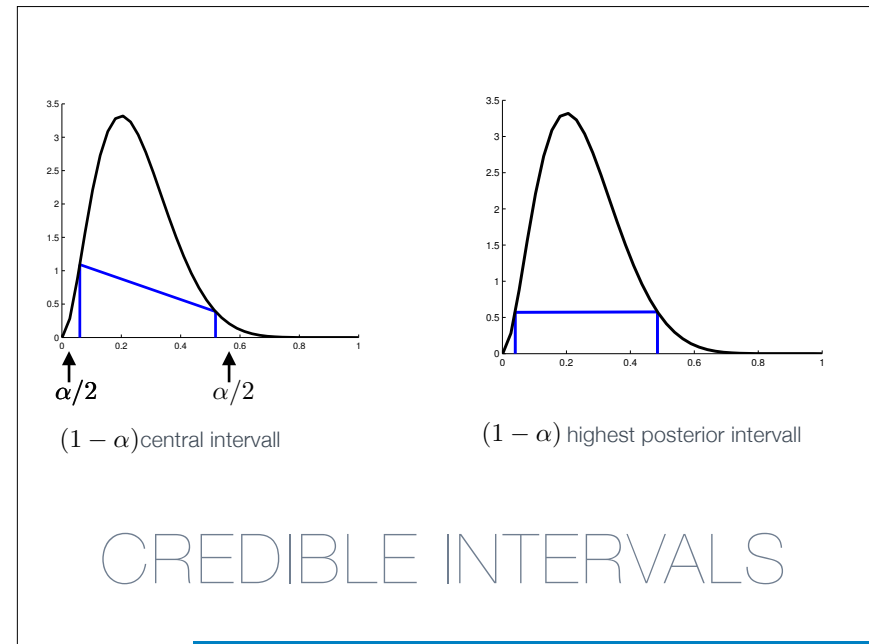
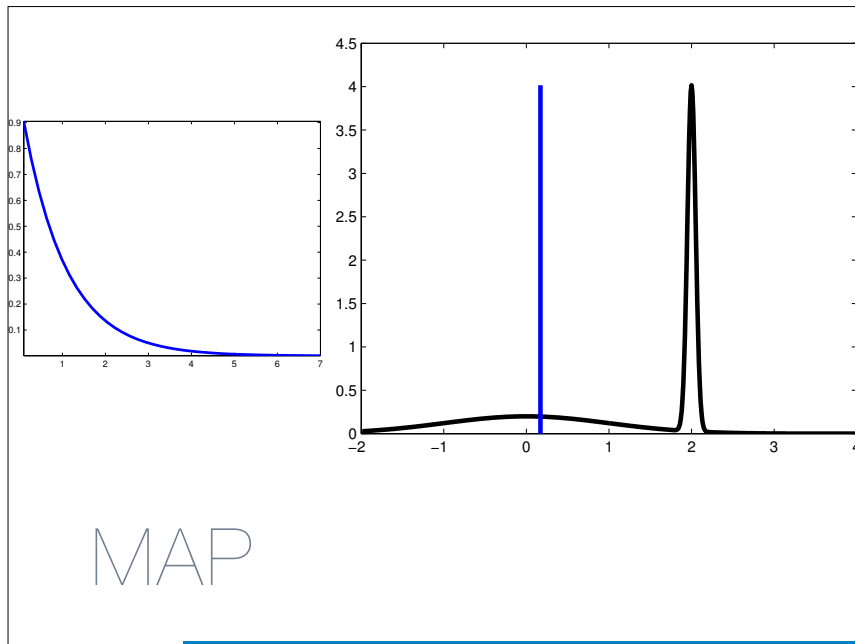
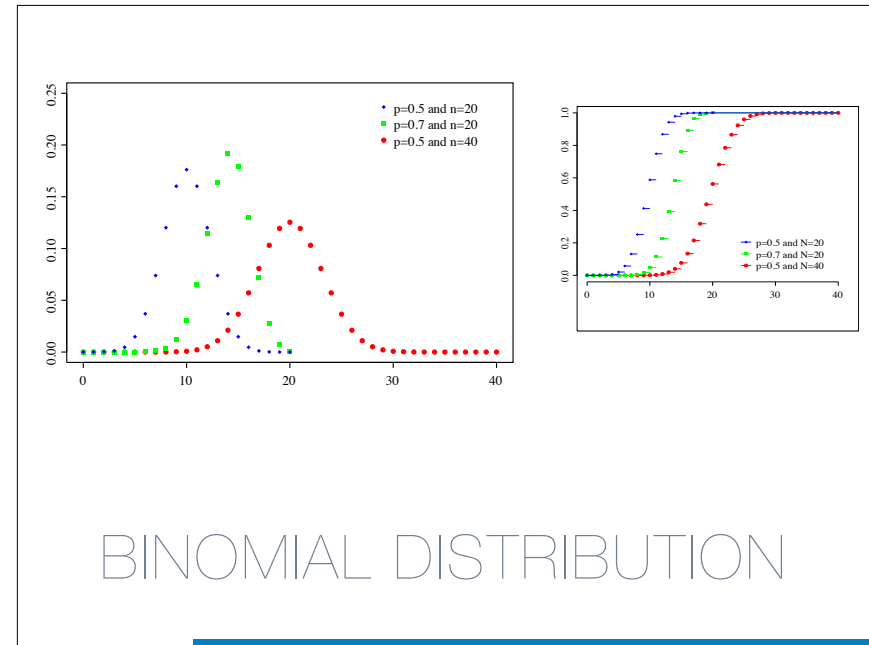


