Behavioral Signatures of Cognitive Noise*

Ryan Oprea Michael Woodford

Abstract

The cognitive noise hypothesis is an increasingly popular explanation for a wide range of puzzling regularities in human behavior. Using simple, familiar tools and minimal assumptions drawn from neuroscience, cognitive noise models provide a parsimonious, unified explanation for a number of the most important anomalies documented in behavioral economics. The hypothesis proposes that limited cognitive resources force the brain to represent information imprecisely, generating coarser decisions, stochastic behavior and subjective uncertainty. Because brains deal with this kind of imprecision in a broadly Bayesian manner, imprecision leads to a number of well-known behavioral biases that may in fact be optimal given that imprecision. Because brains also allocate scarce cognitive resources adaptively, these biases and their severity are impacted by seemingly-irrelevant aspects of the choice environment, providing an explanation for a number of well-known context and experience effects that are otherwise difficult to explain. We review the basic structure of these models, their behavioral consequences, and the growing body of empirical evidence supporting them.

*Oprea: Haas School of Business, UC Berkeley, roprea@berkeley.edu. Woodford: Department of Economics, Columbia University, mw2230@columbia.edu. Prepared for the *Annual Review of Economics*. Please cite as: *Annual Review of Economics* 19, forthcoming, https://doi.org/10.1146/annurev-economics-081226-081014. We thank Ben Enke, Larry Samuelson, and Ferdinand Vieider for helpful comments on earlier drafts, while absolving them of responsibility for any remaining misinterpretations.

1 Introduction

In this essay we provide an introduction to the *cognitive noise hypothesis* – a promising theory of bounded rationality, rooted in neuroscience, that (i) parsimoniously explains a number of otherwise puzzling behaviors using familiar tools from statistical decision theory, and (ii) has received significant empirical support in recent years. The simplicity, portability and technical familiarity of cognitive noise models makes them a promising approach for improving both the psychological realism and the predictive and explanatory power of economic theory.

Cognitive noise models are motivated by three basic facts about human behavior that are well-documented both in the field and in the lab.

1. People's decisions are often **imprecise**, leading people to fail to make distinctions that matter, to make internally inconsistent choices, and to be uncertain about the quality of their own judgments. For example, the same experimental subject will not make the same choice with regard to a simple gamble, when presented repeatedly with the same choice (Mosteller and Nogee, 1951); even expert wine judges score the very same wine differently, in a blind test, when it is poured for them three times in a single sitting (Hodgson, 2008); used-car buyers treat an odometer reading of 79,900 miles as far better than an identical car at 80,000, reading a magnitude based only on the leading digit (Lacetera, Pope and Sydnor, 2012); and in a wide range of economic decisions people tend to express uncertainty about the correctness of their choices (Enke et al., 2026).

2. People's decisions are often systematically **biased**: relative to normative benchmarks, decision-makers over-respond to some things that matter for a choice, and under-respond to others. As examples of over-responding, bettors over-back longshots at the track (Snowberg and Wolfers, 2010) and homeowners overpay to eliminate small risks by choosing low deductibles (Sydnor, 2010), in each case behaving as if they overweight a small probability.. On the under-responding side, the very same bettors under-back favorites (Snowberg and Wolfers, 2010), reacting too weakly to a large probability; and stated valuations barely move with the *scale* of what is at stake — willingness to pay to save 2,000, 20,000, or 200,000 birds is essentially the same (Desvousges et al., 1993); more generally, stated values are notoriously insensitive to scope (Kahneman and Knetsch, 1992).

3. People's decisions are often **unstable**, influenced by seemingly-superficial details of the choice environment and their own seemingly-irrelevant prior experiences. For example, people's judgments drift with their own history — investors who came of age in a bear market shun stocks for decades (Malmendier and Nagel, 2011), and households rush to buy flood insurance after a nearby flood, only to let it lapse over the years that follow (Gallagher, 2014) — and with the mere framing of a choice — fingerprint examiners reverse their own identifications given an irrelevant cue (Dror, Charlton and Péron, 2006), and far more people register as organ donors under opt-out than opt-in defaults (Johnson and Goldstein, 2003).

The cognitive noise hypothesis proposes that these three puzzling features of human behavior are fundamentally linked, and jointly explained by simple models built using conceptual tools familiar to any economist. Using only minimal deviations from standard economic models — deviations inspired by long literatures in psychology and neuroscience — these models can parsimoniously explain a number of key anomalies from behavioral economics and generate novel, testable predictions — predictions supported in a number of recent empirical studies.

Cognitive noise models rest on three basic theses that map roughly to the three motivating facts discussed above.

1. As we discuss in Section 2, people's judgments are imprecise because precision is costly for the brain, requiring significant metabolic and neuronal resources to store information accurately, to process it optimally and to prevent interference from other cognitive processes. As a result, people's choices are best modeled as responses, not to precise information, but rather to **imprecise representations** that are subject to random fluctuations – a fact that immediately leads to coarse, noisy and often meta-cognitively uncertain judgments.

2. As we discuss in Section 3, if decisions respond optimally (i.e., in a roughly **Bayesian** way) to these noisy representations, the result will usually not be random fluctuation around the correct response, but a response that is **systematically biased.** Moreover, it will often be efficient for decisions to take account of additional available information, that would be irrelevant for a decision maker with access to a more precise representation of the decision-relevant information. Hence many of the most important biases documented in behavioral economics can be interpreted as Bayesian responses to noisy internal representations.

3. As we discuss in Section 4, there is significant evidence that the brain does not assign precision arbitrarily, but instead does so **adaptively**. As a result, not all information is represented with equal noise, and therefore not all judgments are biased in the same way. The degree of imprecision, and the way in which imprecise representations are interpreted, are sensitive to the brain's expectations, past experiences, and prior beliefs, and to the familiarity and complexity of the tasks it tackles. This creates immediate scope for many of the seemingly-irrational context effects documented in the literature.

As we emphasize throughout, all three of these theses have significant empirical support in neuroscience, psychology and behavioral economics, some of it going back more than a century. Figure 1 offers a schematic summary of the cognitive noise hypothesis.

The cognitive noise literature builds on and connects with other important approaches to modeling cognition. It has obvious connections to theoretical perspectives that similarly view the brain as an efficient organ and people's "mistakes" as solutions to constrained optimization problems. These include literatures on bounded rationality (Simon, 1955), ecological rationality (Gigerenzer and Goldstein, 1996), resource rationality (Lieder and Griffiths, 2020), evolutionary models of preferences (Robson, 2001; Page, 2022; Netzer, 2009), rational inattention (Sims, 2003; Mackowiak, Matejka and Wiederholt, 2023), and especially the "efficiently irrational" program in neuroeconomics, with which it shares many ideas and themes (Glimcher, 2011; 2022). It has important connections to non-Bayesian approaches that also emphasize cognitive noise, such as decision by sampling (Stewart, Chater and Brown, 2006) and drift-diffusion models (Ratcliff and McKoon, 2008; Krajbich, Armel and Rangel, 2010); and we believe that it provides an important complement to other approaches that emphasize imprecision, such as sparse representations (Gabaix, 2014; 2019), salience (Bordalo, Gennaioli and Shleifer, 2012; 2022), and memory effects (Bordalo, Gennaioli and Shleifer, 2020).

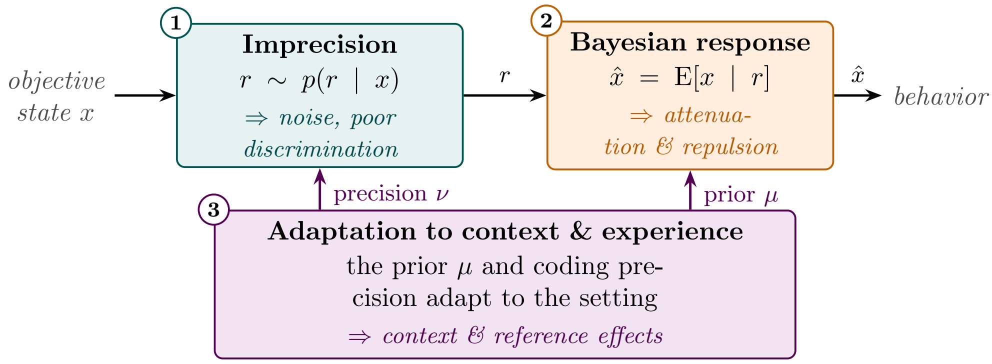

Figure 1: **Three facets of cognitive noise models.** Each facet couples a modeling assumption with a behavioral consequence. **(1)** Cognition is imprecise: the decision situation x is encoded by a noisy representation r, producing stochastic choice and imperfect discrimination between states. **(2)** Behavior responds optimally to that imprecision: decoding r to the posterior mean $\hat{x} = \mathrm{E}[x \mid r]$ creates systematic biases (both attenuation and repulsion). **(3)** The prior μ and coding precision are not static but adapt to context, experience, cues and expectations, so that the same x yields different behavior across settings.

In the following three sections (Sections 2, 3 and 4) we motivate, model and provide evidence for each of these three tendencies in turn, and the theses that the cognitive noise hypothesis advances to understand them. Our goal is to give the reader a clear sense of the ideas behind these models, and some illustrative evidence in favor of them. We cannot attempt anything like a comprehensive review of this large and growing literature, let alone the many currents of research in psychology, neuroscience and neuroeconomics that motivate it, but we hope that readers will be motivated to dig deeper into these topics. In Section 5, we close with brief reflections on where further work is needed.

2 Cognitive Noise

The first observation motivating noisy cognition models is that people often (especially when not using formal rules like the rules of mathematics) make decisions and judgments that are *imprecise*. That is, unlike ideal rational actors, people often do not make perfectly sharp distinctions between alternative situations. In other words, they often make decisions based on "coarse" or "compressed" representations of the situations they face.

For example, Olenski et al. (2020) find that heart-attack patients admitted in the two weeks after turning 80 were substantially less likely to undergo coronary-artery bypass surgery than those admitted just before turning 80 — physicians encode age in coarse decade bins and make high-stakes decisions on that basis. List et al. (2023) show, using a field experiment on millions of Lyft riders, that trip conversions fall discontinuously at dollar thresholds, and Strulov-Shlain (2023) finds similar patterns in retail scanner data, suggesting consumers treat \$3.99 as meaningfully below \$4.00. All point to people imprecisely coarsening highly relevant information when making choices.

Not only are decisions often imprecise, they are also often *noisy*, based on random factors that make judgments inconsistent across repetitions. For instance, Hodgson (2008) shows that expert judges at a major wine competition, served the same wine multiple times in a blind tasting, score the wine differently by a median of about four points on a 100-point scale, with a pooled standard deviation of 3.6 points and a 95% spread near 14 points — almost the entire range of medals awarded. For decades, marketing researchers have modeled brand choice as a random-utility process (e.g., Guadagni and Little, 1983), reflecting the substantial trial-to-trial randomness in observed choices — a type of evidence that has motivated economists' use of random utility models (McFadden, 1974) to account for unexplained variability in choices when estimating preferences. This kind of stochasticity in judgment has been documented for over a century in controlled laboratory experiments in psychology, where even the simplest perceptual judgments exhibit randomness (Fechner, 1860; Thurstone, 1927). Experimental economists have found similar failures of test–retest reliability ever since the earliest measurements of utility (e.g., Mosteller and Nogee 1951).

Finally, people are often *uncertain* about the quality of their own judgments — metacognitively aware that they are unsure what response to give, even when the task should not involve any facts not available to them. Enke et al. (2026) find that in laboratory experiments, most people express some doubt — "cognitive uncertainty" — about the correctness of the choices that they make in dozens of distinct economic tasks. Reported uncertainty furthermore correlates with inconsistency of the subject's choice across repetitions, as well as with reduced sensitivity to variation in objective conditions (Enke and Graeber 2023), indicating some awareness of the degree of precision of one's mental operations.

Underlying all three of these observations is a single idea: that people often act not on the basis of the decision situation itself, but of an imperfect internal *representation* of it, one that the brain must construct using limited resources. The neuroscience literature has shown that random noise is a pervasive feature of neural processing at all stages, from the early processing of sensory stimuli to the circuits that generate motor responses ((Faisal, Selen and Wolpert, 2008)); and in some cases, random variation in recorded neural activity is found to explain random variation in expressed judgments about stimuli (e.g., Newsome, Britten and Movshon, 1989; Britten et al., 1996). The effects of such noise can be mitigated by having a larger number of neurons represent a given quantity, so that randomness of the overall signal is reduced; but this approach is metabolically expensive, so that the nervous system cannot afford to represent everything with high precision. The general principle seems equally relevant to the representation of more abstract quantitities such as the strength of economic incentives.

2.1 Modeling Cognitive Noise

The noisy cognition literature models this imprecision by positing that a decision maker's internal representation of their situation is a random variable r drawn from a distribution $p(r \mid x)$ that depends on the objective condition x. In a simple but widely used formulation due to Thurstone (1927),

$$r \; = \; f(x) + \varepsilon, \qquad \varepsilon \sim N(0, \nu^2), \tag{1}$$

where the function f maps the objective situation into a location in an abstract "psychological space" (here assumed to be the real line), and ν indexes the imprecision of the representation. How differently two different situations x and x' will be treated by the decision maker is determined by how near the values $f(x)$ and $f(x')$ are to one another (relative to the size of ν); or more generally, by how much the distributions $p(r \mid x)$ and $p(r \mid x')$ overlap. In the specification (1), the single parameter ν indicates both the degree of imprecision (i.e., limited discrimination between nearby cases) and the degree of randomness of trial-by-trial responses. The modeling approach follows a century-and-a-half-old literature in psychophysics (Fechner, 1860) that has modeled sensory imprecision this way with great empirical success.

The claim here is *not* that people always misperceive the quantities they are shown — e.g., that they misread the numbers on the screen in experiments.[1] Rather, it is that even when basic data are apprehended (through symbols) in an initially precise way, the subsequent operations using the information to reach a conclusion are often not precise calculations (such as the conscious use of rules of arithmetic). Cognitive noise models apply to *complex* settings in which several pieces of information must be interpreted and integrated, so that mere recognition of the numbers on the screen is not enough to allow a response — one must think about what they mean, and weigh them against one another.

Indeed, even when data are recognized precisely — because they are presented as familiar number symbols — it is not easy for people to deploy decision rules that perform precise operations on those precise symbolic representations, except in special cases. They generally find it easier to form intuitive judgments through an unconscious process that takes imprecise "semantic" representations as inputs, one that draws on evolutionarily older brain structures that judge quantity without language or symbols. Even after schooling we continue to work more readily with semantic representations of quantity — the "approximate number system" — than with exact numbers, and intuitive processes inherit the imprecision of these representations. (See (Dehaene, 2003), (Nieder, 2016), and (Khaw, Li and Woodford, 2026).) It is the imprecision of the semantic representation of the quantity x that (1) is intended to specify.

In general, choice problems will require a *vector* of features to be encoded, not just a single scalar magnitude. In this case r can also involve noise in operations through which several pieces of information are integrated.

2.2 Measuring Cognitive Noise

The noisy cognition literature uses several methods to measure cognitive noise. A first approach is familiar to any empirical economist. Applied work routinely models choice as containing a random component and estimates its magnitude, typically finding that the estimated magnitude is not small and drives a large share of the variation in observed choice in a wide range of settings. Thus, in an important sense economics has been measuring noise in choice for half a century (McFadden, 1974; Berry, Levinsohn and Pakes, 1995). It is common to interpret this not as noise in the decision maker but as a limitation of the *analyst* — unobserved attributes of the options, unobserved characteristics of the chooser, or stable heterogeneity in tastes (Ben-Akiva and Lerman, 1985). However, similar stochastic choice models have been applied to experimental data in which such explanations are much less plausible, and finds similar evidence of noise. For instance, Hey and Orme (1994) found that subjects reverse a substantial fraction of identical, repeated lottery choices so that no deterministic theory of preference can fit

[1] The same is true, in our view, of models of "rational inattention" (Sims, 2003; Caplin and Dean, 2015; Mackowiak, Matejka and Wiederholt, 2023).

their data without a stochastic error term, concluding that expected utility *plus noise* accounts for behavior about as well as its more complicated generalizations. Careful econometric comparisons confirm that this inconsistency is genuine trial-to-trial noise, and suggest that how this noise is modeled can matter more for prediction than the preference model itself (Ballinger and Wilcox, 1997; Wilcox, 2008).

