fig. 1: (cont.) **b.** Choice probabilities under priors of different widths in a numerosity discrimination task[43] (top row) and in a risky choice task[3] (bottom row), as a function of the absolute difference between the two options (left column), of this difference divided by the prior width raised to the exponent $3/4$ (middle column), and divided by the prior width (right column). The first column shows that participants are less precise with the wider priors, in absolute terms. The last column shows that, in *relative* terms, they are more precise with the wider priors. The middle column shows that their imprecision scales as the prior width to the power $3/4$. **c.** Generalized Weber law: the prior (top panel) determines the discrimination threshold (bottom panel; predictions (solid lines) are derived from priors via Eq. 12 with $\gamma = 2$). Priors and data analysis from Ref. 26; data for acoustic frequency, auditory modulation frequency, and spatial frequency from Refs. 62–64 (crosses) and 65–67 (dots). **d.** Wei-Stocker law[68]: the discrimination threshold for visual orientations[64] (middle panel) increases where the bias[69] (bottom panel) is positive, and it decreases where the bias is negative. The discrimination threshold moreover reflects the prior[70] (top panel) and is lowest for the more frequent orientations, consistent with efficient coding. **e.** Bias as a function of the centered, normalized stimulus $(x - \mu)/w$ in estimation tasks, with numerosity[43] and motion direction[45] where the prior width, w, is manipulated (top row), and with distance and angle[71] where the prior center, μ, is manipulated (bottom row).

($m = 2$, $p = 1$, and thus $\gamma_{class} = 2$; see Box 2). The data are thus consistent both with MI maximization, and with the minimization of a classification-task loss, since $\gamma_{MI} = \gamma_{class} = 2$ (experimental reports, unfortunately, rarely specify the actual reward structure). Disentangling these explanations would require manipulating experimentally the loss function and observing the resulting perceptual decisions. Some results mentioned above suggest that the encoding is task-sensitive[43,50–53], but more direct tests are warranted to ascertain whether representations are generally adaptive, or whether some sensory systems invariably maximize the MI, regardless of the context.

Bias of the Bayesian mean

Imprecision generates variability, but also biases: a Bayesian decoder, typically, is biased by the prior toward more frequent stimuli, but it is also biased by the specifics of the encoding. Consider the Bayesian-mean estimator (which is optimal for estimation and choice tasks). In the small-noise regime, when a series of independent (but not necessarily identically-distributed) signals encode the stimulus x, the bias of the Bayesian mean, $\mathbb{E}[\hat{x} - x|x]$, is approximately[29]

$$bias(x) = \frac{1}{I(x)}\left[\frac{f'(x)}{f(x)} - \frac{I'(x)}{I(x)}\right]. \tag{19}$$

A generalization exists for the optimal p-norm loss decoder, under assumptions on the encoding[74]. The scaling term is the inverse Fisher information: more precise representations yield smaller biases. The term f'/f is a 'prior effect' which pulls estimates toward more frequent stimuli. Perhaps less intuitive, the 'encoding effect', $-I'/I$, pulls the Bayesian mean toward the less-precisely encoded stimuli. We illustrate this in Fig. I (Box 1).

As we have seen, if the encoding is optimally adapted to the prior and to the task, then the Fisher information is proportional to a power γ_{task} of the prior (Eq. 12), and thus the encoding effect is proportional, but opposite, to the prior effect, i.e., $-I'/I = -\gamma_{task}(f'/f)$. Depending on γ_{task}, i.e., on the specifics of the task, the encoding effect may dominate, thus accounting for 'anti-Bayesian' biases[24], directed toward the less probable stimuli. Furthermore, the bias depends on the variations of the Fisher information, and therefore on the variations of the imprecision; we now turn to this relationship.

• Law #3: Wei & Stocker law relating bias and imprecision.

The expressions of the bias (Eq. 19) and of the optimal Fisher information (Eq. 12) imply a relationship between the bias and the derivative of the imprecision. Before detailing this relationship, we note that for a given sensory modality, different behavioral measures may be determined with different tasks, e.g., biases with estimation tasks, and JNDs with discrimination tasks. But as noted above, these measures depend on the tasks themselves. Thus in order to make a precise statement about the relation between bias and imprecision, we consider the case in which the former, $bias_\gamma$, is measured in a task with exponent γ, while the latter, $\sigma_{\gamma'}$, is measured in a task with exponent γ'. Equations 12 and 19 imply:

$$bias_\gamma(x) = \mathcal{C}\frac{\mathrm{d}}{\mathrm{d}x}\,\sigma_{\gamma'}(x)^{\,2\gamma/\gamma'}, \tag{20}$$

where $\mathcal{C}$ is a constant (see Supplemental Information).

