

DIAGNOSTIC REASONING

Christopher A. Klipstein

Objectives:

After the teaching session on diagnostic reasoning, students should be able to:

- 1) explain in plain English the concepts of prevalence, sensitivity, specificity, positive predictive value, negative predictive value, positive likelihood ratio, negative likelihood ratio, pretest probability and posttest probability
- 2) explain the interrelationship between all of the above concepts (e.g., calculate each value when given several of the other values; describe which values change with prevalence of disease and which can be considered characteristics of the test itself)
- 3) convert from probability to odds and from odds to probability
- 4) determine which test has the best ability to "rule in" or "rule out" a given disease when given sensitivity, specificity and likelihood ratios
- 5) calculate likelihood ratios when given information in a 2 x 2 table format
- 6) calculate likelihood ratios when given sensitivity and specificity
- 7) calculate posttest probability when given pretest probability and likelihood ratios (by using longhand math and also by using Fagan nomogram)
- 8) calculate posttest probability when given pretest probability and several likelihood ratios (linking several LRs together)

Other Suggested Materials:

Chapter 3 (Diagnosis) from Clinical Epidemiology: The Essentials, R. Fletcher and S. Fletcher, Lippincott Williams and Wilkins. (Textbook from 2nd year Clinical Epidemiology course)

<http://gim.unmc.edu/dxtests/Default.htm> Interpreting Diagnostic Tests (tutorial by Thomas Tape at U of Nebraska)

<http://www.cebm.net/downloads.asp> Teaching About Diagnosis (a big PowerPoint by Tom Sensky at Oxford)

<http://www.cebm.utoronto.ca/glossary/lrs.htm#top> Lists some likelihood ratios for common tests

<http://www.jr2.ox.ac.uk/bandolier/booth/diagnos/Liketab.html> Some more likelihood ratios

Introductory Thoughts:

Why do clinicians order tests? What are clinicians hoping to be able to do with the test results?

You are reading about a new test. The journal article reports that the test has a sensitivity of 80%, a specificity of 90%, a positive predictive value of 88% and a negative predictive value of 82%.

The prevalence of disease in the population on which you wish to use this test is 1%. Is this new test a good test to use? How good is the test? How much does this test change your pretest probability of disease? Which number(s) help you answer all these questions?

By the end of the teaching session, you should be able to answer all of these important questions.

The relationship between prevalence, sensitivity, specificity, positive predictive value, negative predictive value, and likelihood ratios

Imagine there is a new test for diagnosing HIV. It is a baseball hat. It is white at baseline. If an HIV negative person puts on the hat, it remains white. If an HIV positive person (determined by the “gold standard” of a positive HIV viral load) puts on the hat, the hat turns red.

a) The hat was first tested at an Infectious Disease clinic where the following results were obtained:

	HIV Viral Load Positive	HIV Viral Load Negative
Hat Red (Positive)	160	20
Hat White (Negative)	40	180

From the results, please determine the prevalence of HIV viremia in this clinic, along with the sensitivity, specificity, positive predictive value and negative predictive value of the hat test in this clinic.

b) The hat test was then used at an inner-city jail where the following results were obtained:

	HIV Viral Load Positive	HIV Viral Load Negative
Hat Red (Positive)	80	30
Hat White (Negative)	20	270

From the results, please determine the prevalence of HIV viremia in this jail, along with the sensitivity, specificity, positive predictive value and negative predictive value of the hat test in this jail.

c) The hat was then used to screen all students in a school district. The following results were obtained:

	HIV Viral Load Positive	HIV Viral Load Negative
Hat Red (Positive)	64	792
Hat White (Negative)	16	7128


From the results, please determine the prevalence of HIV viremia in this school district, along with the sensitivity, specificity, positive predictive value and negative predictive value of the hat test in this school district.

Which of the statistical properties of the hat test remain constant in all the above scenarios (i.e., which are characteristics of the test itself) and which change depending on the prevalence of disease? Which numbers tell you how “good” the test is?

Likelihood Ratios

Let’s return to our initial two questions: why do clinicians order tests and what are they hoping to do with the results?

When we decide to order a diagnostic test, we want to know which test (or tests) will best help us “rule in” or “rule out” a disease in a given patient. Using the language of clinical epidemiology, we take an initial assessment of the likelihood of disease (“pretest probability”), do a test, then use the test results to shift our suspicion of disease one way or the other, thereby determining a final assessment of the likelihood of disease (“posttest probability”). Take a look at the diagram below, which graphically illustrates this process of “revising the probability of disease”.



