


Basic Bayes: I

Steven Goodman, MD, PhD
Johns Hopkins University
Baltimore, USA
sgoodman@jhmi.edu

A short research quiz

A well done study is reported on a new electrical stimulator for pain control, and the authors state that it has turned out, somewhat surprisingly (i.e. they thought this would have no more than a 25% chance of being true before the experiment), to be effective in reducing migraine pain, risk $\square=15\%$, 95% CI: 0 to 30%, $p=0.05$. The probability that this association is real is:

-  a.) $< 75\%$
- b.) 75% to 94.99...%
- c.) $\geq 95\%$

S. Goodman, FDA Bayes Workshop

Things I won't say

- That if we turn to Bayesian methods, all our problems will go away.
- That the only “right thinkers” in the statistics world are Bayesian.
- That the Bayesian approach doesn't have difficulties.

S. Goodman, FDA Bayes Workshop

Things I will say

- That if we turn to Bayesian methods, difficult issues will be discussed in the right way by the right people.
- Some of the dilemmas that FDA decision-makers face are artifacts of the statistical methods they use, and not due to demands of the scientific method.
- That the Bayesian perspective provides the best way to think about evidence.

S. Goodman, FDA Bayes Workshop

Things identified as cancer risks (Altman and Simon, JNCI, 1992)

- Electric Razors
- Broken Arms (in women)
- Fluorescent lights
- Allergies
- Breeding reindeer
- Being a waiter
- Owning a pet bird
- Being short
- Being tall
- Hot dogs

Having a refrigerator!!

S. Goodman, FDA Bayes Workshop

A16 THE NEW YORK TIMES NATIONAL WEDNESDAY, JANUARY 6, 1999

Magnets Lessen Foot Pain Of Diabetics, a Study Finds

By HOLCOMB B. NOBLE

In one of the first scientific studies of the centuries-old and highly debated use of magnets for treatment of medical disorders, a New York neurologist reported today that he had significantly lessened the foot pain that afflicts millions of diabetics.

Dr. Michael I. Weintraub, a clinical professor of neurology at New York Medical College, emphasized that his study was small, involving only 24 patients, and must be regarded as preliminary to much more clinical research. But he said that the early results were clear and that the treatment ought to be put to use immediately, provided the correct magnets are used and the treatment is limited to the types of pain that have been studied.

The study, which appears in this month's *American Journal of*

facts have established that they actually work.

In November 1997, reporting in the Archives of Physical and Rehabilitation Medicine, Dr. Carlos Vallbona of the Baylor College of Medicine, Houston said that he applied low-intensity magnets to his own knee pain and that the pain was gone in minutes. He then did a small study of patients with post-polio-syndrome pain. One group was exposed to small magnets, the other to sham magnets. The patients with real magnets reported a 50 percent reduction in pain, while the others reported less than 10 percent.

hands and feet.

In July 1997, Dr. ... a four-month study of ... diabetic and nondiabetic ... Twenty-four patients with ... foot pain caused by diabetes, multiple myeloma, uremia, ischemia, lupus and alcoholism were enrolled in a randomized placebo study.

In the first month, each patient was given a pad equipped with ...

“We have no idea how or why the magnets work.”

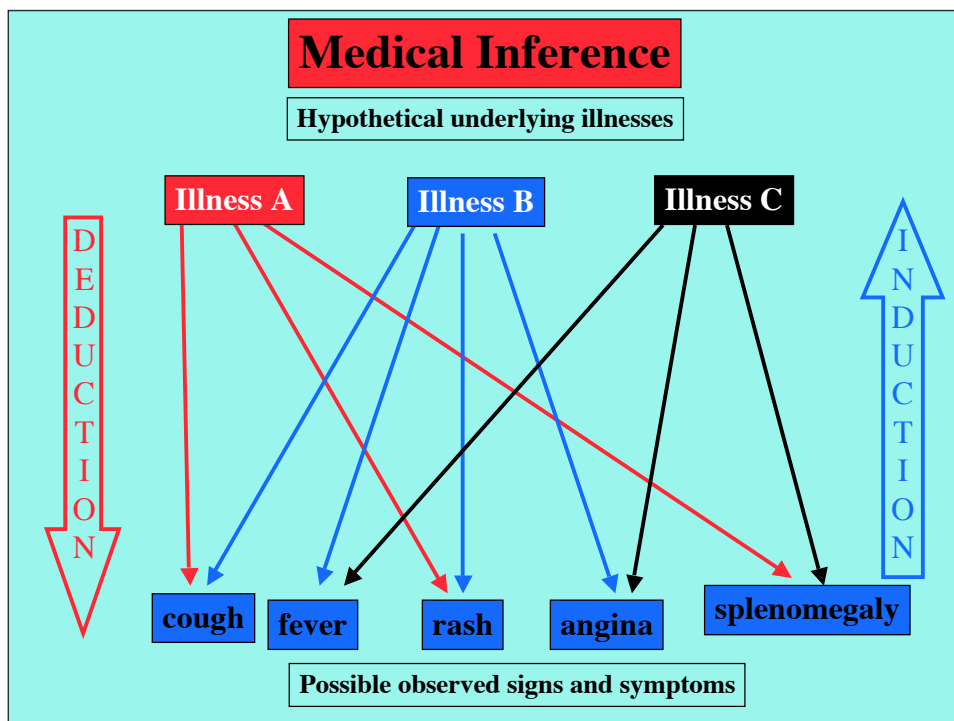
“A real breakthrough...”