Wei & Stocker demonstrate that this relationship enjoys qualitative experimental support across many sensory modalities[68]. Specifically, they find that alternations in the sign of the bias coincide with alternations in the variations of the imprecision—which themselves coincide with alternations in the variations of the prior (see Fig. 1d). They formulate a version of this law, $bias(x) \propto \frac{\mathrm{d}}{\mathrm{d}x}\sigma^2(x)$, which implicitly assumes that bias and imprecision are measured with the same task, although the results they report rely on many studies with different experimental protocols.

Empirically, the data suggest a positive coefficient for most modalities, though not all[68]. But if the encoding is optimized to a squared-error loss in the task measuring the bias, then Eq. 20 predicts a negative coefficient instead (see Supplemental Information). This suggests that participants may not minimize the squared error, in these studies. We note that in none of them is the reward tied to the estimation error, and in most cases participants receive no feedback (which would allow them to learn the objective of the task). Most of these tasks are matching tasks ('method of adjustments'); recent work suggests that behavior in such tasks is better captured by an objective that includes a categorical-loss component, in addition to an error term[75]. This may yield a different optimal allocation of precision, and thereby a different sign of the coefficient. Another possibility, as noted above, is that some low-level sensory systems implement a fixed objective, such as MI maximization, rather than adapting to the task. Distinguishing between these explanations will require carefully designed experiments in which the objective is explicitly manipulated and controlled.

• Law #4: Quantile invariance of the bias.

Finally, we note a 'quantile invariance' property of the bias, for location-scale families of priors parameterized by a location, μ, and a scale, w. In estimation tasks, the bias is then a function of the quantile of the magnitude, and not of the specific prior. Specifically, Eqs. 12 and 19 entail:

$$bias_{\mu,w}(x) = \beta(\tfrac{x-\mu}{w}), \tag{21}$$

where $\beta(.)$ does not depend on the prior.

Consider two priors with different scales, $w \neq w'$. Quantile invariance implies that the bias for a magnitude $2w$ away from the mean, under the first prior, equals the bias for a magnitude $2w'$ away from the mean, under the second prior. This property appears verified in numerosity and orientation estimation data[43,45] (Fig. 1e, top row). This invariance is not trivial, and follows from the endogenous adjustment of precision (modeled as the ability to choose the number of samples, n). If the imprecision increased linearly with the scale of the prior, the bias for a given quantile would also increase linearly; instead, the sublinear increase (Law #1) makes the bias invariant across priors.

Similarly, with priors that differ only in location ($\mu \neq \mu'$ but $w = w'$), the bias is predicted to depend only on the quantile (equivalently, on the relative location $x - \mu$). This is roughly consistent with magnitude estimation experiments that test ranges of equal size, but shifted[76] (Fig. 1e, bottom row).

Flexible neural representations

The endogenous efficient-coding theory has important implications for neural coding. Notably it implies that the neural code is not static: when the prior changes, neural resources should reorganize accordingly. In the brain, most magnitudes are encoded by vast networks of neurons with bell-shaped **tuning curves**, each characterized by its 'preferred stimulus' (i.e., the stimulus to which a neuron maximally reacts). The hypothesis of **distributed range adaptation** posits that the distribution of neurons' preferred stimuli flexibly 'stretches' to match the current prior (Fig. 2a,b). Consistent with this prediction, a recent fMRI study shows that when the prior changes, the receptive fields of numerosity-sensitive parietal populations collectively shift their preferred numerosities, thereby efficiently covering the stimulus space[49] (Fig. 2). These neural changes, moreover, are predictive of corresponding changes in behavior. A similar flexible coding has been reported for the representation of time in the rat's striatal and hippocampal populations[77,78]. Such adaptive coding may be implemented through context-dependent gain modulation in recurrent sensory networks[44].

