2.7 Sensitivity Analysis

Have taken as 0.4 Laura's belief in the probability of Retro's being a goer,

At what probability would Laura choose Retro with No Trial,

What would the value of a completely accurate Trial be then?

Parameterise the probability of Goer as p, so the probability of Fizzer is $1-p$.

If Laura's probability that Retro is a goer is p, then that must be her best guess as to the probability of the event that the Trial says Retro is a goer.

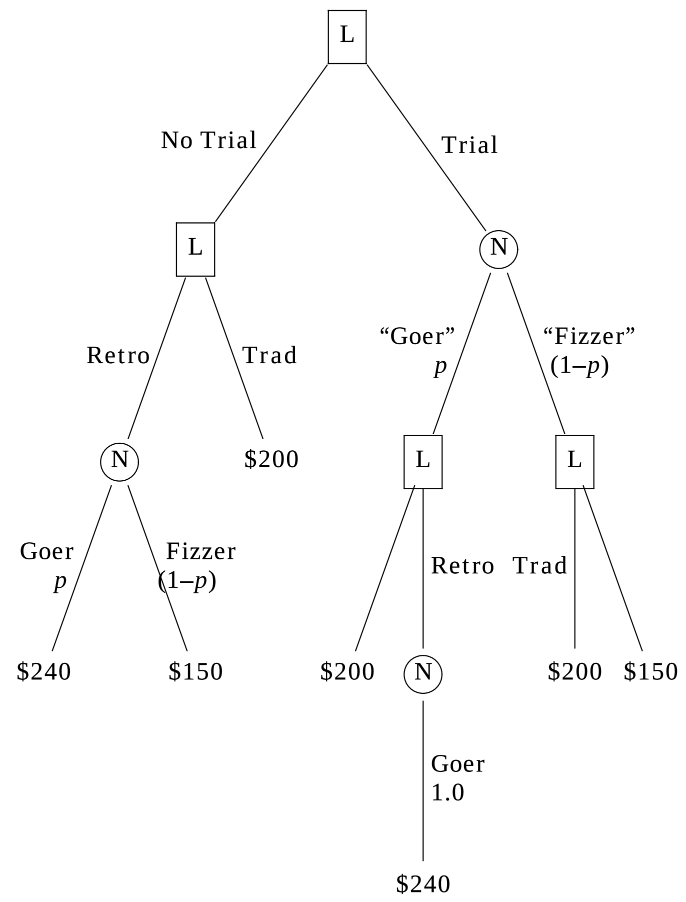

The Shoe Decision with Perfect Information

The expected value of choosing Retro in the absence of a Trial:

$$\$150 + 90p,$$

compared with the unchanged value of $200 of choosing Trad.

On the Trial side of the tree, to be consistent the probability of the test indicating that Retro will be a “Goer” must be p, and a “Fizzer” $1-p$.

The expected value of choosing Trial =

$$200 + 40p.$$

Plotting these expected values as a function of the probability of Goer:

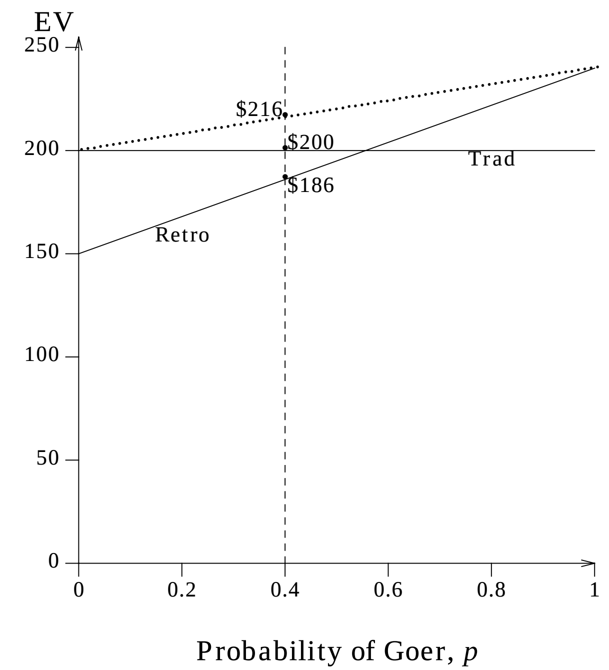

Expected Value against Probability of Goer

2.7.1 The Value of Perfect Information

The dotted line: the expected value of the Trial.