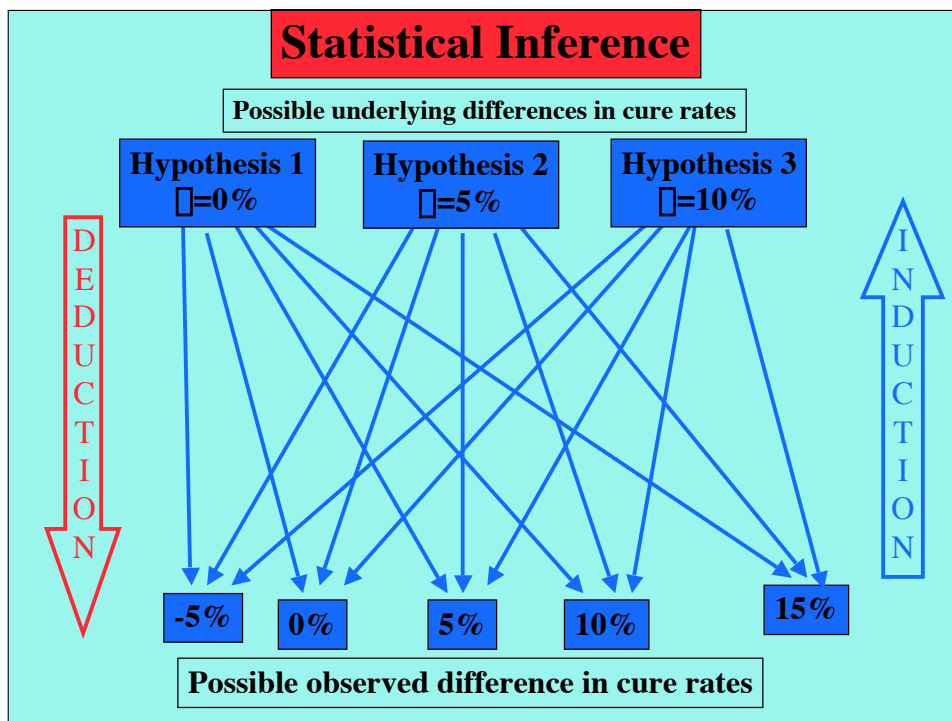
“...the [study] must be regarded as preliminary....”

“But...the early results were clear and... the treatment ought to be put to use immediately.”

A finding that runs counter to many previous studies.