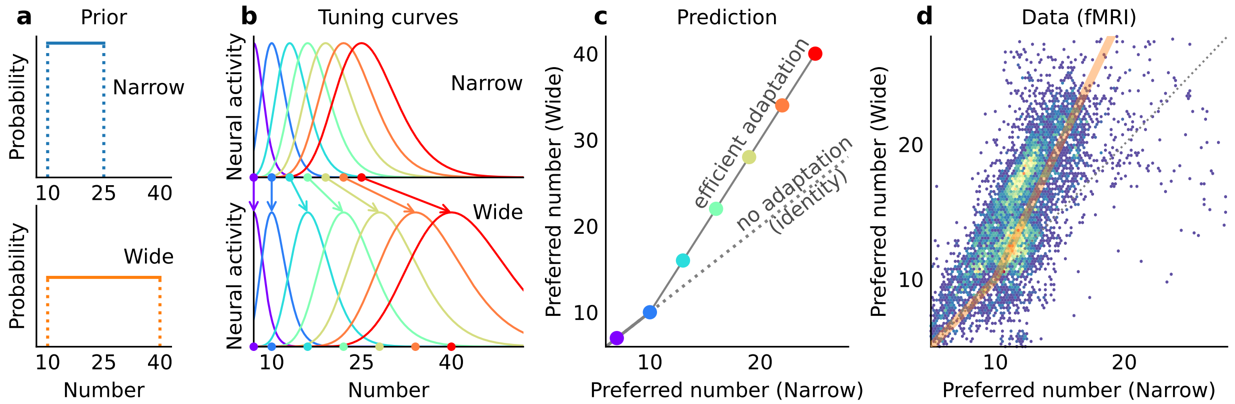

Fig. 2: Distributed range adaptation in neural tuning. The hypothesis of distributed range adaptation posits that when stimuli statistics change from a uniform prior to another (**a**), tuning curves collectively shift to efficiently cover the range of the current prior (**b**). This implies a linear relationship between the neurons' preferred stimuli under the two priors (**c**), which was verified quantitatively in fMRI data collected in a numerosity estimation task[49] (**d**). This illustrates neural populations dynamically implementing efficient coding (Eq. 12).

Box 4: Imprecision shapes behavior

The consequences of imprecision are not limited to noise that averages out. Because reasoning involves nonlinear transformations of information, imprecision generates deviations from optimality in complex ways. We illustrate three cases.

Primacy and recency effects These well-known effects can be explained by a single mechanism of Bayesian estimation. In sequential averaging tasks, participants estimate the average of a sequence of numbers. Assuming they maintain a running estimate of the average which they update every time they see a new number, they need to keep track *of the number of observations seen thus far* to determine how much weight to assign to the new number, in the updated average. If their memory is imperfect, they are only able to imprecisely estimate this count; typically, a Bayesian observer will overestimate small counts and underestimate larger ones (Fig. Ia, left panel). This alone generates primacy and recency effects in the weights assigned to the successive numbers, on the basis of these imprecise counts[79,80] (Fig. Ia, right panel).

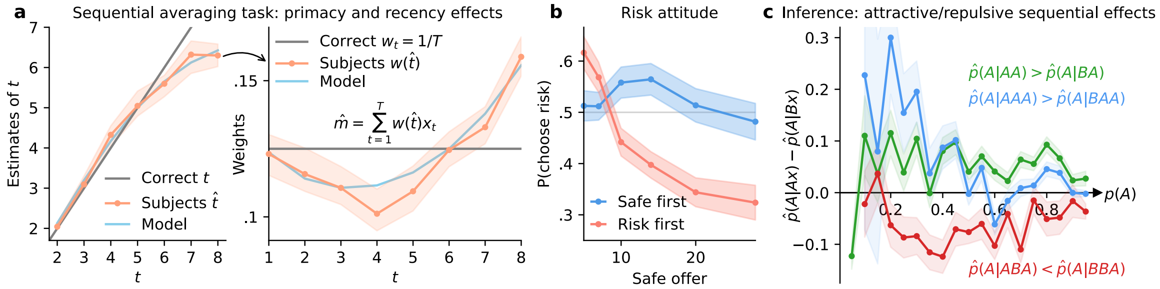

Fig. I: From imprecision to bias. a. Imprecision in estimating the number of observations seen (left) yields primacy and recency effects (right) in a sequential averaging task[79,80]. **b.** The order of presentation of safe and risk options impacts risky choices[7]. **c.** Inference under a cost on the precision of beliefs results in subtle patterns of attractive and repulsive effects of past observations on current predictions[81].

