

Online supplement

We checked the distribution of all outcome measures. Some of these were skew; hence medians (interquartile range) and means (standard deviation) are presented in Table 1. After linear regression, we checked the assumption of normality of the residuals for each model by plotting them on a histogram. Some were normally distributed, despite the outcome variable not being normal. Where the residuals were deviant from normal, we transformed the outcome using the most appropriate transformation to make the outcome normally distributed (or very close to normal) and then repeated the linear regression. For some outcomes there was not a transformation that made the outcome near normal. Where this was the case, we did not attempt to fit a model with a transformed outcome. Where transformation occurred, the models with the transformed and untransformed outcomes were compared using Akaike Information Criteria (AIC). In the majority of cases, the models with the untransformed outcome were the best using this criterion. For others, there was little difference between them so we have presented results from the untransformed models in Tables 2 and 3.