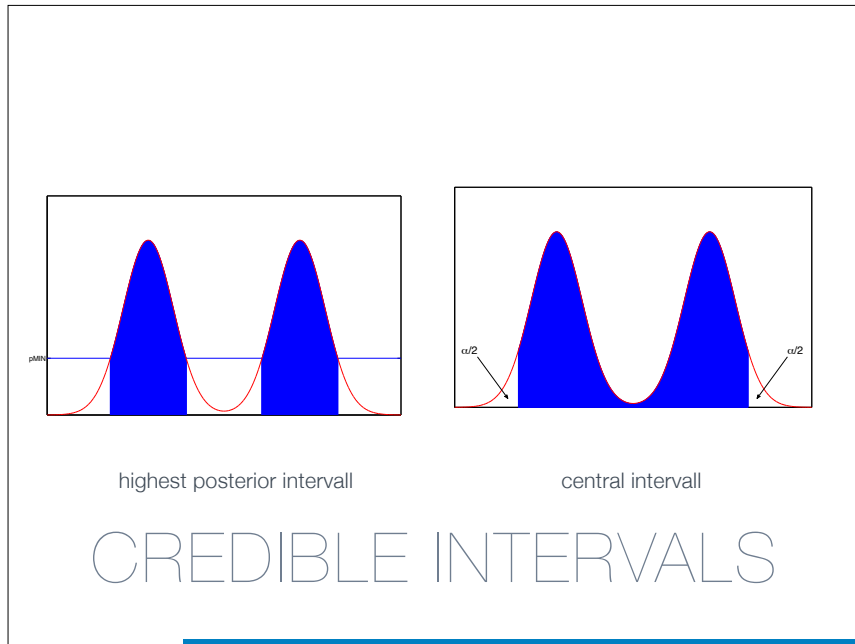


Royal Institute of  
Technology

DD2447 STAT,  
METH. IN CS  
HT 2014

Lecture 4 - Ch. 5,  
Bayesian concepts





# AMAZON SELLERS

- Seller 1 – 90 positive, 10 negative
- Seller 2 – 2 positive, 0 negative
- $\theta_1$  and  $\theta_2$  reliabilities with prior Beta(1,1)
- Uniform on sellers

$$\delta = \theta_1 - \theta_2$$

$$p(\delta > 0 | \mathcal{D}) = \int_0^1 \int_0^1 \mathbb{I}(\theta_1 > \theta_2) \text{Beta}(\theta_1 | y_1 + 1, N_1 - y_1 + 1) \text{Beta}(\theta_2 | y_2 + 1, N_2 - y_2 + 1) d\theta_1 d\theta_2$$

$$p(\delta > 0 | \mathcal{D}) = 0.710,$$



# 5.3 BAYESIAN MODEL SELECTION

- ★ Trade off: low model complexity & underfit vs high model complexity & overfit
- ★ Bayesian: pick MAP model

$$\hat{\mathcal{M}} = \text{argmax}_{\mathcal{M}} p(\mathcal{M} | \mathcal{D})$$

- ★ How to find sweet spot

- ★ Or don't pick one average

- Cross-validation

- ★ If uniform prior, this is ML of marginal


- Posterior probability of models

$$\text{argmax}_{\mathcal{M}} = \int_{\mathcal{M}} p(\mathcal{D} | \theta, \mathcal{M}) p(\theta | \mathcal{M})$$

