ESTIMATING THE SUBJECT BY TREATMENT INTERACTION IN NON-REPLICATED CROSSOVER DIET STUDIES

Matthew Kramer[1], Shirley C. Chen[2], Sarah K. Gebauer[2], and David J. Baer[2]

[1]Biometrical Consulting Service, ARS/USDA, Beltsville, MD 20705
[2]Beltsville Human Nutrition Center, ARS/USDA, Beltsville, MD 20705

Abstract

Researchers in human nutrition commonly refer to the 'consistent' diet effect (i.e. the main effect of diet) and an 'inconsistent' diet effect (i.e. a subject by diet interaction). However, due to the non-replicated designs of most studies, one can only estimate the first part using ANOVA; the latter (interaction) is confounded with the residual noise. In many diet studies, it appears that subjects do respond differently to the same diet, so the subject by diet interaction may be large. In a search of over 40,000 published human nutrition studies, most using a crossover design, we found that in none was a subject by diet interaction effect estimated. For this paper, we examined LDL-cholesterol data from a non-replicated crossover study with four diets, the typical American diet, with and without added plant sterols, and a cholesterol-lowering Step-1 diet, with and without sterols. We also examined LDL-cholesterol data from a second crossover study with some replications with three diets, representing the daily supplement of 0, 1 or 2 servings of pistachio nuts. These two data sets were chosen because experience suggested that LDL-cholesterol responses to diet tend to be subject-specific. The second data set, with some replication, allowed us to estimate the subject by diet interaction term in a traditional ANOVA framework. One approach to estimating an interaction effect in non-replicated studies is through the use of a multiplicative decomposition of the interaction (sometimes called AMMI—additive main effects, multiplicative interaction). In this type of analysis, residuals, formed after estimated main effects are subtracted from the data, are arrayed in a matrix with diets as columns and subjects as rows. A singular value decomposition of the matrix is performed and the first, or first and second, principal component(s) are used as estimates of the interaction, and can be tested for significance using approximate F-tests. Using the R gnm package, we found large and significant subject by diet interaction effects in both data sets; estimates of the interaction in the second data set were similar to interaction estimates from traditional ANOVA. Of an additional 26 dependent variables from the first and a third data set (the latter investigating the effect of mild alcohol consumption on blood variables), 19 had significant subject by diet interactions, based on the AMMI methodology. These results suggest that the subject by diet interaction is often important and should not be ignored when analyzing data obtained from non-replicated crossover designs—the AMMI methodology works well and is readily available in statistical software packages.

Key words: nutrition, crossover design, subject-treatment interaction

1. Introduction

The use of crossover designs for human nutrition diet studies were discussed in the 1930's with the writings of T.W. Simpson (Jones and Kenward 1989). Early on in this literature, many of the peculiarities characteristic of crossover designs were investigated, such as period and carry-over effects (Jones and Kenward 1989) but there appears to have been little interest in estimating subject-diet interactions. Traditionally these designs are not replicated (each subject goes through a series of different diets without repeating any of them) and analyzed in an ANOVA framework. Without replication, the subject-diet interaction is confounded with the residual term. If this interaction term is large, tests on the main effects in a linear model become more conservative.

There appears to be some recognition that the subject-diet interaction may be large and should be accounted for in clinical trials (Hauck et al. 2000). Some sophisticated statistical techniques have been developed for this problem, i.e. those of Ghosh and Crosby (2005). Their idea was to group subjects that respond similarly across the diets, then to treat these groups as if they were replicates of the same individual. One can then estimate both a subject-diet interaction term and a residual. If statistics is part science and part art, the part that is art involves deciding how to group subjects. They show that it is possible to use fairly objective criteria for a single dependent variable. However, our own experience, using a clustering algorithm and multiple dependent variables, is that subject group compositions change for different collections of dependent variables, and this impacts the estimates of the interaction and residual terms.

A different approach, and one that has apparently not previously been explored by researchers in the nutrition field, is to use a multiplicative decomposition of a 'residual' formed by subtracting estimated main effects from the data. This 'residual' contains both the subject-diet interaction as well as within-subject noise (error). A multiplicative decomposition of this 'residual', using, say, principal components, can be used to extract the signal, i.e. the subject-diet interaction. One can use approximate F-tests to determine how many principal components are required (typically one or two), resulting in a partition of this 'residual' term into a part that can be attributed to subject-diet interaction and a part that represents within-subject variability.

Considering that the use of applying principal components to the residuals of a non-replicated two factor ANOVA to estimate an interaction term has been available since 1968 (Gollob 1968), and special cases of it even earlier, such as Tukey's (1949) one d.f. test for interaction, it is remarkable that a search through the literature (over 40,000 studies in the last 10 years) for subject-diet interactions in human nutrition studies yielded no "hits". We searched issues of the *Journal of Nutrition* and the *American Journal of Clinical Nutrition* for articles (both primary research and review articles) from 2000 to early 2010 with "crossover" (or a similar term indicating the study used a crossover design) and phrases indicating that the authors were aware of a potential subject by diet interaction (e.g. "subject by diet interaction", "subject by treatment interaction", "repeated treatments", "inconsistent diet/treatment effect", etc.). Finding no "hits" at all was surprising because researchers seem to be well aware of the heterogeneity in responsiveness to dietary interventions (Rideout 2011), and many of the studies were done in collaboration with statisticians. While our search techniques may not have been

optimal, the lack of "hits" is consistent with our subjective opinion that these researchers generally do not recognize that a subject by diet interaction term is missing from their analyses.

