

Foundations of Statistics – Frequentist and Bayesian

Mary Parker, University of Texas at Austin and Austin Community College, 2005
parker@math.utexas.edu

“Statistics is the science of information gathering, especially when the information arrives in little pieces instead of big ones.” – Bradley Efron

This is a very broad definition. How could we possibly come up with a structured way of doing this? Actually, a variety of structures have been developed. The purpose of this talk is to give you some understanding of more than one of these structures.

If you had a statistics course in college, it probably described the “frequentist” approach to statistics. A few of you might possibly have had a second or later course that also did some Bayesian statistics. In my opinion, and in the opinion of many academic and working statisticians, statistical practice in the world is noticeably changing. The courses for most undergraduates haven’t begun to catch up yet!

Typical Conclusions		
	Frequentist	Bayesian
Estimation	I have 95% confidence that the population mean is between 12.7 and 14.5 mcg/liter.	There is a 95% probability that the population mean is in the interval 136.2 g to 139.6 g.
Hypothesis Testing	If H_0 is true, we would get a result as extreme as our data only 3.2% of the time. Since that is smaller than 5%, we would reject H_0 at the 5% level. These data provide significant evidence for the alternative hypothesis.	The odds in favor of H_0 against H_A are 1 to 3.

Those of you who have taken statistics courses may recognize the first set of these as typical conclusions from frequentist statistics. The second set lists typical conclusions from a Bayesian perspective. I think most of us would agree that the conclusions in the second set are easier for most readers to understand.

So why don’t we all do Bayesian statistics?

Short answer: There are two reasons.

1. The calculations needed for Bayesian statistics can be overwhelming.
2. The structure requires a “prior distribution” on the parameter of interest. If you use a different prior you will obtain different results and this “lack of objectivity” makes some people uncomfortable.

In the last twenty years or so, advances in computing have made it possible to do calculations that earlier would have been completely impractical. This has generated a resurgence of interest in Bayesian methods and it is a rapidly growing field. Mostly that hasn’t reached the undergraduate curriculum, especially the introductory courses.

My goal in this talk is to help you understand the basic philosophical differences between frequentist and Bayesian statistics. My examples are quite simplified, and don’t do justice to the most interesting applications of these fields. They are chosen to illustrate the mathematics used to derive these conclusions. I encourage each of you to browse in some of the easier references, particularly those on the web, to see additional examples.

Frequentist Estimation – An Informal Introduction

Example: Mothers' breast milk includes calcium. Some of that calcium comes from the food the mothers eat and some comes from the mothers' bones. Researchers measured the percent change in calcium content in the spines of a random sample of 47 mothers during the three months of breast-feeding immediately after the baby was born. We want to use these data to estimate the mean (average) percent of change in the bone calcium level for the population of nursing mothers during the first three months of breast feeding.

From the data, we find that the sample mean is $\bar{X} = -3.6$ and the standard deviation is 2.5.

Of course, our point estimate of the mean percent of change in the bone calcium level is -3.6 . Generally, however, we are also interested in some measure of variability. For instance, we might want an interval estimator, to say "We estimate that the mean percent of change in the bone calcium level is between ___ and ___."

In frequentist statistics, inferences such as this are based solely on the sampling distribution of the statistic. In this example, the appropriate estimator is the sample mean.

To construct the sampling distribution of the sample mean when $n = 47$:

1. Take all possible samples of size 47 from the population.
2. For each sample, compute the sample mean, \bar{X}_{47} .
3. Consider the distribution of all those different values of \bar{X}_{47} . That is the sampling distribution of \bar{X}_{47} .