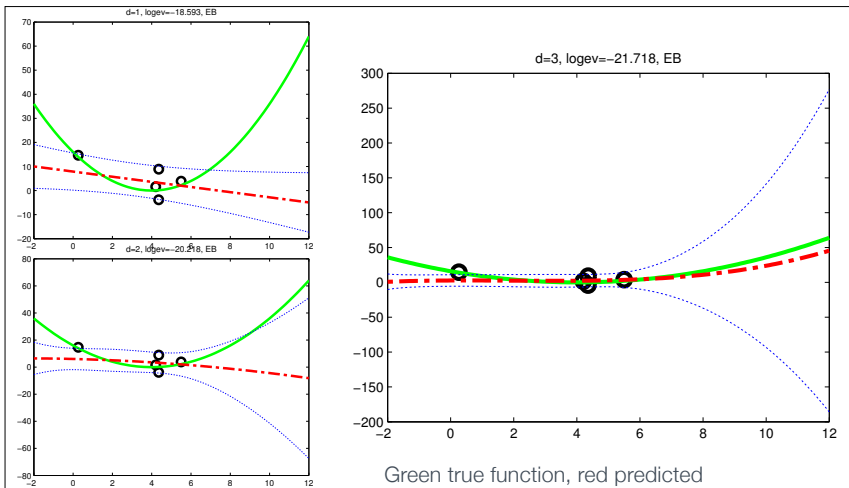
$$p(\mathcal{M} | \mathcal{D}) = \frac{p(\mathcal{D} | \mathcal{M}) p(\mathcal{M})}{p(\mathcal{D})}$$

# 5.3.1 BAYESIAN OCCAM'S RAZOR

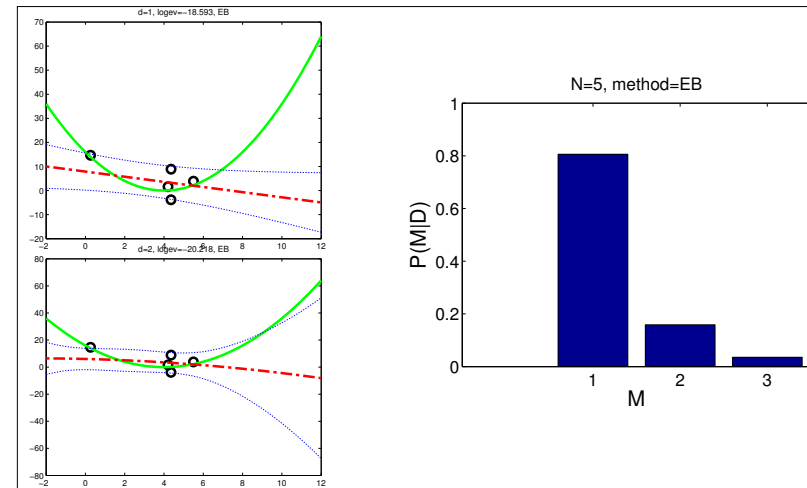
- ★ Marginalizing protects against overfitting
- ★ If  $M' \subset M$  (nested models), then
 
$$p(\mathcal{D}|\theta_{ML}^{M'}) \geq p(\mathcal{D}|\theta_{ML}^M)$$
- ★ Also true for MAP with uniform prior (& others)
- ★ More possible parameters choices for  $M$ , than  $M'$ .
- ★ By marginalized likelihood, these are taken into account



