Markov Chain Monte Carlo

Last modified: 2026-11-12 22:02:32 (PST)

This page explains why most posterior distributions must be simulated, introduces Monte Carlo integration and Markov chains, describes the Metropolis–Hastings and Gibbs samplers, shows how to check a sampler's output, and presents a criterion for comparing models [1, Ch. 13]. It builds on the priors and posteriors defined on the Bayesian Inference page. For a video introduction, see Richard McElreath's lecture *MCMC and Item Response Models* (Statistical Rethinking 2026, Lecture A08).

Foundations of MCMC

When the posterior has a known closed form, as in the Beta-Bernoulli example, we can compute its summaries exactly or sample from it directly. In most real-world models, however, the marginal likelihood $\mathbb{p}(y)$ that normalizes the posterior cannot be computed, so neither can the posterior's summaries.~Markov chain Monte Carlo methods get around this problem: instead of independent draws from the posterior, they generate a correlated sequence of draws whose distribution converges to the posterior [1, Ch. 13], p. 287.

Definition 0.1 (Markov chain Monte Carlo)**.** A **Markov chain Monte Carlo (MCMC)** method for a target distribution generates a sequence of draws $\underset{\sim}{\theta}^{(1)}, \underset{\sim}{\theta}^{(2)}, \ldots$, each drawn from a distribution that depends on the previous draw, such that the distribution of $\theta^{(t)}$ converges to the target as $t \to \infty$.
Source: ~Wikipedia contributors [2].

Example 0.1 (A sequence that converges to the standard Gaussian)**.** Let $\theta^{(t+1)} = 0.5\,\theta^{(t)} + \varepsilon_t$, with $\varepsilon_t \sim_{\text{iid}} \mathrm{N}(0, 0.75)$ independent of the past. Each draw depends on the previous one. If $\theta^{(t)}$ has mean m_t and variance v_t, then $m_{t+1} = 0.5\,m_t$ and $v_{t+1} = 0.25\,v_t + 0.75$, so from any starting value $m_t \to 0$ and $v_t \to 1$: the draws' distribution converges to $\mathrm{N}(0, 1)$. This sequence is an MCMC method for sampling the standard Gaussian distribution; its draws are correlated, unlike independent draws.

Why the normalizing constant is hard to compute

With a parameter vector $\underset{\sim}{\theta}$ in the parameter space Θ, the normalizing constant of the posterior is

$$\mathbb{p}(\underset{\sim}{y}) = \int_{\Theta} \mathbb{p}(\underset{\sim}{y} \mid \underset{\sim}{\theta})\, \mathbb{p}(\underset{\sim}{\theta})\, d\underset{\sim}{\theta},$$

an integral over every dimension of $\underset{\sim}{\theta}$ that typically has no closed form. Computing it numerically becomes infeasible as the dimension of $\underset{\sim}{\theta}$ grows [1, Ch. 13], p. 287. MCMC sidesteps the problem: MCMC samplers use the posterior only through ratios $\mathbb{p}(\underset{\sim}{\theta}^* \mid \underset{\sim}{y})/\mathbb{p}(\underset{\sim}{\theta} \mid \underset{\sim}{y})$, in which the normalizing constant cancels.

Monte Carlo integration

Definition 0.2 (Monte Carlo estimate)**.** Given draws $\underset{\sim}{\theta}^{(1)}, \ldots, \underset{\sim}{\theta}^{(M)} \sim_{\text{iid}} \mathbb{p}(\underset{\sim}{\theta} \mid y)$ and a function g, the **Monte Carlo estimate** of the posterior expectation $\mathbb{E}[g(\underset{\sim}{\theta}) \mid \underset{\sim}{y}]$ is

$$\frac{1}{M} \sum_{m=1}^{M} g\big(\underset{\sim}{\theta}^{(m)}\big).$$

Source: Wikipedia contributors [3].

By the law of large numbers, a Monte Carlo estimate converges to the posterior expectation it estimates as the number of draws M grows [1, Ch. 13], p. 288. The posterior mean ($g(\theta) = \theta$) and posterior probabilities ($g(\theta) = 1_{\theta \le c}$, an indicator) are posterior expectations, so they are estimated by averages of the draws; posterior quantiles, and so credible-interval endpoints, are estimated by the corresponding quantiles of the draws.

Example 0.2 (Monte Carlo posterior summaries for a Bernoulli model)**.** The posterior in the Beta-Bernoulli example is $\text{Beta}(56, 37)$. We draw $M = 5{,}000$ values from it, and compare Monte Carlo estimates with the exact values:

```r
set.seed(42)
pi_draws <- stats::rbeta(n = 5000, shape1 = 56, shape2 = 37)
rbind(
  exact = c(
    mean = 56 / (56 + 37),
    lower = stats::qbeta(0.025, 56, 37),
    upper = stats::qbeta(0.975, 56, 37)
  ),
  monte_carlo = c(
    mean(pi_draws),
    stats::quantile(pi_draws, c(0.025, 0.975))
  )
) |>
  round(4)
#>                mean  lower  upper
#> exact       0.6022 0.5014 0.6988
#> monte_carlo 0.6022 0.5011 0.7013
```

The two agree to about two decimal places, and the agreement improves as M grows.

Markov chains

Definition 0.3 (Markov chain)**.** A sequence of random variables $\underset{\sim}{\theta}^{(1)}, \underset{\sim}{\theta}^{(2)}, \ldots$ is a **Markov chain** if, for every t, the conditional distribution of $\underset{\sim}{\theta}^{(t+1)}$ given all the earlier values depends only on $\underset{\sim}{\theta}^{(t)}$:

$$\mathbb{p}\left(\underset{\sim}{\theta}^{(t+1)} \mid \underset{\sim}{\theta}^{(t)}, \underset{\sim}{\theta}^{(t-1)}, \ldots, \underset{\sim}{\theta}^{(1)}\right) = \mathbb{p}\left(\underset{\sim}{\theta}^{(t+1)} \mid \underset{\sim}{\theta}^{(t)}\right).$$

Source: Wikipedia contributors [4].