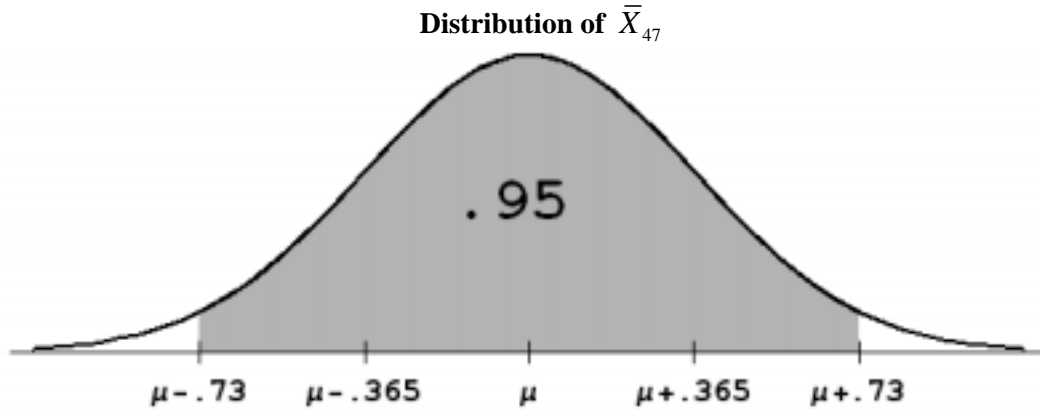
We use some results from probability theory to say that, since $n = 47$ is a large enough sample (from a population for which we call the mean μ and standard deviation σ), the distribution of \bar{X}_{47} , which has mean $\mu_{\bar{X}}$ and standard deviation $\sigma_{\bar{X}}$,

- is approximately normal,
- has mean equal to the mean of the original population $\mu_{\bar{X}} = \mu$, and
- has standard deviation which can be approximated by $\sigma_{\bar{X}} = \frac{\sigma}{\sqrt{n}}$.

Further, we typically will choose to estimate the standard deviation of the population, σ , by the standard deviation of the sample s , so we conclude that, for the nursing mothers example, the sampling distribution for our \bar{X}_{47} is approximately normal, with mean equal to the mean of the original population

and standard deviation approximated by $s_{\bar{X}} = \frac{2.5}{\sqrt{47}} = 0.365$.

It is a fact about the shape of normal distributions that the probability of getting a score within two standard deviations of the mean is very close to 95%. Following is a sketch of the sampling distribution of \bar{X}_{47} , labeled with the mean and the estimated standard deviation from the data.



To obtain an interval estimate for the population mean, μ , we make a conceptual shift (that is often rather murky for our students.)

$\mu \pm 0.73$ includes 95% of the possible \bar{X}_{47} 's (a probability statement about the dist'n of \bar{X}_{47})

shifts to

one particular $\bar{X}_{47} \pm 0.73$ gives a 95% confidence interval for μ

The computational formula we see in our textbooks is $\bar{X} \pm 2.0 \cdot \frac{s}{\sqrt{n}}$. When we plug in the values for our data, we have $-3.6 \pm 2.0 \cdot 0.365 = -3.6 \pm 0.73$.

We interpret that statement: "We have 95% confidence that the actual mean μ of the percent change in blood calcium level for nursing mothers is in the interval between -4.33% and -2.87% ."

Notice here that our conclusion is NOT a probability statement. It is based on a probability statement, but since it is not a statement about a random quantity, it is not a probability. The random quantities in this discussion are the individual X values and all of the possible \bar{X}_{47} values. The quantity μ is not random. It is the fixed, but unknown, mean of the population.

In an elementary statistics class, when we teach this, we encourage students to think of the possibility of doing many different possible samples of size 47, find the \bar{X}_{47} for each one and then use that particular value to form a 95% confidence interval. So we can find numerous different 95% confidence intervals – one for each different possible sample of size 47. The 95% confidence means that, when we do this many, many times, approximately 95% of the resulting intervals will actually include the true population mean.

From the “typical conclusions” on page 1, you may have noticed the following:

1. The frequentist conclusions are based on $P(\text{data}|\text{parameter})$
2. The Bayesian conclusions are based on $P(\text{parameter}|\text{data})$

In the two methods, we observe the data in the same way, and model the distribution of the data in the same way. The differences lie in what we do with the distributions after that. We use Bayes Theorem.

Introduction to Bayes Theorem (The “audience participation part of this talk!”)