At $p = 0.4$,

- the expected value of Retro is $186,
- the value of Trad is $200,
- and the value of Trial is $216,

The value of perfect information = the improvement in expected value with the Trial,

= the difference between the dotted line and the next highest value, whether Trad or Retro.

When Laura is certain about the outcome with Retro, the value of reducing uncertainty is zero

She is certain twice: when she

- knows that Retro is a Goer ($p = 1.0$), or
- is certain that Retro is a Fizzer ($p = 0.0$).

The cross-over probability $\hat{p}$ at which choosing Retro has a higher expected value than choosing Trad is 0.556.

This probability corresponds to the highest value of perfect information.

2.7.2 Which variables are most crucial?

Have considered the decision's sensitivity to a single variable, the probability that Retro is a Goer.

But some uncertainty about the payoffs of Trad and Retro under the two possibilities.

Which is the most critical variable on which to perform a sensitivity analysis?

Ronald Howard (1988, in package) has suggested that, holding all other variables at their most likely values, one by one each variable be taken from its lowest likely value to its highest, and the effect of this on the optimand (the variable being maximised or minimised) be plotted.

He suggests a *tornado plot*, with the variable with the greatest effect on top and that with the least on the bottom.

Those variables which can push the maximand lowest are the ones that should be subject to a sensitivity analysis.

See Clemen (1996) for further discussion of this topic.

2.8 The Value of Imperfect Information

But what if the Trial is *not* 100%-reliable?

We need to calculate two things:

- Laura's probability that the unreliable test will say "Goer",
- and the *conditional probability* of Retro being a goer if that's what the test indicates.

(With a 100%-reliable test, the former probability is 0.4 and the latter is 1.0.)

The following tree models Laura's decision:

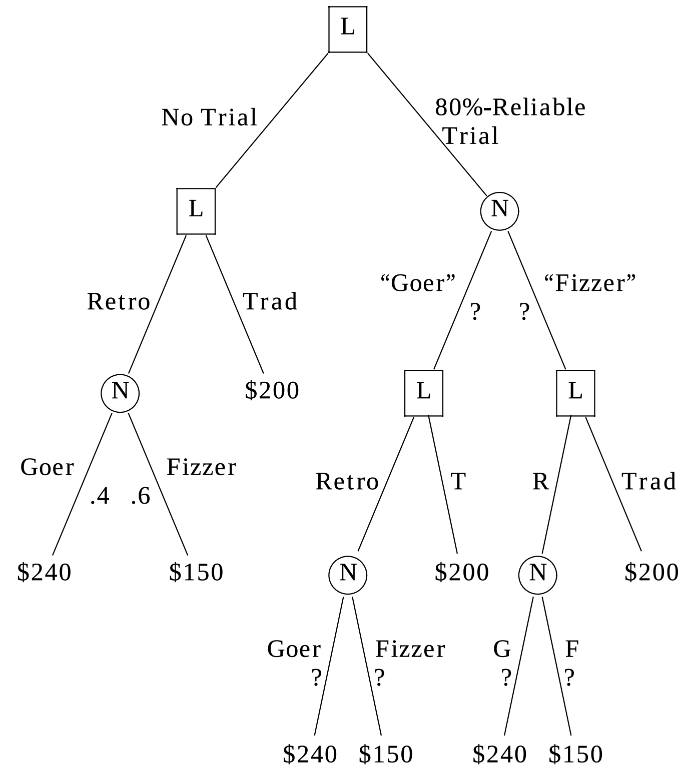

What probabilities do the question marks represent?

To answer this question we need to flip the tree to calculate the conditional probabilities.

2.9 Conditional Probability

Before we continue with Laura's decision, we shall examine a diagnostic problem that, unfortunately, is all too common these days.

If a diagnostic test is not 100% reliable, then what does a positive test (or a negative test) mean in terms of being ill?

Assume that cancer testing is 98% accurate:

- If you have cancer, then the test will show "Positive" 98% of the time, and
- if you do not have cancer, then the test will show "Negative" 98% of the time.
- You have read that 0.5%, or 1 in 200, of the population actually have cancer.

