Bayesian Data Analysis in R

Kassie Fronczyk
Institute for Defense Analyses

28 October, 2016

Fundamentals of Bayesian Analysis

Case Study: Littoral Combat Ship

Case Study: MaxxPro LWB Ambulance

Case Study: Bio-chemical Detection System

Discussion

Bayesian methods are commonly used and becoming more widely accepted

- Applications
  - FAA/ USAF in estimating probability of success of launch vehicles
  - Delphi Automotive for new fuel injection systems
  - Science-based Stockpile Stewardship program at LANL for nuclear warheads
  - Army for estimating reliability of new aircraft systems
  - FDA for approving new medical devices
- Recent High Profile Successes:
  - During the search for Air France 447 (2009-2011), black box location
  - The Coast Guard in 2013 found the missing fisherman, John Aldridge
  - Use in defining the search area for Malaysian Airlines flight MH370 in 2014 by Australian government

Classical versus Bayesian Statistics

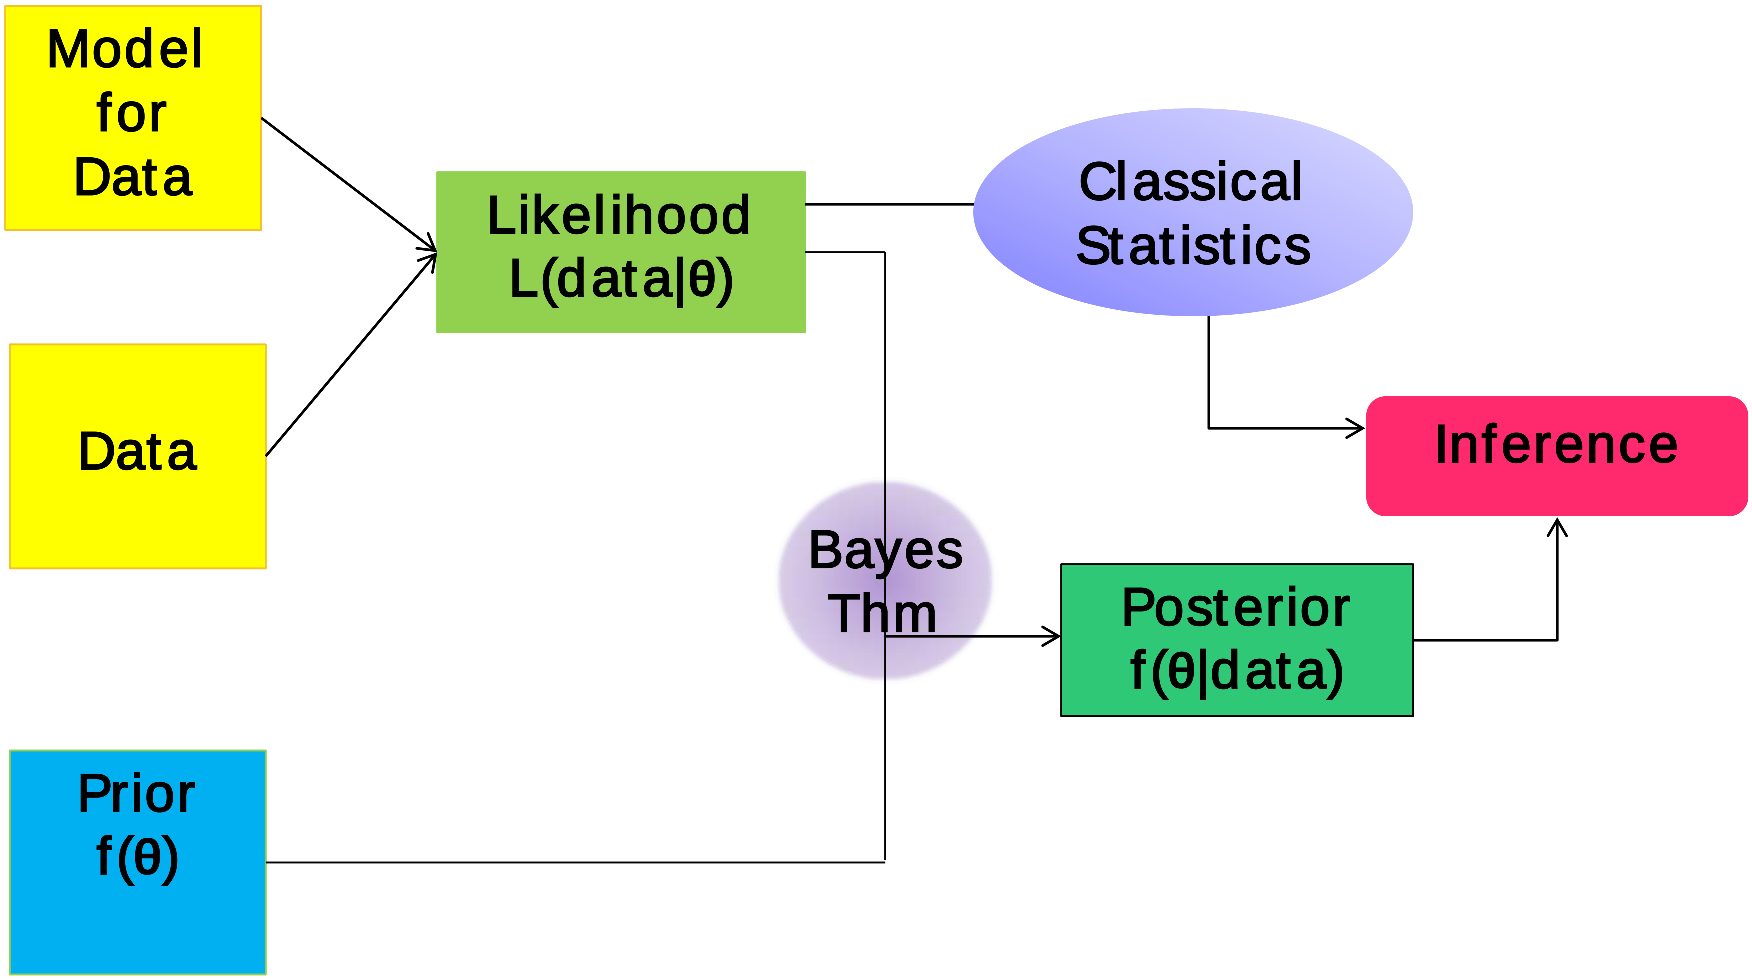

Bayesian Statistics 101

We have a system comprised of 2 components: Component 1 and Component 2.

For each of the two components, 10 pass/ fail tests are administered and results are recorded. Component 1 fails twice and Component 2 fails zero times.

- We can calculate the reliability of each component, R_1 and R_2.
- We also want an assessment of the system reliability, assuming the components work in series.

$$R_{system} = R_1 * R_2 = \left(1 - \frac{2}{10}\right) * \left(1 - \frac{0}{10}\right) = 0.8 * 1 = 0.8$$

**For the purposes of the next few slides, focus on Component 1.

Bayesian Statistics 101

Your Bayesian analysis is just 3 steps away:

1. Construct prior from prior information
2. Construct likelihood from test data
3. Estimate posterior distribution using Bayes Theorem

NOTE: these are not trivial steps! They require thought and understanding of both the system and the statistics!

Bayesian Statistics 101: Priors

The prior distribution of the reliability, $f_{prior}(R)$, is constructed from previous data or expert knowledge. This is your first assessment of the system.

Say Component 1 was previously tested and failed 3 out of 40 tests: use Beta distribution.

$$f_{prior}(R_1) \propto R_1^{\,n_p p}(1-R_1)^{n_p(1-p)}$$

- p is the reliability estimate and $n_p \geq 0$ weights the relevance of the prior test data.

Careful thought should always be put into the prior distribution!

Bayesian Statistics 101: Likelihood

Tests are performed and the resulting test data is used in the likelihood function, L(data|R).

This is the same as in Classical Statistics!

The binary test data of Component 1 follows a Binomial distribution with probability of a pass of R_1

$$L(data|R_1) \propto R_1^{s1}(1-R_1)^{f1}$$

s_1 is the number of successes and f_1 is the number of failures from Component 1.

