

Teacher's Corner

What Happens When a $1 \times 1 \times r$ Die is Rolled?

Peter K. DUNN

This article examines the probabilities of outcomes from rolling dice with the dimension $1 \times 1 \times r$ for various values of r . Experiments were conducted by school students and university students. The results of the experiments are given and the probabilities examined using a generalized linear model. Notes are also made about the value of the experiment in teaching the groups of students.

KEY WORDS: Binomial; Experimentation; Dice; Generalized linear model; Probability.

1. INTRODUCTION

Dice have been used for centuries in gambling. Probabilities of various outcomes have been studied since about 1654 when Blaise Pascal and Pierre de Fermat began discussions. All of these discussions naturally focus on the use of standard cubic dice. But what happens if we consider dice of dimensions $1 \times 1 \times r$? (Note that r can be considered the aspect ratio of the die.)

This problem was recently posed to a group of Year 11 and 12 students (ages about 16 and 17) and then a few weeks later to a small group of undergraduate university students studying generalized linear models. A picture of the dice in question (see Figure 1) was presented to the students and the students were asked a series of questions. Note that both 1×1 faces on the die are labeled with a 6. The students then performed a set of experiments on some dice constructed for various values of r .

This article has two points of focus: First, the results of rolling the nonstandard die and the associated probabilities are discussed. Second, the approach of the two different student groups is examined. Initially, the Year 11 and 12 students will be considered followed by the university students. Some results will then be presented.

2. YEAR 11 AND YEAR 12 STUDENTS

I was asked to conduct a one-hour session for Year 11 and 12 school children at a Maths Day conducted by a local school. The Maths Day consisted of four one-hour activities throughout

Peter K. Dunn is a Lecturer in Statistics in the Faculty of Sciences, University of Southern Queensland, Toowoomba QLD Australia, 4350 (E-mail: dunn@usq.edu.au). The author acknowledges the workmanship of Oliver Kinder in producing the three sets of dice used in the experiment and the dedication of the students in rolling so many dice. The author also thanks the reviewers for some important corrections and suggestions that improved the article.

a day, conducted by willing school teachers and university staff (namely myself).

As a statistician, I was keen to be involved as previous Maths Days (and their precursor, Maths Camps) had been devoid of any statistical activities; I was hoping to change that. The result is the activity discussed in this article, prompted by a talk by Dave Griffiths which I had the pleasure of attending at the Australian Statistics Conference in Adelaide in 2000.

2.1 Preliminaries

Before approaching the problem of rolling a 6 with nonstandard dice, I tried to raise a bit of interest in the old (and potentially uninteresting) problem of rolling dice. I began by examining some data of actual dice rolls from history [using data found in Hand, Daly, Lunn, McConway, and Ostrowki (1996); dataset 131 shows Wolf's data, the frequency of each face in 20,000 rolls; dataset 263 gives Weldon's dice data, showing counts of various outcomes]. We discussed sampling errors and computed some simple probabilities.

I then showed them Efron's dice, invented by Bradley Efron. The dice consist of four cubes as shown in Figure 2. With Efron's dice $\Pr(A \text{ beats } B) = 2/3$, $\Pr(B \text{ beats } C) = 2/3$, $\Pr(C \text{ beats } D) = 2/3$, but also $\Pr(D \text{ beats } A) = 2/3$. The students were rather interested in this outcome.

Having created some interest, the students were then asked about the probabilities of various faces showing up for dice with dimensions $1 \times 1 \times r$. After showing a picture of the die for clarity (see Figure 1), they were asked to consider the plot of r against the probability of rolling a 6, say p . Many students were not sure, though some proposed linear solutions while others suggested curved and sigmoid shaped solutions.

As a guide, students were asked to consider some special cases: When $r = 1$ (standard cubic dice), as $r \rightarrow 0$ and as $r \rightarrow \infty$. They were then asked to make some guesses for $r = 2$, $r = 0.75$, and $r = 0.5$. After discussing these issues, most students were happy to accept that a linear relationship was unlikely.




