Elizabeth Mosher

Department of Mathematical Sciences
Montana State University

April 24, 2014

A writing project submitted in partial fulfillment
of the requirements for the degree

Master of Science in Statistics

APPROVAL

of a writing project submitted by

Elizabeth Mosher

This writing project has been read by the writing project advisor and has been found to be satisfactory regarding content, English usage, format, citations, bibliographic style, and consistency, and is ready for submission to the Statistics Faculty.

4/24/2014
Date

Dr. John Steve Cherry
Writing Project Advisor

4/24/2014
Date

Dr. Megan D. Higgs
Writing Projects Coordinator

Bokashi Compost Experiment Analysis

Elizabeth Mosher

April 24th 2014

Contents

1 Background

The Bokashi method of decomposition is a widely used method in which some mixture containing microorganisms is added to food waste to assist in its decomposition in an anaerobic (or oxygen-limited) environment. After the food has had time to decompose, usually within some air-tight container, it is transferred into the ground where it will finish the decomposition process. The final product is a nutrient-rich soil which will be used for growing new plants. Inoculants may be added to the decomposing food waste to assist in this process. Common inoculants include both a mixture of effective microorganisms (EM) and soil which contains microorganisms vital to the decomposition process.

A group of undergraduate students in the Land Resources and Environmental Sciences department at Montana State University conducted a four month long experiment consisting of three phases. Phase I allowed for food waste to decompose within an oxygen-limited environment (buckets) under varying conditions. Phase II began when the food waste within these buckets was transferred to a field where it was buried and allowed to decompose in the ground (separated out by the same treatment conditions). In Phase II of the experiment the weather was not conducive to effective decomposition, and as a result the data collected were not informative. Therefore, Phase II analyses will not be discussed. Finally, in Phase III the decomposed food was dug up and used to grow lettuce plants under varying conditions. The main objective is to assess evidence of a difference in decomposition characteristics among the three treatment groups; compost inoculated with EM's, compost inoculated with soil, and compost not inoculated.

2 Experimental Setup

2.1 Phase I

In Phase I, tons of food left over from two cafeterias at Montana State University was obtained and mixed together as thoroughly as possible. The researchers packed 26 buckets with 12 inches of food, while at the same time adding assigned treatments for each bucket. The buckets were assigned either the EM inoculum, the soil inoculum, or no inoculum, and then assigned a specific food-packing setup (Table 1). There were six control buckets that did not contain any food waste. Three of these buckets contained only the EM inoculum and the other three buckets contained only the soil inoculum.

Amount of food waste	Bokashi Bran inoculum (EM's)	Soil Inoculum	No Inoculum
3 inch layers (x4 per bucket)	A1 (5 replicates)	A2 (5 replicates)	C1 (6 replicates)
12 inch layer (x1 per bucket)	B1 (5 replicates)	B2 (5 replicates)	
No Food	C2 (3 replicates)	C3 (3 replicates)	

Table 1. Different treatments separated by the inoculum added to the compost (columns), and the method of adding the inoculum (rows).

The treatment combinations were:

- A1 and A2: Ten buckets of A1 and A2 (five each) consisted of a lasagna of 3 inches of food, followed by a layer of the inoculum (EM's or soil), repeated 4 times to get a total of 12 inches of food.
- B1 and B2: Ten buckets of B1 and B2 (five each) consisted of 12 inches of food and a top layer of the inoculum (EM's or soil).
- C1: Six buckets consisted of 12 inches of food with no inoculum.
- C2: Three buckets consisted of only one layer of EM's.
- C3: Three buckets consisted of only a layer of soil.

Each bucket that contained an inoculum (EM's or soil) contained the same amount of that inoculum. For example, a bucket containing the B1 treatment had a set amount of EM's which was sprinkled on top of the food waste and a bucket containing treatment A1 sprinkled that same amount of EM's split over the four layers. Also, the amount of EM's sprinkled on the food waste was the same as the amount of soil sprinkled on the food waste.

After these buckets were filled with the different treatments, they were put into a second, larger bucket. This larger bucket was sealed off except for a small hole in the lid of the bucket that was covered by tape and a hole in the bottom of the bucket which was plugged by a rubber stopper. The tape was only removed for a brief moment every day when inserting a syringe to gather air from the head-space of the bucket. The rubber stopper was never removed.

The response variable in Phase I is the amount of CO_2 in ppm emitted into the head-space of the bucket, and the researchers are interested in assessing evidence for a difference in mean CO_2 among treatments over the 30 day period of time. The researchers are interested in comparing treatments A1 to A2, B1 to B2, A1 to C1, A2 to C1, B1 to C1, and B2 to C1.

During a 30 day period of time, 28 measurements of CO_2 were taken on the 32 buckets at around the same time on each day. The CO_2 was taken from the head space of the bucket by filling a syringe with air (an equal amount of air for each bucket) and using a machine (Licor Gas Analyzer) to calculate peak CO_2 in ppm in the sample of air.

2.2 Phase II

After the 29 days of lab incubation (Phase I) Phase II began. The researchers went to the Townes Harvest Garden field and dug 26 equal sized holes in which they placed a pipe whose purpose was to separate the soil and compost inside the pipe from soil outside the pipe. There was one pipe for each of the 26 buckets containing treatments A1, A2, B1, B2, and C1.

The soil that was dug up to create the holes was thoroughly mixed and the pipes were then filled with an equal layer of soil, compost from a treatment, and a top layer of soil (a "sandwich" of soil from the ground and compost from the buckets). Note that there was no bottom to these pipes so that seepage from the compost was able to seep into the ground below the pipe and bugs and worms were able to crawl up through the ground into the material within the pipe. A sixth treatment was added to the study at this time which consisted only of soil as a control. This control was placed in four pipes, bringing the total number of pipes to 30. The pipes were randomly assigned a treatment bucket using a randomization plot in Excel.

The response variable in Phase II is the amount of CO_2 in ppm emitted into the head-space of the pipe. The researchers are interested in assessing evidence for a difference in mean CO_2 among treatments over the one month period of time. They are interested in comparing treatments A1 to A2, B1 to B2, A1 to C1, A2 to C1, B1 to C1, and B2 to C1.

For one month the compost was allowed to decompose in the ground and CO_2 measurements were collected from the head-space of 22 of the pipes ten times within 36 days. These ten measurements were not equally spaced in time. On the remaining four pipes there was a machine that measured CO_2 in the head-space of the pipes once every hour for two weeks.

2.3 Phase III

After one month, the soil and compost were dug up and the compost was separated out from the soil as best as possible. Unfortunately there was a lot of rain while the food was composting in the ground so there was much less decomposition than expected. Because of this, the compost from each tube was laid out to dry separately and then ground up and consolidated for each treatment (A1, A2, B1, B2, and C1).

A pile of Sunshine brand soil was selected to be used along with the dried, ground up compost to grow lettuce plants. Two samples of 150 grams of the consolidated, well-mixed compost were arbitrarily chosen from each treatment (ten total samples). For the two samples from each treatment, one sample of compost was thoroughly mixed with 1350 grams of non-autoclaved Sunshine brand soil, while the other sample of 150 grams of compost was thoroughly mixed with 1350 grams of the autoclaved Sunshine brand soil. This was done for all 5 treatments (A1, A2, B1, B2, and C1).

Autoclaving is a process in which the soil is heated to a very hot temperature using super-heated steam in order to kill all micro-organisms. The autoclaved soil was used as a method of trying to control for any microorganisms living within the soil, and their effect on plant growth. A sixth treatment was added to the study at this time consisting of only the Sunshine soil with no inoculum for both the autoclaved and non-autoclaved soil (Table 2).

		Original Treatments					
		A1	A2	B1	B2	C1	Soil (S)
Type of Sunshine Soil Used	Non-Autoclaved	1	2	3	4	5	6
	Autoclaved	1A	2A	3A	4A	5A	6A

Table 2. Different treatments separated by the inoculum added to the soil in the pots (columns), and the soil used, either autoclaved or non-autoclaved soil (rows) to create the 12 mixtures.

Lettuce plants were grown in a greenhouse. Twelve rows (blocks) were created on two benches (West Bench and East Bench), containing six rows (blocks) each. The blocks were constructed under the assumption that lettuce plants within a row will receive an equal amount of sunlight, and the lettuce plants between rows will receive different amounts of sunlight. In other words, the growing conditions within a row will be less variable then the conditions between rows. Each row contained 12 pots and each of the 12 mixtures was randomly assigned to a pot within each row (there was no replication of mixtures within blocks) (Figure 1).

