eman ta zabal zazu

Universidad del País Vasco   Euskal Herriko Unibertsitatea

Activity 3

1 Synopsis.

What we are set about to do. This activity will illustrate the use of statistical testing to decide whether data supports (or not) a given hypothesis. In this activity we will examine the issue of patents on pharmaceutical compounds, a matter which has stirred up considerable debate in recent years.

What you need to know. In order to benefit from this activity you need a working knowledge of R, such as you may have gained from previous seminars and activities, and a good understanding on how to perform a chi-square, Fisher's exact test, and related tests.

2 Context.

2.1 The intellectual property debate.

In the last ten years the debate over intellectual property has been increasingly heated, as big discographic companies press their interests while consumers claim what they see as their right to download or exchange music on the Internet without restraints. This has not been the only battle front: the purchase of Motorola by Google in the summer of 2011 has been seen as preparation for an all-out patent war over smart phones technology, with Apple as the other main contender. Even more recently, Apple obtained an injunction preventing Samsung from selling their tablets in some countries.

In Spain, the ill-fated "Sinde law" and the opposition it has aroused shows that the controversy is well alive and we are not anywhere near a consensus on the issue.

Intellectual property protection is a matter of great interest for an economist. Is it necessary for progress, or rather does it hinder progress? [1] is a passionate, well researched attempt to answer this question. One of the aspects they address is patent regulation in the pharmaceutical industry. The data next come from that book.

2.2 Pharmaceutical innovation.

In [1] the authors claim that intellectual property protection in the pharma industry does more harm than good. They present in Table 1 evidence from Italy, were legislation was enacted in 1980 affording considerable more protection to pharma intellectual property.

Period	Italy	Rest of world	% Italy s/total
1961 – 1980	119	1163	9.28
1980 – 1983	8	100	7.41

Table 1: Number of innovations in Italy and the rest of the World in the pharma industry. Source: [1], p. 251.

It is their claim that when the pharma industry became more protected in Italy, their rate of innovation actually *declined*. Italy made 9.28% of all pharma innovations prior to 1980, but only 7.41% afterwards.

3 Questions.

1. Propose a model for the data in Table 1.

   Answer: We may think of the periods 1961–1980 and 1980–1983 as two different "populations", and then our task would be to check whether those two populations are homogeneous with respect to the distribution of discoveries between Italy and the rest of the world.

   We might also think that we sample a whole "population" of discoveries 1961–1983, although the first approach seems a little bit more natural. The results would be the same either way.

   Remarks: What is important is that we regard the Italy/Rest of the world shares of discoveries as random, and we try to answer a question about the *parameters* of the underlying random model ("Was the *probability* of a pharmaceutical innovation happening in Italy constant over the period 1961–1983 (H_0), or did it decline in 1980–1983? (H_a)") In other words, we regard what happened as one of the many (random) outcomes that might have been seen, and ask ourselves whether what we have actually seen is consistent with H_0.

   Statistical testing or even the notions of homogeneity or independence make no sense dissociated from the statistical model. From a purely descriptive point of view, we can only state that the share of Italian innovations declined fro 9.28% to 7.41%, period. It is only our belief that things might have happened in different (random) ways which makes meaningful to ask questions *not about the actual numbers*, but *about the parameters of the underlying model.*

2. Assuming the data has been generated by the model of your choice, what are the best (maximum likelihood) fits for each cell?

   Answer: Under both the product multinomial (or homogeneity) and multinomial (independence) models, the MLE are the same, and easily computed. If we use R's function `chisq.test` the result gives the answer in the component `expected`:

```
> farma <- matrix(c(119,8,1163,100),2,2)
> dimnames(farma) <- list(c("1961-1980","1980-1983"),
                          c("Italy","RestWorld"))
> farma <- as.table(farma)
> result1 <- chisq.test(farma)
> result1$expected
```

```
                Italy  RestWorld
1961-1980 117.132374 1164.86763
1980-1983   9.867626   98.13237
```

(do `help(chisq.test)` to see the documentation). We can use the (more general) function `loglin` as follows:

```
> result2 <- loglin(farma,margin=list(1,2),fit=TRUE)
2 iterations: deviation 0
> result2
$lrt
[1] 0.446106

$pearson
[1] 0.4217988

$df
[1] 1

$margin
$margin[[1]]
[1] 1

$margin[[2]]
[1] 2


$fit
                 Italy   RestWorld
1961-1980  117.132374 1164.867626
1980-1983    9.867626   98.132374
```

The values in `fit` are the expected values, and match those computed bu `chisq.test`.

