## Estimation

## Estimation

- All probability models depend on parameters. E.g.,

Binomial depends on probability of success $\pi$. Normal depends on mean $\mu$, standard deviation $\sigma$.

- Parameters are properties of the "population" and are typically unknown.
- The process of taking a sample of data to make inferences about these parameters is referred to as "estimation".
- There are a number of different estimation methods ... we will study two estimation methods:

Maximum likelihood (ML)
Bayes


1


## Maximum Likelihood

Fisher (1922) invented this general method.
Problem: Unknown model parameters, $\theta$.
Set-up: Write the probability of the data, $\boldsymbol{Y}$, in terms of the model parameter and the data, $P(Y, \theta)$.

Solution: Choose as your estimate the value of the unknown parameter that makes your data look as likely as possible. Pick $\hat{\theta}$ that maximizes the probability of the observed data.

The estimator $\hat{\theta}$ is called the maximum likelihood estimator (MLE).

$$
y \sim N(\mu, 1)
$$

$$
y=10
$$



## Maximum Likelihood - Example

Data: $Y_{i}=0 / 1$ for $i=1,2, \ldots . n$ (independent)
Model: $Z=\sum_{i} Y_{i} \sim \operatorname{Binomial(n)}$
Probability: Let's fix the number in the sample at $n=20$. The resulting model for Z is
Binomial with size 20 and success probability $\pi$.
The probability distribution function is:

$$
P(Z ; \pi)=\binom{20}{Z} \pi^{Z}(1-\pi)^{(20-Z)}
$$

where $Z$ is the variable and $\pi$ is fixed.

The likelihood function is the same function:

$$
L(\pi ; Z)=\binom{20}{(Z)} \pi^{Z}(1-\pi)^{(20-Z)}
$$

except now $\pi$ is the variable and $\mathbf{Z}$ is fixed.

## Maximum Likelihood - Example

Two ways to look at this:

- Fix $\pi$ and look at the probability of different values of $Z$ :

- Fix $Z$ and look at the probabiitity under different values of $\pi$ (this is called the likelihood function):

| $Z=3$ |  |
| :---: | :---: |
| $\pi$ |  |
| 0.01 |  |
| 0.05 |  |
| 0.10 |  |

## Maximum Likelihood - Example

If you observe the data $Z=3$ then the likelihood function is shown in the plots below:
$P(Z=3)$ as function of $\pi$

$\log P(Z=3)$ as function of $\pi$



## Maximum Likelihood - Example

- We can use elementary calculus (an oxymoron?) to find the maximum of the (log) likelihood function:

$$
\begin{aligned}
\frac{d \log L}{d \pi} & =0 \\
\frac{d}{d \pi} \frac{Z \log \pi+(20-Z) \log (1-\pi)}{\frac{Z}{\pi}-\frac{(20-Z)}{1-\pi}} & =0 \\
\hat{\pi} & =\frac{Z}{20}
\end{aligned}
$$

- Not surprisingly, the likelihood in this example is maximized at the observed proportion, 3/20.
- Sometimes (e.g. this example) the MLE has a simple closed form. In more complex problems, numerical optimization is used.
- Computers can find these maximum values!


## Maximum Likelihood - Notation

$\mathrm{L}(\theta)=$ Likelihood as a function of the unknown parameter, $\theta$.
$l(\theta)=\log (\mathrm{L}(\theta))$, the $\log$-likelihood.
Usually more convenient to work with analytically and numerically.
$\mathrm{S}(\theta)=\mathrm{d} l(\theta) / \mathrm{d} \theta=$ the "score".
Set $\mathrm{d} l(\theta) / \mathrm{d} \theta=0$ and solve for $\theta$ to find the MLE.
$\mathrm{I}(\theta)=-\mathrm{d}^{2} l(\theta) / \mathrm{d} \theta^{2}=$ the "information".
If evaluated at the MLE, then $-\mathrm{d}^{2} l(\theta) / \mathrm{d} \theta^{2}$ is referred to as the observed information; $\mathrm{E}\left(-\mathrm{d}^{2} l(\theta) / \mathrm{d} \theta^{2}\right)$ is referred to as the expected or Fisher information.