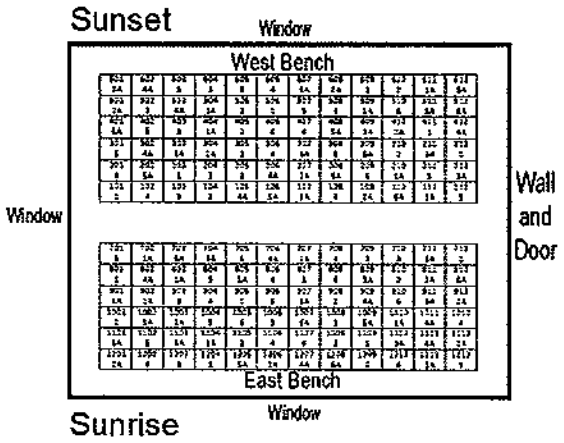

Figure 1. Layout of the greenhouse in which pots were separated into 12 blocks (rows) on two benches (6 blocks per bench), as well as the location of the windows and behavior of the sun during the day.

The response variables in Phase III are the number of seeds germinated, the heights of the lettuce plants over time and the biomass of the lettuce plant at the end of the 52 days. The researchers are interested in assessing evidence of a difference in mean dry weight (biomass) among the six treatment groups (A1, A2, B1, B2, C1, and Soil) and if this differs between autoclaved and non-autoclaved soil. The comparisons of interest are A1 to A2, B1 to B2, A1 to C1, A2 to C1, B1 to C1, and B2 to C1 as well as A1 to S (soil with no additive), A2 to S, B1 to S, B2 to S, and C1 to S.

Four lettuce seeds were planted in each pot and the number of germinated seeds was recorded 6 times over 26 days. There were some pots that did not germinate any seeds and this was recorded as a 0 for all three response variables.

After 26 days, the tallest of the germinated lettuce plants was left in the pot and all other germinated lettuce plants were removed in order to allow the biggest lettuce plant to grow without having to compete for soil nutrients, water, etc. The biggest lettuce plant's height was measured from the base of the lettuce plant (at the soil) to the tip of it's longest leaf. The measurements were taken 12 times over 35 days (equally spaced in time), including its final height at the end of a total of 52 days. After 52 days the lettuce plants were taken out of the pots and immediately weighed to record its wet weight. They were then dried and weighed again to get the dry weight. The part of the lettuce plant that was weighed for both the wet and dry weight was the above-ground biomass of the lettuce plants.

3 Preliminary Plots

3.1 Phase I: CO_2

CO_2 measurements for the five treatments A1, A2, B1, B2 and C1 for each bucket were plotted over time (Figure 2).

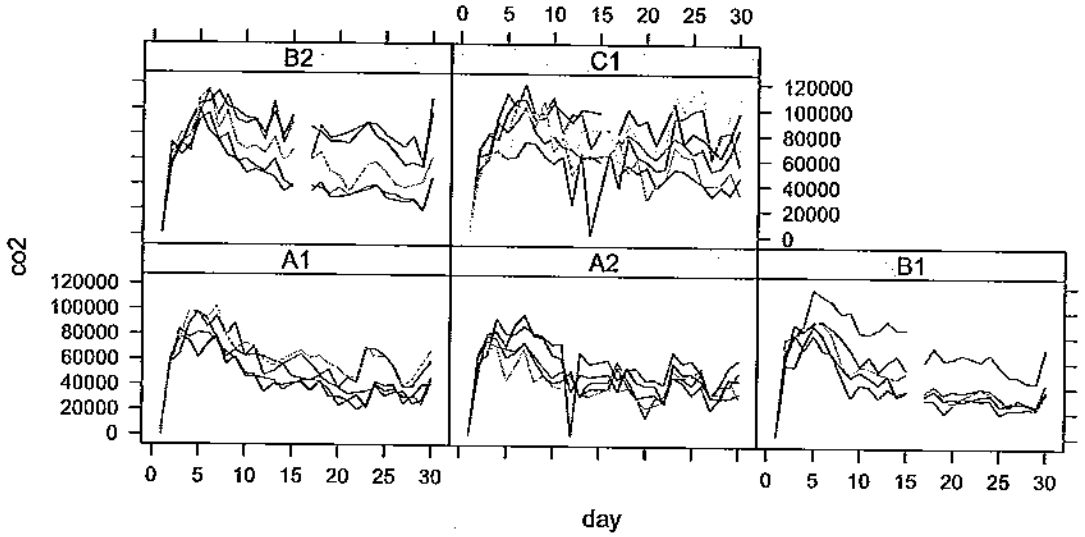

Figure 2. Change of CO_2 over 30 days for five treatments (A1, A2, B1, B2 and C1) for each bucket within each treatment.

Missing values were noted (Figure 2). The researchers explained that the missing values in the data set were all from the same day and this was because no measurements were taken on that day. We will assume that these values are missing completely at random.

There were also a few CO_2 measurements that appeared to be possible outliers. These were judged to be invalid based on the knowledge of the researchers and were removed. (Table 3).

Index	Treatment	Bucket #	Date	Day	Reason for Exclusion
1	A1	2	4/26/13	14	Outlier
2	A2	1	4/24/13	12	Outlier
3	A2	2	5/2/13	20	Outlier
4	B1	1	4/21/13	16	Missing
5	B1	2	4/21/13	16	Missing
6	B1	3	4/21/13	16	Missing
7	B1	4	4/21/13	16	Missing
8	B1	5	4/21/13	16	Missing
9	B2	1	4/21/13	16	Missing
10	B2	2	4/21/13	16	Missing
11	B2	3	4/21/13	16	Missing
12	B2	4	4/21/13	16	Missing
13	B2	5	4/21/13	16	Missing
14	C1	1	4/21/13	16	Missing
15	C1	2	4/21/13	16	Missing
16	C1	3	4/21/13	16	Missing
17	C1	4	4/24/13	12	Outlier
18	C1	4	4/26/13	14	Outlier
19	C1	4	4/29/13	17	Outlier
20	C1	5	5/2/13	20	Outlier

Table 3. All excluded CO_2 measurements.

Seven outliers were removed and the data were re-plotted (Figure 3).

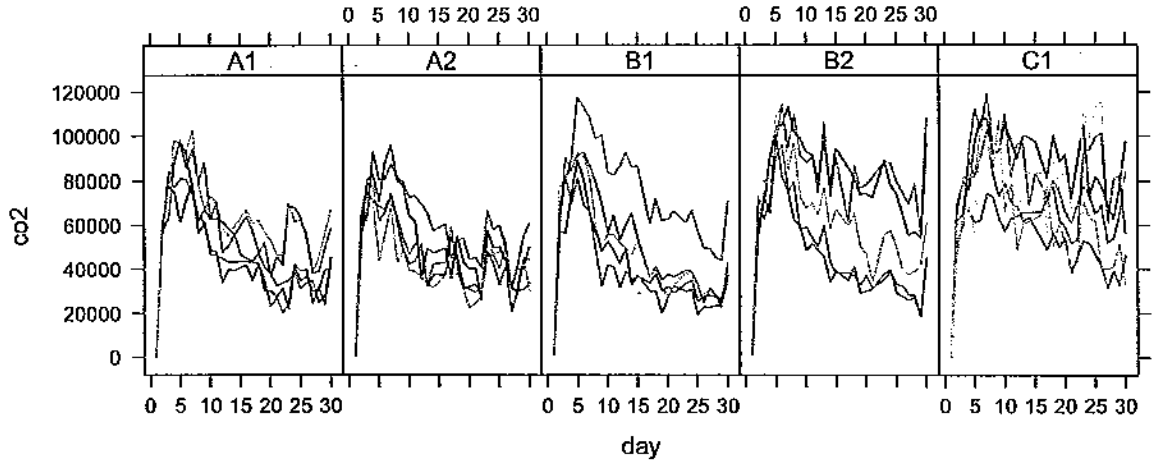

Figure 3. Change of CO_2 over 30 days for five treatments (A1, A2, B1, B2 and C1) for each bucket within each treatment excluding seven outliers.