Figure 1. A noncubic die. Notice that the 6 is on both 1×1 faces on the die.




Figure 2. Efron's dice. For these dice, on average A beats B, B beats C, C beats D and D beats A.

I then suggested an experiment be conducted to find the form of the relationship (fortunately, I had suitable equipment with me).

The relationship is interesting; it is not symmetric about $r = 1$, for example, since the probability of obtaining a 6 then is $1/3$ (remembering the two 1×1 faces are both marked with a 6), which is not halfway between the limits of 0 and 1.

2.2 Experimentation

The idea of an experiment was initially treated with enthusiasm by approximately half the students. Nonetheless, the discussion of how to conduct the actual experiment was engaged in by all students with an increasing degree of interest. The purpose of the experimentation was to promote discussion of some statistical concepts using a concrete example, and to then have the students conduct a simple experiment.

Some of the questions posed to the students included:

1. What values of r should we use, and why?
2. How many times should we roll each die of side r ? Should there be a different number of rolls depending on r ? Explain!
3. What other factors are there that might affect the answers besides the value of r ?
4. How should the data be reported?
5. What steps should be taken to conduct the experiment?
6. How should the dice be rolled?

Naturally, the questions did not have simple correct answers, and many students seemed (initially at least) to be reluctant to answer because they did not know the "right" answer. The data were briefly analyzed; we leave this discussion until Section 4. First, we discuss the above questions and the student responses.

2.2.1 What Values of r to Use?

The first question was approached with some trepidation. One student, on seeing the box of tricks I had with me, offered the solution "Whatever you have brought along." Some students thought just a few values of r would be sufficient, while others were keen for many values of r . After some discussion, most seemed to agree that lots of dice with values of r near 1 would be more helpful than having lots of dice with larger and larger values of r (such as $r = 2$, $r = 3$, $r = 4$, and so on). Students had previously deduced, not using this notation, that $p \rightarrow 0$ as $r \rightarrow \infty$ and $p \rightarrow 1$ as $r \rightarrow 0$. It was also noted that if r was too small (or too large), the probability of rolling a 6 would be effectively one (or zero) and the notion of an experiment with a random outcome would then fail.

Having never done the experiment myself, I was not sure of useful values of r to use either; I was also unsure what constituted

values of r that were too large or too small to remove the random outcome element. I did have, however, a selection of dice with me which I revealed at this time. The props consisted of three sets with dice having the following (approximate) values of r : 0.25, 0.5, 0.75, 0.85, 0.9, 1, 1.1, 1.15, 1.25, 1.5, 1.75, and 2. A unit was about 5cm (2 inches). I divided the students into three groups for conducting the experiment.

2.2.2 How Many Times Should Each Die be Rolled?

The second question was probably the most keenly discussed of all, and from that point discussion was more open. Some students were of the belief that no rolls were necessary for "small" and "large" values of r since "everyone knows what the answer will be" (despite some of their previous answers!). Other students, however, were keen to roll dice numerous times in the hope that an unusual event would appear, and probabilities very close to, but not exactly, 0 or 1 could be estimated with some accuracy. The students finally decided that 30 rolls for each value of r would be used.

Although 30 rolls may not seem a large number, it should be remembered that each of the 12 die would be rolled 30 times for a total of 360 rolls. There were also pragmatic issues: the potential for boredom was real, and the discussion thus far had taken about 30 minutes so there was about 30 minutes left. The sample of size 30 therefore represented a compromise between completing the experiment in a reasonable time and having enough data points to make some sensible conclusions.

2.2.3 What Other Factors Might Affect the Results?

Other factors that might affect the outcome were then discussed; I was fishing for some kind of randomization and experimental design. The students found it easy to accept that all other variables should be kept as constant as possible—in particular, the same person in each group should be responsible for rolling. After some discussion, I was able to persuade the students that randomly choosing a die to roll might be preferred over each group starting (for example) with the largest value of r and systematically working through the set.

