In this third video on multilevel modelling we will consider the random coefficient model. This extends the random intercept model, by relaxing the assumption that the relationship between our dependent variable and the explanatory variables is the same across all groups. This can be used for continuous or categorical explanatory variables. We will also consider the use of cross level interactions to examine how the strength of an individual level relationship may be moderated by an included group level effect.

Previously we considered the relationship between weight and height. Here we assumed that a one unit increase in weight, was associated with the same increase in height across all families, but we allowed the point where the line crossed the y axis to vary across families. But what if the relationship between weight and height is not uniform across all families. The random coefficient model allows this by including an additional set of residuals, UJ, for each explanatory variable. In our case we have one explanatory variable, X1, so we include one new set of residuals U1J. So, in family one the residual is positive, meaning the relationship between weight and height is stronger than the overall average by beta 1 plus U11; And in family two the residual is negative, meaning the relationship between weight and height is weaker than the overall average, beta 1 minus U12. Note that we also now attach the subscript naught to the original U residuals to show that these now relate specifically to the intercept. Note also that the intercept residual is now dependent on where we choose to make the value of x = 0, and refers only to variation at this point. It's therefore often advisable to centre explanatory variables when we include them in our models.

Now we have two sets of group level residuals, one associated with the intercept U0J, and one associated with the coefficient U1J. These residuals are assumed bivariate normal, with mean zero, and variance is summarized by variance covariance matrix. The matrix is composed of an intercept variance Sigma Squared U0, which accounts for variability in the intercept, a coefficient variance Sigma Squared U1, which accounts for variability in the effect of weight on height, and a covariant term between the two sets of residuals U01. This covariance term tells us how the coefficient and intercept residuals are linked. So, a positive value tells us that in general, groups with a positive intercept residual tend to have a positive coefficient residual and groups with a negative intercept residual tend to have a negative coefficient residual. A negative value tells us that in general groups with a positive intercept residual tend to have a negative coefficient residual, and groups were the negative intercept residual tend to have a positive coefficient residual. But the interpretation of this covariance term requires care and is dependent on the relationship between Y and X. For example, when the relationship between Y and X is positive, a positive covariance means that when the intercept residual is positive, and hence the intercept is higher than average, there will also be a positive coefficient residual. This means a stronger than average coefficient, because you are adding a positive residual to the positive slope. In contrast a negative covariance term means that a higher than average intercept is associated with a weaker than average coefficient, because you're adding a negative residual to the positive slope. When the coefficient is negative the opposite happens. A positive covariance means that when the intercept is higher than average, you're adding a positive residual to a negative coefficient, making it weaker; and a negative covariance means that when the intercept is higher than average, you're adding a negative residual to a negative coefficient, making it stronger in the negative direction.

So, returning to our fear of crime example, we allow the positive effects of victim status to vary randomly across areas in model 3. This allows for the possibility that the high levels of fear amongst victims of crime, may not be so apparent in some areas, whilst in other areas the difference between victims and non-victims may be bigger than average. Here we identify positive variance term, suggesting that there are differences in the magnitude of the effect of being a victim of crime, on levels of fear across neighbourhoods. Here we identify a significant variance term suggesting that there are differences in the magnitude of the effect of being a victim of crime, on levels of fear across neighbourhoods. We also identify a negative covariance term, which suggests that in areas where the intercept is higher than average, the gap between victims and non-victims will be smaller than average. Remember a negative covariance means that positive intercept residuals tend to go with negative coefficient residuals, and as the intercept refers to a non-victim, this means, then in areas where non-victims tend to report more fear of crime, there are less notable differences between victims and non-victims. In other words, everyone is more fearful in these areas.

We can represent this graphically by considering the differences between victims and non-victims for four sampled neighbourhoods. These areas have been ranked based on the levels of fear of non-victims, the red circles, from lowest on the left to highest on the right. Here we can see that in areas where non-victims tend to report more fear, towards the right of the graph, there is a smaller increase in fear amongst victims of crime. Finally, we can incorporate level interactions to more directly model how individual level relationships are moderated by features at the group level. This is the exact same logic as interaction effects in single level regression models. So, the interaction between X1 and X2 is just X1 by X2. The only difference here is that X1 and X2 can be from different levels of analysis. So, considering fear of crime, model for extends our analysis by also including the interaction term between individual victim status and the area crime rates. This is identified as significant and negative. This suggests that for victims of crime, the positive association between the crime rate and fear is weaker than it is for non-victims. Alternatively, we could say than in areas where the crime rate is higher than average, the difference in fear between victims and non-victims will be smaller than average. Everyone will be more fearful. Both explanations are technically valid in this context. Of course, this model needs extension, but it provides a basic example of the multilevel model.

So, to sum up, in this video we've introduced the random coefficients model. This relaxes the assumption of a random intercept model that the relationship between X and Y will be the same in every group. Instead we allow a residual difference in the magnitude of the coefficient for each group, and estimate their variance and covariance with the intercept. We can also use cross level interactions to more directly model the connections between individual and group effects. Although these models have been introduced in relation to a continues outcome, a two-level model and a single random coefficient, this then generalizes neatly to non-normal data, more than two levels and multiple random terms.