Bayes Theorem

$$f(\theta|x) = \frac{L(x|\theta)f(\theta)}{\int L(x|\theta)f(\theta)}$$

$$\propto L(x|\theta)f(\theta)$$

Bayesian Statistics 101: Posterior Distribution

Bayes theorem is used to find the posterior reliability distribution, $f_{posterior}(R|data)$. The posterior distribution is the product of the prior distribution and the likelihood function for all subsystems in the unit

For our small example, choosing the Beta distribution as a prior is ideal for a few reasons: it ensures that *R* is between (0, 1) and it is the "conjugate" prior for the Binomial distribution (i.e. the math works out easily)

$$f_{posterior}(R_1) \propto R_1^{s1}(1-R_1)^{f1} R_1{}^{\boldsymbol{n_p}p}(1-R_1)^{\boldsymbol{n_p}(1-p)}$$
$$\propto R_1^{s1+\boldsymbol{n_p}p}(1-R_1)^{f1+\boldsymbol{n_p}(1-p)}$$

Bayesian Statistics 101: Posterior Distribution

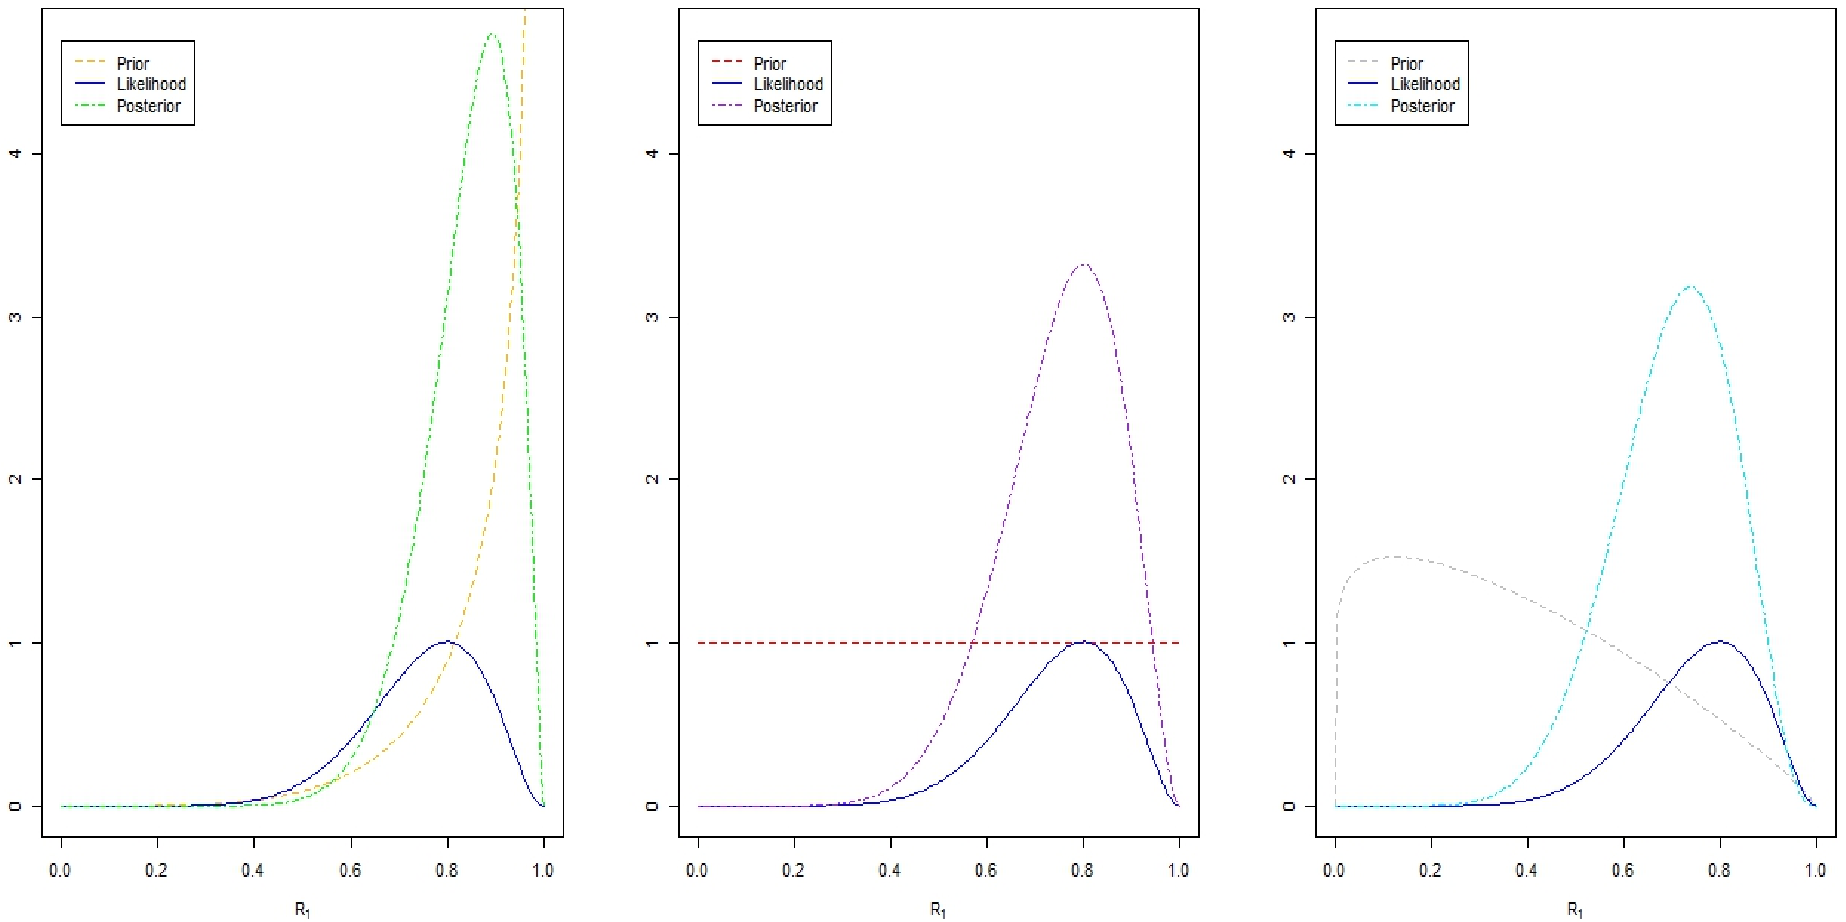

0.86 (0.76, 0.95) | 0.75 (0.58, 0.89) | 0.70 (0.54, 0.85)

Classical Estimate: 0.8 (0.55, 0.95)

Bayesian Statistics 101: Conjugate Priors

Likelihood	Parameter	Prior	Posterior
Binomial(s+f, R)	$0 \leq R \leq 1$	Beta(a,b) $a > 0, b > 0$	Beta(a',b') $a' = a + s$ $b' = b + f$
Poisson(λ)	$\lambda > 0$	Gamma(a,b) $a > 0, b > 0$	Gamma(a',b') $a' = a + n$ $b' = b + \Sigma t$
Exponential(λ)	$\lambda > 0$	Gamma(a,b) $a > 0, b > 0$	Gamma(a',b') $a' = a + n$ $b' = b + \Sigma t$

**For more examples, see Conjugate Priors Wikipedia page or "Bayesian Reliability" pg 48

When should we think about using Bayesian techniques?

- To obtain interval estimates (credible intervals) when there are zero failures
  - Mean time between failure for short tests or for highly reliable systems
  - Interval estimates in kill-chain analysis where zero failures occur at any point along the kill-chain
- If you are assessing a complex system mission reliability
  - LCS Example - Confidence intervals are not straightforward to obtain using frequentist methods, impossible with zero failures in any sub-system
- If there is relevant prior information to be incorporated in your analysis – this may include previous developmental (or operational) test data, engineering analyses, or information from modeling and simulation.
  - MaxxPro LWB Ambulance Example
  - BDS Example

Fundamentals of Bayesian Analysis

Case Study: Littoral Combat Ship

Case Study: MaxxPro LWB Ambulance

Case Study: Bio-chemical Detection System

Discussion

Case Study: LCS

Freedom Class

Independence Class

- The Littoral Combat Ship (LCS) are a new family of surface ships.
- The Capability Development Document for LCS provides a reliability requirements for four functional areas
  - Sea Frame Operations
  - <u>Core Mission</u>
  - Mission Package Support
  - SUW Mission Package

Case Study: LCS

The Capability Development Document for LCS provides a reliability threshold for Core Mission functional area.

Critical Subsystem	Total System Operating Time	Operational Mission Failures
Total Ship Computing Environment (full-time)	4500 hours	1
Sea Sensors and Controls (underway)	2000 hours	3
Communications (full-time)	4500 hours	0
Sea Engagement Weapons (on-demand)	11 missions	2

The target reliability for Core Mission is 0.80 in 720 hours.

➢Assume the functional area is a series system: system is up if all subsystems are up.

Case Study: LCS – Prior Assumptions

On-demand system

- Assume no belief in the relevance of prior knowledge, $n_p = 0$

Continuous systems

- The Gamma prior parameter a is set to 1, giving large variance. To ensure the 50th percentile is set at λ_{50}=1/ $MTBF_{guess}$, choose b=log(2)×$MTBF_{guess}$
- $MTBF_{guess}$ chosen by solving the reliability function at the requirement

Guiding Principles in Prior Selection:

- *Start with the properties of the parameter of interest*
- *Decide on what prior information to use*
- *Allow for the analysis to change freely based on the data observed*
- *Priors specified at the system level, as opposed to mission level – check impact on system prior*

Case Study: LCS – Prior Assumptions

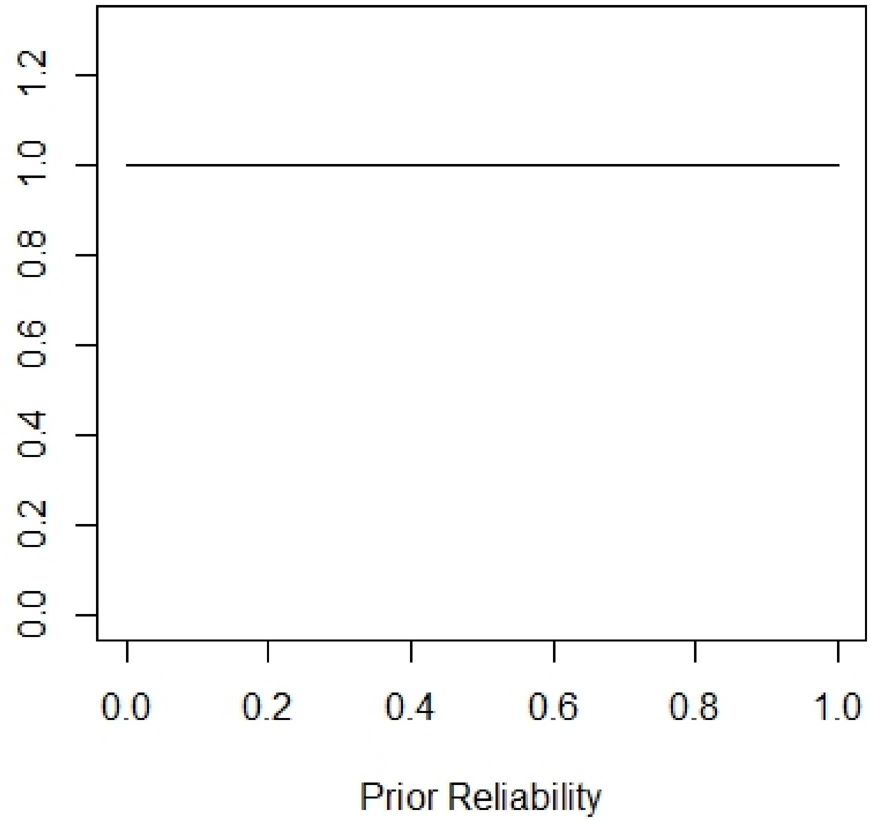

Case Study: LCS – Prior Assumptions

Possible Prior 1: Too much information, if the prior probability of a value is 0, no amount of data can move the posterior MTBF there!