These issues appeared to be received better in theory than in practice; in practice, some students were systematically working through the sets of dice until I reminded them of the discussion on randomization. The groups did, however, stick with the one person rolling the dice. However, in general it would be untrue to say the same person rolled the dice in similar fashion; there were some very creative rolling techniques in place to avoid the potential problems mentioned in Section 2.2.6.

One further issue discussed was how to ensure each face had a chance of appearing face-up. This is discussed more in Section 2.2.6.

2.2.4 How Should the Data be Reported?

The students answered this question easily, but I believe it is an important question to ask. If it had not been asked—even though it was easily answered—I suspect the data recording would have been insufficient or haphazard. Because the students had decided to use 30 rolls for each value of r , however, reporting the number of rolls was less prone to error than if different numbers of rolls were chosen for various r .

2.2.5 What Steps Should be Taken to Conduct the Experiment?

To address this question, we simply discussed some practical issues: Who would roll? Who would record? How would dice be randomly selected? There were (conveniently) three groups of four, and it was decided that each group would divide into pairs and do half the dice each (time was scarce).

2.2.6 How Should the Dice be Rolled?

An important issue to discuss is how to ensure each face had a chance of appearing face-up. Unless the students were made aware of this problem, the natural tendency, especially for the larger values of r , was to roll the dice out of the hand in a systematic manner giving the faces with a six little chance of appearing; naturally this must be warned against.

This question is of particular importance because a die rolled on a side with $r \neq 1$ will (almost) never show a 6 (unless perhaps r is very small), and a die rolled on a side length 1 will show a 6 with probability 0.5 (unless perhaps r is very large). Gentle rolling to give each face a reasonable chance to appear uppermost should be the aim; extreme rolling techniques approaching those mentioned above should be avoided. A simple demonstration of rolling dice in these extreme cases is generally sufficient to identify and explain the problem. Indeed, attention can be drawn by taking the cubic die and claiming the probability of rolling a 6 is 0.5, and then demonstrating this by rolling the die a number of times on a appropriate edge.

3. UNIVERSITY STUDENTS

A few weeks after the session with the school children, after returning to the corridors of academia, I was teaching a course on generalized linear models to a group of three undergraduate students. In discussions with the students (during a sidetrack), the dice experiment was mentioned. The students seemed very interested in the problem and proceeded to speculate on the relationship between r and p . Because of their interest, I then decided to use the experiment as an assignment question. The response of the students exceeded my expectations.

Similar questions to those used for the school students (see Section 2.2) were initially asked of the university students. Pleasingly, the response was quite enthusiastic. I encouraged the students to talk among themselves about the questions and tried to keep myself from interjecting (except to keep them on-task). It should be noted that the students were not aware which size dice I had available until after the questions were answered.

The students were asked to consider the plot of r against the probability of rolling a 6, p . Thankfully, none of the students proposed linear solutions; indeed, they all suggested curved and sigmoid shaped solutions between the limits of $p = 0$ and $p = 1$.

One more formal suggestion was that the probability of obtaining a 6 might be related to the ratio of the surface areas. Because the total surface area is $2 + 4r$, the probability of obtaining a 6 then might be $2/(2 + 4r) = 1/(1 + 2r)$. Interestingly, the shape of this graph is certainly not correct even though it satisfies the limits conditions as $r \rightarrow 0$ and $r \rightarrow 1$, and also that $r = 1$ produces the correct probability of $1/3$.

Similar questions to those asked of the school students were asked of the university students. Only three warrant any further discussion.

3.1 What Values of r to Use?

There was much discussion on this topic. The students finally decided that numerous values of r near 1 would be beneficial, while fewer values of r between 0 and 1 would be beneficial. There was more discussion on values of $r > 1$, since this is an unbounded region.