Now you test positive. What is the probability that you have cancer?

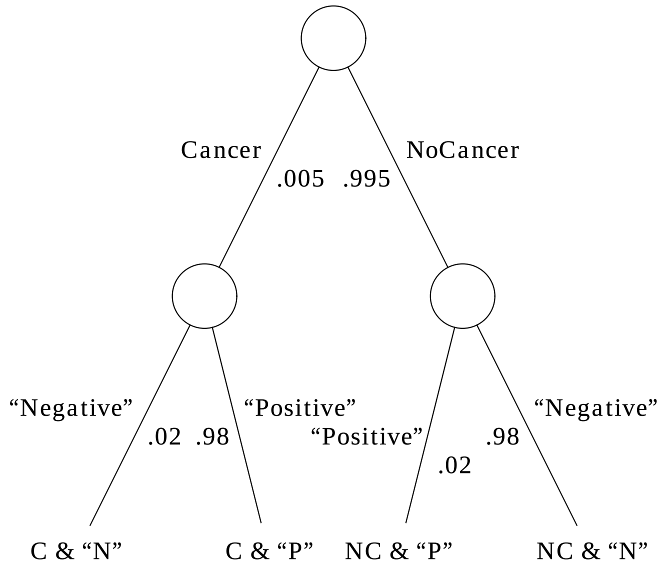

Cancer Testing

C & "N": false negative
NC & "P": false positive

We want the *conditional probability* of Cancer, given that the test says “Positive”:

$$\text{prob}(\text{Cancer} \mid \text{“Positive”})$$

Assume that 10,000 tests performed.

Using the probabilities in the tree, can calculate the expected numbers at the ends of the tree:

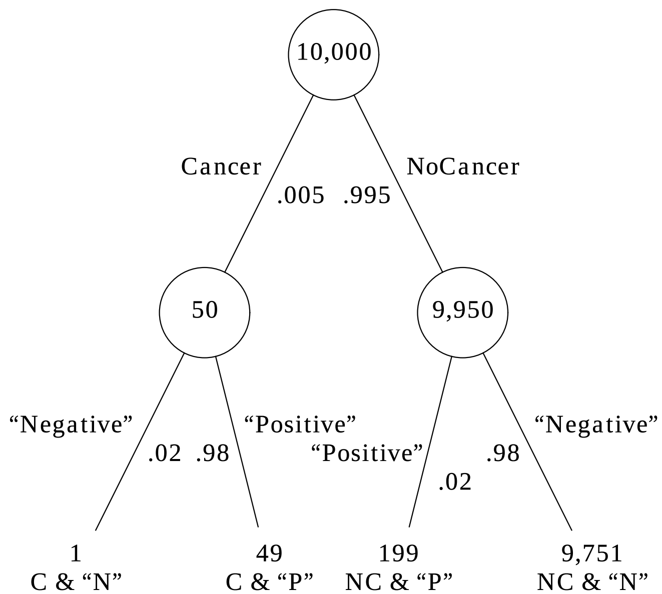

- Of the 10,000 tests, on average 50 people will have cancer.
- Of these, 98% will test positive and so, on average,
- there will be 49 positive tests.
- Of the (on average) 9,950 cancer-free people, 2%, or 199, test positive.
- Thus there are a total of 248 positive tests, of which 199 are false positives.
- So the probability of a positive test indicating cancer is only 49/248, or only 19.8%.

  or $\text{prob}(\text{Cancer} \mid \text{“Positive”}) = \dfrac{49}{49 + 199} =$ 19.8%
- Surprisingly, only 20% of those testing positive will actually have cancer.

But trees like the one above are not very useful:

- Useless to screen *after* we know that someone has cancer.
- Need to know what the likelihood is
  - — that a positive test means that the person has cancer, or
  - — that a negative test means that the person does not.
- That is, we want to know what the probability is that the patient has cancer, given that the test is positive (or negative).
- Or, in Laura's case, we want to know what the probability of Retro being successful is, given that the (imperfect) market trial suggests that it will be successful.

