Biostats VHM 801 Course
Fall 2020, Atlantic Veterinary College, PEI
Henrik Stryhn

Solution to Final Exam, December 2020

The solution is more detailed than expected, in particular by including solutions for almost all of the analysis specifications in Question 2.

Question 1

Denote by X_i and Y_i the antibody concentrations for subject (person) i before and 4 weeks after immunization, respectively, where $i = 1, \ldots, 20$. Furthermore, denote by $D_i = Y_i - X_i$ the difference (or gain, or increase) in antibody concentration from before to after immunization.

Subquestion a)

The design is a matched pairs (or two paired samples) design, with observations before and after immunization corresponding to a pair. Most naturally, the analysis would focus on the differences. The t-test from the paper would therefore most naturally be for the single sample of D_i's. The corresponding statistical model is that $D_1, \ldots, D_{20}$ are a sample (i.i.d.) from $\mathrm{N}(\mu_D, \sigma)$. A t-test for the hypothesis $\mathrm{H}_0 : \mu_D = 0$ is $t = \bar{D}/s_D$. The Minitab output (for the variable `post-pre`) shows this test to have a value of $t = 1.85$ and $P = 0.080$; this must be the reported result in the article (in fact, the value was rounded down correctly from $t = 1.8498$). Based on the information provided, it would seem that also a 2-sample t-test could have given this result; however its $t = 1.746$ should have been rounded to 1.7 with one decimal.

Subquestion b)

By the matched pairs design, the variable of principal interest is the differences (D_i). Descriptive analyses of the `pre` and `post` variables are of comparatively minor interest, because the analysis is based on the differences. A simple graphical display for the 17 values is a dotplot, shown in detail below but which can be sketched fairly quickly even without entering data into software. Naturally such a display is incomplete by not including the three masked values, but it should still give a fairly good impression of the distribution, and the descriptive statistics provided show that the three extra values are all inside the range from the 17 values shown.

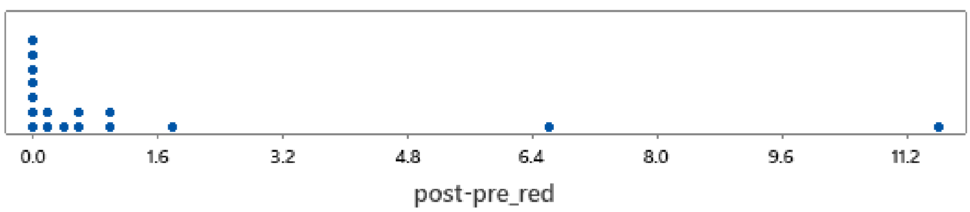

The dotplot for `post-pre` shows that assuming a normal distribution for the differences, as was done for the above t-test, is very questionable. The distribution looks quite strange, with a cluster of points around zero and two very large values. These values look "different" than the others but there is nothing to indicate them to be errors (two of the patients had very high antibody concentrations post immunization). Clearly, the distribution is right-skewed, which is also evident from the descriptive statistics (median $\ll$ mean, inter-quartiles far from symmetric around the median). Even without the largest one or two observations, there might remain considerable right-skewness in the distribution.

Subquestion c)

In view of the discussion above, alternatives to a normal distribution analysis are needed. Note that the robustness of the t-procedures would not cover such an extreme situation as the present. To transform the data (differences) seems difficult, with many values around zero and even one value less than zero (excluding log and square-root transformations). It might be an option to transform the original values and then form differences on transformed scale. Such analyses can not be carried out on the basis of the given Minitab output.

Therefore, the most natural approach might be a non-parametric analysis. Two possibilities exist for testing H_0: median = 0: the sign test and the Wilcoxon signed rank test, and they are both included in the Minitab listings. The Wilcoxon signed rank test assumes the distribution to be symmetric around its median, which seems little meaningful here because the asymmetry is the most striking deviation from a normal distribution. The sign test remains as the only valid option. Among the 20 differences, 8 of them are zero and are discarded for the sign test; among the 12 remaining differences, 11 of them are positive, which with a two-sided alternative hypothesis (H_a: median $\neq 0$) yields a P-value of $P = 0.0063$. Therefore, there is clear evidence that the antibody concentration increases after immunization. The estimated median increase is 0.1. It would be desirable to supplement this estimate with a confidence interval, but requires an additional Minitab analysis.

Subquestion d)

The following critique may be given of the quote from the paper:

- To use a t-test is not justified because the data clearly are not normally distributed, and the sample size is not large enough to compensate for such deviations from a normal distribution.
- The P-value and conclusion obtained is wrong, because there is clear evidence/significance against the distributions before and after immunization being the same.
- To give a P-value of 0.08 as ">0.05" violates the recommendation to state the exact P-value instead of its significance category.
- Moreover, the reporting disregards the fact that the P-value is in fact quite close to the 5% significance limit.
- A P-value without an estimate and a confidence interval tells nothing about the size and direction of the effect.
- Phrasing the conclusion as "no significant increase" gives the impression that a one-sided alternative was used, which was not the case.

Question 2

The results presented in the solution correspond to randomly selecting 500 observations using Minitab's `Sample from Columns` menu. The seed (Base) was set at 20201214. Before the random selection, the nine rows with a negative BMI-value (-1) were discarded, because these rows cannot contribute information to analyses for relationships with BMI. To keep these values as valid quantitative observations for BMI is totally wrong.

As a descriptive analysis of the outcome would be included in any of the specified analyses, we start out with this part. With the hint of a potential useful transformation of BMI to its inverse, we include a Graphical summary for this variable here as well.