There was an increase in CO_2 output during the first five days of decomposition followed by a gradual decline in CO_2 output (Figure 3). The researchers believed this increase in CO_2 output was either caused by initial oxygen in the bucket being used for aerobic respiration (which results in more CO_2 output than anaerobic respiration) or to initial materials in the food waste being broken down by the microorganisms. They ultimately decided that it would be of interest to keep the first five days in the analysis. Based on this decision the two periods of time were separated and analyzed separately.

3.1.1 Phase I-A: First 5 Days of Decomposition

CO_2 measurements for the first five days were plotted (Figure 4).

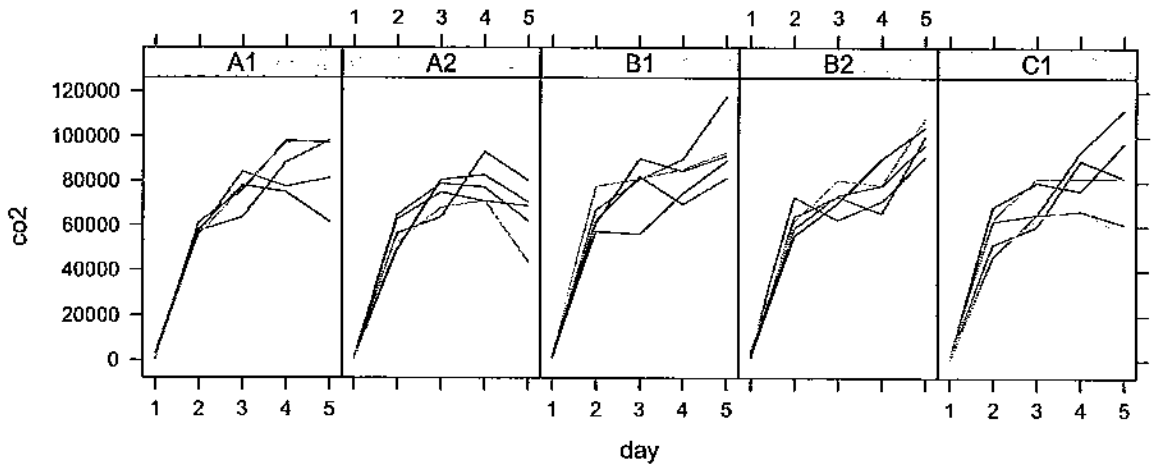

Figure 4. Change of CO_2 over the first 5 days for five treatments (A1, A2, B1, B2 and C1) for each bucket within each treatment.

There appears to be a quadratic trend in CO_2 measurements over the first five days, but the CO_2 measurement for Day 1 (the day the food was added to the buckets) was very small and then on Day 2 the CO_2 measurements drastically increased due to decomposition (Figure 4). Day 1 CO_2 measurements were not affected by the treatments and were removed. After removing data from Day 1, CO_2 measurements appear to increase linearly until Day 5 except for treatment A2 which appears to have peaked a day earlier than the others. However, there could be a slight quadratic trend over the four day period of time. There is also large variability between buckets within each treatment (Figure 5).

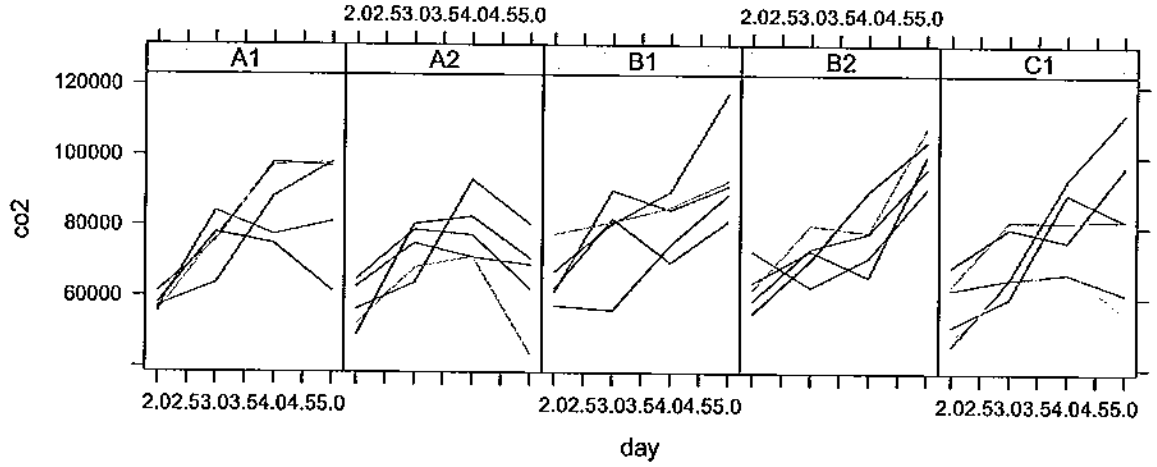

Figure 5. Change of CO_2 over Day 2 to Day 5 for five treatments (A1, A2, B1, B2 and C1) for each bucket within each treatment.

3.1.2 Phase I-B: Last 25 Days of Decomposition

CO_2 measurements for the last 25 days were plotted (Figure 6).

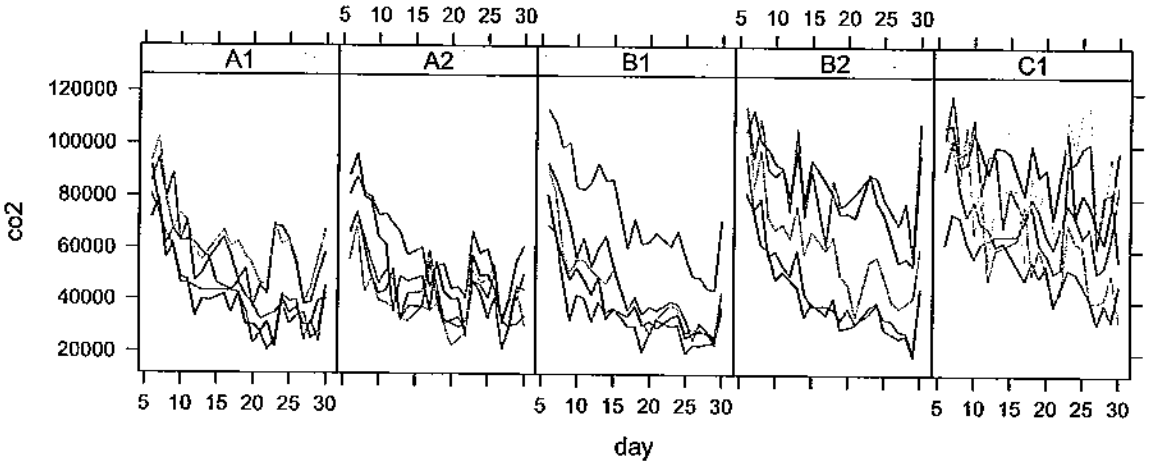

Figure 6. Change of CO_2 over the last 25 days for five treatments (A1, A2, B1, B2 and C1) for each bucket within each treatment.

There appears to be a slight quadratic trend of the CO_2 measurements over the last 25 days of Phase I. There also appears to be a jump in CO_2 measurements on Day 23 and Day 30 (Figure 6). The researchers believed that there was an assignable cause for the large jump in the CO_2 measurements on Day 23, but did not mention anything regarding Day 30. It is also clear that there is a lot of variability between buckets within each treatment.

3.2 Phase III: Biomass

3.2.1 Zeros in the Data

A bar chart and contingency table were created to help visualize the number of zero's in the data set for biomass (Figure 7). These zero's represent an empty pot in which lettuce plants had failed to germinate and grow.

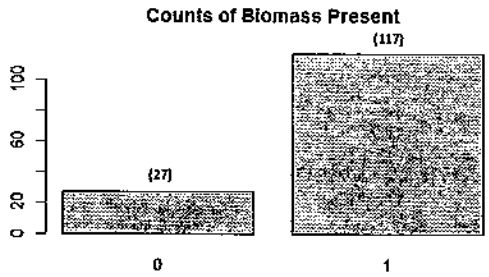

Figure 7. Counts of biomass recorded where 0 = no lettuce plant present and 1 = lettuce plant present and biomass recorded.

Two separate analyses were conducted. A model of the probability of germination was fit and a separate analysis of biomass was carried out for those plants that germinated.

3.2.2 Biomass Without Zeros in the Data

Biomass was plotted by treatment separated by autoclaving (Figure 8).