Bayes Theorem enables us to “turn around” conditional probabilities.

$$P(\theta | X) = \frac{P(X | \theta) \cdot P(\theta)}{P(X)}$$

Example:

Enzyme immunoassay (EIA) tests are used to screen blood for the presence of antibodies to HIV, the virus that causes AIDS. Antibodies indicate the presence of the virus. The test is quite accurate – about 99% accurate. Specifically, it has 99.85% probability of detecting antibodies when they are present and 99.4% probability of a negative result when antibodies are not present. It is also known that about 1% of the overall population has the antibodies.

When a person has a positive result on this test, how likely is it that they actually carry the antibodies?

Solution: Call A the condition of having the antibodies and $+$ and $-$ the conditions of having a positive and negative result, respectively.

$$\begin{array}{lll} P(+ | A) = & P(+ | \text{not } A) = & P(A) = \\ P(- | A) = & P(- | \text{not } A) = & P(\text{not } A) = \end{array}$$

Write the question as a conditional probability. $P(A | +)$

Use Bayes Theorem to “turn around” the probabilities.

$$P(A | +) = \quad \text{(write Bayes Thm here.)}$$

$$P(A | +) = \quad \text{(expand the denominator so that you can compute it)}$$

$$P(A | +) = \quad \text{(plug in numbers)}$$

(Look to the end of this handout to see the worked-out solution.)

Most people find that this is a surprising result. It’s not even close to 99%, which was the (correctly) reported accuracy of the tests.

Bayesian Estimation – An Informal Introduction

Example:


I take a coin out of my pocket and I want to estimate the probability of heads when it is tossed. I am only able to toss it 10 times. When I do that, I get seven heads. I ask three statisticians to help me decide on an estimator of p , the probability of heads for that coin.

Case 1. Sue, a frequentist statistician, used $\hat{p} = \frac{X}{10} = 0.7$.

Case 2. Jose, who doesn't feel comfortable with this estimator, says that he already has an idea that p is close to 0.5, but he also wants to use the data to help estimate it. How can he blend his prior ideas and the data to get an estimate? Answer: Bayesian statistics.

Jose makes a sketch of his prior belief about p . He thinks it is very unlikely that p is 0 or 1, and quite likely that it is somewhere pretty close to 0.5. He graphs his belief.

Jose's drawing:



Then he notices that the graph corresponds to a particular probability distribution, which is Beta(5,5). So this is called his prior distribution for p .

Recall that, for a beta distribution with parameters a and b , we have these formulas for the mean and variance.


$$\text{Mean} = \frac{a}{a+b} \quad \text{Variance} = \frac{ab}{(a+b)^2(a+b+1)}$$

Thus the mean is $\frac{5}{5+5} = 0.5$ and the variance is $\frac{5 \cdot 5}{10^2 \cdot 11} = 0.022727$, so the standard deviation is 0.15.


Then, by some magic! (i.e. Bayes Theorem), he combines the data and his prior distribution and gets that the distribution of p , given the data, is a Beta(12,8). This is called the posterior distribution of p . Using those same beta distribution formulas, this mean now is 0.6 and the variance is 0.01143, so the standard deviation is 0.107.

Jose's posterior distribution

Beta(12,8)




Case 3: Vicki, who is very sure that coins are unbiased, has a prior distribution like Jose's, but much narrower. There's a much higher probability on values very close to 0.5. She graphs her belief. Vicki's prior distribution:



She notices that this corresponds to a particular probability distribution, which is Beta(138,138), so that is her prior distribution of p . So the mean is 0.5, the variance is 0.0009, and the standard deviation is 0.03. Notice that her standard deviation is much smaller than Jose's.

Now, she also uses Bayes Theorem to combine the data and her prior distribution and finds that her posterior distribution of p is a Beta(145,141). So her posterior mean is 0.507, variance is 0.0008709, and standard deviation is 0.0295.

Vicki's posterior distribution. Beta(145, 141)



Jose and Vicki are both doing Bayesian estimation. Both of them decide to use the mean of the posterior distribution of the parameter as their estimator.