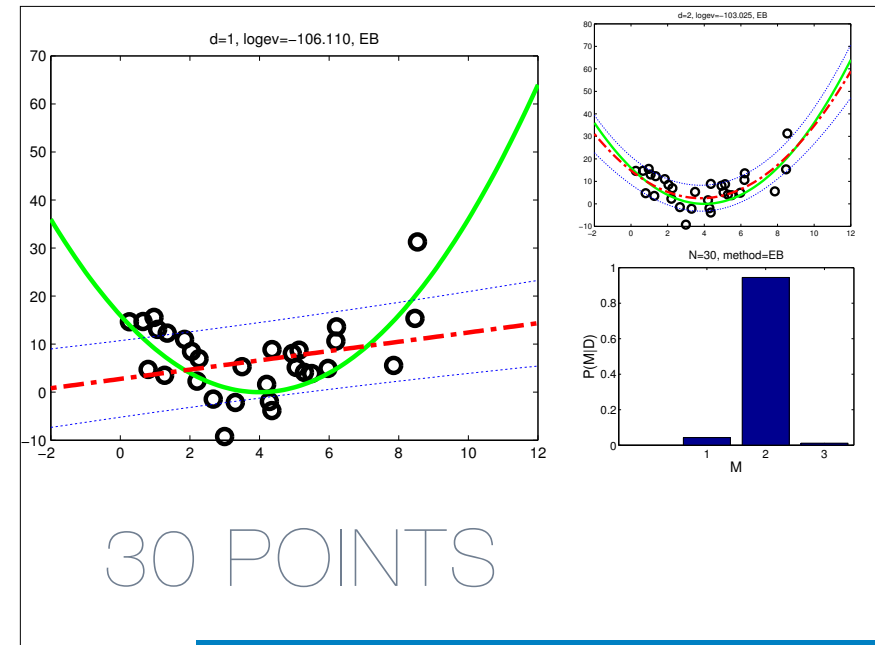
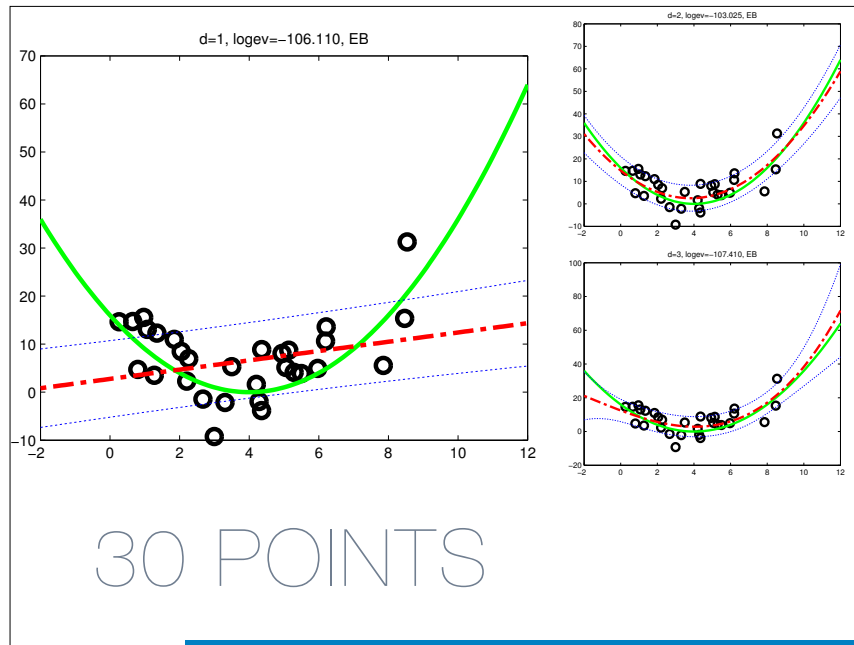
MODELS – BENEFITS OF COMPLEX VS. LESS SO



- Fitting degree 1-3 to 5 points.



- Fitting degree 1-3 to 5 points & posterior



DO 5.3.2, IN YOUR OWN WAY

$\mathcal{D} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$


BAYESIAN INFORMATION CRITERIA (BIC)

$$\text{BIC}(\mathcal{D}, \mathcal{M}) = \underbrace{\log p(\mathcal{D}|\theta_{\text{ML}}^{\mathcal{M}})}_{\propto N} - \underbrace{\text{dof}(\mathcal{M}) \log(N)}_{\propto \log N}$$

$$\log p(\mathcal{D}|\theta_{\text{ML}}^{\mathcal{M}}) = \sum_{n \in N} \log p(\mathbf{x}_n|\theta_{\text{ML}}^{\mathcal{M}})$$

- Computing the marginal likelihood often hard
- BIC score is an approximation of it
- dof - degrees of freedom  $\approx$  number of parameters
- Popular approach
- Above example dof=1,2,3

# TESTING A COINS FAIRNESS



$N_1 = \# \text{ heads}$   
 $N_0 = \# \text{ tails}$

★ Two models (hypothesis):  $M_0$  fair coin  $\Theta=1/2$ ,  $M_1$   $\Theta$  is Beta(1,1)

