Bayesian analysis of Gaussian samples

Marshall Bradley

Bayesian probability theory is used to analyze random data drawn from a Gaussian probability distribution with unknown mean and standard deviation. The analysis begins with the determination of the parameters of a single sample and then proceeds to the joint analysis of a pair of samples. The Bayesian evidence is used to characterise the probabilities of hypotheses about the relationships between the means and standard deviations of the two samples. Of particular interest is the probability that the two samples have different means.

Theoretical analysis of a single sample

We begin with the analysis of a single sample. Our hypothesis is that on initial hypothesis $\mathcal{I}$ we are drawing data from a Gaussian distribution with unknown mean μ and standard deviation σ. Our desire is to make inferences about μ and σ based upon prior knowledge and any data that we might observe. If x is a data value then the sampling distribution is

$$h(x \mid \mu, \sigma) = \frac{1}{(2\pi)^{1/2}} \exp\left\{-\frac{(x-\mu)^2}{2\sigma^2}\right\}, \quad -\infty < x < \infty.$$

On initial hypothesis $\mathcal{I}$ we assume that the unknown parameter μ is uniformly distributed on the interval (μ_a, μ_b) and that the unknown parameter σ follows a Jeffreys prior (log uniform) on the interval (σ_a, σ_b). Thus in explicit form our prior probability distributions are

$$f(\mu) = \frac{1}{\mu_b - \mu_a}, \ \mu_a < \mu < \mu_b, \quad g(\sigma) = \frac{1}{\log(\sigma_b/\sigma_a)} \frac{1}{\sigma}, \quad \sigma_a < \sigma < \sigma_b.$$

The likelihood of obtaining the independent data sample D consisting of the n values $x_1\, x_2 \ldots x_n$ is

$$L(D \mid \mu, \sigma) = \prod_{i=1}^{n} h(x_i \mid \mu, \sigma) = \frac{1}{(2\pi)^{n/2}\,\sigma^n} \exp\left\{-\frac{1}{2\sigma^2} \sum_{i=1}^{n} (x_i - \mu)^2\right\}.$$

If we define the statistics $\bar{x}$ and s of the data $x_1\, x_2 \ldots x_n$ via

$$n\,\bar{x} = \sum_{i=1}^{n} x_i \ \text{ and } \ (n-1)\,s^2 = \sum_{i=1}^{n} (x - \bar{x})^2,$$

then the sum of squared deviations of the data from the mean μ can be written in the form

$$\sum_{i=1}^{n} (x_i - \mu)^2 = \sum_{i=1}^{n} [(x_i - \bar{x}) + (\bar{x} - \mu)]^2 = (n-1)\,s^2 + n(\mu - \bar{x})^2$$

since $\sum_{i=1}^{n} (x_i - \bar{x}) = 0$.

The quantities $\bar{x}$, s and n are the sufficient statistics of the data. The likelihood in terms of the sufficient statistics is

$$L(\bar{x}, s, n \mid \mu, \sigma) = \frac{1}{(2\pi)^{n/2}\sigma^n} \exp\left\{-\frac{1}{2\sigma^2}\left[(n-1)s^2 + n(\mu-\bar{x})^2\right]\right\}.$$

Bayes theorem tells us that the posterior distribution of the parameters (μ, σ) given the data D is proportional to the product of prior and likelihood:

$$P(\mu, \sigma \mid D) = \frac{1}{E} f(\mu) g(\sigma) L(\bar{x}, s, n \mid \mu, \sigma)$$

where the evidence E is a normalization constant

$$E = \int_{\sigma_a}^{\sigma_b} \int_{\mu_a}^{\mu_b} f(\mu) g(\sigma) L(\bar{x}, s, n \mid \mu, \sigma) \, d\mu \, d\sigma .$$

In explicit terms the evidence is

$$E = \frac{1}{\mu_b-\mu_a} \frac{1}{\log(\sigma_b/\sigma_a)} \frac{1}{(2\pi)^{n/2}} \int_{\sigma_a}^{\sigma_b} \int_{\mu_a}^{\mu_b} \sigma^{-(n+1)} \exp\left\{-\frac{1}{2\sigma^2}\left[(n-1)s^2 + n(\mu-\bar{x})^2\right]\right\} d\mu \, d\sigma.$$

The evidence integral can be evaluated analytically but the evaluation is tedious and does not provide any great physical insight. A more productive path is to observe that the posterior parameter distribution is

$$P(\mu, \sigma \mid D) \propto \sigma^{-(n+1)} \exp\left\{-\frac{1}{2\sigma^2}\left[(n-1)s^2 + n(\mu-\bar{x})^2\right]\right\}$$