In other words, we want information in the form:

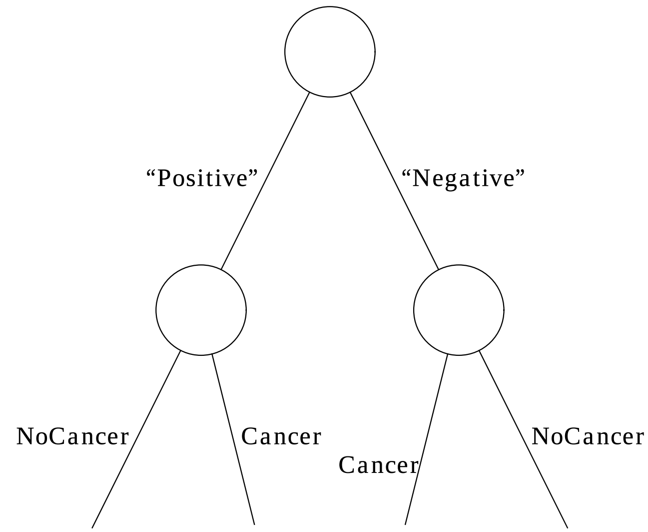

What does it mean to follow along the leftmost branch of the tree on p. 2-2? — that the person has cancer and at the same time the test is negative.

The probability of this combined event — the joint probability — is $0.005 \times 0.02 = 0.0001$

The joint probability of the next most leftmost path is $0.005 \times 0.98 = 0.0049$, the probability of the combined event of the person having cancer and testing positive.

But this combined event is the same as the combined event "tests positive and has cancer", the order doesn't matter.

Hence the joint probability must be the same, 0.0049.

Similarly, each of the paths on the last tree is equivalent to a path on the earlier trees.

The probabilities for each path are shown in the following two trees:

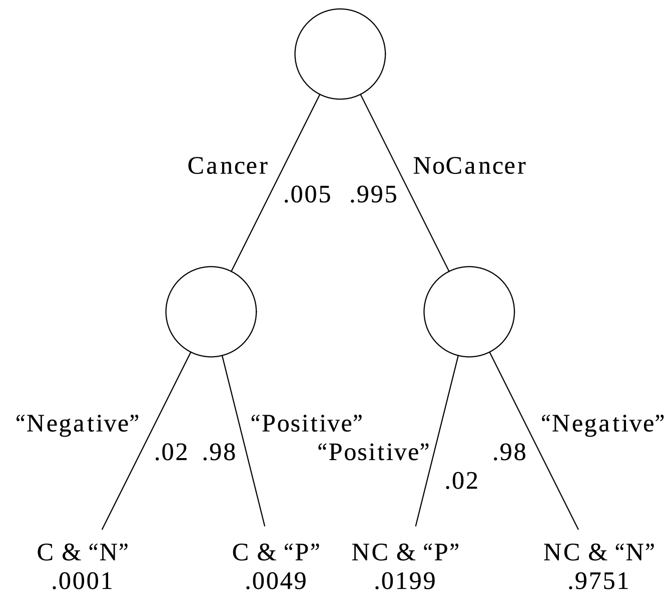

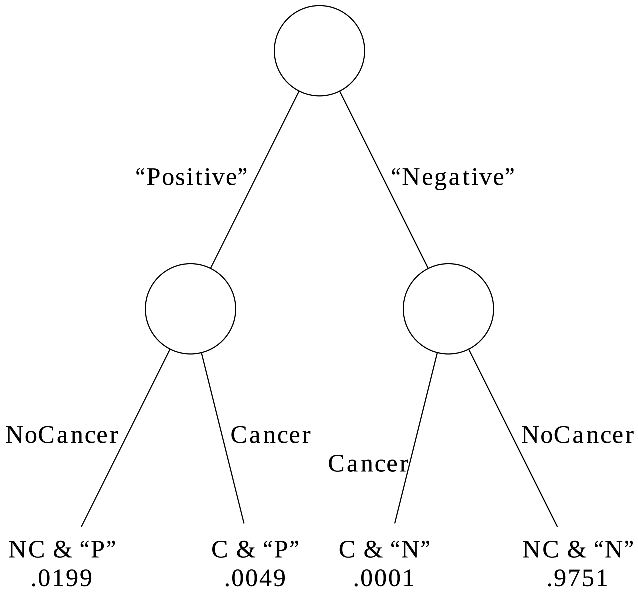

Now comes the crucial step.

