

The *t*-statistics for interactions involving `tint.Q` are 0.46, -0.15, 1.34 and 1.10. The output can be simplified by omitting these interactions.

None of the main effects and interactions involving `agegp` and `sex` are significant at the conventional 5% level, though `agegp` comes close. On the other hand, the interaction terms (`tint.L:agegpOlder`, `targethicon:agegpOlder`, `tint.L:targethicon` and `tint.L:sexm`) that are statistically significant stand out much less clearly in Figures 2.12A and 2.12B.

This may seem inconsistent with Figures 2.12A and 2.12B, where it is the older males who seem to have the longer times. To resolve this apparent inconsistency, observe that

- Comparisons that relate to `agegp` and `sex` are made relative to variation between individuals. Standard errors for such comparisons, in the output, are in the range 0.23 - 0.32, in each case with 22 degrees of freedom. (There are 9 younger and 4 older females, against 4 younger and 9 older males.¹⁰)
- Comparisons between levels of `tint` or `target` are made several times for each of the 26 individuals, and are relatively consistent from one individual to another. Standard errors for these comparisons are small – in the range 0.042 - 0.058.

Statistical variation cannot be convincingly ruled out as the explanation for the effects that stand out most strongly in the graphs. The graphs are not designed to highlight the consistency with which individuals respond to differences between levels of tinting and target contrast.

10.5 A Generalized Linear Mixed Model

Consider again the `moths` data of Subsection 8.4.2. The analysis in Subsection 8.4.2 assumed a quasipoisson error, which uses a constant multiplier for the Poisson variance. It may be better to assume a random between transects error that is additive on the scale of the linear predictor. The model incorporates a term that allows for normally distributed random variation, additional to the poisson variation at each observation. Technically, this is an example of the use of “observation level random effects”.

The attempt to fit a model that uses the default log link generates (lme4.1.1-7), if data for the habitat Bank is included, a warning that the model is nearly unidentifiable. This problem is avoided if a square root link is used.

The code is:

```
moths$transect <- 1:41 # Each row is from a different transect
moths$habitat <- relevel(moths$habitat, ref="Lowerside")
A.glmmer <- glmer(A~habitat+sqrt(meters)+(1|transect),
                 family=poisson(link=sqrt), data=moths)
print(summary(A.glmmer), show.resid=FALSE, correlation=FALSE)
```

Output is:

```
. . . . .
      AIC      BIC  logLik deviance df.resid
      213      230    -96     193      31
```

¹⁰`subs <- with(tinting, match(unique(id), id)); with(tinting, table(sex[subs], agegp[subs]))`

Random effects:

Groups	Name	Variance	Std.Dev.
transect	(Intercept)	0.319	0.564

Number of obs: 41, groups: transect, 41

Fixed effects:

	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	1.7322	0.3513	4.93	8.2e-07
habitatBank	-2.0415	0.9377	-2.18	0.029
habitatDisturbed	-1.0359	0.4071	-2.54	0.011
habitatNEsoak	-0.7319	0.4323	-1.69	0.090
habitatNWsoak	2.6787	0.5101	5.25	1.5e-07
habitatSEsoak	0.1178	0.3923	0.30	0.764
habitatSWsoak	0.3900	0.5260	0.74	0.458
habitatUpperside	-0.3135	0.7549	-0.42	0.678
sqrt(meters)	0.0675	0.0631	1.07	0.285

The Poisson component of the variance, on the square root scale of the linear predictor, is 0.25. The observation level random effect, labelled `transect` in the above output, increases this by 0.319 to 0.569, i.e., by a factor of 2.28. Compare this with the increase by a factor of 2.7 for the quasipoisson model.

Now compare, between the quasipoisson model and the observation level random effects model, predicted values for habitats and standard errors of difference from `Lowerside`:¹¹

	fit-glm	glm-SE	fit=glmer	glmer-SE
Lowerside	2.13	0.00	1.99	0.00
Bank	0.00	0.86	0.00	0.95
Disturbed	1.07	0.41	0.86	0.40
NEsoak	1.53	0.43	1.42	0.41
NWsoak	4.86	0.54	4.72	0.51
SEsoak	2.30	0.41	2.18	0.39
SWsoak	2.58	0.54	2.53	0.51
Upperside	2.37	0.45	2.34	0.43

Observe that the standard errors for comparisons with `Lowerside` are similar for the two models. The fitted values in the the observation level random effects model are pulled in towards zero, relative to the quasipoisson model.

It is left as an exercise for the reader to compare the plots of residuals versus fitted values between the two models.

```

11A1quasi.glm <- glm(A`habitat, data=moths, family=quasipoisson(link=sqrt))
A1.glmmer <- glmer(A`habitat+(1|transect), data=moths, family=poisson(link=sqrt))
Cglm <- coef(summary(A1quasi.glm))
Cglmer <- coef(summary(A1.glmmer))
fitboth <- cbind("fit-glm"=Cglm[1,1]+c(0, Cglm[-1,1]), "glm-SE"=c(0, Cglm[-1,2]),
                "fit=glmer"=Cglmer[1,1]+c(0, Cglmer[-1,1]), "glmer-SE"=c(0, Cglmer[-1,2]))
rownames(fitboth)[-1] <- substr(rownames(fitboth)[-1],8)
rownames(fitboth)[1] <- "Lowerside"
round(fitboth, 2) # NB, all SEs are for the difference from 'Lowerside'

```