

Example: Demand function using STATA

二人以上の世帯のうち勤労者世帯（全国）

year
y = 実収入（一月当たり，実質データ）
q1 = 穀類支出額（一年当たり，実質データ）
p1 = 穀類価格（相対価格=穀類CPI/総合CPI）
p2 = 魚介類価格（相対価格=魚介類CPI/総合CPI）
p3 = 肉類価格（相対価格=肉類CPI/総合CPI）

| year | y | q1 | p1 | p2 | p3 |
|------|--------|--------|----------|----------|----------|
| 2000 | 567865 | 7087.0 | 1.043390 | 0.884965 | 0.818365 |
| 2001 | 561722 | 6993.1 | 1.032520 | 0.886179 | 0.822154 |
| 2002 | 553768 | 6934.4 | 1.031800 | 0.891282 | 0.834872 |
| 2003 | 539928 | 6816.8 | 1.050410 | 0.876543 | 0.843621 |
| 2004 | 547006 | 6651.6 | 1.089510 | 0.865226 | 0.868313 |
| 2005 | 541367 | 6615.8 | 1.020640 | 0.862745 | 0.887513 |
| 2006 | 540863 | 6523.7 | 1.000000 | 0.878601 | 0.891975 |
| 2007 | 543994 | 6680.5 | 0.994856 | 0.886831 | 0.908436 |
| 2008 | 541821 | 6494.7 | 1.043610 | 0.894523 | 0.932049 |
| 2009 | 533154 | 6477.3 | 1.066870 | 0.898148 | 0.934156 |
| 2010 | 539577 | 6458.2 | 1.040420 | 0.889119 | 0.924352 |
| 2011 | 529750 | 6448.4 | 1.025960 | 0.894081 | 0.925234 |
| 2012 | 538988 | 6377.6 | 1.057170 | 0.904366 | 0.917879 |
| 2013 | 542018 | 6360.7 | 1.047620 | 0.909938 | 0.916149 |
| 2014 | 523953 | 6174.6 | 1.016130 | 0.971774 | 0.960685 |
| 2015 | 525669 | 6268.0 | 1.000000 | 1.000000 | 1.000000 |
| 2016 | 527501 | 6244.8 | 1.018020 | 1.019020 | 1.017020 |

2017 531693 6106.6 1.027890 1.066730 1.025900

```
. tsset year
      time variable: year, 2000 to 2017
             delta: 1 unit
```

```
. reg q1 y p1 p2 p3 if year>2000.5
```

| Source | SS | df | MS | Number of obs | = | 17 |
|----------|------------|----|------------|---------------|---|--------|
| Model | 913640.443 | 4 | 228410.111 | F(4, 12) | = | 25.83 |
| Residual | 106100.077 | 12 | 8841.67308 | Prob > F | = | 0.0000 |
| | | | | R-squared | = | 0.8960 |
| | | | | Adj R-squared | = | 0.8613 |
| Total | 1019740.52 | 16 | 63733.7825 | Root MSE | = | 94.03 |

| q1 | Coef. | Std. Err. | t | P> t | [95% Conf. Interval] |
|-------|-----------|-----------|-------|-------|----------------------|
| y | .0067843 | .0045443 | 1.49 | 0.161 | -.003117 .0166856 |
| p1 | -1128.834 | 998.7698 | -1.13 | 0.280 | -3304.966 1047.299 |
| p2 | 356.8095 | 806.2301 | 0.44 | 0.666 | -1399.815 2113.434 |
| p3 | -3442.221 | 1130.078 | -3.05 | 0.010 | -5904.448 -979.9931 |
| _cons | 6850.563 | 3179.316 | 2.15 | 0.052 | -76.57278 13777.7 |

```
. gmm (q1-{b0}-{b1}*y-{b2}*p1-{b3}*p2-{b4}*p3) if year>2000.5, instruments(y p1
> p2 p3)
```

Step 1

Iteration 0: GMM criterion $Q(b) = 42400764$
Iteration 1: GMM criterion $Q(b) = 6.781e-12$
Iteration 2: GMM criterion $Q(b) = 6.781e-12$ (backed up)

Step 2

Iteration 0: GMM criterion $Q(b) = 1.966e-15$
Iteration 1: GMM criterion $Q(b) = 1.963e-15$ (backed up)

convergence not achieved

The Gauss-Newton stopping criterion has been met but missing standard errors indicate some of the parameters are not identified.