One student suggested values of r up to 5; others were happy to restrict to values up to $r = 2$.

3.2 How Many Times Should Each Die be Rolled?

The second question was again the most keenly discussed of all. Initially, the students' answers were ad hoc like the school students, until one student suggested determining the size of the sample necessary to achieve given accuracy using, as found in numerous introductory statistics classes,

$$n = (z_{\alpha/2}^*)^2 p(1-p)/B^2,$$

where $z_{\alpha/2}^*$ is the z -score for an $100(1-\alpha)\%$ confidence interval; p is the proportion; and B is the bound placed on the estimate of p . For a 95% confidence interval with $p = 0.5$, the conclusion was that around 380 rolls would be needed for a margin of error of ± 0.05 .

This was an interesting exercise as the students grappled with the issue of estimating the necessary sample size when the value of p was unknown (and indeed, was the quantity being estimated). All the students knew the formula and this potential problem, but I believe they all better understand the implication having encountered it themselves.

3.3 What Other Factors Might Affect the Results?

The students quickly realized the potential for experimental design techniques to be used. They all, thankfully, identified randomization as important (especially if they were to roll each of the 12 dice 380 times!). They also nominated controlling variation as important, and explained how the experiment should be conducted so that each roll was as similar as possible. The important issues from Section 2.2.6 were discussed at this point.

One student raised an interesting issue after the experiment had been completed: He noted that in rolling dice with r near 0 or for large r that there was a human inclination (for him at least) to try and roll the "unlikely" event. Another raised the issue, after experimentation had been completed also, of the physical size of the dice. For dice with small r , there is a small physical area on which the dice could balance; in contrast, the smaller surface areas of the dice with larger r still had a reasonable surface area on which to land.

All the students did exceptionally well in this assignment, and each spent many hours (probably far too many) rolling dice to get accurate data on the probabilities. Their analyses were thorough and thoughtful (and also submitted late, incidentally).

4. RESULTS

4.1 Initial analysis

The results are given in Table 1 and plotted in Figure 3. The first three groups are for the school students, and the last three "groups" are the university students (the letters correspond to the students first initials). Note that the last student actually mea-

Table 1. The Results of Throwing $1 \times 1 \times r$ Dice; Both 1×1 Faces Were Marked with a 6. The first three groups were collected by Year 11 and 12 school children; the final three were collected by three university students. Note that the last student actually measured the aspect ratio, r .

Ratio	Group 1		Group 2		Group 3		Group D		Group S		Group A		
r	s_1	n_1	s_2	n_2	s_3	n_3	s_D	n_D	s_S	n_S	r_A	s_A	n_A
0.25	30	30	27	30	30	30	294	300	97	100	0.22	266	270
0.50	28	30	26	30	29	30	234	300	84	100	0.45	208	270
0.75	18	30	20	30	24	30	167	300	68	100	0.73	170	270
0.85	16	30	16	30	23	30	151	300	48	100	0.86	139	270
0.90	19	30	13	30	15	30	144	300	48	100	0.90	116	270
1.00	12	30	8	30	12	30	99	300	29	100	1.00	93	270
1.10	9	30	8	30	11	30	62	300	20	100	1.12	80	270
1.15	4	30	5	30	7	30	79	300	24	100	1.16	61	270
1.25	6	30	5	30	4	30	58	300	16	100	1.24	48	270
1.50	1	30	1	30	2	30	14	300	5	100	1.50	24	270
1.75	0	30	2	30	1	30	5	300	4	100	1.76	14	270
2.00	0	30	0	30	0	30	4	300	0	100	2.00	10	270

sured the aspect ratio, r , rather than taking the (approximate) values I had supplied. The school students performed calculations on their calculators, and some plotted the data using a graphics calculator; the university students used the R software (see Ihaka and Gentleman 1996) for analysis. The results at $r = 1$ offer a kind of check against bias: the expected proportion is $1/3$, while the overall observed proportion is $253/780 \approx 0.33289$, suggesting that there is no large bias in the results.