A second approach is to measure *cognitive* noise by measuring *behavioral* noise. The intuition is that, to the degree people's representations are noisy, the decisions they base on these representations will be noisy as well.[2] Behavioral noise is typically measured via test-retest reliability: when people fail to make the same choice on encountering the same task twice, the inconsistency is treated as a measure of noise. For example, a recent meta-analysis of more than seventy thousand test-retest correlations from roughly half a million respondents finds that behavioral, choice-based measures of risk preference have an average test-retest reliability of only about 0.25 (Bagaïni et al., 2025). Randomness of exactly this kind has, moreover, been a central object of measurement in the psychophysics literature dating back to (Fechner, 1860).

A third, somewhat different approach is to simply ask subjects how likely they think it is that they chose correctly. To the degree that people are metacognitively aware of the imprecision of their cognition, these self-reports can be interpreted as a measure of cognitive noise. There is now significant evidence that subjects have this kind of metacognitive awareness in many economics-style tasks. Enke, Graeber and Oprea (2023) study the correlation between cognitive uncertainty (CU) and objective performance in 15 standard cognitive tasks and show that CU is correlated with objective mistakes in the majority of them; the exceptions seem to be tasks in which mistakes rest on highly misleading intuitions. Enke et al. (2026) study an additional eight economic decision-making tasks with objectively correct answers and find that CU is predictive of mistakes in all of them. An advantage of this approach is that it allows a measure of the degree of cognitive noise to be collected on an individual trial, rather than only after the same decision problem is presented to the same subject multiple times.

2.3 Manipulating Cognitive Noise

The degree of noise in internal representations determines the severity of the biases we will discuss in sections 3 and 4). Because of this, many of the model's predictions concerning how context influences behavior are driven by the way context is predicted to influence the quality of internal representations (the parameter ν in (1)). Understanding what should influence ν determines both how we apply the theory and also our approach to testing it.

The quality of information processing depends both on the *supply* of cognitive resources and the cognitive resources that a given task *demands*.[3] On the *supply* side, less cognitively sophisticated decision makers can be expected to be noisier, while experience and expertise can be expected to lower noise, because learning allows the brain to rep-

[2] An important caution is that behavioral noise need not be monotonic in cognitive noise: very high levels of cognitive noise should imply low, rather than high, variability of behavior, if people learn to respond optimally to it. Vieider (2026) shows that this distinctive non-monotonicity can be seen in experimental data — strong evidence in favor of Bayesian cognitive noise models.

[3] These are "Simon's Scissors": "[h]uman rational behavior is shaped by a scissors whose two blades are the structure of task environments and the computational capabilities of the actor" (Simon, 1990). A task's demand for resources provides a natural definition of its *complexity* (Oprea, 2024).

resent a familiar problem with less contamination from extraneous detail. Noise can be expected to rise under time pressure, cognitive load, and distraction, all of which tie up the resources needed for precise decision making, or when the relevant information must be held in memory before being used, since stored information is retrieved with greater noise.[4] Conversely, in some tasks we can expect stronger incentives for accuracy to lower noise by inducing greater effort or costlier decision making approaches (Dean and Neligh, 2023; Prat-Carrabin and Woodford, 2026).

On the *demand* side, more complex tasks — those requiring more information to be held in memory and combined or that require more intensive computations (Aragones et al., 2005) — can be expected to result in noisier representations (Oprea, 2026; Enke et al., 2026; Enke, Graeber and Oprea, 2025; Enke and Shubatt, 2023). Seemingly superficial features of a problem *format* can matter: information shown in more familiar, intuitive, or easier-to-process forms can be expected to generate less noise (Oprea and Vieider, 2026b), as can tasks whose context or magnitudes are *familiar*, recruiting procedural memory (Garagnani and Vieider, 2025). All of these variables have been found to influence test-retest reliability, expressed cognitive uncertainty and structural estimates of noise (Enke and Graeber, 2023; Enke, Graeber and Oprea, 2025; Dean and Neligh, 2023). We should expect these comparative statics to be ceteris-paribus, mattering most at intermediate levels of the driving variable: incentives will change little in a task that is simply overwhelming, and cognitive load will matter little in one that is sufficiently trivial.

Other reasons for cognitive precision to vary are discussed more later. First, cognitive noise varies not only across tasks but also across parameter values of a given task. One reason is that some tasks become more difficult for some parameter values than others (e.g., lotteries are easy to value when they pay off for sure; Enke et al., 2026). A subtler reason is that there is significant evidence that the brain allocates precision adaptively based on prior experience (Heng, Woodford and Polania, 2020; Frydman and Jin, 2022), and responds involuntarily to subtle aspects of the task environment that influence attention, such as salience (Bordalo, Gennaioli and Shleifer, 2012; 2013; 2022). We discuss implications of these variations in precision in more detail below.

3 Optimal Bias

The second observation motivating cognitive noise models is that people's decisions are not only imprecise, but also systematically biased in a variety of ways. Over the last half century, behavioral economists have documented — both in the field and in the lab — a striking variety of biases afflicting belief formation, decision under risk, intertemporal choice, statistical reasoning, financial decisions and a host of other economically important decisions (DellaVigna, 2009; Dhami, 2016). These biases are often catalogued as a zoo of domain-specific anomalies, each with its own (sometimes unclear) scope of application.

The cognitive noise literature proposes that many of these biases may be deeply related, deriving ultimately from the cognitive imprecision discussed in Section 2 above. Instead

[4]For example, experimental evidence shows that estimation biases can be modified by manipulating time delays, both in comparisons of sequentially presented sensory magnitudes and in choices between sequentially presented gambles (Ashourian and Loewenstein, 2011; Hollander et al., 2025).

of deriving from domain-specific motivational factors, they reflect domain-general cognitive factors that arise repeatedly in different decision contexts. To the degree that this is true, there may be more conceptual unity to behavioral economics than is often supposed.

To understand why, it is important to note that the noisy representations mentioned above require *interpretation.* Even scalar physical magnitudes are represented in the brain ("encoded") by something very high-dimensional, a pattern of activity in a large population of neurons (Pouget, Dayan and Zemel, 2000; Ma et al., 2006). Hence the noisy representation r is not just an erroneous report of the true value x (a value "measured with error"). It need not even be a point in the same space as the possible values of x, and in a realistic model it generally is not.[5] So one can't assume that people "believe their distorted perception to be the truth" but otherwise behave like an optimal decision maker.

Instead, an *inference* must be drawn from the representation to determine how it is sensible to act; the representation must be "decoded." A key assumption of cognitive noise models is that the brain does this in an optimal way, relative to some distribution of possible true situations and to the decision maker's objectives. This requires that the information in the noisy representation be interpreted in the light of prior knowledge, beliefs and expectations in an approximately Bayesian manner.

Bayesian decoding is optimal, but nonetheless generally produces biased judgments and decisions relative to the ones that a rational actor would make in the absence of cognitive noise. These Bayesian biases turn out to match, and therefore provide an explanation for, a striking array of seemingly distinct biases documented in the behavioral literature. Here we emphasize two broad types of biases that flow directly from Bayesian statistics. In Section 3.2, we consider Bayesian compression or regression bias, and in Section 3.3 we discuss additional, subtler biases that occur when representational precision is not uniform across parameter values specifying a decision problem.

It is important to emphasize that when we use Bayesian calculations to derive behavioral predictions, we do *not* assume that humans know Bayes' rule, perform mental calculations, or even that they are natural Bayesians in deliberative statistical judgments. In fact, as we will see below, noisy cognition models have been used to explain *why* people depart from Bayes' rule in the directions that they do when asked to make explicit statistical inferences. Our thesis, instead, is that the human brain is adapted to make use of its noisy information in a relatively efficient way. Bayesian reasoning explains, to the analyst of human behavior, why particular strategies are optimal, even if people do not consciously follow this reasoning when making intuitive judgments. This is to say little more than (i) that the brain tends to supplement low-quality information with *other sources of information* that it also has access to when guiding choice, and (ii) that the brain does this more intensively the lower the quality (i.e., the noisier) the first information is. Although many papers in this tradition use explicit Bayesian calculations to obtain precise predictions, exact Bayesianism isn't necessary to derive most of the

[5]In (1) we assume that r is one-dimensional, but this is purely for pedagogical purposes. The crucial feature of the model is what it implies about the degree of overlap between the distributions $p(r|x)$ associated with different objective situations x, what is sometimes called the "representational geometry" implied by the noisy coding Kriegeskorte and Kievit (2013). This can be defined independently from any assumption about the space in which r lies, though a minimum dimension for the embedding space may be needed (e.g., Wei and Woodford (2025)).

important implications of these models — often, principles (i) and (ii) will suffice.

The idea that the brain deals with imprecision in a roughly Bayesian fashion follows a robust literature in perceptual neuroscience on what is known as the *Bayesian brain* hypothesis (Knill and Pouget, 2004; Ma et al., 2006). As discussed above, over a century of perceptual research has established that the brain encodes stimuli noisily. When people combine these noisy signals (across the senses or across time), they weight each in proportion to its reliability, exactly as optimal inference would prescribe (Körding and Wolpert, 2004). Likewise, when they estimate the speed or magnitude of a stimulus, their percepts are pulled toward the values they most often encounter in their environment, as if taking account of a Bayesian prior(Stocker and Simoncelli, 2006; Petzschner, Glasauer and Stephan, 2015). This behavior has a plausible basis in the firing of populations of neurons, whose activity can encode not merely a point estimate but an entire probability distribution (Ma et al., 2006).

3.1 Bayesian decoding

As discussed above, while a decision process based on a noisy representation is necessarily suboptimal, cognitive noise models often assume that this is the *only* respect in which behavior fails to be optimal; that is, the rule that determines a subject's action or response as a function of the representation r is the one that maximizes the subject's average reward, given the joint distribution of representations r and objective (payoff-relevant) states x. (Here "reward" need not be purely financial; it is whatever the decision process is adapted to achieve more of.) This requires an inference to be drawn (at least implicitly) about the state x that may have given rise to the representation r — the problem of "decoding" the representation.

In general, a given representation r can be produced by any of a variety of possible objective situations, so that it is not possible for a decision process to guarantee an optimal choice in every situation x, considered in isolation. One can at best ask for the choice to be the best one *on average,* averaging over possible objective situations weighted by the frequencies with which one expects to encounter them. This means that optimal decoding (and hence optimality of the decision rule) can be defined only relative to a particular *prior distribution* $p(x)$ over possible situations x. The optimal decision rule for a particular prior will furthermore be the one that, for any representation r, chooses the action $a(r)$ that maximizes expected reward under the *posterior distribution* $p(x\,|r)$ implied by the prior and the noisy coding scheme, computed using Bayes' rule.

As an example, suppose that the decision maker's reward is a quadratic function of a and x, so that under full information the optimal action $a^*(x)$ in any situation will be a linear function of x. When there is uncertainty about the value of x, a "certainty equivalence" result holds: the optimal action is still $a^*(\hat{x})$, where the optimal estimate of the state is simply the posterior mean,

$$\hat{x}(r) \;\equiv\; \mathrm{E}[x\,|r]. \tag{2}$$

If we also assume linear-Gaussian encoding — equation (1) with $f(x) = x$ — paired with a Gaussian prior $x \sim N(\mu, \sigma^2)$, this takes an especially transparent form: the posterior mean is a precision-weighted average of the representation and the prior mean,

$$\hat{x}(r) \;=\; \gamma\, r + (1-\gamma)\,\mu, \qquad \gamma \;\equiv\; \frac{\sigma^2}{\sigma^2+\nu^2}, \tag{3}$$

so that the estimate is shrunk toward the prior mean μ — and shrunk more strongly, the noisier the representation (i.e., the larger is ν and hence the smaller is γ). The optimal decision rule will then choose the action $a^*(\hat{x}(r))$.

Conditional on the true situation x, both the decoded value $\hat{x}$ and action choice a will be random variables, as both depend on r. Note also that in general, $\hat{x}$ will not be an unbiased estimate ($\mathrm{E}[\hat{x}\,|x] \neq x$), and neither will the optimal action choice $a^*(\hat{x}(r))$ in the presence of noise equal the ideal action choice $a^*(x)$ on average. We now discuss the biases implied by Bayesian optimality more generally.

3.2 Regression bias

A characteristic implication of Bayesian optimality is that people's responses to variation in imprecisely represented parameters will be *smaller* (at least overall) than the optimal response to a precisely represented parameter would be. In the linear-Gaussian example (3), a unit increase in x increases the average value of the estimate $\hat{x}$ by only $\gamma < 1$ units, with greater compression the noisier the encoding. Hence all small enough values of x are over-estimated on average, while all large enough values are under-estimated; estimates are on average regressed toward the prior mean of x.

These strong conclusions depend on special features of the linear-Gaussian example. But one can show quite generally that if $\hat{x}$ is a Bayesian posterior mean estimate (2), then a linear regression of observed values of $\hat{x}$ on the associated true values of x should have an asymptotic slope coefficient

$$\hat{\gamma} \;\equiv\; \frac{\mathrm{cov}(\hat{x},\, x)}{\mathrm{var}(x)} \;=\; 1 - \frac{\mathrm{E}[\mathrm{var}(x|r)]}{\mathrm{var}(x)} \tag{4}$$

satisfying $0 < \hat{\gamma} < 1$.[6] Thus a change in x moves behavior in the *same direction* as the change in the optimal response, on average, though by a smaller amount. This under-response to variables that ought to matter for the decision is what Enke et al. (2026) call "behavioral attenuation."

Equation (4) also implies that for a given degree of prior uncertainty about x, the regression coefficient $\hat{\gamma}$ must be smaller the greater the average posterior uncertainty about x conditional on the internal representation. Thus greater noise in the representation implies stronger regression bias (at least overall).[7] This allows a variety of empirical tests of the hypothesis: greater bias should co-move with trial-to-trial variability, and (to the extent people are metacognitively aware of their own imprecision) with expressed cognitive uncertainty. Bias should similarly be responsive in predictable ways to manipulations to the determinants of cognitive noise discussed in Section 2.3.

Systematic regression bias of this kind has been observed in a wide range of perceptual domains in the psychophysics literature. As Petzschner, Glasauer and Stephan (2015) review, experiments in which subjects must estimate extensive magnitudes such as lengths, distances, areas, size of angles, weight, force, or time duration all reveal a common pattern of over-estimation of smaller magnitudes and under-estimation of larger

[6]See the Supplementary Information for a derivation. This formula reduces to the value of γ given in (3) in the linear-Gaussian example.

[7]See Vaeth (2026) for further analysis, in a more special class of encoding models.

ones.[8] The same pattern is observed in the case of estimates of the numerosity of items in a visual array, as first noted by Jevons (1871).[9] Finally, in these kinds of perceptual tasks, the strength of estimation biases tends to vary with the degree of imprecision, as the Bayesian model predicts: regression bias is stronger when a moving image is shown at lower visual contrast (Stocker and Simoncelli, 2006), and Xiang et al. (2021) find that both trial-to-trial variability and regression bias rise on the trials where subjects express greater uncertainty or are allowed less time to view the stimulus.