Risk attitude When offered a 'safe' (certain) option or a 'risk' one (a lottery), participants' choices depend on whether the risk option is presented first or second. If it is presented first, participants are more risk-seeking when values of options are low, and more risk-averse when they are high (Fig. Ib). This is readily explained if one assumes that the option presented first is less well remembered, and is thus estimated with a stronger 'central tendency', inflating lower values and depreciating higher ones[7]. More generally, many patterns of risk attitude are explained by models of cognitive imprecision, supported by neurophysiological evidence[1–3,5–7,82].

Inference Inferring the process generating observations enables making better predictions. With sequences of Bernoulli signals, Bayesian beliefs converge, but the predictions of human participants fluctuate with the recent history of observations. Counter-intuitively, predictions in some cases are biased away from past observations (Fig. Ic), a repulsive effect captured by a model of inference under a cost on the precision of representations[81]. More generally, constraints on the computational resources required to represent beliefs and conduct Bayesian inference seem to underlie many seemingly suboptimal behavioral patterns[83–90].

Concluding Remarks

All cognitive models make assumptions about representations. Many posit precise representations and build elaborate suboptimal decision processes on top of them. Given that representations are necessarily imperfect, a sensible alternative is to consider instead optimal reasoning on the basis of imperfect representations. Box 4 illustrates successful examples of this approach, where long-known behavioral patterns (e.g., primacy and recency effects) are readily explained by *optimal* decision-making over *imprecise* representations[91].

Under this view, understanding cognition begins with identifying representations. We have argued that representations themselves are optimal, considering the context and resource constraints—with the immediate implication that representations are context-dependent. Explicitly identifying the constraints is crucial, and turns qualitative ideas into quantitative, testable predictions. Here we have done so by drawing on the neuroscientific literature on neural representations of magnitudes. The resulting theory is thus grounded in observable neural activity (Fig. 2), and has implications for both cognition and neural coding.

More broadly, combining first principles with empirical input from the neuroscience of perception provides a path toward a neurocognitive theory of representations, itself a foundation for a theory of human decision-making. We venture that this approach will remain fruitful as it is extended beyond simple magnitudes to richer forms of information (see Outstanding Questions).

Outstanding Questions

How does the framework extend to multiple or multi-dimensional variables? In higher dimensions, representational resources must be shared across variables: how is precision jointly allocated, and how are trade-offs resolved?

On what timescales do representations adapt to the prior and to the task? In some settings a linear encoding is optimal (e.g., numerosity estimation with uniform priors), yet a logarithmic encoding, presumably adapted to a long-term prior, provides a better fit of behavioral data. This suggests that some components of the encoding may be slow to adapt. One possibility is that the structure of the encoding across stimuli (here, $I_1(x)$) is relatively stable, while its overall precision (n) adjusts more rapidly.

How are beliefs represented under resource constraints? What objective governs their representation? What are the behavioral consequences of imprecise beliefs, and how does such imprecision in beliefs affect inference and learning?

What are the resource constraints, and what is the representational geometry for cognitive variables that are not simple magnitudes but instead structured objects in richer spaces (e.g., semantic categories, relational knowledge)?

How do encoding constraints relate to, or interact with, memory limitations? How do they shape the fidelity and structure of stored representations?

We have interpreted n as the number of samples, but what does it correspond to mechanistically, and can it be measured? It may reflect, for instance, the duration of sampling or the number of neurons recruited.

Do early sensory representations adapt to task demands, or do they implement a fixed objective such as mutual information maximization for a long-term prior? Where along the processing hierarchy do representations become task-dependent?

Glossary

Bias Systematic deviation of estimates from the true value ($\mathbb{E}[\hat{x} - x|x]$). It arises from prior expectations, but also from variations in the encoding precision (Eq. 19 and Box 1).

Choice task A task in which two magnitudes are presented and the observer selects one, with a reward proportional to the chosen magnitude (Box 2).

Classification task A task in which the observer identifies the greater of two magnitudes and receives a fixed reward for a correct response, independently of the magnitude difference (Box 2).

Distributed range adaptation A form of adaptive neural coding in populations of sensory neurons, in which the distribution of neurons' preferred stimuli scales to match the current prior, so that the population efficiently covers the relevant range of stimuli (Fig. 2).

Efficient coding The principle that representations are optimized under constraints, typically allocating more precision to more relevant or more frequent stimuli (Eq. 12, Box 3).

Estimation task A task in which a single magnitude is presented and the observer reports it as accurately as possible (Box 2).