A new wrinkle has been introduced with the advent of mixed models software, human nutrition experiments are now often analyzed with subjects as a random effect (thus a subject-diet interaction would also be a random effect). Although perhaps unrecognized by researchers, ignoring effects in mixed models does not necessarily have the same outcome as ignoring them in linear models with only fixed effects. While ignoring a fixed subject-diet interaction term in a linear model makes the test on diet more conservative, ignoring a random subject-diet interaction term in a mixed model can make the test on diet excessively liberal, as demonstrated in Boykin et al. (2011). In a small mixed model simulation, mimicking the kind of data collected in non-replicated human nutrition studies, we found that, while among-subject variance was accurately estimated, the estimated 'residual' variance was about 20% less than the sum of the true variance contributions of error and subject-diet interaction. Thus the estimated variability in an experiment was, on average, less than that actually present, leading to inflated F-statistics, consistent with results from Boykin et al. (2011).

In this paper, we work through the estimation of the subject-diet interaction in a non-replicated human nutrition study. We do a similar analysis for a second study which differed in an important way; there was some replication, allowing us to directly compare a model with a multiplicative decomposition of the subject-diet interaction with one analyzed using traditional ANOVA methods. This is important because a multiplicative decomposition of the interaction using principal components may not capture the kind of subject-diet interaction estimated using traditional ANOVA (had there been sufficient replication), the way most researchers would recognize and understand this interaction. We then briefly discuss results from the first data set and a third data set, both with many different dependent variables, to see how often a subject-diet interaction occurred.

2. Description and analysis of dataset 1

An analysis of these data was published by Chen et al. (2009). The object of the study was to determine if the main effects of diet (typical American diet versus recommended cholesterol-lowering Step-1 diet) and incorporation of plant sterols (believed to lower LDL-cholesterol, 0 and 3.3 g/day) were additive. A number of blood compounds were measured on 22 adult men and women, we only discuss results for LDL-cholesterol. Each subject went through each of the four diets, each lasting 23 days, with no washout period. As is typical of crossover designs, different cohorts of the subjects were on different diets at each of the four periods. Measures are means of two samples, from day 22 and day 24 of each period. Baseline (pre-experiment) measurements were taken during the week prior to the beginning of the experiment. The original data were analyzed without transformation, in this analysis we took natural logs. Figure 1 gives scatter plots representing each subject on each of the four diets; the blue line in each plot represents no change from baseline. What is obvious from this figure is the large effect of adding plant sterols to a diet (2nd column).

In a traditional linear model context, with all factors fixed, the first step is to subtract the main effects of diet and subject to form 'residuals', shown in Fig. 2. These 'residuals' contain the confounded within-subject error and diet-treatment interaction effects. The AMMI (additive
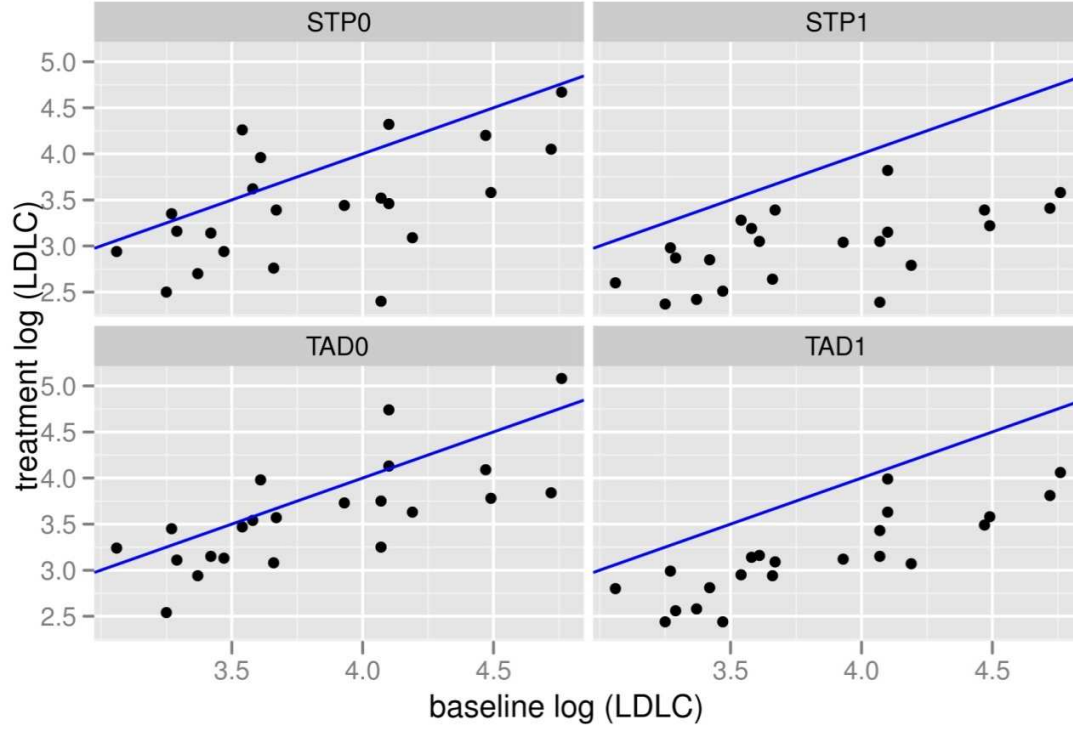
Figure 1. Scatter plots of pre-experiment baseline log LDL-cholesterol values versus log LDL-cholesterol values on each of the four diets. The blue line represents no change from baseline. STP = Step1 diet, TAD = typical American diet, 0 = no plant sterols added to diet, 1 = plant sterols added to diet.

