# Bayesian Inference

IE 5545: Decision Analysis
Prof. Krishnamurthy Iyer*

May 2, 2026

A natural question that the Bayesian decision making framework raises is how does the decision maker arrive at her belief $\mu$ over the uncertain states. This can a surprisingly difficult question to address. While developing the framework, we stated that the decision maker "uses all the information at her disposal" to come up with the belief. While we will not try to model this process, we will instead seek to answer a simpler question: how does the (rational, Bayesian) decision maker incorporate new information in forming (or updating) her beliefs? In other words, if the decision maker already has formed a belief, and before making the decision, comes across a new piece of information, how should she learn from this new information in order to form new and updated beliefs? This learning process is known as Bayesian inference.

So, suppose we have a decision maker with some incomplete information, captured by the set of states $\Theta$. Let $\bar{\theta}$ denote the true (unknown) state, and the decision maker's belief about the state is captured by $\mu \in \Delta(\Theta)$:

$$\mu(\theta) = \mathbf{P}(\bar{\theta} = \theta).$$

The decision maker receives some new information about the true state $\bar{\theta}$, and we want to understand how her belief gets updated. The first question to understand is how do we represent "information about the true state $\bar{\theta}$".

## 1 Representing information

To represent information, we will use the notion of a *signal structure* (also known as information structure). We motivate this notion using the following example.

**Example** (Seller contd.). Consider the setting with a seller who is uncertain about whether their item will be a hit or a flop. We have $\Theta = \{\mathsf{hit}, \mathsf{flop}\}$ with the seller's belief given by

$$\mu(\mathsf{hit}) = \mathbf{P}(\bar{\theta} = \mathsf{hit}) = \frac{2}{3} \qquad\qquad \mu(\mathsf{flop}) = \mathbf{P}(\bar{\theta} = \mathsf{flop}) = \frac{1}{3}.$$

Suppose the seller seeks to gather more information about true state $\bar{\theta}$. One way to do this is to run a focus group, where the item is given to some users who later share their opinion about the item. In the simplest case, suppose the seller seeks a single user's opinion about whether they "liked" or "disliked" the item.

To learn from this user's opinion, we will need to know how the user forms their opinion about whether they like or dislike the item. For instance, if the user is such that they like all items

*©2026 Krishnamurthy Iyer

(regardless of whether the item is a hit or a flop), then knowing they like this item does not give us any information. At the other extreme, if the user is such that they only like items that will become a hit, then knowing that they like this item perfectly informs us that it will indeed be a hit.

We can represent the information from the user by stating (1) what different opinions they could have (i.e., $\{\mathsf{like}, \mathsf{dislike}\}$), and (2) how the true state affects their opinion, i.e., the probability $\mathbf{P}(\mathsf{like}|\mathsf{hit})$ with which they like an item if it is going to be a hit, and the probability $\mathbf{P}(\mathsf{like}|\mathsf{flop})$ with which they like an item if it is going to be a flop. Together, these two quantities constitute the *signal structure*, with the user's actual opinion referred to as the *signal*.

In general, we can represent some information about an unknown state $\bar{\theta} \in \Theta$ through the notion of a *signal structure*.

**Definition 1.** *A signal structure for a state $\bar{\theta} \in \Theta$ is defined as a pair $(\mathcal{X}, \mathbb{P})$ containing*

1. *the set $\mathcal{X} = \{x_1, \dots, x_L\}$ of possible values taken by a signal $\bar{x}$, and*
2. *the probabilities ("likelihoods") $\mathbb{P}(\bar{x} = x_\ell | \bar{\theta} = \theta_k)$ with which the signal takes a value $x_\ell \in \mathcal{X}$ in any state $\theta_k \in \Theta$.*

We now illustrate this concept with an extended example.

**Example** (Information about a dice roll)**.** Suppose we are uncertain about the roll of a fair dice. Perhaps we have bet on the outcome of this roll. Thus, the state is the outcome of the dice roll, and we have $\bar{\theta} \in \Theta = \{1, 2, 3, 4, 5, 6\}$, with $\mu$ putting equal weights $(1/6)$ on each $\theta \in \Theta$. The following lists different information we may obtain about the outcome of the dice roll:

1. Suppose the information tells us only whether the dice roll is even or odd. We can represent this by the signal structure with set of signals $\mathcal{X} = \{\mathsf{odd}, \mathsf{even}\}$, along with the likelihoods
$$\begin{aligned} &\mathbb{P}(\bar{x} = \mathsf{odd}|\bar{\theta} = 1) = 1, \qquad &&\mathbb{P}(\bar{x} = \mathsf{odd}|\bar{\theta} = 3) = 1, \qquad &&\mathbb{P}(\bar{x} = \mathsf{odd}|\bar{\theta} = 5) = 1 \\ &\mathbb{P}(\bar{x} = \mathsf{even}|\bar{\theta} = 2) = 1, \qquad &&\mathbb{P}(\bar{x} = \mathsf{even}|\bar{\theta} = 4) = 1, \qquad &&\mathbb{P}(\bar{x} = \mathsf{even}|\bar{\theta} = 6) = 1. \end{aligned}$$

2. Suppose the information tells us only whether the dice roll is greater than or equal to 3. This can be represented by the signal structure with $\mathcal{X} = \{\text{"}\geq\text{"}, \text{"}<\text{"}\}$, and likelihoods
$$\mathbb{P}(\bar{x} = \text{"}\geq\text{"}|\bar{\theta} = \theta) = \begin{cases} 1 & \text{if } \theta \in \{3, 4, 5, 6\}; \\ 0 & \text{if } \theta \in \{1, 2\}. \end{cases}$$

3. Suppose the information fully reveals the dice roll. This is represented by the signal structure with $\mathcal{X} = \{1, 2, 3, 4, 5, 6\}$ and likelihoods
$$\mathbb{P}(\bar{x} = x|\bar{\theta} = \theta) = 1, \quad \text{if and only if } x = \theta.$$

4. Suppose the information reveals nothing about the dice roll. This can be represented by the signal structure with $\mathcal{X} = \{\text{"}\mathsf{blah}\text{"}\}$ with $\mathbb{P}(\bar{x} = \text{"}\mathsf{blah}\text{"}|\bar{\theta} = \theta) = 1$ for all $\theta \in \Theta$.

5. Suppose a friend has seen the dice roll, and plans to share with us information about it in this way: the friend will first toss a *biased* coin (with the bias towards heads of 0.75), If the coin comes up heads, the friend will whistle if the dice roll is odd, and hum if its even. If instead the coin comes up tails, then the friend will whistle if the dice roll is prime ($\{2, 3, 5\}$) and hum otherwise. We only hear the whistle or the hum, and do not observe the coin toss.

This information from the friend can be represented by the signal structure $(\mathcal{X}, \mathbf{P})$, where $\mathcal{X} = \{\mathsf{whistle}, \mathsf{hum}\}$, with likelihoods

$$\begin{aligned}
\mathbb{P}(\bar{x} = \mathsf{whistle}|\bar{\theta} = 1) &= 0.75(1) + 0.25(0) = 0.75 \\
\mathbb{P}(\bar{x} = \mathsf{whistle}|\bar{\theta} = 2) &= 0.75(0) + 0.25(1) = 0.25 \\
\mathbb{P}(\bar{x} = \mathsf{whistle}|\bar{\theta} = 3) &= 0.75(1) + 0.25(1) = 1 \\
\mathbb{P}(\bar{x} = \mathsf{whistle}|\bar{\theta} = 4) &= 0.75(0) + 0.25(0) = 0 \\
\mathbb{P}(\bar{x} = \mathsf{whistle}|\bar{\theta} = 5) &= 0.75(1) + 0.25(1) = 1 \\
\mathbb{P}(\bar{x} = \mathsf{whistle}|\bar{\theta} = 6) &= 0.75(0) + 0.25(0) = 0.
\end{aligned}$$

## 2 Learning from information

Now that we have a way to represent information, we return to the question of how a (Bayesian) decision maker should update her beliefs when she gets some new information.

**Example** (Pricing (contd.))**.** Recall again the seller who is uncertain about whether their item will be a hit or a flop. We have $\Theta = \{\mathsf{hit}, \mathsf{flop}\}$ with the seller's belief given by

$$\mu(\mathsf{hit}) = \mathbf{P}(\bar{\theta} = \mathsf{hit}) = \frac{2}{3} \qquad\qquad \mu(\mathsf{flop}) = \mathbf{P}(\bar{\theta} = \mathsf{flop}) = \frac{1}{3}.$$

To get more information, the seller seeks a single user's opinion about whether they "liked" or "disliked" the item. Suppose this information can be represented by the following signal structure: $\mathcal{X} = \{\mathsf{like}, \mathsf{dislike}\}$ with likelihoods

$$\mathbb{P}(\bar{x} = \mathsf{like}|\bar{\theta} = \mathsf{hit}) = \frac{2}{3} \qquad\qquad \mathbb{P}(\bar{x} = \mathsf{like}|\bar{\theta} = \mathsf{flop}) = \frac{1}{5}.$$

In other words, the user is fairly attuned to the "pulse of the nation", and has a high chance of liking the item if it is going to be a hit, whereas if the item is going to be a flop, has a very low chance of liking it.

