Ordinal regression_Part 3: Proportional odds assumption

Author: Dr Heini Väisänen

Transcript of: https://youtu.be/iPDwMrw2D30

This video is part of an NCRM Online Resource

Hi and welcome to the third video of the ordinal regression model series at NCRM. My name is Dr Heini Vaisanen and I'm a lecturer at the University of Southampton. In this last video we will talk about the proportional odds assumption, which is a very important assumption that we are making every time that we run ordinal logistic regression models. We will talk about what this assumption means and how we can test or examine whether that assumption holds. I will also compare multinomial and ordinal models, so I hope that you watch the multinomial logistic model videos as well or are otherwise familiar with this method, because it also helps with understanding the proportional assumption and it helps examine that assumption a little bit, like you will see.

So what is this proportional odds assumption that I keep mentioning? Well when we, whenever we run an ordinal model, we assume that there is a common slope for the effect of any of the explanatory variables that we have included in our model, that applies for each of the cumulative logits that we are fitting, and that is seen by the fact that you only have one set of odds ratios whenever you run an ordinal model. If you compare that to a multinomial model where you would have different sets of odds ratios for every equation that you are modelling, you can see that the ordinal model is very different, you only have one set of odds ratios, the only thing that changes is the intercept. As long as you can make this assumption, that the slopes are the same for every cut point of your ordered outcome, ordinal models are great because they are more parsimonious than the multinomial model, and you only have to deal with one set of estimates for your odds ratios. However, we do need to check if this assumption holds. If it doesn't, then your estimates from your ordinal model will not be very reliable and you will have to use multinomial model to estimate better results. One way of doing this is to formally test this statistically and a test that can do this is called 'test of parallel lines', so essentially it is testing whether the slopes in your model actually are the same, so whether we can actually apply the same odds ratio across the different categories of the outcome or not, and whenever we conduct this test our null hypothesis are saying that the slopes do not differ and the alternative hypothesis is saying that the slopes are different.

Essentially what this means, is that when you're conducting this test, if you get a large p-value as a result it means that we fail to reject null hypothesis and we conclude that the proportional odds assumption is reasonable and that we can use the ordinal model. However, if you get a small p-value it means that we do reject null hypothesis, that the slopes do not differ and we conclude that proportional odds assumption is not reasonable and we cannot use the ordinal model. So while this is very logical when you're thinking about the different hypothesis that we're making, it is sort of the other way around as you normally do in statistical testing. Normally you hope for small p-values, but here you actually hope for a big p-value. There are different ways of conducting this test, SPSS for instance, very readily provides this option, when you're running an ordinal model, in Stata is not included in the o-logit command that is the default command for ordinal models in Stata, but you can download another command called 'omodel' and in the computer workshops materials I show you how to do that, which then conducts this test for you. In terms of examples, I'm continuing to use the same example as I did in the second video, so I'm not going to go through that in a lot of detail, we're still interested in how gender and education are associated with how worried people are about their homes being broken into, and our outcome has four categories ranging from, one 'not at all worried', to four 'very worried', so we have four categories, ordered categories of the outcome.

The model that you can see on this slide is the same as the model that you saw in the second video, but what's different here is that now we have the result of the test of parallel lines, and you can see that the p-value that is associated with the test statistic that's reported on this slide is very small, which means that we do reject null hypothesis of the slopes not being different, which means that in our model where we use gender and education as our explanatory variable, the assumption of proportional odds was not reasonable. I would like to say before we move forward that not everyone is the fan of this test and there are other ways that you can, that are slightly more informal that you can use to examine the assumption of whether you have met the parallel lines assumption or not, and here is where multinomial models can help us and I will talk you through how that works.

So you might remember that in a multinomial logit model, we model response probabilities. So we want to know what is the probability that an observation or a respondent is in a given category, in that exact category, which is different from ordinal models where we model cumulative probabilities. In a multinomial model we take one of the outcome variables categories as the reference and then we compare all the other categories to that one category, which is again different from ordinal models where we are using the ordering of the categories to build our model. In a multinomial model we then estimate pairwise contrasts between each response category and its reference, and since we have, this essentially means that we run different equations for every different pairwise contrast. So if we have four categories of the outcome, like we do in our example, we will end up with three different equations, that each have a different intercept and different odds ratios or slopes for your estimates.

So this means that unlike in an ordinal model, we are not making any assumptions about whether the slopes are the same or different, across the different categories of the outcome. They can be parallel or they can be very different from each other and that means that in every case, even if the assumption of proportional odds holds, multinomial model will always give you more precise estimates than an ordinal model, and that's because we are making fewer assumptions.

You might be wondering why we would ever bother with an ordinal model if we get better answers or more precise answers from a multinomial model. Well the reason is that in statistics we normally are trying to find the best possible model that is as parsimonious as possible. So we want the number of parameters in our model to be as small as possible and that's why we test statistically, significance and think carefully which models, which variables we want to include, because we don't want to have too many variables in our model and also that goes for your modelling strategy and so if there is a model that can get good estimates with a smaller number of parameters, you'd normally go for that, even if that means that there is a slight loss of position. Another advantage of the ordinal model compared to a multinomial one, is that you have common effects that you can interpret across the different levels of core areas. So because you only have this one set of estimates, you say, you know this applies to all the different categories for the outcome, versus the multinomial model, you kind of have to compare every pair by the equation separately.