Fisher information A measure of the precision of an encoding (Box 1). In the small-noise regime, it is inversely related to the variance of estimates, and it is an important component of the bias (Eq. 19).

Mutual information A measure of statistical dependence between stimulus and representation, often used as a general-purpose objective in efficient-coding models (Box 2).

Prior The probability distribution over stimuli ($f(x)$), reflecting their frequency or expected occurrence in a given context.

Tuning curve The average response of a neuron as a function of the stimulus it encodes (Fig. 2).

Weber's law The empirical observation that discrimination thresholds scale approximately proportionally to stimulus magnitude, implying constant relative precision.

Acknowledgments

This research was supported by the Kempner Institute for the Study of Natural and Artificial Intelligence and a Polymath Award from Schmidt Sciences (S.J.G.).

References

[1] Mel Win Khaw, Ziang Li, and Michael Woodford. Cognitive Imprecision and Small-Stakes Risk Aversion. *The Review of Economic Studies*, 88(4):1979–2013, August 2020.

[2] Mel Win Khaw, Ziang Li, and Michael Woodford. Cognitive Imprecision and Stake-Dependent Risk Attitudes. September 2022.

[3] Cary Frydman and Lawrence J Jin. Efficient Coding and Risky Choice. *The Quarterly Journal of Economics*, 137(1):161–213, December 2021.

[4] Ryan Oprea. Decisions under Risk Are Decisions under Complexity. *American Economic Review*, 114(12):3789–3811, December 2024.

[5] Miguel Barretto-García, Gilles de Hollander, Marcus Grueschow, Rafael Polanía, Michael Woodford, and Christian C. Ruff. Individual risk attitudes arise from noise in neurocognitive magnitude representations. *Nature Human Behaviour*, 7(September), 2023.

[6] Gilles de Hollander, Marius Moisa, and Christian C. Ruff. Risk preferences causally rely on parietal magnitude representations: Evidence from combined TMS-fMRI, January 2025.

[7] Gilles de Hollander, Marcus Grueschow, Franciszek Hennel, and Christian C. Ruff. Rapid Changes in Risk Attitudes Originate from Bayesian Inference on Parietal Magnitude Representations, October 2025.

[8] Christopher A. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690, April 2003.

[9] Michael Woodford. Inattentive Valuation and Reference-Dependent Choice. *Unpublished Manuscript, Columbia University*, page 89, 2012.

[10] Andrew Caplin and Mark Dean. Revealed Preference, Rational Inattention, and Costly Information Acquisition. *American Economic Review*, 105(7):2183–2203, July 2015.

[11] Filip Matêjka and Alisdair McKay. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–298, 2015.

[12] Mark Dean and Nathaniel Neligh. Experimental Tests of Rational Inattention. *Journal of Political Economy*, 131(12):3415–3461, December 2023.

[13] Xavier Gabaix. A sparsity-based model of bounded rationality. *Quarterly Journal of Economics*, 129(4):1661–1710, 2014.

[14] Xavier Gabaix. Behavioral inattention. In *Handbook of Behavioral Economics: Applications and Foundations 1*, volume 2, pages 261–343. Elsevier, 2019.

[15] Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Salience Theory of Choice Under Risk. *The Quarterly Journal of Economics*, 127(3):1243–1285, August 2012.

[16] Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Salience and Consumer Choice. *Journal of Political Economy*, 121(5):803–843, 2013.

[17] Benjamin Enke and Thomas Graeber. Cognitive Uncertainty. *The Quarterly Journal of Economics*, 138(4):2021–2067, September 2023.

[18] Benjamin Enke. The Cognitive Turn in Behavioral Economics. 2024.

[19] H. B. Barlow. Possible Principles Underlying the Transformations of Sensory Messages. In Walter A. Rosenblith, editor, *Sensory Communication*, chapter 13, pages 217–234. The MIT Press, Cambridge, MA, sep 1961.

[20] Nicolas Brunel and Jean-Pierre Nadal. Optimal tuning curves for neurons spiking as a Poisson process. *Proceedings of the ESANN Conference*, 1997.

[21] Mark D. McDonnell and Nigel G. Stocks. Maximally Informative Stimuli and Tuning Curves for Sigmoidal Rate-Coding Neurons and Populations. *Physical Review Letters*, 101(5):058103, aug 2008.

[22] Deep Ganguli and Eero P Simoncelli. Implicit encoding of prior probabilities in optimal neural populations. In J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23*, pages 658—-666. Curran Associates, Inc., 2010.