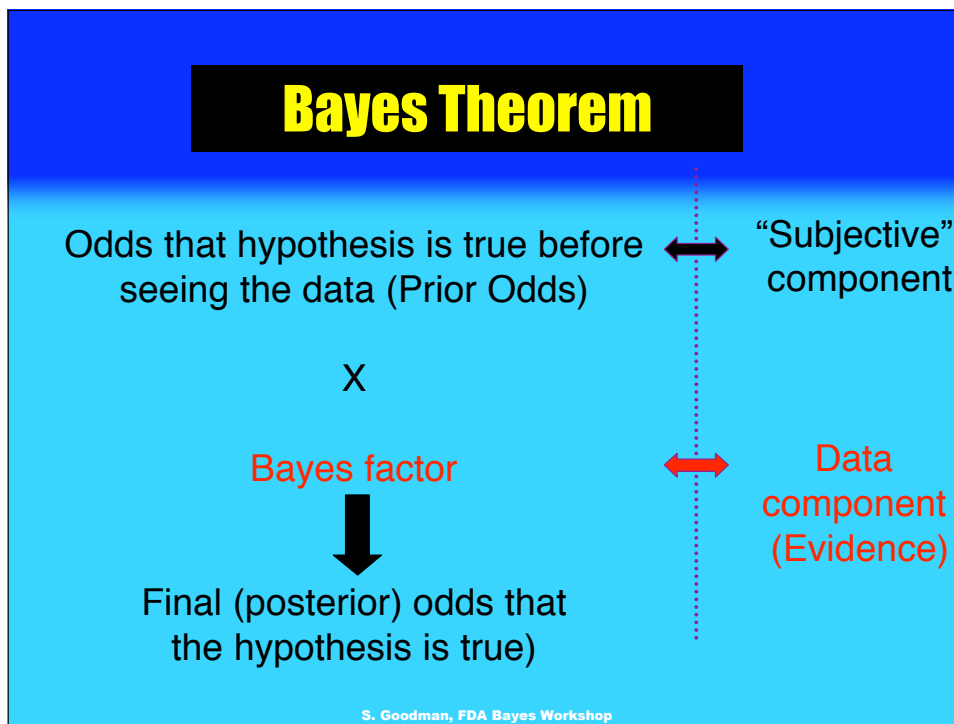
Inference





Statistical inference

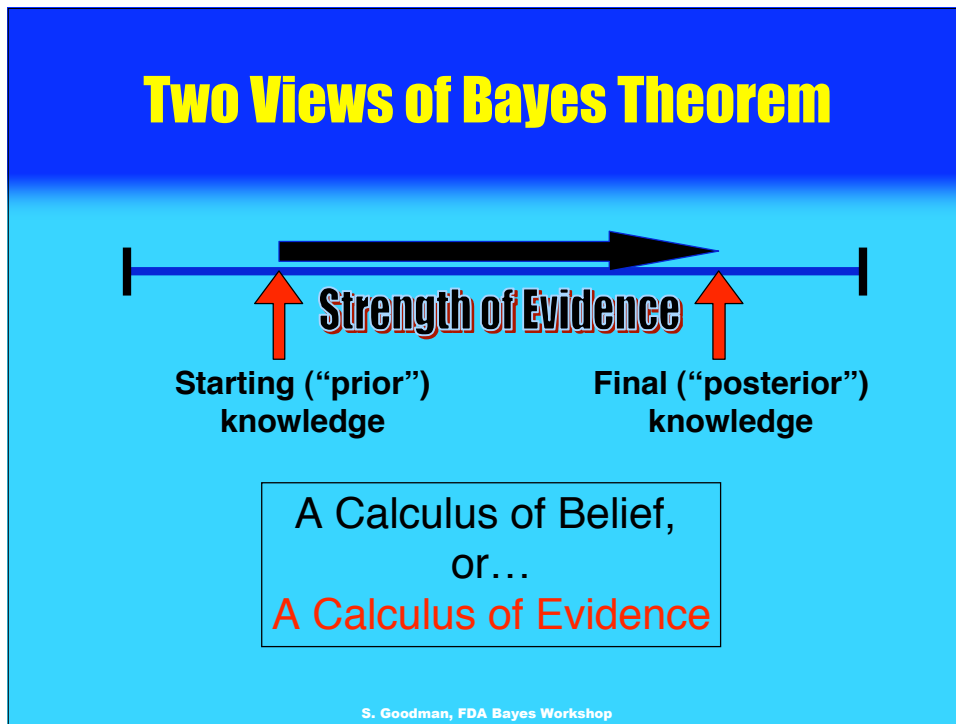
- There is only one formal, coherent calculus of statistical inference: Bayes Theorem.
- “Traditional” statistical rules of inference are a collection of principles and conventions to avoid errors over the long run. *They do not tell us how likely our claims are to be true, nor do they easily apply to individual results.*



Bayes Theorem

$$\underbrace{\frac{\Pr(H_0 | \text{Data})}{\Pr(H_1 | \text{Data})}}_{\text{Post-test Odds}} = \underbrace{\frac{\Pr(H_0)}{\Pr(H_1)}}_{\text{Pre-test Odds}} \square \underbrace{\frac{\Pr(\text{Data} | H_0)}{\Pr(\text{Data} | H_1)}}_{\substack{\text{Likelihood Ratio} \\ \text{a.k.a. Bayes factor}}}$$