GMM estimation

Number of parameters = 5
Number of moments = 5
Initial weight matrix: Unadjusted Number of obs = 17
GMM weight matrix: Robust

| | Coef. | Robust Std. Err. | z | P> z | [95% Conf. Interval] | |
|-----|-----------|---------------------|-------|-------|----------------------|----------|
| /b0 | 6850.563 | 17645.71 | 0.39 | 0.698 | -27734.4 | 41435.53 |
| /b1 | .0067843 | .0282325 | 0.24 | 0.810 | -.0485504 | .062119 |
| /b2 | -1128.834 | 1057.915 | -1.07 | 0.286 | -3202.309 | 944.6415 |
| /b3 | 356.8095 | 1565.86 | 0.23 | 0.820 | -2712.219 | 3425.838 |
| /b4 | -3442.221 | 5085.561 | -0.68 | 0.498 | -13409.74 | 6525.296 |

Instruments for equation 1: y p1 p2 p3 _cons
Warning: convergence not achieved

```
. gmm (q1-{b0}-{b1}*y-{b2}*p1-{b3}*p2-{b4}*p3) if year>2000.5, instruments(p1 p2  
> p3 1.p1 1.p2 1.p3)
```

Step 1

```
Iteration 0: GMM criterion Q(b) = 42404066  
Iteration 1: GMM criterion Q(b) = 2790.3146  
Iteration 2: GMM criterion Q(b) = 2790.3146
```

Step 2

```
Iteration 0: GMM criterion Q(b) = .3201826  
Iteration 1: GMM criterion Q(b) = .2469289  
Iteration 2: GMM criterion Q(b) = .2469289
```

GMM estimation

```
Number of parameters = 5  
Number of moments = 7  
Initial weight matrix: Unadjusted  
GMM weight matrix: Robust  
Number of obs = 17
```

| | Coef. | Robust Std. Err. | z | P> z | [95% Conf. Interval] | |
|-----|-----------|---------------------|-------|-------|----------------------|----------|
| /b0 | -1192.466 | 4669.012 | -0.26 | 0.798 | -10343.56 | 7958.63 |
| /b1 | .0186312 | .0067682 | 2.75 | 0.006 | .0053657 | .0318967 |

| | | | | | | |
|-----|-----------|----------|-------|-------|-----------|----------|
| /b2 | -1016.864 | 780.979 | -1.30 | 0.193 | -2547.554 | 513.8271 |
| /b3 | -905.5585 | 598.0885 | -1.51 | 0.130 | -2077.79 | 266.6734 |
| /b4 | -499.8064 | 1147.985 | -0.44 | 0.663 | -2749.815 | 1750.202 |

 Instruments for equation 1: p1 p2 p3 L.p1 L.p2 L.p3 _cons

. estat overid

Test of overidentifying restriction:

Hansen's J chi2(2) = 4.19779 (p = 0.1226)

Going back to:

Introduction to Causal Inference (因果推論)

We have shown the following theorem:

● **Theorem:**

Assume that D_i is independent of Y_i , given X_i , for all i .

Suppose that D_i , Y_i and X_i are identically distributed.

Then, we have the following theorem:

$$\frac{1}{n} \sum_{i=1}^n \left(\frac{D_i Y_i}{\pi(X_i)} - \frac{(1 - D_i) Y_i}{1 - \pi(X_i)} \right) \xrightarrow{p} E(Y^1 - Y^0)$$

In practice, we have to show the following asymptotic normality:

● **Theorem:**

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n \left(\left(\frac{D_i Y_i}{\pi(X_i)} - \frac{(1 - D_i) Y_i}{1 - \pi(X_i)} \right) - \mu \right) \xrightarrow{d} N(0, \sigma^2)$$

where $\mu = E(Y^1 - Y^0)$ and $\sigma^2 = V(Y^1 - Y^0)$.