main effects, multiplicative interaction) model does a principal components decomposition of the 'residuals', after they have been arrayed into a subject (row) by diet (column) matrix. Typically, the first (or first and second) principal component(s) are used to capture the interaction; the remainder of the variance goes to within-subject error. The AMMI model can be written as

$$y_{ijk} - \left(\hat{\mu} + \widehat{\boldsymbol{\beta}}\boldsymbol{x}_j + \hat{\tau}_i + \hat{\gamma}_j \right) = \sum_r \lambda_r v_{ir} \delta_{jr} + \epsilon_{ijk},$$

where y are LDL-cholesterol data, i indexes diets, j indexes subjects, k indexes diet repeats for subject j, μ the overall mean, $\boldsymbol{\beta}$ the vector of slopes for covariates $\boldsymbol{x}$, τ the overall diet effect on LDL, γ the subject effect, λ the singular value for component r, v the eigenvalue score for diet i and component r, δ the eigenvalue score for subject j and component r, and ϵ random error. For the data and model just described, $k = 1$ (i.e. no repeats), γ (the subject effect) is considered to be a fixed effect, and there are no covariates. Typical covariates that are considered for these kinds of studies are gender, age, body-mass index (BMI), design effects (e.g. period, carryover,
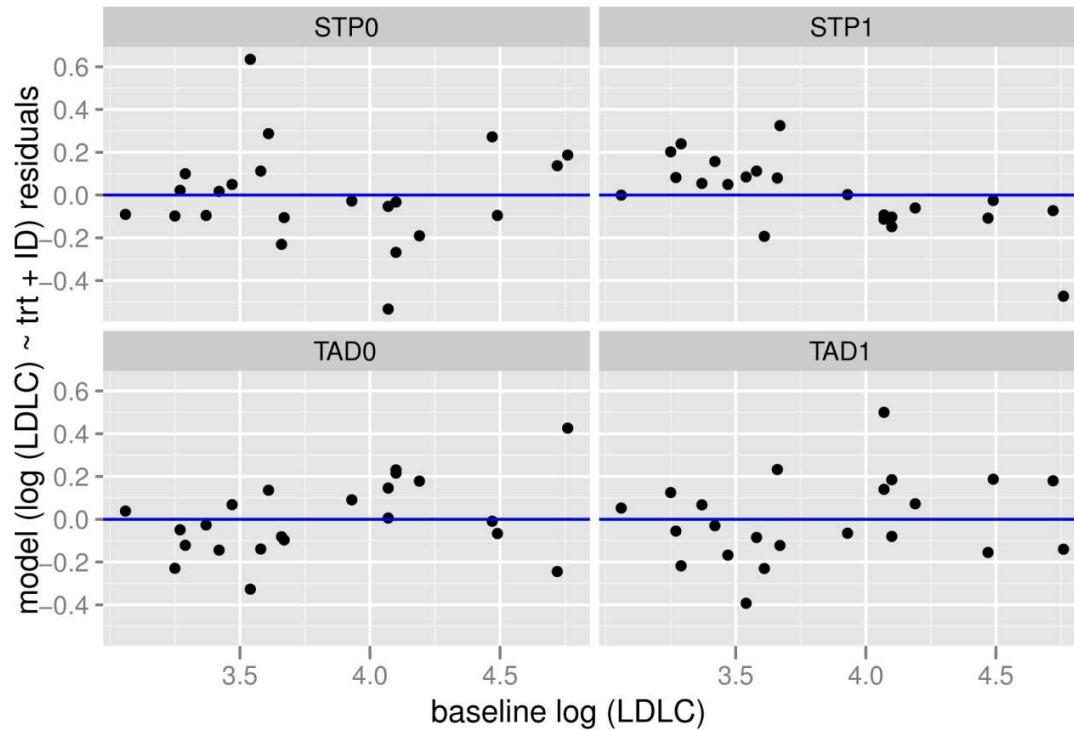
Figure 2. Scatter plots of baseline log LDL-cholesterol values versus 'residuals', formed by subtracting the diet and subject main effects from the data. See Fig. 1 caption for explanation of plot titles. The horizontal blue line represents the mean of the 'residuals', zero, useful to judge if baseline values are predictive of diet effects, after removing the subject main effect. There is some indication that the model somewhat under predicts LDL-cholesterol for subjects with lower baseline values on the STP1 diet.

sequence), and, if γ is considered as a random effect, the subject's baseline measurement. For this data set, none of these candidate covariates appeared useful.