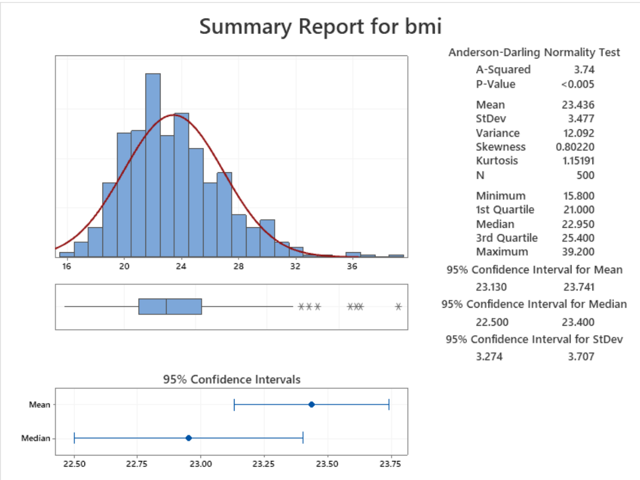

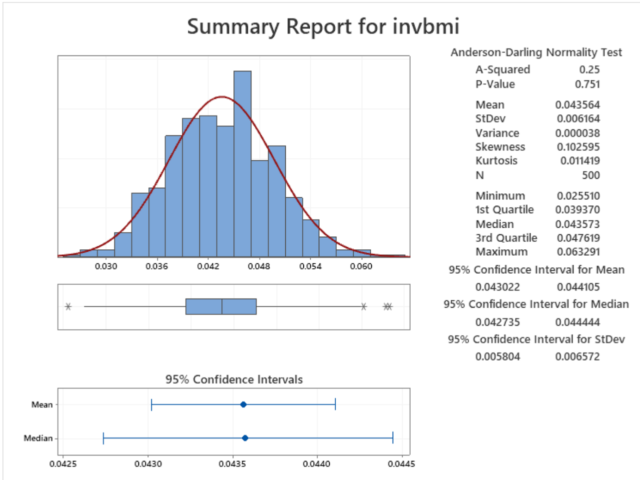

It is seen that the distribution of `bmi` is markedly right-skewed with a longer right tail (and consequently some "suspected outliers" that are hardly real outliers), whereas `invbmi` has a reasonable symmetric and bell-shaped distribution (as indicated by the high P-value for the A-D normality test). Although the assumptions of normality in the subsequent models will *never* be for the overall distribution, it seems clear that normality assumptions for subgroups or residuals might be easier to achieve on transformed scale. We will return to these assessments in the following. The descriptive statistics for the grouped BMI-variable are just the counts and proportions for the four categories (which could also be displayed in a bar or pie chart). Not too surprisingly, the second group is by far the largest.

Group	1	2	3	4	Total
count	20	382	125	23	500
proportion	0.040	0.664	0.250	0.046	1

We next go through the first eight analysis specifications; the multiple regression model of **i.** is considered beyond the main part of the course and was included only for completeness (and in case some students wanted the extra challenge). In order to avoid breaking the flow of the presentation, all Minitab output is deferred to an appendix.

a.: `bmi` versus age

The natural choice for describing the relationship between these two quantitative variables is simple linear regression. Correlation is an option as well, but offers no advantages over linear regression.

Objective: To estimate a linear relationship between age and BMI.

Model: $\texttt{bmi}_i = \beta_0 + \beta_1 \cdot \texttt{age}_i + \varepsilon_i$, where the errors ε_i are i.i.d. and $\mathrm{N}(0, \sigma)$.

Model checking: The residual plots at the two scales for BMI show that both the normality assumption (for the errors/residuals) and the assumption of equal variances are somewhat problematic for `bmi`, so it seems natural to carry out the analysis on inverse BMI scale. Note that the direction of the association is reversed on that scale, because the transformation reverses the order of the observations. Due to the large number of points in the plots, it is difficult to visually assess whether the relation is truly linear, and it is quite noisy anyway, with R^2-values as low as 13%. (Methods beyond VHM 801 will show that the relation between `invbmi` and age can be improved substantially by allowing for non-linearity.)

Estimation (inverse BMI scale): $\hat{\beta}_0 = 0.049$, $\hat{\beta}_1 = -0.000129$, $\hat{\sigma} = s = 0.00576$, and the 95% confidence intervals for intercept and slope:

$$\beta_0 : \ (0.0478, 0.0505), \quad \beta_1 : (-0.000159, -0.000100).$$

The confidence interval for the slope is very far from zero (note that the small parameter values are entirely due to the scales of age and `invbmi`). Note also that these estimates and CIs cannot be backtransformed to BMI scale.

Test: The relevant hypothesis is $\mathrm{H}_0 : \beta_1 = 0$, most naturally against a two-sided alternative, and the t-test gives $t = -8.60$ and $P < 0.001$, and thus strong significance against H_0, which therefore can be said to be strongly incompatible with the data.

Conclusion: The data give very clear evidence of a negative relationship between `invbmi` and age (and consequently a positive relationship between BMI and age); in other words, BMI seems to increase with age. The estimated linear relationship at inverse BMI scale corresponds to a non-linear relationship between BMI and age, which could be graphed based on the equation or the predicted values: $\widehat{\texttt{bmi}} = 1/(0.049 - 0.000129 \cdot \texttt{age})$.

b.: `bmi` versus gender

The natural choice for comparing BMI between women and men is a t-test for two independent samples. A non-parametric test seems less obvious because the robustness of t-distribution procedures combined with the large sample sizes should take care of any concerns about non-normality. A one-way ANOVA with two groups makes the unnecessary assumption that the two groups have equal standard deviations, and has the additional drawback that the robustness of F-statistics are not as well described as for t-statistics.

Objective: To compare the mean (or median) BMI between women and men.

Model: Among the men, `bmi`-values are i.i.d. and $\mathrm{N}(\mu_1, \sigma_1)$, and independently hereof the `bmi`-values among the women are i.i.d. and $\mathrm{N}(\mu_2, \sigma_2)$.

Model checking: The graphical summaries of BMI-values for men and women show that both distributions are right-skewed, with a single quite large extreme value (39.2) among the men, and several less extreme values in the tail of the distribution for women. Considering sample sizes of about 250 in both groups, t-procedures should be fine (as discussed above); the impact of the single extreme value on results is minimal (not shown).