We have to assume each individual is mutually independent and identically distributed.

That is, the i th individual does not influence j th individual for $i \neq j$, and it is distributed with same mean μ and same variance σ^2 for all i .

$\left(\frac{D_i Y_i}{\pi(X_i)} - \frac{(1 - D_i) Y_i}{1 - \pi(X_i)} \right)$ is taken as one random variable with mean μ and variance σ^2 .

Then, the central limit theorem is applied.

The estimate of σ^2 is given by:

$$\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n \left(\frac{D_i Y_i}{\pi(X_i)} - \frac{(1 - D_i) Y_i}{1 - \pi(X_i)} \right)^2 - n \left(\frac{1}{n} \sum_{i=1}^n \left(\frac{D_i Y_i}{\pi(X_i)} - \frac{(1 - D_i) Y_i}{1 - \pi(X_i)} \right) \right)^2$$

(*) For random variable X , $V(X) = X((X - \mu)^2) = E(X^2) - \mu^2$, where $\mu = E(X)$.

(**) $\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 = \frac{1}{n} \sum_{i=1}^n x_i^2 - n\bar{x}^2$, where $\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$.

Therefore, in practice, the estimator of the average treatment effect is distributed as follows:

$$\frac{1}{n} \sum_{i=1}^n \left(\frac{D_i Y_i}{\pi(X_i)} - \frac{(1 - D_i) Y_i}{1 - \pi(X_i)} \right) \sim N\left(\mu, \frac{\hat{\sigma}^2}{n}\right)$$

for large n .

We can derive the confidence interval for μ .

We can test the null hypothesis $H_0 : \mu = 0$ (i.e., no treatment effect in average).

6 Bayesian Estimation

Bayes' procedure is briefly discussed (see Zellner (1971), Bernardo and Smith (1994), O'Hagan (1994), Hogg and Craig (1995) and so on for further discussion).

6.1 Elements of Bayesian Inference

When we have the random sample (X_1, X_2, \dots, X_n) , consider estimating the unknown parameter θ . The maximum likelihood estimator is introduced for estimation of the parameter. Suppose that X_1, X_2, \dots, X_n are mutually independently distributed and X_i has a probability density function $f(x; \theta)$, where θ is the unknown parameter to be estimated. The joint density of X_1, X_2, \dots, X_n is given by:

$$f(x_1, x_2, \dots, x_n; \theta) = \prod_{i=1}^n f(x_i; \theta),$$

which is called the likelihood function, denoted by $l(\theta) = f(x_1, x_2, \dots, x_n; \theta)$.

In Bayes' estimation, the parameter is taken as a random variable, say Θ , where a prior information on Θ is taken into account for estimation. The joint density function (or the likelihood function) is regarded as the conditional density function of X_1, X_2, \dots, X_n given $\Theta = \theta$. Therefore, we write the likelihood function as the conditional density $f(x_1, x_2, \dots, x_n|\theta)$. The probability density function of Θ is called the **prior probability density function** and given by $f_\theta(\theta)$. The conditional probability density function, $f_{\theta|x}(\theta|x_1, x_2, \dots, x_n)$, have to be obtained, which is represented as:

$$f_{\theta|x}(\theta|x_1, x_2, \dots, x_n) = \frac{f(x_1, x_2, \dots, x_n|\theta)f_\theta(\theta)}{\int f(x_1, x_2, \dots, x_n|\theta)f_\theta(\theta) d\theta} \\ \propto f(x_1, x_2, \dots, x_n|\theta)f_\theta(\theta).$$

The relationship in the first equality is known as Bayes' formula. The conditional probability density function of Θ given $X_1 = x_1, X_2 = x_2, \dots, X_n = x_n$, i.e., $f_{\theta|x}(\theta|x_1, x_2, \dots, x_n)$, is called the **posterior probability density function**, which is propor-

tional to the product of the likelihood function and the prior density function.