If a test is positive, it must be that

- either the person has cancer (with probability of 0.0049 for the joint event "has cancer, tests positive")
- or that he doesn't (with joint probability of 0.0199 for the joint event "tests positive, no cancer");

the total probability of a positive test is the sum of these, 0.0248.

Similarly, the probability of a negative test is 0.9751 + 0.0001 = 0.9752.

Fill in the probabilities of the final two chance nodes, by simple arithmetic.

If the probability of a positive test is 0.0248 and the probability of a the path "Positive" and Cancer is 0.0049, then the *conditional probability* that the person has cancer *once his test is positive* is 0.0049/0.0248 = 0.198.

Similarly, we can calculate the rest of the tree:

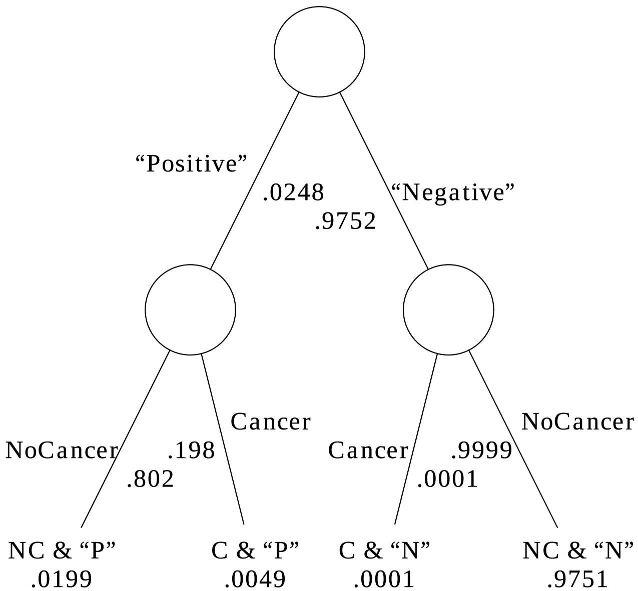

The tree is in the form that is of use to us;

the process is called *decision-tree flipping*, the intuitive version of Bayes formula.

Tree flipping is simple: there's no formula to forget, and it works for more complex situations.

Or assume a population of 10,000 people, and calculate the numbers. The conditional probailities are easily calculated.

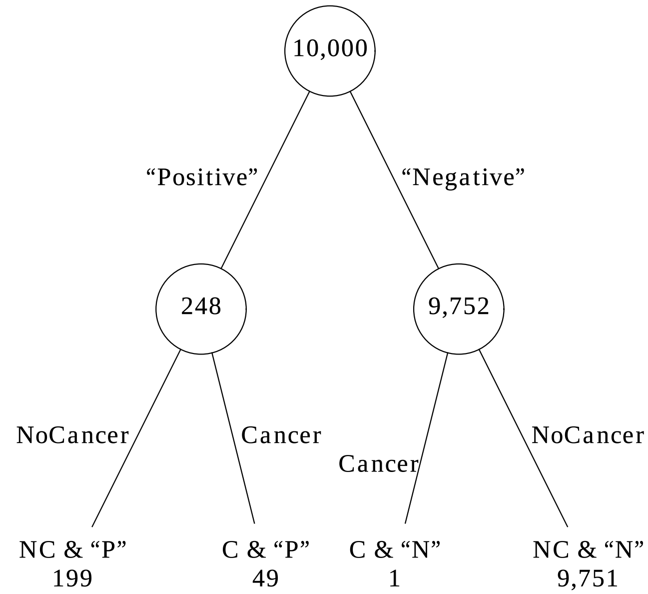

Joint probabilities:

Prob ("P" & NC)	=	$^{199}/_{10000}$	=	0.0199
Prob ("P" & C)	=	$^{49}/_{10000}$	=	0.0049
Prob ("N" & C)	=	$^{1}/_{10000}$	=	0.0001
Prob ("N" & NC)	=	$^{9751}/_{10000}$	=	0.9751
				1.0000

Conditional probabilities:

Prob (NC | "P")	=	$^{199}/_{248}$	=	0.8024
Prob (C | "P")	=	$^{49}/_{248}$	=	0.1976
				1.0000
Prob (NC | "N")	=	$^{9751}/_{9752}$	=	0.9999
Prob (C | "N")	=	$^{1}/_{9752}$	=	0.0001
				1.0000

2.10 Exercise 2: Laura and the Shoe Decision (continued)

Laura decided to employ the Acme Marketing Company. Unfortunately, they are only 80% reliable:

- given that Retro is a Fizzer, Acme will say otherwise 20% of the time, and
- given that Retro is a Goer, Acme will say otherwise 20% of the time.