S. Goodman, FDA Bayes Workshop

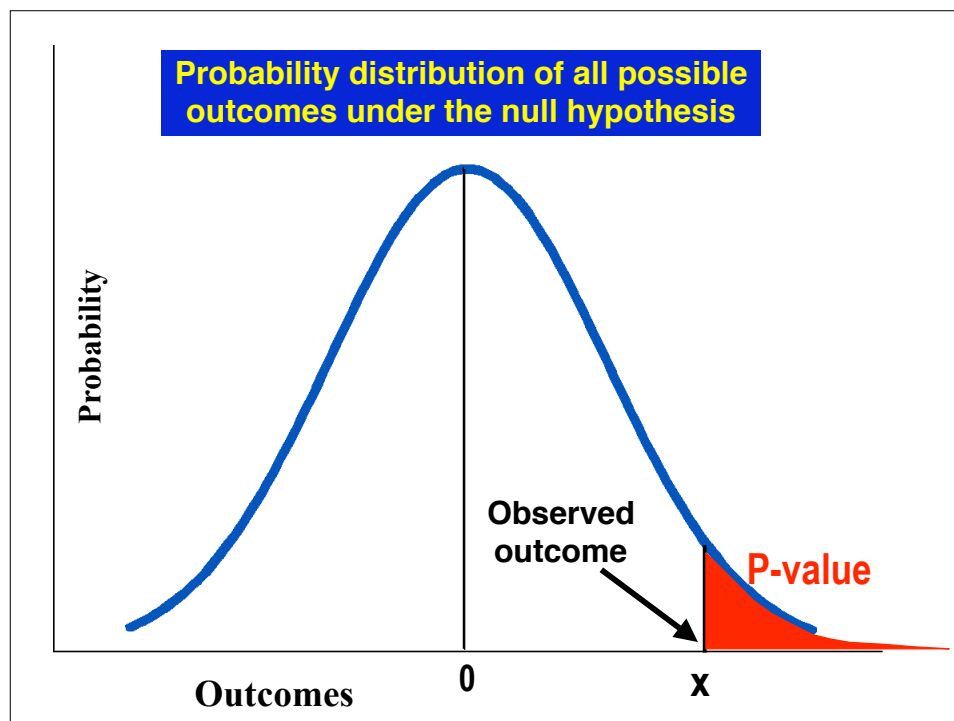


P-values

RA Fisher on statistical education

“I am quite sure it is only personal contact with ... the natural sciences that is capable to keep straight the thought of mathematically-minded people...I think it is worse in this country [the USA] than in most, though I may be wrong. Certainly there is grave confusion of thought. We are quite in danger of sending highly trained and intelligent young men out into the world with tables of erroneous numbers under their arms, and with a dense fog in the place where their brains ought to be. In this century, of course, they will be working on guided missiles and advising the medical profession on the control of disease, and there is no limit to the extent to which they could impede every sort of national effort.” 1958

S. Goodman, FDA Bayes Workshop



Meaning of the p-value

Probability?

Plausibility?

Possibility?

Publish!!!

S. Goodman, FDA Bayes Workshop

The P-value is...

- The probability of getting a result as or more extreme than the observed result, if the null hypothesis (of chance) were true.
- Since the p-value is calculated *assuming the null hypothesis to be true*, it cannot represent the *probability of the truth of the null hypothesis*.

S. Goodman, FDA Bayes Workshop

The P-value is not...

- “The probability of the null hypothesis.”
- “The probability that you will make a Type I error if you reject the null hypothesis.”
- “The probability that the observed data occurred by chance.”
- “The probability of the observed data under the null hypothesis.”

S. Goodman, FDA Bayes Workshop

FDA Discussion

(Fisher, CCT, 20:16-39,1999)

L. Moyé, MD, PhD

“What we have to wrestle with is how to interpret p-values for secondary endpoints in a trial which frankly was negative for the primary. ...In a trial with a positive endpoint...you haven't spent all of the alpha on that primary endpoint, and so you have some alpha to spend on secondary endpoints....In a trial with a negative finding for the primary endpoint, you have no more alpha to spend for the secondary endpoints.”

S. Goodman, FDA Bayes Workshop

FDA Discussion, cont.

(Fisher, CCT, 20:16-39,1999)

Dr. Lipicky: What are the p-values needed for the secondary endpoints? ...Certainly we're not talking 0.05 anymore. ...You're out of this 0.05 stuff and I would have like to have seen what you thought was significant and at what level...

What p-value tells you that it's there study after study?

Dr. Konstam: ...what kind of statistical correction would you have to do that survival data given the fact that it's not a specified endpoint? I have no idea how to do that from a mathematical viewpoint.

S. Goodman, FDA Bayes Workshop

Likelihood

Definition of Likelihood

- The degree to which a hypothesis predicts the data (probability) is proportional to the support that the data gives the hypothesis (likelihood).

