Chocolate Chip Cookies as a Teaching Aid

Herbert K. H. Lee*

Abstract

Getting and retaining the attention of students in an introductory statistics course can be a challenge, and poor motivation or outright fear of mathematical concepts can hinder learning. By using an example as familiar and comforting as chocolate chip cookies, the instructor can make a variety of statistical concepts come to life for the students, greatly enhancing learning.

Keywords: Variability; Cookies; Teaching Statistics; Active Learning

1 Introduction

Teaching introductory statistics at a university can be a major challenge because most of the students in the classroom do not want to be there, and are only taking the course because it is a requirement for their major or it fulfills a distributional requirement. Many students have already heard stories from their friends or even their parents about how hard and unpleasant statistics will be. The challenge is amplified when teaching in a larger lecture format, because of the limited opportunities for personal contact. So it is quite helpful to convince the students to buy-in to the course, to get their interest and to connect the course to concepts that they are comfortable with (Singer and Willett, 1990; Snee, 1993). The use of hands-on activities to facilitate active learning has been well established (Garfield and Ahlgren, 1988; Gnanadesikan et al., 1997).

Many introductory textbooks use examples from sports, biology, and economics in an attempt to relate the subject matter to areas of interest to the students, and discussions of

*Herbert Lee is Associate Professor, Department of Applied Mathematics and Statistics, University of California, Santa Cruz, 95064.

the benefits are present in the literature (Kvam and Sokol, 2004, for example). However, I have found that none of these examples are of interest to the whole class, often leaving one-third to one-half of the class uninterested in the examples and still uninterested in the course. Even worse, there can be gender biases in the relative interest in these examples. However, I have found that almost all students like chocolate chip cookies. Other instructors have also documented success in using snack food to motivate statistical education, including M&M's (Dyck and Gee, 1998), Reese's Pieces (Rossman and Chance, 1999), Hershey's Kisses (Richardson and Haller, 2002), and chewing gum (Richardson et al., 2005), although these tend to occur as isolated examples of specific concepts in a course. From a pedagogical standpoint, cookies can go beyond being a gimmick in a single class to becoming a recurring theme that can demonstrate many of the key concepts in a basic statistics course, allowing students to better see the connections between the different topics. The rest of this article will discuss how the following topics can all be related to cookies:

- variability
- inter-rater agreement and measurement error
- random distributions, the Poisson distribution, extreme values
- exploratory data analysis: displays of data, outliers
- sampling distributions
- hypothesis testing: one sample t-tests, two sample t-tests, and analysis of variance
- Bayesian statistics: prior elicitation and prior sensitivity

2 Bringing Cookies into the Classroom

By connecting statistical concepts directly to cookies, I have been able to gain the interest of students, even in larger classes. Once the students have found a reason to pay attention, subsequent teaching is much easier. In particular, what is needed is an eye-opening example that the students can relate to. What better than chocolate chip cookies? Consider a bag of mass-produced cookies, such as the brand-name Chips Ahoy or Keebler, or your favorite local store brand. These cookies are all produced to fairly high quality control standards. They are all about the same size. Most students expect that the cookies will all have about the same number of chips. That's what mass-production is all about, right?

Most statistics textbooks have a chapter on exploratory data analysis and graphical displays near the beginning of the book. This is a perfect time to bring in the cookies. I bring in enough cookies for everyone. We discuss the cookies and discuss variability, but the students consistently underestimate the variability. I then pass around the cookies, telling them to each take one cookie and count the chips. Of course, you can't see all of the chips from the outside, so you need to eat the cookie to find all of the chips. The counting sometimes leads to discussions of "what counts as a chip" because of broken chips, and can lead to discussions of inter-rater variability.

In any case, I have found that the students suddenly become engaged in the course. Almost everyone likes eating cookies. After they have had a chance to count, I have them tell me how many they found and I enter the numbers into my laptop. The students at first tend to think that this is just going to be a long boring process in a large class, but they are soon surprised to hear the numbers coming from their classmates. The variability is far beyond what they expect. For example, in one class we counted 101 cookies with the number