Biases of the same kind are potentially equally important for understanding economic behavior. In particular, regression bias can explain a number of seemingly distinct behavioral anomalies emphasized in the behavioral economics literature, and recent research has provided significant evidence in favor of this interpretation. Here we give as examples some of the most salient biases in behavioral economics (see Enke et al. (2026) for dozens more).

1. Probability weighting and the fourfold pattern. An important example is the "fourfold pattern" of risk attitudes documented by Tversky and Kahneman (1992), which they call "the most distinctive implication of prospect theory." An experiment elicits the certain amount C that a subject regards as equivalent to a gamble paying X with probability p (and zero otherwise). Modal responses are risk averse for high-probability gains and low-probability losses, but risk seeking for low-probability gains and high-probability losses. Tversky and Kahneman interpret this fourfold pattern as indicating that gambles are valued in proportion not to the probability p but to a nonlinear transformation $w(p)$ — the "probability weighting function" — with $w(p) > p$ for small probabilities and $w(p) < p$ for larger ones.

This distortion can be rationalized as an example of regression bias, reflecting optimal Bayesian inference about the value of the gamble if the judgment is based on a noisy representation of the value of p (or of the expected value of the gamble).[10] A number of streams of evidence support this interpretation. Apparent probability weighting and the fourfold pattern are (i) highly predicted by cognitive uncertainty and strongly associated with behavioral noise (Enke and Graeber, 2023); (ii) disappear when subjects are forced to sample from lotteries before choosing between them (an effort to increase the precision of internal representations, Oprea and Vieider (2026a)); intensifies when the lotteries are presented as complex compound lotteries (Enke and Graeber, 2023); (iii) diminishes with the cognitive sophistication of the subject (Choi et al., 2022); (iv) disappears with framing manipulations designed to alter cognitive noise (Oprea and Vieider, 2026b); (v) waxes and wanes depending on prior experience with probabilities in a way anticipated by forces discussed in Section 4 (Frydman and Jin, 2025) (vi) and appears also in deterministic arithmetic tasks that resemble lotteries – but in which no risk preferences can actually be at work (Oprea, 2024). All of these findings are consistent with this central regularity in behavioral economics reflecting the frictions described by cognitive noise

[8]This well-documented pattern has been given a variety of names, including "regression bias" (Petzschner, Glasauer and Stephan, 2015), "contraction bias" (Ashourian and Loewenstein, 2011), "conservatism" (Hilbert, 2012), and "central-tendency effect" (Xiang et al., 2021).

[9]For more recent illustrations, see Anobile, Cicchini and Burr (2012) and Xiang et al. (2021).

[10]Bayesian explanations of this form for the probability weighting function of prospect theory have been offered by Fennell and Baddeley (2012), Steiner and Stewart (2016), Zhang, Ren and Maloney (2020), Khaw, Li and Woodford (2021; 2026), Enke and Graeber (2023), Vieider (2024), Frydman and Jin (2025), and Bedi, de Hollander and Ruff (2025).

models, and inconsistent with risk-preference-based interpretations.[11]

2. Hyperbolic time discounting. One of the most important findings from behavioral economics is that experimental subjects choosing between receiving a smaller amount of money sooner versus a larger amount later depart from standard time-consistent, exponential discounting. Instead, researchers often find that the factor by which later payments are discounted appears to be a quasi-hyperbolic function of the time delay, rather than decreasing exponentially with time. This can also be viewed as an example of regression bias, in which short delays are estimated on average to be longer than they actually are, while long delays are estimated to be shorter than they actually are. This leads to exaggerated impatience over short horizons and exaggerated patience over longer horizons (Vieider, 2025; Gabaix and Laibson, 2022; Gershman and Bhui, 2020; Ebert and Prelec, 2007). Consistent with this interpretation, Enke, Graeber and Oprea (2025) find that measured hyperbolic discounting is highly associated with measured cognitive uncertainty and behavioral noise, increases sharply with manipulations to the complexity of the description of time delays, and appears in tasks that involve no delay but require the same kind of complex recursive reasoning demanded by time delays. Oprea and Vieider (2026b) finds that manipulating the framing or familiarity of the quantities used in intertermporal choice tasks can make hyperbolic discounting nearly disappear at the median.[12]

3. Errors in probabilistic and statistical judgment. Some of the oldest and best-established anomalies in behavioral economics and psychology concern the formation of beliefs from statistical information: base-rate neglect, conservatism in belief updating, and over-inference from weak evidence. Adaptation to noisy representations provides a unified explanation for all of these biases. Under-response to the degree to which the prior probabilities of two hypotheses are unequal ("base-rate neglect"), and under-response to the degree to which evidence is more likely under one hypothesis than the other ("conservatism") are both examples of behavioral attenuation, and can be understood as optimal Bayesian responses to noise in the internal representation of the prior (in the first case) or in the strength of the evidence (in the second). Consistent with this, Enke and Graeber (2023) show that the sensitivity of beliefs to both the prior and the evidence falls substantially as elicited cognitive uncertainty rises. They also show that making the problem noisier (using a compound or ambiguous lottery) compresses beliefs further.

Problems differ in the way in which one should expect the likelihood of the presented evidence to be encoded, and this results in correspondingly different biases. In the experiment of DuCharme (1970), subjects are shown a series of heights drawn from one of two height distributions (men vs. women), and must assign a probability to the distribution. In such a case, subjects have little reason to be able to sharply discriminate between draws that are slightly more likely to be men's heights and those slightly more likely to be women's; one should thus expect the strength of evidence to be encoded as

[11] The cognitive noise interpretation can also explain the difference in apparent risk preferences obtained using different methods of preference elicitation, as discussed in Section 4.2 below.

[12] Findings in this literature suggests that hyperbolic discounting is due not to true declining impatience, but instead to an internal inconsistency called subadditive discounting Read (2001); recent evidence suggests that subadditive discount in turn is largely driven by cognitive noise-induced regression bias (Enke, Graeber and Oprea, 2025; Oprea and Vieider, 2026b).

a variable that could be of either sign, with regression bias shading the estimate toward less influence of the evidence in either direction ("conservatism").

But in another classic type of experiment, subjects are told that two possible urns each contain balls of only two types, with one of the urns containing more of one type than the other does. In a case of this kind, when a ball is drawn it is quite obvious that this event is more likely under one hypothesis than the other, though a subject may have an imprecise sense of exactly *how much* it is more likely. In this case, the *sign* of the belief adjustment should be clear, while there may be imprecision in estimating *how much* to adjust in that direction. In this case, regression bias in a subject's estimate of the diagnosticity of the evidence leads to *over-response* to weak evidence (weaker than usual), together with under-inference from stronger evidence (more diagnostic than usual), as Augenblick, Lazarus and Thaler (2025) document.

4. Over- and under-confidence. A long-established finding in the belief calibration literature is the *hard-easy effect:* people over-estimate their performance on difficult tasks and under-estimate it on easy ones (e.g., Erev, Wallsten and Budescu, 1994). Moore and Healy (2008) pair this with a second, seemingly opposite regularity: on those same hard tasks people rank their *relative* performance too low, and on easy tasks too high (Kruger, 1999). They further show that both biases follow from a single Bayesian model rooted in cognitive imprecision.

If people access only a noisy signal of their own performance, their estimate should regresses toward a prior: pulled *up* (over-estimation) on a hard task, and pulled *down* (under-estimation) on an easy one. The reversal in relative judgments follows from the same model, if people know even less about others than about themselves, so that their estimates of others' scores are noisier and regress *more* strongly toward the same prior. On a hard task, then, a person inflates her estimate of her own low score, but inflates her guess about others' scores *even more*, thus ranking herself *below* average at the same time that she is over-estimates her absolute performance; similarly, but in reverse, on an easy task. Moore and Healy (2008) elicit subjects' full belief distributions over their own score and find that less precise beliefs are associated with larger biases, just as a noisy-cognition model predicts.

These examples cite only some of the best-known biases that appear to be generated (at least in part) by regression bias in response to cognitive noise. Enke et al. (2026) run 30 distinct experiments, covering many kinds of decisions besides these — saving, effort provision, budget allocation, portfolio allocation, and public-goods provision, among others. They show in nearly all of them that (i) most people are cognitively uncertain and (ii) more cognitively uncertain subjects exhibit significantly more "behavioral attenuation" — less sensitivity to variation in payoff-relevant parameters.

3.3 Bayesian repulsion

Result (4) tells us the value of a linear regression coefficient, but the relationship between x and the average estimate $m(x) \equiv \mathrm{E}[\hat{x}\,|x]$ need not be linear (though in our solved example (3) it is). Thus (4) doesn't tell us the local effect of varying the parameter x at a particular point in its range. In the linear-Gaussian example, $m'(x) = \gamma < 1$ globally, but this strong conclusion depends on a special feature of that example — that encoding precision is uniform over the range of x.

When precision is non-uniform, the bias implied by Bayesian decoding is more complicated.[13] In a high-precision asymptotic limit, Prat-Carrabin and Woodford (2022) show that the bias is approximately

$$b(x) \;\equiv\; m(x) - x \;\approx\; \frac{1}{I(x)}\left[\frac{p'(x)}{p(x)} - \frac{I'(x)}{I(x)}\right], \tag{5}$$

where $p(x)$ is the prior density and $I(x)$ is the Fisher information implied by the noisy encoding, a local measure of the precision of representation at each value of x.

The first term inside the brackets corresponds to the regression bias discussed above: when the Fisher information is constant (as in the linear-Gaussian example), the bias is in the direction of increasing prior density (toward the prior mean, for a Gaussian prior). But when precision is not uniform, there is an additional bias term, indicating repulsion of the estimated value $\hat{x}$ from the direction in which values are encoded with greater precision (thus where $I(x)$ is greater). This second effect can outweigh the pull toward greater prior probability if the non-uniformity of precision is strong enough, giving rise to supposedly "anti-Bayesian" biases in some perceptual domains — effects actually predicted by a Bayesian model in which the prior is calibrated to the environmental frequency distribution of stimuli and the non-uniformity of encoding precision is calibrated to behavioral and neural evidence (Wei and Stocker, 2015).

Moreover, even when $\log p(x)$ is concave (as with a Gaussian prior), so that $p'(x)/p(x)$ is decreasing in x, $b(x)$ need not be globally decreasing: it can be locally increasing where the precision of encoding changes sufficiently rapidly, so that $m(x)$ increases *more* than one-for-one with x over some range. Optimal adaptation to cognitive noise can thus produce *excess sensitivity* to parameter variation (locally), and not just attenuation.

Several important biases from behavioral economics illustrate the effects of non-uniform cognitive noise.

1. Diminishing sensitivity to magnitudes. Extensive magnitudes—length, area, weight, duration, number — are recognized as positive with near-certainty, yet nearby magnitudess are discriminated ever less precisely the larger they get. A classic statement is Weber's Law, according to which discriminability depends on the ratio, or percentage difference, of two magnitudes rather than their absolute difference (Fechner, 1860). A similar logarithmic compression is visible in the neural code for number, the "mental number line" (Dehaene, 2003; Nieder, 2016).

This inhomogeneity can be captured by supposing that a magnitude x has an internal representation of the form (1) with $f(x) = \log x$, so that

$$r \;\sim\; N(\log x,\; \nu^2), \tag{6}$$

where ν determines the percentage difference required for a given degree of discriminability. If the prior is also log-normal, $\log x \sim N(\mu, \sigma^2)$, the optimal (posterior-mean) estimate is a power function of the true magnitude,

$$m(x) \;\propto\; x^{\gamma}, \tag{7}$$

with $0 < \gamma < 1$ again given by (3) — now an elasticity rather than a slope (Khaw, Li and Woodford, 2021). In this case $m(x)$ is strictly concave, with excess sensitivity

[13] See the Supplementary Information for additional discussion.

($m'(x) > 1$) for small enough x, but increasingly greater attenuation ($m'(x) < 1$) when x is larger. The relationship is predicted to be more concave (smaller γ) if coding precision falls. An important example is given by numerosity estimation: estimated number is a concave function of true number, and reducing precision makes it more concave (Anobile, Cicchini and Burr, 2012; Prat-Carrabin and Woodford, 2026).

Diminishing sensitivity is also an important regularity in behavioral economics, as perhaps first noted by Kahneman and Tversky (1979). Enke et al. (2026) document diminishing sensitivity directly across a wide range of economic tasks, and show that elicited cognitive uncertainty rises with the distance of an intrinsically positive decision parameter from zero, finding exactly the fall in coding precision that (6) first proposed as an explanation for Weber's Law in sensory domains. As cognitive noise models predict, they find that locally measured cognitive uncertainty strongly predicts locally measured sensitivity to parameters. The next three applications — small-stakes risk-aversion, gain-loss asymmetry, and present bias — can each be considered more specific applications of this general regularity.

2. Small-stakes risk aversion. Kahneman and Tversky (1979) motivated their nonlinear value function by analogy with the diminishing sensitivity observed in sensory domains; Khaw, Li and Woodford (2021) propose that, just as with perceptual illusions, the nonlinear way behavioral rules respond to monetary payoffs can be explained as optimal adaptation to noisy coding of the amounts. If monetary amounts are encoded logarithmically as in (6), an optimal decision rule responds on average to a payoff x as if its value were the concave function (7). This can explain small-stakes risk aversion, and also why apparent risk aversion is greater for subjects whose choices are noisier: a larger ν in (6) implies both greater trial-to-trial variation in valuations and a smaller γ in (7), so that average valuations depart farther from risk-neutrality.[14]

3. Stake-size effects. Measured risk attitudes also shift systematically with the stakes involved. Khaw, Li and Woodford (2026) have subjects report a willingness to pay for a large set of gambles offering an amount X with probability p, and find that the average relative risk premium $\rho \equiv \log(WTP/pX)$ declines roughly linearly in $\log|X|$, for both gains ($X > 0$) and losses ($X < 0$). This regularity is exactly what an optimal bid based on a noisy representation (6) of the payoff implies, with the predicted elasticity of ρ with respect to $|X|$ equal to $\gamma - 1 < 0$. The elasticity further varies with p in a way that makes sense given the reduced incentive for encoding precision when p is smaller. Oprea and Vieider (2026b) find "stakes effects" on apparent risk preferences from simply changing the denomination of currency in lotteries; this suggests that the precision of encoding of payoffs depends on the range of numbers used, and that it is not even uniform in *percentage* terms — that infrequently encountered numbers are encoded in a particularly noisy way.[15]

4. Excess sensitivity to losses. It is well-established that in choices under risk,

[14] These two aspects of choice under risk are shown to be correlated across subjects in Khaw, Li and Woodford (2021) and Barretto-Garcia et al. (2023). The latter paper also presents evidence from brain scans that subjects with less precise number coding in the brain are less accurate in numerosity comparisons, make noisier choices under risk, and exhibit greater risk aversion.

[15] These authors find similar results for the "absolute magnitude effect" in delay discounting (Thaler, 1981; Frederick, Loewenstein and O'Donoghue, 2002), suggesting that stakes effects there are also driven by noisier encoding of infrequently encountered numbers.

people respond differently to prospective gains and losses of the same size. In choices between an even chance of gaining X or losing the same amount and getting zero for sure, subjects are more likely to take the gamble when X is small, but to decline it when X is large. If for a given loss L, we let $G^*(L)$ be the size of gain required for indifference between the gamble and zero for sure, Bouchouicha, Li and Vieider (2025) find that $G^*(L) = L^\eta$ with $\eta > 1$ — a pattern that they attribute to greater *sensitivity* to the size of losses, rather than "loss aversion" in the traditional sense. This is what a decision rule optimally adapted to cognitive noise implies if the sizes of both G and L are encoded as in (6), but *losses are encoded with greater precision*: a larger γ for losses implies $\eta = \gamma_{loss}/\gamma_{gain} > 1$. The account is parsimonious — it assumes the same weight on gains and losses in the objective that is maximized — and is consistent with documented evidence of greater neural (Tom et al., 2007) and attentional (Pachur et al., 2018; Hirmas, Engelmann and van der Weele, 2024) sensitivity to losses.[16]

5. Present bias. It is well-established that people place disproportionate weight on immediate rewards relative to ones that are even slightly delayed. This can be explained as a repulsion effect. If delays are represented with a scheme like (6), as with many other extensive magnitudes, the shortest delays should be effectively repelled from zero and perceived as disproportionately long. This produces a premium on immediacy that dissipates once some delay is already present. Vieider (2025) develops this theoretically, deriving present bias from the noisy near-zero coding of delay, and distinguishing this effect from the hyperbolic compression discussed in Section 3.2.1. Oprea and Vieider (2026b) provide experimental evidence, showing that manipulations aimed at reducing relative noise in internal representations of delays can sharply reduce measured present bias, eliminating it for the median subject.