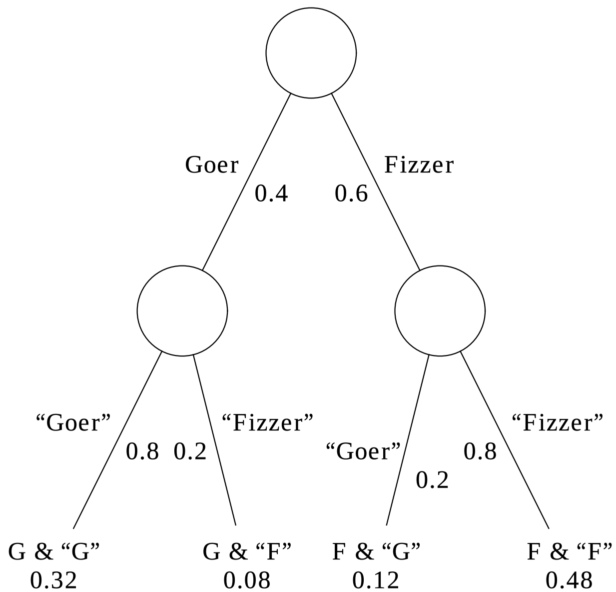

Market Testing

"Flip" the above tree, to determine the chance of Retro being a Fizzer, given that the unreliable test indicates that it will be, etc.

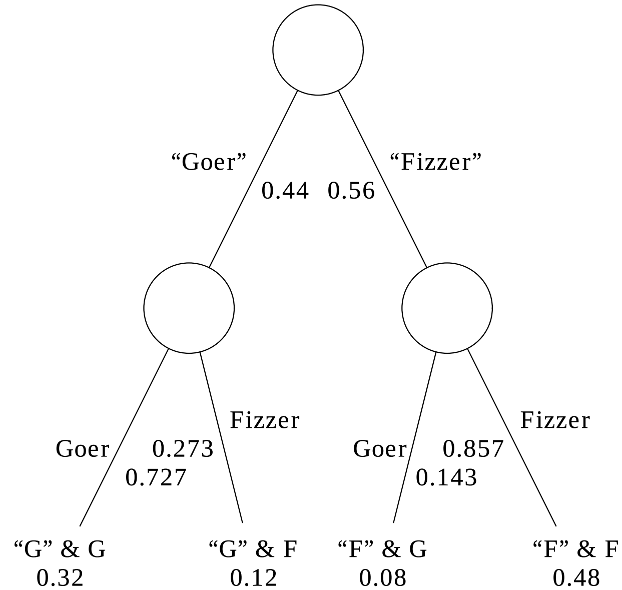

Market Testing

- the conditional probability of Retro being a Goer given that Acme says it's a "Fizzer" is 0.08/(0.08+0.48) or 1/7 or 0.143;
- the conditional probability of Retro being a Goer given that Acme says it's a "Goer" is 0.32/(0.32+0.12) or 8/11 or 0.727.
- based upon Laura's prior belief that Retro is a Goer with a probability of 40%, she expects that with probability 0.32 + 0.12 = 0.44 Acme will say "Goer".

We can now replace the question marks in the decision tree In Section 2.7 which allows us to solve the decision problem, with expected values.

From the sensitivity graph in Section 2.6 the crossover probability is greater than 0.44:

- if Acme says "Goer", which Laura expects will happen with probability of 0.44, then she will choose Retro.
- Her expected payoff is $(150 + 90p) \times 1000 =$ \$215,430, with p = the conditional probability that Retro is a Goer, given that Acme said "Goer" = 8/11 or 0.727.
- If Acme says "Fizzer", which Laura expects will happen with probability of 0.56, then she will choose Trad, with a payoff of \$200,000.
- Her expected payoff with Acme's imperfect information is thus $0.56 \times$ \$200,000 + $0.44 \times$ \$215,430 = \$206,789.
- Her expected payoff without this information is \$200,000, since she chooses Trad.
- Thus the expected value to Laura of 80%-reliable information is \$6,789.