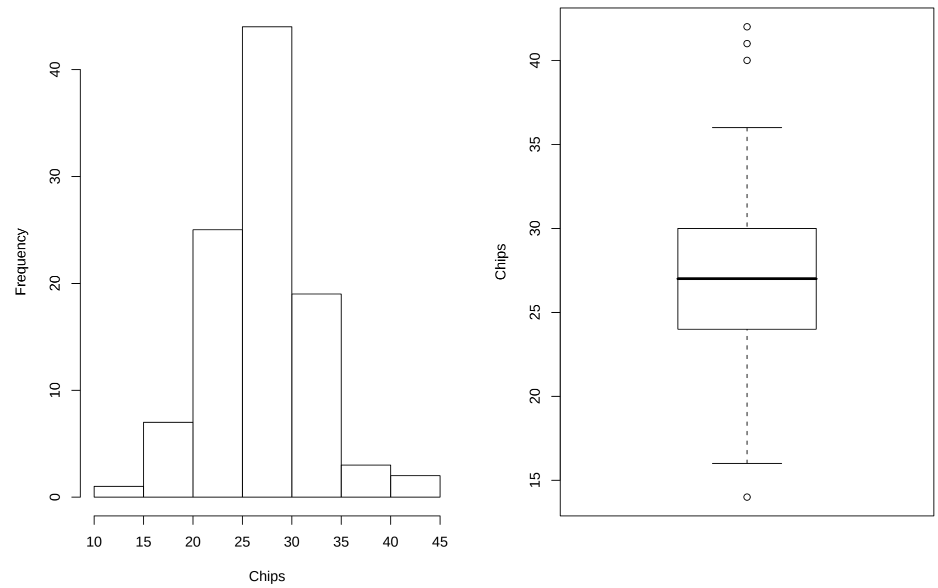

Figure 1: Histogram and boxplot of the number of chips in each of 124 cookies.

of chips ranging from 14 to 42 (see Figure 1). Until you stop and think about it, that seems like an awful lot of variability. But if you consider a large batch of well-mixed dough, the number of chips per cookie should be approximately Poisson. In fact, that range is similar to what one could reasonably expect from draws from a Poisson distribution with mean 27.2 (the empirical mean). In an advanced class, this can lead to further discussions of the Poisson distribution and of the distribution of extremes. In an introductory class, this dataset also serves as a good example for graphical displays such as histograms, stem-and-leaf diagrams, and boxplots. It usually generates datapoints that a standard statistical package will mark as potential outliers (see Figure 1), leading to a discussion of what it means to be an outlier.

This is a paradigm-shifting moment for most students. They think they understand cookies, yet suddenly realize that there is far more variability than they expected. They may call into question previous cookie eating experiences. They may wonder if their friends have been getting more chips than they have. But they suddenly have a very tangible example of how variability can affect them personally. And now they have a reason to care about their statistics class. As documented by Sowey (2001), such a "striking demonstration" provokes excitement in the topic and facilitates learning. At this point I can impress upon them that statistics is the study of randomness and variability.

3 Inference and Beyond with Cookies

With the taste of cookies still fresh, I can now preview why they should pay attention to the rest of the course. I ask them how they would decide if a manufacturer is putting enough chips in their cookies. If the manufacturer claims an average of 27 chips per cookie, how would they decide if the manufacturer is truthful, using only a highly variable sample of cookies? Next we can consider comparing brands. Given how much difference there is in the counts for a single brand, how could we possibly hope to say if one brand of cookies has more chips than another? The answer, as we know, is that we need statistical methodology.

The concept of a sampling distribution can be difficult for the students to grasp until they have a live example (Gourgey, 2000), and the cookies work wonderfully. After personally counting one bag, the students can easily envision that a different bag would have different counts, even if it came from the same population. Indeed, in a larger class, more than one bag is necessary for everyone to get a cookie, so the data can even be recorded by bag.

The data are clearly relevant when we get to one-sample t-tests. And when we get to two-sample t-tests, I bring in a different brand of cookies. Confronted with the natural variability, the students understand that without statistical tools, they have no way to decide if a difference in means between brands is just random variability or if it is a real difference. Thus they can see why the t-test is necessary, and they can relate the parts of the t statistic to the cookie data.

The running example goes further. I bring in a third brand when I introduce one-way ANOVA. By having a continuing example, it makes the fact that ANOVA is a generalization of a t-test more tangible for the students. And by now, the concept of comparing variability within brands to variability across brands is familiar to the students.

After the successes I have had in introductory classes, I have started using cookies in many more of my classes. As previously stated, they can be used in discussions of the Poisson distribution, or of extremes. Inter-rater variability in counting is another topic. One can compare empirical and theoretical distribution functions.