6. Inverse-S probability weighting. A more complex nonlinear bias arises for variables that necessarily lie in $[0, 1]$, such as the probability of an event or a relative frequency. Average estimates of the fraction are commonly an inverse-S function of its true value p: excess sensitivity ($m'(p) > 1$) near 0 and 1, and relative insensitivity ($m'(p) < 1$) for intermediate values. Zhang and Maloney (2012) show this is often well approximated by a linear-in-log-odds form,

$$\log \frac{m(p)}{1 - m(p)} \;=\; \alpha + \gamma \cdot \log \frac{p}{1 - p}, \tag{8}$$

with $0 < \gamma < 1$. This too reflects optimal adaptation to noise, greatest for intermediate p — and indeed people discriminate more extreme probabilities more accurately (Frydman and Jin, 2025).

The form (8) represents an approximation to the Bayesian posterior-mean estimate in a model of log-odds encoding,

$$r \;\sim\; N\Big(\log \tfrac{p}{1-p},\; \nu^2\Big), \qquad \log \tfrac{p}{1-p} \;\sim\; N\big(\mu,\; \sigma^2\big)\,, \tag{9}$$

with γ again from (3), so that greater imprecision sharpens the inverse-S (Khaw, Li and Woodford, 2026). Enke et al. (2026) find bias of a similar shape — again stronger under greater cognitive uncertainty — in their "Recall" and "Signal Aggregation" tasks,

[16] For an alternative noisy-coding account of gain-loss asymmetry, with a different assumption about the noisy representation, see Villas-Boas (2024).

where the varying input is an objective proportion with a correct answer, so that the inverse-S clearly reflects a bias rather than an aspect of preferences. It thus appears to be a general feature of responses to information about fractions or proportions, rather than something specific to risk.

Our examples here have all involved repulsion from boundary values, but Bayesian repulsion is also possible when intermediate values are recognizable with particular precision. In the case of estimates of proportions, a more complex pattern of bias (with repulsion from 0, 1 *and also* from 0.5, and two disjoint regions over which there is attenuation) is observed in cases where proportions are displayed (for example, using a pie chart) in a way that makes deviations from exactly 1/2 especially salient (Hollands and Dyre, 2000). This confirms that the relevant principle involves repulsion from values where discrimination is more acute, rather than repulsion from boundary values as such.[17]

4 Context and Experience

The third observation motivating cognitive models is that people are not only imprecise (Section 2) and biased (Section 3); their behavior is also strongly influenced by payoff-irrelevant factors that should not matter in rational-choice models. Seemingly irrelevant information can have significant effects on behavior. For instance, Hartzmark and Shue (2018) find that the market responds less favorably to a firm's earnings surprise when the earnings surprises of other firms announced the previous day happened to be strong — as though today's news is judged in contrast to yesterday's — and Genesove and Mayer (2001) find that the price a person paid for a house has a strong effect on the price at which they are later willing to sell it.

People's decisions are also strongly influenced by the way that information is presented. For example, Frydman and Wang (2020) find that when a brokerage makes a stock's gains and losses more salient — without providing any new information — the likelihood that investors will sell the stock changes. Tversky and Kahneman (1981) show that people choose systematically differently between identical public-health interventions depending on whether the outcomes are described in terms of lives saved or lives lost – a celebrated example of a framing effect. Finally, people's behavior is strongly influenced by past experiences: Malmendier and Nagel (2011) show that the macroeconomic conditions a person has experienced, even in their distant past, significantly affect their willingness to take financial risk. As with biases, these effects have typically been catalogued as a menagerie of separate irrationalities: anchoring, framing, default effects, contrast, reference dependence, peer effects etc.

Noisy cognition models create immediate scope for these kinds of context effects, and thus offer the prospect of a unified explanation. Once one recognizes that the payoff-relevant information, that should suffice to determine a decision according to a rational-choice model, is likely to be encoded only imprecisely, it ceases to be obvious that people's decisions will always be better if they can be insulated from all "extraneous" information. Aspects of the context in which the decision is made can provide information that can augment the imperfect information about "fundamentals." Moreover, sensitivity to context (including experience far in the past) can be useful not only in interpreting the noisy observations about one's current problem that one has made, but can also help

[17] See Polania, Woodford and Ruff (2019) for an example of repulsion from an intermediate region where coding precision is high in the case of economic valuations.

to direct one's limited attention to aspects of the situation that are likely to be worth encoding.

Context effects and experience effects are pervasive in the case of sensory judgments (e.g., Laming (2011)), and the neuroscience and psychology literatures have shown that many of these can be understood as adaptive responses to noisy coding (Schwartz, Hsu and Dayan, 2007). First, sensory features are not only encoded with noise, but are typically encoded in *relative* terms — in ways that depend on contrasts with similar features, nearby in either space or time; and context-sensitive coding of this kind, rather than a "bug," can be argued to make more efficient use of limited cognitive resources by making representations less redundant.[18] And the type of encoding that is efficient will depend on statistical regularities of the environment. Such regularities must be learned as an organism adapts to new environments, so that efficient coding should also be shaped by experience.

And second, noisy sensory information is *interpreted* in the light of both immediate context and more distant experiences, for reasons that make sense in a Bayesian framework — though again this depends on the statistical regularities of the environment, that need to be learned. We believe that similar principles can help us to make sense of many of the kinds of context effects observed in the case of economic decisions.

Our review of such effects will have to be brief, focusing on two ways in which they might be adaptive under noisy coding: the adaptive use of context in *decoding* noisy representations, and also in the way situations are *encoded*.

4.1 Contextual decoding

A first reason for context effects in a noisy-coding framework is our observation in Section 3 that once representations are imprecise, optimal decoding must be relative to a *prior*. Importantly, there is no reason to assume that this prior will be informed *only* by information that would be used by an ideal rational decision maker. Even in artificial lab settings, the priors subjects bring may be adapted not to the lab task but to serving them well (on average) in everyday life. Here we sketch some of the observed behaviors that may well be interpretable as examples of the use of contextual information in Bayesian decoding.

1. Experience effects. In perception, there is considerable evidence that responses to a given physical stimulus vary with the distribution of stimuli to which a subject has been exposed in a given context. Hollingworth's "central tendency of judgment" (Hollingworth, 1910) asserts that the value toward which magnitude estimates regress depends on the average stimulus magnitude encountered in that environment; the crossover point from over- to under-estimation moves with the range of magnitudes recently encountered (Petzschner, Glasauer and Stephan, 2015; Xiang et al., 2021). One should expect this in a Bayesian model of regression bias (Section 3.1), if the prior in the light of which people interpret their noisy signals is adaptively changing to match the distribution of magnitudes encountered in one's current environment.

There is evidence that economic biases adapt in a similar way. Bouchouicha, Li and

[18] Representations of the value of economic options in the brain have been argued to be context-sensitive in a similar way, and perhaps on similar grounds (Padoa-Schioppa, 2009; Louie, Khaw and Glimcher, 2013; Conen and Padoa-Schioppa, 2019; Glimcher, 2022).

Vieider (2025) ask subjects to choose between a mixed gamble (equal chance of a gain G or a loss L) and nothing for sure, following an adaptation phase in which the gains they are exposed to are larger or smaller on average, depending on the treatment; the test set is identical across treatments. They find that subjects who were exposed to larger gains make more risk-seeking choices — requiring a lower G to accept, for a given L — consistent with their decisions being based on a noisy representation r_G, decoded using a prior that is shifted by the distribution of gains seen during adaptation.

Charles, Frydman and Kilic (2026) find similar adaptation in willingness-to-pay bids,[19] and show how this can reflect Bayesian updating of the prior used on each trial. The experience effects of Malmendier and Nagel (2011) can be given a similar Bayesian interpretation, if personally experienced events are encoded with less noise than other information sources.

2. Social learning and peer effects. Priors can be affected not only by one's own past experience but also those of others. Humans are social learners and partly for this reason, peer effects are observed in many decisions when a decision maker can observe others' choices before committing to their own. A natural interpretation is that observing others shifts one's prior about the value of a given option, which makes sense if the options that others face are similar to one's own. Moutoussis, Dolan and Dayan (2016) give exactly this Bayesian account of peer effects in young people's intertemporal choices: subjects adjust toward the peer's observed behavior but less than all the way, the effect is stronger for those whose own choices are noisier, and it declines with age (i.e., experience) — all as expected if observing a peer provides information.

Graeber and Enke (2026) document this peer assimilation directly and provide evidence for its informational basis: across incentivized effort, investment, and belief-updating tasks, a comparison point that conveys information about the right *decision* — as another's observed choice does — pulls one's own choice toward it. And the pull weakens as subjects resolve their own uncertainty about the problem, and weakens substantially when the comparators are transparently randomized so as to carry no information.

3. Sunk costs. The "sunk-cost fallacy" — persisting with a plan chosen in the past even when it is no longer optimal today — is usually read as an inability to reason correctly about rewards, but as Baliga and Ely (2011) note, it can be an optimal rule when a decision maker remembers their past choice but not the information behind it. If current information alone is insufficient to determine the best action, and memory of what one once knew is imperfect — both potential consequences of cognitive noise — then awareness of one's past choice can convey useful information.

4. Contrasts, attractors and decoys. A similar information channel produces subtle effects when the contextual cue is a *comparator*: a specific rival value, present in the situation or cued by it against which an option is assessed. For example, workers who learn that comparable colleagues earn more, supply less effort and attend less (Breza, Kaur and Shamdasani, 2018) and report lower satisfaction and greater intention to search for other jobs (Card et al., 2012); here the higher comparison wage lowers, rather than raises, the assessment of one's own wage. The "asymmetric dominance effect" in marketing is similar: adding a "decoy good" dominated by good A but not by good B raises preference for A over B (Huber, Payne and Puto, 1982).

[19] See also Khaw, Glimcher and Louie (2017).

Whether a comparison pushes behavior toward or away from the comparator turns on what the comparator is potentially informative about. When the comparator relates to the *fundamentals* of one's situation, it produces *contrast*: learning that someone else's job pays more leads one to assess one's own job offer as less attractive, and seeing that a good is very evidently more valuable than a decoy leads one to assess that good as more attractive (while one's evaluation of goods less easily compared with the decoy may be little affected).[20] When the comparison bears instead on the right *decision* — what others chose or received — it produces *assimilation*: a cue to the correct choice draws one towards that choice.

Graeber and Enke (2026) demonstrate both of these effects experimentally, and find two results. First, the same quantity has opposite signed effects depending on its role: a comparison workload *contrasts* when it is a parameter of one's problem but *assimilates* when workload is what must be chosen. Second, both effects shrink when a subject's uncertainty about the problem is reduced, and shrink further when the comparator is randomized to carry no information. A systematic classification of more than 100 prior experimental and field studies confirms that this input-versus-output distinction predicts the sign of comparison effects throughout the literature.Contrast and assimilation effects can thus both be explained as implications of noisy cognition, driven by the fact that people use comparisons as sources of information.

4. Reference points. Reference-point effects can be viewed similarly, as cases in which comparison of a policy-relevant parameter to some other value (the reference point) provides additional information. Dean et al. (2026) explain a wide range of such effects in Bayesian terms: if one is able to see how the relevant parameter ranks against the reference far more precisely than one can recognize its absolute value, awareness of the comparison with the reference should influence one's posterior regarding the relevant parameter.

The particular value used as a reference point is then an aspect of context that can influence subjective valuations; and efficiency considerations can help to explain which reference points should be used in a given context. Bordalo, Gennaioli and Shleifer (2020) argue that values encountered in the past bias current choices, because they are brought to mind by similarity-based recall. Dean et al. (2026) show that this can be regarded as an efficient selection principle: among candidate references, the one whose comparison is most informative about the current problem is the one most worth recalling, and under natural assumptions that most-informative reference takes precisely the similarity-based form.

4.2 Contextual encoding

We have also emphasized above that the biases implied by cognitive noise models depend on the degree of precision of the decision maker's internal representation. In multidimensional problems (essentially all economic problems), one can say further that this depends on the degree of precision with which *each* of the many features of the situation are encoded. Context can therefore be expected to influence biases to the extent that it influences either the overall amount of attention that is devoted to precisely recognizing the current situation, or the particular features of the situation that are more

[20] Bayesian explanations of the asymmetric dominance effect are developed in detail in Natenzon (2019) and Shubatt and Yang (2026).

precisely recognized. Here we illustrate some of the effects that can be understood in this way.

1. Incentives for precision. First, the overall precision with which aspects of one's environment are represented may respond to the strength of incentives.[21] Dean and Neligh (2023) find that the accuracy of subjects' responses increases when the financial incentive for a correct answer is larger. More generally, the same piece of information may be encoded with different precision when other features of the decision problem change how much precise encoding of that parameter matters for one's payoff. Khaw, Li and Woodford (2026) use this idea to explain the variable stake-size effects in lottery valuation depending on the probability with which a lottery pays off: the incentive for precise encoding of the payoff amount is smaller when the probability of a payoff is smaller, and the less precise encoding increases the regression effect in subjects' responses to the stated payoff.[22]

2. Adaptation to environmental statistics. The incentives for precise encoding of a parameter can also be modulated by changes in the distribution of values encountered for it. Frydman and Jin (2022) compare choices between a gamble (offering X with probability p) and a certain amount C in two different settings: under a "high-volatility" treatment, X and C are drawn from a distribution with three times the standard deviation used in the "low-volatility" treatment. They observe noisier choice (probabilities biased toward 50 percent) on identical decision problems when they occur in the high-volatility context.[23] This is puzzling from the standpoint of a random-utility model if one expects the distribution of utilities associated with a gamble to depend only on the features of that particular gamble. It can be explained, instead, if decisions are random because they are based on noisy representations of the payoffs on a given trial, and the precision of encoding of these quantities varies with the distribution of values that may need to be encoded.

Context-dependent precision makes sense, in turn, if there are costs of a more informative representation (in an information-theoretic sense), since representing payoffs with a given precision (in terms of the value of ν in (6), for example) requires a communication channel with greater capacity when the range of payoffs that the brain must represent is wider. Prat-Carrabin and Woodford (2026) show that the degree to which choices become noisier in the Frydman-Jin experiment can be explained by a model of "efficient coding" in which the endogenous precision of number encoding in a given environment trades off the gains from more accurate choices against an assumed cost of increased information transmission.[24]

When encoding precision adjusts optimally in this way, the standard deviation of coding noise ν increases when the prior uncertainty σ about the parameter increases, but less

[21] The amount by which precision should increase when incentives increase depends on the assumed cost of precision; Caplin and Dean (2015) discuss the measurement of such cost functions.

[22] Khaw *et al.* obtain a good quantitative fit using the same cost function for precision that Prat-Carrabin and Woodford (2026) find explains the endogenous variation of number-representation precision in numerosity-estimation and number-averaging experiments.

[23] Frydman and Nunnari (forthcoming) similarly find an attenuated response to payoff variation in a coordination game when the range of payoffs is greater.