Definition 0.4 (Stationary distribution)**.** A distribution π is a **stationary distribution** of a Markov chain if, whenever $\underset{\sim}{\theta}^{(t)}$ has distribution π, $\underset{\sim}{\theta}^{(t+1)}$ also has distribution π.
Source: Wikipedia contributors [5].

Example 0.3 (A two-state weather chain)**.** Each day is dry or wet. A dry day is followed by a dry day with probability 0.9, and a wet day is followed by a dry day with probability 0.5, whatever the weather on earlier days. So the daily weather is a Markov chain.
Let π_{dry} be the probability that a day is dry. For that probability to be the same the next day,

$$\begin{aligned} \pi_{\text{dry}} &= 0.9\,\pi_{\text{dry}} + 0.5\,(1 - \pi_{\text{dry}}) && \text{(tomorrow is dry after a dry or a wet day)} \\ &= 0.5 + 0.4\,\pi_{\text{dry}} && \text{(collecting terms)}, \end{aligned}$$

so $0.6\,\pi_{\text{dry}} = 0.5$ and $\pi_{\text{dry}} = 5/6$. The stationary distribution is dry with probability $5/6$ and wet with probability $1/6$. A simulated chain, started on a wet day, spends about that fraction of its days dry:

```
set.seed(7)
n_days <- 10000
dry <- logical(n_days)
for (t in seq(2, n_days)) {
  p_dry <- if (dry[t - 1]) 0.9 else 0.5
  dry[t] <- stats::runif(1) < p_dry
}
c(simulated = mean(dry), stationary = 5 / 6) |> round(3)
#>  simulated stationary
#>      0.833      0.833
```

Example 0.4 (A first-order autoregressive chain)**.** Let ρ and $\sigma > 0$ be constants, let $\varepsilon^{(2)}, \varepsilon^{(3)}, \ldots \sim_{\text{iid}} \mathrm{N}(0, \sigma^2)$, and let the starting value $\theta^{(1)}$ be independent of them. For each $t \geq 1$, let

$$\theta^{(t+1)} \stackrel{\text{def}}{=} \rho\,\theta^{(t)} + \varepsilon^{(t+1)}. \tag{1}$$

The sequence $\theta^{(1)}, \theta^{(2)}, \ldots$ is a **first-order autoregressive** process. Each of $\theta^{(1)}, \ldots, \theta^{(t)}$ is a function of $\theta^{(1)}, \varepsilon^{(2)}, \ldots, \varepsilon^{(t)}$, so $\varepsilon^{(t+1)}$ is independent of all of them. So, given $\theta^{(t)} = s$ and any earlier values, $\theta^{(t+1)} = \rho\,s + \varepsilon^{(t+1)}$ has the $\mathrm{N}(\rho\,s, \sigma^2)$ distribution, which depends only on the current state s. The sequence is therefore a Markov chain. Unlike the chain of Example 0.3, its states can be any real number, so it has a continuous state space, like the chains that MCMC builds for a continuous parameter, whose values lie in the parameter space.

Exercise 0.1 (Stationary distribution of the autoregressive chain)**.** For the chain of Example 0.4 with $|\rho| < 1$, find the variance $v > 0$ for which $\mathrm{N}(0, v)$ is a stationary distribution.

Solution

Solution 0.1. Suppose $\theta^{(t)} \sim \mathrm{N}(0, v)$. By Example 0.4, $\varepsilon^{(t+1)}$ is independent of $\theta^{(t)}$, so $\theta^{(t+1)} = \rho\,\theta^{(t)} + \varepsilon^{(t+1)}$, a linear combination of independent Gaussian random variables, is Gaussian [6], Corollary 4.6.10. By linearity of expectation, its mean is

$$
\begin{aligned}
\mathbb{E}\left[\theta^{(t+1)}\right] &= \mathbb{E}\left[\rho\,\theta^{(t)} + \varepsilon^{(t+1)}\right] && \text{(definition of the chain)} \\
&= \rho\,\mathbb{E}\left[\theta^{(t)}\right] + \mathbb{E}\left[\varepsilon^{(t+1)}\right] && \text{(linearity of expectation)} \\
&= \rho \cdot 0 + 0 && \text{(substituting the two means)} \\
&= 0 && \text{(arithmetic).}
\end{aligned} \tag{2}
$$

Independent random variables are uncorrelated, so $\operatorname{Cov}\left(\theta^{(t)}, \varepsilon^{(t+1)}\right) = 0$, and by the variance of a sum of two random variables, the variance is