I use the cookies as an example for prior elicitation in a graduate course in Bayesian statistics. I have discovered that most students don't have a good idea of what the mean should be, which leads to an interesting discussion of forming a prior. We can then see how the prior is updated with cookie data using Bayes' Theorem, and see what the relative influences of the prior and the data are. Because the students do both the prior elicitation and the data collection, the whole problem becomes much more real for them.

4 Assessment

Assessment of teaching effectiveness is an important consideration. One important metric is student comprehension, which is classically measured through quizzes and exams, although projects and dynamic assessment techniques have possible advantages (Chance, 1997). Another important aspect is that of student perception, for which the traditional approach is that of teaching evaluations filled out by the students at the end of the course. Here our evaluations have three open-ended questions that ask "Please comment on how the instructor's teaching helped your learning in this course", "Please suggest how the instructor's teaching might improve", and "Other Comments". From those questions alone, I get a number of responses that show the value of the cookies, for example: "I especially felt that his examples were relevant and helpful (such as the cookies!)"; "I really liked the hands on examples, e.g. the chocolate chip counting in the three different types of cookies. It made the material fun and easy to understand."; and "Loved the examples – good connections to other subjects and brought 'statistics to life' with cookie examples." The students thus respond positively to active learning methods. They also provide comments showing that such examples do make a difference in gaining their attention and assisting with learning a subject about which they had initial apprehensions: "The approach for teaching this usually 'dry' class is very good. Cookies in class are very good."; "Cookies were a great way to get students to come to class, get into the material, and understand it."; "I hate math and I was dreading the idea of takings stats. This instructor made the class bearable and I actually learned what I needed to."; and "He made a dull subject (statistics) interesting ... I enjoyed this class and I didn't think that I would."

Comments under the improvement question rarely relate to the cookies, with the memorable exceptions being "have more examples with healthier types of food" and "Vegan friendly class experiments, please."

5 Conclusions

One of the goals of a teacher of introductory statistics is to have their students think about statistical concepts in their lives outside of the classroom. By using the common and supposedly familiar chocolate chip cookie, one can demonstrate variability in an eye-opening way to the students. This demonstration causes them to look at cookies differently, and thus brings an important statistical concept directly into their lives. It is then much easier for them to apply these concepts to other examples in their lives. And now that they have a personally relevant example, they are much more likely to be interested in the rest of the course.

References

Chance, B. L. (1997). "Experiences with Authentic Assessment Techniques in an Introductory Statistics Course." *Journal of Statistics Education*, 5, 3.

Dyck, J. L. and Gee, N. A. (1998). "A Sweet Way to Teach Students About the Sampling Distribution of the Mean." *Teaching of Psychology*, 25, 3, 192–195.

Garfield, J. and Ahlgren, A. (1988). "Difficulties in Learning Basic Concepts in Probability and Statistics: Implications for Research." *Journal for Research in Mathematics Education*, 19, 1, 44–63.

Gnanadesikan, M., Scheaffer, R. L., Watkins, A. E., and Witmer, J. A. (1997). "An Activity-Based Statistics Course." *Journal of Statistics Education*, 5, 2.

Gourgey, A. F. (2000). "A Classroom Simulation Based on Political Polling To Help Students Understand Sampling Distributions." *Journal of Statistics Education*, 8, 3.

Kvam, P. H. and Sokol, J. (2004). "Teaching statistics with sports examples." In *INFORMS Transactions on Education*, vol. 5. Http://ite.pubs.informs.org/Vol5No1/KvamSokol/.

Richardson, M. and Haller, S. (2002). "What is the Probability of a Kiss? (It's Not What You Think)." *Journal of Statistics Education*, 10, 3.

Richardson, M., Rogness, N., and Gajewski, B. (2005). "4 out of 5 Students Surveyed Would Recommend this Activity (Comparing Chewing Gum Flavor Durations)." *Journal of Statistics Education*, 13, 3.

Rossman, A. J. and Chance, B. L. (1999). "Teaching the Reasoning of Statistical Inference: A 'Top Ten' List." *College Mathematics Journal*, 30, 4, 297–305.

Singer, J. D. and Willett, J. B. (1990). "Improving the Teaching of Applied Statistics: Putting the Data Back Into Data Analysis." *The American Statistician*, 44, 223–230.

Snee, R. (1993). "What's Missing in Statistical Education." *The American Statistician*, 47, 2, 149–154.

Sowey, E. R. (2001). "Striking Demonstrations in Teaching Statistics." *Journal of Statistics Education*, 9, 1.