This distribution is large only when σ is near the sample standard deviation s and when μ is near the sample mean $\bar{x}$. Provided that s is much greater than σ_a and much less than σ_b, then the marginal distribution for the mean μ is

$$P(\mu \mid D) \propto \int_{\sigma_a}^{\sigma_b} \sigma^{-(n+1)} \exp\left\{-\frac{1}{2\sigma^2}\left[(n-1)s^2 + n(\mu-\bar{x})^2\right]\right\} d\sigma \propto \left\{(n-1)s^2 + n(\mu-\bar{x})^2\right\}^{-n/2}$$

This is a Student t distribution. For $n = 2$ it is a Cauchy distribution with very fat tails. If n is large then

$$P(\mu \mid D) \propto \exp\left\{-\frac{n(\mu-\bar{x})^2}{2s^2}\right\}$$

Thus the posterior distribution of μ is approximately Gaussian distributed with mean $\bar{x}$ and standard deviation $s/\sqrt{n}$. For n small this approximation is very poor.

Provided $\mu_a < \bar{x} - 2s/\sqrt{n}$ and $\mu_b > \bar{x} - 2s/\sqrt{n}$ then

$$P(\sigma \mid D) \propto \int_0^\infty \sigma^{-(n+1)} \exp\left\{-\frac{1}{2\sigma^2}\left[(n-1)s^2 + n(\mu-\bar{x})^2\right]\right\} d\mu \propto \sigma^{-n} \exp\left\{-\frac{1}{2\sigma^2}(n-1)s^2\right\}.$$

This is an inverse gamma distribution. Note that for $n = 1$ this reduces to

$$P(\sigma \mid D) \propto \sigma^{-1}$$

This is just the prior and all is as it should be.

In terms of a Mathematica function the unnormalized posterior is :

```
In[∘]:= Clear[μ, σ, xbar, s, n, μa, μb, σa, σb];
postUnormalized[μ_, σ_, xbar_, s_, n_, μa_, μb_, σa_, σb_] :=
   1/(μb - μa) 1/Log[σb / σa] 1/(2 π)^(n/2) 1/σ^(n+1) Exp[-((n - 1) s^2 + n (μ - xbar)^2)/(2 σ^2)]
```

The unnormalized μ posterior and σ posteriors are respectively:

In[∘]:=
```
postUμ[μ_, xbar_, s_, n_] := ((n - 1) s^2 + n (μ - xbar)^2)^(-n/2)
postUσ[σ_, s_, n_] := σ^-n Exp[-((n - 1) s^2)/(2 σ^2)]
```

When normalized these marginal probability distributions are :

In[∘]:=
```
postμ[μ_, xbar_, s_, n_] := 1/(((-1+n)^(-1/2-n/2) Sqrt[n] Sqrt[π] s^(1-n) Gamma[(1+n)/2])/Gamma[1+n/2]) ((n - 1) s^2 + n (μ - xbar)^2)^(-n/2)

postσ[σ_, s_, n_] := 1/(2^(1/2 (-3+n)) (-1 + n)^(1/2-n/2) s^(1-n) Gamma[1/2 (-1 + n)]) σ^-n Exp[-((n - 1) s^2)/(2 σ^2)]
```

We check to make sure that they are properly normalized :

In[∘]:=
```
check1 = Integrate[postσ[σ, s, n], {σ, 0, ∞}, Assumptions → {s > 0, n ≥ 2}];
check2 = Integrate[postμ[μ, xbar, s, n],
    {μ, -∞, ∞}, Assumptions → {s > 0, n ≥ 2, (μ - xbar)^2 > 0}];
{check1, check2}
```

Out[∘]=
```
{1, 1}
```

Visual comparison of two different samples

We consider a pair of river sediment toxicity samples from Gregory (2005) . Ultimately we want to make statements about the probability that these samples are somehow different. Here we restrict ourselves to visual comparisons using the results from the previous section.

The data and our assumed ranges of the prior parameters μ and σ are:

In[∘]:=
```
x = {13.2, 13.8, 8.7, 9.0, 8.6, 9.9, 14.2, 9.7, 10.7, 8.3, 8.5, 9.2};
y = {8.9, 9.1, 8.3, 6.0, 7.7, 9.9, 9.9, 8.9};
μa = 7.0; μb = 12.0;
σa = 1.0; σb = 4.0;
```

The prior probability density functions $f(\mu)$ and $g(\sigma)$ are:

In[ ]:=

```
gpriorμ = Plot[PDF[UniformDistribution[{μa, μb}], μ], {μ, μa - 1, μb + 1},
    Axes → False, Frame → True, FrameLabel → {…}, Filling → Axis];
priorσ = Plot[1/Log[σb / σa] 1/σ If[σ < σa, 0, 1] × If[σ < σb, 1, 0], {σ, σa - 1, σb + 1},
    Axes → False, Frame → True, FrameLabel → {…}, Filling → Axis];
GraphicsRow[{gpriorμ, priorσ}]
```

Out[ ]=

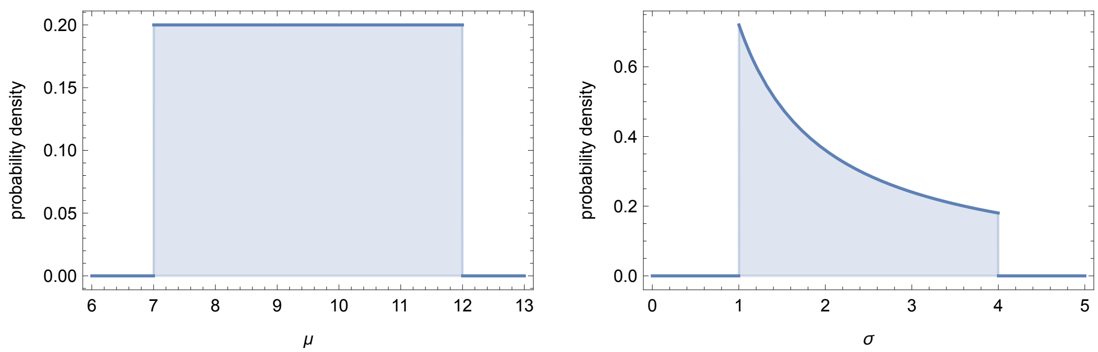

We define some basic statistics and plot the posterior marginal probability density functions:

In[ ]:=

```
xbarx = Mean[x]; sx = StandardDeviation[x]; nx = Length[x];
xbary = Mean[y]; sy = StandardDeviation[y]; ny = Length[y];
gμ = Plot[{Callout[postμ[μ, xbarx, sx, nx], "x", Above],
     Callout[postμ[μ, xbary, sy, ny], "y", Above]},
    {μ, μa, μb}, Axes → False, Frame → True, FrameLabel → {…}];
gσ =
  Plot[{Callout[postσ[σ, sx, nx], "x", Above], Callout[postσ[σ, sy, ny], "y", Above]},
    {σ, σa, σb}, Axes → False, Frame → True, FrameLabel → {…}];
GraphicsRow[{gμ, gσ}]
```

Out[ ]=

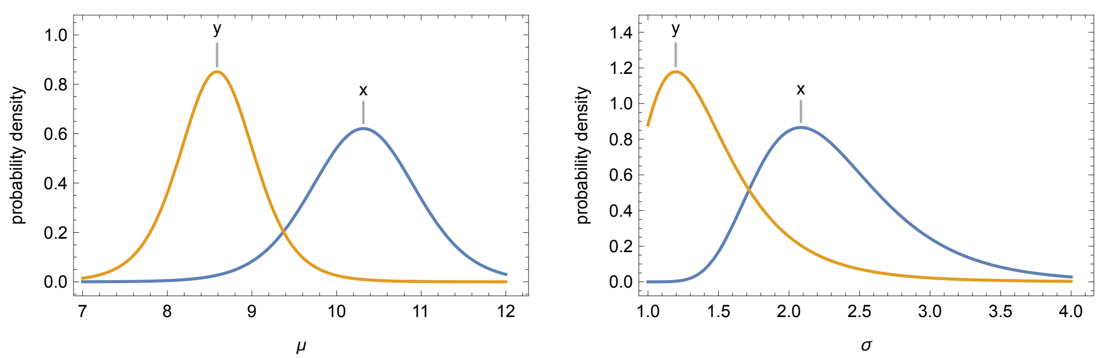

The following code will compute the posterior probability mass function, probability mass marginals and the evidence for a single data set:

```
In[]:= bayes2d[x_, μa_, μb_, Nμ_, σa_, σb_, Nσ_] :=
  Module[{xbar, s, n, dσ, dμ, logposterior, logterm,
    maxlog, dim, posterior, norm, logevidence, σmarginal, μmarginal},
   xbar = Mean[x]; s = StandardDeviation[x]; n = Length[x];
   dσ = (σb - σa)/(Nσ - 1); dμ = (μb - μa)/(Nμ - 1);
   logposterior = Table[
     logterm = Log[1/(μb - μa)] + Log[1/Log[σb / σa]] +
       n * Log[1/Sqrt[2 π]] - (n + 1) Log[σ] - ((n - 1) s^2 + n (μ - xbar)^2)/(2 σ^2);
     logterm, {σ, σa, σb, dσ}, {μ, μa, μb, dμ}];
   maxlog = Max[logposterior];
   dim = Dimensions[logposterior];
   posterior = Chop@Exp[ArrayReshape[Flatten[logposterior] - maxlog, dim]];
   norm = Total[Flatten[posterior]];
   posterior = posterior / norm;
   logevidence = Log[norm] + maxlog + Log[dμ * dσ];
   σmarginal = Map[Total, posterior];
   μmarginal = Map[Total, Transpose[posterior]];
   {μmarginal, σmarginal, posterior, logevidence}]
```