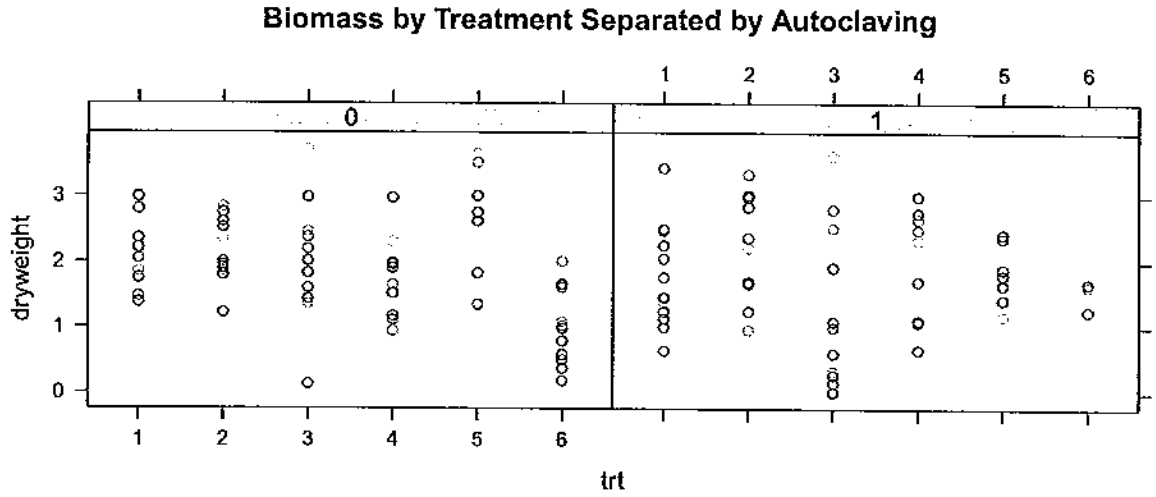

Figure 8. Biomass measurements for six treatments, separated out by whether or not the soil used was autoclaved (0 = non-autoclaved and 1 = autoclaved).

The treatment axis on the plot above is coded as follows:

- Treatment 1 = A1
- Treatment 2 = A2
- Treatment 3 = B1
- Treatment 4 = B2
- Treatment 5 = C1
- Treatment 6 = S

It is not clear that there are any treatment differences or any differences between plants grown in autoclaved or non-autoclaved soil. However, it appears that the variability is fairly constant (Figure 8).

4 Analysis and Model Selection

4.1 Phase I: CO_2

4.1.1 Phase I-A: First 5 Days of Decomposition

Explanatory variables considered are:

- *Treatment*
- *Cday*: Centered day, to account for a change in CO_2 measurements over the 25 day period of time.
- *Treatment* $*$ *Cday*: The interaction term between treatment and centered day to allow different linear trends of CO_2 over time depending on the treatment.
- $Cday^2$: A quadratic term for centered day to allow for a quadratic trend of CO_2 over the 25 day period of time.
- *Treatment*$*Cday^2$: The interaction term between treatment and $cday^2$ to allow different quadratic trends of CO_2 over time depending on the treatment.

Day^2 was included due to hints of a quadratic relationship. *Day* was centered to control for collinearity with Day^2. The following three models were fit in R and compared:
Model 1: $CO_2 = trt + cday + trt * cday + cday^2 + trt * cday^2$
Model 2: $CO_2 = trt + cday + trt * cday + cday^2$
Model 3: $CO_2 = trt + cday + trt * cday$

Model	R^2	AIC
Model 1	0.5835	2238.397
Model 2	0.5241	2244.262
Model 3	0.5032	2246.751

Table 4. R^2 and AIC values from the three fitted models.

Model 1 was chosen based on the smaller AIC value (Table 4). This model yielded an $R^2 = 0.5835$, meaning that this model explains about 58% of the variability of the data. Diagnostic plots revealed no major concerns and were judged acceptable (Figure 9).

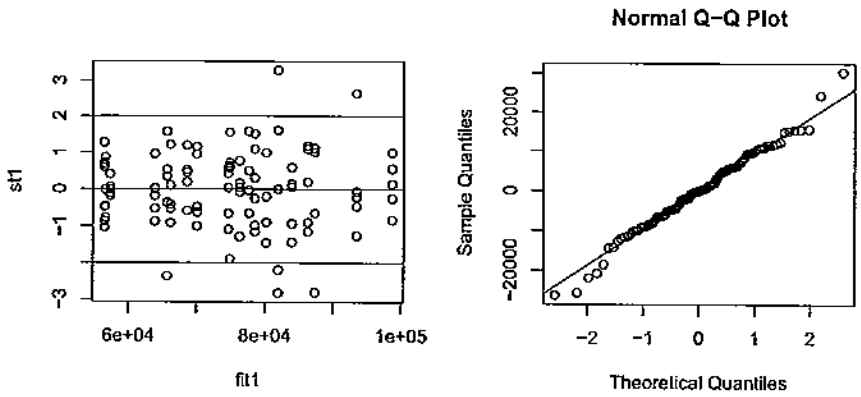

Figure 9. Studentized Residuals vs. Fitted Values plot (left) and Normal Q-Q plot (right).

4.1.2 Phase I-B: Last 25 Days of Decomposition

The preliminary plots of the last 25 days of decomposition indicate a possible issue with the CO_2 measurements on Day 23 and on Day 30. In both cases the amount of CO_2 measured on these two days from almost every bucket appears to be much larger than the measurements surrounding these days. Models were fit both including the CO_2 measurements from these two days and without them. The models were compared to assess how influential the CO_2 measurements from both of these days were on the fit of the model in order to assess if these CO_2 measurements should be removed.

Model Including Day 23 and Day 30 CO_2 Measurements

The same explanatory variables were considered for these data as for the analysis for the first five days of decomposition. The following three models were fit in R and compared:

Model 1: $CO_2 = trt + cday + trt * cday + cday^2 + trt * cday^2$
Model 2: $CO_2 = trt + cday + trt * cday + cday^2$
Model 3: $CO_2 = trt + cday + trt * cday$

Model	R^2	AIC
Model 1	0.5132	13375.42
Model 2	0.5101	13371.29
Model 3	0.4558	13432.46

Table 5. R^2 and AIC values from the three fitted models.

Model 2 was chosen based on the smaller AIC value (Table 5). This model yielded an $R^2 = 0.5101$. Diagnostic plots were judged acceptable (Figure 10).

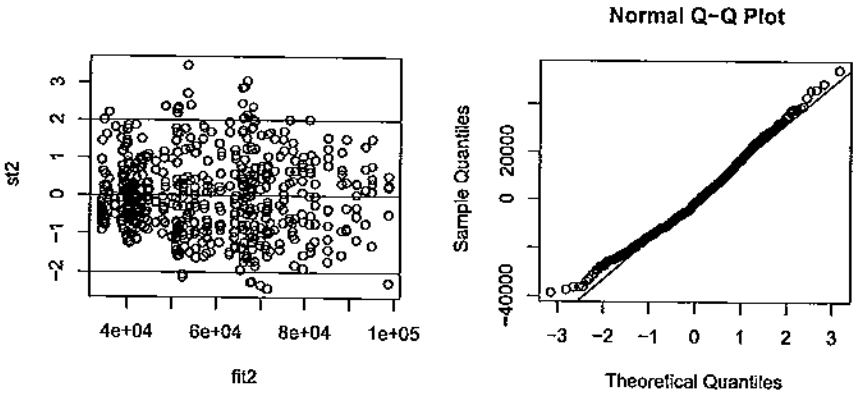

Figure 10. Studentized Residuals vs. Fitted Values plot (left) and Normal Q-Q plot (right).

Residuals vs. day were plotted to check for a temporal pattern of CO_2 over time. There appears to be a slight pattern of the residuals over the 30 days but not an obvious pattern (Figure 11). It was decided to not attempt to account for temporal correlation and to ignore the repeated measures taken on each bucket over the 30 day period of time due to the exploratory nature of this analysis and the statistical knowledge of the researchers.

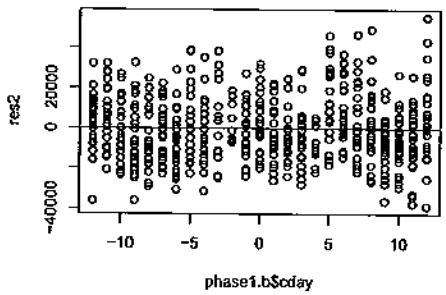

Figure 11. Residuals vs. Centered Day.