Figure 3 may help to understand in a graphical way what the principal components decomposition is doing. Essentially, the first principal component rotates and scales the 'residuals' in such a way that, at least for the STP0 and TAD1 diets, the 'residuals' can be replaced by a line. In the traditional way of accounting for degrees of freedom (e.g. Gollob 1968), this first principal component 'costs' 23 degrees of freedom, the formula is: (num. trts – 1) + (num. subj. – 1) + (2 * i – 1), where i is the ith component, here 1. For these data, the second principal component also appears useful (Fig. 4). It appears to capture the interaction effects for the other two diets, TAD0 and STP1, at a 'cost' of 21 d.f. More recent work on multiplicative decomposition models has determined that Gollob's (1968) degrees of freedom produce somewhat liberal F-tests, i.e. too many multiplicative terms can be retained. Several alternate methods to calculate degrees of freedom have been proposed (reviewed in Dias and Krzanowski
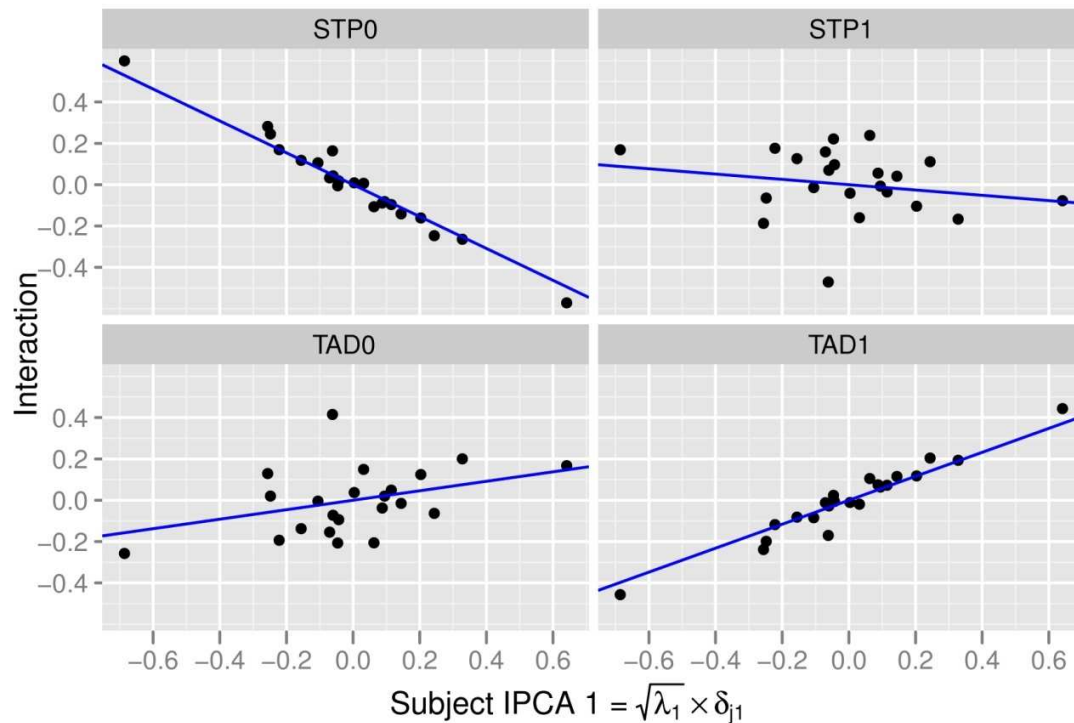
Figure 3. 'Residual' values (created by subtracting main effects from the data, so including subject-diet interaction), on the *y* axis, plotted against the rotation and scaling performed by the first principal component, on the *x* axis.

2003), many favor Mandel's (1971) estimates (obtained through simulation), though these may be too large for the first few principal components. For our data, the method used to obtain degrees of freedom does not change overall conclusions, though it does affect the number of components needed by the model.