## $\operatorname{Var}(\theta)=\mathrm{I}^{-1}(\theta)$ (in most cases)

## Maximum Likelihood - Example

$$
\begin{aligned}
& L(\pi)=\binom{20}{Z} \pi^{Z_{(1-\pi)}}{ }^{(20-Z)} \\
& 1(\pi)=Z \log (\pi)+(20-Z) \log (1-\pi) \\
& S(\pi)=\frac{Z}{\pi}-\frac{(20-Z)}{1-\pi} \Rightarrow \hat{\pi}=\frac{Z}{20} \\
& I(\pi)=\frac{Z}{\pi^{2}}+\frac{(20-Z)}{(1-\pi)^{2}} \\
& E(I(\pi))=\frac{20 \pi}{\pi^{2}}+\frac{(20-20 \pi)}{(1-\pi)^{2}} \\
& \left.=\frac{20}{\pi(1-\pi)}\right) \mathrm{n} \\
& \operatorname{Var}(\hat{\Pi})=\frac{1}{I}=
\end{aligned}
$$

(note: constant dropped from I $(\pi)$ )

## Numerical Optimization

- In complex problems it may not be possible to find the MLE analytically; in that case we use numerical optimization to search for the value of $\theta$ that maximizes the likelihood
- A common problem with maximum likelihood estimation is accidentally finding a local maximum instead of a global one; solution is to try multiple starting values



## Comments:

always based on a probability model for the data.

- Maximum likelihood is the "best" method of estimation for any situation that you are willing to write down a probability model (so generally does not apply to nonparametric problems).
- Maximum likelihood can be used even when there are multiple unknown parameters, in which case $\theta$ are multiple unknown parameters, in which case $\theta$

$$
\text { (ie. } \theta_{0}, \theta_{1}, \mathrm{~K}, \theta_{p} \text { ). }
$$

has several components
single most likely value of $\theta$ ). In lecture 5 we will

- The MLE is a "point estimate" (i.e gives the single most likely value of $\theta$ ). In lecture 5 we will learn about interval estimates, which describe a range of values which are likely to include the true value of $\theta$. We combine the MLE and $\operatorname{Var}(\theta)$ to generate these intervals.


## Model Comparisons

Q: Suppose we have two alternative models for the data; in each case we use maximum likelihood to estimate the parameters. How do we decide which model fits the data "better"?
Q: Suppose we have two alternative models for the data; in each case we use maximum
likelihood to estimate the parameters. How do we decide which model fits the data "better"?

A: First thought - compare the likelihoods.

A common approach is to "penalize" the
-
$\Leftrightarrow$
complex model.

- How to choose?
as a special case (1 parameter) of a penalized likelihood.


## Example - LOD scores

Suppose we have a sample of size N gametes in which the number of recombinants ( R ) and nonrecombinants ( $\mathrm{N}-\mathrm{R}$ ) for two loci can be counted. Let $\theta$ be the recombination fraction between the two loci. Then the probability of the data can be modeled using the binomial distribution:


The situation of no linkage corresponds to $\theta=0.5$, so we can express the models as
$\rightarrow$ Model 1: $\theta=0.5$

$\rightarrow$ Model 2: $\theta$ anywhere between 0 and 0.5

## Example - LOD scores

Model 1: The situation of no linkage corresponds to $\theta=0.5$. If we substitute this into the likelihood equation, we get

$$
\begin{aligned}
\log _{10} L_{1} & =R \log _{10} 0.5+(N-R) \log _{10} 0.5 \\
& =N \log _{10} 05
\end{aligned}
$$