Which statistical property of a diagnostic test do we use to determine the posttest probability of disease when given a pretest probability and a specific test result? The answer: the likelihood ratio.

Clinicians hope that the posttest probability is different from the initial pretest probability in a clinically meaningful way (e.g., before the test we did not think the patient had the disease and after the test we do think the patient has the disease). Which statistical property of a diagnostic test tells us how much a particular test result changes our pretest probability? The answer: the likelihood ratio.

Clearly we need to master the concept of likelihood ratios since they help us do what we want tests to do (change a pretest probability into a posttest probability) and they tell us how “powerful” a test is (how big of a difference a given test result will make between the pretest probability and the posttest probability, implying how likely the test result is to make a difference that is clinically meaningful).

So what exactly is a likelihood ratio and how do we determine a given test's likelihood ratios?

The likelihood ratio is actually a fairly simple concept but sometimes it takes a bit of explaining for the concept to make sense. I will therefore describe the concept of a likelihood ratio in several different but somewhat repetitive ways.

The likelihood ratio (LR) is the likelihood (or probability) that a given test result would be expected in a patient with the target disorder compared to the likelihood (or probability) that the same result would be expected in a patient without the target disorder. The “formula” for calculating the likelihood ratio is:

$$\text{The likelihood ratio} = \frac{\text{the probability of a **given test result** in an individual **with** a disease}}{\text{the probability of **the same test result** in an individual **without** a disease}}$$

To state the obvious, the likelihood ratio is a “ratio of likelihoods” for a given test. The first likelihood is the probability that a given test result occurs among people with disease (this is the numerator in the formula). The second likelihood is the probability that the same test result occurs among people without disease (this is the denominator in the formula). The ratio of these two likelihoods is the likelihood ratio.

Because tests results can be positive or negative, there are at least two likelihood ratios for each test. The “positive likelihood ratio” (LR+) tells us how much a positive test result increases the probability of disease, while the “negative likelihood ratio” (LR-) tells us how much a negative test result decreases the probability of disease.

The positive and negative likelihood ratios are:

$$\text{LR+} = \frac{\text{probability of a **positive** test in an individual **with** a disease}}{\text{probability of a **positive** test in an individual **without** a disease}}$$

$$\text{LR-} = \frac{\text{probability of a **negative** test in an individual **with** a disease}}{\text{probability of a **negative** test in an individual **without** a disease}}$$

Let's consider an example:

In a study of the ability of a rapid antigen test to diagnose strep pharyngitis, 90% of patients with strep pharyngitis had a positive rapid antigen test, while only 5% of those without strep pharyngitis had a positive test. The LR+ for the ability of this rapid antigen test to diagnose strep pharyngitis is:

$$\text{LR+} = 90\% / 5\% = 18$$

This LR+ of 18 means that a positive rapid antigen test is 18 times more likely to happen in a patient with strep than a positive rapid strep test is to happen in a patient without strep. Another way to say this is that a patient with strep pharyngitis is 18 times more likely to get a positive rapid antigen test than a patient without strep pharyngitis is to get a positive rapid antigen test. In still other words, for any given patient, a positive rapid strep test is 18 times more likely to be a true positive test (the patient does indeed have strep) than a false positive test (the test is positive but the patient actually does not have strep). This last way of describing the meaning of a likelihood ratio is usually the easiest to comprehend.

Please determine the positive likelihood ratio (LR+) and the negative likelihood ratio (LR-) of the hat test for each of the three scenarios on the previous page.

As the hat test example demonstrates, while the positive predictive value and the negative predictive value of a test will change as the prevalence of disease in the population being studied changes, the likelihood ratios (along with sensitivity and specificity) do not change. Likelihood ratios (along with sensitivity and specificity) can be considered to be properties of the test itself that do not change with the prevalence of disease.

So how do we actually use these likelihood ratios?

As noted above, one of the most important clinical uses of likelihood ratios is the ability of likelihood ratios to take a pretest probability and change it into a posttest probability. This can be done in two ways: using longhand math and using the Fagan nomogram.

Longhand Math

Likelihood ratios, as a ratio of two probabilities, are actually odds. To use likelihood ratios in a mathematical formula, all the other numbers need to be odds as well. The formula for converting pretest odds to posttest odds is:

$$\text{pretest odds} \times \text{likelihood ratio} = \text{post test odds}$$

Clinicians are usually more comfortable thinking in terms of probabilities rather than odds. In order to use the longhand method of turning pretest probabilities into posttest probabilities, one needs to know how to convert probabilities into odds and vice versa. Use the following questions to help understand the concepts for making these conversions.