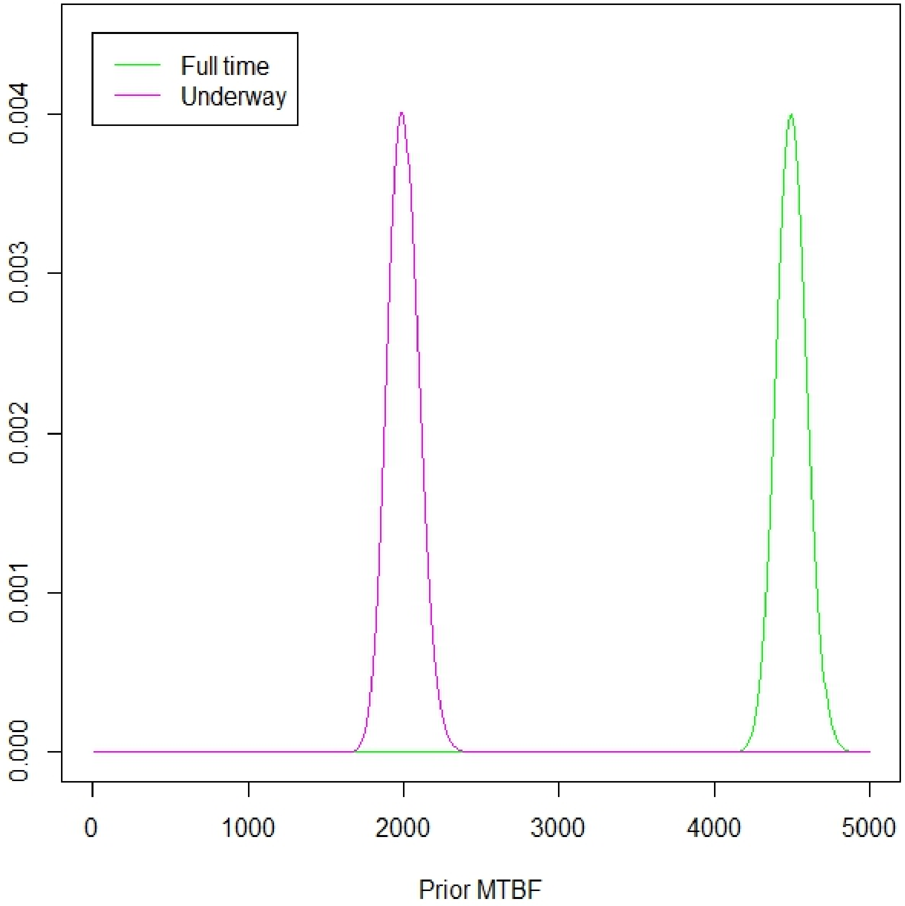

Possible Prior 2: Too little information, for continuous measures, flat priors can be problematic when there are few or zero failures.

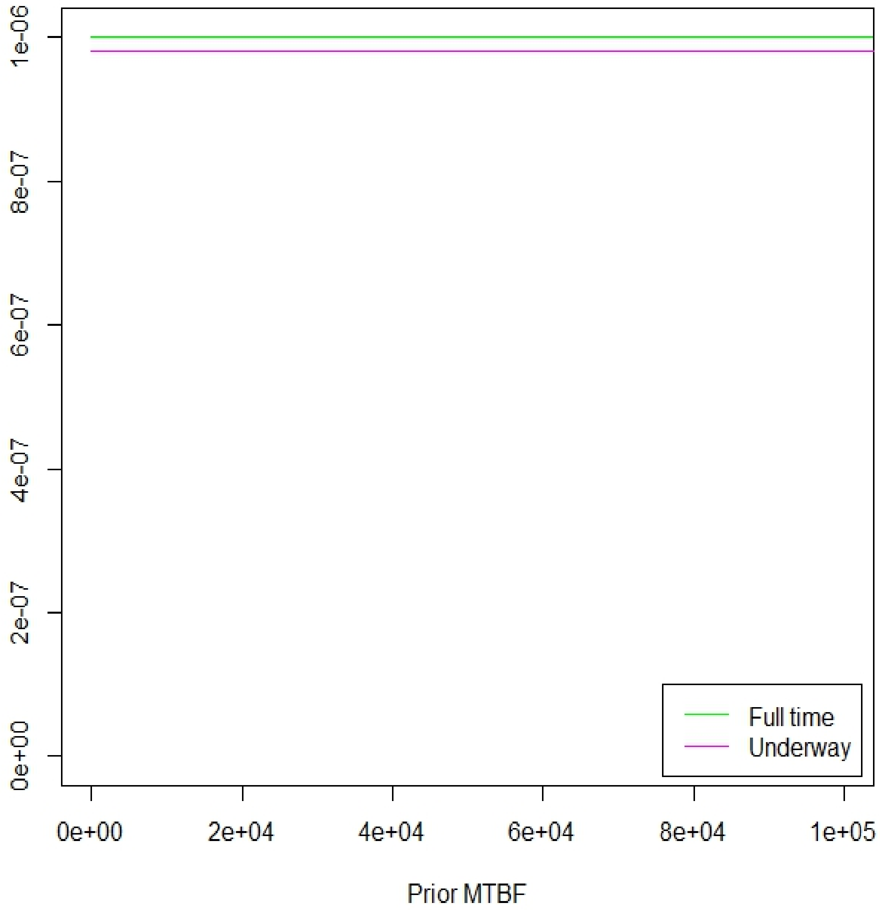

Case Study: LCS – Prior Assumptions

Possible Prior 3: Bounding the prior in case there are few failures, but ensuring enough flexibility for the data to speak for itself.