This model has 0 (free) parameters.
Model 2: The log-likelihood when $\theta$ is unrestricted is

$$
\log _{10} L_{2}=R \log _{10} \theta+(N-R) \log _{10}(1-\theta)
$$

this model has 1 parameter.
Taking the derivative and solving for $\theta$ gives

$$
\hat{\theta}=\frac{R}{N}
$$

If we substitute this back into the log-likelingod, we get ...

$$
\log _{10} L_{2}=R \log _{10} R / N^{+(N-R) \log _{10}(1-R / N)}
$$

## Example - LOD scores

The LOD score is

$$
\begin{aligned}
\mathrm{LOD} & =\left(\log _{10} \mathrm{~L}_{2}-\log _{10} \mathrm{~L}_{1}\right) \\
& =R \log _{10}\left(\frac{R}{N-R}\right)+N \log _{10}\left(\frac{N-R}{0.5 N}\right.
\end{aligned}
$$

Large values of the LOD score ( $>3$ ) are considered evidence of linkage
(i.e. the penalty is 3 ).
(As we will see, this is a pretty big hurdle to overcome.)

## Example - LOD scores

E.g. $\mathrm{N}=50$ and $\mathrm{R}=18$
$\hat{\boldsymbol{\theta}}=18 / 50=36 \%$
$\log _{10} \mathrm{~L}_{1}=-15.0$
$\log _{10} \mathrm{~L}_{2}=-14.2$
LOD $=-14.2-(-15.0)=0.8$
$\Rightarrow$ No evidence of linkage; conclude $\theta=.5$

## Model Comparisons - AIC, BIC

AIC - Akaike's Information Criterion
BIC - Bayes Information Criterion

$$
\begin{aligned}
& \mathrm{AIC}=2 \ell(\theta)-2 \mathrm{~K} \\
& \mathrm{BIC}=2 \ell(\theta)-4 \log (\mathrm{n}) \\
& \mathrm{k}=\# \text { parameters }
\end{aligned}
$$

- Use to compare a series of models. Pick the model with the largest AIC or BIC
- Larger model $\Rightarrow$ larger likelihood (typically)
- Therefore, "penalize" the likelihood for each added parameter
- AIC tries to find the model that would have the minimum prediction error on a new set of data.
- BIC tries to find the model with the highest "posterior probability" given the data
- Typically, BIC is more conservative (picks smaller models)


## Model Comparisons - AIC, BIC

Example - Recombinant ( $\mathrm{N}=50, \mathrm{R}=18$ )
$\log (\mathrm{L} 1)=-34.66$
$\theta=.5$
$\log (\mathrm{L} 2)=-32.67 \quad \theta \underset{h_{e} \perp \text { ne en }}{\text { (natural logs now) }} 0-.5$
Lb: $\theta=.5$ して: $\theta$ arb
ATC $-2 * 34.66=\mathbf{- 6 9 . 3 2} \quad-2 * 32.67-2 \quad=\underline{\mathbf{6 7 . 3 4}}$
BIC $-2 * 34.66=\mathbf{- 6 9 . 3 2}-2 * 32.67-\log (50)=\mathbf{- 6 9 . 2 5}$
$\mathrm{AIC} \Rightarrow$ pick $\theta=.36$ $\theta=\frac{18}{50}=.3$
$\mathrm{BIC} \Rightarrow$ pick $\theta=.36$ ( but almost tied)

## Bayes Estimation

Recall Bayes theorem (written in terms of data X and parameter $\theta$ ):

Notice the change in perspective - $\theta$ is now treated as a random variable instead of a fixed number.
$\mathrm{P}(\mathrm{X} \mid \theta)$ is the likelihood function, as before.
$\mathrm{P}(\theta)$ is called the prior distribution of $\theta$.
$\mathrm{P}(\theta \mid \mathrm{X})$ is called the posterior distribution of $\theta$
Based on $P(\theta \mid X)$ we can define a number of possible estimators of $\theta$. A commonly used estimate is the maximum a posteriori (MAP) estimate:

$$
\hat{\theta}_{\mathrm{MAP}}=\max _{\theta} \cdot(\theta \mid \mathrm{X})
$$

We can also use $P(\theta \mid X)$ to define "credible" intervals for $\theta$.