Summary:

Sue's estimate of the probability of heads: 0.700
 Jose's estimate of the probability of heads: 0.600
 Vicki's estimate of the probability of heads: 0.507

Now, Jennifer offers you a bet. You pick one of these values. She chooses one of the other two. An impartial person tosses the coin 1000 times, and get a sample proportion of heads. If that sample proportion is closer to Jennifer's value, you pay her \$25. If it is closer to yours, she pays you \$25. Which value would you choose?

Overview:

Jose and Vicki are both doing Bayesian estimation. But they get different answers. This is one of the reasons that frequentist statisticians criticize Bayesian methods. It doesn't seem objective enough. Different people can get different estimates based on the same data, they say. And of course, Jose and Vicki did get different estimators. But, were they really using the same data? Well, not if we consider data in the broad sense. If their prior beliefs are considered data, then these are not based on the same data. And, if their prior beliefs are not considered data, then we are back to frequentist statistics. Would any of you be willing to take up Jennifer's bet and choose Sue's frequentist statistics estimator of 0.7? (If so, I'll be willing to bet with you.)

What are some of the difficulties in the Bayesian approach?

1. Quantifying prior beliefs into probability distributions is not simple. First, we haven't all thought much about our prior beliefs about most things, and, even if we have some beliefs, those aren't usually condensed into a probability distribution on a parameter.
2. We might not agree with colleagues on the prior distribution.
3. Even if we can find a formula for the distribution describing our prior beliefs about the parameter, actually doing the probability calculations to find the posterior distribution using Bayes Theorem may be more complex than we can do in closed form. Until people had powerful computing easily available, this was a major obstacle to using Bayesian analysis.

The Bayesian statisticians have some answers for most of this, which will become clearer as we develop the theory.

Mathematical Theory:

How do we do the mathematics to get the posterior distribution?

Use Bayes Theorem, which is a theorem to "turn around" conditional probabilities.

$$P(A|B) = \frac{P(B|A) \cdot P(A)}{P(B)}$$

From our data, we get $f(\text{data}|p)$, which can also be denoted as $L(p)$ and called the likelihood function. We want $h(p|\text{data})$, which is the posterior distribution of p . We also need $g(p)$, which is our prior distribution of p .

For theoretical purposes, (but we often manage to avoid using it), we also need the distribution of the data, not conditioned on p . That means $k(data) = \int_{\text{domain}} f(data|p)dp$, or, if the distribution of p is discrete, $k(data) = \sum_p f(data|p)$. In these, we are integrating or summing over all possible values of p .

So, using Bayes Theorem, we get

$$h(p|data) = \frac{f(data|p) \cdot g(p)}{k(data)}$$

Now, since $k(data)$ is not a function of p , and since we are going to fix our data, then the denominator is a constant, as we calculate $h(p|data)$. Since the only difference that a constant multiple makes in specifying a probability distribution is to make the density function integrate (or sum) to 1 over the entire domain, then we can know the form of the distribution even before we bother with computing the denominator. And, once we know the form of the distribution, we could figure out the constant multiple just by determining what it would take to have the integral over the entire domain be 1. So, in practice, people use this form of Bayes Theorem, where the symbol \propto means “proportional to”, i.e. “is a constant multiple of” and we call the right hand side of this the “unnormalized posterior density”.

$$h(p|data) \propto f(data|p) \cdot g(p)$$

In fact, when we do problems using this, we often neglect the constant multiples (considering p as the variable and the rest as constants) on both $f(data|p)$ and $g(p)$, since if we are going to gloss over one constant multiple, then there is little point in keeping track of the rest until we start trying to really get the constant multiples figured out. And then, when we do need to find the constant multiple, if we recognize the kernel of the distribution (the factors which include the variable) as one of our standard distributions, we can use that to determine the constant multiple.

When you read about Bayesian statistics, you’ll see the statement that Bayesians use the data only through the likelihood function. Recall that the likelihood function is just the same as the joint density function of the data, given p , reinterpreted as a function of p . $L(p) = f(data|p)$

$$h(p|data) \propto f(data|p) \cdot g(p) = L(p) \cdot g(p)$$

Back to the example:

Here, $L(p) = f(data|p) \propto \prod_{i=1}^{10} p^{x_i} (1-p)^{(1-x_i)} = p^7 (1-p)^3$

which is from a binomial distribution, in “unnormalized form”, i.e. without the constant multiple, considering it as a function of p and defined on $p \in [0,1]$.