6.1.1 Bayesian Point Estimate

Thus, the Bayesian approach yields the posterior probability density function for Θ . To obtain a point estimate of Θ , we introduce a loss function, denoted by $L(\Theta, \hat{\theta})$, where $\hat{\theta}$ indicates a point estimate depending on $X_1 = x_1, X_2 = x_2, \dots, X_n = x_n$. Since Θ is considered to be random, $L(\Theta, \hat{\theta})$ is also random. One solution which yields point estimates is to find the value of Θ that minimizes the mathematical expectation of the **loss function**, i.e.,

$$\min_{\hat{\theta}} E(L(\Theta, \hat{\theta})) = \min_{\hat{\theta}} \int L(\theta, \hat{\theta}) f_{\theta|x}(\theta|x_1, x_2, \dots, x_n) d\theta,$$

where the absolute value of $E(L(\Theta, \hat{\theta}))$ is assumed to be finite.

Now we specify the loss function as: $L(\Theta, \hat{\theta}) = (\Theta - \hat{\theta})' A (\Theta - \hat{\theta})$, which is called the **quadratic loss function**, where A is a known nonstochastic positive definite sym-

metric matrix. Then, the solution gives us the posterior mean of Θ , i.e.,

$$\hat{\theta} = E(\Theta) = \int \theta f_{\theta|x}(\theta|x_1, x_2, \dots, x_n) d\theta.$$

An alternative loss function is given by: $L(\Theta, \hat{\theta}) = |\Theta - \hat{\theta}|$, which is called the **absolute error loss function**, where both Θ and $\hat{\theta}$ are assumed to be scalars. Then, the median of the posterior probability density function is an optimal point estimate of θ , i.e.,

$$\hat{\theta} = \text{median of the posterior probability density function.}$$

We have shown two Bayesian point estimates. Hereafter, the quadratic loss function is adopted for estimation. That is, the posterior mean of Θ is taken as the point estimate.

6.1.2 Bayesian Interval for Parameter

Given that the posterior probability density function $f_{\theta|x}(\theta|x_1, x_2, \dots, x_n)$ has been obtained, it is possible to compute the probability that the parameter Θ lies in a particular subregion, R , of the parameter space. That is, we may compute the following probability:

$$P(\Theta \in R) = \int_R f_{\theta|x}(\theta|x_1, x_2, \dots, x_n) d\theta.$$

When the above probability is set to be $1 - \alpha$, it is possible to find the region that satisfies $P(\Theta \in R) = 1 - \alpha$, which region is not necessarily unique. In the case where Θ is a scalar, one possibility to determine the unique region $R = \{\Theta|a < \Theta < b\}$ is to obtain a and b by minimizing the distance $b - a$ subject to $\int_a^b f_{\theta|x}(\theta|x_1, x_2, \dots, x_n) d\theta = 1 - \alpha$. By solving this minimization problem, determining a and b such that $\int_a^b f_{\theta|x}(\theta|x_1, x_2, \dots, x_n) d\theta = 1 - \alpha$ and $f_{\theta|x}(a|x_1, x_2, \dots, x_n) = f_{\theta|x}(b|x_1, x_2, \dots, x_n)$ leads to the shortest interval with probability $1 - \alpha$.

6.1.3 Prior Probability Density Function

We discuss a little bit about the prior probability density function. In the case where we know any information about the parameter θ beforehand, the plausible estimate of the parameter might be obtained if the parameter θ is estimated by including the prior information. For example, if we know that Θ is normally distributed with mean θ_0 and variance Σ_0 , the prior density $f_\theta(\theta)$ is given by $N(\theta_0, \Sigma_0)$, where θ_0 and Σ_0 are known. On the contrary, we have the case where we do not know any prior information about the parameter θ . In this case, we may take the prior density as:

$$f_\theta(\theta) \propto \text{constant},$$

where the prior density of Θ is assumed to be uniform, which prior is called the **improper prior**, the **noninformative prior**, the **flat prior** or the **diffuse prior**. Then,

the posterior density is given by:

$$f_{\theta|x}(\theta|x_1, x_2, \dots, x_n) \propto f(x_1, x_2, \dots, x_n|\theta).$$

That is, the posterior density function is proportional to the likelihood function.