## Bayes Estimation

## Comments:

- The MAP estimator is a very simple Bayes estimator. More generally, Bayes estimators minimize a "loss function" - a penalty based on how far $\theta$ is from $\theta$ (e.g. Loss $=(\theta-\theta) \uparrow 2$ ).
- The Bayesian procedure provides a convenient way of combining external information or previous data (through the prior distribution) with the current data (through the likelihood) to create a new estimate.
- As N increases, the data (through the likelihood) overwhelms the prior and Bayes estimator typically converges to the MLE
- Controversy arises when $\mathrm{P}(\theta)$ is used to incorporate subjective beliefs or opinions.
- If the prior distribution $P(\theta)$ is simply that $\theta$ is uniformly distributed over all possible values, this is called an "uninformative" prior, and the MAP is the same as the MLE.


## Bayes Estimation

## Example

Suppose a man is known to have transmitted allele A1 to his child at a locus that has only two alleles: A1 and A2. What is his most likely genotype?

Soln. Let X represent the paternal allele in the child and let $\theta$ represent the man's genotype:

$$
\begin{aligned}
\mathrm{X} & =\mathrm{A} 1 \\
\theta & =\{\mathrm{A} 1 \mathrm{~A} 1, \mathrm{~A} 1 \mathrm{~A} 2, \mathrm{~A} 2 \mathrm{~A} 2\}
\end{aligned}
$$

We can write the likelihood function as:

$$
\begin{align*}
& \mathrm{P}(\mathrm{X} \mid \theta=\mathrm{A} 1 \mathrm{~A} 1)=1 \\
& \mathrm{P}(\mathrm{X} \mid \theta=\mathrm{A} 1 \mathrm{~A} 2)=.5 \\
& \mathrm{P}(\mathrm{X} \mid \theta=\mathrm{A} 2 \mathrm{~A} 2)=0
\end{align*}
$$

Therefore, the MLE is $\theta=\mathrm{A} 1 \mathrm{~A} 1$.

## Bayes Estimation

Suppose, however, that we know that the frequency of the A1 allele in the general population is only $1 \%$. Assuming HW equilibrium we have

$$
\begin{aligned}
& \mathrm{P}(\theta=\mathrm{A} 1 \mathrm{~A} 1)=.0001 \\
& \mathrm{P}(\theta=\mathrm{A} 1 \mathrm{~A} 2)=.0198 \\
& \mathrm{P}(\theta=\mathrm{A} 2 \mathrm{~A} 2)=.9801
\end{aligned}
$$

$\overbrace{}^{*} p^{r}$

This leads to the posterior distribution

$$
\begin{aligned}
& \mathrm{P}(\theta=\mathrm{A} 1 \mathrm{~A} 1 \mid \mathrm{X}) \\
& \quad=\mathrm{P}(\mathrm{X} \mid \theta=\mathrm{A} 1 \mathrm{~A} 1) \mathrm{P}(\theta=\mathrm{A} 1 \mathrm{~A} 1) / \mathrm{P}(\mathrm{X}) \\
& \quad=1^{*} .0001 / .01=.01 \\
& \mathrm{P}(\theta=\mathrm{A} 1 \mathrm{~A} 2 \mid \mathrm{X}) \\
& \quad=\mathrm{P}(\mathrm{X} \mid \theta=\mathrm{A} 1 \mathrm{~A} 2) \mathrm{P}(\theta=\mathrm{A} 1 \mathrm{~A} 2) / \mathrm{P}(\mathrm{X}) \\
& \quad=.5^{*} .0198 / .01=.99 \\
& \mathrm{P}(\theta=\mathrm{A} 2 \mathrm{~A} 2 \mid \mathrm{X})=0
\end{aligned}
$$

So the Bayesian MAP estimator is $\theta=\mathrm{A} 1 \mathrm{~A} 2$.