•  $M_0$  likelihood  $p(\mathcal{D}|\mathcal{M}_0) = 2^{-N}$

•  $M_1$  marginal likelihood

$$p(\mathcal{D}|\mathcal{M}_1) = \int_{\theta} p(\mathcal{D}|\theta)p(\theta)d\theta = B(N_1 + 1, N_0 + 1)/B(1, 1)$$

$$p(\mathcal{D}|\mathcal{M}) = \int \int p(\mathcal{D}|\mathbf{w})p(\mathbf{w}|\alpha, \mathcal{M})p(\alpha|\mathcal{M})d\mathbf{w}d\theta$$

## HOW TO CHOSE HYPER PARAMETERS

$$\hat{\alpha} = \operatorname{argmax}_{\alpha} p(\mathcal{D}|\alpha, \mathcal{M})$$

$$= \operatorname{argmax}_{\alpha} \int_{\mathbf{w}} p(\mathcal{D}|\mathbf{w})p(\mathbf{w}|\alpha, \mathcal{M})d\mathbf{w}$$

$$p(\mathcal{D}|\mathcal{M}) \approx \int p(\mathcal{D}|\mathbf{w})p(\mathbf{w}|\hat{\alpha}, \mathcal{M})d\mathbf{w}$$

- ★ You chose
- ★ Prior on the prior
- ★ The higher in hierarchy, the less effect of parameters
- ★ Empirical Bayes: estimate the level 2 parameters

# LEVELS OF BAYESIANISM

Method	Definition
Maximum likelihood	$\hat{\theta} = \operatorname{argmax}_{\theta} p(\mathcal{D} \theta)$
MAP estimation	$\hat{\theta} = \operatorname{argmax}_{\theta} p(\mathcal{D} \theta)p(\theta \eta)$
ML-II (Empirical Bayes)	$\hat{\eta} = \operatorname{argmax}_{\eta} \int p(\mathcal{D} \theta)p(\theta \eta)d\theta = \operatorname{argmax}_{\eta} p(\mathcal{D} \eta)$
MAP-II	$\hat{\eta} = \operatorname{argmax}_{\eta} \int p(\mathcal{D} \theta)p(\theta \eta)p(\eta)d\theta = \operatorname{argmax}_{\eta} p(\mathcal{D} \eta)p(\eta)$
Full Bayes	$p(\theta, \eta \mathcal{D}) \propto p(\mathcal{D} \theta)p(\theta \eta)p(\eta)$

$$BF_{\mathcal{M}, \mathcal{M}'} := \frac{p(\mathcal{D}|\mathcal{M})}{p(\mathcal{D}|\mathcal{M}')} = \frac{p(\mathcal{M}|\mathcal{D})/p(\mathcal{M})}{p(\mathcal{M}'|\mathcal{D})/p(\mathcal{M}'')}$$

# BAYES FACTORS

★ Ratio between marginals

- Natural way to compare models

★ But what do they mean?

★ “However, ultimately our goal is to convert our beliefs into actions.”

Bayes factor $BF(1,0)$	Interpretation
$BF < \frac{1}{10}$	Decisive evidence for $M_0$
$BF < \frac{1}{3}$	Strong evidence for $M_0$
$\frac{1}{10} < BF < \frac{1}{3}$	Moderate evidence for $M_0$
$\frac{1}{3} < BF < 1$	Weak evidence for $M_0$
$1 < BF < 3$	Weak evidence for $M_1$
$3 < BF < 10$	Moderate evidence for $M_1$
$BF > 10$	Strong evidence for $M_1$
$BF > 100$	Decisive evidence for $M_1$

Table 5.1 Jeffreys' scale of evidence for interpreting Bayes factors.

$$p(\theta) = \frac{1}{2} \text{Beta}(\theta|20, 20) + \frac{1}{2} \text{Beta}(\theta|30, 10)$$

MIXTURE OF  
CONJUGATE IS  
CONJUGATE  
PRIOR


In general

$$p(\theta) = \sum_k \underbrace{p(\theta|z=k)}_{\text{conjugate}} \underbrace{p(z=k)}_{\text{mixing weights}}$$

Gives posterior

$$p(\theta|\mathcal{D}) = \sum_k p(\theta|z=k, \mathcal{D}) p(z=k|\mathcal{D})$$

- Can approximate any prior
- Say, likelihood  $\text{Ber}(\theta)$
- Mixing with prior weights gives posterior weights