Example: Suppose that X_1, X_2, \dots, X_n are mutually independently, identically and normally distributed with mean μ and variance σ^2 . Then, the likelihood function is given by:

$$\begin{aligned} f(x_1, x_2, \dots, x_n|\theta) &= \prod_{i=1}^n f(x_i; \theta) = \prod_{i=1}^n (2\pi\sigma^2)^{-1/2} \exp\left(-\frac{1}{2\sigma^2}(x_i - \mu)^2\right) \\ &= (2\pi\sigma^2)^{-n/2} \exp\left(-\frac{1}{2\sigma^2} \sum_{i=1}^n (x_i - \mu)^2\right), \end{aligned}$$

where θ indicates μ in this case. For simplicity of discussion, we assume that σ^2 is known. Therefore, we focus on μ .

Now, we consider two prior density functions for μ . One is noninformative and another is normal, i.e.,

- (i) **Noninformative Prior:** $f_{\theta}(\mu) \propto \text{constant}$, where μ is uniformly distributed.
- (ii) **Normal Prior:** $f_{\theta}(\mu) = (2\pi\sigma_0^2)^{-1/2} \exp\left(-\frac{1}{2\sigma_0^2}(\mu - \mu_0)^2\right)$, where μ_0 and σ_0 are assumed to be known.

For each prior density, we obtain the posterior distributions as follows:

- (i) When the prior density is noninformative, the posterior density function is:

$$\begin{aligned} f_{\theta|x}(\mu|x_1, x_2, \dots, x_n) \\ \propto f(x_1, x_2, \dots, x_n|\mu) &= (2\pi\sigma^2)^{-n/2} \exp\left(-\frac{1}{2\sigma^2} \sum_{i=1}^n (x_i - \mu)^2\right) \\ &= (2\pi\sigma^2)^{-n/2} \exp\left(-\frac{1}{2\sigma^2} \sum_{i=1}^n (x_i - \bar{x})^2 - \frac{1}{2\sigma^2} n(\bar{x} - \mu)^2\right) \end{aligned}$$

$$\begin{aligned}
&= (2\pi\sigma^2)^{-n/2} \exp\left(-\frac{(n-1)s^2}{2\sigma^2} - \frac{1}{2\sigma^2/n}(\mu - \bar{x})^2\right) \\
&\propto \exp\left(-\frac{1}{2\sigma^2/n}(\mu - \bar{x})^2\right)
\end{aligned}$$

where $\bar{x} = \sum_{i=1}^n x_i/n$ and $s^2 = \sum_{i=1}^n (x_i - \bar{x})^2/(n-1)$. Thus, the posterior density of μ represents the normal distribution with mean \bar{x} and variance σ^2/n . Since under the quadratic loss function the point estimate of μ is given by the posterior mean, \bar{x} gives us Bayes' point estimate. The Bayesian interval estimate of μ is: $(\bar{x} - z_{\alpha/2}\sigma/\sqrt{n}, \bar{x} + z_{\alpha/2}\sigma/\sqrt{n})$, because from the posterior density function we have $P\left(\left|\frac{\mu - \bar{x}}{\sigma/\sqrt{n}}\right| < z_{\alpha/2}\right) = 1 - \alpha$.