Estimation: $\hat{\mu}_1 = 24.31$, $\hat{\sigma}_1 = 3.30$, $\hat{\mu}_2 = 22.61$, $\hat{\sigma}_2 = 3.45$, from which we get the difference:

$$\hat{\mu}_1 - \hat{\mu}_2 = 1.70, \quad 95\% \text{ CI for } \mu_1 - \mu_2: \ (1.11, 2.29).$$

The confidence interval for the difference is very far from zero, showing that the data indicate men to have a larger mean BMI than women.

Test: Equal means is expressed as $H_0 : \mu_1 = \mu_2$, most naturally against the two-sided alternative $H_a : \mu_1 \neq \mu_2$, and the t-test gives $t = 5.64$ and $P < 0.001$ — strong significance against H_0, which therefore can be said to be strongly incompatible with the data.

Conclusion: The data give very clear evidence that mean BMI is higher for men than women.

c.: `bmi` versus `energy`

As for **a.**, the natural choice for describing the relationship between these two quantitative variables is simple linear regression. The present analysis is summarized in more condensed form (refer to **a.** for extra details).

Objective: To estimate a linear relationship between energy (for convenience rescaled by dividing all values by 1000) and BMI.

Model: $\texttt{bmi}_i = \beta_0 + \beta_1 \cdot \texttt{energy}_i/1000 + \varepsilon_i$, where the errors ε_i are i.i.d. and $N(0, \sigma)$.

Model checking: The residual plots for the model at inverse BMI scale look better, for the same reasons as above. At both scales the predictive ability of the regression is minimal, with R^2-values almost zero ($R^2 = 0.2\%$ and 0.1%), and the large scatter of points almost eliminates the question of linearity.

Estimation (inverse BMI scale): $\hat{\beta}_0 = 0.0430$, $\hat{\beta}_1 = 0.000054$, $\hat{\sigma} = s = 0.00617$, and the 95% confidence intervals for intercept and slope:

$$\beta_0: \ (0.0415, 0.0445), \quad \beta_1: (-0.000072, 0.000182).$$

Test: The relevant hypothesis is $H_0 : \beta_1 = 0$, most naturally against a two-sided alternative, and the t-test gives $t = 0.83$ and $P = 0.41$, and thus no significance against H_0, which therefore can be said to be absolutely compatible with the data.

Conclusion: The estimated slope was positive, indicating a negative relationship between `bmi` and `energy`, but there was no evidence against a slope of zero, corresponding to no linear relation, and the confidence interval for the slope included a wide range of both negative and positive values. We conclude that the data fail to convincingly indicate a relationship between BMI and daily energy intake.

d.: `bmi` versus `urban`

The natural choice for comparing BMI among the three categories of `urban` (corresponding to independent samples) is a one-way ANOVA, possibly carried out non-parametrically (i.e., Kruskal-Wallis test), if concerns about normality assumptions cannot be resolved.

Objective: To compare the mean (or median) BMI among residents in the three urbanity groups.

Model: $\texttt{bmi}_{ij} = \mu_i + \varepsilon_{ij}$, where $i = 1, 2, 3 \sim$ group, $j = 1, \ldots, n_i$ and the errors ε_{ij} are i.i.d. and $\mathrm{N}(0, \sigma)$.

Model checking: Normal plots for the three groups of `urban` show curved relations corresponding to right-skewness in all three groups, all with evidence against normality by the A-D test. The three distributions have similar standard deviations (easily meeting the max/min rule), and may have sufficiently similar shapes for a non-parametric analysis focused on medians. As however the normal plots for `invbmi` are straight with non-significant normality tests for all three groups, it seems natural to conduct the analysis on inverse BMI scale.

Estimation (inverse BMI scale): $\hat{\mu}_1 = 0.0439$, $\hat{\mu}_2 = 0.0433$, $\hat{\mu}_3 = 0.0436$, $\hat{\sigma} = 0.00617$, with 95% confidence intervals centered around the group means with margin of errors ranging from 0.00083 to 0.00114, thereby making them all overlapping with the three estimates inside all intervals.

Test: Equal means is expressed as $\mathrm{H}_0 : \mu_1 = \mu_2 = \mu_3$, with the alternative H_a : not H_0, and the F-test of the ANOVA table equals $F(2, 497) = 0.46$ with $P = 0.63$, and thus no significance against H_0, which therefore can be said to be absolutely compatible with the data. As well, the predictive ability of the model is very low ($R^2 = 0.19\%$). Equal means on transformed scale corresponds to equal medians on original scale, because backtransformed means can be interpreted as medians.

Conclusion: We conclude that the data fail to convincingly indicate any differences in mean inverse BMI, and hence in median BMI, between the urbanity groups.

e.: `bmi` versus social group

As for **d.**, the natural choice for comparing BMI among the five social groups is a one-way ANOVA. The present analysis is summarized in more condensed form (refer to **d.** for extra details).

Objective: To compare the mean (or median) BMI among people from the five social groups.

Model: $\texttt{invbmi}_{ij} = \mu_i + \varepsilon_{ij}$, where $i = 1, 2, 3, 4, 5 \sim$ group, $j = 1, \ldots, n_i$ and the errors ε_{ij} are i.i.d. and $\mathrm{N}(0, \sigma)$.

Model checking: Normal plots for the five groups of `socgrp` show curved relations corresponding to right-skewness in all four of the five groups, all with evidence against normality by the A-D test. Analysis on inverse BMI scale seems most natural.

Estimation (inverse BMI scale): $\hat{\mu}_1 = 0.0450$, $\hat{\mu}_2 = 0.0441$, $\hat{\mu}_3 = 0.0422$, $\hat{\mu}_4 = 0.0438$, $\hat{\mu}_5 = 0.0434$, $\hat{\sigma} = 0.00615$, with 95% confidence intervals centered around the group means with margin of errors ranging from 0.00084 to 0.00193, thereby making them all overlapping, and with only the estimate for group 3 outside all other intervals.