The forecast calls for an 80% chance of rain tomorrow. How many times more likely is it to rain tomorrow than to be clear?

A friend has a bag of marbles. There are 100 marbles in the bag. Some of them are red and the rest are white. Your friend tells you to reach into the bag and pull out one marble. Before you do this, your friend tells you that you are three times as likely to pull out a red marble than a white marble. How many red marbles are in the bag? What percent of the marbles are red? What are your chances of pulling out a red marble on your first attempt?

For those of you who prefer formulas, the formulas for converting probabilities into odds and vice versa are:

For a probability of X %, the odds are $\frac{X}{100 - X}$ For an odds of A: B, the probability is $\frac{A}{A + B}$

So, for a probability of 80%, the odds are $\frac{80}{(100 - 80)} = \frac{80}{20} = \frac{4}{1}$ (also written as 4 : 1 odds)

And for an odds of 3 : 1 (also written as 3 / 1 odds), the probability is $\frac{3}{3 + 1} = \frac{3}{4} = 0.75$ (or 75%)

To use the longhand method of using likelihood ratios to convert pretest probabilities into posttest probabilities, first convert the pretest probability into pretest odds, then multiply these pretest odds by the appropriate likelihood ratio to get posttest odds, and then convert these posttest odds into posttest probabilities.


Try a couple of examples in your teaching session.

- 1) You are seeing a woman with a swollen calf. Based on her history and exam she has a 5% pre-ultrasound probability of a DVT ([see JAMA, April 8, 1998, pgs 1094-1099](#)). The LR+ of an ultrasound for DVT is around 20 and the LR- is 0.05. The ultrasound tech calls you and says the ultrasound is positive for a DVT. Does this woman have a DVT? What is her post-ultrasound probability of DVT when the ultrasound is positive? What if the ultrasound is negative? What is the posttest probability if her pretest probability was 80% and the ultrasound is negative?

- 2) You are seeing an elderly gentleman with persistent abdominal and back pain. He has lost 10 pounds over the past 2 months. Assume your pretest (pre-CT scan) probability of pancreatic cancer is around 20%. The LR+ of a CT scan for pancreatic cancer is around 8 and the LR- is around 0.2. What is the post-CT scan probability that this man has pancreatic cancer if the CT scan shows a pancreatic mass? What if the CT scan does not reveal a pancreatic mass?

Fagan nomogram

The Fagan nomogram allows you to turn pretest probabilities into posttest probabilities without needing to convert into odds. To use the nomogram, simply line up the pretest probability on the left with the appropriate likelihood ratio in the center and read off the posttest probability on the right. Try using the nomogram for the two example problems above.



You can use the nomogram to play around with the different variables that go into the nomogram. This can be done forwards or backwards.

For example:

If you have a patient with a pretest probability of a pulmonary embolism of 20%, how good of a test (i.e., how big of a likelihood ratio) do you need to get a posttest probability above 50%?

If you need a posttest probability above 80% before you would use a certain therapy, and the test you want to use has a LR+ of 5, what is the minimal pretest probability you need to allow a positive result on the test get you above the treatment threshold?

Again, likelihood ratios indicate how powerful a test is (by how much a given diagnostic test result will raise or lower the pretest probability of the target disorder). A likelihood ratio of 1 means that the posttest probability will be exactly the same as the pretest probability. Likelihood ratios greater than 1 increase the probability that the target disorder is present, and the higher the LR the greater this increase. The best test to “rule in” a disease is the one with the largest positive likelihood ratio (LR+). Conversely, likelihood ratios less than 1 decrease the probability of the target disorder, and the smaller the LR, the greater the decrease in probability. The best test to “rule out” a disease is the one with the smallest negative likelihood ratio (LR-).

Use the nomogram above to see for yourself that the farther the likelihood ratio is from 1, the larger the change will be from the pretest to the posttest probability. One can group likelihood ratios into magnitudes:

LR+	> 10	makes large and often conclusive increases in the likelihood of disease
LR+	5-10	makes moderate increases
LR+	2-5	makes small increases
LR+	1-2	makes insignificant increases
LR-	0.5-1.0	makes insignificant decreases
LR-	0.2-0.5	makes small decreases
LR-	0.1-0.2	makes moderate decreases
LR-	< 0.1	makes large and often conclusive decreases in the likelihood of disease

Using likelihood ratios for tests with multilevel results

Some tests have dichotomous results, meaning the results are categorized only as positive or negative. Other tests can have several magnitudes of results. One can determine likelihood ratios for these different magnitudes of results. This allows us to utilize the full power of extreme test results. For example:

Military recruits are routinely screened for HIV infection with an HIV ELISA test. The prevalence of HIV in this population is roughly one in a thousand.