Model Not Including Day 23 and Day 30 CO_2 Measurements

The CO_2 measurements for Day 23 and Day 30 were removed and the same process was repeated fitting the same three models including the same explanatory variables.

Model	R^2	AIC
Model 1	0.5505	12178.36
Model 2	0.5490	12172.29
Model 3	0.5065	12219.73

Table 6. R^2 and AIC values from the three fitted models.

Model 2 was chosen based on the smaller AIC value (Table 6). This model yielded an $R^2 = 0.5490$. Diagnostic plots were judged acceptable (Figure 12).

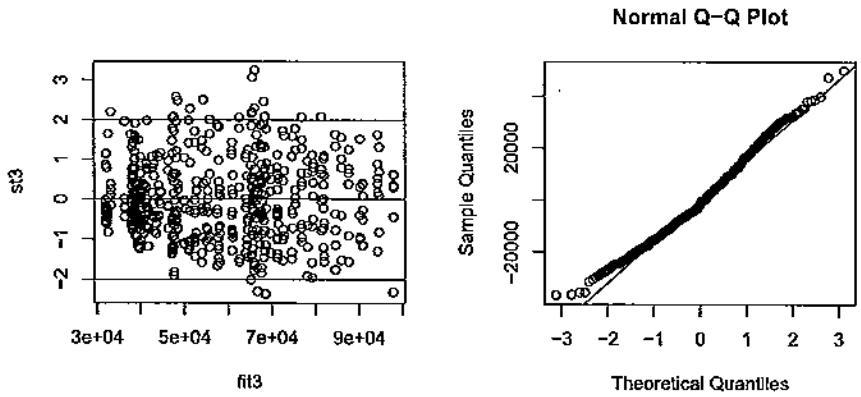

Figure 12. Studentized Residuals vs. Fitted Values plot (right) and Normal Q-Q plot (right).

Comparing the Models

Model 2, $CO_2 = trt + cday + trt * cday + cday^2$, was judged best whether Day 23 and Day 30 were excluded or not.

Both fitted models (CO_2 measurements included and excluded) appear to have fairly normal residuals and similar residual vs. fitted value plots. It does appear looking at the residual vs fitted value plots that the model fit excluding the CO_2 measurements for Day 23 and Day 30 has slightly smaller residual values overall compared to the model fit including the CO_2 measurements for Day 23 and Day 30, but this is to be expected (Figures 11 and 14).

Finally, comparing the two R^2 values from the two models, we see that the model fit excluding CO_2 measurements for Day 23 and Day 30 has a slightly larger R^2 value (Table 7). This tells me that the model excluding the CO_2 measurements from these two days explains slightly more of the variability than the model including these measurements.

Model	R^2
With Day 23 and Day 30	0.5101
Without Day 23 and Day 30	0.5490

Table 7. R^2 values from the model fit including CO_2 measurements from Day 23 and Day 30 and the model fit without these measurements.

The model fit including the CO_2 measurements from Day 23 and Day 30 was chosen because removing these measurements did not largely effect the fit of the model and the researchers were wary about removing any CO_2 measurements.

4.1.3 Treating Day as a Discrete Variable

Treatment comparisons would be easier if *day* was treated as a factor. Some implications of treating day as a categorical variable instead of a continuous variable include not being able to discuss the effect of time (day) on CO_2 output which is not of interest to the researchers and a loss of 16 degrees of freedom. However, there are over 500 degrees of freedom associated with the error term, which makes a difference of 16 degrees of freedom small in comparison.

Also, the results from above comparing the models fit both including and excluding the CO_2 measurements from Day 23 and Day 30 may indicate that estimating/accounting for CO_2 for each individual day may be more appropriate and fit the data better then modeling the trend of CO_2 over time treating day as a continuous variable. In order to assess how well including day as a discrete variable would model the data, the following two models were fit and compared. These models include the same explanatory variables as above except that *day* will now be a discrete variable that will not need to be centered.

Model 1: $CO_2 = trt + day + trt * day$
Model 2: $CO_2 = trt + day$

Model	R^2	AIC
Model 1	0.5941	13476.21
Model 2	0.5531	13352.11

Table 8. R^2 and AIC values from the two fitted models.

Model 2 was chosen based on the smaller AIC value (Table 8). This model yielded an $R^2 = 0.5531$. Diagnostic plots were judged to be acceptable (Figure 13).

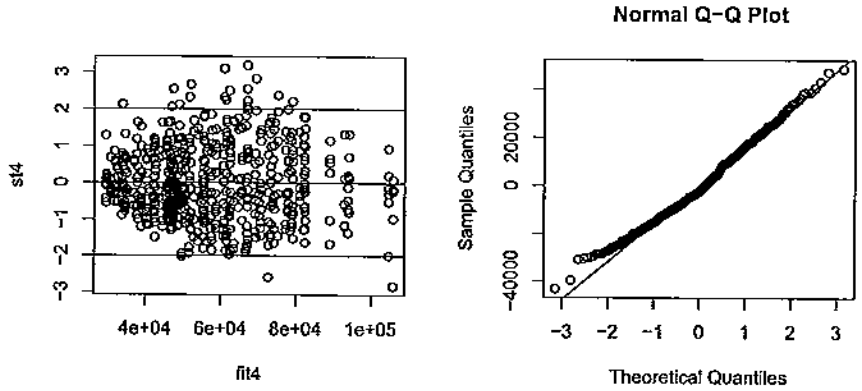

Figure 13. Studentized Residuals vs. Fitted Values plot (right) and Normal Q-Q plot (right).

4.1.4 Choosing a Model

In order to compare the fit of the model treating day as a discrete variable to the model chosen treating day as a continuous variable, I decided to compare both the R^2 values as well as their AIC values. After considering the implications and taking into account the limited statistical background of the researchers, the model fit treating *day* as a discrete variable was chosen because it yielded a slightly larger R^2 value and smaller AIC value (Table 9).

Model	R^2	Adjusted R^2	AIC
Treating *cday* as Continuous	0.5101	0.5018	13371.29
Treating *day* as Discrete	0.5531	0.5312	13352.11

Table 9. R^2, Adjusted R^2 and AIC values of the two models (day as continuous vs. day as discrete).

4.1.5 Final Model and Interpretation

Based on the results above, it was decided to model the CO_2 measurements for the full 30 days using this model. The same procedure was used to decide which model to use including the same models to compare as well as the same explanatory variables as described earlier.

Model 1: $CO_2 = trt + day + trt * day$
Model 2: $CO_2 = trt + day$

Model	R^2	AIC
Model 1	0.6233	15744.88
Model 2	0.5437	15665.93

Table 10. R^2 and AIC values from the two fitted models.

Model 2 was chosen based on the smaller AIC value (Table 10). This model yielded an $R^2 = 0.5437$. Diagnostic plots were judged to be acceptable (Figure 14).

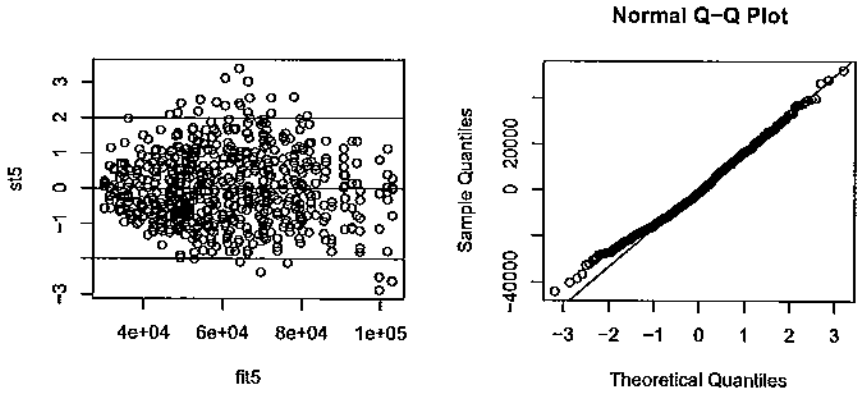

Figure 14. Studentized Residuals vs. Fitted Values plot (right) and Normal Q-Q plot (right).