The evidence associated with data set *x* and our model together with our choice of priors is:

```
In[]:= Nμ = 150; Nσ = 160;
{μmarginal, σmarginal, posterior, logevidence} = bayes2d[x, μa, μb, Nμ, σa, σb, Nσ];

In[]:= logevidence
Out[]=
-27.9783
```

We can also evaluate the evidence by direct numerical integration of the unnormalized posterior:

```
In[]:= Log@NIntegrate[
  postUnormalized[μ, σ, xbarx, sx, nx, μa, μb, σa, σb], {μ, μa, μb}, {σ, σa, σb}]
Out[]=
-27.9791
```

If we widen the priors then the evidence will diminish unless we capture significantly more likelihood .

This can be seen by examining the following computation :

In[∘]:=
```
μA = μa / 2; μB = 2 μb; σA = σa / 2; σB = 2 σb;
Log@NIntegrate[
  postUnormalized[μ, σ, xbarx, sx, nx, μA, μB, σA, σB], {μ, μA, μB}, {σ, σA, σB}]
```
Out[∘]=
```
-30.0603
```

Our posterior probability density functions look like the following :

In[∘]:=
```
compx = bayes2d[x, μa, μb, Nμ, σa, σb, Nσ]; posteriorx = compx〚3〛;
compy = bayes2d[y, μa, μb, Nσ, σa, σb, Nσ]; posteriory = compy〚3〛;
g1 = ReliefPlot[posteriorx, PlotRange → All,
   LightingAngle → None, DataRange → {{μa, μb}, {σa, σb}}, AspectRatio → 1,
   FrameTicks → True, FrameLabel → {⋯}, PlotLabel → "…"];
g2 = ReliefPlot[posteriory, PlotRange → All,
  LightingAngle → None, DataRange → {{μa, μb}, {σa, σb}}, AspectRatio → 1,
  FrameTicks → True, FrameLabel → {⋯}, PlotLabel → "…"];
GraphicsRow[{g1, g2}, ImageSize → 500]
```
Out[∘]=

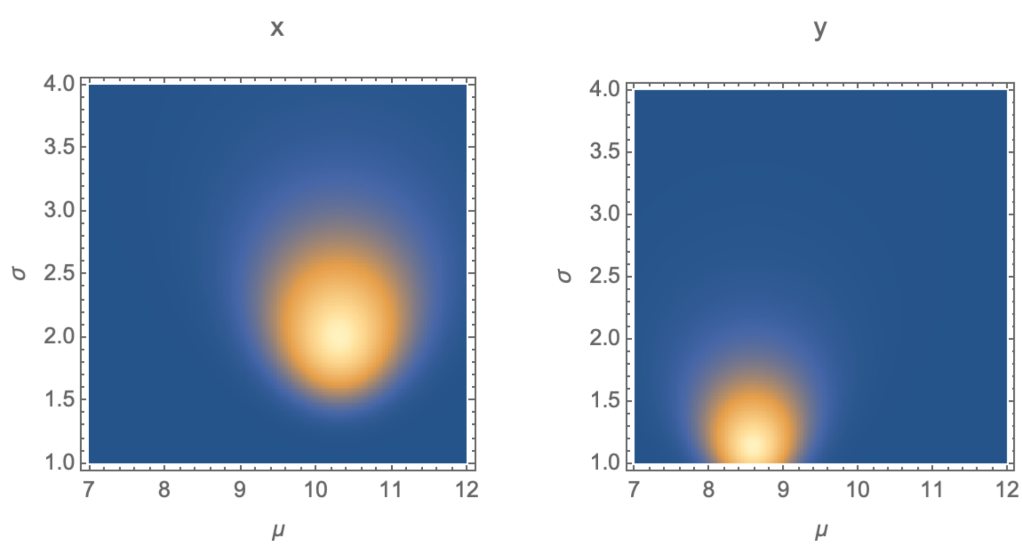

They are clearly different . But how different?