If $\Pr(\text{Data} \mid \text{Hypothesis}) = p$

Then

Likelihood (Hypothesis \mid Data) = $c \times p$

where c = arbitrary constant

S. Goodman, FDA Bayes Workshop

Bayes Theorem

$$\underbrace{\frac{\Pr(H_0 \mid \text{Data})}{\Pr(H_1 \mid \text{Data})}}_{\text{Post-test Odds}} = \underbrace{\frac{\Pr(H_0)}{\Pr(H_1)}}_{\text{Pre-test Odds}} \underbrace{\frac{\Pr(\text{Data} \mid H_0)}{\Pr(\text{Data} \mid H_1)}}_{\text{Likelihood Ratio}}$$
$$\frac{L(H_0 \mid \text{Data})}{L(H_1 \mid \text{Data})}$$

S. Goodman, FDA Bayes Workshop

Bayes factor vs. P-value

P-value	Bayes factor
Non-comparative	Comparative
Observed + hypothetical data	Only observed data
Alternative hypothesis implicit, partly data-defined	Alternative hypothesis explicit, pre-defined
Evidence only negative	Evidence negative or positive
Sensitive to stopping rules	Insensitive to stopping rules
No formal justification or interpretation	Formal justification and interpretation

S. Goodman, FDA Bayes Workshop

Calibrating LRs

Strength of Evidence	LR	Final probability when prior probability =		
		25%	50%	75%
Zero	1	25	50	75
Moderate	5	62	83	94
Mod/Strong	10	77	91	97
Strong	20	83	95	98
Very Strong	40	93	97.5	99
Very Strong	80	96	99	99.6

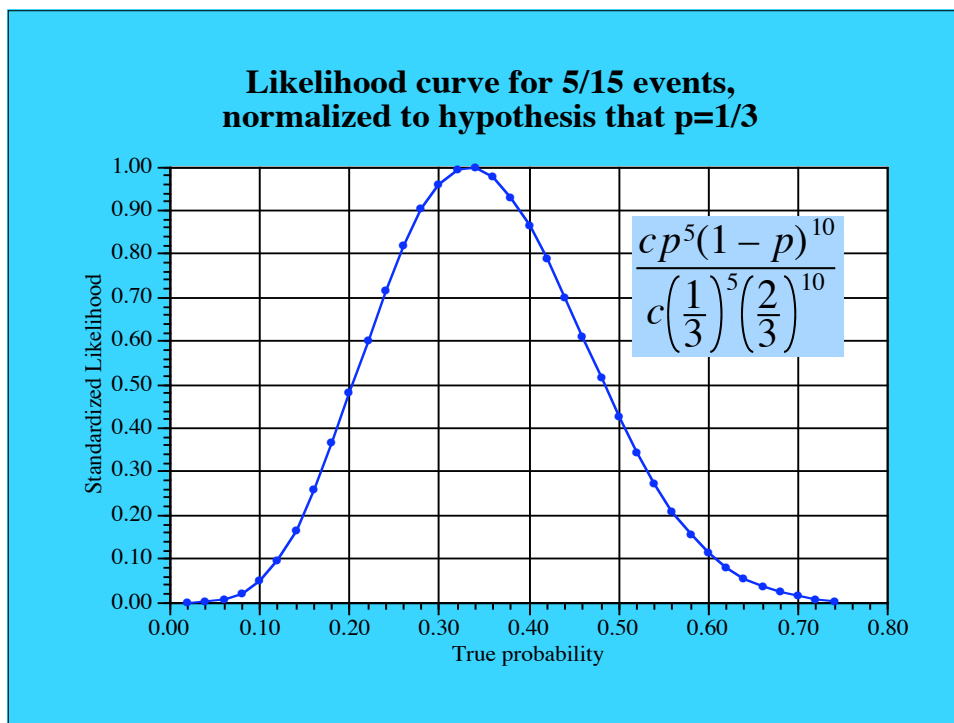
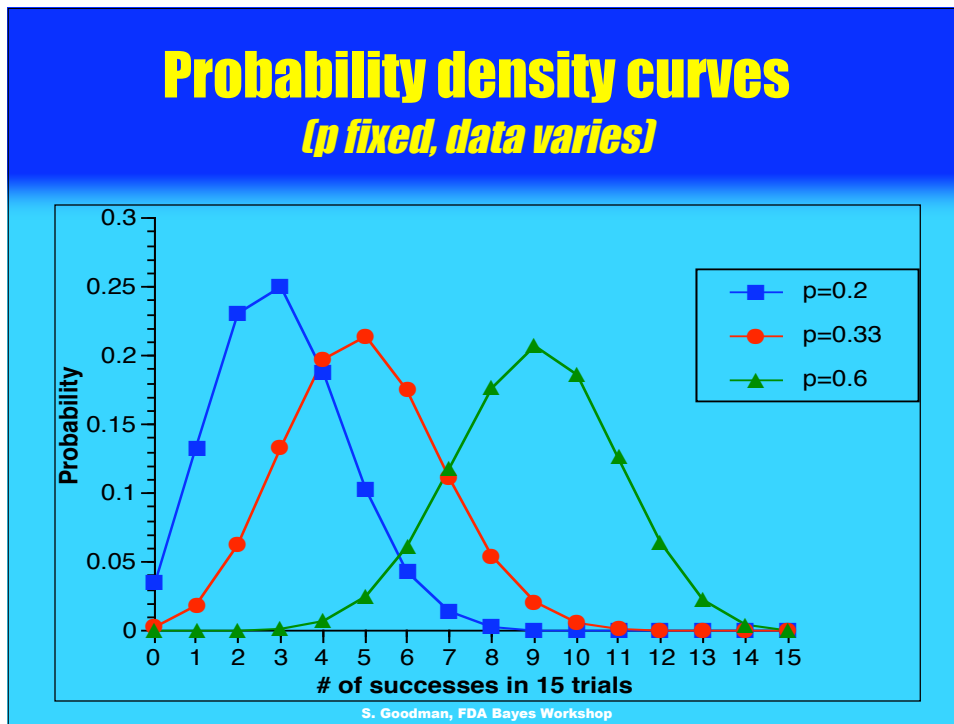
S. Goodman, FDA Bayes Workshop