So my recommendation would be that if you have an outcome variable with more than two categories that are ordered, you start with an ordinal model and see if that works and see if your assumption of proportional odds are being held or not, and you only consider a multinomial model if the proportional odds assumption is invalid.

Okay, so now I will show you, using the same example that we've been looking at in the second video and in here as well, which is the example of being worried about one's home being broken into and whether gender and education are associated with that. I will run a multinomial model and using that data, and then compare what it looks like, compared to the ordinal model that we ran earlier. So here are the results for the first equation. So we decided to use 'not at all worried' as the reference category, so that's the lowest category of the audit outcome, and the first equation is comparing those who were 'not very worried', so the second category of the ordered outcome, to those who were 'not at all worried', and we can see similarly to what we, to our results from the ordinal model, that women are more likely to be in the more worried categories than men, and as education increases then does the likelihood of being less worried.

The similar pattern for 'fairly' versus 'not at all worried', again women are more likely to be fairly worried, rather than not at all when we compare to men, and those with some education, when we compare to no education, are more likely to be fairly, rather than not at all worried.

Then finally 'very worried' versus 'not at all worried', again women more likely to be very worried than men, and when it comes to education, actually here is slightly different from the other equations, so those with higher levels of education are less likely to be in the 'very worried' category, than in the 'not at all worried' category, compared to those with no education. But as you can see it's already kind of complicated to say what is going on, it's much more complicated than it was when we were interpreting the results in the second video, and that's because we have these three different equations with different slopes and intercepts every time.

If you want to see what's going on with your proportional odds assumption, one thing that you can do is you can calculate predictive probabilities for a range of different combinations in your data set, or in your model, and compare the probabilities from a multinomial model, which we know is always more precise than the ordinal model, and compared those to the probabilities that you get from the ordinal model, and that can give you some clues about whether your proportional assumption is working or not. Here I've taken women only, so I haven't calculated any probabilities for men in this case, and I've calculated probabilities for the different combinations of education and the four different categories of the outcome, so the level of worry that these respondents have, and we can start by looking at the no education category. You can see that the probabilities of being in the different categories of the outcome among the no education category, are quite different between the multinomial and the ordinal model. So for instance, the multinomial model says that about 16% of those with no education are in the 'not at all worried' category among women, whereas in the ordinal model that's only 11.4%. So the ordinal model here is underestimating the probability of being in this category for those with no education. Similarly if you look at 'very worried', the multinomial model is estimating that about 18% of women or 18% probability of women with no education of being very worried, whereas the ordinal model is again underestimating the proportion of women in this category saying that it's about 13%. So there are quite big differences in pretty much all of the categories for this group of people. For the two middle categories there aren't that many big differences, there are some differences, but they're not huge, but then again when we look at the degree category we can see that there are bigger differences here. So again the multinomial model says that there are about 13% of women or 12.6% of women who are not at all worried in the degree category of education, whereas the ordinal model overestimates this percentage saying that it's 15%. If we look at women with degree education and very worried category, the multinomial model says that it's about 7% probability of being in that category, whereas the ordinal model over estimates this to be about 10%. This is something that you see when you have an ordinal model that doesn't fit very well, so an ordinal model where we are violating the proportional odds assumption, and we kind of knew that already based on our test results as well, but this is another way of getting more information what is going on, and since we know that the multinomial results will always be more precise than the ordinal results, what we would like to see is predictive probabilities that are very close to each other, they don't have to be exactly the same, but if they are close and then that means that your model is probably doing better in terms of the parallel lines assumption, than if they are very different, and in here they are quite different so our model is not doing so well.

If we had calculated the same probabilities for the example that we were using the very first ordinal video, which was looking at the likelihood of applying for a postgraduate education, based on whether your parents are educated or not, that would have been a model where their proportional assumption was reasonable, and here are predicted probabilities for the different categories of the outcome. So how likely are you to apply for postgraduate education, unlikely, somewhat likely, very likely, and then the binary variable for parents not educated and parents educated. The predictive probabilities that you see in the cells in the upper row are from a multinomial model and the ones in the lower row inside parentheses are from an ordinal model, and as you can see they are very close, for four categories they are actually exactly the same and there are two categories, so parents educated and somewhat likely and very likely categories where there are small differences, but they are quite small, so we can see that our ordinal model is doing a good job and estimating well what the probabilities of being in these different categories are.

So to sum up this slightly complicated lecture you can either use ordinal or multinomial models whenever you have a multi-category outcome where the categories are ordered, and the differences between the models are that an ordinal model is more parsimonious than the multinomial model, but the downside is that it makes the proportional odds assumption which doesn't always hold, and that's why you need to test whether it holds, whether it is by using the test of parallel lines or comparing the predicted probabilities between multinomial and original models and seeing how close they are, or maybe doing both, but one way or another you need to examine this. If the assumption is not met, then you will have to use a multinomial model, but if the assumption is met then it's better to use an ordinal model because it is more parsimonious.

That's all from me today, but I do encourage you to take a look at the computer workshop materials because they will also give you a bit more information about how to actually use these models hands-on. Thank you.