Given this signal structure, suppose the seller finds out that the user indeed likes the item, i.e., $\bar{x} = \mathsf{like}$. This information will lead the seller to form an updated belief about whether the item will be a hit or a flop. Denote this belief by $\mu_{\mathsf{new}} \in \Delta(\Theta)$. How does this new belief $\mu_{\mathsf{new}}$ relate to her old belief $\mu$?

To answer this question, we observe that $\mu_{\mathsf{new}}$ can be thought of as a conditional probability:

$$\mu_{\mathsf{new}}(\mathsf{hit}) = \mathbf{P}(\bar{\theta} = \mathsf{hit}|\bar{x} = \mathsf{like}).$$

From the definition of conditional probability, we have

$$\mathbf{P}(\bar{\theta} = \mathsf{hit}|\bar{x} = \mathsf{like}) = \frac{\mathbf{P}(\bar{\theta} = \mathsf{hit}, \bar{x} = \mathsf{like})}{\mathbf{P}(\bar{x} = \mathsf{like})} = \frac{\mathbf{P}(\bar{\theta} = \mathsf{hit}) \cdot \mathbf{P}(\bar{x} = \mathsf{like}|\bar{\theta} = \mathsf{hit})}{\mathbf{P}(\bar{x} = \mathsf{like})}.$$

(The last expression is an instance of the Bayes' rule.) The denominator can be found using the law of total probability:

$$\mathbf{P}(\bar{x} = \mathsf{like}) = \mathbf{P}(\bar{\theta} = \mathsf{hit}) \cdot \mathbf{P}(\bar{x} = \mathsf{like}|\bar{\theta} = \mathsf{hit}) + \mathbf{P}(\bar{\theta} = \mathsf{flop}) \cdot \mathbf{P}(\bar{x} = \mathsf{like}|\bar{\theta} = \mathsf{flop}).$$

Putting this all together, we get

$$\mathbf{P}(\bar{\theta} = \mathsf{hit}|\bar{x} = \mathsf{like}) = \frac{\mathbf{P}(\bar{\theta} = \mathsf{hit}) \cdot \mathbf{P}(\bar{x} = \mathsf{like}|\bar{\theta} = \mathsf{hit})}{\mathbf{P}(\bar{\theta} = \mathsf{hit}) \cdot \mathbf{P}(\bar{x} = \mathsf{like}|\bar{\theta} = \mathsf{hit}) + \mathbf{P}(\bar{\theta} = \mathsf{flop}) \cdot \mathbf{P}(\bar{x} = \mathsf{like}|\bar{\theta} = \mathsf{flop})}.$$

In the above expression, we know that $\mathbf{P}(\bar{\theta} = \mathsf{hit}) = \mu(\mathsf{hit})$, the old belief. Similarly, $\mathbf{P}(\bar{\theta} = \mathsf{flop}) = \mu(\mathsf{flop})$. Furthermore, we know $\mathbf{P}(\bar{x} = \mathsf{like}|\bar{\theta} = \mathsf{hit})$ and $\mathbf{P}(\bar{x} = \mathsf{like}|\bar{\theta} = \mathsf{flop})$ from the likelihoods. Thus, we obtain

$$\mu_{\mathsf{new}}(\mathsf{hit}) = \frac{\mu(\mathsf{hit}) \cdot \mathbb{P}(\bar{x} = \mathsf{like}|\bar{\theta} = \mathsf{hit})}{\mu(\mathsf{hit}) \cdot \mathbb{P}(\bar{x} = \mathsf{like}|\bar{\theta} = \mathsf{hit}) + \mu(\mathsf{flop}) \cdot \mathbb{P}(\bar{x} = \mathsf{like}|\bar{\theta} = \mathsf{flop})}. \tag{1}$$

The equation (1) relates the new belief $\mu_{\mathsf{new}}$ to the old belief $\mu$ using the likelihoods from the signal structure $\mathbb{P}(\bar{x} = x|\bar{\theta} = \theta)$. Thus, it captures how the seller should update the belief. Since this expression uses Bayes' rule, this is referred to as Bayesian belief update.