There is software available to perform this analysis in both SAS (Lee and Johnson 2006; SAS macros are available at http://www.k-state.edu/stats/facultypages/ammi_macros.htm, along with a manual; by default degrees of freedom are calculated using Mandel's (1971) method, though Gollob's (1968) method is also available) and R (e.g. package gnm; only Gollob's (1968) degrees of freedom estimates are available). We show code and output from the R software (R Development Core Team 2011), since R was used for the analysis of the data presented here, but also include test results based on Mandel's (1971) degrees of freedom. The gnm (generalized nonlinear models) package (Turner and Firth 2011) fits a broad spectrum of models, including AMMI models. Of potential interest to users of AMMI models, this package can also fit models from non-normal distributions, e.g. Poisson and binomial, so may be of use for count data. We know of no other software that will fit native count data with an AMMI model. It does not (yet) fit random effects.
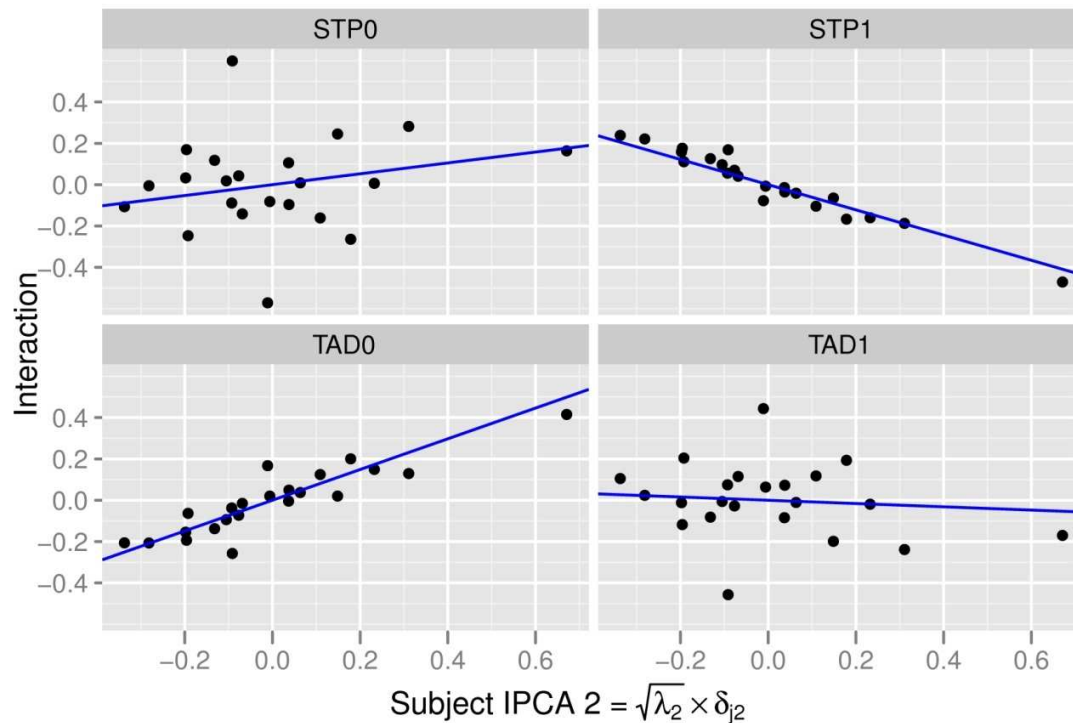
Figure 4. 'Residual' values (created by subtracting main effects from the data, so including subject-diet interaction), on the *y* axis, plotted against the rotation and scaling performed by the second principal component, on the *x* axis.

The following code was used to fit the data after the `gnm` package was installed and loaded, and the `lp` dataset read in.

```
maineffects1 <- gnm(LDLC ~ trt + ID, data=lp)
bilinear1 <- update(maineffects1, . ~ . + Mult(trt, lp$ID))
bilinear2 <- update(maineffects1, . ~ . + instances (Mult(trt, lp$ID), 2))
anova(maineffects1, bilinear1, bilinear2, test = "F")
```

The first line fits a basic ANOVA model. The second line updates the ANOVA model with a multiplicative interaction term using the first principal component only. The third line updates the ANOVA model with a multiplicative interaction term composed of the first and second principal components. The fourth line tells R to display *F*-test results, testing whether the increased complexity of the models with the multiplicative interaction term(s) significantly improve(s) the model, output below.

```
Analysis of Deviance Table
  Resid. Df Resid. Dev Df Deviance       F     Pr(>F)
1        63     3.2190
2        40     1.3372 23   1.8818 5.3934 0.0002225 ***
3        19     0.2882 21   1.0490 3.2930 0.0057560 **
```

Based on this output, both the first and second principal components appear to be necessary. D.E. Johnson kindly supplied us with the degrees of freedom using Mandel's simulation method, which results in the revised table below (changes are bolded).

```
Analysis of Deviance Table
  Resid. Df Resid. Dev Df Deviance       F     Pr(>F)
1        63     3.2190
2        40     1.3372 31   1.8818 2.5275   0.0454 *
3        12     0.2882 20   1.0490 2.1839   0.0832
```

The reallocation of degrees of freedom (from the residual to the principal components) changes both the *F*-values and the *p*-values (now only the first principal component is significant at $\alpha = 0.05$, though we retain both for the discussion that follows).

To examine the relative contributions to the total variance, as sums of squares, from each of the terms in the `bilinear2` model (with two principal components), the command, `anova (bilinear2),` was run, with the following output.

```
                            Df Deviance Resid. Df Resid. Dev
NULL                                          87    27.9587
trt                          3   4.8605       84    23.0983
ID                          21  19.8793       63     3.2190
Mult(trt, lp$ID, inst = 1)  23   1.8818       40     1.3372
Mult(trt, lp$ID, inst = 2)  21   1.0490       19     0.2882
```

The columns titled "Deviance" give the sums of squares (pseudo sums of squares for multiplicative interaction components and estimated within-subject variation). The largest source of variation is among subjects (ID). For the two sources that relate to treatment, the inconsistent (interaction, summing over the two principal components) diet effect is about 60% as large as the consistent (main) diet effect. The pseudo sums of squares for the multiplicative interaction terms can also be obtained directly from the singular value decomposition; they are the squares of the λ's. Note that the 'residual' sum of squares (data minus main effects) is 3.2190, only 0.2882 of that (9%) is within-subject error, so ignoring the interaction would make a test on the diet main effect much more conservative, even with the loss of degrees of freedom to the interaction.