## Exercise: redo assuming the man has 2 children who both have the A1 paternal allele.

## Summary

- Maximum likelihood is a method of estimating parameters from data
- ML requires you to write a probability model for the data
- MLE's may be found analytically or numerically
- (Inverse of the negative of the) second derivative of the log-likelihood gives variance of estimates
- Comparison of log-likelihoods allows us to choose between alternative models
- Bayesian procedures allow us to incorporate additional information about the parameters in the form of prior data, external information or personal beliefs.


## Problem 1

Suppose we are interested in estimating the recombination fraction, $\theta$, from the following experiment. We do a series of crosses: $\mathrm{AB} / \mathrm{ab} x$ $\mathrm{AB} / \mathrm{ab}$ and measure the frequency of the various phases in the gametes (assume we can do this). If the recombination fraction is $\theta$ then we expect the following probabilities (sorry, I can't explain these...):
phase probability (*4)
$\mathrm{AB} \quad 3-2 \theta+\theta^{2}$
$\mathrm{Ab} \quad 2 \theta-\theta^{2}$
$\mathrm{aB} \quad 2 \theta-\theta^{2}$
ab
$1-2 \theta+\theta^{2}$
Suppose we observe $(A B, A b, a B, a a)=(125,18,20,34)$. Use maximum likelihood to estimate $\theta$.

## Solution to problem 1

$$
\begin{aligned}
& \operatorname{Pr}(\text { data } \mid \theta) \propto\left(3-2 \theta+\theta^{2}\right)^{\mathrm{AB}}\left(2 \theta-\theta^{2}\right)^{\mathrm{Ab}}\left(2 \theta-\theta^{2}\right)^{\mathrm{aB}}\left(1-2 \theta+\theta^{2}\right)^{\mathrm{ab}} \\
& l(\theta)=\mathrm{AB} \log \left(3-2 \theta+\theta^{2}\right)+(\mathrm{Ab}+\mathrm{aB}) \log \left(2 \theta-\theta^{2}\right)+\mathrm{ab} \log \left(1-2 \theta+\theta^{2}\right) \\
& \frac{d \mathrm{l}(\theta)}{d \theta}=\frac{2 A B(\theta-1)}{3-2 \theta+\theta^{2}}+\frac{2(A b+a B)(1-\theta)}{2 \theta-\theta^{2}}+\frac{2 a b(\theta-1)}{1-2 \theta+\theta^{2}}=0
\end{aligned}
$$

## Numerical solution gives $\boldsymbol{\theta}=.21$

$$
\frac{d^{2} l(\theta)}{d \theta^{2}}=\frac{A B\left(1+2 \theta-\theta^{2}\right)}{\left[3-2 \theta+\theta^{2}\right]^{2}}-\frac{(A b+a B)}{\theta^{2}}-\frac{a b}{(1-\theta)^{2}}
$$

$$
\begin{aligned}
& I=E(-d \uparrow 2 \ell(\theta) / d \theta \uparrow 2)=-\mathrm{N} *(1+2 \theta-\theta \uparrow 2 / 3-2 \theta+\theta \uparrow \\
& =\mathrm{N}^{*} 16.6
\end{aligned}
$$

$\operatorname{Var}(\theta)=1 / 213.6=.00468$

## Problem 2

Every human being can be classified into one of four blood groups: O, $\mathrm{A}, \mathrm{B}, \mathrm{AB}$. Inheritance of these blood groups is controlled by 1 gene with 3 alleles: $\mathrm{O}, \mathrm{A}$ and B where O is recessive to A and B . Suppose the frequency of these alleles is $r, p$, and $q$, respectively $(p+q+r=1)$. If we observe $(\mathrm{O}, \mathrm{A}, \mathrm{B}, \mathrm{AB})=(176,182,60,17)$ use maximum likelihood to estimate $\mathrm{r}, \mathrm{p}$ and q .