Test: Equal means is expressed as $H_0 : \mu_1 = \ldots = \mu_5$, with the alternative H_a : not H_0, and the F-test of the ANOVA table equals $F(4, 495) = 1.73$ with $P = 0.14$, and thus no significance against H_0, which therefore can be said to be compatible with the data, although the P-value is in an intermediate range. Pairwise comparisons make little sense with a non-significant overall test (when adjusting for ($k = 10$) multiple comparisons, they will all be non-significant), so we may just note the lower `invbmi` mean (and hence higher `bmi` median) for group 3 as potentially of interest for further study.

Conclusion: We conclude that the data fail to convincingly indicate any differences in mean inverse BMI, and hence in median BMI, between the social groups, even if the median for social group 3 was observed to be a bit higher than the rest.

f.: `bmi` versus gender and social group

With two categorical explanatory variables (factors), the natural choice is a two-way ANOVA. We already noted the one-way ANOVA for social groups to be preferably carried out on inverse BMI scale, so this will be our default here as well.

Objective: To compare the mean (or median) BMI among women and men from the five social groups.

Model: $\texttt{invbmi}_{ijk} = \mu + \alpha_i + \beta_j + (\alpha\beta)_{ij} + \varepsilon_{ijk}$, where $i = 1, 2 \sim$ gender, $j = 1, \ldots, 5 \sim$ social group, $k = 1, \ldots, n_{ij}$ and the errors ε_{ijk} are i.i.d. and $N(0, \sigma)$.

Model checking: With group sizes for the combined gender × `socgrp` groups ranging from 4 to 127, model checking should be based on the residuals. The normal plot looks straight, and the groupwise standard deviations are all close ($s_{\max}/s_{\min} = 0.006837/0.004532 = 1.5 < 2$, not shown).

Estimation (inverse BMI scale): The the pooled standard deviation equals $\hat{\sigma} = 0.00615$; the choice of further estimates will be informed by the conclusions from the ANOVA table.

Test: Three hypotheses are tested in the ANOVA table: main effects of gender and social group, as well as their interaction. The latter is completely non-significant ($F(4, 490) = 0.49$, $P = 0.74$), giving no indication that differences among social groups would depend on gender, or vice versa. In the absence of a substantial interaction, we should turn to the two main effects. Because of the strong unbalancedness of the design, it is better to refit the model without the interaction (this part is really beyond what has been discussed in VHM 801). In the reduced model, we find gender to be strongly significant ($F(1, 494) = 46.4$, $P < 0.001$), and also the overall test for social groups to be clearly significant ($F(4, 494) = 3.62$, $P = 0.006$). We get estimated means and standard errors from the software (note: due to the unbalancedness, these are not simple means), and for 95% confidence intervals we use $t^* = 1.965$ from $t(494)$ (also obtained from software), and backtransform to BMI scale to give estimates and intervals for medians:

$$\begin{aligned}
\text{men:}\quad & \texttt{invbmi}: & 0.00421 \pm 0.00076, \quad & \texttt{bmi}: & 23.8\ (23.3, 24.2),\\
\text{women:}\quad & \texttt{invbmi}: & 0.00459 \pm 0.00087, \quad & \texttt{bmi}: & 21.8\ (21.4, 22.2),\\
\text{social group 1:}\quad & \texttt{invbmi}: & 0.00465 \pm 0.00190, \quad & \texttt{bmi}: & 21.5\ (20.7, 22.4),\\
\text{social group 2:}\quad & \texttt{invbmi}: & 0.00445 \pm 0.00155, \quad & \texttt{bmi}: & 22.5\ (21.7, 23.3),\\
\text{social group 3:}\quad & \texttt{invbmi}: & 0.00424 \pm 0.00132, \quad & \texttt{bmi}: & 23.6\ (22.9, 24.3),\\
\text{social group 4:}\quad & \texttt{invbmi}: & 0.00433 \pm 0.00082, \quad & \texttt{bmi}: & 23.1\ (22.7, 23.5),\\
\text{social group 5:}\quad & \texttt{invbmi}: & 0.00431 \pm 0.00104, \quad & \texttt{bmi}: & 23.2\ (22.6, 23.8).
\end{aligned}$$

By the F-test for gender, the data support higher BMI for men than women. Among the social groups, the estimate for group 1 is lowest and for group 3 is highest with groups 4-5 not much lower. Pairwise comparisons between social groups with Bonferroni adjustment for multiple testing show that social group 1 is significantly lower than groups $3 - 5$, whereas no differences between groups $2 - 5$ can be established.

Conclusion: We conclude that there is no evidence against simultaneous effects of gender and social groups being additive (on inverse BMI scale). Furthermore, the data give strong evidence of higher BMI for men than women, and also suggest that people in social group 1 on the average have higher BMI levels than in the lowest three social groups.

g.: `bmigrp` versus gender

With both outcome and explanatory variables categorical, the data can be summarized into a two-way table of counts, and therefore invites analysis by the Pearson X^2-test. Both variables were observed as response variables, so it is most natural to view the data as one sample from a population (with two response variables); in the course, it was termed a Model II.

Objective: To explore any dependence between gender and BMI groups.

Model: The counts N_{ij} of people in gender category i and BMI group j, where $i = 1, 2 \sim$ gender and $j = 1, 2, 3, 4 \sim$ BMI group, follow a multinomial distribution $(N, (p_{ij}))$ with $N = 500$ and the probabilities summing to 1, $\sum_{ij} p_{ij} = 1$.

Model checking: Nothing really to do here, if we believe the data correspond to a multinomial setting. The assumed independence and the homogeneity of the probability within groups could be invalidated, but apart from including extra relevant grouping variables in the model (thereby taking the analysis beyond VHM 801) there is nothing we can do about it.

Estimation: Considering the study objective, it seems most natural to focus on the conditional probabilities of BMI groups within each gender; these are displayed in the two-way cross-tabulation of the data (Appendix). Further discussion after the test.