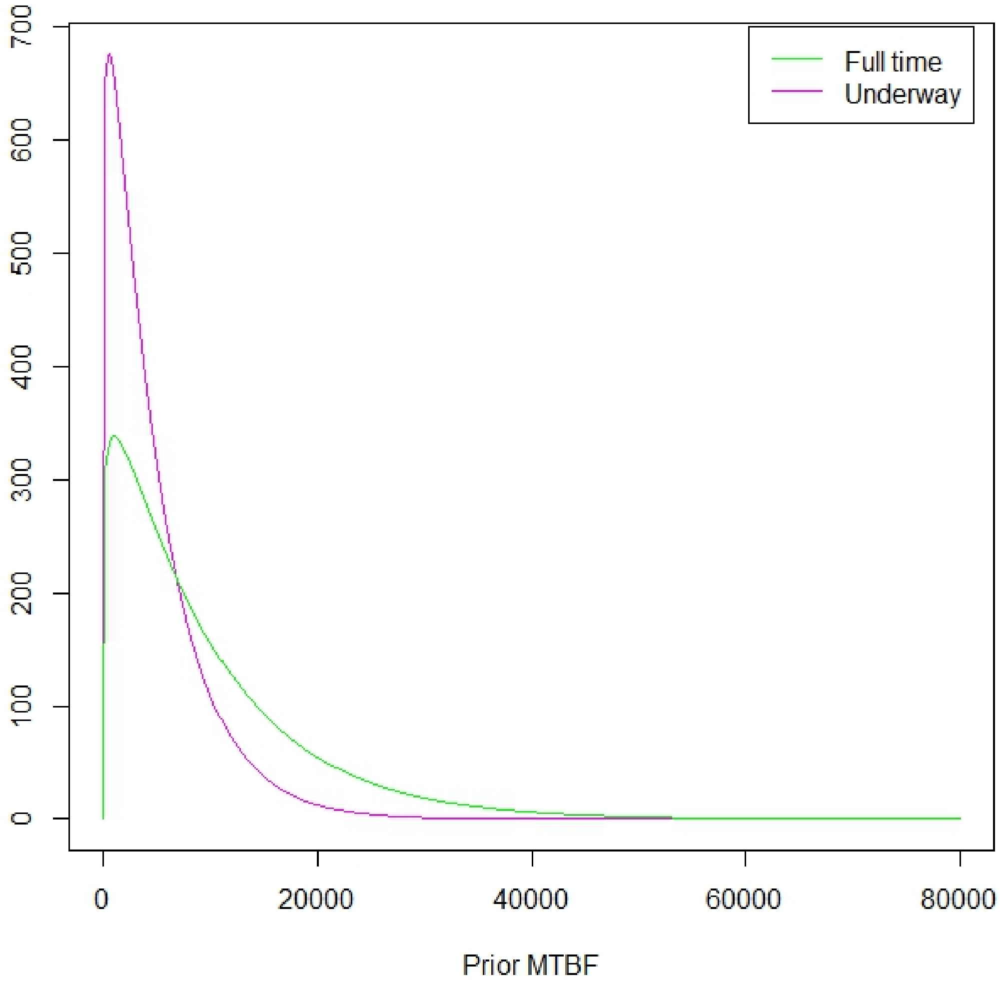

Case Study: LCS – Prior Assumptions

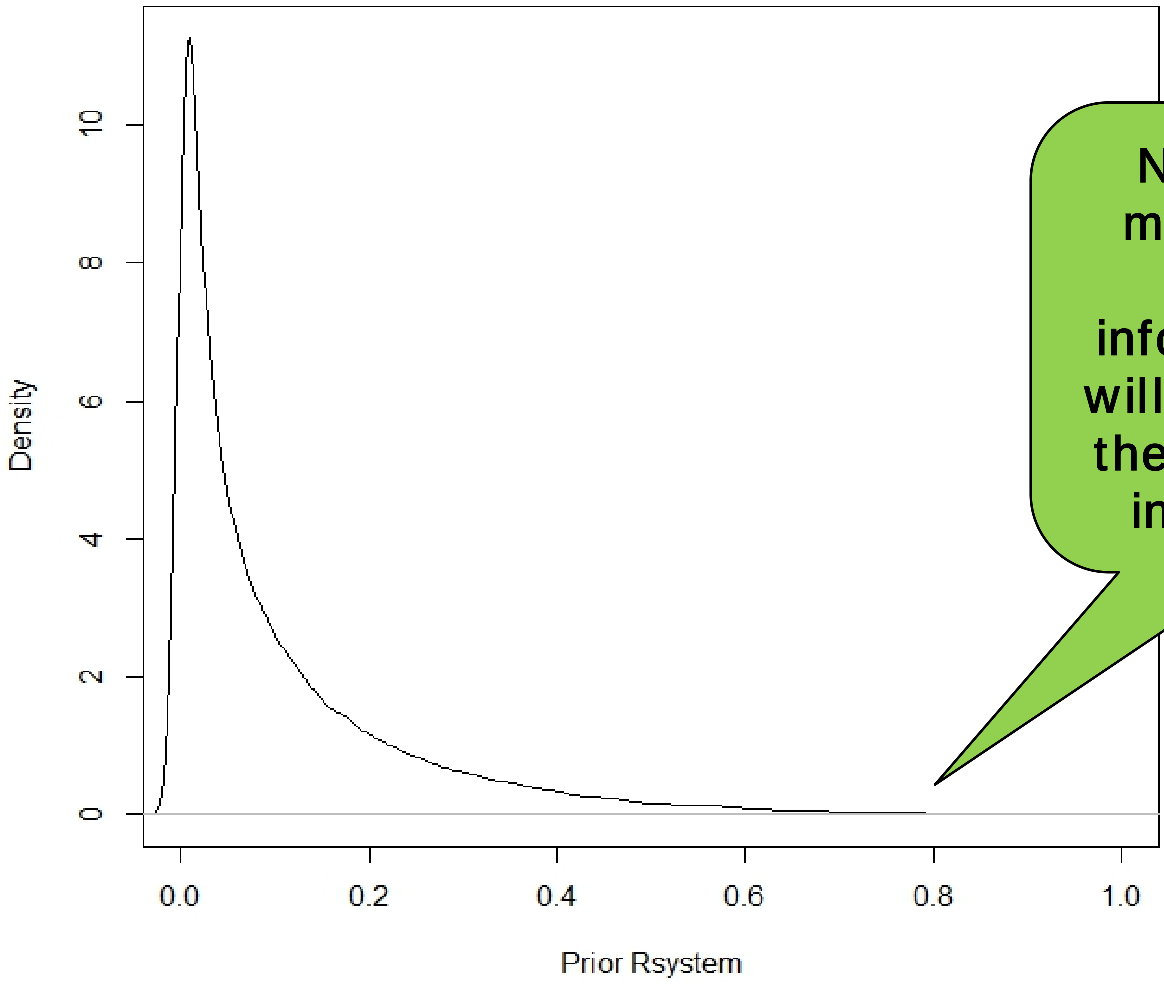

Note the core mission prior is somewhat informative – We will want to check the impact of this in the analysis

Case Study: LCS – Posterior Computation

- Even when we choose independent, conjugate priors for each subsystem or component, the combined system reliability does not have an analytic solution for its distribution!
- Simulate values from the posterior distribution via MCMC

Basic MCMC algorithm

- Generate a value from each component/ subsystem reliability parameter posterior distribution.
- If necessary, calculate reliability at time t.
- Combine the component/ subsystem reliabilities through the expression determined by the system structure.
- Repeat many, many times.

Caste Study: LCS – Posterior Computation

```
> for(i in 1:B){
+
+        lambdaTC[i]=rgamma(1,TSCE+aT,TSCET+bT)
+
+        lambdaSSC[i]=rgamma(1,SensCont+aS,SensContT+bS)
+
+        lambdaCOMM[i]=rgamma(1,Comm+aC,CommT+bC)
+
+        pSEW[i]=rbeta(1,SEWs + pguess*pweight,SEWf+(1-pguess)*pweight)
+
+        Rsys[i]=exp(-720*lambdaTC[i])*exp(-720*lambdaSSC[i])*exp(-720*lambdaCOMM[i])*pSEW[i]
+ }
```

Case Study: LCS - Results

	Classical MTBOMF	Classical Reliability at 720hrs	Bayesian MTBOMF	Bayesian Reliability at 720hrs
TSCE	4500 hrs (1156 hrs, 42710 hrs)	0.85 (0.54,0.98)	3630 hrs (1179 hrs, 6753 hrs)	0.73 (0.54,0.90)
SSC	667 hrs (299 hrs, 1814 hrs)	0.33 (0.09,0.67)	697 hrs (332 hrs, 1172 hrs)	0.31 (0.11,0.54)
Comm	> 2796 hrs	> 0.77*	10320 hrs (1721 hrs, 18210 hrs)	0.83 (0.66,0.96)
SEW		0.82 (0.58,0.95)		0.77 (0.62,0.91)
Core Mission		?????		0.15 (0.05, 0.27)

Many ways to think about calculating this, none of which are particularly satisfactory

Note the impact of the prior is greater in the one failure system

Full Mission mean is comparable with the simple point estimate

TSCE: Total Ship Computing Environment
SSC: Sea Sensors and Controls
Comm: Communications
SEW: Sea Engagement Weapons

Case Study: LCS - Results

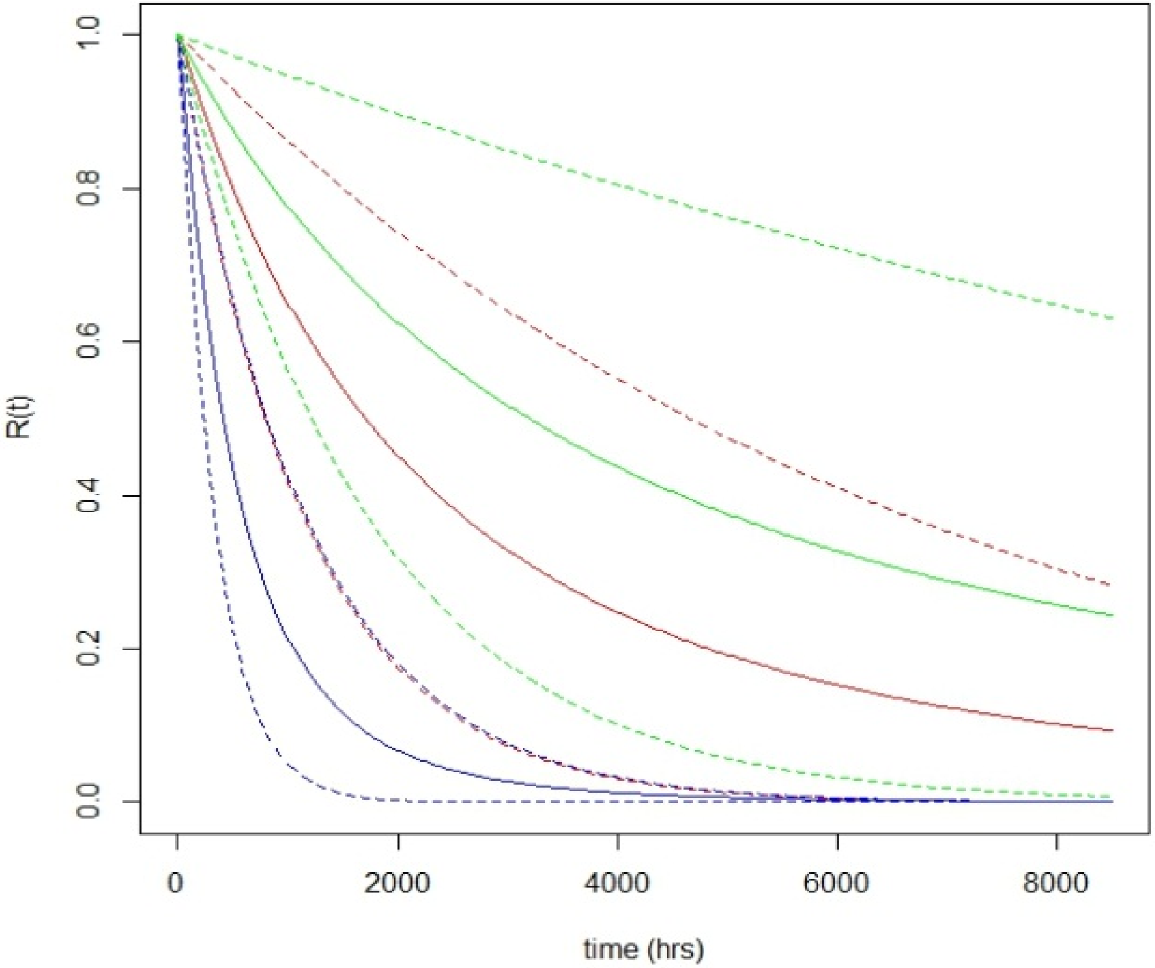

Posterior reliability as a function of time for TSCE (red), SSC (blue), and Communications (green)

Case Study: LCS - Results

Posterior mean and 80% intervals for each subsystem and the total system reliability over 15 days (light blue) and 30 days (dark blue) for the notional example.