Also, for Jose, his prior distribution is $g(p) = \frac{\Gamma(10)}{\Gamma(5) \cdot \Gamma(5)} p^{5-1} (1-p)^{5-1} \propto p^{5-1} (1-p)^{5-1}$

so $h(p|data) = f(data|p) \cdot g(p) \propto p^{7+5-1} (1-p)^{3+5-1}$

Since the domain of h is all real numbers between 0 and 1, it is a continuous distribution, so it fits the form of a Beta distribution, with parameters 12 and 8. Thus, Jose has a posterior distribution on p which is Beta(12,8), as we said before, and is used to calculate his estimator, which is the mean of the posterior distribution on p .

We calculate Vicki’s posterior distribution on p in exactly the same manner.

We can use these posterior distributions on p to make probability statements about p as well. Typically Bayesians would compute a HPD (highest posterior density) interval with the required probability.

Back to the Theory:

Since we are using this theory to find a posterior distribution on the parameter, we can use that distribution to find probability intervals for the parameter. In a frequentist statistics course, we learn that parameters aren't random variables, however, so when we find intervals for parameters, we called them confidence intervals to distinguish them from probability intervals. The major conceptual difference between Bayesian statistics and frequentist statistics is that, in Bayesian statistics, we consider the parameters to be random, so one consequence of that is that we can write probability intervals for parameters.

Free the Parameters!

-- Bayesian Liberation Front

Back to the Overview: (How they chose their prior distributions)

No doubt it will have occurred to you that it was very convenient that both Jose's and Vicki's prior beliefs about the parameter fit into a beta distribution so nicely. It was particularly convenient since the data, given the parameter, fit a binomial distribution and so the product of those had such a simple form. Surely, you say, that doesn't always happen in real life! Probably not.

But, in real life, most people don't have fully developed prior beliefs about most of the interesting questions. In fact, you can usually elicit from a person what the average (mean) of their prior belief is, and, after a fair amount of coaxing, you might be able to get their standard deviation. There are convenient forms for prior distributions in a number of cases, based on the distribution of the data. These are called "conjugate priors". If the data is binomial, the conjugate prior is a beta distribution. If the data is normal, the conjugate prior is a normal distribution. So, in practice, you often elicit the mean and standard deviation of a person's prior belief, notice what form the distribution of the data has, and decide to use a conjugate prior. You then use the mean and standard deviation to solve for the parameters in your conjugate prior.

That's how Jose and Vicki did the problem here. The data was binomial, so they needed a beta prior distribution. Recall that, for a beta distribution with parameters a and b , we have these formulas for the mean and variance.

$$\text{Mean} = \frac{a}{a+b} \quad \text{Variance} = \frac{ab}{(a+b)^2(a+b+1)}$$

Each of them wanted a mean of 0.5, so they needed the two parameters to be the same. Then, the smaller that they wanted the standard deviation, the larger the parameters had to be. So each decided on a standard deviation that seemed to reflect his(her) beliefs. Jose chose 0.15 and then wrote

$$(0.15)^2 = \frac{a \cdot a}{(2a)^2 \cdot (2a + 1)} \text{ and solved for } a. \text{ He reduced this to a linear equation and found } a = 5.0556.$$

Since his standard deviation was only an approximation, it seemed most reasonable to just round this off

to use integer parameters. Vicki did a similar calculation, starting from her prior belief, which she characterized as mean 0.5 and standard deviation 0.03.

Notice that, when we look at the forms of the data distribution and the prior distribution, they are similar, and so Jose's prior distribution was the equivalent to starting your problem with 4 heads and 4 tails. Vicki's was equivalent to starting with 137 heads and 137 tails.

Exercise:

How would your prior belief compare with Jose's or Vicki's? Decide on a mean and standard deviation that reflect your beliefs. Then compute which beta distribution reflects those beliefs. Then compute your posterior distribution, based on your prior and the data. What is the mean of your posterior distribution? That is your Bayesian estimate for p , using a conjugate prior distribution that best fits your beliefs. Would you rather make a bet on that, or on one of the three estimates given here?