Test: The hypothesis of interest is H_0 : independence between gender and BMI group, or when formulated in terms of the conditional probabilities, that the probabilities for BMI groups are the same for both genders. The Pearson X^2-test is strongly significant ($X^2(3) = 30.3$, $P < 0.001$), and the conditions for use of the χ^2-distribution are met (all expected counts above 5, the lowest being 9.72). The test gives gives evidence of dependence between gender and BMI group, and we find the largest the contributions to the X^2-value for BMI group 3, where $\hat{p}_{13} = 0.342$ and $\hat{p}_{23} = 0.163$, corresponding to an overrepresentation of men in BMI group 3. Confidence intervals for the estimated proportions can be added, e.g. classical/normal approximation 95% CIs for p_{13} : 0.342 ± 0.060 and for p_{23} : 0.163 ± 0.045. Conversely, women are overrepresented in BMI groups 1 and 2.

Conclusion: We conclude that the data give evidence of a dependence between gender and BMI group, reflecting once more that men have higher BMI values (and here are more strongly represented in the higher BMI groups). It is maybe also interesting to note that women are strongly overrepresented in the lowest BMI group.

h.: `bmigrp` versus social group

As for **g.**, the specification invites a two-way table analysis, based on a Model II, and a X^2-test. The present analysis is summarized in more condensed form (refer to **g.** for extra details).

Model: We use a similar multinomial distribution for the counts N_{ij} of people in social group i and BMI group j where $i = 1, \ldots, 5 \sim$ social group and $j = 1, 2, 3, 4 \sim$ BMI group.

Model checking: The assumed homogeneity in probabilities within groups contrasts the finding from **g.** that differences exist in the BMI group probabilities between genders. As before, it would however take us beyond VHM 801 to account for that.

Estimation: Considering the study objective, it seems most natural to focus on the conditional probabilities of BMI groups within each social group; these are displayed in the two-way cross-tabulation of the data (Appendix).

Test: The hypothesis of interest is H_0 : independence between social group and BMI group, or when formulated in terms of the conditional probabilities, that the probabilities for BMI groups are the same across all social groups. The Pearson X^2-test is totally non-significant ($X^2(12) = 10.9$, $P = 0.55$), but here the conditions for use of the χ^2-distribution are not met with 7 cells out of 20 ($\sim 35\%$) having expected counts below 5 (however none below 1). Because the concern is that cells with low expected counts give unduly high contributions to the X^2-statistic, the non-significance of the test effectively cancels the concern about not meeting the condition. The data seem to be compatible with independence between social group and BMI group.

Conclusion: We conclude that the data give no evidence of a dependence between social group and BMI group, when looking at these two variables on their own. As noted above (and in previous questions), this conclusion might change if gender was accounted for as well, so we have to state our conclusion as corresponding to only looking at these two variables.

Appendix: Minitab listings and graphs for the analyses in **a.-h.**

a.: `bmi` versus age

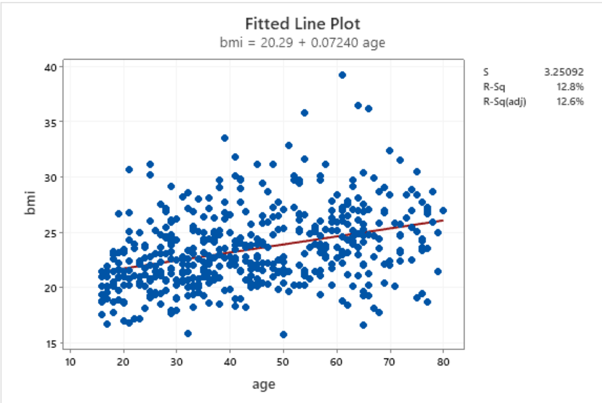

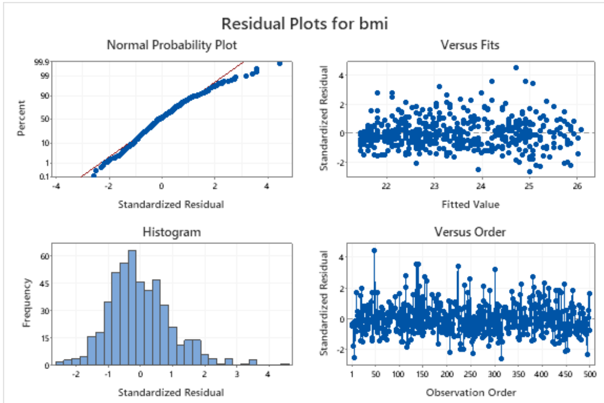

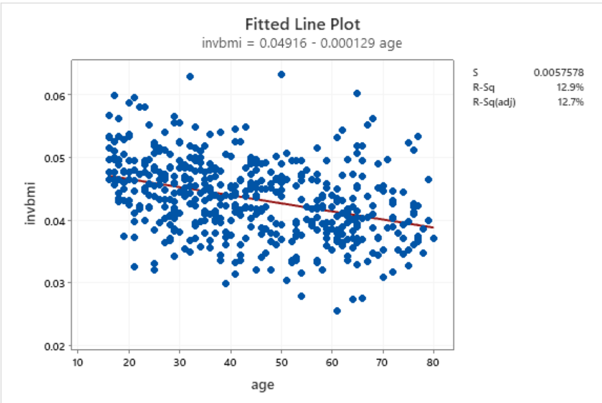

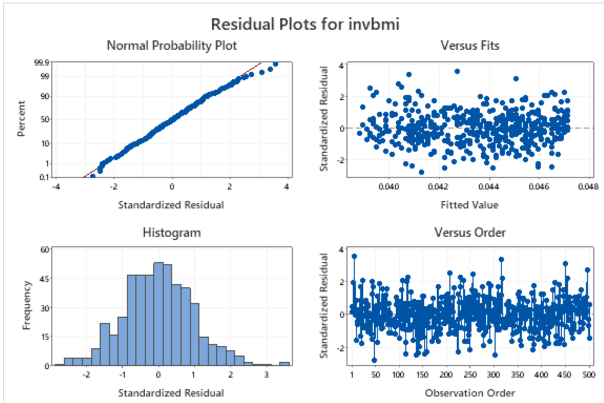

Coefficients

Term	Coef	SE Coef	95% CI	T-Value	P-Value	VIF
Constant	0.049162	0.000700	(0.047786, 0.050538)	70.20	0.000	
age	-0.000129	0.000015	(-0.000159, -0.000100)	-8.60	0.000	1.00