A useful graphical tool to understand which subjects (or diets) are poorly fit by the interaction term is a biplot. As an example, the biplot for the first principal component is given in Fig. 5. Subjects and diets further from the zero line are less well captured by the first principal component, in this plot these are subjects 1806 and 1855, and diets STP0 and TAD1 (recall that
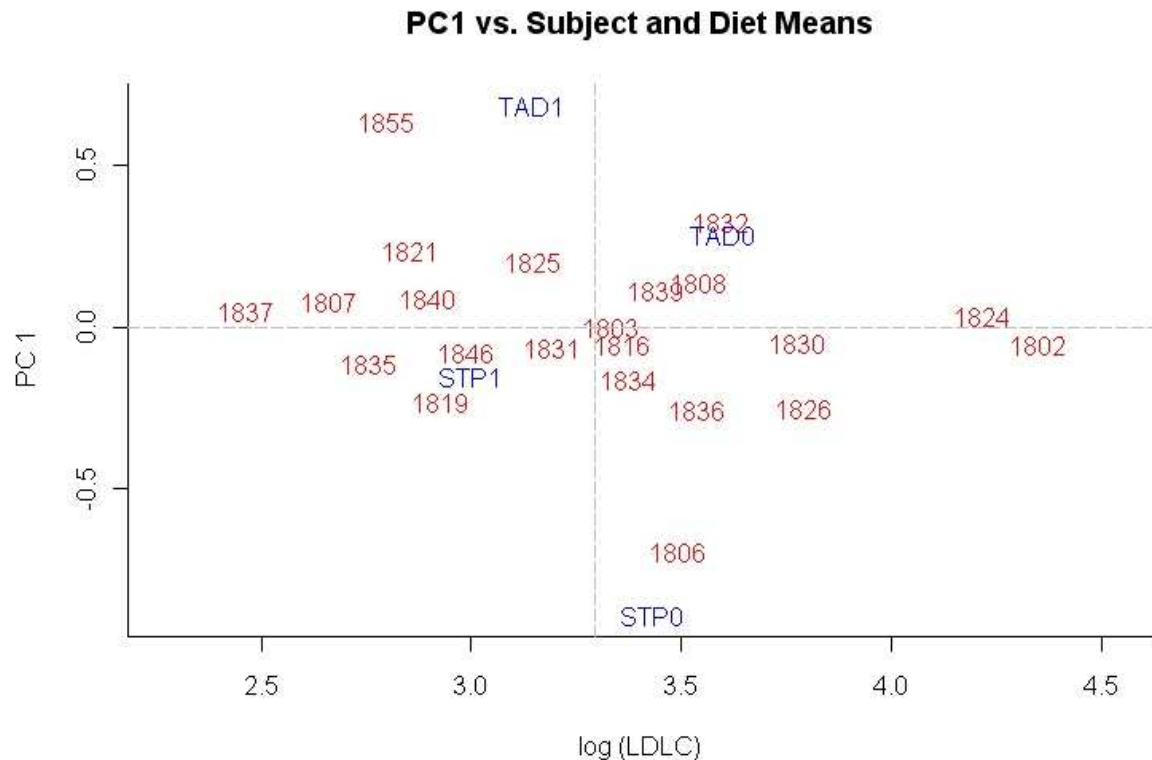
Figure 5. Biplot giving the values on the first principal component (*y*-axis) for each subject and diet and their average log (LDL-cholesterol) values on the *x*-axis. Points further from zero on the *y*-axis are less well captured by the subject-diet interaction, as represented by the first principal component.

these are the two diets whose interaction effects were not represented well with the first principal component).

Since mixed models are often used to analyze these kinds of data, it is worth investigating how estimates change by considering subjects to be random effects. Interactions between a fixed and random effect are typically considered also to be random effects. However, since it was not clear to us how to conduct a singular value decomposition in a mixed models framework (though Gogel et al. 1995 do present an algorithm for estimating such a model), what we did instead was create the 'residual' array from the residuals of a mixed model where diet was a fixed effect and subjects a random effect, baseline LDL-cholesterol was also included as a fixed effect covariate. Differences between the all fixed effects model and the random effects model were small (Table 1 gives sums of squares). These results are probably typical for these kinds of data sets as long as the number of subjects is sufficiently large to provide an accurate estimate of the random among-subject variance (i.e. there is little shrinkage), and the baseline covariate is included (which will account for much of the among subject variability as a fixed effect).

Table 1. Sums of squares (pseudo sums of squares for the interaction components) from a random effects model and from a fixed effects model, using the same LDL-cholesterol data set with four diets.

Effect	SS-random effects model	SS-fixed effects model
Diet	4.8605	4.8605
PC1	1.8855	1.8818
PC2	1.0755	1.0490

3. Description and analysis of dataset 2

For this study, there was deliberate replication (i.e. subjects each replicated one of three diets) so that a direct comparison could be made between using traditional ANOVA to estimate the subject-diet interaction and using a multiplicative interaction from a singular value decomposition (principal components) to estimate the subject-diet interaction. Since a multiplicative decomposition of the interaction does not necessarily capture the same effect as the interaction in a traditional ANOVA, we wanted to verify that a multiplicative decomposition and a traditional ANOVA analysis done using the same human nutrition dataset would yield similar estimates of the magnitude of the interaction.