MIXTURE OF 2 BETAS

## CONSERVED DNA SEQUENCES

```

pap  CATTAGACCGATCTTTTATGCTGT-AAATCAATTGCCATGATGTTTTTATCTGAGTACCCTCTTGCTATTAGTGTTTT
foo  CATTAGACCGATCTTTTATGCTGT-AAATCAATTCACCATGATGTTTTTATCTGAGTGTATTCTTGTTGTTTGTGTTTT
sfa  GATTTTAAACCGATCTTTTATACTGA-ATATTCATGCTTATACAGTATTATAAATCTAAAACGCCAATCCACTCGGATATA
afa  GATTATAACCGATCTTTTATTCAC-ACATGAATAATATCCCGTTATAATTTCTGATTGTATTCTTTTTGTGTTATCTG
daa  AATTATAACCGATCTTTTATTCAGC-ATATGAATAATATCCAGTCATATATCTGATTGTATTCTTTTTGTGTTATATG
clp  GTTTTAGAACGATCTTTTATCTGIGATTTTGTGTTTTTTGGTGGTTTTTGTGTTGTTGTTGTTTTGTGTTTTTAT
fae  TGAATAACCGATCTTTTATTTGTGATTTTTCGTGTTTTGGTGTGTTGCTGTGCTTTTGTGTTGTTGTTTTT
pef  TGCTATACCGATCTTTTATTCGCATTATGTATGATAA-ATAGACGTTTAAACCTGTTTTCTTGCTATTTTGTGGTGAAC
    
```

- Functional elements are more conserved than non-functional
- Dependence between consecutive positions, which we ignore

## TWO MODELS – THAT CAN BE TESTED


$t$  – column of nucleotides

$z_t = 1$  if conserved;  $z_t = 0$  if not conserved

$$p(N_t|z_t) = \int_{\mathcal{M}} p(N_t|\theta_t) p(\theta_t|z_t) d\theta_t$$

$$p(\theta_t|z_t = 1) = \frac{1}{4} [\text{Dir}(\theta|10, 1, 1, 1) + \dots + \text{Dir}(\theta|1, 1, 1, 10)]$$

$$p(\theta_t|z_t = 0) = \text{Dir}(\theta|1, 1, 1, 1)$$



## HIERARCHICAL BAYES: EX. CANCER RATES

- \* Model  $\text{Bin}(x|\Theta, N)$
- \* Alternatives:
  - cities independent
  - tie  $\Theta$
  - common prior  $\text{Beta}(\Theta|a,b)$
- \* Problem: small city, poor estimate
- \* Tie – pool data and use MLE
  - but we expect differences

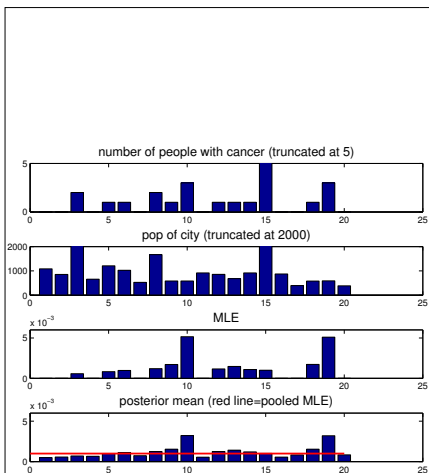
$$\mathcal{D} := \{x_1, \dots, x_S\}$$

Full joint

$$p(\mathcal{D}, \theta, \overset{(a,b)}{\eta} | N) = p(\mathcal{D}, \theta | \eta, N) p(\eta) = p(\eta) \prod_{i \in [S]} \text{Bin}(x_i | N_i, \theta_i) \text{Beta}(\theta_i | \eta)$$

## EX. CANCER RATES – CONTINUED

- \*  $x_i$  cancer deaths in city  $i$
- \*  $N_i$  population size in city  $i$
- \* Common prior  $\text{Beta}(a,b)$
- \*  $\eta = (a,b)$  must be inferred
  - otherwise (given  $\eta=(a,b)$ ) cities are independent

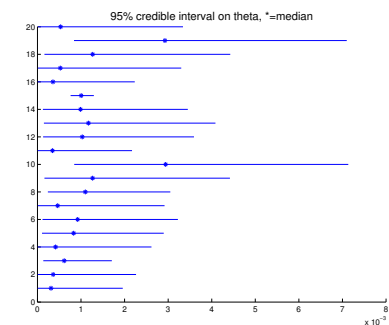



Means of marginals of joint posterior

## HIERARCHICAL BAYES: EX. CANCER RATES

- \* Model  $\text{Bin}(x|\Theta, N)$
- \* Alternatives:
  - cities independent
  - tie  $\Theta$
  - common prior  $\text{Beta}(\Theta|a,b)$