## Solution to problem 2

First, we use basic genetics to find the probability of the observed phenotypes in terms of the unknown parameters. Assuming random mating, we have:

| Genotype | prob. | Phenotype | prob. |
| :--- | :--- | :--- | :--- |
| OO | $\mathrm{r}^{2}$ | O | $\mathrm{r}^{2}$ |
| AA | $\mathrm{p}^{2}$ |  |  |
| AO | 2 pr | A | $\mathrm{p}^{2}+2 \mathrm{pr}$ |
| BB | $\mathrm{q}^{2}$ |  |  |
| BO | 2 qr | B | $\mathrm{q}^{2}+2 \mathrm{qr}$ |
| AB | 2 pq | AB | 2 pq |

$\operatorname{Pr}($ data $\mid \theta) \propto\left(\mathrm{r}^{2}\right)^{\mathrm{O}}\left(\mathrm{p}^{2}+2 \mathrm{pr}\right)^{\mathrm{A}}\left(\mathrm{q}^{2}+2 \mathrm{qr}\right)^{\mathrm{B}}(2 \mathrm{pq})^{\mathrm{AB}}$
$l(p, q, r)=2 \mathrm{O} \log (\mathrm{r})+\mathrm{A} \log \left(\mathrm{p}^{2}+2 \mathrm{pr}\right)+\mathrm{B} \log \left(\mathrm{q}^{2}+2 \mathrm{qr}\right)+\mathrm{AB} \log (\mathrm{p})+\mathrm{AB} \log (\mathrm{q})$
To estimate $\mathrm{p}, \mathrm{q}$ and r , we need to maximize $l(p, q, r)$ subject to the constraint $\mathrm{p}+\mathrm{q}+\mathrm{r}=1$. This constraint makes the problem a bit harder .... one approach is to just put $\mathrm{r}=1-\mathrm{p}-\mathrm{q}$ in the likelihood so we have just 2 parameters ... p and q. Then

$$
\begin{gathered}
d l / d p=-2 O / r+2 A r / p(2 r+p)-2 B q / q(2 r+q)+A \beta / p \\
d l / d q=-2 O / r-2 A p / p(2 r+p)+2 B r / q(2 r+q)+A
\end{gathered}
$$

For $(\mathrm{O}, \mathrm{A}, \mathrm{B}, \mathrm{AB})=(176,182,60,17)$, this gives

$$
\mathrm{p}=.264 \quad \mathrm{q}=.093 \quad \mathrm{r}=.642
$$

Further analysis would take $2^{\text {nd }}$ derivatives to find the information and, therefore, the variances of the estimates.

## Problem 3

Suppose we have the following simple pedigree.


Define the phenotype of person i as $\mathrm{H}_{\mathrm{i}}$ and the genotype as $\mathrm{G}_{\mathrm{iH}}$ How can we use maximum likelihood to estimate parameters of the penetrance function, $\operatorname{Pr}(\mathrm{H} \mid \mathrm{G} ; \theta)$ ?

## Solution to problem 3

- If we knew all the genotypes the problem would be "easy". We would simply write down the log-likelihood and maximize it numerically or analytically:

$$
l(\theta)=\sum_{i} \log \operatorname{Pr}\left(H_{i} \mid G_{i}\right)
$$

- If we don't know the genotypes (only data are the phenotypes), then we must maximize

$$
l(\theta)=\log \operatorname{Pr}(H)
$$

where H represents the collection of all 6 phenotypes. The general idea is to use the total probability rule to write

$$
\begin{aligned}
\operatorname{Pr}(H) & =\sum_{G} \operatorname{Pr}(H \mid G) \operatorname{Pr}(G) \\
& =\sum_{G_{1}, G_{2}, G_{3}, G_{4}, G_{5}, G_{6}}\left\{\prod_{i} \operatorname{Pr}\left(H_{i} \mid G_{i}\right)\right\} \operatorname{Pr}\left(G_{1}, G_{2}, G_{3}, G_{4}, G_{5}, G_{6}\right)
\end{aligned}
$$