$$
\begin{aligned}
\operatorname{Var}\left(\theta^{(t+1)}\right) &= \operatorname{Var}\left(\rho\,\theta^{(t)} + 1 \cdot \varepsilon^{(t+1)}\right) && \text{(definition of the chain)} \\
&= \rho^2 \operatorname{Var}\left(\theta^{(t)}\right) + 1^2 \operatorname{Var}\left(\varepsilon^{(t+1)}\right) + 2(\rho \cdot 1)\operatorname{Cov}\left(\theta^{(t)}, \varepsilon^{(t+1)}\right) && \text{(variance of a sum of two random varia} \\
&= \rho^2 \operatorname{Var}\left(\theta^{(t)}\right) + 1^2 \operatorname{Var}\left(\varepsilon^{(t+1)}\right) + 2(\rho \cdot 1) \cdot 0 && \text{(the covariance is 0)} \\
&= \rho^2 \operatorname{Var}\left(\theta^{(t)}\right) + 1^2 \operatorname{Var}\left(\varepsilon^{(t+1)}\right) + 0 && (a \cdot 0 = 0) \\
&= \rho^2 \operatorname{Var}\left(\theta^{(t)}\right) + 1^2 \operatorname{Var}\left(\varepsilon^{(t+1)}\right) && (a + 0 = a) \\
&= \rho^2 \operatorname{Var}\left(\theta^{(t)}\right) + \operatorname{Var}\left(\varepsilon^{(t+1)}\right) && (1^2 = 1) \\
&= \rho^2 v + \operatorname{Var}\left(\varepsilon^{(t+1)}\right) && \text{(substituting } \operatorname{Var}\left(\theta^{(t)}\right) = v\text{)} \\
&= \rho^2 v + \sigma^2 && \text{(substituting } \operatorname{Var}\left(\varepsilon^{(t+1)}\right) = \sigma^2\text{).}
\end{aligned} \tag{3}
$$

So $\theta^{(t+1)} \sim \mathrm{N}(0, \rho^2 v + \sigma^2)$, and $\mathrm{N}(0, v)$ is stationary exactly when $\rho^2 v + \sigma^2 = v$. Solving for v:

$$
\begin{aligned}
v &= \rho^2 v + \sigma^2 && \text{(stationarity condition)} \\
v - \rho^2 v &= \sigma^2 && \text{(subtracting } \rho^2 v \text{ from both sides)} \\
(1 - \rho^2)\, v &= \sigma^2 && \text{(factoring out } v\text{)} \\
v &= \frac{\sigma^2}{1 - \rho^2} && \text{(dividing both sides by } 1 - \rho^2\text{, which is positive because } |\rho| < 1\text{).}
\end{aligned} \tag{4}
$$

The stationary distribution is $\mathrm{N}(0, \sigma^2/(1-\rho^2))$. If instead $|\rho| \geq 1$, then $(1-\rho^2)\, v \leq 0 < \sigma^2$ for every $v > 0$, so no $\mathrm{N}(0, v)$ distribution is stationary.
With $\rho = 0.9$ and $\sigma = 1$, the stationary variance is

$$
\begin{aligned}
1/(1 - 0.9^2) &= 1/(1 - 0.81) \\
&= 1/0.19 \\
&\approx 5.26.
\end{aligned}
$$

A simulated chain started far from 0, at $\theta^{(1)} = 10$, has mean near 0 and variance near that value once its first 1,000 values are discarded:

```r
set.seed(11)
n_iter <- 100000
rho <- 0.9
theta <- numeric(n_iter)
theta[1] <- 10
for (t in seq(2, n_iter)) {
  theta[t] <- rho * theta[t - 1] + stats::rnorm(1)
}
kept <- theta[-(1:1000)]
c(
  mean = mean(kept),
  variance = stats::var(kept),
  stationary_variance = 1 / (1 - rho^2)
) |>
  round(2)
#>                mean            variance stationary_variance
#>                0.05                5.07                5.26
```

Source

That a linear combination of independent normal random variables is normal is Casella and Berger [6], Corollary 4.6.10.

Theorem 0.1 (Stationary distribution of the autoregressive chain)**.** *For the chain of Example 0.4 with $|\rho| < 1$, the normal distribution with mean 0 and variance $\sigma^2/(1-\rho^2)$ is a stationary distribution:*

$$\mathrm{N}\left(0, \frac{\sigma^2}{1-\rho^2}\right). \tag{5}$$

Proof

Proof. This result is the solution to Exercise 0.1. □

Under conditions on its transition probabilities (irreducibility, aperiodicity and positive recurrence; for a continuous parameter, Harris positive recurrence), a Markov chain has a unique stationary distribution and the distribution of $\underset{\sim}{\theta}^{(t)}$ converges to it [7, Sec. 11.2], p. 279, [8, Sec. 6.6.1], Theorem 6.51, p. 234. Sample means along the chain also converge to expectations under that distribution, even though successive values are correlated [8, Sec. 6.7.1], Theorem 6.63, p. 241, [8, Sec. 7.2], p. 269. MCMC algorithms construct a Markov chain whose stationary distribution is the posterior $\mathbb{p}(\underset{\sim}{\theta} \mid \underset{\sim}{y})$ [7, Ch. 11], p. 275, so that Monte Carlo estimates

can be computed from the chain's values in place of independent draws.

MCMC samplers

The Metropolis–Hastings algorithm

Definition 0.1 (Proposal distribution)**.** In a Metropolis–Hastings sampler, a **proposal distribution** $q(\cdot \mid \underset{\sim}{\theta})$ is a conditional distribution from which a candidate value $\underset{\sim}{\theta}^*$ is drawn, given the current value $\underset{\sim}{\theta}$.

Definition 0.2 (Symmetric proposal)**.** A proposal distribution is **symmetric** if $q(\underset{\sim}{\theta}^* \mid \underset{\sim}{\theta}) = q(\underset{\sim}{\theta} \mid \underset{\sim}{\theta}^*)$ for every $\underset{\sim}{\theta}$ and $\underset{\sim}{\theta}^*$.

Example 0.1 (A Gaussian random-walk proposal)**.** For a scalar parameter π, the proposal $\pi^* \mid \pi \sim \mathrm{N}(\pi, 0.05^2)$ draws a candidate within about ± 0.1 of the current value π. It is symmetric: $q(\pi^* \mid \pi) = q(\pi \mid \pi^*)$, because the Gaussian density depends on $\pi^* - \pi$ only through its square.