To complete the example, plugging the values of $\mu$ and $\mathbb{P}$, we obtain

$$\mu_{\mathsf{new}}(\mathsf{hit}) = \frac{(2/3) \cdot (2/3)}{(2/3) \cdot (2/3) + (1/3) \cdot (1/5)} = \frac{20}{23}.$$

Thus, after hearing that the user likes the item, the seller should now believe that there is a $(20/23)$ chance that the item will be a hit.

We can generalize the above discussion as follows. Recall the decision maker is uncertain about the state $\bar{\theta}$, which can take values in the set $\Theta$. The decision maker's belief (before receiving any new information) is given by $\mu \in \Delta(\Theta)$.

Suppose the information the decision maker receives is represented by the signal structure $(\mathcal{X}, \mathbb{P})$. After seeing the signal $\bar{x} = x_\ell \in \mathcal{X}$, the decision maker updates her belief to $\mu_{\mathsf{new}}$, which is obtained as the conditional distribution of $\bar{\theta}$, given $\bar{x} = x_\ell$. In other words, for any $\theta_k \in \Theta$, we have

$$\begin{aligned}
\mu_{\mathsf{new}}(\theta_k) &= \mathbf{P}(\bar{\theta} = \theta_k|\bar{x} = x_\ell) \\
&= \frac{\mathbf{P}(\bar{\theta} = \theta_k) \cdot \mathbf{P}(\bar{x} = x_\ell|\bar{\theta} = \theta_k)}{\mathbf{P}(\bar{x} = x_\ell)} && \text{(Bayes' rule)} \\
&= \frac{\mathbf{P}(\bar{\theta} = \theta_k) \cdot \mathbf{P}(\bar{x} = x_\ell|\bar{\theta} = \theta_k)}{\sum_{\theta \in \theta} \mathbf{P}(\bar{\theta} = \theta) \cdot \mathbf{P}(\bar{x} = x_\ell|\bar{\theta} = \theta)} && \text{(Law of total probability)} \\
&= \frac{\mu(\theta_k) \cdot \mathbb{P}(\bar{x} = x_\ell|\bar{\theta} = \theta_k)}{\sum_{\theta \in \theta} \mu(\theta) \cdot \mathbb{P}(\bar{x} = x_\ell|\bar{\theta} = \theta)}. && (2)
\end{aligned}$$

Here, we have used the Bayes' rule in the second equality and the law of total probability in the third equality. The final equality follows from the fact that $\mu(\theta) = \mathbf{P}(\bar{\theta} = \theta)$ for each $\theta \in \Theta$, and the definition of the likelihoods.

Since the denominator of (2) does not depend on $\theta_k$ and only depends on $x_\ell$, we can equivalently write the equation as[1]

$$\mu_{\mathsf{new}}(\theta_k) \propto \mu(\theta_k) \cdot \mathbb{P}(\bar{x} = x_\ell|\bar{\theta} = \theta_k).$$

This is the fundamental relationship of the Bayesian philosophy. The decision maker's belief $\mu$ before receiving the information in the signal $\bar{x}$ is referred to as her *prior belief*, whereas the

[1]Here $x_k \propto y_k$ means $x_k$ is proportional to $y_k$, i.e., $x_k = Cy_k$ for each $k$, for some $C \in \mathbb{R}$.

belief $\mu_{\mathsf{new}}$ after receiving the information $\bar{x} = x_\ell$ is referred to as her *posterior belief*. Thus, the fundamental relationship connects the decision maker's posterior belief to her prior belief using the likelihoods of the signal structure (and the realized signal $x_\ell$). This can be stated more succinctly as

$$\text{posterior belief} \propto (\text{prior belief}) \cdot (\text{likelihood})$$

**Remark 1.** *If either $\bar{\theta}$ or $\bar{x}$ is a continuous random variable taking values in an interval, then a similar expression relating the prior belief, posterior belief and the likelihoods holds, if we replace probabilities with densities for any continuous random variable.*

**Example** (Retail promotion (contd.))**.** Consider an online retailer who must decide whether to offer a user a discount on an item being sold on their website. The regular price for the item is given by $r = \$10$. If the discount is offered, the price is $d = \$8$. The discount has the effect of increasing the customer's probability of purchasing the item: if the customer purchases the item at a regular price with probability $\theta \in [0,1]$, then with the discount, she will purchase it with probability $H(\theta) = \frac{1+\theta}{2} \geq \theta$. The retailer wants to maximize her expected revenue $R$, where $R = r \cdot \theta$ if no discount is offered, and $R = d \cdot H(\theta)$ if a discount is offered. However, the retailer does not know the probability $\theta$ with which the user will purchase an item without a discount.