[23] Deep Ganguli and Eero P. Simoncelli. Efficient Sensory Encoding and Bayesian Inference with Heterogeneous Neural Populations. *Neural Computation*, 26(10):2103–2134, October 2014.

[24] Xue-Xin Wei and Alan A. Stocker. A Bayesian observer model constrained by efficient coding can explain 'anti-Bayesian' percepts. *Nature Neuroscience*, 18(10):1509–1517, October 2015.

[25] Xue-Xin Wei and Alan A Stocker. Mutual Information, Fisher Information, and Efficient Coding. *Neural Computation*, 326:305–326, 2016.

[26] Deep Ganguli and Eero P. Simoncelli. Neural and perceptual signatures of efficient sensory coding. *ArXiv e-prints*, pages 1–24, February 2016.

[27] Zhuo Wang, Alan A. Stocker, and Daniel D. Lee. Efficient Neural Codes That Minimize Lp Reconstruction Error. *Neural Computation*, 28(12):2656–2686, dec 2016.

[28] Michael Morais and Jonathan W Pillow. Power-law efficient neural codes provide general link between perceptual bias and discriminability. *Advances in Neural Information Processing Systems 31*, 2(1):5076–5085, 2018.

[29] Arthur Prat-Carrabin and Michael Woodford. Bias and variance of the Bayesian-mean decoder. In *Advances in Neural Information Processing Systems*, volume 34, pages 23793–23805. Curran Associates, Inc., 2021.

[30] Ling Qi Zhang and Alan A. Stocker. Prior Expectations in Visual Speed Perception Predict Encoding Characteristics of Neurons in Area MT. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 42(14):2951–2962, 2022.

[31] 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, 2015.

[32] 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:e1, February 2020.

[33] Rahul Bhui, Lucy Lai, and Samuel J. Gershman. Resource-rational decision making. *Current Opinion in Behavioral Sciences*, 41:15–21, 2021.

[34] Thomas F. Icard. *Resource Rationality*. MIT Press, 2026.

[35] Wei Ji Ma and Michael Woodford. Multiple conceptions of resource rationality. *Behavioral and Brain Sciences*, 43:e15, March 2020.

[36] Bayes Procedures. In A. W. van der Vaart, editor, *Asymptotic Statistics*, Cambridge Series in Statistical and Probabilistic Mathematics, pages 138–152. Cambridge University Press, Cambridge, 1998.

[37] Yang Xiang, Thomas Graeber, Benjamin Enke, and Samuel J. Gershman. Confidence and central tendency in perceptual judgment. *Attention, Perception, & Psychophysics*, 83(7):3024–3034, October 2021.

[38] Bertrand S. Clarke and Andrew R. Barron. Jeffreys' prior is asymptotically least favorable under entropy risk. *Journal of Statistical Planning and Inference*, 41(1):37–60, August 1994.

[39] Arthur Prat-Carrabin and Michael Woodford. Efficient coding of numbers explains decision bias and noise. *Nature Human Behaviour*, pages 845–848, may 2022.

[40] S B Laughlin, R R de Ruyter Van Steveninck, and J C Anderson. The metabolic cost of neural information. *Nature neuroscience*, 1(1):36–41, 1998.

[41] Andrea Hasenstaub, Stephani Otte, Edward Callaway, and Terrence J. Sejnowski. Metabolic cost as a unifying principle governing neuronal biophysics. *Proceedings of the National Academy of Sciences of the United States of America*, 107(27):12329–12334, 2010.

[42] Biswa Sengupta, Martin Stemmler, Simon B. Laughlin, and Jeremy E. Niven. Action potential energy efficiency varies among neuron types in vertebrates and invertebrates. *PLoS Computational Biology*, 6(7):35, 2010.

[43] Arthur Prat-Carrabin and Michael Woodford. Endogenous Precision of the Number Sense. *eLife*, 13, March 2026.

[44] Arthur Prat-Carrabin, Maximilian V. Harl, and Samuel J. Gershman. Fast efficient coding and sensory adaptation in gain-adaptive recurrent networks, February 2026.

[45] Steeve Laquitaine and Justin L. Gardner. A Switching Observer for Human Perceptual Estimation. *Neuron*, 97(2):462–474.e6, January 2018.

[46] Long Ni and Alan A. Stocker. Efficient sensory encoding predicts robust averaging. *Cognition*, 232:105334, March 2023.