Analysis of two independent data samples

Suppose we have two data sets x and y. Let $\overline{x}$, s_x and m_x denote the sample mean, sample standard deviation and number of data samples in data set x. Similarly let $\overline{y}$, s_y and m_y denote the sample mean, sample standard deviation and number of data samples in data set y. If μ_x and σ_x denote the mean and standard deviation that data set x is drawn from and similarly if μ_y and σ_y denote the mean and standard deviation that data set y is drawn from, then the likelihood of obtaining data sets x and y jointly is proportions to the quantity

$$L(x, y \mid \mu_x, \mu_y, \sigma_x, \sigma_y) = L(\bar{x}, s_x, m_x \mid \mu_x, \sigma_x)\, L(\bar{y}, s_y, m_y \mid \mu_y, \sigma_y)$$

where

$$L(\bar{x}, s_x, m_x \mid \mu_x, \sigma_x) = \frac{1}{(2\pi)^{m_x/2}\, \sigma_x^{m_x}} \exp\left\{-\frac{1}{2\sigma_x^2}\left[(m_x - 1)\, s_x^2 + m_x(\mu_x - \bar{x})^2\right]\right\}$$

$$L(\bar{y}, s_y, m_y \mid \mu_y, \sigma_y) = \frac{1}{(2\pi)^{m_y/2}\, \sigma_y^{m_y}} \exp\left\{-\frac{1}{2\sigma_y^2}\left[(m_y - 1)\, s_y^2 + m_y(\mu_y - \bar{y})^2\right]\right\}.$$

In Mathematica terms this likelihood function is:

```
In[ ]:= Clear[μx, σx, μy, σy, xbar, sx, mx, ybar, sy, my];
L[μx_, σx_, μy_, σy_, xbar_, sx_, mx_, ybar_, sy_, my_] := (1/(Sqrt[2 π] σx))^mx
  Exp[-((mx - 1) sx^2 + mx (xbar - μx)^2)/(2 σx^2)] (1/(Sqrt[2 π] σy))^my Exp[-((my - 1) sy^2 + my (ybar - μy)^2)/(2 σy^2)]
```

We consider four hypotheses for the two data sets x and y.

H_1: Same means and standard deviations
H_2: Different means but same standard deviations
H_3: Same mean but different standard deviations
H_4: Different means and different standard deviations

The probability of H_j given the measured data D is

$$P(H_j \mid D) = \frac{P(D \mid H_j)\, P(H_j)}{P(D)} = \frac{P(D \mid H_j)\, P(H_j)}{\sum_{j=1}^{4} P(D \mid H_j)\, P(H_j)}$$

where $P(H_j)$ is the probability of a hypothesis on initial information. It is the case that

$$P(D \mid H_j) = E_j$$

where E_j is the Bayesian evidence in support of a hypothesis. In specific terms these evidences are as follows:

$$E_1 = \int_{\sigma_a}^{\sigma_b}\int_{\mu_a}^{\mu_b} f(\mu)\, g(\sigma)\; L(x, y \mid \mu_x, \mu_x, \sigma_x, \sigma_x)\; d\mu_x\, d\sigma_x$$

$$E_2 = \int_{\sigma_a}^{\sigma_b}\int_{\mu_a}^{\mu_b}\int_{\mu_a}^{\mu_b} f(\mu_x)\, f(\mu_y)\, g(\sigma)\, L(x, y \mid \mu_x, \mu_y, \sigma_x, \sigma_x)\, d\mu_x\, d\mu_y\, d\sigma_x$$

$$E_3 = \int_{\sigma_a}^{\sigma_b}\int_{\sigma_a}^{\sigma_b}\int_{\mu_a}^{\mu_b} f(\mu)\, g(\sigma_x)\, g(\sigma_y)\, L(x, y \mid \mu_x, \mu_x, \sigma_x, \sigma_y)\, d\mu_x\, d\sigma_x\, d\sigma_y$$

$$E_4 = \int_{\sigma_a}^{\sigma_b}\int_{\sigma_a}^{\sigma_b}\int_{\mu_a}^{\mu_b}\int_{\mu_a}^{\mu_b} f(\mu_x)\, f(\mu_y)\, g(\sigma_x)\, g(\sigma_y)\, L(x, y \mid \mu_x, \mu_y, \sigma_x, \sigma_y)\, d\mu_x\, d\mu_y\, d\sigma_x\, d\sigma_y$$

We will assume that the hypotheses have initially the same probability. Then

$$P(H_j \mid D) = \frac{E_j}{E_1 + E_2 + E_3 + E_4}.$$

In order to evaluate the evidence we need to compute integrals of dimensions 2, 3, 3 and 4. This can be done using a combination of analytic and numerical techniques. The approach we follow here is adapted from Gregory (2005).