[24] The cost function for precision that they use is the same one that they show can account quantitatively for similar effects of changing the range of numbers used in an experimental treatment on the precision of numerosity estimates.

than in proportion to the increase in σ.[25] This means that regression effects should be weaker (the value of γ implied by (3) is closer to 1) when prior uncertainty is greater. Prat-Carrabin and Woodford (2026) find this to be true for numerosity estimates; Mackowiak and Wiederholt (2009) propose a similar model of endogenous precision to explain why the prices set by firms respond more rapidly to more volatile shocks.

3. Salience and attention. In a complex environment, certain features of the environment (said to be more salient) draw more attention, giving them greater influence on choice (Li and Camerer, 2022). Among the best-established determinants of salience is the degree to which a feature contrasts with its surroundings or with what one would normally expect (Itti and Baldi, 2009). More salient features are assigned more cognitive resources to represent them, resulting in more precise representation (Rust and Cohen, 2022) This is an important way in which context affects the precision of representation of particular features of a situation. A decision rule that is optimally adapted to the degree of noise with which different features are represented should then be more sensitive to the values of such features, for reasons discussed in Section 3.2.

This provides a foundation for the proposal by Bordalo, Gennaioli and Shleifer (2012; 2013; 2022) that decisions put greater weight (relative to a correct assessment) on the attributes of an option that are most salient, where salience depends on how much that attribute contrasts with the corresponding attributes of the other options in the consideration set.[26] For a lottery ticket, for example, the outcome in which one wins is far more salient than other outcomes, because what one gets then is so unlike what one gets from other options — resulting in over-valuation. Because salience depends on the contrast between an item's attributes and those it is compared with, it is yet another source of context effects, including apparent preference reversals across elicitation methods. Experimental support for the predicted effects is provided by Dertwinkel-Kalt et al. (2017; 2022), Frydman and Mormann (2018), Dertwinkel-Kalt and Köster (2020; 2025), and Bondi et al. (2026).

4. Experience and familiarity. An important puzzle in the literature on choice under risk is the so-called "description-experience gap." People show evidence of anomalies like probability weighting when given descriptions of the primitives of lotteries (decisions from description), but instead show the opposite tendencies (over-sensitivity to probabilities in their choices) when they instead learn about lotteries by sampling outcomes (decisions from experience) (Hertwig and Erev, 2009). Noisy cognition models offer an explanation, because experience with the consequences of an option can improve the precision with which those consequences are internally represented: experience lowers the coding noise ν, raising the evidence weight γ in (3), hence pulling estimates away from the prior and toward the observed parameter. As discussed above, probability weighting can arise from a noisy internal representation of the form (9) decoded by Bayesian shrinkage. Under this account, because ν falls as more samples from the distribution are observed, forcing subjects to sample increases precision and so reduces the scope for probability distortions. Oprea and Vieider (2026a) show that by forcing subjects to observe large balanced samples from lotteries before choosing, both probability weighting and the description-experience gap largely disappear. Relatedly, Bohren et al. (2025)

[25] The decrease in ν/σ means that coding resources are increased, owing to the increased marginal value of information transmission about a more volatile parameter.

[26] See also Köszegi and Szeidl (2013) for a closely related proposal.

elicit mental representations directly and associate the gap with attention and memory limits, reinforcing the view that such anomalies reflect distorted representations rather than preferences.

5. Elicitation effects. The procedure used to elicit a decision can influence both the allocation of attention and the priors brought to the interpretation of noisy representations, so that the same underlying preferences can yield different behavioral patterns under different methods. For example, a cognitive noise model can explain the long-known fact (Hershey and Schoemaker, 1985) that eliciting the certainty equivalent of a gamble and eliciting the "probability equivalent" of a certain amount yield differing apparent risk attitudes. The departures from risk neutrality in both cases are examples of behavioral attenuation: in each case the elicited quantity is insufficiently sensitive to variation in the stated one. And Enke et al. (2026) confirm that both departures are stronger on trials where greater cognitive uncertainty is expressed.

Similarly, Bouchouicha et al. (2026) find that the fourfold pattern of risk attitudes is robust when values are elicited as certainty equivalents from a choice list, but that only a two-fold pattern survives under simple binary choice (see also Harbaugh, Krause and Vesterlund, 2010). They note that certainty-equivalent elicitation encourages attention to different aspects of the lottery to be valued than what is necessary for choice between the lottery and a specific alternative. To test this, they run a novel binary-choice task in which subjects must first evaluate the lottery (as in a choice-list task), and find that this causes binary-choice behavior to become much more consistent with the valuations elicited in choice-list tasks, as predicted by their cognitive noise explanation.[27] This kind of evidence suggests that noisy cognition may help make sense of "the risk elicitation puzzle" (Pedroni et al., 2017): the tendency of measured risk preferences to correlate only weakly across elicitation procedures.

6. Framing effects. "Framing effects" — like the lives-saved/lives-lost problem of Tversky and Kahneman (1981), in which logically equivalent descriptions of a problem produce systematically different behavior — are easily generated by noisy coding models, if the description of a problem determines which variables are encoded and how precisely. Woodford (2012) shows how this can reflect efficient coding, if framing changes one's prior about the cases that it is more likely to be valuable to distinguish. See Woodford (2012) and Oprea and Vieider (2026b) for applications to classic anomalies of risk and time preference.

5 Conclusion

Even this brief review has, we hope, made it evident that the cognitive noise hypothesis shows great promise as a unifying principle to explain a wide range of phenomena in behavioral economics. But the theory remains incomplete in important respects. While it seems clear that mental representations are imprecise, a better understanding of the determinants of that imprecision will allow us to make sharper predictions. And while the hypothesis of optimal decoding clearly allows us to understand the general character of many observed biases, the process through which people learn to adapt their behavior to the noise in their cognitive processing and to the statistics of their environment deserves further study. Such adaptation surely requires a sufficient degree of experience — but

[27] See also related discussion in Khaw, Li and Woodford (2026).

how much experience, and in what ways should we expect responses to be suboptimal while experience remains inadequate? Moreover, our emphasis on the imprecision with which people recognize payoff-relevant states means that we can hardly assume it to be trivial for people to accurately assess the degree of noise in their mental representations, or the degree of variability in their environment. Even with extensive experience, it may be more realistic to expect only approximate optimality of the way in which mental representations are decoded (Aridor, Azeredo da Silveira and Woodford, 2025).

Thus further study of processes of learning and adaptation — both how more efficient encoding approaches are learned, and how more accurate decoding rules are learned — is one of the most crucial requirements for a more complete theory, that can be applied in practice in a more flexible way. As with the work surveyed here, insights from other domains, such as the neuroscience of perception, seem likely to be helpful. Ideas developed by cognitive economists about the determinants of attention allocation, the way in which memory works, the way in which behavior is shaped by experience, and so on, also provide clues to the kinds of phenomena that a more complete theory needs to explain, though these have not yet been given resource-rational foundations. Given the rate at which the work reviewed here has developed over the last five years, we are optimistic that even greater understanding will be possible before long.

References

Anobile, Giovanni, Guido Marco Cicchini, and David C. Burr. 2012. "Linear Mapping of Numbers onto Space Requires Attention." *Cognition*, 122: 454–459.

Aragones, Enriqueta, Itzhak Gilboa, Andrew Postlewaite, and David Schmeidler. 2005. "Fact-Free Learning." *American Economic Review*, 95(5): 1355–1368.

Aridor, Guy, Rava Azeredo da Silveira, and Michael Woodford. 2025. "Information-Constrained Coordination of Economic Behavior." *Journal of Economic Dynamics and Control*, 172: 104985.

Ashourian, Paymon, and Yonatan Loewenstein. 2011. "Bayesian Inference Underlies the Contraction Bias in Delayed Comparison Tasks." *PLoS ONE*, 6(5): e19551.

Augenblick, Ned, Eben Lazarus, and Michael Thaler. 2025. "Overinference from Weak Signals and Underinference from Strong Signals." *Quarterly Journal of Economics*, 140: 335–401.

Bagaïni, Alexandra, Yunrui Liu, Mradula Kapoor, Gayoung Son, Paul-Christian Bürkner, Loreen Tisdall, and Rui Mata. 2025. "A Systematic Review and Meta-analyses of the Temporal Stability and Convergent Validity of Risk Preference Measures." *Nature Human Behaviour*, 9: 700–712.

Baliga, Sandeep, and Jeffrey C. Ely. 2011. "Mnemonomics: The Sunk Cost Fallacy as a Memory Kludge." *American Economic Journal: Microeconomics*, 3(4): 35–67.

Barretto-Garcia, Miguel, Gilles de Hollander, Marcus Grueschow, Rafael Polania, Michael Woodford, and Christian C. Ruff. 2023. "Individual Risk Attitudes Arise from Noise in Neurocognitive Magnitude Representations." *Nature Human Behaviour*, 7(9): 1551–1567.

Bedi, Saurabh, Gilles de Hollander, and Christian Ruff. 2025. "Probability Weighting Arises from Boundary Repulsions of Cognitive Noise." Preprint, bioRxiv, posted September 11, 2025.

Ben-Akiva, Moshe, and Steven R. Lerman. 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand.* Cambridge, MA:MIT Press.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica*, 63(4): 841–890.

Bohren, J. Aislinn, Josh Hascher, Alex Imas, Michael Ungeheuer, and Martin Weber. 2025. "Mental Models of Information: The Role of Attention and Memory." Working paper, SSRN 4731525.

Bondi, Tommaso, Dániel Csaba, Evan Friedman, and Salvatore Nunnari. 2026. "Complexity and Range Effects." Working paper, Cornell University.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012. "Salience Theory of Choice Under Risk." *Quarterly Journal of Economics*, 127: 1243–1285.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2013. "Salience and Consumer Choice." *Journal of Political Economy*, 121: 803–843.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2020. "Memory, Attention, and Choice." *Quarterly Journal of Economics*, 135: 1399–1442.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2022. "Salience." *Annual Review of Economics*, 14: 521–544.

Bouchouicha, Ranoua, Ryan Oprea, Ferdinand M. Vieider, and Jilong Wu. 2026. "Cognitive Frictions and Canonical Patterns in Risk-Taking." Working paper, RISL$\alpha\beta$, Ghent University.

Bouchouicha, Ranoua, Yuchi Li, and Ferdinand M. Vieider. 2025. "Loss-Sensitivity versus Loss-Aversion." Working paper, RISL$\alpha\beta$, Ghent University.

Breza, Emily, Supreet Kaur, and Yogita Shamdasani. 2018. "The Morale Effects of Pay Inequality." *Quarterly Journal of Economics*, 133(2): 611–663.

Britten, Kenneth H., William T. Newsome, Michael N. Shadlen, Simona Celebrini, and J. Anthony Movshon. 1996. "A Relationship between Behavioral Choice and the Visual Responses of Neurons in Macaque MT." *Visual Neuroscience*, 13(1): 87–100.

Caplin, Andrew, and Mark Dean. 2015. "Revealed Preference, Rational Inattention, and Costly Information Acquisition." *American Economic Review*, 105: 2183–2203.

Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez. 2012. "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction." *American Economic Review*, 102(6): 2981–3003.

Charles, Constantin, Cary Frydman, and Mete Kilic. 2026. "History-Dependent Cognitive Defaults." Working paper.

Choi, Syngjoo, Jeongbin Kim, Eungik Lee, and Jungmin Lee. 2022. "Probability Weighting and Cognitive Ability." *Management Science*, 68: 5201–5215.

Conen, Katherine E., and Camillo Padoa-Schioppa. 2019. "Partial Adaptation to the Value Range in the Macaque Orbitofrontal Cortex." *Journal of Neuroscience*, 39(18): 3498–3513.

Dean, Mark, and Nate Neligh. 2023. "Experimental Tests of Rational Inattention." *Journal of Political Economy*, 131: 3415–3461.

Dean, Mark, Ben Enke, Thomas Graeber, and Pietro Ortoleva. 2026. "Reference Points as Information." Working paper, Columbia University.

Dehaene, Stanislas. 2003. "The Neural Basis of the Weber–Fechner Law: A Logarithmic Mental Number Line." *Trends in Cognitive Sciences*, 7(4): 145–147.

DellaVigna, Stefano. 2009. "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature*, 47(2): 315–372.

Dertwinkel-Kalt, Markus, and Mats Köster. 2020. "Salience and Skewness Preferences." *Journal of the European Economic Association*, 18: 2057–2107.

Dertwinkel-Kalt, Markus, and Mats Köster. 2025. "Salient Cues and Complexity." *Management Science*, 71: 428–469.

Dertwinkel-Kalt, Markus, Holger Gerhardt, Gerhard Reiner, Frederik Schwerter, and Louis Strang. 2022. "Concentration Bias in Intertemporal Choice." *Review of Economic Studies*, 89: 1314–1334.

Dertwinkel-Kalt, Markus, Katrin Köhler, Mirjam R.J. Lange, and Tobias Wenzel. 2017. "Demand Shifts Due to Salience Effects: Experimental Evidence." *Journal of the European Economic Association*, 15: 626–653.

Desvousges, William H., F. Reed Johnson, Richard W. Dunford, Kevin J. Boyle, Sara P. Hudson, and K. Nicole Wilson. 1993. "Measuring Natural Resource Damages with Contingent Valuation: Tests of Validity and Reliability." In *Contingent Valuation: A Critical Assessment.* Vol. 220 of *Contributions to Economic Analysis*, , ed. Jerry A. Hausman, 91–164. Amsterdam:North-Holland.

Dhami, Sanjit. 2016. *The Foundations of Behavioral Economic Analysis.* Oxford:Oxford University Press.

Dror, Itiel E., David Charlton, and Ailsa E. Péron. 2006. "Contextual Information Renders Experts Vulnerable to Making Erroneous Identifications." *Forensic Science International*, 156(1): 74–78.

DuCharme, Wesley M. 1970. "Response Bias Explanation of Conservative Human Influence." *Journal of Experimental Psychology*, 85: 66–74.

Ebert, Jane E. J., and Drazen Prelec. 2007. "The Fragility of Time: Time-Insensitivity and Valuation of the Near and Far Future." *Management Science*, 53(9): 1423–1438.

Enke, Benjamin, and Cassidy Shubatt. 2023. "Quantifying Lottery Choice Complexity." *NBER Working Paper No. 31677.*

Enke, Benjamin, and Thomas Graeber. 2023. "Cognitive Uncertainty." *Quarterly Journal of Economics*, 138: 2021–2067.

Enke, Benjamin, Thomas Graeber, and Ryan Oprea. 2023. "Confidence, Self-Selection, and Bias in the Aggregate." *American Economic Review*, 113(7): 1933–1966.

Enke, Benjamin, Thomas Graeber, and Ryan Oprea. 2025. "Complexity and Time." *Journal of the European Economic Association*, 23: 1838–1867.

Enke, Benjamin, Thomas Graeber, Ryan Oprea, and Jeffrey Yang. 2026. "Behavioral Attenuation." Working paper, Harvard University.

Erev, Ido, Thomas S. Wallsten, and David V. Budescu. 1994. "Simultaneous Over- and Underconfidence: The Role of Error in Judgment Processes." *Psychological Review*, 101(3): 519–527.

Faisal, A. Aldo, Luc P.J. Selen, and Daniel M. Wolpert. 2008. "Noise in the Nervous System." *Nature Reviews Neuroscience*, 9: 292–303.

Fechner, Gustav Theodor. 1860. *Elemente der Psychophysik.* Leipzig:Breitkopf und Härtel.

Fennell, John, and Roland Baddeley. 2012. "Uncertainty Plus Prior Equals Rational Bias: An Intuitive Bayesian Probability Weighting Function." *Psychological Review*, 119: 878–887.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue. 2002. "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature*, 40(2): 351–401.