A Tukey's multiple comparison procedure (MCP) was performed to assess differences in mean CO_2 output between the different treatments including 95% family-wise confidence intervals using the glht function in the multcomp package (Table 11).

					95% Confidence Interval	
Hypotheses	Estimate	SE	t-value	Pr(>|t|)	lower	upper
A2 − A1 = 0	-3787.5053	1903	-1.991	0.271	-8991.0479	1416.0372
B1 − A1 = 0	-2217.1855	1926	-1.151	0.779	-7484.1717	3049.8006
B2 - A1 = 0	10354.6737	1926	5.377	< 0.0001	5087.6876	15621.6599
C1 - A1 = 0	19424.2924	1840	10.556	< 0.0001	14392.0557	24456.5291
B1 − A2 = 0	1570.3198	1934	0.812	0.927	-3717.9154	6858.5550
B2 − A2 = 0	14142.1790	1934	7.314	< 0.0001	8853.9438	19430.4142
C1 − A2 = 0	23211.7977	1844	12.588	< 0.0001	18168.7315	28254.8639
B2 − B1 = 0	12571.8593	1922	6.540	< 0.0001	7314.4013	17829.3172
C1 − B1 = 0	21641.4779	1856	11.663	< 0.0001	16566.8605	26716.0953
C1 − B2 = 0	9069.6186	1856	4.888	< 0.0001	3995.0013	14144.2360

Table 11. Results from Tukey's MCP including 95% family-wise confidence intervals.

The following are results from Tukey's MCP for the comparisons of interest:

- Treatments A1 vs. A2: There is no evidence of a difference in mean CO_2 output (p-value = 0.271). We estimate the mean difference to be -3787.5053 ppm (SE = 1903 ppm) and a 95% confidence interval of (-8991.0479, 1416.0372).
- Treatments B1 vs. B2: There is strong evidence of a difference in mean CO_2 output (p-value < 0.0001). We estimate the mean difference to be 12571.8593 ppm (SE = 1922 ppm) and a 95% confidence interval of (7314.4013, 17829.3172).
- Treatments A1 vs. C1: There is strong evidence of a difference in mean CO_2 output (p-value < 0.0001). We estimate the mean difference to be 19424.2924 ppm (SE = 1840 ppm) and a 95% confidence interval of (14392.0557, 24456.5291).
- Treatments A2 vs. C1: There is strong evidence of a difference in mean CO_2 output (p-value < 0.0001). We estimate the mean difference to be 23211.7977 ppm (SE = 1844 ppm) and a 95% confidence interval of (18168.7315, 28254.8639).
- Treatments B1 vs. C1: There is strong evidence of a difference in mean CO_2 output (p-value < 0.0001). We estimate the mean difference to be 21641.4779 ppm (SE = 1856 ppm) and a 95% confidence interval of (16566.8605, 26716.0953).
- Treatments B2 vs. C1: There is strong evidence of a difference in mean CO_2 output (p-value < 0.0001). We estimate the mean difference to be 9069.6186 ppm (SE = 1856 ppm) and a 95% confidence interval of (3995.0013, 14144.2360).

Note that these are family-wise confidence intervals. We can be approximately 95% confident all of the intervals capture the true difference in mean CO_2 output.

4.2 Phase III: Biomass

4.2.1 Logit Model: Modeling Probability of No Seeds Germinating

A logistic regression model was fit to assess the probability of a zero (no plant grown in pot) in order to use the non-zero data when modeling biomass and height. Explanatory variables considered are:

- *Treatment*: To assess if the probability of getting a zero (no plant) differs between treatments (A1, A2, B1, B2, C1 and S).
- *Autoclave*: To assess if the probability of getting a zero (no plant) differs between soil that was autoclaved and non-autoclaved.
- *Row*: To account for the design of the experiment in which the researchers decided to block by rows.
- *Treatment* $*$ *Autoclave*: The interaction term between treatment and autoclave to allow the probability of a zero (no plant) to differ by treatment depending on which type of soil was used (autoclaved or non-autoclaved).

The following model was fit in R:

$\text{logit}(\pi) = trt + autoclave + row + trt * autoclave$

There were numerical difficulties in fitting the model in R. This issue could be accounted for by excluding any interaction between treatment and autoclave, however it is clear by looking at the data that an interaction exists (Tables 12 and 13). This issue may also be addressed using some other statistical methods, such as a Bayesian approach, but nothing has been done with these data yet.

In general, there appears to be a larger number of pots that yielded zero germinated seeds with the autoclaved soil than the non-autoclaved soil. It also appears that this difference is mostly seen in treatment 6, the soil with no added compost (Tables 12 and 13).

	Autoclaved Soil						
	A1	A2	B1	B2	C1	S	Total
Zero seeds germinated in pot	0	2	2	3	4	9	20
At least one seed germinated in pot	12	10	10	9	8	3	52

Table 12. Counts of pots that yielded either no seeds germinated or at least one seed germinated separated by treatments for autoclaved soil.

	Non-Autoclaved Soil						
	A1	A2	B1	B2	C1	S	Total
Zero seeds germinated in pot	1	1	1	2	2	0	7
At least one seed germinated in pot	11	11	11	10	10	12	65

Table 13. Counts of pots that yielded either no seeds germinated or at least one seed germinated separated by treatments for non-autoclaved soil.

	Autoclaved Soil				
	Number of Germinated Seeds				
Treatment	0	1	2	3	4
1 (A1)	0	2	4	3	3
2 (A2)	2	2	2	5	1
3 (B1)	2	5	2	2	1
4 (B2)	3	3	5	0	1
5 (C1)	4	3	4	1	0
6 (S)	9	0	0	1	2
Total	20	15	17	12	8

Table 14. Number of pots that germinated a particular number of seeds for the six treatments for autoclaved soil.

	Non-Autoclaved Soil				
	Number of Germinated Seeds				
Treatment	0	1	2	3	4
1 (A1)	1	5	4	2	0
2 (A2)	1	2	2	5	2
3 (B1)	1	7	3	0	1
4 (B2)	2	3	5	2	0
5 (C1)	2	2	3	5	0
6 (S)	0	1	5	4	2
Total	7	20	22	18	5

Table 15. Number of pots that germinated a particular number of seeds for the six treatments for non-autoclaved soil.

4.2.2 Modeling Biomass Without Zeros

All values of zero for biomass were removed and the explanatory variables considered are the same used for the logistic regression stated previously. The following two models were fit in R and compared:

Model 1: $Biomass = trt + autoclave + row + trt * autoclave$
Model 2: $Biomass = trt + autoclave + row$

Model	R^2	AIC
Model 1	0.2783	281.2053
Model 2	0.2310	278.6217

Table 16. R^2 and AIC values from the two fitted models.

Model 2 was chosen based on the smaller AIC value (Table 16). This model yielded an $R^2 = 0.2310$. Diagnostic plots were judged to be acceptable (Figure 15).

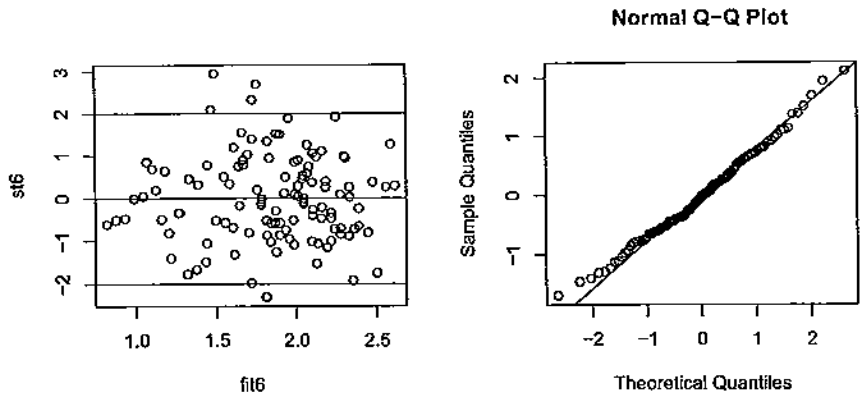

Figure 15. Studentized Residuals vs. Fitted Values plot (right) and Normal Q-Q plot (right).

There is some evidence of an affect due to autoclaving (p-value = 0.0717). We estimate that mean biomass is 0.2635 grams less for plants grown in autoclaved soil than in non-autoclaved soil (SE = 0.1449 grams) with a 95% confidence interval of (-0.5505, 0.0236) (Table 17).