[47] Jiang Mao, Constantin A. Rothkopf, and Alan A. Stocker. Adaptation optimizes sensory encoding for future stimuli. *PLOS Computational Biology*, 21(1):e1012746, January 2025.

[48] Joseph A. Heng, Michael Woodford, and Rafael Polania. Efficient sampling and noisy decisions. *eLife*, 9:e54962, September 2020.

[49] Arthur Prat-Carrabin, Gilles de Hollander, Saurabh Bedi, Samuel J. Gershman, and Christian C. Ruff. Distributed range adaptation in human parietal encoding of numbers, March 2026.

[50] Jonathan Schaffner, Sherry Dongqi Bao, Philippe N. Tobler, Todd A. Hare, and Rafael Polania. Sensory perception relies on fitness-maximizing codes. *Nature Human Behaviour*, 7(7):1135–1151, 2023.

[51] John T. Serences. Value-Based Modulations in Human Visual Cortex. *Neuron*, 60(6):1169–1181, December 2008.

[52] Arezoo Pooresmaeili, Thomas H. B. FitzGerald, Dominik R. Bach, Ulf Toelch, Florian Ostendorf, and Raymond J. Dolan. Cross-modal effects of value on perceptual acuity and stimulus encoding. *Proceedings of the National Academy of Sciences*, 111(42):15244–15249, October 2014.

[53] Burkhard Pleger, Felix Blankenburg, Christian C. Ruff, Jon Driver, and Raymond J. Dolan. Reward Facilitates Tactile Judgments and Modulates Hemodynamic Responses in Human Primary Somatosensory Cortex. *Journal of Neuroscience*, 28(33):8161–8168, August 2008.

[54] Tolga Çukur, Shinji Nishimoto, Alexander G. Huth, and Jack L. Gallant. Attention during natural vision warps semantic representation across the human brain. *Nature Neuroscience*, 16(6):763–770, June 2013.

[55] Alexandra O. Constantinescu, Jill X. O'Reilly, and Timothy E. J. Behrens. Organizing conceptual knowledge in humans with a gridlike code. *Science*, 352(6292):1464–1468, June 2016.

[56] Mark K. Ho, David Abel, Carlos G. Correa, Michael L. Littman, Jonathan D. Cohen, and Thomas L. Griffiths. People construct simplified mental representations to plan. *Nature*, 606(7912):129–136, 2022.

[57] Camillo Padoa-Schioppa. Range-adapting representation of economic value in the orbitofrontal cortex. *Journal of Neuroscience*, 29(44):14004–14014, 2009.

[58] Antonio Rangel and John A. Clithero. Value normalization in decision making: Theory and evidence. *Current Opinion in Neurobiology*, 22(6):970–981, 2012.

[59] 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(7):e1002607, July 2012.

[60] Sophie Bavard and Stefano Palminteri. The functional form of value normalization in human reinforcement learning. *eLife*, 12:e83891, July 2023.

[61] Samuel J. Gershman and Rahul Bhui. Rationally inattentive intertemporal choice. *Nature Communications*, 11(1):3365, July 2020.

[62] B. C. J. Moore. Frequency difference limens for short-duration tones. *The Journal of the Acoustical Society of America*, 54(3):610–619, September 1973.

[63] C. Formby. Differential sensitivity to tonal frequency and to the rate of amplitude modulation of broadband noise by normally hearing listeners. *The Journal of the Acoustical Society of America*, 78(1):70–77, July 1985.

[64] Terry Caelli, Hans Brettel, Ingo Rentschler, and Rudi Hilz. Discrimination thresholds in the two-dimensional spatial frequency domain. *Vision Research*, 23(2):129–133, January 1983.

[65] Craig C. Wier, Walt Jesteadt, and David M. Green. Frequency discrimination as a function of frequency and sensation level. *The Journal of the Acoustical Society of America*, 61(1):178–184, January 1977.

[66] J. Lemańska, A. P. Sęk, and E. B. Skrodzka. Discrimination of the amplitude modulation rate. *Archives of Acoustics*, 27(1), 2002.

[67] D. Regan, S. Bartol, T. J. Murray, and K. I. Beverley. Spatial Frequency Discrimination in Normal Vision and in Patients with Multiple Sclerosis. *Brain*, 105(4):735–754, 1982.

[68] 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, 2017.

[69] V. de Gardelle, S. Kouider, and J. Sackur. An oblique illusion modulated by visibility: Non-monotonic sensory integration in orientation processing. *Journal of Vision*, 10(10):6–6, August 2010.