3. It is extremely unlikely that those fitted values entirely agree with the observe values. The further the observed values are from your fits, the strongest the evidence against your model. Propose an statistic summarizing the deviation of observations from fitted values.

   Answer: If we are not concerned with a particular alternative, a test statistic such as $\sum_i (O_i - E_i)^2$ (where O_i is the number of observations in cell i and E_i the fitted number for the same cell) would be adequate. In particular,

$$Z = \sum_i \frac{(O_i - E_i)^2}{E_i}$$

   could be used (with the advantage that we know its asymptotic distribution).

4. Is there evidence that the rate of innovation in Italy in the post-1980 period *was different*?

   An adequate test for this hypothesis would be the χ^2 goodness-of-fit test; the value of the test statistic Z can be obtained from the output of `loglin` above, or from:

```
> chisq.test(farma)
        Pearson's Chi-squared test with Yates' continuity
        correction

data:  farma
X-squared = 0.2262, df = 1, p-value = 0.6344
```

Note that the result does not quite match the output of `loglin` because `chisq.test` uses a continuity correction as a default. If we omit that correction,

```
> chisq.test(farma,correct=FALSE)
        Pearson's Chi-squared test

data:  farma
X-squared = 0.4218, df = 1, p-value = 0.516
```

we obtain the same results (compare 0.4218 to 0.4217988 above, similar except for rounding). The `chisq.test` function gives the p-value right away; we see that it is by no means small, so we would not reject H_0. At any reasonable significance level, there seems to be no difference between the pre- and post 1980 years.

If we want to reproduce the results of `chisq.test` using the output from `loglin`, we can type:

```
> 1 - pchisq(result2$pearson,df=result2$df)
[1] 0.5160408
```

5. Is there evidence that it actually *declined*?

Answer: The chi-square test is not useful with an alternative hypothesis, as it is "alternative blind"; it merely tests departure from the null hypothesis in any direction. One possibility would be to let p_1 and p_2 be the probability that an innovation in (respectively) the pre- and post-1980 period was made by Italy, and then test $H_0 : p_1 = p_2$ versus $H_a : p_1 > p_2$. Another possibility is to use Fisher's exact test with an alternative, as described next.

The odds ratio in a 2×2 table is defined as,

$$\text{OR} = \frac{p_1/(1-p_1)}{p_2/(1-p_2)};$$

the hypothesis of homogeneity is equivalent to $H_0 : \text{OR} = 1$. We can perform Fisher's exact test specifying alternatives $\text{OR} > 1$ or $\text{OR} < 1$. In our particular case, we are interested in the alternative $p_2 < p_1$ or $\text{OR} > 1$. We can do:

```
> fisher.test(farma,alternative="greater")
        Fisher's Exact Test for Count Data

data:  farma
p-value = 0.3288
alternative hypothesis: true odds ratio is greater than 1
95 percent confidence interval:
```

```
 0.6678711       Inf
sample estimates:
odds ratio
  1.278811
```

We see that the p-value is still quite large, 0.3288. The sample odds ratio is,

```
>  (119 / 1163) / (8 / 100)
[1] 1.27902
```

larger than one, but not large enough to reject the null hypothesis.

6. Are the previous two questions *the same* question? Why or why not?

   Answer: No, clearly not. For instance, if the table had been

```
> farma2
          Italy RestWorld
1961-1980   119      1163
1980-1983    25        83
```

the results of the test would have been:

```
> chisq.test(farma2,correct=FALSE)
        Pearson's Chi-squared test

data:  farma2
X-squared = 20.6222, df = 1, p-value = 5.594e-06
```

This is highly significant, but could hardly be understood as evidence of a decline in the Italian innovation rate, as the proportion of innovations made by Italy has actually increased from the pre-1980 level (25/108 = 0.2314 is far greater than 119/1282 = 0.0928).

In conclusion, the perceived decline in the Italian rate of innovation is far from a strong evidence; from a statistical view point, it cannot be substantiated.

It is important that you develop skills so you can assess evidence in cases such as the one presented.

References

[1] M. Boldrin and D.K. Levine. *Against Intellectual Monopoly*. Cambridge Univ. Press, 2010. Freely available at http://www.dklevine.com/papers/imbookfinalall.pdf.