	Estimate	SE	t-value	p-value	95% CI	
Autoclave1	-0.2635	0.1449	-1.819	0.0717	-0.5505	0.0236

Table 17. Output for the autoclave1 term in the fitted model.

A Tukey's MCP was performed to assess differences in mean biomass between the different treatments. The glht function in the multcomp package in R was used (Table 18).

					95% Confidence Interval	
Hypotheses	Estimate	SE	t-value	Pr(>|t|)	lower	upper
2 - 1 = 0	0.1695	0.2283	0.742	0.9760	-0.4922	0.8311
3 - 1 = 0	-0.3511	0.2254	-1.558	0.6264	-1.0043	0.3021
4 - 1 = 0	-0.1728	0.2346	-0.737	0.9768	-0.8529	0.5072
5 - 1 = 0	0.1441	0.2465	0.585	0.9918	-0.5702	0.8584
6 - 1 = 0	-1.0028	0.2559	-3.918	0.0021	-1.7446	-0.2610
3 - 2 = 0	-0.5206	0.2331	-2.234	0.2299	-1.1961	0.1550
4 - 2 = 0	-0.3423	0.2418	-1.416	0.7162	-1.0430	0.3584
5 - 2 = 0	-0.0253	0.2534	-0.100	1.0000	-0.7599	0.7092
6 - 2 = 0	-1.1723	0.2616	-4.481	< 0.001	-1.9305	-0.4140
4 - 3 = 0	0.1783	0.2393	0.745	0.9756	-0.5153	0.8718
5 - 3 = 0	0.4952	0.2509	1.974	0.3627	-0.2320	1.2224
6 - 3 = 0	-0.6517	0.2602	-2.504	0.1312	-1.4059	0.1025
5 - 4 = 0	0.3170	0.2592	1.223	0.8238	-0.4342	1.0681
6 - 4 = 0	-0.8300	0.2671	-3.108	0.0282	-1.6040	-0.0560
6 - 5 = 0	-1.1469	0.2785	-4.119	0.0010	-1.9540	-0.3398

Table 18. Results from Tukey's MCP including 95% family-wise confidence intervals.

The following are results from Tukey's MCP for the comparisons of interest:

- Treatments 1 vs. 2 (A1 VS. A2): There is no evidence of a difference in mean biomass (p-value = 0.9760). We estimate the mean difference to be 0.1695 grams (SE = 0.2283 grams) and a 95% confidence interval of (-0.4922, 0.8311).
- Treatments 3 vs. 4 (B1 VS. B2): There is no evidence of a difference in mean biomass (p-value = 0.9756). We estimate the mean difference to be 0.1783 grams (SE = 0.2393 grams) and a 95% confidence interval of (-0.5153, 0.8718).
- Treatments 1 vs. 5 (A1 VS. C1): There is no evidence of a difference in mean biomass (p-value = 0.9918). We estimate the mean difference to be 0.1441 grams (SE = 0.2465 grams) and a 95% confidence interval of (-0.5702, 0.8584).
- Treatments 2 vs. 5 (A2 VS. C1): There is no evidence of a difference in mean biomass (p-value = 1.0000). We estimate the mean difference to be -0.0253 grams (SE = 0.2524 grams) and a 95% confidence interval of (-0.7599, 0.7092).
- Treatments 3 vs. 5 (B1 VS. C1): There is no evidence of a difference in mean biomass (p-value = 0.3627). We estimate the mean difference to be 0.4952 grams (SE = 0.2509 grams) and a 95% confidence interval of (-0.2320, 1.2224).
- Treatments 4 vs. 5 (B2 VS. C1): There is no evidence of a difference in mean biomass (p-value = 0.8238). We estimate the mean difference to be 0.3170 grams (SE = 0.2592 grams) and a 95% confidence interval of (-0.4342, 1.0681).
- Treatments 1 vs. 6 (A1 VS. S): There is strong evidence of a difference in mean biomass (p-value = 0.0021). We estimate the mean difference to be -1.0028 grams (SE = 0.2559 grams) and a 95% confidence interval of (-1.7446, -0.2610).
- Treatments 2 vs. 6 (A2 VS. S): There is strong evidence of a difference in mean biomass (p-value < 0.0010). We estimate the mean difference to be -1.1723 grams (SE = 0.2616 grams) and a 95% confidence interval of (-1.9305, -0.4140).
- Treatments 3 vs. 6 (B1 VS. S): There is no evidence of a difference in mean biomass (p-value = 0.1312). We estimate the mean difference to be -0.6517 grams (SE = 0.2602 grams) and a 95% confidence interval of (-1.4059, 0.1025).
- Treatments 4 vs. 6 (B2 VS. S): There is strong evidence of a difference in mean biomass (p-value = 0.0282). We estimate the mean difference to be -0.8300 grams (SE = 0.2671 grams) and a 95% confidence interval of (-1.6040, -0.0560).
- Treatments 5 vs. 6 (C1 VS. S): There is strong evidence of a difference in mean biomass (p-value = 0.0010). We estimate the mean difference to be -1.1469 grams (SE = 0.2785 grams) and a 95% confidence interval of (-1.9540, -0.3398).

Note that these are family-wise confidence intervals. We can be approximately 95% confident all of the intervals capture the true difference in mean biomass.

5 Conclusion/Scope of Inference

The researchers indicated at the beginning of the study that they did not know what to expect. Further, there were issues with the study design and subsequent measurements of responses. Accordingly, this analysis should be treated as exploratory.

5.1 Phase I

For Phase I, the buckets of food were not randomly assigned a treatment, so we cannot claim causation but rather discuss an association between treatment and CO_2 output. Therefore, we can say that there is a strong association between lower CO_2 measurements and treatment B1 when compared to treatment B2, treatment A1 when compared to treatment C1, treatment A2 when compared to treatment C1, treatment B1 when compared to treatment C1, and treatment B2 when compared to treatment C1.

The samples of food used in the buckets for decomposition were not randomly sampled from all cafeterias on Montana State University's campus but were taken from two cafeterias on MSU's campus that had leftover food available. Since this food was then thoroughly mixed and the samples to be used in the buckets were arbitrarily selected from this pile to fill the buckets, we can only infer our results for CO_2 output back to the sample of food used for decomposition within the buckets from these two cafeterias on MSU's campus.

5.2 Phase III

For Phase III, since the treatments were randomly assigned to a pot of 22 grams of soil, we can claim causation between treatment and the response variable of interest. Since the experimental unit is a pot of 22 grams of soil, and since the soil was chosen from one brand and was not a random sample from this brand of soil as it was taken arbitrarily from the soil on hand, we can only infer back to the sample of soil from this brand used in this experiment. This means that we can conclude that adding in compost from treatments A1, A2, B1, and C1 caused more biomass to grow when compared individually to the pots only filled with soil while holding autoclaved/non-autoclaved soil constant and after accounting for the blocking effect of row in the model.