The test characteristics of an HIV ELISA test are described in the table below. (Please note that these numbers are from an older version of the HIV ELISA. The current ELISA is much improved. Since the numbers below demonstrate a teaching point, they were not updated in this handout.) Using the Western Blot as a gold standard, one study looked at 200 patients with a high risk of disease.

HIV ELISA	Western Blot Positive	Western Blot Negative
Positive	85	12
Negative	15	88
Total # Patients	100	100

A recruit comes to your office. He feels entirely well. His risk factor profile puts him at the average risk for recruits (1/1000 or 0.1%). He tells you he had his routine test for HIV and “they told me the test was positive”. What is the LR + for a “positive” ELISA test and what is this recruit’s posttest probability of HIV?

The recruit then remembers, “Oh yeah, they said it was strongly positive”. Using the information in the table below, what is the LR+ for a “strongly positive” ELISA test and what is this recruit’s posttest probability of HIV now?

HIV ELISA	Western Blot	Western Blot
	Positive	Negative
Strongly Positive	55	1
Moderately Positive	25	6
Weakly Positive	5	5
Weakly Negative	5	2
Moderately Negative	9	20
Strongly Negative	1	66
Total # Patients	100	100

Note that one can “collapse” the results of the second table into the results of the first table by combining all the “positive” results (regardless of strength) together and all the “negative” results together.

What if this recruit actually told you that he used IV drugs and had shared needles? If your pretest probability of HIV is now say 10%, what is the posttest probability if he gets a positive ELISA? What if it is a strongly positive ELISA? Do the same two calculations again assuming your pretest probability of HIV is 50%.

This example highlights the extra information one can get by using tests that have multilevel (instead of just dichotomous positive/negative) results. Likelihood ratios can be determined for each level of test result. You can then use longhand math or the nomogram to determine the posttest probability of disease for a specific test level result.

This example also reminds us that when you have a very low pretest probability of disease, even a positive result on a very good test leaves you with a posttest probability that is still quite low.

Another example:

You are seeing a 72-year-old woman in clinic. She was feeling well at her last visit and she left clinic with a packet of screening fecal occult blood tests. Today she reports a several month history of fatigue and the nurse tells you that 2 of the 3 FOBT samples she submitted last week were guaiac positive. What is your pretest (pre-lab) probability that this woman has iron deficiency anemia?

Her hematocrit is 32% with an MCV of 78fL and a ferritin of 16 ng/L. The “normal” range for ferritin at UNC Hospitals is 3-151 ng/L. According to her labs, is this patient iron deficient?

The results of an overview to determine the test characteristics of ferritin to detect iron deficiency anemia are summarized in the table below. The overview involved several thousand patients and used histiologic examination of iron on bone marrow aspiration as the gold standard.

Serum ferritin	Iron Deficient Patients	Not Iron Deficient Patients	Likelihood Ratio
> 100 ng/L	48	1320	0.08
45 – 100 ng/L	76	398	0.54
35 – 45 ng/L	36	43	1.83
25 – 35 ng/L	58	50	2.54
15 – 25 ng/L	117	29	8.83
< 15 ng/L	474	20	51.58

Use your pretest probability that this woman has iron deficiency anemia and the likelihood ratio from the table above to determine your posttest (post-lab) probability that this woman has iron deficiency anemia.

Why do you think the UNC lab uses a single cutoff point for an abnormal ferritin? (I do not know why they do this.)

Using likelihood ratios with tests done in series

When several tests are done in series, likelihood ratios can be used to determine a final posttest probability at the end of the series of tests. The posttest odds after the first test becomes the pretest odds for the second test. (Technically speaking, this concept only works if all the test results are independent of one another— do not worry about this for now.)