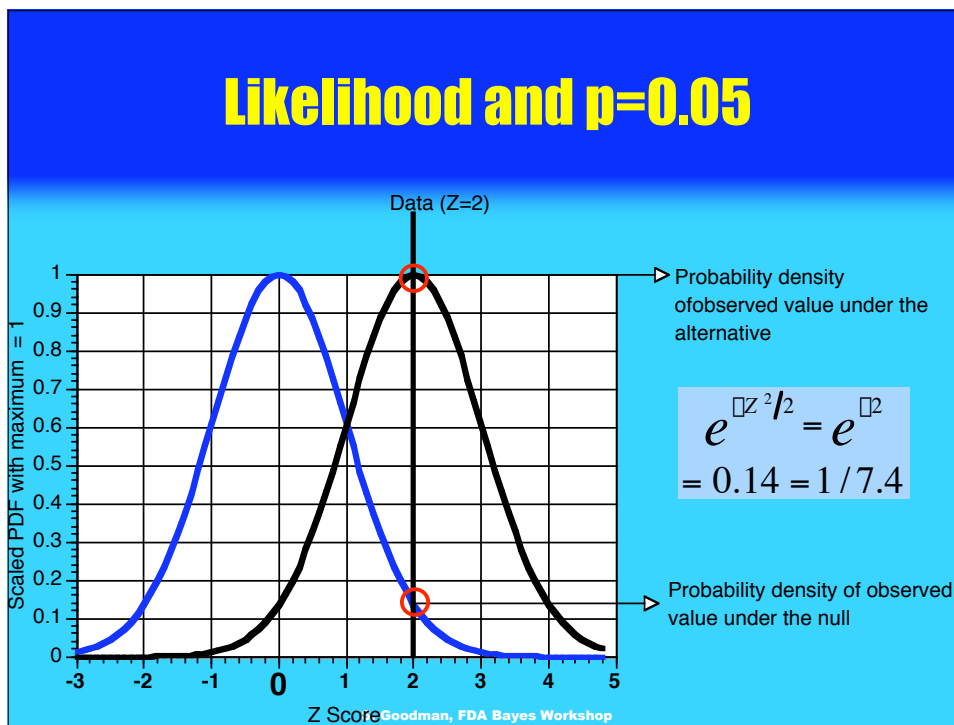
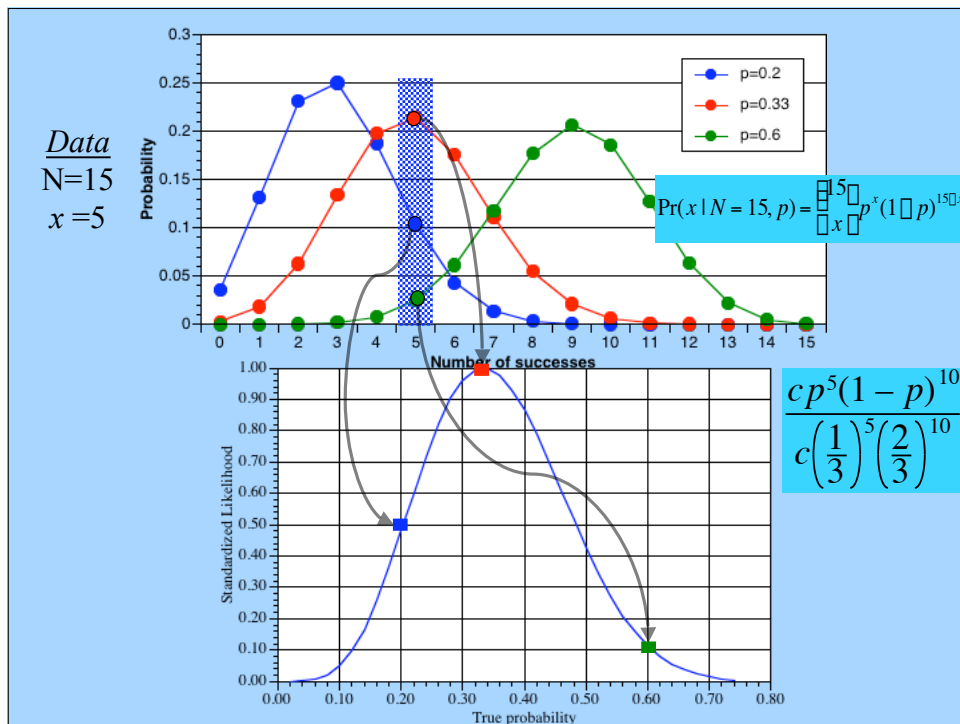
Examples of hypotheses