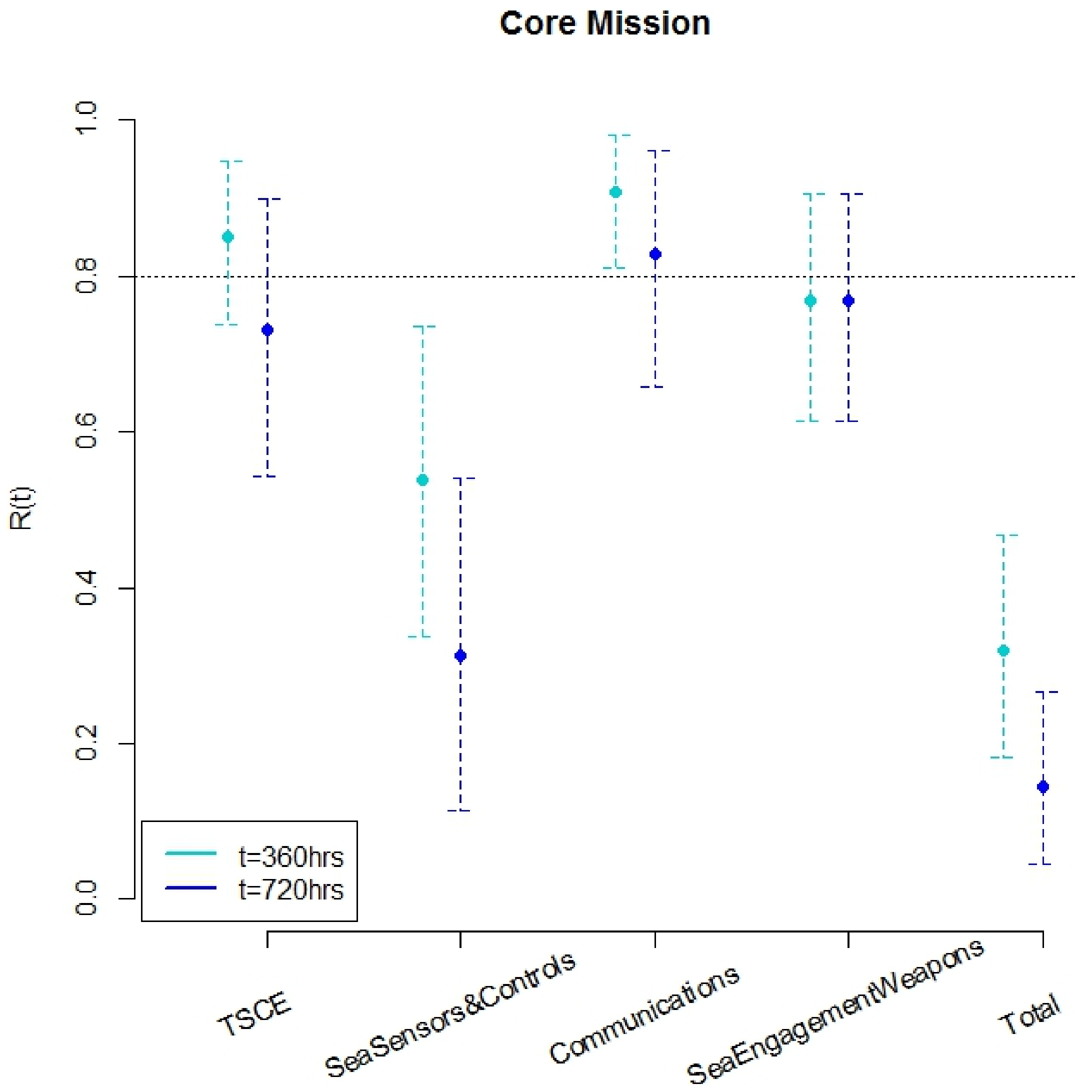

Case Study: LCS – Value of Bayesian Statistics

- Avoids unrealistic reliability estimates when there are no observed failures.
- In our notional example (zero failures for the Communications system), the Bayesian approach helped us solve an otherwise intractable problem.
- Obtaining interval estimates is straightforward for system reliability
  - Frequentist methods would have to employ the Delta method, Normal approximations, or bootstrapping.
- Flexibility in developing system models
  - We used a series system for the core mission reliability
  - Many other system models are possible and we can still get full system reliability estimates with intervals.

Fundamentals of Bayesian Analysis

Case Study: Littoral Combat Ship

Case Study: MaxxPro LWB Ambulance

Case Study: Bio-chemical Detection System

Discussion

Case Study: MaxxPro LWB Ambulance

LWB Ambulance with 72" tall Soldier

Litter rail and trolley system

- Primary mission of an ambulance-equipped unit is medical evacuation.
- Three medical Soldiers crew the vehicle: driver, vehicle commander, and medic/gunner.
- Holds two litter patients, one litter patient and two ambulatory patients, or four ambulatory patients.
- The ambulance has a rail, trolley, and hoist system for litter loading/unloading, and medical support equipment (monitoring equipment, intravenous management system, oxygen concentrators, etc.)

Case Study: MaxxPro LWB Ambulance

- MRAP vehicles have a reliability requirement of at least 600 Mean Miles Between Operational Mission Failures (MMBOMF)
- 1 Vehicle available for DT to drive 1025 miles
- 2 Vehicles available in the LUT that drove a total of 3026 miles

- There were four OMFs in DT and one OMF in LUT
  - 3 flat tires and 1 air conditioner failure in DT
  - 1 flat tire in LUT

- Flat tires during missions result in OMFs, because the LWB Ambulance does not carry a usable jack and spare tire

Case Study: MaxxPro - Model

Model:

Modeling Vehicle Failure Miles

$$X_{1DTl} \sim Exp(\lambda)$$
$$X_{iLUTl} \sim Exp(\gamma\lambda)$$

Estimating MMBF

Change from DT to LUT

$$i = 1,2 \; for \; vehicle$$
$$l = 1,2,\dots,n_{ij} \; for \; observed \; failures$$

$$\lambda \sim Gamma(a_\lambda, b_\lambda), \gamma \sim Gamma(a_\gamma, b_\gamma),$$
$$a_\lambda = b_\lambda = a_\gamma = b_\gamma = 0.001$$

Posteriors:

$$\lambda | \gamma, X \sim Gamma(n_{1\boldsymbol{DT}} + n_{1\boldsymbol{LUT}} + n_{2\boldsymbol{LUT}} + a_\lambda, t_{1\boldsymbol{DT}c} + \gamma(t_{1\boldsymbol{LUT}c} + t_{2\boldsymbol{LUT}c}) + b_\lambda)$$
$$\gamma | \lambda, X \sim Gamma(n_{1\boldsymbol{LUT}} + n_{2\boldsymbol{LUT}} + a_\gamma, \lambda(t_{1\boldsymbol{LUT}c} + t_{2\boldsymbol{LUT}c}) + b_\gamma)$$

Case Study: MaxxPro – Posterior Computation

```
> for(i in 2:B){
+
+   lambda[i] = rgamma(1, DTFails+LUTFails+alam, DTMiles+gamma[i-1]*LUTMiles+blam)
+
+   gamma[i] = rgamma(1, LUTFails+agam, lambda[i]*LUTMiles+bgam)
+
+ }
>
```