Definition 0.3 (Acceptance ratio)**.** For a target posterior $\mathbb{p}(\underset{\sim}{\theta} \mid y) \propto \mathbb{p}(y \mid \underset{\sim}{\theta})\,\mathbb{p}(\underset{\sim}{\theta})$, a proposal distribution q, a current value $\underset{\sim}{\theta}^{(t)}$, and a candidate $\underset{\sim}{\theta}^*$, the **acceptance ratio** is

$$\alpha \overset{\text{def}}{=} \frac{\mathbb{p}(y \mid \underset{\sim}{\theta}^*)\,\mathbb{p}(\underset{\sim}{\theta}^*)}{\mathbb{p}(y \mid \underset{\sim}{\theta}^{(t)})\,\mathbb{p}(\underset{\sim}{\theta}^{(t)})} \cdot \frac{q(\underset{\sim}{\theta}^{(t)} \mid \underset{\sim}{\theta}^*)}{q(\underset{\sim}{\theta}^* \mid \underset{\sim}{\theta}^{(t)})}.$$

Example 0.2 (Acceptance ratio for a Bernoulli posterior)**.** With $r = 55$ successes in $n = 91$ trials, a uniform prior, and the symmetric proposal of Example 0.1, the q terms cancel and the prior is constant, so moving from $\pi^{(t)} = 0.5$ to $\pi^* = 0.6$ gives

$$\begin{aligned} \alpha &= \frac{0.6^{55}\,0.4^{36}}{0.5^{55}\,0.5^{36}} && \text{(Bernoulli likelihood ratio)} \\ &= 1.2^{55}\,0.8^{36} && \text{(combine the powers)} \\ &\approx 7.35 && \text{(arithmetic)} \end{aligned}$$

The reverse move, from 0.6 to 0.5, has $\alpha \approx 1/7.35 \approx 0.136$.

Definition 0.4 (Metropolis–Hastings algorithm)**.** Given a target posterior $\mathbb{p}(\underset{\sim}{\theta} \mid \underset{\sim}{y})$, a proposal distribution $q(\cdot \mid \underset{\sim}{\theta})$, and a starting value $\underset{\sim}{\theta}^{(1)}$, the **Metropolis–Hastings algorithm** produces $\underset{\sim}{\theta}^{(t+1)}$ from $\underset{\sim}{\theta}^{(t)}$ as follows [1, Ch. 13], p. 291:

1. Draw a candidate $\underset{\sim}{\theta}^{*} \sim q(\cdot \mid \underset{\sim}{\theta}^{(t)})$.
2. Compute the acceptance ratio α.
3. With probability $\min(1, \alpha)$, set $\underset{\sim}{\theta}^{(t+1)} = \underset{\sim}{\theta}^{*}$ (accept the candidate); otherwise set $\underset{\sim}{\theta}^{(t+1)} = \underset{\sim}{\theta}^{(t)}$ (reject it).

Source: Wikipedia contributors [9].

The first factor of α is the ratio $\mathbb{p}(\underset{\sim}{\theta}^{*} \mid \underset{\sim}{y})/\mathbb{p}(\underset{\sim}{\theta}^{(t)} \mid \underset{\sim}{y})$ with the normalizing constant $\mathbb{p}(\underset{\sim}{y})$ canceled (the posterior equation), so the algorithm needs only the unnormalized posterior. When the proposal is symmetric, as for a random walk $\underset{\sim}{\theta}^{*} = \underset{\sim}{\theta}^{(t)} + \varepsilon$ with ε drawn from a distribution symmetric about 0, the second factor equals 1. Provided the proposal can reach every region where the posterior is positive, the resulting chain is a Markov chain whose stationary distribution is the posterior [7, Sec. 11.2], pp. 279-280, [8, Sec. 7.3.1], Theorem 7.2, p. 272.

Example 0.3 (Metropolis–Hastings for a Bernoulli model)**.** We sample the posterior of the Beta-Bernoulli example ($r = 55$ successes in $n = 91$ trials, uniform prior) with a Gaussian random-walk proposal $\pi^{*} = \pi^{(t)} + \varepsilon$, $\varepsilon \sim \mathrm{N}(0, 0.05^2)$. The proposal is symmetric, and the uniform prior density is 1 on $(0, 1)$, so the log of the acceptance ratio is the difference of log-likelihoods, and a proposal outside $(0, 1)$ is always rejected.

```r
log_post_unnorm <- function(pi, r, n) {
  if (pi <= 0 || pi >= 1) {
    return(-Inf)
  }
  r * log(pi) + (n - r) * log(1 - pi)
}

mh_bernoulli <- function(n_iter, r, n, pi_init, step_sd = 0.05) {
  chain <- numeric(n_iter)
  chain[1] <- pi_init
  for (t in seq(2, n_iter)) {
    pi_curr <- chain[t - 1]
    pi_prop <- pi_curr + stats::rnorm(1, sd = step_sd)
    log_alpha <- log_post_unnorm(pi_prop, r, n) -
      log_post_unnorm(pi_curr, r, n)
    accept <- log(stats::runif(1)) < log_alpha
    chain[t] <- if (accept) pi_prop else pi_curr
  }
  chain
}

set.seed(1)
mh_chain <- mh_bernoulli(n_iter = 5000, r = 55, n = 91, pi_init = 0.5)
c(
  acceptance_rate = mean(diff(mh_chain) != 0),
  mh_mean = mean(mh_chain),
  exact_mean = 56 / 93
) |>
  round(3)
#> acceptance_rate         mh_mean      exact_mean
#>           0.710           0.601           0.602
```

Comparing $\log u$, for u uniform on $(0, 1)$, with $\log \alpha$ accepts with probability $\min(1, \alpha)$, and avoids computing α itself, which can overflow or underflow. Every accepted proposal changes the chain's value, so the fraction of iterations that change estimates the acceptance rate. The chain's sample mean is close to the exact posterior mean $56/93$.