b.: `bmi` versus gender

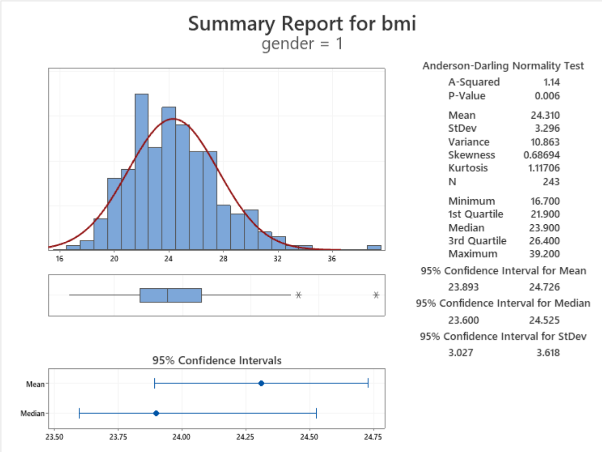

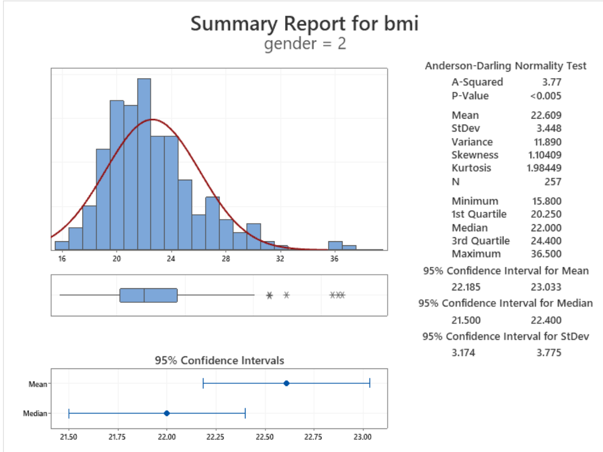

SUBSET OF EXAM.MTW

Two-Sample T-Test and CI: bmi, gender

Method

μ_1: population mean of bmi when gender = 1

μ_2: population mean of bmi when gender = 2

Difference: $\mu_1 - \mu_2$

Equal variances are not assumed for this analysis.

Descriptive Statistics: bmi

gender	N	Mean	StDev	SE Mean
1	243	24.31	3.30	0.21
2	257	22.61	3.45	0.22

Estimation for Difference

Difference	95% CI for Difference
1.701	(1.108, 2.294)

Test

Null hypothesis H_0: $\mu_1 - \mu_2 = 0$

Alternative hypothesis H_1: $\mu_1 - \mu_2 \neq 0$

T-Value	DF	P-Value
5.64	497	0.000

c.: `bmi` versus `energy`

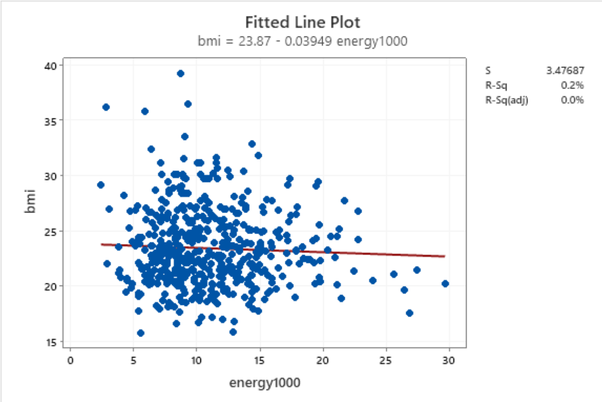

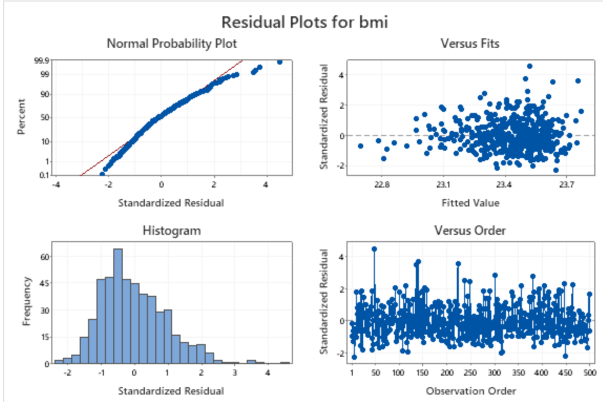

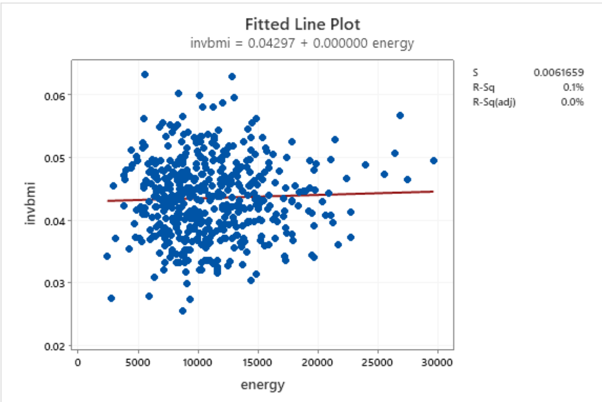

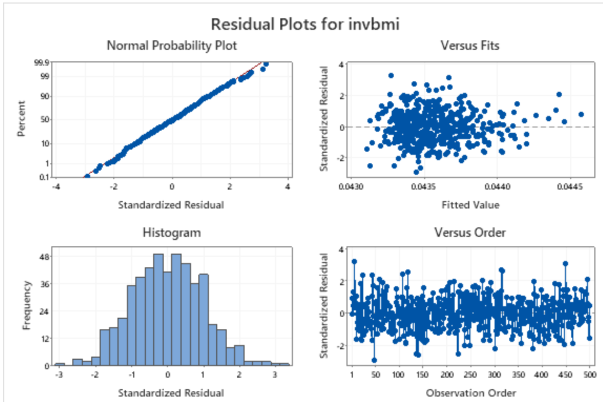

Coefficients

Term	Coef	SE Coef	95% CI	T-Value	P-Value	VIF
Constant	0.042972	0.000764	(0.041472, 0.044472)	56.28	0.000	
energy1000	0.000054	0.000065	(-0.000074, 0.000182)	0.83	0.406	1.00