Perhaps you fall at one of the extremes – either that you believe $p = 0.5$ with a standard deviation of 0, or that you have no prior belief about p at all, and so want to go entirely by the data. The Bayesian method can accommodate both of those. In the former case, Bayes Theorem will give you the same posterior distribution as the prior distribution – the probability of zero elsewhere than 0.5 wipes out any other possibility, even from the data. For problems of the latter sort, Bayesians might use a “flat prior” which gives equal probability to all possible values of the parameter.

Back to the Overview:

Everyone is a Bayesian in some situations. I don't think I could find anyone to actually place a bet on the estimate of 0.7 in the example. So does that mean that we are all really Bayesians, and there is no point to learning frequentist statistics?

Of course, the basic issue in this example is that we had a quite small set of additional data and most of us feel that we have quite a lot of prior information about whether a coin is likely to be fair. It is an example “made to order” to make frequentist analysis look bad. In other situations where prior information is not agreed upon, then Bayesian analysis looks less attractive.

How do Bayesians deal with these issues?

The main criticism of Bayesian analysis is that it is subjective because the conclusion depends on the prior distribution and different people will have different prior distributions.

One approach to counter the criticism of subjectivity is to use “flat priors” which give equal weight to all possible values of the parameter. This is a very common approach in contemporary Bayesian methods. But there are some complications. One is that, when the parameter space is infinite (e. g. the whole real line) then a flat prior is not actually a distribution because it has an infinite integral rather than integrating to 1. Another complication is that we may be interested in a function of the parameter rather than the parameter itself, and whether you take a flat prior on the parameter or on the function of the parameter makes a difference in the result. Nevertheless, much work has been and is being done in this area.

Some people call this “objective Bayesian analysis.” The general term is “noninformative priors” and some names of these types of prior distributions are default priors, Jeffrey's prior, maximum entropy priors, and reference priors.

Many Bayesian statisticians consider “**subjective Bayesian analysis**” to be the heart of the Bayesian method and work in areas where it is feasible to specify prior distributions reasonably well. Among other things, they have developed techniques for eliciting such priors. (That is a harder task than it might appear at first glance!)

Another approach, called **“robust Bayesian analysis”** works with **“informative” priors but allows for a family of prior distributions** and a range of possible priors in the analysis, and so avoids some of the problems of trying to be quite specific. It is not unusual for a group of researchers to be able to agree on a reasonable range of possible values for the parameters in the prior distribution and so this removes some of the objection to the subjectivity.

“Empirical Bayes” is a technique that allows one to use data from somewhat similar situations to “borrow strength” and produce better estimators in the individual situations.

Computation:

Perhaps it has occurred to you that when we use a prior distribution other than a beta distribution on this problem, we can certainly write an expression for the posterior density as we did here, but actually doing any calculations with this density function, such as finding the mean, variance, or any probability calculations could be very difficult in general if the posterior density wasn't one of our usual distributions.

There are many functions which can't be integrated in closed form and so numerical techniques are needed. In high-dimensional problems, numerical integration is an even more complicated problem than in low-dimensional problems. Markov Chain Monte Carlo (MCMC) techniques were introduced in the late 1980's and use a computationally-intensive simulation method to replace exact integrals. The aim of MCMC is to use simulation to draw a large sample from the full posterior distribution and, using that sample, estimate the needed values from the posterior distribution. Some of those are means, medians, probability intervals, etc. There are several approaches to MCMC.

What does the statistics community say?

When I was in graduate school in the 1970's and 1980's, we discussed the “Bayesian-frequentist” controversy. People took sides and made impassioned arguments. Many graduate schools had no courses involving Bayesian statistics and few courses mentioned it. My advisor was working in empirical Bayes areas and I learned various interesting things, such that insurance companies had been using some of these estimators since the early 1900's without any formal proof that they were good estimators. Why do you suppose it was the insurance companies in particular?