6 R Code Appendix

```
setwd("C:/Users/Elizabeth/Desktop/Bokashi Experiment/Consulting")
phase1 <- read.csv("co2-phase1a.csv", head = TRUE)
phase1 <- subset(phase1, phase1$trt != "C2")
phase1 <- subset(phase1, phase1$trt != "C3")

xyplot(co2 ~ day | trt, group = bucket, data = phase1, type = "l")

phase1.data <- read.csv("co2-phase1.csv", head=TRUE)
phase1.data <- subset(phase1.data, phase1.data$trt!="C2")
phase1.data <- subset(phase1.data, phase1.data$trt!="C3")
phase1.data <- subset(phase1.data, phase1.data$co2!="NA")
phase1.data$inc <- ifelse(phase1.data$day==1 | phase1.data$day==2
      | phase1.data$day==3 | phase1.data$day==4 | phase1.data$day==5, 1,0)
phase1.data$trt <- factor(phase1.data$trt)

xyplot(co2 ~ day | trt, group = bucket, data = phase1.data, type="l")

phase1.a <- subset(phase1.data, phase1.data$inc == 1)
xyplot(co2 ~ day | trt, group = bucket, data = phase1.a, type = "l")

phase1.a <- subset(phase1.a, phase1.a$day != 1)
xyplot(co2 ~ day | trt, group = bucket, data = phase1.a, type = "l")

phase1.b <- subset(phase1.data, phase1.data$inc == 0)
xyplot(co2 ~ day | trt, group = bucket, data = phase1.b, type = "l")

phase3.w.data <- read.csv("phase3-weight.csv", head = TRUE)
phase3.w.data$trt <- factor(phase3.w.data$trt)
phase3.w.data$autoclave <- factor(phase3.w.data$autoclave)
phase3.w.data$units <- interaction(phase3.w.data$row, phase3.w.data$pot)
phase3.w.data$ind <- ifelse(phase3.w.data$dryweight == 0, 0, 1)
phase3.w.data <- read.csv("phase3-weight.csv", head=TRUE)
phase3.w.data$trt <- factor(phase3.w.data$trt)
phase3.w.data$autoclave <- factor(phase3.w.data$autoclave)
phase3.w.data <- subset(phase3.w.data, phase3.w.data$dryweight!=0)

xyplot(dryweight ~ trt|autoclave, data = phase3.w.data, groups = row,
       main="Biomass by Treatment Separated by Autoclaving")

phase1.a$trt <- factor(phase1.a$trt)

phase1.a$cday <- phase1.a$day - 3.5

co2a.lm1 <- lm(co2 ~ trt + cday + I(cday^2) + trt*cday + trt*I(cday^2),
               data = phase1.a)
co2a.lm2 <- lm(co2 ~ trt + cday + I(cday^2) + trt*cday, data = phase1.a)
co2a.lm3 <- lm(co2 ~ trt + cday + trt*cday, data = phase1.a)

r.s1 <- summary(co2a.lm1)$r.squared
r.s2 <- summary(co2a.lm2)$r.squared
r.s3 <- summary(co2a.lm3)$r.squared

c(r.s1, r.s2, r.s3)
c(AIC(co2a.lm1), AIC(co2a.lm2), AIC(co2a.lm3))
```

```
res1 <- residuals(co2a.lm1)
fit1 <- fitted(co2a.lm1)

par(mfrow = c(1, 2))
st1 <- studres(co2a.lm1)
plot(fit1, st1)
abline(h = 0)
abline(h = -2)
abline(h = 2)

qqnorm(res1)
qqline(res1)

phase1.b <- subset(phase1.data, phase1.data$inc == 0)

## center day ##

phase1.b$cday <- phase1.b$day - 18
phase1.b$trt <- factor(phase1.b$trt)

co2b.lm1 <- lm(co2 ~ trt + cday + trt*cday + I(cday^2) + trt*I(cday^2),
               data = phase1.b)

co2b.lm2 <- lm(co2 ~ trt + cday + trt*cday + I(cday^2), data = phase1.b)

co2b.lm3 <- lm(co2 ~ trt + cday + trt*cday, data = phase1.b)

r.s1b <- summary(co2b.lm1)$r.squared
r.s2b <- summary(co2b.lm2)$r.squared
r.s3b <- summary(co2b.lm3)$r.squared

c(r.s1b, r.s2b, r.s3b)
c(AIC(co2b.lm1), AIC(co2b.lm2), AIC(co2b.lm3))

res2 <- residuals(co2b.lm2)
fit2 <- fitted(co2b.lm2)

par(mfrow = c(1, 2))
st2 <- studres(co2b.lm2)
plot(fit2, st2)
abline(h = 0)
abline(h = -2)
abline(h = 2)

qqnorm(res2)
qqline(res2)
plot(phase1.b$cday, res2)
abline(h = 0)

phase1.bb <- subset(phase1.b, phase1.b$day != 23)
phase1.bb <- subset(phase1.bb, phase1.bb$day != 30)

## model- full model ##
co2bb.lm1 <- lm(co2 ~ trt + cday + trt*cday + I(cday^2) + trt*I(cday^2),
                data = phase1.bb)

co2bb.lm2 <- lm(co2 ~ trt + cday + trt*cday + I(cday^2), data = phase1.bb)

co2bb.lm3 <- lm(co2 ~ trt + cday + trt*cday, data = phase1.bb)
```

```
r.s1bb <- summary(co2bb.lm1)$r.squared
r.s2bb <- summary(co2bb.lm2)$r.squared
r.s3bb <- summary(co2bb.lm3)$r.squared

c(r.s1bb, r.s2bb, r.s3bb)
c(AIC(co2bb.lm1), AIC(co2bb.lm2), AIC(co2bb.lm3))

res3 <- residuals(co2bb.lm2)
fit3 <- fitted(co2bb.lm2)

par(mfrow = c(1, 2))
st3 <- studres(co2bb.lm2)
plot(fit3, st3)
abline(h = 0)
abline(h = -2)
abline(h = 2)

qqnorm(res3)
qqline(res3)

phase1.b2 <- phase1.b
phase1.b2$day <- as.factor(phase1.b2$day)

co2b2.lm1 <- lm(co2 ~ trt * day, data = phase1.b2)

co2b2.lm2 <- lm(co2 ~ trt + day, data = phase1.b2)

r.sb1 <- summary(co2b2.lm1)$r.squared
r.sb2 <- summary(co2b2.lm2)$r.squared

c(r.sb1, r.sb2)
c(AIC(co2b2.lm1), AIC(co2b2.lm2))

res4 <- residuals(co2b2.lm2)
fit4 <- fitted(co2b2.lm2)

par(mfrow = c(1, 2))
st4 <- studres(co2b2.lm2)
plot(fit4, st4)
abline(h = 0)
abline(h = -2)
abline(h = 2)

qqnorm(res4)
qqline(res4)

setwd("C:/Users/Elizabeth/Desktop/Bokashi Experiment/Consulting")
phase1 <- read.csv("co2-phase1.csv", head = TRUE)

phase1.data <- phase1[, 1:5]
phase1.data <- phase1.data[1:751, ]
phase1.data <- subset(phase1.data, !is.na(co2))

phase1.data <- subset(phase1.data, phase1.data$day != 1)
phase1.data$trt <- factor(phase1.data$trt)
phase1.data$day <- factor(phase1.data$day)
```

```
co2.lm1 <- lm(co2 ~ trt * day, data = phase1.data)

co2.lm2 <- lm(co2 ~ trt + day, data = phase1.data)

r.1 <- summary(co2.lm1)$r.squared
r.2 <- summary(co2.lm2)$r.squared

c(r.1, r.2)
c(AIC(co2.lm1), AIC(co2.lm2))

res5 <- residuals(co2.lm2)
fit5 <- fitted(co2.lm2)

par(mfrow = c(1, 2))
st5 <- studres(co2.lm2)
plot(fit5, st5)
abline(h = 0)
abline(h = -2)
abline(h = 2)

qqnorm(res5)
qqline(res5)

mcp1 <- glht(co2.lm2, linfct = mcp(trt = "Tukey"))
summary(mcp1)
confint(mcp1)

germ.data <- read.csv("germ-log.csv", head = TRUE)
germ.data$germin <- ifelse(germ.data$germ == "y", 1, 0)
germ.data$trt <- as.factor(germ.data$trt)
germ.data$autoclave <- as.factor(germ.data$autoclave)
germ.data$row <- as.factor(germ.data$row)

germ.fit1 <- glm(germin ~ trt * autoclave + row, data = germ.data,
                 family = binomial(link = "logit"))

phase3.w.data <- read.csv("phase3-weight.csv", head=TRUE)
phase3.w.data$trt <- factor(phase3.w.data$trt)
phase3.w.data$autoclave <- factor(phase3.w.data$autoclave)
phase3.w.data <- subset(phase3.w.data, phase3.w.data$dryweight!=0)

weight.lm1 <- lm(dryweight ~ trt + autoclave + row + trt*autoclave,
                 data = phase3.w.data)

weight.lm2 <- lm(dryweight ~ trt + autoclave + row, data = phase3.w.data)

r.w1 <- summary(weight.lm1)$r.squared
r.w2 <- summary(weight.lm2)$r.squared

c(r.w1, r.w2)
c(AIC(weight.lm1), AIC(weight.lm2))

res6 <- residuals(weight.lm2)
fit6 <- fitted(weight.lm2)

par(mfrow = c(1, 2))
st6 <- studres(weight.lm2)
plot(fit6, st6)
abline(h = 0)
```

```
abline(h = -2)
abline(h = 2)

qqnorm(res6)
qqline(res6)

mcp2 <- glht(weight.lm2, linfct = mcp(trt = "Tukey"))
summary(mcp2)
confint(mcp2)
```