The 16 subjects were stratified by gender and initial body mass index (BMI), and randomly assigned to a treatment sequence, consisting of three amounts of pistachios fed as a part of a controlled diet for three weeks each: 0 ounces of pistachios per day (control), 1.5 ounces of pistachios per day, and 3 ounces of pistachios per day. Each subject participated in two different diets; with one diet repeated twice (one subject did not repeat a diet). On the last day of each treatment period, subjects provided a fasted blood sample. There was a break of approximately one week in between each treatment period.

The study was conducted as an incomplete block design experiment; each subject was a block. Due to the study design (each subject participated in two different diets, with one diet repeated), there were missing cells in the matrix of subjects by treatments, which is not allowed for a singular value decomposition. We did two sets of analysis to circumvent this problem. One set was composed of three analyses based on two-diet subsets (so the matrices had no missing values). The second used imputation for the missing cells, based on a mixed analysis of the original data to calculate the subject by treatment variance. The 16 missing cells were imputed using a BLUP estimate to which 16 samples from a normal distribution, with mean zero and variance based on the estimated subject by treatment variance, were added. A mixed model analysis of this augmented data set produced parameter estimates very close to the original ones. The residuals from this augmented data set (omitting the subject by treatment variance term from the model) were then arrayed in a 16 by 3 matrix (with replicate subject-treatment combination

residuals averaged) for the SVD. Similar to the analysis for dataset 1, carryover, period, and other candidate covariate effects were small, so not included in the analyses. We considered subject to be a random effect, so formed our 'residual' matrices from a mixed models residual, as described above for dataset 1. The variance estimates for the interaction terms calculated with both methods is given in Table 2.

Variance estimates for the two-diet subsets are generally not as similar to each other as they were for the augmented data set, likely due to both the small number of subjects in the subsets and that the two-diet subsets are attempting to capture the subject-diet interaction with only one principal component.

From dataset 1 and another published study involving moderate alcohol consumption (Baer et al. 2002), both using non-replicated crossover designs, we had 26 dependent variables available. Using the same methods described above for mixed models, we found that 19 (73%) had significant subject-diet interactions, based on a multiplicative decomposition of the 'residual' matrix. This suggests that subject-diet interactions are widespread in nutrition studies.

Table 2. Estimates of the subject-diet interaction variance using a mixed model (column 2) or a multiplicative decomposition of the interaction using principal components (column 3).

	Subject-diet variance estimate using mixed model	Variance estimate using a multiplicative decomposition	Number of principal components retained	Number of subjects
Complete (incomplete block design)	0.04260			16
Control vs. 1 serving	0.00380	0.01049	1	5
Control vs. 2 servings	0.40976	0.16071	1	5
2 vs. 1 serving	0.02592	0.02138	1	6
Complete + augmented with imputed data	0.04052	0.04634	2	16

4. Conclusions

Despite the fact that essentially no nutrition studies using crossover designs test for a subject-diet interaction effect, researchers in the field generally understand that not every subject responds to a diet in the same way. We have formally demonstrated that a subject-diet interaction exists for many of the dependent variables in three datasets. It is not clear to us why, when smaller effects, such as period and carryover, are fussed over and routinely estimated, this potentially much more important interaction effect is ignored. Perhaps the non-replicated nature of the design (making it more difficult to estimate this interaction effect) and the general absence of discussion on treatment-subject interaction in texts on crossover designs contribute to researchers not recognizing that this term is missing from their analyses. Researchers in this field, however, often collaborate with statisticians (who 'should know better'), and the field is both sufficiently mature and sufficiently well populated by publications that this finding is overdue.

Using a multiplicative decomposition to estimate this interaction term works well for crossover designs (as it does for agricultural field trials, where it is commonly employed, see Gauch 1992) and does not require that subjects repeat diets, which is cost-prohibitive for most nutrition studies. Thus, the only additional burden on the researchers is at the analysis stage. Given that software is available in commonly used statistical packages and that the analysis is itself straightforward, this is a rather small burden for such a potentially large benefit, and we feel that nutrition researchers (and other researchers using crossover designs, such as in clinical trials) should routinely adopt this methodology for their analyses.

Of particular interest may be how much variability is accounted for by the diet main effect (consistent diet effect) versus the subject-diet interaction (inconsistent diet effect). In our datasets we found that the two parts were similar for LDL-cholesterol, but there were other blood variables where most of the diet effect was in one or the other of the two parts. If the interaction component is at all large, one must be careful about prescriptive diets, which may affect different people in different ways. Clearly, if a diet's effect is mostly expressed in the subject-diet interaction, one cannot give general guidelines.

If the subject-diet interaction is large, it may be possible to find groups of subjects that behave similarly, which may provide clues about the influence of genetic or environmental backgrounds. This is necessary as we move toward individualized diet recommendations.

5. Summary

Human nutrition diet studies using crossover designs started in the 1930's. While researchers in this field commonly refer to the 'consistent' diet effect (i.e., a main effect) and an 'inconsistent' diet effect (i.e., a subject by diet interaction), due to the non-replicated designs of most studies, one can only estimate the first part using ANOVA. The latter (interaction) is confounded with the residual noise. There are important consequences that result from not separating the interaction from residual noise. The first is that the error term is estimated to be too large, making significance tests excessively conservative in a traditional ANOVA framework. In a mixed models framework, with subjects random, ignoring the interaction has the opposite effect, the test on main effects can be too liberal. In any case, as the field is moving towards

individualized diet prescriptions, ignoring this interaction ignores a subject specific response to a diet, and one that, if investigated, might reveal interesting environmental or genetic background effects. In many diet studies, it appears that subjects do respond differently to the same diet, so the subject by diet interaction may be large. However, in a search through over 40,000 published human nutrition studies, most using a crossover design, we found that in none was a subject by diet interaction effect estimated.