Frydman, Cary, and Baolian Wang. 2020. "The Impact of Salience on Investor Behavior: Evidence from a Natural Experiment." *Journal of Finance*, 75(1): 229–276.

Frydman, Cary, and Lawrence J. Jin. 2022. "Efficient Coding and Risky Choice." *Quarterly Journal of Economics*, 137: 161–213.

Frydman, Cary, and Lawrence J. Jin. 2025. "On the Source and Instability of Probability Weighting." Working paper, University of Southern California.

Frydman, Cary, and Milica Mormann. 2018. "The Role of Salience in Choice under Risk: An Experimental Investigation." Working paper, University of Southern California.

Frydman, Cary, and Salvatore Nunnari. forthcoming. "Coordination with Cognitive Noise." *American Economic Journal: Microeconomics.* Forthcoming.

Gabaix, Xavier. 2014. "A Sparsity-Based Model of Bounded Rationality." *Quarterly Journal of Economics*, 129(4): 1661–1710.

Gabaix, Xavier. 2019. "Behavioral Inattention." In *Handbook of Behavioral Economics: Applications and Foundations.* Vol. 2, , ed. B. Douglas Bernheim, Stefano DellaVigna and David Laibson, 261–343. North-Holland.

Gabaix, Xavier, and David Laibson. 2022. "Myopia and Discounting." National Bureau of Economic Research Working Paper 23254.

Gallagher, Justin. 2014. "Learning about an Infrequent Event: Evidence from Flood Insurance Take-Up in the United States." *American Economic Journal: Applied Economics*, 6(3): 206–233.

Garagnani, Michele, and Ferdinand M. Vieider. 2025. "Economic Consequences of Numerical Adaptation." *Psychological Science*, 36(6): 407–420.

Genesove, David, and Christopher Mayer. 2001. "Loss Aversion and Seller Behavior: Evidence from the Housing Market." *Quarterly Journal of Economics*, 116(4): 1233–1260.

Gershman, Samuel J., and Rahul Bhui. 2020. "Rationally Inattentive Intertemporal Choice." *Nature Communications*, 11: 3365.

Gigerenzer, Gerd, and Daniel G. Goldstein. 1996. "Reasoning the Fast and Frugal Way: Models of Bounded Rationality." *Psychological Review*, 103(4): 650–669.

Glimcher, Paul W. 2011. *Foundations of Neuroeconomic Analysis.* New York:Oxford University Press.

Glimcher, Paul W. 2022. "Efficiently Irrational: Deciphering the Riddle of Human Choice." *Trends in Cognitive Sciences*, 26(8): 669–687.

Graeber, Thomas ⓡ, and Benjamin Enke. 2026. "Comparisons." Working paper.

Harbaugh, William T., Kate Krause, and Lise Vesterlund. 2010. "The Fourfold Pattern of Risk Attitudes in Choice and Pricing Tasks." *Economic Journal*, 120: 595–611.

Hartzmark, Samuel M., and Kelly Shue. 2018. "A Tough Act to Follow: Contrast Effects in Financial Markets." *Journal of Finance*, 73(3): 1567–1613.

Heng, Joseph A., Michael Woodford, and Rafael Polania. 2020. "Efficient Sampling and Noisy Decisions." *eLife*, 9: e54962.

Hershey, John C., and Paul J.H. Schoemaker. 1985. "Probability versus Certainty Equivalence Methods in Utility Measurement: Are They Equivalent?" *Management Science*, 31: 1213–1231.

Hertwig, Ralph, and Ido Erev. 2009. "The Description-Experience Gap in Risky Choice." *Trends in Cognitive Sciences*, 13: 517–523.

Hilbert, Martin. 2012. "Toward a Synthesis of Cognitive Biases: How Noisy Information Processing Can Bias Human Decision Making." *Psychological Bulletin*, 138: 211–237.

Hirmas, Alejandro, Jan B. Engelmann, and Joël van der Weele. 2024. "Individual and Contextual Effects of Attention in Risky Choice." *Experimental Economics*, 27: 1211–1238.

Hodgson, Robert T. 2008. "An Examination of Judge Reliability at a Major U.S. Wine Competition." *Journal of Wine Economics*, 3(2): 105–113.

Hollander, Gilles De, Marcus Grueschow, Franciszek Hennel, and Christian C. Ruff. 2025. "Rapid Changes in Risk Preferences Originate From Bayesian Inference on Parietal Magnitude Representations." Preprint, bioRxiv, posted October 22, 2025.

Hollands, J.G., and Brian P. Dyre. 2000. "Bias in Proportion Judgments: The Cyclical Power Model." *Psychological Review*, 107(3): 500–524.

Hollingworth, Harold L. 1910. "The Central Tendency of Judgment." *Journal of Philosophy*, 7: 461–469.

Huber, Joel, John W. Payne, and Christopher Puto. 1982. "Adding Asymmetrically Dominated Alternatives: Violations of Regularity and the Similarity Hypothesis." *Journal of Consumer Research*, 9: 90–98.

Itti, Laurent, and Pierre Baldi. 2009. "Bayesian Surprise Attracts Human Attention." *Vision Research*, 49: 1295–1306.

Jevons, W. Stanley. 1871. "The Power of Numerical Discrimination." *Nature*, 3: 281–282.

Johnson, Eric J., and Daniel Goldstein. 2003. "Do Defaults Save Lives?" *Science*, 302(5649): 1338–1339.

Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica*, 47: 263–291.

Kahneman, Daniel, and Jack L. Knetsch. 1992. "Valuing Public Goods: The Purchase of Moral Satisfaction." *Journal of Environmental Economics and Management*, 22(1): 57–70.

Khaw, Mel W., Paul W. Glimcher, and Kenway Louie. 2017. "Normalized Value Coding Explains Dynamic Adaptation in the Human Valuation Process." *Proceedings of the National Academy of Sciences*, 114(48): 12696–12701.

Khaw, Mel W., Ziang Li, and Michael Woodford. 2021. "Cognitive Imprecision and Small-Stakes Risk Aversion." *Review of Economic Studies*, 88: 1979–2013.

Khaw, Mel W., Ziang Li, and Michael Woodford. 2026. "Cognitive Imprecision and Stake-Dependent Risk Attitudes." National Bureau of Economic Research NBER Working Paper 30417.

Knill, David C., and Alexandre Pouget. 2004. "The Bayesian Brain: The Role of Uncertainty in Neural Coding and Computation." *Trends in Neurosciences*, 27(12): 712–719.

Körding, Konrad P., and Daniel M. Wolpert. 2004. "Bayesian Integration in Sensorimotor Learning." *Nature*, 427(6971): 244–247.

Krajbich, Ian, Carrie Armel, and Antonio Rangel. 2010. "Visual Fixations and the Computation and Comparison of Value in Simple Choice." *Nature Neuroscience*, 13(10): 1292–1298.

Kriegeskorte, Nioklaus, and Rogier A. Kievit. 2013. "Representational Geometry: Integrating Cognition, Computation, and the Brain." *Trends in Cognitive Sciences*, 17(8): 401–412.

Kruger, Justin. 1999. "Lake Wobegon Be Gone! The "Below-Average Effect" and the Egocentric Nature of Comparative Ability Judgments." *Journal of Personality and Social Psychology*, 77(2): 221–232.

Köszegi, Botond, and Adam Szeidl. 2013. "A Model of Focusing in Economic Choice." *Quarterly Journal of Economics*, 128: 53–104.

Lacetera, Nicola, Devin G. Pope, and Justin R. Sydnor. 2012. "Heuristic Thinking and Limited Attention in the Car Market." *American Economic Review*, 102(5): 2206–2236.

Laming, Donald. 2011. *Human Judgment: The Eye of the Beholder.* Cengage Learning.

Lieder, Falk, and Thomas L. Griffiths. 2020. "Resource-Rational Analysis: Understanding Human Cognition as the Optimal Use of Limited Computational Resources." *Behavioral and Brain Sciences*, 43: e1.

List, John A., Ian Muir, Devin G. Pope, and Gregory Sun. 2023. "Left-Digit Bias at Lyft." *Review of Economic Studies*, 90(6): 3186–3237.

Li, Xiaomin, and Colin F. Camerer. 2022. "Predictable Effects of Visual Salience in Experimental Decisions and Games." *Quarterly Journal of Economics*, 137: 1849–1900.

Louie, Kenway, Mel W. Khaw, and Paul W. Glimcher. 2013. "Normalization Is a General Neural Mechanism for Context-Dependent Decision Making." *Proceedings of the National Academy of Sciences*, 110(15): 6139–6144.

Mackowiak, Bartosz, and Mirko Wiederholt. 2009. "Optimal Sticky Prices under Rational Inattention." *American Economic Review*, 99: 769–803.

Mackowiak, Bartosz, Filip Matejka, and Mirko Wiederholt. 2023. "Rational Inattention: A Review." *Journal of Economic Literature*, 61: 226–273.

Malmendier, Ulrike, and Stefan Nagel. 2011. "Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?" *Quarterly Journal of Economics*, 126(1): 373–416.

Ma, Wei Ji, Jeffrey M. Beck, Peter E. Latham, and Alexandre Pouget. 2006. "Bayesian Inference with Probabilistic Population Codes." *Nature Neuroscience*, 9: 1432–1438.

McFadden, Daniel. 1974. "Conditional Logit Analysis of Qualitative Choice Behavior." In *Frontiers in Econometrics.* , ed. Paul Zarembka, 105–142. New York:Academic Press.

Moore, Don A., and Paul J. Healy. 2008. "The Trouble with Overconfidence." *Psychological Review*, 115: 502–517.

Mosteller, Frederick, and Philip Nogee. 1951. "An Experimental Measurement of Utility." *Journal of Political Economy*, 59(5): 371–404.

Moutoussis, Michael, Raymond J. Dolan, and Peter Dayan. 2016. "How People Use Social Information to Find out What to Want in the Paradigmatic Case of Intertemporal Preferences." *PLoS Computational Biology*, 12: e1004965.

Natenzon, Paulo. 2019. "Random Choice and Learning." *Journal of Political Economy*, 127: 419–457.

Netzer, Nick. 2009. "Evolution of Time Preferences and Attitudes toward Risk." *American Economic Review*, 99(3): 937–955.

Newsome, William T., Kenneth H. Britten, and J. Anthony Movshon. 1989. "Neuronal Correlates of a Perceptual Decision." *Nature*, 341(6237): 52–54.

Nieder, Andreas. 2016. "The Neuronal Code for Number." *Nature Reviews Neuroscience*, 17(6): 366–382.

Olenski, Andrew R., André Zimerman, Stephen Coussens, and Anupam B. Jena. 2020. "Behavioral Heuristics in Coronary-Artery Bypass Graft Surgery." *New England Journal of Medicine*, 382(8): 778–779.

Oprea, Ryan. 2024. "Decisions Under Risk Are Decisions Under Complexity." *American Economic Review*, 114: 3789–3811.

Oprea, Ryan. 2026. "Complexity and Its Measurement." In *Handbook of Experimental Methods in the Social Sciences.* Forthcoming.

Oprea, Ryan, and Ferdinand M. Vieider. 2026a. "Minding the Gap: On the Origins of Probability Weighting and the Description-Experience Gap." Working paper, UC Berkeley.

Oprea, Ryan, and Ferdinand Vieider. 2026b. "Information Processing and the Anomalies of Risk and Time." Working paper, UC Berkeley.

Pachur, Thorsten, Michael Schulte-Mecklenbeck, Ryan O. Murphy, and Ralph Hertwig. 2018. "Prospect Theory Reflects Selective Allocation of Attention." *Journal of Experimental Psychology: General*, 147: 147–169.

Padoa-Schioppa, Camillo. 2009. "Range-Adapting Representation of Economic Value in the Orbitofrontal Cortex." *Journal of Neuroscience*, 29(44): 14004–14014.

Page, Lionel. 2022. *Optimally Irrational: The Good Reasons We Behave the Way We Do.* Cambridge University Press.

Pedroni, Andreas, Renato Frey, Adrian Bruhin, Gilles Dutilh, Ralph Hertwig, and Jörg Rieskamp. 2017. "The Risk Elicitation Puzzle." *Nature Human Behaviour*, 1(11): 803–809.

Petzschner, Frederike H., Stefan Glasauer, and Klaas E. Stephan. 2015. "A Bayesian Perspective on Magnitude Estimation." *Trends in Cognitive Sciences*, 19: 285–293.

Polania, Rafael, Michael Woodford, and Christian C. Ruff. 2019. "Efficient Coding of Subjective Value." *Nature Neuroscience*, 22: 134–142.

Pouget, Alexandre, Peter Dayan, and Richard Zemel. 2000. "Information Processing with Population Codes." *Nature Reviews Neuroscience*, 1: 125–132.

Prat-Carrabin, Arthur, and Michael Woodford. 2022. "Bias and Variance of the Bayesian-Mean Decoder." In *Advances in Neural Information Processing Systems.* Vol. 34, , ed. M. Ranzato et al. Curran Associates.

Prat-Carrabin, Arthur, and Michael Woodford. 2026. "Endogenous Precision of the Number Sense." *eLife*, 13: RP101277.

Ratcliff, Roger, and Gail McKoon. 2008. "The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks." *Neural Computation*, 20(4): 873–922.

Read, Daniel. 2001. "Is Time-Discounting Hyperbolic or Subadditive?" *Journal of Risk and Uncertainty*, 23: 5–32.

Robson, Arthur. 2001. "The Biological Basis of Economic Behavior." *Journal of Economic Literature*, 29: 11–33.

Rust, Nicole C., and Marlene R. Cohen. 2022. "Priority Coding in the Visual System." *Nature Reviews Neuroscience*, 23: 376–388.

Schwartz, Odelia, Anne Hsu, and Peter Dayan. 2007. "Space and Time in Visual Context." *Nature Reviews Neuroscience*, 8: 522–535.

Shubatt, Cassidy, and Jeffrey Yang. 2026. "Valuations under Tradeoff Complexity." Working paper, Harvard University.

Simon, Herbert A. 1955. "A Behavioral Model of Rational Choice." *Quarterly Journal of Economics*, 69(1): 99–118.

Simon, Herbert A. 1990. "Invariants of Human Behavior." *Annual Review of Psychology*, 41: 1–19.

Sims, Christopher A. 2003. "Implications of rational inattention." *Journal of Monetary Economics*, 50: 665–690.

Snowberg, Erik, and Justin Wolfers. 2010. "Explaining the Favorite-Long Shot Bias: Is It Risk-Love or Misperceptions?" *Journal of Political Economy*, 118(4): 723–746.

Steiner, Jakub, and Colin Stewart. 2016. "Perceiving Prospects Properly." *American Economic Review*, 106: 1601–1631.

Stewart, Neil, Nick Chater, and Gordon D. A. Brown. 2006. "Decision by Sampling." *Cognitive Psychology*, 53(1): 1–26.

Stocker, Alan A., and Eero P. Simoncelli. 2006. "Noise Characteristics and Prior Expectations in Human Visual Speed Perception." *Nature Neuroscience*, 9: 578–585.

Strulov-Shlain, Avner. 2023. "More Than a Penny's Worth: Left-Digit Bias and Firm Pricing." *Review of Economic Studies*, 90(5): 2612–2645.

Sydnor, Justin. 2010. "(Over)insuring Modest Risks." *American Economic Journal: Applied Economics*, 2(4): 177–199.

Thaler, Richard. 1981. "Some Empirical Evidence on Dynamic Inconsistency." *Economics Letters*, 8(3): 201–207.

Thurstone, Leon L. 1927. "A Law of Comparative Judgment." *Psychological Review*, 34: 273–286.

Tom, Sabrina M., Craig R. Fox, Christopher Trepel, and Russell A. Poldrack. 2007. "The Neural Basis of Loss Aversion in Decision-Making Under Risk." *Science*, 315: 515–518.