In the 1990's to now, most of the graduate courses don't mention Bayesian statistics. It is becoming more common for graduate schools to have one or two courses in Bayesian methods. From the leading theoretical statisticians, I hear “one should have many techniques in your ‘toolkit’ and use the ones appropriate for the particular problem.” One such person is Bradley Efron, the current President of the American Statistical Association. I found a couple of interesting articles of his on the web and those are listed in the references.

Introduction to Bayes Theorem (Full explanation)

Bayes Theorem: Enables us to “turn around” conditional probabilities.

$$P(\theta | X) = \frac{P(X | \theta) \cdot P(\theta)}{P(X)}$$

Example:

Enzyme immunoassay (EIA) tests are used to screen blood for the presence of antibodies to HIV, the virus that causes AIDS. Antibodies indicate the presence of the virus. The test is quite accurate – about 99% accurate. Specifically, it has 99.85% probability of detecting antibodies when they are present and 99.4% probability of a negative result when antibodies are not present. It is also known that about 1% of the overall population has the antibodies

When a person has a positive result on this test, how likely is it that they actually carry the antibodies?

Solution: Call A the condition of having the antibodies and $+$ and $-$ the conditions of having a positive and negative result, respectively.

$$P(+ | A) = 0.9985$$

$$P(+ | \text{not } A) = 0.0060$$

$$P(A) = 0.01$$

$$P(- | A) = 0.0015$$

$$P(- | \text{not } A) = 0.9940$$

$$P(\text{not } A) = 0.99$$

Write the question as a conditional probability. $P(A | +)$

Use Bayes Theorem to “turn around” the probabilities.

$$P(A | +) = \frac{P(+ | A) \cdot P(A)}{P(+)}$$

$$P(A | +) = \frac{P(+ | A) \cdot P(A)}{P(+ | A) \cdot P(A) + P(+ | \text{not } A) \cdot P(\text{not } A)}$$

$$P(A | +) = \frac{0.9985 \cdot 0.01}{0.9985 \cdot 0.01 + 0.0060 \cdot 0.99} = 0.627$$

Most people find that this is a surprising result. It’s not even close to 99%, which was the (correctly) reported accuracy of the tests.

Additional Resources:

These are generally organized, within categories, by the order in which I think you might want to pursue them, with the easiest first! (Go to the website listed at the top of this page for links to the URLs below.)

Overview of Bayesian statistics:

Stangl, Dalene, *An Introduction to Bayesian Methods for Teachers in Secondary Education*. 1999.

<http://www.isds.duke.edu/~dalene/talks/ncssm/sld001.htm>

Yudkowsky, E. *An Intuitive Explanation of Bayesian Reasoning*.

<http://yudkowsky.net/bayes/bayes.html>

Lee, P. M. *Bayesian Statistics, an Introduction*. 1989. Edward Arnold. Distributed in US by Oxford U Press. (second ed. 1997)

Berger, J. O. "Bayesian Analysis: A Look at Today and Thoughts of Tomorrow." p. 275-290 *Statistics in the 21st Century*, 2002. Chapman and Hall.

Gelman, A., Carlin, J.b., Stern, H. S., and Rubin, D. B. *Bayesian Data Analysis, Second Edition*. 2003. Chapman and Hall.

Use a Bayesian application yourself!

Bayesian Spam Filtering. http://email.about.com/cs/bayesianfilters/a/bayesian_filter.htm

Elementary Statistics with Bayesian methods:

Albert, J, Rossman, A. *Workshop Statistics: Discovery with Data. A Bayesian Approach*. 2001 Key College Publishing.

Berry, Donald. *Statistics, A Bayesian Perspective*, 1996 Wadsworth.

Comparison of Bayesian and Frequentist methods:

Efron, Bradley, Interview. <http://www.mhhe.com/business/opsci/bstat/efron.mhtml>

Efron, Bradley, *Bayesian, Frequentists, and Physicists*.

<http://www-stat.stanford.edu/~brad/papers/physics.pdf>

Barnett, Vic. *Comparative Statistical Inference*. 1999 Wiley.

Teaching Statistics:

Consortium for the Advancement of Undergraduate Statistics Education: <http://www.causeweb.org/>
(many resources)