[70] Ahna R Girshick, Michael S Landy, and Eero P Simoncelli. Cardinal rules: Visual orientation perception reflects knowledge of environmental statistics. *Nature Neuroscience*, 14(7):926–932, July 2011.

[71] Frederike H. Petzschner and Stefan Glasauer. Iterative Bayesian estimation as an explanation for range and regression effects: A study on human path integration. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 31(47):17220–9, 2011.

[72] Jingyang Zhou, Lyndon R. Duong, and Eero P. Simoncelli. A unified framework for perceived magnitude and discriminability of sensory stimuli. *Proceedings of the National Academy of Sciences*, 121(25):e2312293121, June 2024.

[73] Steven T. Piantadosi. A rational analysis of the approximate number system. *Psychonomic Bulletin & Review*, 23(3):877–886, June 2016.

[74] Michael Hahn and Xue-Xin Wei. A unifying theory explains seemingly contradictory biases in perceptual estimation. *Nature Neuroscience*, 27(4):793–804, April 2024.

[75] Jiang Mao and Alan A. Stocker. Sensory perception is a holistic inference process. *Psychological Review*, 131(4):858–890, July 2024.

[76] Frederike H. Petzschner, Stefan Glasauer, and Klaas E. Stephan. A Bayesian perspective on magnitude estimation. *Trends in Cognitive Sciences*, 19(5):285–293, May 2015.

[77] Gustavo B. M. Mello, Sofia Soares, and Joseph J. Paton. A Scalable Population Code for Time in the Striatum. *Current Biology*, 25(9):1113–1122, May 2015.

[78] Akihiro Shimbo, Ei Ichi Izawa, and Shigeyoshi Fujisawa. Scalable representation of time in the hippocampus. *Science Advances*, 7(6), 2021.

[79] Verena Clarmann Von Clarenau, Stefan Appelhoff, Thorsten Pachur, and Bernhard Spitzer. Over- and underweighting of extreme values in decisions from sequential samples. *Journal of Experimental Psychology: General*, 153(3):814–826, March 2024.

[80] Arthur Prat-Carrabin and Michael Woodford. Imprecise counting of observations in averaging tasks predicts primacy and recency effects, October 2024.

[81] Arthur Prat-Carrabin, Florent Meyniel, and Rava Azeredo da Silveira. Resource-rational account of sequential effects in human prediction. *eLife*, 13:1–42, January 2024.

[82] Saurabh Bedi, Gilles de Hollander, and Christian C. Ruff. Probability weighting arises from boundary repulsions of cognitive noise, September 2025.

[83] Thomas F Icard and Noah D. Goodman. A Resource-Rational Approach to the Causal Frame Problem. *Proceedings of the 37th Annual Meeting of the Cognitive Science Society*, pages 962–967, 2015.

[84] Arthur Prat-Carrabin, Florent Meyniel, Misha Tsodyks, and Rava Azeredo da Silveira. Biases and variability from costly Bayesian inference. *Entropy*, 23(5):603, May 2021.

[85] Arthur Prat-Carrabin, Robert C. Wilson, Jonathan D. Cohen, and Rava Azeredo da Silveira. Human inference in changing environments with temporal structure. *Psychological Review*, 128(5):879–912, October 2021.

[86] Rava Azeredo da Silveira and Michael Woodford. Noisy Memory and Over-Reaction to News. *AEA Papers and Proceedings*, 109:557–561, May 2019.

[87] Rava Azeredo Da Silveira, Yeji Sung, and Michael Woodford. Optimally Imprecise Memory and Biased Forecasts. *American Economic Review*, 114(10):3075–3118, October 2024.

[88] Arthur Prat-Carrabin and Michael Woodford. Imprecise probabilistic inference from sequential data. *Psychological Review*, April 2024.

[89] Jian-Qiao Zhu, Adam N. Sanborn, and Nick Chater. The Bayesian sampler: Generic Bayesian inference causes incoherence in human probability judgments. *Psychological Review*, 127(5):719–748, October 2020.

[90] Jian-Qiao Zhu and Thomas L. Griffiths. Computation-limited Bayesian updating: A resource-rational analysis of approximate Bayesian inference. *Psychological Review*, June 2025.

[91] Michael Woodford. Modeling imprecision in perception, valuation, and choice. *Annual Review of Economics*, 12:579–601, 2020.