Trace plots and burn-in

Definition 0.5 (Trace plot)**.** A **trace plot** of an MCMC chain plots each sampled value of a parameter against its iteration number.

Definition 0.6 (Burn-in)**.** The **burn-in** period is the initial segment of an MCMC chain that is discarded before posterior summaries are computed [1, Ch. 13], p. 301. For a chain $\theta^{(1)}, \ldots, \theta^{(T)}$ and a burn-in of B draws, the summaries use only the remaining draws $\underset{\sim}{\theta}^{(B+1)}, \ldots, \underset{\sim}{\theta}^{(T)}$.

A chain started far from where the posterior puts its probability takes some iterations to get there, and until it does, its values are not draws from the posterior. Burn-in removes those iterations. Its length is chosen by inspecting trace plots of chains started from different values: the burn-in should end after the chains have stopped drifting and are moving around the same region.

Example 0.4 (Burn-in for the Bernoulli sampler)**.** Continuing Example 0.3, we start two more chains at the extreme values 0.05 and 0.95. Figure 1 shows their first 200 iterations.

```r
set.seed(2)
chain_low <- mh_bernoulli(n_iter = 5000, r = 55, n = 91, pi_init = 0.05)
chain_high <- mh_bernoulli(n_iter = 5000, r = 55, n = 91, pi_init = 0.95)
shown <- 1:200
traces <- tibble::tibble(
  iteration = rep(shown, times = 2),
  pi = c(chain_low[shown], chain_high[shown]),
  start = rep(c("0.05", "0.95"), each = length(shown))
)
ggplot2::ggplot(traces) +
  ggplot2::aes(x = iteration, y = pi, colour = start) +
  ggplot2::geom_line() +
  ggplot2::labs(
    y = expression("Success probability, " * pi),
    colour = "starting value"
  )
```

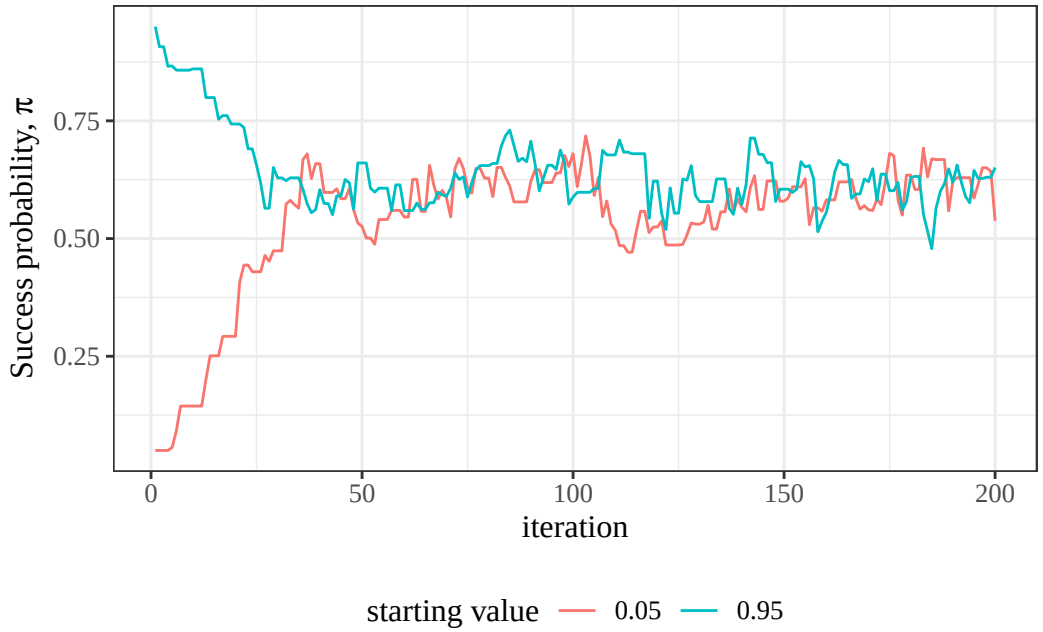


Figure 1: Trace plots of the first 200 iterations of two Metropolis–Hastings chains for the Bernoulli model, started at $\pi = 0.05$ and $\pi = 0.95$.

Both chains reach the region around 0.6 within a few dozen iterations, and after that the two are indistinguishable, so discarding the first 500 iterations of each leaves a generous

margin.

The Gibbs sampler

Definition 0.7 (Full conditional distribution)**.** For a parameter vector $\underset{\sim}{\theta} = (\theta_1, \dots, \theta_K)$, the **full conditional distribution** of the component θ_k is its conditional distribution given the data and all the other components, $\mathbb{p}(\theta_k \mid \underset{\sim}{\theta}_{-k}, \underset{\sim}{y})$, where $\underset{\sim}{\theta}_{-k}$ denotes every component of $\underset{\sim}{\theta}$ except θ_k.

Definition 0.8 (Gibbs sampler)**.** The **Gibbs sampler** produces $\underset{\sim}{\theta}^{(t+1)}$ from $\underset{\sim}{\theta}^{(t)}$ by updating one component at a time: for $k = 1, \dots, K$ in turn, it draws $\theta_k^{(t+1)}$ from the full conditional distribution of θ_k, with each other component set to its most recent value [1, Ch. 13], p. 293.
Source: Wikipedia contributors [10].

The Gibbs sampler is a special case of the Metropolis–Hastings algorithm, one component at a time, whose proposal is the full conditional itself; with that proposal the acceptance ratio is always 1, so every draw is accepted [1, Ch. 13], p. 293. It needs a way to draw from each full conditional. JAGS ("Just Another Gibbs Sampler") builds a Gibbs sampler from a model's description, choosing a sampling method for each full conditional [1, Ch. 13].