- Cure rate = 15% (Simple)
- Cure rate $>$ 15% (Composite)
- Treatment difference = 0 (Simple)
- Treatment is beneficial (Composite)
- Treatment is harmful (Composite)

S. Goodman, FDA Bayes Workshop

Understanding Likelihood Functions



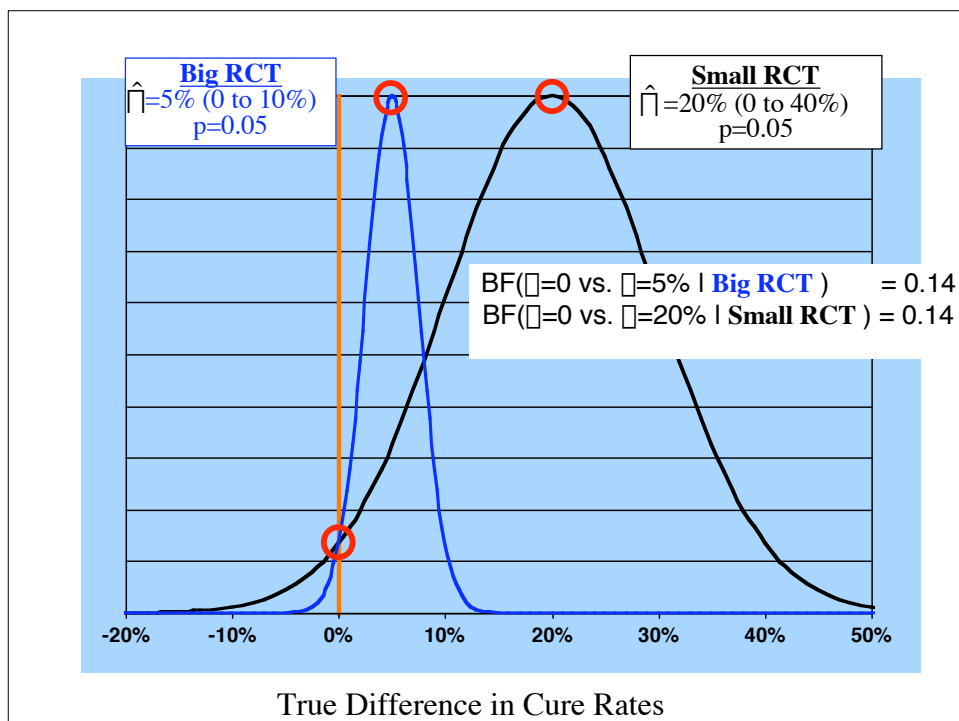




Standardized gaussian likelihood

$$\frac{L(\theta = 0 \mid \bar{x})}{\text{Max}_{\theta} L(\theta \mid \bar{x})} = e^{-z^2/2}$$

- The ratio of the data's probability under the null hypothesis versus the hypothesis that the observed effect is the true one.
- The smallest possible likelihood ratio (or Bayes Factor) for the null hypothesis versus any other hypothesis.