After tree-flipping:

- Laura’s *conditional probability* that Retro is a Goer, given that Acme has states that it will be a Goer, is 8/11, or 0.727.
- Laura’s probability that Acme will state that Retro is a Goer is 0.44.

Laura’s full decision tree:

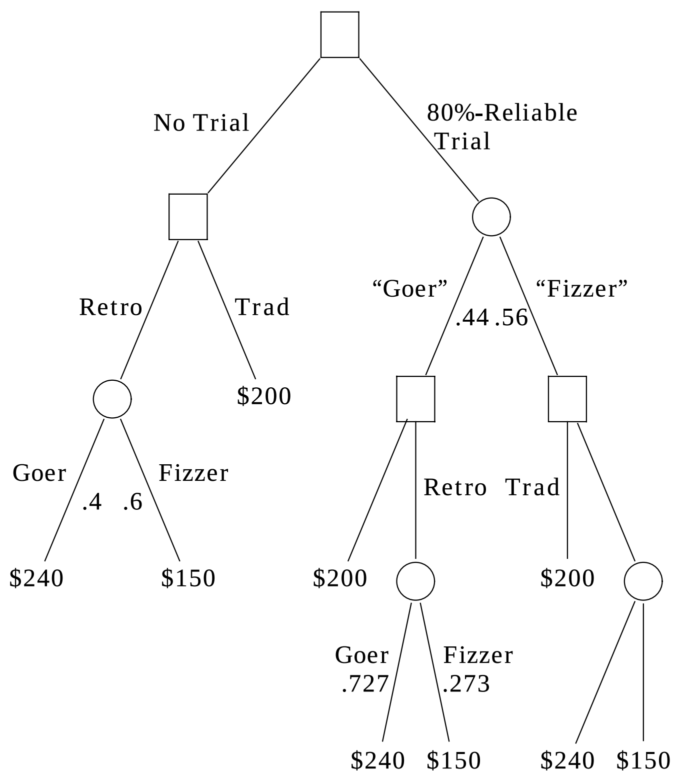

So the following tree (from p. 2-8) models Laura’s decision:

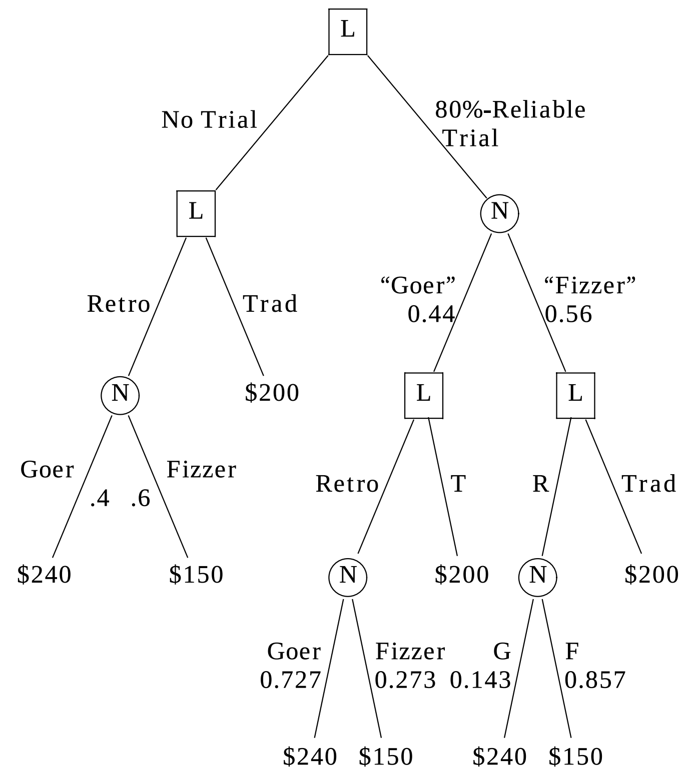

EV with the 80%-test = $206,790
EV without the test = $200,000

∴ EV of the 80%-reliable information = $6,790