Example 0.5 (Gibbs sampling from a correlated Gaussian target)**.** Let the target be the bivariate Gaussian distribution of (θ_1, θ_2) with means 0, variances 1 and correlation ρ, whose density is proportional to $\exp\left\{-\frac{\theta_1^2 - 2\rho\theta_1\theta_2 + \theta_2^2}{2(1-\rho^2)}\right\}$ [6], Definition 4.5.10. As a function of θ_1,

$$\begin{aligned}
\mathbb{p}(\theta_1 \mid \theta_2) &\propto \exp\left\{-\frac{\theta_1^2 - 2\rho\theta_1\theta_2}{2(1-\rho^2)}\right\} && \text{(dropping the factor with } \theta_2^2 \text{ only)} \\
&= \exp\left\{-\frac{(\theta_1 - \rho\theta_2)^2 - \rho^2\theta_2^2}{2(1-\rho^2)}\right\} && \text{(completing the square)} \\
&\propto \exp\left\{-\frac{(\theta_1 - \rho\theta_2)^2}{2(1-\rho^2)}\right\} && \text{(dropping the factor with } \theta_2^2 \text{ only),}
\end{aligned}$$

so the full conditional of θ_1 is $\mathrm{N}(\rho\theta_2,\, 1-\rho^2)$, and by symmetry that of θ_2 is $\mathrm{N}(\rho\theta_1,\, 1-\rho^2)$. We run the Gibbs sampler with $\rho = 0$ and with $\rho = 0.99$:

```r
gibbs_bvn <- function(n_iter, rho) {
  draws <- matrix(NA_real_, nrow = n_iter, ncol = 2)
  theta <- c(0, 0)
  cond_sd <- sqrt(1 - rho^2)
  for (t in seq_len(n_iter)) {
    theta[1] <- stats::rnorm(1, mean = rho * theta[2], sd = cond_sd)
    theta[2] <- stats::rnorm(1, mean = rho * theta[1], sd = cond_sd)
    draws[t, ] <- theta
  }
  draws
}

set.seed(3)
draws_indep <- gibbs_bvn(n_iter = 2000, rho = 0)
draws_corr <- gibbs_bvn(n_iter = 2000, rho = 0.99)
lag1_autocorrelation <- function(x) stats::cor(x[-1], x[-length(x)])
c(
  rho_0 = lag1_autocorrelation(draws_indep[, 1]),
  rho_0.99 = lag1_autocorrelation(draws_corr[, 1])
) |>
  round(3)
#>    rho_0 rho_0.99
#>    0.020    0.971
```

With $\rho = 0$, successive draws of θ_1 are nearly uncorrelated. With $\rho = 0.99$, each full conditional has standard deviation $\sqrt{1 - 0.99^2} \approx 0.14$, so each update moves only a short way along the narrow ridge where the target puts its probability, and successive draws are almost perfectly correlated: 2,000 such draws carry far less information about the target than 2,000 independent ones.

Checking and improving MCMC

Because MCMC draws are correlated, and a chain may take many iterations to reach its stationary distribution, we must check whether the chains have converged before using them for inference [1, Ch. 13], p. 302. Trace plots give a visual check; the potential scale reduction factor gives a numerical one, by comparing several chains started from different values.

Definition 0.1 (Potential scale reduction factor)**.** Suppose m chains each contribute n draws of a parameter θ after burn-in, with chain sample means $\bar{\theta}_1, \ldots, \bar{\theta}_m$, overall sample

mean $\bar{\theta}$, and within-chain sample variances $s_1^2, \ldots, s_m^2$. Let

$$W \stackrel{\text{def}}{=} \frac{1}{m}\sum_{j=1}^{m} s_j^2 \qquad \text{(within-chain variance)}$$

$$B \stackrel{\text{def}}{=} \frac{n}{m-1}\sum_{j=1}^{m}(\bar{\theta}_j - \bar{\theta})^2 \qquad \text{(between-chain variance)}$$

$$\hat{V} \stackrel{\text{def}}{=} \frac{n-1}{n}W + \frac{1}{n}B \qquad \text{(pooled variance estimate).}$$

The **Gelman–Rubin potential scale reduction factor** is

$$\hat{R} \stackrel{\text{def}}{=} \sqrt{\hat{V}/W}.$$

If the chains have all converged to the posterior, $\hat{V}$ estimates the posterior variance and W approaches it as n grows, so $\hat{R}$ is close to 1. If the chains are still exploring different regions, the chain means differ, B is large, and $\hat{R}$ exceeds 1 [7, Sec. 11.4], pp. 283-285. Definition 0.1 follows Gelman et al. [7, Sec. 11.4], pp. 284-285, except that they first split each chain in half and compute $\hat{R}$ over the half-chains, so that $\hat{R}$ also checks each chain for stationarity [7, Sec. 11.4], p. 285, footnote 2. The statistic $\hat{R}$ is also called the **Gelman–Rubin statistic**; it formally assesses whether several chains have converged [1, Ch. 13], p. 306. Software reports $\hat{R}$ as `psrf` or `Rhat`; `coda::gelman.diag()` also applies a small-sample correction, so its value differs slightly from Definition 0.1's. A value of $\hat{R}$ near 1 does not prove convergence: chains that are all stuck in the same wrong region also agree.