## CANCER RATES CREDIBLE INTERVAL



# 5.7 BAYESIAN DECISION THEORY

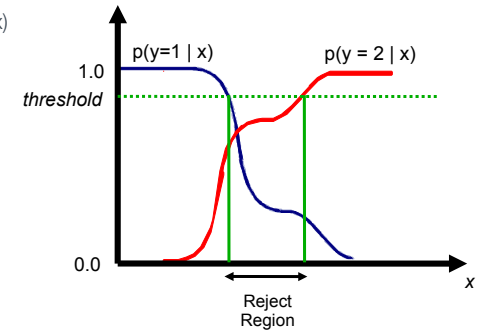
- Given  $x$  we chose an action  $a$
- Economics utility  $U(y,a) = -L(y,a)$
- Loss  $L(y,a)$  measured compared to hidden state/param./class  $y$
- Optimal decision procedure
 
$$\delta(\mathbf{x}) = \operatorname{argmin}_a E[L(y, a)]$$
- Bayesian approach expected posterior loss
 
$$\rho(a|\mathbf{x}) = \operatorname{argmin}_{p(y|\mathbf{x})} E[L(y, a)]$$

$$= \sum_y p(y|\mathbf{x}) L(y, a)$$

← if discrete
- misclassification  $L(y,a) = I(y \neq a)$
- squared  $L(y,a) = (y-a)^2$

# MAP MINIMIZES 0-1 LOSS

- $L(y,a) = I(y \neq a)$
- $\rho(a|x) = p(a \neq y|x) = 1 - p(a|x)$
- Hence,
  - maximizing  $p(a|x)$
  - minimizes  $\rho(a|x)$
- Read about reject option



$$\rho(a|\mathbf{x}) = E[(y - a)^2 | \mathbf{x}] = E[y^2 | \mathbf{x}] - 2aE[y | \mathbf{x}] + a^2$$

$$\frac{\partial \rho(a|\mathbf{x})}{\partial a} = 2a - 2E[y | \mathbf{x}]$$

$$\frac{\partial \rho(a|\mathbf{x})}{\partial a} = 0$$

$$a = E[y | \mathbf{x}] = \int y p(y|\mathbf{x}) d$$

POSTERIOR MEAN MINIMIZES QUADRATIC ( $L_2$ ) LOSS

- Squared loss  $L(y,a) = (y-a)^2$
- Assuming continuous

# TRUE AND FALSE, POSITIVE AND NEGATIVE

		Truth		$\Sigma$
		1	0	
Estimate	1	TP	FP	$N_+ = TP + FP$
	0	FN	TN	$N_- = FN + TN$
$\Sigma$		$N_+ = TP + FN$	$N_- = FP + TN$	$N = TP + FP + FN + TN$

Table 5.2 Quantities derivable from a confusion matrix.  $N_+$  is the true number of positives,  $\hat{N}_+$  is the "called" number of positives,  $N_-$  is the true number of negatives,  $\hat{N}_-$  is the "called" number of negatives.


$\hat{y}$	$y = 1$		$y = 0$	
	1	$TP/N_+ = \text{TPR} = \text{sensitivity} = \text{recall}$	$FP/N_- = \text{FPR} = \text{type I}$	
0	$FN/N_+ = \text{FNR} = \text{miss rate} = \text{type II}$	$TN/N_- = \text{FNR} = \text{specificity}$		

Table 5.3 Estimating  $p(\hat{y}|y)$  from a confusion matrix. Abbreviations: FNR = false negative rate, FPR = false positive rate, TNR = true negative rate, TPR = true positive rate.

- Binary decision problem: for  $x \in U$ ,  $x \in C$ ?
- Positives are the ones claimed to be in C
  - true or false depending on membership of C



# RECEIVER OPERATING CHARACTERISTICS (ROC) CURVES



- ★ True positive rate (TPR; recall):  $TP/(TP+FN)=TP/|C|=p(\hat{y}=1 | y=1)$
- ★ False positive rate (FPR):  $FP/(TN+FP)=FP/|U \setminus C|=p(\hat{y}=1 | y=0)$
- ★ Precision:  $TP/(TP+FP)=p(y=1 | \hat{y}=1)$

The end