We have the set of states as $\Theta = [0,1]$ denoting all possible values of the probability with which any user will purchase the item when offered the regular price. Suppose the retailer's *prior belief $\mu$* is specified by the distribution with density $f(u) = 6u(1-u)$ for $u \in [0,1]$.

Now, suppose the retailer obtains some addition information about this specific user in the form of their purchase history. Specifically, in the past, the user was offered the regular price 5 times, of which the user bought the item 3 times. What should be the retailer's posterior belief about the user's purchase probability $\theta$?

To answer this, we first state the signal structure underlying the information provided. We know that user was offered the regular price $N = 5$ times. Of this, they could have bought the item anywhere from 0 times to 5 times. So, we can represent the set of signals as $\mathcal{X} = \{0,1,2,3,4,5\}$, where $\bar{x}$ denotes the number of times the user bought the item. What are the likelihoods $\mathbb{P}$? If we knew the state is $\bar{\theta} = \theta$, then the conditional distribution of $\bar{x}$ is just the $\mathsf{Binomial}(N,\theta)$ distribution, with

$$\mathbb{P}(\bar{x} = k | \bar{\theta} = \theta) = \binom{N}{k} \theta^k (1-\theta)^{N-k},$$

for each $k \in \mathcal{X}$.

Given this signal structure, we can now use the fundamental relationship (or the Bayes' rule) to state what the posterior belief of the retailer will be.

Let $f_{\mathsf{new}}$ denote the density of the posterior belief $\mu_{\mathsf{new}}$ after seeing $\bar{x} = k = 3$. We have using the fundamental relationship,

$$f_{\mathsf{new}}(\theta) \propto f(\theta) \cdot \mathbb{P}(\bar{x} = k | \bar{\theta} = \theta).$$

Letting $C$ denote the constant of proportionality (and substituting $N = 5$ and $k = 3$), we have

$$\begin{aligned} f_{\mathsf{new}}(\theta) &= C \cdot f(\theta) \cdot \mathbb{P}(\bar{x} = k | \bar{\theta} = \theta) \\ &= C \cdot 6\theta(1-\theta) \cdot \binom{5}{3} \theta^3 (1-\theta)^2 \\ &= 60C \cdot \theta^4 (1-\theta)^3. \end{aligned}$$

To find the value of $C$, we use the fact that since $f_{\mathsf{new}}$ is a density over $[0,1]$, we must have $\int_0^1 f_{\mathsf{new}}(\theta)d\theta = 1$. Solving this equation for $C$ and substuting it, we get the posterior belief's density as

$$f_{\mathsf{new}}(\theta) = 280 \cdot \theta^4(1-\theta)^3, \quad \text{for all } \theta \in [0,1].$$

What is the optimal decision of the retailer after receiving this information? We must evaluate the retailer's expected revenue, but using the posterior belief. The retailer's expected revenue is now given by

$$\begin{aligned}\Pi_{\mathsf{new}}(s) &= \mathbf{E}[\pi(\bar{\theta}, s) | \bar{\theta} \sim \mu_{\mathsf{new}}] \\ &= \begin{cases} r \cdot \mathbf{E}[\bar{\theta} | \bar{\theta} \sim \mu_{\mathsf{new}}] & \text{if } s = r; \\ d \cdot \mathbf{E}[H(\bar{\theta}) | \bar{\theta} \sim \mu_{\mathsf{new}}] = d \cdot \left(\frac{1+\mathbf{E}[\bar{\theta}|\bar{\theta} \sim \mu_{\mathsf{new}}]}{2}\right) & \text{if } s = d. \end{cases}\end{aligned}$$

It is straightforward to compute that $\mathbf{E}[\bar{\theta} | \bar{\theta} \sim \mu_{\mathsf{new}}] = \frac{5}{9}$. Thus, $\Pi_{\mathsf{new}}(r) = \frac{50}{9}$ and $\Pi_{\mathsf{new}}(d) = 8(1+\frac{5}{9})/2 = \frac{56}{9}$. Thus, while the gap between the payoff of the two prices has fallen down, it is still optimal for the retailer to offer a discount.