Under hypothesis 1 the samples have the same means and standard deviations. In this case the joint likelihood of the data is:

In[ ]:=
```
likely1 = L[μx, σx, μx, σx, xbar, sx, mx, ybar, sy, my] /.
  {xbar → Mean[x], ybar → Mean[y], mx → Length[x], my → Length[y],
   sx → StandardDeviation[x], sy → StandardDeviation[y]}
```

Out[ ]=

$$\frac{e^{-\frac{11.4688+8\,(8.5875-\mu x)^2}{2\,\sigma x^2}-\frac{52.1367+12\,(10.3167-\mu x)^2}{2\,\sigma x^2}}}{1024\,\pi^{10}\,\sigma x^{20}}$$

The unnormalized posterior is:

In[ ]:=

$$\text{postU1} = \frac{1}{\mu b-\mu a}\;\frac{1}{\text{Log}[\sigma b/\sigma a]}\;\frac{1}{\sigma x}\;\text{likely1}$$

Out[ ]=

$$\frac{1.50444\times 10^{-9}\,e^{-\frac{11.4688+8\,(8.5875-\mu x)^2}{2\,\sigma x^2}-\frac{52.1367+12\,(10.3167-\mu x)^2}{2\,\sigma x^2}}}{\sigma x^{21}}$$

The *μx* integral in the evidence computation can be performed analytically. It is:

In[ ]:=
```
step1 = Integrate[postU1, {μx, μa, μb}]
```

Out[ ]=

$$\frac{e^{-\frac{38.9788}{\sigma x^2}}\left(4.2162\times 10^{-10}\,\text{Erf}\left[\frac{7.51041}{\sigma x}\right]+4.2162\times 10^{-10}\,\text{Erf}\left[\frac{8.30098}{\sigma x}\right]\right)}{\sigma x^{20}}$$

The evidence in support of hypothesis 1 can be found by integrating the results of step1 numerically :

In[ ]:=
```
Log@NIntegrate[step1, {σx, σa, σb}]
```

Out[ ]=

-44.6967

A combined algorithm for making all four computations is :

In[ ]:=
```
BayesianComparisonOfIndependentSamples[x_, y_, μlow_, μhigh_, σlow_, σhigh_] :=
 Module[{L1start, L2start, L3start, L4start, H1, H2, H3, H4,
   m, n, xbar, ybar, sx, sy, μa, μb, σa, σb, μx, μy, σx, σy, H, p},
  (* Case 1: Same means and standard deviations *)
  L1start = (e^(-((-1+m) sx^2+(-1+n) sy^2+m (xbar-μx)^2+n (ybar-μx)^2)/(2 σx^2)) (2 π)^(1/2 (-m-n)) (1/σx)^(1+m+n)) / ((-μa + μb) Log[σb/σa]);
  H1[μx_, σx_] := L1start /. {m → Length[x], n → Length[y],
     xbar → Mean[x], ybar → Mean[y], sx → StandardDeviation[x],
     sy → StandardDeviation[y], μa → μlow, μb → μhigh, σa → σlow, σb → σhigh};
  H1[σx_] := Integrate[H1[μx, σx], {μx, μlow, μhigh}];
  H[1] = NIntegrate[H1[σx], {σx, σlow, σhigh}];
 (* Case2: Different means but same standard deviations *)
  L2start = (e^(-((-1+m) sx^2+(-1+n) sy^2+m (xbar-μx)^2+n (ybar-μy)^2)/(2 σx^2)) (2 π)^(1/2 (-m-n)) (1/σx)^(1+m+n)) / ((μa - μb)^2 Log[σb/σa]);
```

```
  H2[μx_, μy_, σx_] := L2start /. {m → Length[x], n → Length[y],
     xbar → Mean[x], ybar → Mean[y], sx → StandardDeviation[x],
     sy → StandardDeviation[y], μa → μlow, μb → μhigh, σa → σlow, σb → σhigh};
  H2[σx_] := Integrate[H2[μx, μy, σx], {μx, μlow, μhigh}, {μy, μlow, μhigh}];
  H[2] = NIntegrate[H2[σx], {σx, σlow, σhigh}];
(* Case3: Same mean but different standard deviations *)
  L3start =
```