d.: bmi versus urban

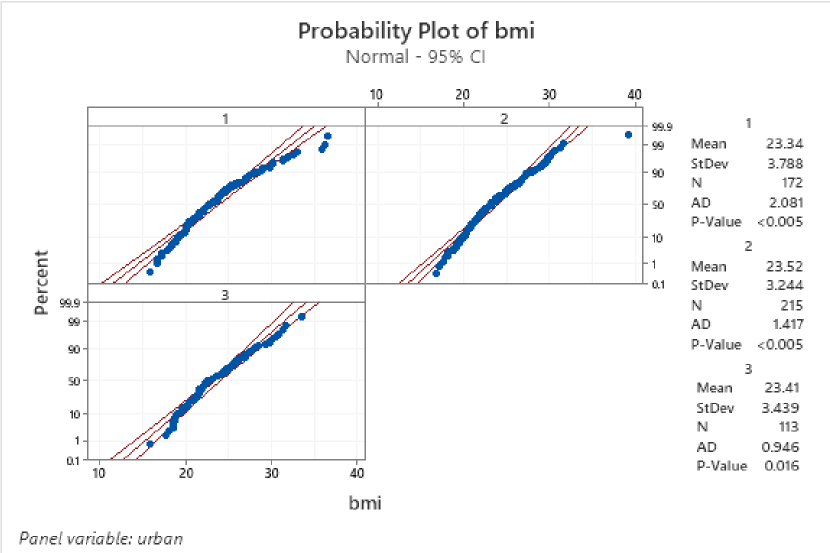

SUBSET OF EXAM.MTW

One-way ANOVA: invbmi versus urban

Method

Null hypothesis	All means are equal
Alternative hypothesis	Not all means are equal
Significance level	α = 0.05

Equal variances were assumed for the analysis.

Factor Information

Factor	Levels	Values
urban	3	1, 2, 3

Analysis of Variance

Source	DF	Adj SS	Adj MS	F-Value	P-Value
urban	2	0.000035	0.000018	0.46	0.631
Error	497	0.018924	0.000038		
Total	499	0.018959			

Model Summary

S	R-sq	R-sq(adj)	R-sq(pred)
0.0061706	0.19%	0.00%	0.00%

Means

urban	N	Mean	StDev	95% CI
1	172	0.043888	0.006653	(0.042963, 0.044812)
2	215	0.043283	0.005715	(0.042456, 0.044110)
3	113	0.043604	0.006244	(0.042464, 0.044745)

Pooled StDev = 0.00617064

e.: `bmi` versus social group

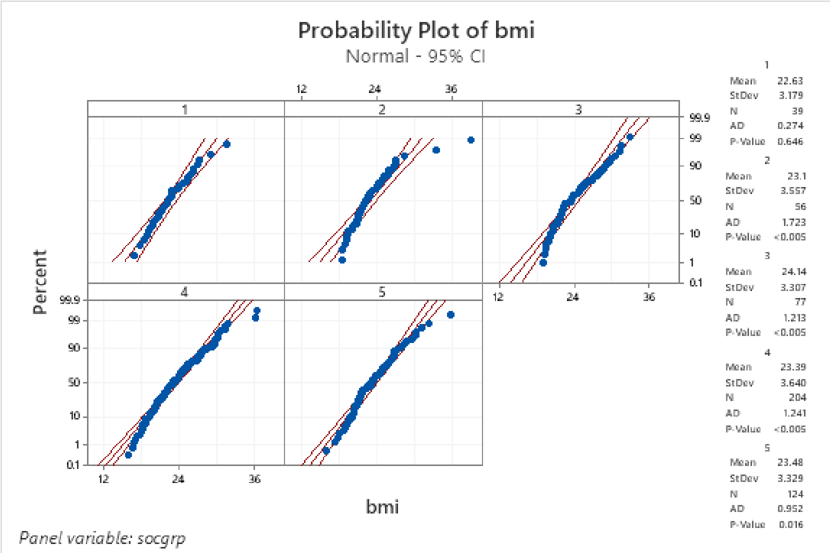

SUBSET OF EXAM.MTW

One-way ANOVA: invbmi versus socgrp

Method

Null hypothesis	All means are equal
Alternative hypothesis	Not all means are equal
Significance level	α = 0.05

Equal variances were assumed for the analysis.

Factor Information

Factor	Levels	Values
socgrp	5	1, 2, 3, 4, 5

Analysis of Variance

Source	DF	Adj SS	Adj MS	F-Value	P-Value
socgrp	4	0.000262	0.000066	1.73	0.141
Error	495	0.018697	0.000038		
Total	499	0.018959			

Model Summary

S	R-sq	R-sq(adj)	R-sq(pred)
0.0061459	1.38%	0.59%	0.00%

Means

socgrp	N	Mean	StDev	95% CI
1	39	0.045023	0.006172	(0.043090, 0.046957)
2	56	0.044122	0.005677	(0.042508, 0.045735)
3	77	0.042162	0.005473	(0.040786, 0.043538)
4	204	0.043753	0.006593	(0.042907, 0.044598)
5	124	0.043412	0.005966	(0.042328, 0.044496)