Further simplification is achieved by writing

$$
\begin{gathered}
\operatorname{Pr}\left(G_{1}, G_{2}, G_{3}, G_{4}, G_{5}, G_{6}\right)=\operatorname{Pr}\left(G_{6} \mid G_{1}, G_{2}, G_{3}, G_{4}, G_{5}\right) \operatorname{Pr}\left(G_{5} \mid G_{1}, G_{2}, G_{3}, G_{4}\right) \operatorname{Pr}\left(G_{4} \mid G_{1}, G_{2}, G_{3}, G_{4}\right) \times \\
\operatorname{Pr}\left(G_{3} \mid G_{1}, G_{2}\right) \operatorname{Pr}\left(G_{2} \mid G_{1}\right) \operatorname{Pr}\left(G_{1}\right)
\end{gathered}
$$

Since the genotype of each individual is determined only by his/her parents

$$
\operatorname{Pr}\left(G_{1}, G_{2}, G_{3}, G_{4}, G_{5}, G_{6}\right)=\operatorname{Pr}\left(G_{6} \mid G_{3}, G_{4}\right) \operatorname{Pr}\left(G_{5} \mid G_{1}, G_{2}\right) \operatorname{Pr}\left(G_{4} \mid G_{1}, G_{2}\right) \operatorname{Pr}\left(G_{3}\right) \operatorname{Pr}\left(G_{2}\right) \operatorname{Pr}\left(G_{1}\right)
$$

Given the inheritance probabilities $\left(\operatorname{Pr}\left(\mathrm{G}_{\mathrm{i}} \mid \mathrm{G}_{\mathrm{j}}, \mathrm{G}_{\mathrm{k}}\right)\right)$ and population frequencies of the genotypes $\left(\operatorname{Pr}\left(\mathrm{G}_{\mathrm{i}}\right)\right.$, we have a fully specified model and can maximize the likelihood using a computer.

## Problem 4

Suppose we wish to estimate the recombination fraction for a particular locus. We observe $\mathrm{N}=50$ and $\mathrm{R}=18$. Several previously published studies of the recombination fraction in nearby loci (that we believe should have similar recombination fractions) have shown recombination fractions between .22 and .44 . We decide to model this prior information as a beta distribution (see
http://en.wikipedia.org/wiki/Beta distribution) with parameters $\mathrm{a}=19$ and $b=40$ :


Find the MLE and Bayesian MAP estimators of the recombination fraction. Also find a $95 \%$ confidence interval (for the MLE) and a 95\% credible interval (for the MAP)

## Solution to problem 4

The data follow a binomial distribution with $\mathrm{N}=50, \mathrm{R}=18$ and the prior information is captured by a beta distribution with parameters $a=19, b=40$ :

$$
\begin{aligned}
& P(\theta)=\frac{\Gamma(a+b)}{\Gamma(a) \Gamma(b)} \theta^{a-1}(1-\theta)^{b-1} \\
& P(X \mid \theta)=\frac{N!}{R!(N-R)!} \theta^{R}(1-\theta)^{N-R}
\end{aligned}
$$

Working through Bayes theorem, we find ...

$$
P(\theta \mid X)=\frac{\Gamma(N+a+b)}{\Gamma(a+R) \Gamma(N-R+b)} \theta^{a+R-1}(1-\theta)^{N-R+b-1}
$$

which is another beta distribution with parameters $(a+R)$ and ( $\mathrm{N}-$ $\mathrm{R}+\mathrm{b})$. The mode of the beta distribution with parameters $\alpha$ and $\beta$ is $(\alpha-1) /(\alpha+\beta-2)$ so

$$
\hat{\theta}_{M A P}=\frac{a+R-1}{N+a+b-2}=\frac{36}{107}=.336
$$

Also, we can find the $2.5^{\text {th }}$ and $97.5^{\text {th }}$ percentiles of the posterior distribution ( $95 \%$ credible interval): [.23-.40]

For comparison the MLE is $18 / 50=0.36$ with a $95 \%$ confidence interval of [.23-.49]