Example 0.1 (Potential scale reduction factor for the Bernoulli sampler)**.**
Continuing Example 0.4, we compute $\hat{R}$ for the two chains, first over their first 50 iterations, which include the burn-in, and then over iterations 501 to 5,000:

```r
psrf <- function(chains) {
  n <- nrow(chains)
  w <- mean(apply(chains, 2, stats::var))
  b <- n * stats::var(colMeans(chains))
  v_hat <- (n - 1) / n * w + b / n
  sqrt(v_hat / w)
}
chains <- cbind(chain_low, chain_high)
c(
  first_50 = psrf(chains[1:50, ]),
  after_burnin = psrf(chains[501:5000, ])
) |>
  round(3)
#>     first_50 after_burnin
#>        1.648        1.001
```

Here `stats::var(colMeans(chains))` is $\frac{1}{m-1}\sum_{j=1}^{m}(\bar{\theta}_j-\bar{\theta})^2$, so `b` is B. While the chains are still approaching the posterior from opposite ends, $\hat{R}$ is well above 1; after burn-in it is essentially 1.

Comparing MCMC estimates to maximum likelihood

When the prior is weak and the sample is moderate or large, the posterior mean from a well-mixed chain and the maximum likelihood estimate typically agree closely, and the posterior standard deviation is close to the frequentist standard error [1, Ch. 13], p. 302. As n grows the likelihood dominates the prior, so the posterior concentrates near the maximum likelihood estimate. Comparing the two is therefore a useful check on an MCMC analysis.

Example 0.2 (Posterior mean and maximum likelihood estimate for a Bernoulli model)**.** With $r = 55$ successes in $n = 91$ trials, the maximum likelihood estimate is $\hat{\pi} = r/n$, with estimated standard error $\sqrt{\hat{\pi}(1-\hat{\pi})/n}$. Under the uniform prior, the posterior is $\text{Beta}(r+1, n-r+1)$ (the Beta-Bernoulli example), whose mean is $(r+1)/(n+2)$ and whose standard deviation is $\sqrt{ab/((a+b)^2(a+b+1))}$ with $a = r+1$ and $b = n-r+1$ [6, Sec. 3.3], p. 107:

```r
r <- 55
n <- 91
pi_hat <- r / n
a <- r + 1
b <- n - r + 1
rbind(
  maximum_likelihood = c(
    estimate = pi_hat,
    spread = sqrt(pi_hat * (1 - pi_hat) / n)
  ),
  posterior = c(a / (a + b), sqrt(a * b / ((a + b)^2 * (a + b + 1))))
) |>
  round(4)
#>                    estimate spread
#> maximum_likelihood   0.6044 0.0513
#> posterior            0.6022 0.0505
```

The two estimates differ only through the prior's pseudo-counts, and so do the two measures of spread; both differences vanish as n grows. The MCMC mean in Example 0.3 agrees with the exact posterior mean, and so with the maximum likelihood estimate to about two decimal places.

A persistent discrepancy between the two is a signal worth investigating: it may reflect a genuinely informative prior, an unconverged chain, or a coding error in the model.

The importance of parameterization

How a model is written affects how well its sampler mixes [1, Ch. 13], p. 299. Two algebraically equivalent parameterizations of the same model can produce chains with very different autocorrelation. Strong posterior correlation between parameters slows a sampler that updates one component at a time, because each update can move only a short way along a narrow, tilted ridge, as Example 0.5 shows with $\rho = 0.99$.

Common remedies include *centering* predictors (subtracting their means), so that the intercept and slopes are less correlated, and *reparameterizing* variance components on a scale on which the posterior is more nearly symmetric. These changes leave the model, and so the scientific conclusions, unchanged; they alter only the geometry the sampler must explore.

Deviance information criterion

To compare models fit to the same data by Bayesian inference, we need a measure that balances goodness of fit against complexity, as Akaike's information criterion does for models fit by maximum likelihood. The standard criterion computed from MCMC output is the deviance information criterion [1, Ch. 13], p. 306.

Definition 0.1 (Deviance of a parameter value)**.** In the deviance information criterion, the **deviance** of a parameter value $\underset{\sim}{\theta}$ is

$$\operatorname{dev}_{\mathrm{DIC}}(\underset{\sim}{\theta}) \stackrel{\mathrm{def}}{=} -2 \log \mathbb{p}(\underset{\sim}{y} \mid \underset{\sim}{\theta}).$$

Remark (Bayesian deviance and the generalized linear model deviance)

Remark 0.1 (Bayesian deviance and the generalized linear model deviance)*.* This deviance is minus twice the log-likelihood. A generalized linear model's deviance subtracts the same quantity for the saturated model, which does not depend on $\underset{\sim}{\theta}$, so the two differ by a constant that cancels in comparisons between models of the same data.

Definition 0.2 (Posterior mean deviance)**.** The **posterior mean deviance** $\overline{\operatorname{dev}_{\mathrm{DIC}}}$ is the posterior mean of the deviance of a parameter value:

$$\overline{\operatorname{dev}_{\mathrm{DIC}}} \stackrel{\mathrm{def}}{=} \mathbb{E}[\operatorname{dev}_{\mathrm{DIC}}(\underset{\sim}{\theta}) \mid \underset{\sim}{y}].$$

Definition 0.3 (Effective number of parameters)**.** Let $\overline{\operatorname{dev}_{\mathrm{DIC}}}$ be the posterior mean deviance, and let $\bar{\underset{\sim}{\theta}} \stackrel{\mathrm{def}}{=} \mathbb{E}[\underset{\sim}{\theta} \mid \underset{\sim}{y}]$ be the posterior mean of the parameters. The **effective number of parameters** is

$$p_D \stackrel{\mathrm{def}}{=} \overline{\operatorname{dev}_{\mathrm{DIC}}} - \operatorname{dev}_{\mathrm{DIC}}(\bar{\underset{\sim}{\theta}}).$$

Definition 0.4 (Deviance information criterion)**.** The **deviance information criterion** is

$$\mathrm{DIC} \stackrel{\mathrm{def}}{=} \operatorname{dev}_{\mathrm{DIC}}(\bar{\underset{\sim}{\theta}}) + 2\, p_D,$$

with $\operatorname{dev}_{\mathrm{DIC}}$ from Definition 0.1 and p_D from Definition 0.3.
Source: Wikipedia contributors [11].