Tversky, Amos, and Daniel Kahneman. 1981. "The Framing of Decisions and the Psychology of Choice." *Science*, 211(4481): 453–458.

Tversky, Amos, and Daniel Kahneman. 1992. "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of Risk and Uncertainty*, 5: 297–323.

Vaeth, Martin. 2026. "Imprecision Attenuates Updating." Working paper, Paris School of Economics.

Vieider, Ferdinand M. 2024. "Decisions Under Uncertainty as Bayesian Inference on Choice Options." *Management Science*, 70: 9014–9030.

Vieider, Ferdinand M. 2025. "Cognitive Foundations of Delay-Discounting." Working paper, RISL$\alpha\beta$, Ghent University.

Vieider, Ferdinand M. 2026. "The Structure of Stochastic Choice." Working paper.

Villas-Boas, J. Miguel. 2024. "Toward an Information-Processing Theory of Loss Aversion." *Marketing Science*, 43: 523–541.

Wei, Xue-Xin, and Alan A. Stocker. 2015. "A Bayesian Observer Model Constrained by Efficient Coding Can Explain 'Anti-Bayesian' Percepts." *Nature Neuroscience*, 18: 1509–1517.

Wei, Xue-Xin, and Michael Woodford. 2025. "Representational Geometry Explains Puzzling Error Distributions in Behavioral Tasks." *Proceedings of the National Academy of Sciences*, 122(4): e2407540122.

Woodford, Michael. 2012. "Prospect Theory as Efficient Perceptual Distortion." *American Economic Review*, 102(3): 41–46.

Xiang, Yang, Thomas Graeber, Benjamin Enke, and Samuel Gershman. 2021. "Bayesian Signatures of Confidence and Central Tendency in Perceptual Judgments." *Attention, Perception & Psychophysics*, 83: 3024–3034.

Zhang, Hang, and Laurence T. Maloney. 2012. "Ubiquitous Log Odds: A Common Representation of Probability and Frequency Distortion in Perception, Action, and Cognition." *Frontiers in Neuroscience*, 6: 1.

Zhang, Hang, Xiangjuan Ren, and Laurence T. Maloney. 2020. "The Bounded Rationality of Probability Distortion." *Proceedings of the National Academy of Sciences*, 117: 22024–22034.

SUPPLEMENTARY INFORMATION: Behavioral Signatures of Cognitive Noise

Ryan Oprea
Haas School of Business

Michael Woodford
Columbia University

August 27, 2026

Here we present further details about the biases that should result if a decision maker's responses are based on the Bayesian posterior mean estimate of the value of some parameter x, conditional on a noisy mental representation r of the decision situation.

1 Covariance of the Bayesian Posterior Mean with the True Value

In this section, we derive formula (4) of the main text. Let $\hat{x}$ be the Bayesian posterior mean estimate of the value of some parameter x,

$$\hat{x}(r) \;\equiv\; \mathrm{E}[x\,|r], \tag{1.1}$$

where r is a noisy mental representation of the decision maker's objective situation. Note that we do not assume that the random variable r must be of the kind specified by equation (1) of the main text. We do not assume that the distribution of r, conditional on the objective situation, must be Gaussian; nor do we assume that the distribution, conditional on the objective situation, must depend only on the true value of x. The variable r may be high-dimensional, and it may depend on (and thus convey information about) a large number of dimensions of a high-dimensional description of the objective situation. Regardless of these details, we can define the conditional expectation (1.1), and this will be a function defined on the support of r. We assume however that x is a scalar-valued parameter, so that $\hat{x}(r)$ is a real number for each possible representation r.

We begin by establishing some general properties of the posterior distribution of values for the parameter x associated with a given representation r. If $\hat{x}(r)$ is the mean of the posterior, then $\hat{x}$ will be a random variable with a conditional expectation

$$\mathrm{E}[\hat{x}] \;=\; \mathrm{E}[x], \tag{1.2}$$

the prior mean value of the parameter. This must be true since

$$\mathrm{E}[\hat{x}] \;=\; \mathrm{E}[\mathrm{E}[x\,|r]] \;=\; \mathrm{E}[x],$$

where the first equality uses the definition (1.1) and the second equality follows from the law of iterated expectations.

Result (1.2) means that a Bayesian posterior mean estimate must be unbiased *on average.* If we let

$$m(x) \equiv \mathrm{E}[\hat{x} \,|x] \qquad (1.3)$$

denote the average estimate when the true parameter value is x, and define the bias in this case as $b(x) \equiv m(x) - x$, then

$$\mathrm{E}[b(x)] = \mathrm{E}[\mathrm{E}[\hat{x} \,|x]] - \mathrm{E}[x] = \mathrm{E}[\hat{x}] - \mathrm{E}[x] = 0.$$

Here the first equality uses the definition (1.3), the second equality follows from the law of iterated expectations, and the final equality follows from (1.2). This result does not imply that the bias must be zero for all x, as we discuss in the main text. But if there are values of x (occurring with positive probability under the prior) for which the parameter is over-estimated on average ($m(x) > x$), then there must also be other values of x (also occurring with positive probability under the prior) for which the parameter is under-estimated on average.

If for any function $h(r)$ we use the notation $\mathrm{E}[h(r) \,|\hat{x}]$ to the mean the expected value of $h(r)$ integrating over the possible representations with the property that $\hat{x}(r) = \hat{x}$, then we also observe that

$$\mathrm{E}[x \,|\hat{x}] = \mathrm{E}[x \,|r] \,|\hat{x}] = \mathrm{E}[\hat{x}(r) \,|\hat{x}] = \hat{x}, \qquad (1.4)$$

for any value of the estimate $\hat{x}$ that is in the support of the marginal distribution for $\hat{x}$. Here the first equality uses the law ot iterated expectations, the second uses the definition (1.1), and the final equality is an identity. Result (1.4) indicates a respect in which Bayesian posterior-mean estimates must be "well-calibrated," even though in general they are not unbiased.

We next consider the variance of the posterior associated with a given representation r. The law of total variance states that for any two random variables x and r, we can decompose the variance of x into two terms,

$$\mathrm{var}(x) = \mathrm{E}[\mathrm{var}(x \,|r)] + \mathrm{var}(\mathrm{E}[x \,|r]). \qquad (1.5)$$

It follows that the average posterior variance associated with a given noisy representation must equal

$$\mathrm{E}[\mathrm{var}(x \,|r)] = \mathrm{var}(x) - \mathrm{var}(\hat{x}) = \mathrm{E}[x^2] - \mathrm{E}[\hat{x}^2]. \qquad (1.6)$$

Here the first equality rearranges (1.5) and simplifies using the notation (1.1); the second equality then follows from the fact that the random variables x and $\hat{x}$ both have the same mean (as established in (1.2)).

We now consider the coefficients $\hat{\alpha}, \hat{\gamma}$ that solve the least-squares problem

$$\min_{\alpha,\gamma} \mathrm{E}[(\hat{x} - \alpha - \gamma x)^2], \qquad (1.7)$$

where the expectation is over the joint distribution for the variables x and $\hat{x}$. Suppose that n decision situations are drawn from the prior distribution, and a noisy representation is drawn for each of them, on the basis of which a posterior mean estimate of the parameter is formed. Then if we let $\{x_i\}$ be the values for the parameter in these instances, and $\{\hat{x}_i\}$ be the corresponding estimates, we can compute an OLS regression of the value of $\hat{x}$ on the true value x (and a constant), using the finite sample $\{x_i, \hat{x}_i\}$. In the limit as $n \to \infty$, the regression coefficients obtained will converge with probability 1 to the values $(\hat{\alpha}, \hat{\gamma})$ that solve problem (1.7). Thus a theoretical characterization of the solution to the problem (1.7) tells us what OLS regression coefficients should be like in the case of a large enough sample.

The solution to problem (1.7) is easily shown to be

$$\hat{\alpha} \;=\; (1-\hat{\gamma})\mathrm{E}[x], \qquad \hat{\gamma} \;=\; \frac{\mathrm{cov}(\hat{x},\, x)}{\mathrm{var}(x)}. \tag{1.8}$$

We also note that

$$\begin{aligned} \mathrm{cov}(\hat{x},\, x) \;&=\; \mathrm{E}[\hat{x}\cdot x] \;-\; \mathrm{E}[x]^2 \\ &=\; \mathrm{E}[\mathrm{E}[\hat{x}\cdot x\,|\hat{x}]] \;-\; \mathrm{E}[x]^2 \;=\; \mathrm{E}[\hat{x}^2] \;-\; \mathrm{E}[x]^2 \\ &=\; \{\mathrm{E}[\hat{x}^2] \;-\; \mathrm{E}[x^2]\} \;+\; \{\mathrm{E}[x^2] \;-\; \mathrm{E}[x]^2\}. \end{aligned} \tag{1.9}$$

Here the first line uses (1.2) to simplify the expression for the covariance; the first equality on the second line uses the law of iterated expectations, and the second equality uses (1.4); and the final line simply adds and subtracts a common term from the expressions in curly brackets. We can then use (1.6) to obtain

$$\mathrm{cov}(\hat{x},\, x) \;=\; \mathrm{var}(x) \;-\; \mathrm{E}[\mathrm{var}(x\,|r)]. \tag{1.10}$$

Substituting (1.10) for $\mathrm{cov}(\hat{x},\, x)$ in the solution (1.8) for $\hat{\gamma}$, we obtain equation (4) in the main text.

Because the posterior variance is necessarily non-negative for each r, (1.10) implies that

$$\mathrm{cov}(\hat{x},\, x) \;\leq\; \mathrm{var}(x),$$

with equality only if with probability 1 there is no uncertainty about x (i.e., a posterior variance of zero) given the representation r. And because $\mathrm{E}[x]$ is the same as $\mathrm{E}[\hat{x}]$ (by (1.2)), (1.9) implies that

$$\mathrm{cov}(\hat{x},\, x) \;=\; \mathrm{E}[\hat{x}^2] \;-\; \mathrm{E}[\hat{x}]^2 \;=\; \mathrm{var}(\hat{x}) \;\geq\; 0,$$

with equality only if the representation r is completely uninformative about x (so that the estimate $\hat{x}$ is equal simply to the prior mean, with probability 1). We can thus show in general that

$$0 \;\leq\; \mathrm{cov}(\hat{x},\, x) \;\leq\; \mathrm{var}(x).$$

Using (1.8), this implies that

$$0 \;\leq \hat{\gamma} \;\leq\; 1. \tag{1.11}$$

Here the lower bound is achieved only if the representation r is completely uninformative about x, while the upper bound is achieved only if the representation r allows the value of x to be reconstructed with perfect precision.

Result (1.11) shows that attenuation is a general consequence of Bayesian optimality in the presence of cognitive noise. The estimated value $\hat{x}$ covaries positively with the true value x, but the slope coefficient of a linear regression of $\hat{x}$ on x is necessarily less than 1.

Result (1.10) also implies that if posterior cognitive uncertainty is increased — in the sense of the average posterior variance $\mathrm{E}[\mathrm{var}(x\,|r)]$ becoming a larger fraction of the prior variance $\mathrm{var}(x)$ — then $\mathrm{cov}(\hat{x}, x)$ will necessarily become a smaller fraction of $\mathrm{var}(x)$, with the implication (using (1.8)) that $\hat{\gamma}$ will necessarily be smaller. In this sense we can show that increased imprecision of the mental representation is necessarily associated with greater attenuation on average.

2 Nonuniformity of Encoding Precision as a Source of Bias

The results of the previous section concern only the overall covariance of the Bayesian estimate $\hat{x}$ with the true value x; the local effect of a change in x need not be qualitatively of the same kind. For example, we have shown above that quite generally, $\hat{\gamma} < 1$; but this does not mean that we must have $m'(x) < 1$ at each possible value of x. Similarly, we have argued that an increase in cognitive uncertainty must be associated with a smaller value of $\hat{\gamma}$; but it is possible for a reduction in encoding precision to increase the size of $m'(x)$ for some values of x. Here we discuss further the nonlinear estimation bias that can result from non-uniform encoding precision over the range of possible values of x.

Let the vector y be a complete description of the objective decision situation (or at least, of the features of the situation that can influence the internal representation r), and let $p(r\,|y)$ indicate the conditional probability of a representation r when the true situation is y.[1] Given an imprecise representation r, the posterior distribution over possible true situations y will then be of the form

$$p(y\,|r) \;=\; \frac{p(y)p(r\,|y)}{p(r)}, \tag{2.12}$$

where $p(y)$ is the prior distribution over possible situations, and $p(r)$ is the implied marginal distribution over possible internal representations.

Let us suppose, as above, that the decision maker's action choice depends on the posterior mean estimate of a scalar quantity $x(y)$.[2] This posterior mean will be given by

$$\hat{x}(r) \;\equiv\; \mathrm{E}[x\,|r] \;=\; \int x(y)p(y\,|r)dy.$$

If we substitute (2.12) into this expression and average over the possible representations r associated with a given objective situation y, we see that the mean estimate in a given

[1] Here we introduce the notation y for the complete description of the situation, as we wish to reserve the notation x for the variable that the decision maker needs to estimate, as in the previous section. We need not assume in general that the distribution of values for r depends solely on the true value of x; it may depend on a more detailed description y of the decision situation.

[2] The relevant quantity x need not be perfectly identifiable even given a complete description y of the observable aspects of the situation; but if it is not, we use the notation $x(y)$ to denote the conditional expectation $\mathrm{E}[x\,|y]$. Then the law of iterated expectations implies that the expectation of $x(y)$ conditional on r will be the same as the expectation of x conditional on r, whether x is a function of y or not.

objective situation will be

$$m(y) \;\equiv\; \mathrm{E}[\hat{x}\,|y] \;=\; \int x(\tilde{y})\Lambda(\tilde{y},\, y)p(\tilde{y})d\tilde{y}, \tag{2.13}$$

where

$$\Lambda(\tilde{y},\, y) \;\equiv\; \int \frac{p(r\,|\tilde{y})p(r\,|y)}{p(r)}dr \tag{2.14}$$

is a kernel function that indicates the degree of similarity of the representations of any two objective situations y and $\tilde{y}$. Here representational similarity depends on the degree of overlap of the two conditional distributions $p(r\,|y)$ and $p(r\,|\tilde{y})$; it is zero in the case of two states such that the supports of the conditional distributions do not overlap at all.

From (2.13) we can see how non-uniformity of encoding precision affects the bias of the Bayesian posterior mean estimate. Consider first the simple case in which y is just a scalar, and corresponds to the value of x; and the noisy representation is given by $r = x + \epsilon$, where ϵ is a mean-zero random variable, drawn independently of the value of x, from a distribution with a symmetric density function $g(\epsilon)$ (i.e., such that $g(-\epsilon) = g(\epsilon)$ for all ϵ). Then the precision of encoding is uniform over the stimulus space, in the sense that the discriminability of any two states x and $\tilde{x}$ on the basis of their noisy representations depends only on the difference $\tilde{x} - x$, regardless of the value of x. If we measure local discriminability using the Fisher information (as in equation (5) in the main text), we find that

$$I(x) \;\equiv\; \mathrm{E}\left[\left(\frac{\partial \log p(r\,|x)}{\partial x}\right)^2 |x\right] \;=\; \int \frac{(g'(\epsilon))^2}{g(\epsilon)}\,d\epsilon \tag{2.15}$$

is the same for all x.