$$\left(e^{-\frac{(-1+m)\,sx^2+m\,(xbar-\mu x)^2}{2\,\sigma x^2}-\frac{(-1+n)\,sy^2+n\,(ybar-\mu x)^2}{2\,\sigma y^2}}\,(2\,\pi)^{\frac{1}{2}\,(-m-n)}\left(\frac{1}{\sigma x}\right)^{1+m}\left(\frac{1}{\sigma y}\right)^{1+n}\right)\Big/\left((-\mu a+\mu b)\,\mathrm{Log}\left[\frac{\sigma b}{\sigma a}\right]^2\right);$$

```
  H3[μx_, σx_, σy_] := L3start /. {m → Length[x], n → Length[y],
     xbar → Mean[x], ybar → Mean[y], sx → StandardDeviation[x],
     sy → StandardDeviation[y], μa → μlow, μb → μhigh, σa → σlow, σb → σhigh};
  H3[σx_, σy_] := Integrate[H3[μx, σx, σy], {μx, μlow, μhigh}];
  H[3] = NIntegrate[H3[σx, σy], {σx, σlow, σhigh}, {σy, σlow, σhigh}];

  (* Case4: Different means and different standard deviations *)
  L4start =
```

$$\left(e^{-\frac{(-1+m)\,sx^2+m\,(xbar-\mu x)^2}{2\,\sigma x^2}-\frac{(-1+n)\,sy^2+n\,(ybar-\mu y)^2}{2\,\sigma y^2}}\,(2\,\pi)^{-\frac{m}{2}-\frac{n}{2}}\left(\frac{1}{\sigma x}\right)^{1+m}\left(\frac{1}{\sigma y}\right)^{1+n}\right)\Big/\left((-\mu a+\mu b)^2\,\mathrm{Log}\left[\frac{\sigma b}{\sigma a}\right]^2\right);$$

```
  H4[μx_, μy_, σx_, σy_] := L4start /. {m → Length[x], n → Length[y],
     xbar → Mean[x], ybar → Mean[y], sx → StandardDeviation[x],
     sy → StandardDeviation[y], μa → μlow, μb → μhigh, σa → σlow, σb → σhigh};
  H4[σx_, σy_] := Integrate[H4[μx, μy, σx, σy], {μx, μlow, μhigh}, {μy, μlow, μhigh}];
  H[4] = NIntegrate[H4[σx, σy], {σx, σlow, σhigh}, {σy, σlow, σhigh}];
  p = {H[1], H[2], H[3], H[4]};
  p = p / Apply[Plus, p];
  (* Return {Prob[Case1],Prob[Case 2],Prob[Case 3],Prob[Case 4]} *)
  p]
```

The algorithm returns probabilities of the four hypotheses.

The probabilities of the four hypotheses in regards to our joint data set are:

In[ ]:=
```
{p1, p2, p3, p4} = BayesianComparisonOfIndependentSamples[x, y, μa, μb, σa, σb]
```

Out[ ]=
```
{0.0975786, 0.289153, 0.14429, 0.468979}
```

The probability of a different mean ($p_{\text{dif}} = p_2 + p_4$) is:

In[ ]:=
```
pdif = p2 + p4
```

Out[ ]=
```
0.758132
```

The odds of a different mean is:

In[ ]:=
```
pdif / (1 - pdif)
```
Out[ ]=
```
3.13448
```

We are now in a position to perform some interesting simulations We begin by generating two samples with different means and standard deviations:

In[ ]:=
```
μxtrue = 9; μytrue = 10.5; σxtrue = 3; σytrue = 1.5; Nsample = 40;
SeedRandom[98 345];
X = RandomVariate[NormalDistribution[μxtrue, σxtrue], {Nsample}];
Y = RandomVariate[NormalDistribution[μytrue, σytrue], {Nsample}];
```

The marginals are assuming wide priors :

In[ ]:=
```
xbarX = Mean[X]; sX = StandardDeviation[X]; nX = Length[X];
xbarY = Mean[Y]; sY = StandardDeviation[Y]; nY = Length[Y];
gμ = Plot[{Callout[postμ[μ, xbarX, sX, nX], "X", Above],
     Callout[postμ[μ, xbarY, sY, nY], "Y", Above]},
    {μ, μa, μb}, Axes → False, Frame → True, FrameLabel → {⋯}];
gσ =
  Plot[{Callout[postσ[σ, sX, nX], "X", Above], Callout[postσ[σ, sY, nY], "Y", Above]},
   {σ, σa, σb}, Axes → False, Frame → True, FrameLabel → {⋯}];
GraphicsRow[{gμ, gσ}]
```