Case Study: MaxxPro - Results

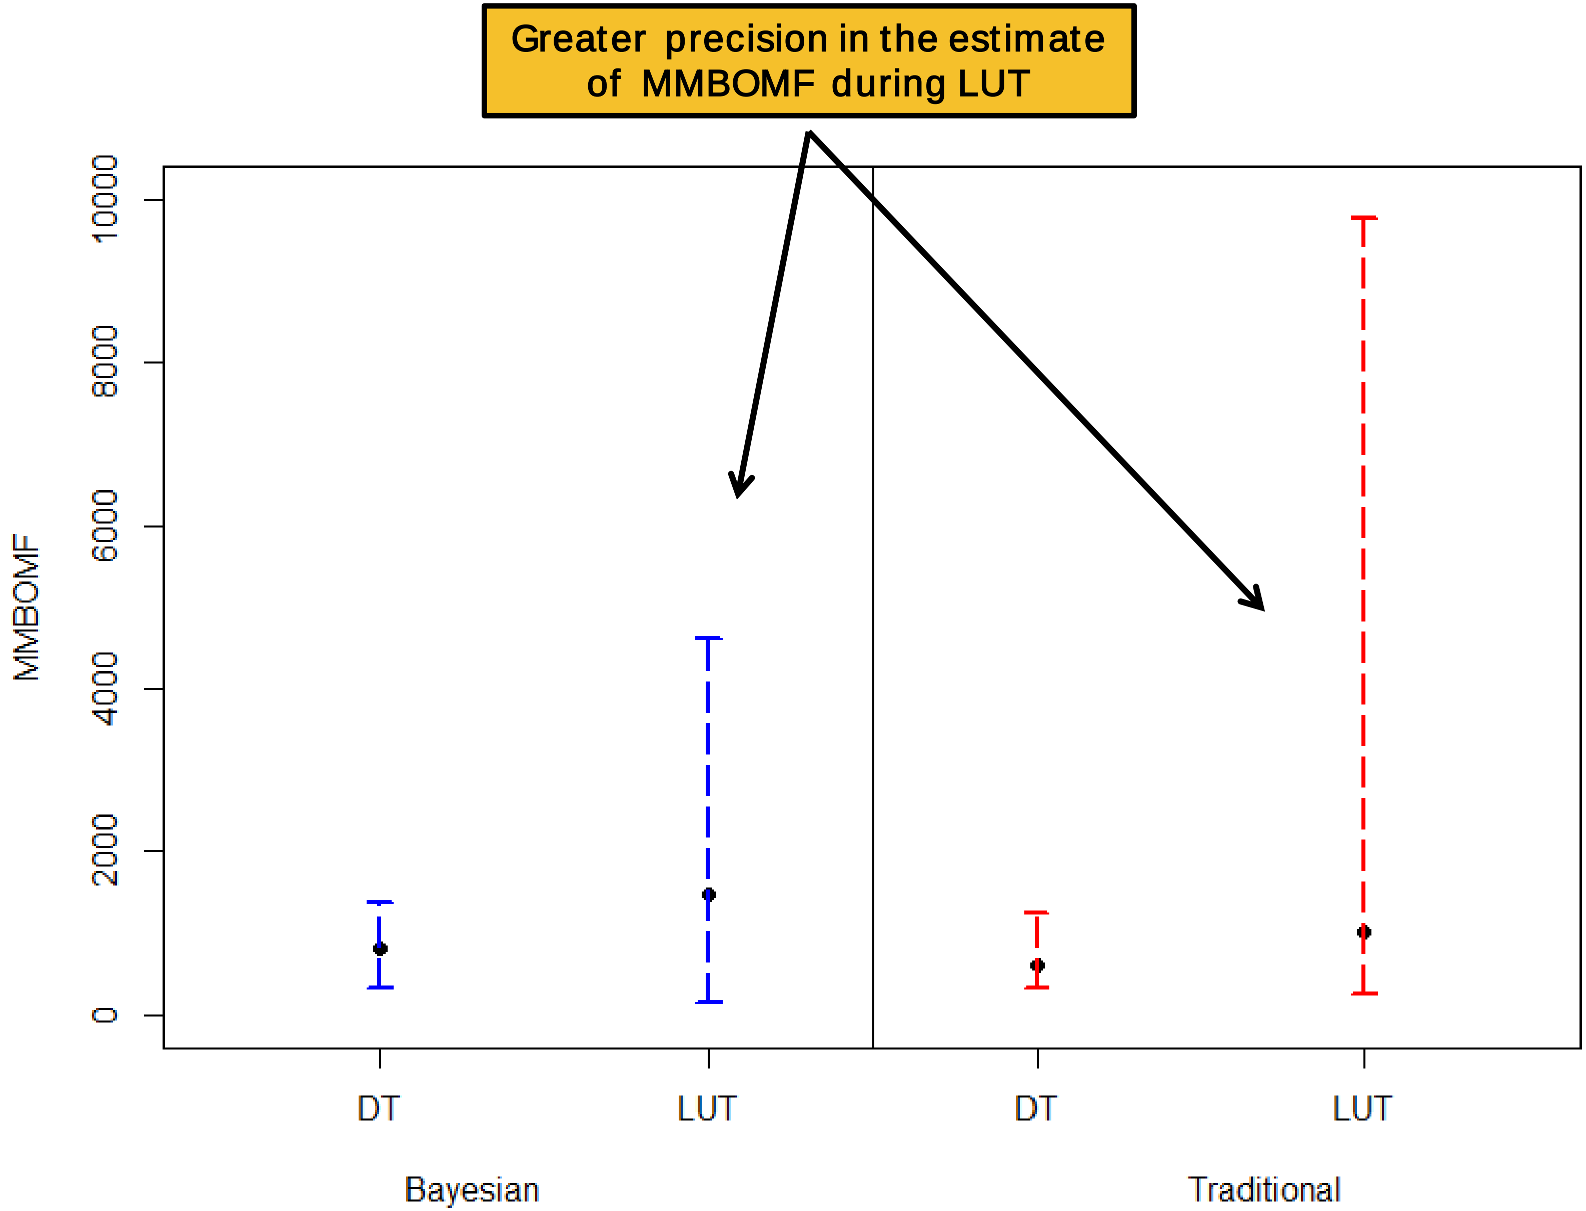

Case Study: MaxxPro - Results

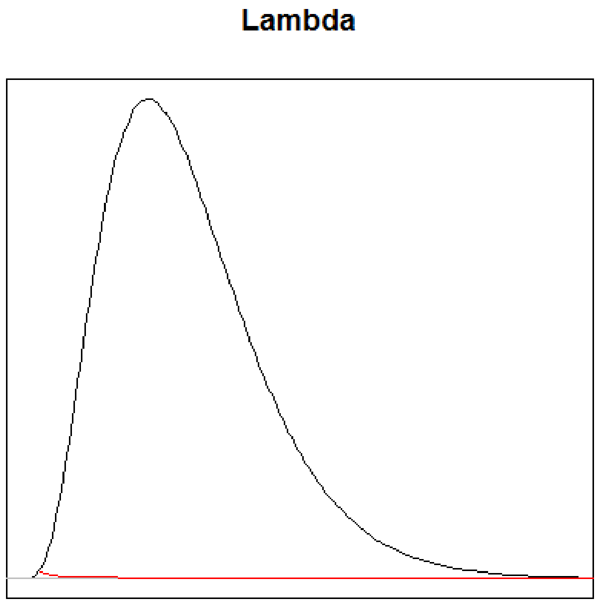

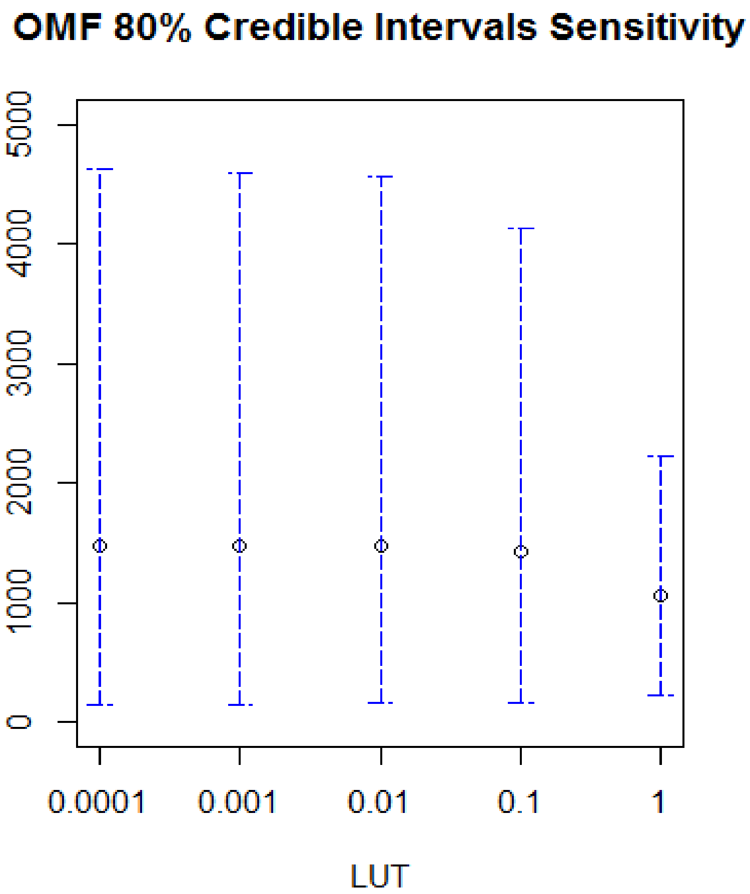

It is always a good idea to check the sensitivity of your prior assumptions on your results!!

Case Study: MaxxPro – Value of Bayesian Statistics

- Instead of assuming that testing is equivalent in DT and LUT, we leverage information across the two phases. The model allows for a change (increase or decrease) in the failure rate from DT to the LUT.
- A sensitivity analysis is performed to assess the robustness of the model. Mean Miles Between Operational Mission Failure (MMBOMF) interval estimates are comparable until we add the equivalent of one failure of information.
- Use of diffuse priors does not include a lot of information, but can be troublesome with few failures.

Fundamentals of Bayesian Analysis

Case Study: Littoral Combat Ship

Case Study: MaxxPro LWB Ambulance

Case Study: Bio-chemical Detection System

Discussion

Case Study: BDS

- The Bio-chemical Detection System analyzes environmental samples and identifies chemical, biological, radiological agents. Each subsystem is comprised of a collection of components of various sensitivity.
- KPP performance requirement for each subsystem: **detect 85-90%** of samples that come into the lab.