(ii) When the prior density is normal, the posterior density function is given by:

$$\begin{aligned}
&f_{\theta|x}(\mu|x_1, x_2, \dots, x_n) \\
&\propto f(x_1, x_2, \dots, x_n|\mu)f_{\theta}(\mu)
\end{aligned}$$

$$\begin{aligned}
&= (2\pi\sigma^2)^{-n/2} \exp\left(-\frac{1}{2\sigma^2} \sum_{i=1}^n (x_i - \mu)^2\right) \\
&\quad \times (2\pi\sigma_0^2)^{-1/2} \exp\left(-\frac{1}{2\sigma_0^2} (\mu - \mu_0)^2\right) \\
&\propto \exp\left(-\frac{1}{2\sigma^2/n} (\mu - \bar{x})^2\right) \times \exp\left(-\frac{1}{2\sigma_0^2} (\mu - \mu_0)^2\right) \\
&\propto \exp\left(-\frac{1}{2} \left(\frac{\sigma_0^2 + \sigma^2/n}{\sigma_0^2\sigma^2/n}\right) \left(\mu - \frac{\bar{x}\sigma_0^2 + \mu_0\sigma^2/n}{\sigma_0^2 + \sigma^2/n}\right)^2\right),
\end{aligned}$$

which indicates that the posterior density of μ is a normal distribution with mean $\frac{\bar{x}\sigma_0^2 + \mu_0\sigma^2/n}{\sigma_0^2 + \sigma^2/n}$ and variance $\left(\frac{\sigma_0^2 + \sigma^2/n}{\sigma_0^2\sigma^2/n}\right)^{-1}$. The posterior mean is rewritten as:

$$\frac{\bar{x}\sigma_0^2 + \mu_0\sigma^2/n}{\sigma_0^2 + \sigma^2/n} = \bar{x}w + \mu_0(1 - w),$$

where $w = \sigma_0^2/(\sigma_0^2 + \sigma^2/n)$. \bar{x} is a maximum likelihood estimate of μ and μ_0 is a prior mean of μ . Thus, the posterior mean is the weighted average of \bar{x} and

μ_0 . As $w \rightarrow 1$, i.e., as $\sigma_0^2 \rightarrow \infty$, the posterior mean approaches \bar{x} , which is equivalent to the posterior mean with the noninformative prior.

6.2 Sampling Methods

6.2.1 Gibbs Sampling

The Gibbs sampler shows how to generate random draws from the unconditional densities under the situation that we can generate random draws from two conditional densities.

Geman and Geman (1984), Tanner and Wong (1987), Gelfand, Hills, Racine-Poon and Smith (1990), Gelfand and Smith (1990), Carlin and Polson (1991), Zeger and Karim (1991), Casella and George (1992), Gamerman (1997) and so on developed the Gibbs sampling theory. Carlin, Polson and Stoffer (1992), Carter and Kohn (1994, 1996) and Geweke and Tanizaki (1999, 2001) applied the Gibbs sampler to the non-

linear and/or non-Gaussian state-space models. There are numerous other applications of the Gibbs sampler. The Gibbs sampling theory is concisely described as follows.

We can deal with more than two random variables, but we consider two random variables X and Y in order to make things easier. Two conditional density functions, $f_{x|y}(x|y)$ and $f_{y|x}(y|x)$, are assumed to be known, which denote the conditional distribution function of X given Y and that of Y given X , respectively. Suppose that we can easily generate random draws of X from $f_{x|y}(x|y)$ and those of Y from $f_{y|x}(y|x)$. However, consider the case where it is not easy to generate random draws from the joint density of X and Y , denoted by $f_{xy}(x, y)$. In order to have the random draws of (X, Y) from the joint density $f_{xy}(x, y)$, we take the following procedure:

- (i) Take the initial value of X as x_{-M} .
- (ii) Given x_{i-1} , generate a random draw of Y , i.e., y_i , from $f(y|x_{i-1})$.

(iii) Given y_i , generate a random draw of X , i.e., x_i , from $f(x|y_i)$.

(iv) Repeat the procedure for $i = -M + 1, -M + 2, \dots, 1$.