The formula simplifies to:

$$\text{pretest odds} \times \text{LR}_{\text{test1}} \times \text{LR}_{\text{test2}} = \text{posttest odds} \quad (\text{you can add as many LR's as there are tests})$$

Some examples:

You are seeing a completely healthy 37-year-old woman whose company requires her to get an “executive physical” every year. This includes a stress thallium test. The exercise (ETT) portion of the test is positive (1.2mm ST depression) and the thallium portion is positive as well (reversible defect). An average 37-year-old asymptomatic woman has a pretest probability of flow limiting coronary artery disease of around 3/1000 (0.3%). The LR+ of an ETT is 6 and the LR+ of the thallium portion is 10. What is her probability of disease after the ETT portion of the test? (Use the Fagan nomogram.) Now, using the post-ETT probability as your pre-thallium probability, what is the probability of disease after the thallium portion of the test?

The next patient in clinic is a 65-year-old man with typical angina. You order a stress thallium test. The exercise (ETT) portion of the test is negative and the thallium portion is negative as well. An average 65-year-old man with typical angina has a pretest probability of flow limiting coronary artery disease of around 94%. The LR– of an ETT is 0.4 and the LR– of the thallium portion is 0.1. What is his probability of disease after the ETT portion of the test? After the thallium portion?

After you have completed all the tests, which patient has the higher probability of disease? Do these test results surprise you?

What would the 65-year-old man’s posttest probability be if the ETT had been positive but the thallium test had been negative?

So if likelihood ratios are so cool, how can I find likelihood ratios for common tests?

First of all, you can use the links provided on the first page of this handout. Second, you can calculate likelihood ratios yourself when you read about a test. Unfortunately, many articles/textbooks provide the sensitivity and specificity of tests but do not provide the likelihood ratios. Fortunately, there is an easy way to calculate likelihood ratios when given sensitivity and specificity.

Let’s go back to the rapid antigen test. Recall that we said that in a study of the ability of this rapid antigen test to diagnose strep pharyngitis, 90% of patients with strep pharyngitis had a positive rapid antigen test, while only 5% of those without strep pharyngitis had a positive test. The LR+ for the ability of this rapid antigen test to diagnose strep pharyngitis was:

$$\text{LR+} = 90\% / 5\% = 18$$

Note that the numerator for the LR+, the probability of an individual with the disease having a positive test, is the same thing (the exact same numbers and the exact same concept) as the sensitivity of the test (90% in this example).

The denominator of the LR+ is the probability of an individual without a disease having a positive test (5% in this example). If one considers *all* the people without strep, and one knows that 5% of these people have a positive test, then the rest of these people (95%) must have a negative test (100% – 5% = 95%). Note that the concept

described in the last half of the previous sentence (the probability of an individual without the disease having a negative test) is the specificity of the test. So the denominator of the LR+ is simply (100% minus the specificity). The same logic should help you understand the “formula” for determining the LR– when given sensitivity and specificity. Note that for the equations below to work one needs to change a sensitivity of 90% and specificity of 95% into the decimals 0.9 and 0.95 respectively.

$$\text{LR+} = \frac{\text{sensitivity}}{1 - \text{specificity}} \quad \text{and} \quad \text{LR-} = \frac{1 - \text{sensitivity}}{\text{specificity}}$$

The likelihood ratio is a way to incorporate the sensitivity and specificity of a test into a single measure that accounts for the balance between a test’s sensitivity and specificity. When you read about a test, and the article/textbook only provides the sensitivity and specificity of the test, you can use the above formulae to calculate quickly the more useful likelihood ratios.

Summary:

After the teaching session you should have a better understanding of the interrelationship between prevalence, sensitivity, specificity, positive predictive values, negative predictive values, positive likelihood ratios, negative likelihood ratios, pretest probability and posttest probability. You should also appreciate some of the clinical uses of likelihood ratios. Some of the advantages of using likelihood ratios include:

- 1) Likelihood ratios (along with sensitivity and specificity) can be considered to be properties of the test itself and do not change with the prevalence of disease. In contrast, the positive and negative values of a test do change as the prevalence of disease in the population being studied changes.
- 2) Likelihood ratios take a pretest probability and change it into a posttest probability (using longhand math or the Fagan nomogram).
- 3) Likelihood ratios can be determined for multiple levels of a given test.
- 4) Likelihood ratios can be used to determine the posttest probability of several tests use in series.
- 5) Likelihood ratios can be used backwards and forwards (using the nomogram) to see how changes in one variable in the nomogram will affect the other variables.

In the session on “Therapeutic Reasoning”, we will learn how to answer the question: What do you do with the post-test probability once you’ve calculated it? We will discuss various thresholds and how they help us understand what it means to “rule in” or “rule out” a disease.