- Multiple Tiers of testing
  - Tier 2 (vendor testing)
    - ✓ 5 components: ~2000 trials
  - Tier 3 (vendor testing)
    - ✓ 8 components: ~3600 trials
  - DT/ OT (government testing)
    - ✓ 80-90 trials on multiple components
    - ✓ Final call made by operator

Case Study: BDS

- DT/ OT: set concentration levels, comparatively small sample size

- Standard logistic regression on the Tier 3 data could be problematic
  - All detections or non-detections

- Bayesian approach with a dispersed prior:

$$logit(P_D) = \beta_1 * conc + \beta_2^{matrix} + \beta_3^{agent}$$
$$(\beta_1, \boldsymbol{\beta_2}, \boldsymbol{\beta_3}) \sim Multivariate\ Normal(\mathbf{0}, \boldsymbol{W})$$

» Explicitly forcing a dependence on concentration.

» Leverage all device runs to learn about each agent/ matrix combination performance curve.

Case Study: BDS – Posterior Computation

```
> post=function(beta1,beta2,beta3){
+                 betavec=c(beta1,beta2,beta3)
+
+                 val=dtnorm(betavec[1],0,10^3,lower=0,log=TRUE)+
+
+                   dmvnorm(betavec[-1],rep(0,length(betavec[-1])),diag(rep(10^3,length(betavec)-1)),log=TRUE)+
+
+                   sum(dbinom(y,1,mylogit(beta1*log(X[,1])+beta2[X[,2]]+beta3*X[,3]),log=TRUE))
+
+
+                 return(val)
+         }
```

Case Study: BDS – Posterior Computation

```
> for(i in 1:size){
+
+               #update beta1
+               cand.beta1=rnorm(1,beta1,0.1)
+               if(cand.beta1 > 0){
+               r = post(cand.beta1,beta2,beta3) - post(beta1,beta2,beta3)
+               u = runif(1) <= exp(r)
+               arate1 = arate1 + u
+               beta1 = cand.beta1*(u==1) + beta1*(u==0)}
+
+               #update beta2
+               for(j in 1:length(unique(X[,2]))){
+                       cand.beta2=beta2
+                       cand.beta2[j]=rnorm(1,beta2[j],0.7)
+                       r = post(beta1,cand.beta2,beta3) - post(beta1,beta2,beta3)
+                       u = runif(1) <= exp(r)
+                       arate2[j] = arate2[j] + u
+                       beta2[j] = cand.beta2[j]*(u==1) + beta2[j]*(u==0)}
+
+               #update beta3
+               for(j in 2:length(unique(X[,3]))){
+                       cand.beta3=beta3
+                       cand.beta3[j]=rnorm(1,beta3[j],0.7)
+                       r = post(beta1,beta2,cand.beta3) - post(beta1,beta2,beta3)
+                       u = runif(1) <= exp(r)
+                       arate3[j] = arate3[j] + u
+                       beta3[j] = cand.beta3[j]*(u==1) + beta3[j]*(u==0)}
+
+
+
+               par[i,] = c(beta1,beta2,beta3) #save current parameter values
+
+       }
```

Case Study: BDS – Posterior Computation

```
#update beta2
for(j in 1:length(unique(X[,2]))){
        cand.beta2[j]=rnorm(1,beta2[j],0.7)
        r = post(beta1,cand.beta2,beta3) - post(beta1,beta2,beta3)
        u = runif(1) <= exp(r)
        arate2[j] = arate2[j] + u
        beta2[j] = cand.beta2[j]*(u==1) + beta2[j]*(u==0)}
```

Case Study: BDS – Posterior Computation

arm library: *bayesglm* is a simple alteration of glm that uses an approximate EM algorithm to update the regression coefficients at each step using an augmented regression to represent the prior information.

```
>
> library(arm)
>
> M1 = bayesglm(y~0+log(X[,1])+as.factor(X[,2])+as.factor(X[,3]),
+ family=binomial(link="logit"),prior.scale=10,prior.df=Inf)
>
```

Case Study: BDS – Posterior Computation

MCMCpack library: functions to perform Bayesian inference using posterior simulation for a number of statistical models, all models return coda mcmc objects that can then be summarized using the coda package.

```
>
> library(MCMCpack)
Loading required package: coda
Loading required package: MASS
##
## Markov Chain Monte Carlo Package (MCMCpack)
## Copyright (C) 2003-2016 Andrew D. Martin, Kevin M. Quinn, and Jong Hee Park
##
## Support provided by the U.S. National Science Foundation
## (Grants SES-0350646 and SES-0350613)
##
> M2 = MCMClogit(y~0+log(X[,1])+as.factor(X[,2])+as.factor(X[,3]),b0=0,B0=0.010)
>
```