From the convergence theory of the Gibbs sampler, as M goes to infinity, we can regard x_1 and y_1 as random draws from $f_{xy}(x, y)$, which is a joint density function of X and Y . M denotes the **burn-in period**, and the first M random draws, (x_i, y_i) for $i = -M + 1, -M + 2, \dots, 0$, are excluded from further consideration. When we want N random draws from $f_{xy}(x, y)$, Step (iv) should be replaced by Step (iv)', which is as follows.

(iv)' Repeat the procedure for $i = -M + 1, -M + 2, \dots, N$.

As in the Metropolis-Hastings algorithm, the algorithm shown in Steps (i) – (iii) and (iv)' is formulated as follows:

$$f_i(u) = \int f^*(u|v)f_{i-1}(v) dv.$$

For convergence of the Gibbs sampler, we need to have the invariant distribution $f(u)$ which satisfies $f_i(u) = f_{i-1}(u) = f(u)$. If we have the reversibility condition shown in equation (3), i.e.,

$$f^*(v|u)f(u) = f^*(u|v)f(v),$$

the random draws based on the Gibbs sampler converge to those from the invariant distribution, which implies that there exists the invariant distribution $f(u)$. Therefore, in the Gibbs sampling algorithm, we have to find the transition distribution, i.e., $f^*(u|v)$. Here, we consider that both u and v are bivariate vectors. That is, $f^*(u|v)$ and $f_i(u)$ denote the bivariate distributions. x_i and y_i are generated from $f_i(u)$ through $f^*(u|v)$, given $f_{i-1}(v)$. Note that $u = (u_1, u_2) = (x_i, y_i)$ is taken while $v = (v_1, v_2) = (x_{i-1}, y_{i-1})$ is set. The transition distribution in the Gibbs sampler is taken as:

$$f^*(u|v) = f_{y|x}(u_2|u_1)f_{x|y}(u_1|v_2)$$

Thus, we can choose $f^*(u|v)$ as shown above. Then, as i goes to infinity, (x_i, y_i) tends in distribution to a random vector whose joint density is $f_{xy}(x, y)$. See, for example, Geman and Geman (1984) and Smith and Roberts (1993).

Furthermore, under the condition that there exists the invariant distribution, the basic result of the Gibbs sampler is as follows:

$$\frac{1}{N} \sum_{i=1}^N g(x_i, y_i) \longrightarrow \mathbb{E}(g(x, y)) = \iint g(x, y) f_{xy}(x, y) dx dy, \quad \text{as } N \longrightarrow \infty,$$

where $g(\cdot, \cdot)$ is a function.

The Gibbs sampler is a powerful tool in a Bayesian framework. Based on the conditional densities, we can generate random draws from the joint density.

Remark 1: We have considered the bivariate case, but it is easily extended to the multivariate cases. That is, it is possible to take multi-dimensional vectors for x and

y. Taking an example, as for the tri-variate random vector (X, Y, Z) , if we generate the i th random draws from $f_{x|yz}(x|y_{i-1}, z_{i-1})$, $f_{y|xz}(y|x_i, z_{i-1})$ and $f_{z|xy}(z|x_i, y_i)$, sequentially, we can obtain the random draws from $f_{xyz}(x, y, z)$.

Remark 2: Let X , Y and Z be the random variables. Take an example of the case where X is highly correlated with Y . If we generate random draws from $f_{x|yz}(x|y, z)$, $f_{y|xz}(y|x, z)$ and $f_{z|xy}(z|x, y)$, it is known that convergence of the Gibbs sampler is slow. In this case, without separating X and Y , random number generation from $f(x, y|z)$ and $f(z|x, y)$ yields better random draws from the joint density $f(x, y, z)$.

Example: X_1, X_2, \dots, X_n are mutually independent with $X_i \sim N(\mu, \sigma^2)$.

Derive Bayesian estimation of μ and σ^2 .

Assume that the prior distributions: $\mu \sim N(\mu_0, \sigma_0^2)$ and $\sigma^2 \sim IG(\alpha_0, \beta_0)$, i.e.,