

Bayesian Decision Making

Study Guide

by Richard J. Young, Ph.D.

richard@deepneuro.ai

Chapter 11 of "The NeuroAI Handbook"

How a rational agent, brain or machine, should bet under uncertainty: hold beliefs as probabilities, update, and act.

One-sentence thesis. A rational agent under uncertainty holds its beliefs as probability distributions, updates them with Bayes' rule as evidence arrives, and acts to maximize expected utility, and a growing body of evidence says the brain does exactly this.

PART 1: Belief as Probability

How Should You Bet Under Uncertainty?

The world never hands you certainty. Sensory signals are corrupted by noise, the same sensation can spring from different causes, and you always know less than you would like. The **Bayesian brain hypothesis** proposes a strikingly clean response: represent what you know not as single guesses but as **probability distributions**, and reason with the calculus of probability.

The payoff of this stance is that it is *normative*, it says what an ideal observer *should* believe, giving us a yardstick to hold human and animal behavior against.

Where behavior matches the Bayesian prediction, we have found a computational principle; where it deviates, we have found something interesting to explain.

You see a shadowy figure in the dark. Person, tree, or trick of the light? Your brain cannot decide from the retinal image alone, so it fuses three things: the **current evidence** (what you see now), your **prior** (trees are common, intruders rare), and **context** (forest or parking garage). A wide distribution means "I am unsure"; a sharp one means "I am confident." Bayesian inference is the rule for combining these into the best possible bet.


PART 2: Updating With Bayes' Rule

Combining What You Knew With What You Just Saw

The engine of all this is one short equation for turning a belief and an observation into an updated belief:

$$P(H | D) = \frac{P(D | H) P(H)}{P(D)}$$

Read it as a recipe. The **prior** $P(H)$ is what you believed before. The **likelihood** $P(D | H)$ is how well each hypothesis predicts the data you saw. Multiply them, normalize by the total **evidence** $P(D)$, and you get the **posterior** $P(H | D)$, your belief after. Prior times likelihood, renormalized, is the whole of it.



The posterior (purple) sits between prior and likelihood but is sharper, more certain, than either. It leans toward the likelihood here because the incoming evidence is more precise than the prior belief.

Two behaviors matter. First, the update is **sequential and self-consistent**: today's posterior is tomorrow's prior, so beliefs converge on the truth as data accumulate, and the prior dominates only while data are scarce. **Conjugate priors** (a Beta prior with a Binomial likelihood yields a Beta posterior) make this updating exact and cheap.

Second, **ignore the base rate at your peril**. A disease affects 1% of people; a test is 95% sensitive and 90% specific. Test positive, and Bayes says your chance of actually having it is only about **8.8%**, because the rare disease's low prior drowns the test's accuracy. The posterior, not the likelihood, is what you should act on.

Two Noisy Senses Beat One


Reach for a coffee cup in daylight and your eyes and your hand both estimate its size. The estimates disagree slightly, both are noisy, so how should the brain fuse them? The Bayesian answer is not to average them equally but to **trust the more reliable one more**, weighting each estimate by its **precision**, the inverse of its variance.

$$\hat{x} = w_1x_1 + w_2x_2, \quad w_i = \frac{1/\sigma_i^2}{1/\sigma_1^2 + 1/\sigma_2^2}$$

The consequence is not just a compromise but an *improvement*: the fused estimate is sharper than either sense alone, because two independent noisy measurements carry more information than one. Ernst and Banks (2002) showed humans combine visual and haptic size cues almost exactly this way, down-weighting whichever sense the experimenter made less reliable, precisely as the equation predicts.

ANALOGY

Two witnesses to an accident, one with sharp eyesight, one who was farther away. A good detective does not average their accounts equally; she leans on the clearer witness while still using the other to narrow things down. Inverse-variance weighting is that instinct written as math.



Optimal cue integration. The combined estimate (purple) is taller, more precise, than either cue, and sits closer to touch because touch is the more reliable (lower-variance) sense.

A Posterior Is Not Yet a Decision


Knowing the state of the world is not the same as knowing what to do about it, that depends on what each outcome is worth. A **utility function** $U(s, a)$ scores the value of action a when the world is in state s . The rational choice maximizes **expected utility**, averaging utility over your posterior beliefs:

$$\mathbb{E}[U(a)] = \sum_s U(s, a) P(s | D)$$

This is why “wait” can beat “treat” even when disease is possible: if the posterior probability is modest and the cost of needless treatment high, waiting carries the higher expected utility. Belief and value must both enter the decision.

Over time, this looks mechanistic. The **drift-diffusion model** casts a choice as evidence *accumulating* stochastically until it hits a decision threshold, which reproduces both the accuracy and the reaction-time distributions of real choices. Neurons in areas **LIP** and **FEF** show exactly this ramping-to-threshold activity. How might a population represent a whole distribution? **Probabilistic population codes** encode the mean in where activity peaks and the uncertainty in how sharp that peak is, while **predictive coding** casts the cortex as a hierarchy sending predictions down and prediction errors up.

The brain only approximates this. Exact Bayesian inference is intractable for real problems, and people show base-rate neglect and overconfidence, so the brain likely runs *approximate* inference, sampling or otherwise, rather than the exact ideal.



The drift-diffusion model: noisy evidence accumulates until it crosses a threshold $\pm\theta$. Stronger evidence crosses sooner (faster, more accurate); the same mechanism sets both choice and reaction time.

Term	Symbol	Meaning	Conjugate pair
Prior	$P(H)$	Belief before data	Beta-Binomial (choices)
Likelihood	$P(D H)$	How data fits H	Gamma-Poisson (spikes)
Posterior	$P(H D)$	Belief after data	Normal-Normal (measurements)

Quick Self-Test

1. Why does the Bayesian brain hypothesis represent knowledge as distributions rather than single best guesses, and what does the width of a distribution mean?
2. Name the four terms in Bayes' rule and, in one line, what each contributes to the posterior.
3. A test is 95% sensitive and 90% specific for a disease with 1% prevalence. Why is a positive result only about 9% likely to be real?
4. In cue integration, why weight each sense by its precision, and why is the fused estimate more certain than either alone?
5. A posterior alone does not pick an action. What else is required, and what quantity does the rational agent maximize?

If You Remember One Thing

Rational behavior under uncertainty is three steps: hold beliefs as probabilities, update them with Bayes' rule (prior times likelihood, renormalized), and choose the action with the highest expected utility. The brain appears to run an approximate version of this loop, which is why the same framework explains a size judgment, a perceptual choice, and a drift-diffusion reaction time.

Richard J. Young, Ph.D. · richard@deepneuro.ai · Senior AI Research Scientist, UnitedHealth Group · UNLV Lee Business School · ORCID 0000-0002-1109-7552