We examined LDL-cholesterol data from a non-replicated crossover study with four diets, the typical American diet, with and without added plant sterols, and the Step-1 diet, with and without sterols. We also examined LDL-cholesterol data from a crossover study with some replications with three diets, representing the daily supplement of 0, 1 or 2 servings of pistachio nuts. These data sets were chosen because experience suggested that LDL-cholesterol responses to diet can be subject-specific. The second data set, with some replication, allowed us to estimate the subject by diet interaction term in a traditional ANOVA framework. The first and a third data set was also examined to see if this interaction term was present for other blood variables.

One approach to estimating an interaction effect in non-replicated studies is through the use of a multiplicative decomposition of the interaction. This approach is commonly used in agricultural field studies under the name AMMI (additive main effects, multiplicative interaction) because, whether or not replication exists, there is a belief that this approach yields better predictions. In this type of analysis, 'residuals', formed after estimated main effects are subtracted from the data, are arrayed in a matrix with diets as columns and subjects as rows. A singular value decomposition of the matrix is performed and the first, or first and second, principal component(s) are used as estimates of the interaction, and can be tested for significance using approximate F-tests. Using the R gnm package to estimate parameters from these data sets with this methodology, we found large and significant subject by diet interaction effects in both data sets, and estimates of the interaction in the second data set were similar to interaction estimates from traditional ANOVA. Of an additional 26 dependent variables from the first and third data set (the latter involving moderate alcohol consumption), 19 had significant subject by diet interactions, based on the AMMI methodology. These results suggest that the subject by diet interaction is often important and should not be ignored when analyzing data obtaining from non-replicated crossover designs. The AMMI methodology appears to be a viable solution to estimate subject-diet interactions in non-replicated crossover designs.

Acknowledgements

References

Baer, D.J., Judd, J.T., Clevidence, B.A., Muesing, R.A., Campbell, W.S., Brown, E.D., Taylor, P.R. 2002. Moderate alcohol consumption lowers risk factors for cardiovascular disease in postmenopausal women fed a controlled diet. Am. J. Clin. Nutr. 75: 593-599.

Boykin, D., Camp, M.J., Johnson, L., Kramer, M., Meek, D., Palmquist, D., Vinyard, B., and West, M. 2011. Generalized linear mixed model estimation using Proc Glimmix: Results from simulations when the data and model match, and when the model is misspecified. *In* Proceedings of the 22nd Annual Conference on Applied Statistics in Agriculture (*ed*. Weixing Song), Kansas State University, April 2000. Pp. 137-170.

Chen, S.C., Judd, J.T., Kramer, M., Meijer, G.W., Clevidence, B.A., Baer, D.J. 2009. Phytosterol intake and dietary fat reduction are independent and additive in their ability to reduce plasma LDL cholesterol. Lipids. 44(3): 273-281.

Dias, C.T. dos S. and Krzanowski, W.J. 2003. Model selection and cross validation in additive main effect and multiplicative interaction models. Crop Science 43: 865-873.

Gauch, H.G. 1992. Statistical Analysis of Regional Yield Trials: AMMI Analysis of Factorial Designs. Elsevier, New York, New York. 278 pp.

Ghosh, S. and Crosby, H.R. 2005. Subject–treatment interactions in crossover trials: performance evaluation of subgrouping methods. J. Statistical Planning and Inference 132: 63–73.

Gollob, H. (1968). A statistical model which combines features of factor analytic and analysis of variance techniques. Psychometrika 33: 73-115.

Gogel, B.J., Cullis, B.R., and Verbyla, A.P. 1995. REML estimation of multiplicative effects in multi-environment variety trials. Biometrics 51: 744-749.

Hauck, W.W., Hyslop, T., Chen, M.-L., Patnaik, R., Williams, R.L., and the FDA Population Individual Bioequivalence Working Group. 2000. Subject-by-Formulation Interaction Bioequivalence: Conceptual and Statistical Issues. Pharmaceutical Research 17: 375-380.

Jones, B. and Kenward, M.G. 1989. Design and analysis of cross-over trials. Chapman and Hall. N.Y.

Mandel, J. 1971. A new analysis of variance model for non-additive data. Technometrics 13: 1-18.

Lee, E.J. and Johnson, D.E. 2006. AMMI macros for multiplicative interaction models. SUGI 31 Proceedings, San Francisco, California, March 26-29, 2006, Paper 049-31, 10 pp.

R Development Core Team. 2011. R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL http://www.R-project.org/

Rideout, T.C. 2011. Getting personal: considering variable interindividual responsiveness to dietary lipid-lowering therapies. Curr. Opin. Lipidol. 22: 37-42.

Tukey, J. 1949. One degree of freedom for non-additivity. Biometrics 5: 232-242.

Turner, H. and Firth, D. 2011. Generalized nonlinear models in R: An overview of the gnm package. R package version 1.0-1. http://CRAN.R-project.org/package=gnm.