If we further assume a prior over values of x that is symmetric around some value μ, then the marginal distribution over values of r will also be symmetric around the value μ, and the kernel

$$\Lambda(\tilde{x},\, x) \;=\; \int \frac{g(r-x)g(r-\tilde{x})}{p(r)}dr$$

will have the property that if $x = \mu$ (the prior mean), $\Lambda(\tilde{x},\, \mu)$ will be a function of $\tilde{x}$ that is symmetric around the value $\tilde{x} = \mu$ — that is, $\Lambda(\mu - \delta,\, \mu) = \Lambda(\mu + \delta,\, \mu)$ for all δ. It follows that $\Lambda(\tilde{x},\, \mu)p(\tilde{x})$ will be symmetric around the value $\tilde{x} = \mu$, so that (2.13) yields

$$m(\mu) \;=\; \int \tilde{x}\Lambda(\tilde{x},\, \mu)p(\tilde{x})d\tilde{x} \;=\; \mu.$$

The estimate will thus be unbiased, when x is equal to its prior mean.

The symmetry that guarantees this result would be broken if the prior were not symmetric around the value of x at which we wish to compute $m(x)$; thus if x is a point other than the prior mean, there will generally be a bias, with $m(x)$ pulled away from x in the direction of greater prior probability, as in our discussion of regression bias.[3] There will also generally be a bias, even when x is equal to the prior mean, if the prior is not symmetric around

[3]See Vaeth (2026) for further analysis of estimation bias in the case of uniform encoding precision and a prior that is symmetric around its mean.

this mean. But there can also be a bias — even when x is equal to the prior mean, and the prior is symmetric around its mean — if the precision of encoding is not uniform. If $\Lambda(\mu + \delta, \mu) > \Lambda(\mu - \delta, \mu)$ for each $\delta > 0$, then (2.13) will imply $m(\mu) > \mu$, even when the prior is symmetric around μ; and this can happen as a result of the precision of encoding decreasing as x increases. The reverse bias can similarly arise if the precision of encoding increases as x increases: the non-uniformity of precision results in a bias in the direction of decreasing encoding precision — or (as stated in the main text) repulsion from the values that are encoded with greater precision.

2.1 A Discrete Example

A simple example may help to clarify the bias resulting from non-uniform precision. Suppose that the set X of possible values of x is discrete, a sequence of equally-spaced values $(x_1, x_2, \ldots, x_N)$ with $x_i - x_{i-1} = d > 0$ for each i, where $N \geq 6$. And let the set of possible representations r be elements of the same discrete grid. Finally, we suppose that in the case of any state $x_2 \leq x \leq x_{N-1}$, the only possible representations r are the three values $(x - d, x, x + d)$.[4] This means that for any state x in the interior set[5]

$$x_3 \;\leq\; x \;\leq\; x_{N-2}, \tag{2.16}$$

the kernel $\Lambda(\tilde{x},\, x)$ can be non-zero only for the five values of $\tilde{x}$ satisfying $x - 2d \;\leq\; \tilde{x} \;\leq\; x + 2d$. For these values of x, (2.13) then reduces to

$$m(x) \;=\; \sum_{j=-2}^{2} (x + jd)\, \Lambda(x + jd,\, x)\, p(x + jd), \tag{2.17}$$

while the definition of the kernel (2.14) reduces to

$$\Lambda(x + jd,\, x) \;=\; \sum_{k=\max(j-1,-1)}^{\min(j+1,1)} \frac{p(r = x + kd\,|x + jd) \cdot p(r = x + kd\,|x)}{p(r = x + kd)}$$

for any of the five values $-2 \leq j \leq 2$.

Let us further suppose that the prior is uniform on the set X: $p(x) = 1/N$ for all $x \in X$. Then for any x in the set (2.16), we will have the same value of $p(x + jd)$ for each of the values of $x + j$ that matter for evaluating (2.17). In this case, for any x in the set (2.16), the prior is "symmetric around x" over the relevant range, and thus there should be no bias in the average estimate $m(x)$ owing to asymmetry of the prior.

We specify the noise in the representation by assuming that $r = x + d$ with probability $e(x)$ for any $x < x_N$, and that $r = x - d$ with probability $e(x)$ for any $x > x_1$, where $e(x)$ is a measure of encoding imprecision at each point in the state space satisfying $0 \leq e(x) \leq 1/2$

[4]Instead, for extreme values of x we assume that there are only two possible representations r: if $x = x_1$, r must equal either x or $x + d$, while if $x = x_N$, r must equal either $x - d$ or x.

[5]Note that under our assumptions, there must be at least two states in this set. Thus if we compute the bias $b(x)$ for the states in this set, we can discuss whether bias is increasing or decreasing as x increases, over this intermediate range.

for all $x_2 \leq x \leq x_{N-1}$ (and $0 \leq e(x) \leq 1$ for the extreme states).[6] It will be useful to introduce the notation

$$\begin{aligned} \Delta(x_i) \;&\equiv\; e(x_i) \,-\, e(x_{i-1}) \qquad \text{for all}\, 2 \leq i \leq N, \\ \Delta^2(x_i) \;&\equiv\; \Delta(x_i) \,-\, \Delta(x_{i-1}) \qquad \text{for all}\, 3 \leq i \leq N \end{aligned}$$

for the first and second differences of the local imprecision measure. Then the unconditional probability of a representation r is given by

$$p(r = x) \;=\; \frac{1}{N}[e(x-d) + (1 - 2e(x)) + e(x+d)] \;=\; \frac{1}{N}[1 \,+\, \Delta^2(x+d)]$$

for any non-extreme value of x.

If we restrict attention to encoding schemes under which $\Delta^2(x) = D$, a constant value (satisfying $-1 < D < N-1$), for all $x_3 \leq x \leq x_N$, then we will have uniform value for $p(r = x)$ for all $x_2 \leq x \leq x_{N-1}$, equal to $\pi \equiv (1+D)/N$.[7] Because calculations are simplified in this case, we shall consider only cases of this kind. We can then establish a relatively direct connection between non-uniformity of encoding precision and asymmetry of the kernel. In particular, (2.17) now implies that

$$b(x) \;\equiv\; m(x) - x \;=\; \frac{d}{1+D}[\Delta(x+d) \,+\, \Delta(x)] \qquad (2.18)$$

for any x in the interior range (2.16). There is thus a positive bias if imprecision is increasing as x increases (specifically, if $e(x+d) > e(x-d)$), and a negative bias if the reverse is true; the bias can be characterized as repulsion from the direction in which encoding is more precise.

As an especially simple case, suppose that $e(x) = a + cx$, for some coefficients a and c such that $e(x)$ falls within the bounds specified above for all $x \in X$.[8] Then $\Delta(x)$ will have the same sign for all x, and the bias will the same for all x in the range (2.16); the sign of this constant bias is determined by the sign of b. There is repulsion from the lowest value of x when x_1 is encoded with the greatest precision, and repulsion from the highest value of x when x_N is encoded with the greatest precision. (If $b = 0$, we again have a case with uniform precision and a symmetric prior, and find that the bias is zero for all x in (2.16).

But we can also consider the case in which $e(x)$ is a quadratic function of x,

$$e(x) \;=\; a \,+\, \frac{D}{2d^2}(x - x^*)^2,$$

[6]Thus for each of the non-extreme values of x, we have $r = x + \epsilon$, where ϵ is again a mean-zero random variable, drawn independently of the value of x, with a symmetric distribution around $\epsilon = 0$: a discrete analog of the case considered above or in the work of Vaeth (2026).

[7]Note that this uniform probability need not equal $1/N$ because the unconditional probabilities of the two most extreme representations are in general different from π.

[8]Note that for any choice of $a \in (0,\, 1/2)$, there will exist a neighborhood of 0 such that the condition is satisfied by any value of c in this neighborhood. Thus $e(x)$ can be either an increasing or decreasing function of x.

for some coefficients a, D, x^* such that $e(x)$ within the bounds specified above for all $x \in X$.[9] We find in this case that $\Delta^2(x) = D$, the same value for all x, with a sign that depends on whether $e(x)$ is a concave or convex function of x. We also find that

$$\Delta(x) \;=\; \frac{D}{d}(x - \frac{1}{2} - x^*),$$

so that (2.18) implies a bias of the form

$$b(x) \;=\; \frac{2D}{1+D}(x - x^*). \tag{2.19}$$

Suppose further in this latter case that $x_1 < x^* < x_N$. Then if $D > 0$ (the convex case), precision will be first increasing in x (for $x < x^*$), then decreasing in x (for $x > x^*$). We find in this case that the bias is negative for all $x < x^*$ in the set (2.16), but becoming less negative as x approaches x^* from below; and positive for all $x > x^*$ in the set (2.16), becoming less positive as x^* is approached from above. Thus in this case we observe repulsion from the value x^*, the point at which encoding precision is greatest. (This occurs despite that fact that, if we assume that x^* is the exact midpoint between x_1 and x_N, this will mean repulsion from, rather than attraction toward, the prior mean.)

If instead $D < 0$ (the concave case), encoding precision is greatest at the two extremes, and minimized at the interior point $x*$. In this case the bias is positive for all $x < x*$ in the set (2.16), but becoming less positive as x approaches $x*$ from below; and negative for all $x > x*$ in the set (2.16), but becoming less negative as $x*$ is approached from above. Thus in this case there is repulsion from both of the extremes (the points of greatest encoding precision) and toward $x*$ (the point of minimum encoding precision).

The quadratic example also illustrates the fact (already noted in the main text) that one need not observe attenuation locally, at all points in the range of variation in the true state; one can instead have *excess sensitivity* locally, in the sense that the change $\Delta m \equiv m(x') - m(x)$ in the average estimate can be *larger* than the change $\Delta x \equiv x' - x$ in the true parameter value. If we let x, x' be two possible states such that $x_3 \le x < x' \le x_{N-2}$ (so that (2.18) applies at both x and x'), then (2.19) implies that

$$\frac{\Delta m}{\Delta x} \;=\; \frac{1+3D}{1+D}.$$

In the case that $D > 0$, this means that we will have (local) excess sensitivity over the entire range of states in the set (2.16). This in turn is because (2.19) implies that if $D > 0$, the bias is an *increasing* function of x over that entire interval.

One might wonder how this is consistent with our demonstration in the previous section that, under quite general assumptions, we must have $\hat{\gamma} < 1$, implying that the overall covariance between x and $b(x)$ must be *negative*. The answer is that the calculation (2.19) only applies over the intermediate range of states (2.16). The bias must instead be strongly positive for x near its lower bound, and strongly negative for x near its upper bound, even

[9] Again we observe that for any choice of $a \in (0,\, 1/2)$ and any choice of $x*$, there will exist a neighborhood of 0 such that the condition is satisfied by any value of D in this neighborhood. Thus $e(x)$ can be either a concave or convex function of x.

when $D > 0$; and the bias at these extreme values is strong enough to change the sign of the overall covariance between x and $b(x)$, and hence to make the regression coefficient $\hat{\gamma}$ less than 1.[10]

2.2 Implications of Logarithmic Encoding

Here we explain further some calculations referred to in section 3.3 of the main text. As an example of a case with a continuum of possible true values x and representations r where an analytical solution is nonetheless possible, suppose that the true value of x is necessarily positive, and is encoded by a representation r drawn from the distribution

$$r \;\sim\; N(\log x,\, \nu^2), \tag{2.20}$$

for some noise variance ν^2 that is independent of x.[11] In this example, the precision with which nearby values of x can be discriminated decreases for larger values of x; for example, the Fisher information defined in equation (2.15) above decreases inversely with the square of x:

$$I(x) \;=\; \frac{1}{\nu^2 x^2}.$$

Suppose further that the decoding of the noisy representation is optimized for a log-normal prior distribution,

$$\log x \;\sim\; N(\mu,\, \sigma^2),$$

for the true value of x. Under these assumptions, the posterior distribution for x conditional on a particular representation r will also be log-normal,

$$\log x\,|r \;\sim\; N(\hat{\mu}(r),\, \hat{\sigma}^2), \tag{2.21}$$

where

$$\hat{\mu}(r) \;\equiv\; \gamma r \,+\, (1-\gamma)\mu, \qquad \gamma \;\equiv\; \frac{\sigma^2}{\sigma^2+\nu^2},$$

and

$$\hat{\sigma}^2 \;\equiv\; \frac{\sigma^2\nu^2}{\sigma^2+\nu^2} \;=\; (1-\gamma)\sigma^2.$$

Note that this definition of γ is the same as in equation (3) of the main text. In fact, if we were to define $y \equiv \log x$, and suppose that the decision maker's goal is to produce a estimate of y so as to minimize the mean squared error (MSE) of that estimate, then this would be equivalent to the linear-Gaussian model discussed in Section 3.1 of the main text, and the optimal Bayesian estimate would be of the form

$$\hat{y}(r) \;=\; \mathrm{E}[y\,|r] \;=\; \gamma r \,+\, (1-\gamma)\mu,$$

[10] For numerical illustrations of this change in the sign of the bias at the extremes, in examples where the precision of encoding is greatest at an interior point, see the calculations reported in Prat-Carrabin and Woodford (2022), and an empirical example in Polania, Woodford and Ruff (2019).

[11] This is a model often used to capture the imprecise discriminability of extensive sensory magnitudes to which "Weber's Law" applies; it has also been used to model imprecise number representation, where the model is sometimes referred to as one involving a logarithmically compressed "mental number line."

as indicated in equations (2) and (3) of the main text. But here we are interested instead in producing a minimum-MSE estimate of x. The resulting biases are different; for while (2.20) implies *uniform* precision in the encoding of different values of y, the nonlinear transformation $x = \exp(y)$ implies non-uniform precision in the encoding of x.

Because the posterior (2.21) is log-normal, rather than normal (as the posterior for y would be), the minimum-MSE estimate is given by

$$\hat{x}(r) \;=\; \mathrm{E}[x\,|r] \;=\; \exp(\hat{\mu}(r) + (1/2)\hat{\sigma}^2). \tag{2.22}$$

(Note that $\hat{x}(r) \neq \exp(\hat{y}(r))$, because of the $(1/2)\hat{\sigma}^2$ term.)

Because $\hat{\mu}(r)$ is a linear function of r, which is normally distributed (conditional on x), result (2.22) implies that $\hat{x}$ is a log-normally distributed random variable (conditional on x). The mean of this distribution is given by

$$\begin{aligned} \mathrm{E}[\hat{x}\,|x] &= \exp(\gamma \mathrm{E}[r|x] + (1-\gamma)\mu + (1/2)\hat{\sigma}^2 + (1/2)\gamma^2 \mathrm{var}(r|x)) \\ &= \exp(\gamma \log x + (1-\gamma)\mu + (1/2)\hat{\sigma}^2 + (1/2)\gamma^2\nu^2) \\ &= A \cdot x^{\gamma}, \end{aligned}$$

where

$$A \;\equiv\; \exp((1-\gamma)\mu + (1/2)\hat{\sigma}^2 + (1/2)\gamma^2\nu^2) \;=\; \exp((1-\gamma)(\mu + (1/2)(1+\gamma)\sigma^2)).$$

This is a solution of the form asserted in equation (7) of the main text. It implies that the value of x will be over-estimated on average for all $0 < x < x^*$, but under-estimated on average for all $x > x^*$, where

$$x^* \;\equiv\; \exp(\mu + (1/2)(1+\gamma)\sigma^2),$$

a quantity that exceeds the prior mean value $\mathrm{E}[x] = \exp(\mu + (1/2)\sigma^2)$. The local sensitivity of the average estimate $m(x)$ to variations in x is given by

$$m'(x) \;=\; A\gamma x^{-(1-\gamma)} \;=\; \gamma(x/x^*)^{-(1-\gamma)}.$$

This is larger than 1 for all small enough values of x (in fact, it becomes unboundedly large as x approaches zero); but it decreases as x increases, and falls below 1 before x reaches the value x^*.

The strong excess sensitivity for small positive values of x reflects strong repulsion from the value $x \to 0$, at which the precision of encoding becomes extreme. As x grows, however, both the precision of encoding and the rate at which the precision of encoding is decreasing become smaller, so that the repulsion effect is weaker. For all large enough values of x, the regression effect dominates, resulting eventually in an attenuated response ($m'(x) < 1$) to further increases in x.