Substituting Definition 0.3 gives an equivalent form:

$$
\begin{aligned}
\mathrm{DIC} &= \mathrm{dev}_{\mathrm{DIC}}(\bar{\underset{\sim}{\theta}}) + 2\big(\overline{\mathrm{dev}_{\mathrm{DIC}}} - \mathrm{dev}_{\mathrm{DIC}}(\bar{\underset{\sim}{\theta}})\big) && \text{(definition of } p_D\text{)} \\
&= 2\overline{\mathrm{dev}_{\mathrm{DIC}}} - \mathrm{dev}_{\mathrm{DIC}}(\bar{\underset{\sim}{\theta}}) && \text{(collecting terms)} \\
&= \overline{\mathrm{dev}_{\mathrm{DIC}}} + \big(\overline{\mathrm{dev}_{\mathrm{DIC}}} - \mathrm{dev}_{\mathrm{DIC}}(\bar{\underset{\sim}{\theta}})\big) && \text{(splitting } 2\overline{\mathrm{dev}_{\mathrm{DIC}}}\text{)} \\
&= \overline{\mathrm{dev}_{\mathrm{DIC}}} + p_D && \text{(definition of } p_D\text{).}
\end{aligned}
$$

The term $\mathrm{dev}_{\mathrm{DIC}}(\bar{\underset{\sim}{\theta}})$ rewards fit, while p_D penalizes complexity: it measures how much the deviance is reduced, on average, by letting the parameters adapt to the data. As with Akaike's criterion, lower DIC is better, and only differences in DIC between models fit to the same data are meaningful [1, Ch. 13], p. 306. Both $\overline{\mathrm{dev}_{\mathrm{DIC}}}$ and $\bar{\underset{\sim}{\theta}}$ are Monte Carlo estimates from the posterior draws, so DIC is a by-product of an MCMC run.

Example 0.1 (DIC for the Bernoulli model)**.** For $r = 55$ successes in $n = 91$ trials, $\mathrm{dev}_{\mathrm{DIC}}(\pi) = -2(r\log\pi + (n-r)\log(1-\pi))$. We estimate $\overline{\mathrm{dev}_{\mathrm{DIC}}}$ and $\bar{\pi}$ from 5,000 draws from the Beta$(56, 37)$ posterior of the Beta-Bernoulli example:

```
deviance_bernoulli <- function(pi, r = 55, n = 91) {
  -2 * (r * log(pi) + (n - r) * log(1 - pi))
}
set.seed(4)
pi_draws <- stats::rbeta(5000, shape1 = 56, shape2 = 37)
d_bar <- mean(deviance_bernoulli(pi_draws))
d_at_mean <- deviance_bernoulli(mean(pi_draws))
c(
  d_bar = d_bar,
  d_at_mean = d_at_mean,
  p_d = d_bar - d_at_mean,
  dic = 2 * d_bar - d_at_mean
) |>
  round(2)
#>     d_bar d_at_mean       p_d       dic
#>    123.13    122.16      0.97    124.10
```

The model has one parameter, and its effective number of parameters p_D is close to 1.

i Remark (Cautions about DIC)

Remark 0.2 (Cautions about DIC). DIC should be used with care: it can behave poorly for models with weakly identified parameters or markedly non-Gaussian posteriors [1, Ch. 13], p. 306.

References

[1] A. J. Dobson and A. G. Barnett, *An introduction to generalized linear models*, 4th ed. CRC press, 2018. doi: 10.1201/9781315182780.

[2] Wikipedia contributors, "Markov chain monte carlo — Wikipedia, the free encyclopedia." 2026. Available: https://en.wikipedia.org/wiki/Markov_chain_Monte_Carlo

[3] Wikipedia contributors, "Monte carlo method — Wikipedia, the free encyclopedia." 2026. Available: https://en.wikipedia.org/wiki/Monte_Carlo_method

[4] Wikipedia contributors, "Markov chain — Wikipedia, the free encyclopedia." 2026. Available: https://en.wikipedia.org/wiki/Markov_chain

[5] Wikipedia contributors, "Stationary distribution — Wikipedia, the free encyclopedia." 2026. Available: https://en.wikipedia.org/wiki/Stationary_distribution

[6] G. Casella and R. Berger, *Statistical inference*, 2nd ed. Cengage Learning, 2002. Available: https://www.cengage.com/c/statistical-inference-2e-casella-berger/9780534243128/

[7] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian data analysis*, 3rd ed. in Chapman & hall/CRC texts in statistical science. Boca Raton, FL: CRC Press, 2013. doi: 10.1201/b16018.

[8] C. P. Robert and G. Casella, *Monte carlo statistical methods*, 2nd ed. in Springer texts in statistics. New York: Springer, 2004. doi: 10.1007/978-1-4757-4145-2.

[9] Wikipedia contributors, "Metropolis–hastings algorithm — Wikipedia, the free encyclopedia." 2026. Available: https://en.wikipedia.org/wiki/Metropolis%E2%80%93Hastings_algorithm

[10] Wikipedia contributors, "Gibbs sampling — Wikipedia, the free encyclopedia." 2026. Available: https://en.wikipedia.org/wiki/Gibbs_sampling

[11] Wikipedia contributors, "Deviance information criterion — Wikipedia, the free encyclopedia." 2026. Available: https://en.wikipedia.org/wiki/Deviance_information_criterion