Out[ ]=

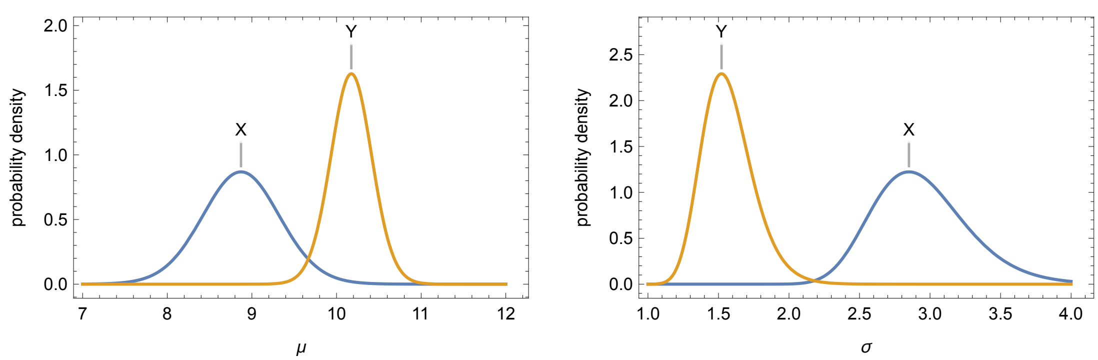

In[∘]:=
```
compx = bayes2d[X, μa, μb, Nμ, σa, σb, Nσ]; posteriorx = compx[[3]];
compy = bayes2d[Y, μa, μb, Nσ, σa, σb, Nσ]; posteriory = compy[[3]];
g1 = ReliefPlot[posteriorx, PlotRange → All,
    LightingAngle → None, DataRange → {{μa, μb}, {σa, σb}}, AspectRatio → 1,
    FrameTicks → True, FrameLabel → {⋯}, PlotLabel → "..."];
g2 = ReliefPlot[posteriory, PlotRange → All,
    LightingAngle → None, DataRange → {{μa, μb}, {σa, σb}}, AspectRatio → 1,
    FrameTicks → True, FrameLabel → {⋯}, PlotLabel → "..."];
GraphicsRow[{g1, g2}, ImageSize → 500]
```

Out[∘]=

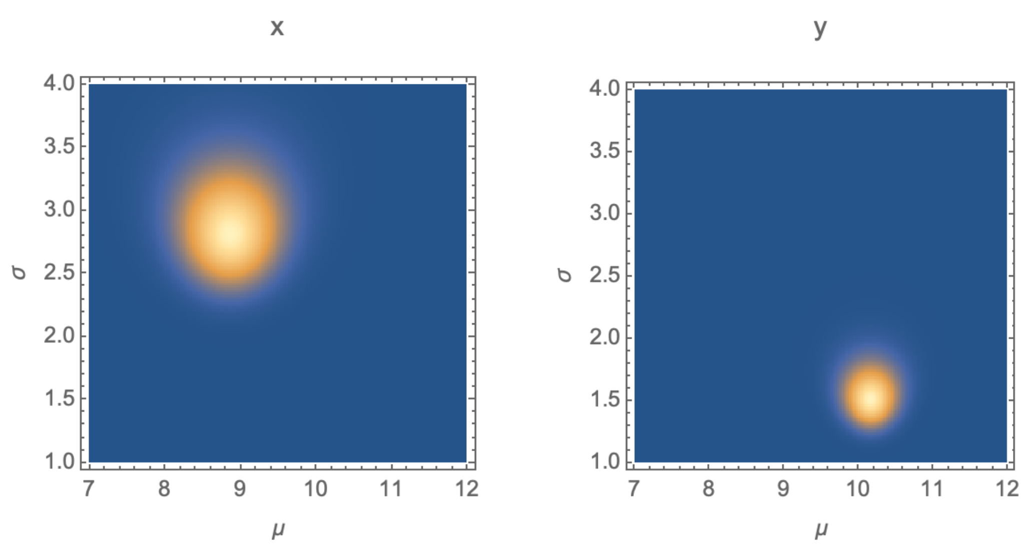

The probabilities of the four hypotheses in regards to our joint data simulated data set set (*X*, *Y*)are:

In[∘]:=
```
{p1, p2, p3, p4} = BayesianComparisonOfIndependentSamples[X, Y, μa, μb, σa, σb]
```

Out[∘]=

```
{0.000367377, 0.00210244, 0.155639, 0.841891}
```

The probability of the hypothesis different means and different standard deviations (H_4) is given by p_4:

In[∘]:=
```
p4
```

Out[∘]=

```
0.841891
```

The probability of different means is:

In[∘]:=
```
p2 + p4
```

Out[∘]=

```
0.843993
```

The probability of different standard deviations is:

In[ ]:=
```
p3 + p4
```
Out[ ]=
```
0.99753
```

References

Our presentation of the analysis of a single data sample drawn from a Gaussian distribution is based upon the work of Sir Harold Jeffreys. Our analysis of a pair of samples is adapted from the work of Phil Gregory.

Gregory, Phil (2005), *Bayesian Logical Data Analysis for the Physical Sciences*, Cambridge University Press.

Jeffreys, Sir Harold (1973), *Scientific Inference*, Third Edition, Cambridge University Press.