S. Goodman, FDA Bayes Workshop



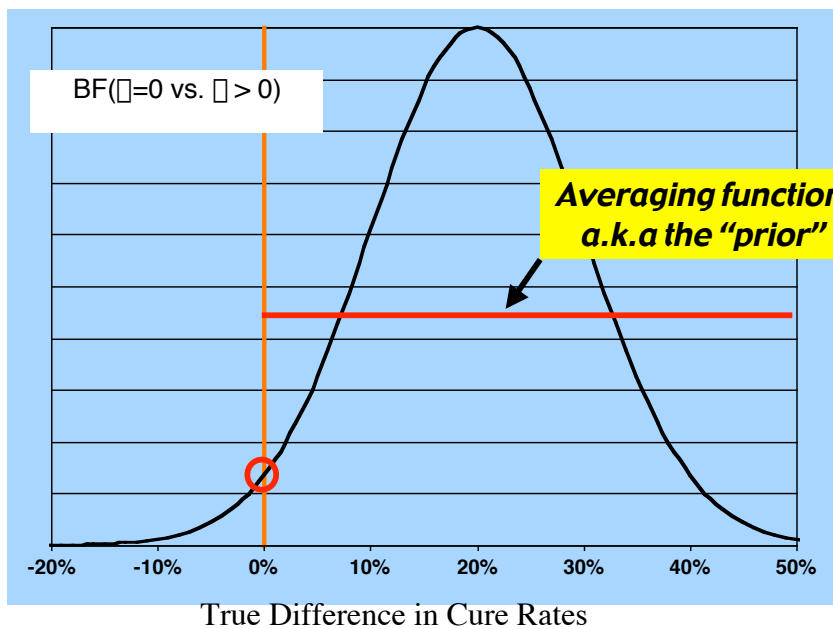


Dependence of Evidence on Alternative Hypothesis

Alternative Hypothesis	Data (P=0.05)	BF (H_0 vs. H_1)
$\delta = 5\%$	Big Trial (5%)	0.14
$\delta = 20\%$	Small Trial (20%)	0.14
$\delta = 5\%$	Small Trial (20%)	0.4
$\delta = 20\%$	Big Trial (5%)	$> 10^6$

S. Goodman, FDA Bayes Workshop

Averaging the likelihood: The Bayes Factor



P-values: Bayesian Translations

P-value (Z-score)	Minimum Bayes factor	-e p ln(p)	Strength of evidence	Decrease in probability of the null hypothesis, %	
				From	To no less than
0.10 (1.64)	.26	.6	Weak	75 50 17	44 21 5
0.05 (1.96)	.15	.4	Moderate	75 50 26	31 13 5
0.03 (2.17)	.1	.3	Moderate	75 50 33	22 9 5
0.01 (2.58)	.04	.13	Moderate to strong	75 50 60	10 3.5 5
0.001 (3.28)	.005	.02	Strong to very strong	75 50 92	1 0.5 5

Stopping Rules

Stopping Rule “Paradox”

The probability of misleading evidence (small p-value) approaches 100% as # of looks \rightarrow .

$$\Pr(p < \alpha) | H_0 \rightarrow 1$$

But!!

$$\Pr(BF < \alpha | H_0) \leq \alpha$$

The Type I error rate has a relationship to evidential strength, *but only when the evidence is measured properly.*

S. Goodman, FDA Bayes Workshop

P-values = Data dredging

- The high Type I error rate with multiple looks is created by is produced when we summarize the likelihood curve at a point determined by the data.
- Using p-values is like data dredging, in that we measure the evidence for a data-suggested hypothesis instead of averaging the evidence over pre-specified simple hypotheses.
- The optimal averaging is done with a Bayesian prior. ***This is why Bayesian methods can be viewed as a “calculus of evidence” as well as a “calculus of belief.”***

S. Goodman, FDA Bayes Workshop

What FDA Needs to Know About Bayesian Statistics

- That Bayes theorem has separable data and belief components, and can be viewed as a calculus of evidence, not just belief.
- That likelihood-based evidential measures have very attractive frequentist properties, as well as a sound theoretical foundation and intuitive interpretations.
- That standard inferential methods represent evidence inappropriately, and produce unnecessary rigidity in design and interpretation.

S. Goodman, FDA Bayes Workshop

And...

- That the use of Bayesian evidential measures can have an impact far beyond the (sometimes) different numbers they produce; they can affect how we talk about evidence, and who participates in that dialogue.

S. Goodman, FDA Bayes Workshop

Final thoughts

“What used to be called judgment is now called prejudice, and what used to be called prejudice is now called the null hypothesis....it is dangerous nonsense (dressed up as ‘the scientific method’) and will cause much trouble before it is widely appreciated as such.”

A.W.F. Edwards (1972)

S. Goodman, FDA Bayes Workshop