- So we do not need to know the value of the normalizing constant B to create this Markov chain!

<u>Note:</u> The stationary probabilities $\{\pi(j)\}$ of the chain will be limiting probabilities if

<u>Example 1:</u> Suppose we want a sample from a discrete distribution with pmf

$$p(x) = \frac{c}{x^4}, \quad x = 1, 2, \ldots$$

for some constant c.

- Since a ______________ distribution has support on $\{0, 1, 2, \ldots\}$, we could use a __________________ as the proposal distribution.

- Which __________________ distribution?

Maybe use the one whose parameter equals

That is,

Given state i, we will accept each proposed candidate j with probability

- In our implementation, in deciding whether to accept the candidate j, we will generate $U \sim Unif(0,1)$ and accept j if

- See R code for implementation.

Example 2: The standard Gumbel distribution (often used to model (standardized) annual maximum rainfall amounts) has pdf

- It has mean ≈ 0.577 and variance ≈ 1.6.
- Generate a sample from the standard Gumbel distribution.
- Because a ____________ distribution has support on $(-\infty, \infty)$, we can use a ____________ as the proposal distribution.
- Which one?

- When the proposal distribution is symmetric around the current state, then $\alpha(i,j)$ does not depend on the proposal distribution!

- See implementation in R with several choices of proposal variance.
- Note the acceptance probability changes with the proposal distribution's variance.
- This acceptance probability will affect how fast the chain converges to its limiting (stationary) probabilities.
- It is recommended to choose the proposal distribution so that the acceptance rate is

The Gibbs Sampler

- The Gibbs sampler is a special case of the M-H algorithm that is designed to sample from multivariate distributions.
- Let $\underset{\sim}{X} = (x_1, \ldots, x_n)$ be a random vector with pmf

where $g(\underset{\sim}{x})$ is known but the constant c may not be known.

- To use the Gibbs sampler, we must be able to sample from the <u>full</u> <u>conditional</u> <u>distributions</u>:

- Given a current state $\underset{\sim}{x} = (x_1, \ldots, x_n)$, we propose a candidate vector $\underset{\sim}{y}$ in this way:

* We randomly choose one of the coordinates, and for coordinate i that is chosen, we generate x from the corresponding full conditional:

- Then set

- Then use the M-H algorithm with

Then our acceptance ratio is

- So the Gibbs sampler is a case of the M-H algorithm where the candidate is <u>always</u> accepted.

Example 3: We sample from a bivariate normal density with $\mu_1 = \mu_2 = 0$, $\sigma_1^2 = \sigma_2^2 = 1$, and $\rho = 0.5$. The joint pdf is:

- The full conditionals $f(x_1 | x_2)$ and $f(x_2 | x_1)$ can be derived and found to be univariate normals:

- With MCMC methods, it may take a while for the chain to converge to the stationary distribution, so in practice we often discard the first few hundred (or thousand?) sampled values as "burn-in". See R example.