Pooled StDev = 0.00614592

f.: `bmi` versus gender and social group

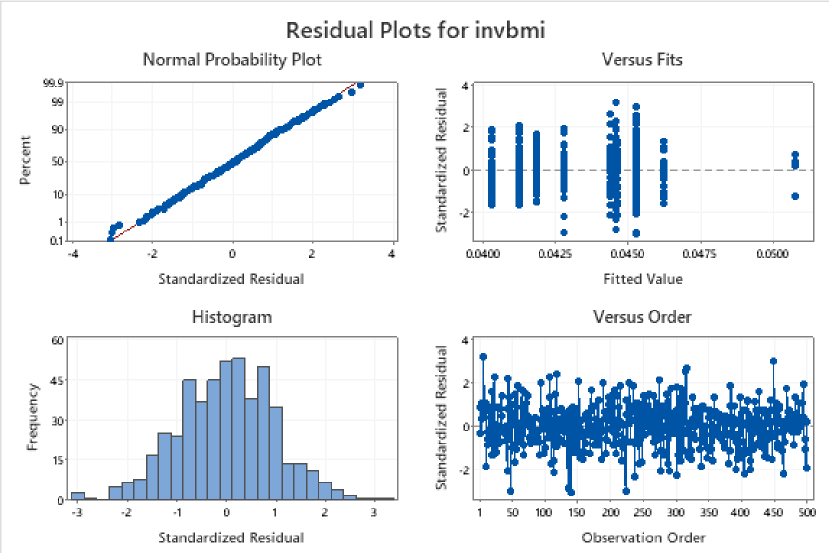

Analysis of Variance

Source	DF	Adj SS	Adj MS	F-Value	P-Value
gender	1	0.000957	0.000957	27.54	0.000
socgrp	4	0.000386	0.000096	2.78	0.027
gender*socgrp	4	0.000068	0.000017	0.49	0.744
Error	490	0.017025	0.000035		
Total	499	0.018959			

Analysis of Variance

Source	DF	Adj SS	Adj MS	F-Value	P-Value
gender	1	0.001604	0.001604	46.37	0.000
socgrp	4	0.000501	0.000125	3.62	0.006
Error	494	0.017093	0.000035		
Lack-of-Fit	4	0.000068	0.000017	0.49	0.744
Pure Error	490	0.017025	0.000035		
Total	499	0.018959			

Means

Term	Fitted Mean	SE Mean
gender		
1	0.042105	0.000388
2	0.045855	0.000444
socgrp		
1	0.046514	0.000967
2	0.044523	0.000788
3	0.042430	0.000671
4	0.043293	0.000417
5	0.043140	0.000530

Bonferroni Simultaneous Tests for Differences of Means

Difference of socgrp Levels	Difference of Means	SE of Difference	Simultaneous 95% CI	T-Value	Adjusted P-Value
2 - 1	-0.00199	0.00124	(-0.00548, 0.00150)	-1.61	1.000
3 - 1	-0.00408	0.00117	(-0.00738, -0.00079)	-3.49	0.005
4 - 1	-0.00322	0.00107	(-0.00623, -0.00021)	-3.02	0.027
5 - 1	-0.00337	0.00111	(-0.00651, -0.00024)	-3.04	0.025
3 - 2	-0.00209	0.00103	(-0.00501, 0.00082)	-2.03	0.433
4 - 2	-0.001230	0.000896	(-0.003758, 0.001297)	-1.37	1.000
5 - 2	-0.001384	0.000952	(-0.004069, 0.001301)	-1.45	1.000
4 - 3	0.000864	0.000794	(-0.001375, 0.003102)	1.09	1.000
5 - 3	0.000710	0.000857	(-0.001707, 0.003127)	0.83	1.000
5 - 4	-0.000153	0.000670	(-0.002044, 0.001737)	-0.23	1.000

Individual confidence level = 99.50%

g.: bmigrp versus gender

SUBSET OF EXAM.MTW

Tabulated Statistics: gender, bmigrp

Rows: gender   Columns: bmigrp

	1	2	3	4	All
1	3	143	83	14	243
	1.23	58.85	34.16	5.76	100.00
	9.72	161.35	60.75	11.18	
	4.646	2.087	8.149	0.712	
2	17	189	42	9	257
	6.61	73.54	16.34	3.50	100.00
	10.28	170.65	64.25	11.82	
	4.393	1.974	7.705	0.674	
All	20	332	125	23	500
	4.00	66.40	25.00	4.60	100.00

Cell Contents
Count
% of Row
Expected count
Contribution to Chi-square

Chi-Square Test

	Chi-Square	DF	P-Value
Pearson	30.340	3	0.000
Likelihood Ratio	31.615	3	0.000

h.: `bmigrp` versus social group

SUBSET OF EXAM.MTW

Tabulated Statistics: socgrp, bmigrp

Rows: socgrp Columns: bmigrp

	1	2	3	4	All
1	3	26	9	1	39
	7.69	66.67	23.08	2.56	100.00
	1.560	25.896	9.750	1.794	
	1.3292	0.0004	0.0577	0.3514	
2	1	43	10	2	56
	1.79	76.79	17.86	3.57	100.00
	2.240	37.184	14.000	2.576	
	0.6864	0.9097	1.1429	0.1288	
3	0	50	23	4	77
	0.00	64.94	29.87	5.19	100.00
	3.080	51.128	19.250	3.542	
	3.0800	0.0249	0.7305	0.0592	
4	12	132	50	10	204
	5.88	64.71	24.51	4.90	100.00
	8.160	135.456	51.000	9.384	
	1.8071	0.0882	0.0196	0.0404	
5	4	81	33	6	124
	3.23	65.32	26.61	4.84	100.00
	4.960	82.336	31.000	5.704	
	0.1858	0.0217	0.1290	0.0154	
All	20	332	125	23	500
	4.00	66.40	25.00	4.60	100.00

Cell Contents
Count
% of Row
Expected count
Contribution to Chi-square

Chi-Square Test

	Chi-Square	DF	P-Value
Pearson	10.808	12	0.545
Likelihood Ratio	13.681	12	0.322

7 cell(s) with expected counts less than 5.