Case Study: BDS - Results

Component: Chemical Detector 1
Agent: Chemical A
Matrix: Soil, Swab

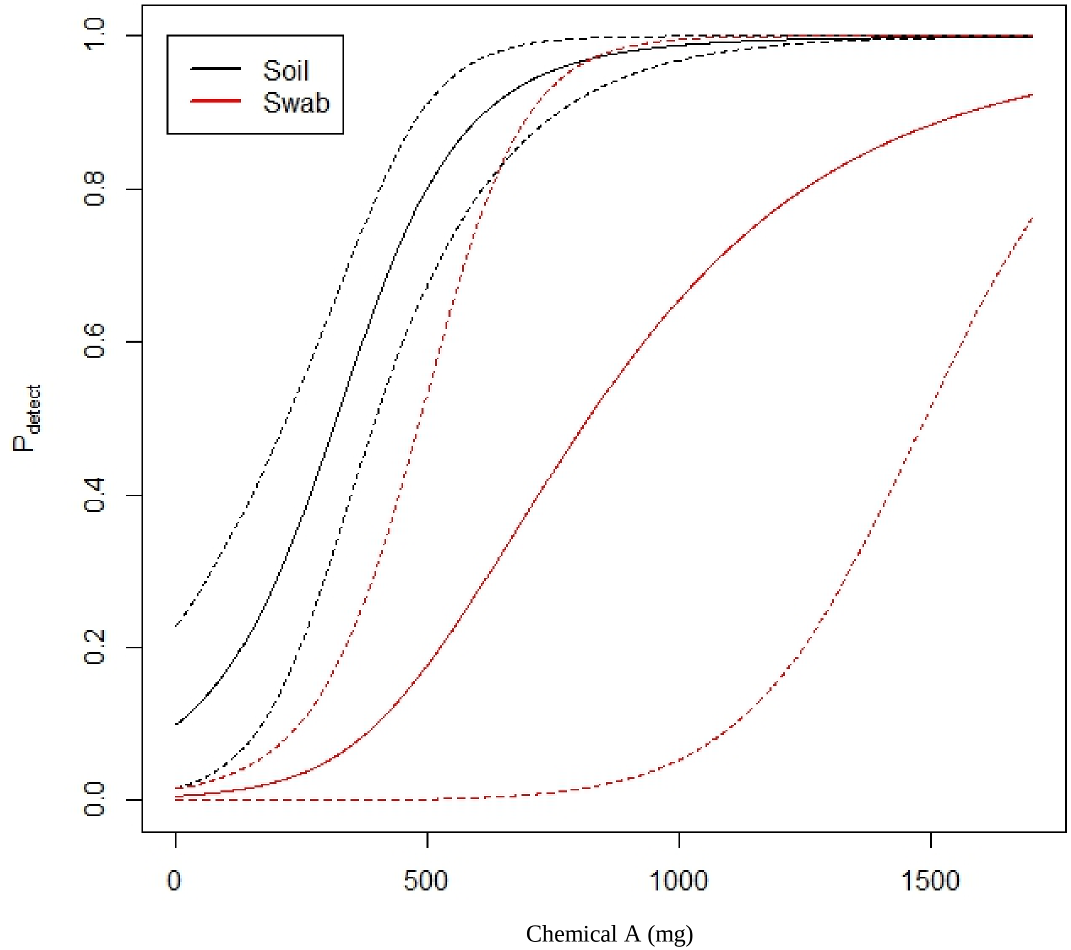

Case Study: BDS - Results

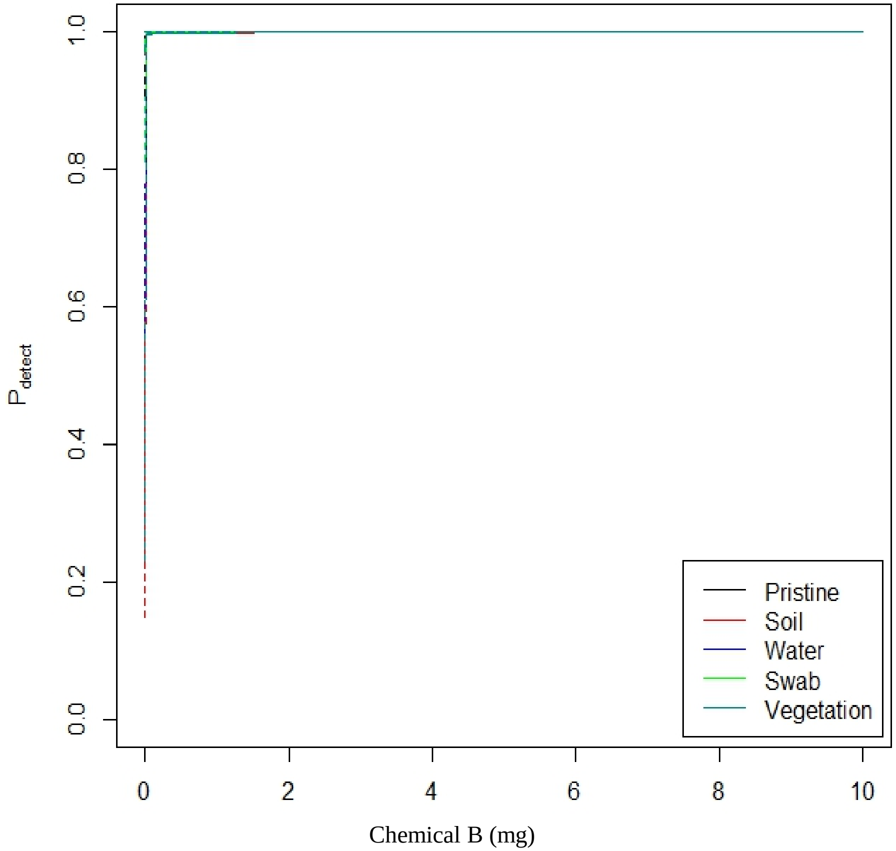

Chemical Detector 2

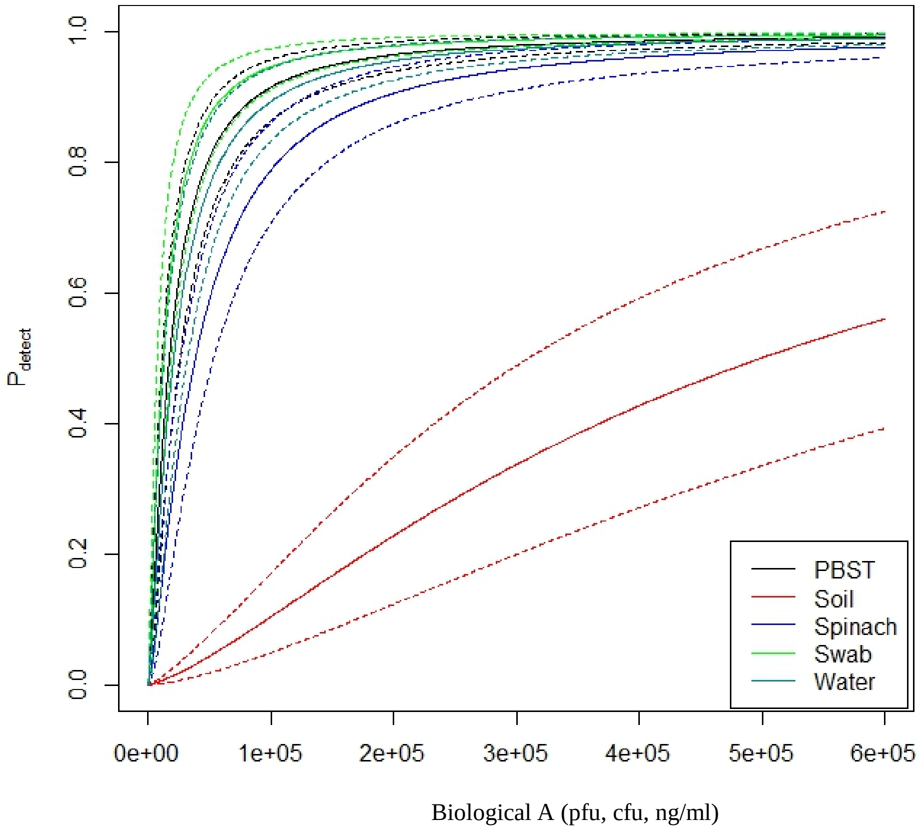

Biological Detector 1

Case Study: BDS - Results

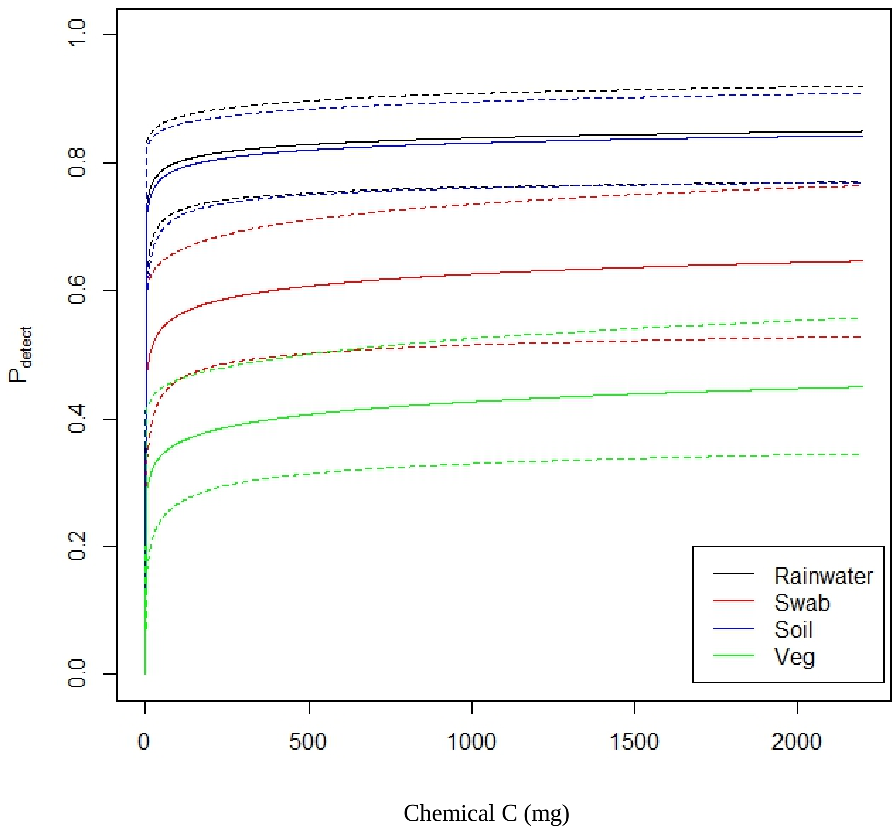

Chemical Detector 3

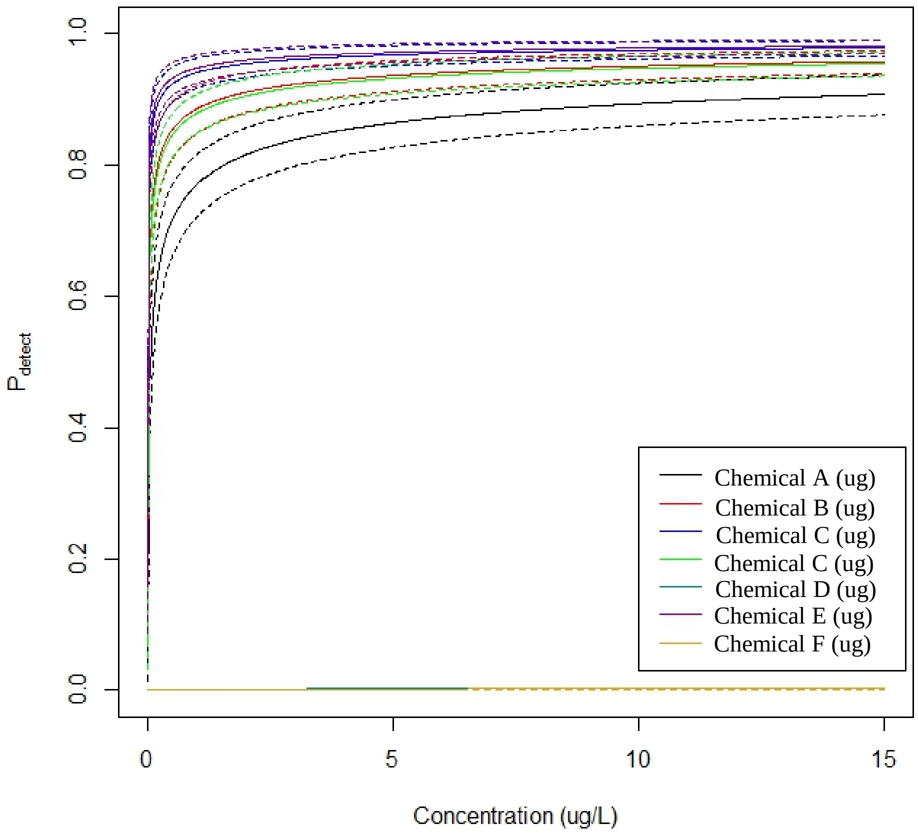

Chemical Detector 4

Case Study: BDS – Value of Bayesian Statistics

- Tier 2 and Tier 3 produced a lot of data which we can leverage to make informed decisions.
- By knowing the concentrations of agents within various matrices that each component can detect, we can determine concentrations that the system of devices might be easy or difficult for the operators to identify in DT/ OT.
- This analysis can serve as the basis for the analysis of the DT/ OT data.

Fundamentals of Bayesian Analysis

Case Study: Littoral Combat Ship

Case Study: MaxxPro LWB Ambulance

Case Study: Bio-chemical Detection System

Discussion

Discussion: When Is it Worth the Effort?

- ✓ Inclusion of prior information from prior testing, modeling and simulation, or engineering analyses only when it is relevant to the current test. We do not want to bias the results.
- ✓ Even when including prior information, the prior must have enough variability to allow the estimates to move away from what was previously seen if the data support such values.
- ✓ We can use very flexible models for many types of test data (e.g. kill chains, complex system structures, linking EFFs to SA) and obtain estimates more readily than with the frequentist paradigm. The model and assumptions have to make sense for the test at hand.

Discussion: Other Resources

For other R packages that provide easy to implement tools and short but informative how to guides with examples, see

https://cran.r-project.org/web/views/Bayesian.html

As with any new statistical method, it is important to have an expert review your work the first few times you apply these techniques.
There are many ways to accidentally do bad statistics!