Note that the university students did not have the patience to roll the calculated 380 rolls for each of the 12 dice. That the students rolled from 1,200 to 3,600 dice in total, however, is a testament to their patience and reflects the high level of enthusiasm and interest that the students had for the experiment.

Combining all the data for the various values of r (for Group A, the results were combined with the closest nominal value of r used by the other groups), estimates of the probabilities can be found. A plot of the probabilities is shown in Figure 4 and the probabilities are given in Table 2. From the plot it becomes apparent that some dice with $r < 0.25$ could be useful to fix the left tail of the plot.

4.2 Generalized Linear Models: Background

More sophisticated analyses are possible. The university students were asked to find a model for predicting the probability of rolling a 6, say p , from the aspect ratio r . Because the response variable is a proportion, all three chose to use a generalized linear

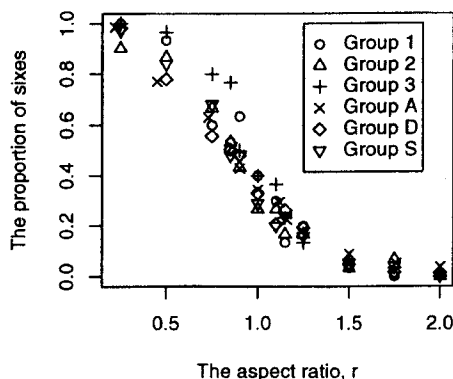


Figure 3. The probabilities of rolling a six for various values of r for the die in Figure 1.




Figure 4. The probabilities of rolling a six for various values of r for the die in Figure 1. The data in Table 1 have been combined.

model (GLM) based on the binomial distribution (see, e.g., Dobson 1983 or McCullagh and Nelder 1989). Standard regression and appropriate transformation could also be used, for example.

Generalized linear models, as proposed by Nelder and Wedderburn (1989), enable fitting models to a wide range of data types. These models are based on the family of distributions called exponential dispersion models, or EDMs, of which the binomial is a member (as are the normal, gamma, and Poisson distributions, for example). Generalized linear models consist of two components:

1. The response variable, y_i (\hat{p}_i is used for binomial data), comes from an EDM with mean μ (or $0 < p < 1$) and dispersion parameter ϕ (nominally, $\phi = 1$ for the binomial distribution).

Table 2. The Probabilities of Rolling a 6 for Various Values of r for the Die in Figure 1. The data in Table 1 have been combined (r was used rather than r_A with Group A for this purpose).

Aspect ratio r	Prob. of rolling 6	Aspect ratio r	Prob. of rolling 6
0.25	0.98	1.10	0.25
0.50	0.80	1.15	0.24
0.75	0.61	1.25	0.18
0.85	0.52	1.50	0.06
0.90	0.47	1.75	0.03
1.00	0.33	2.00	0.02

Table 3. The Deviance for Various Binomial Generalized Linear Models Fitted Without the Interaction (top line) and With the Interaction (bottom line).

	Link function used		
	Logit link	Probit	Comp. log-log
r + Group	141.6766	168.1534	121.1826
r + Group + r:Group	111.2969	131.6479	93.8946

- The expected values of the y_i, μ_i ($0 < p_i < 1$ is used for the binomial distribution), are related to the covariates \mathbf{x}_i through a monotonic, differentiable link function $g(\cdot)$ so that

$$g(\mu_i) = \mathbf{x}_i^T \boldsymbol{\beta},$$

where $\boldsymbol{\beta}$ is a vector of unknown regression coefficients.

Often, the linear component $\mathbf{x}_i^T \boldsymbol{\beta}$ is given the symbol η_i (the linear predictor), so that

$$g(\mu_i) = \eta_i = \mathbf{x}_i^T \boldsymbol{\beta}.$$

There are three link functions commonly used for binomial GLMs to map $0 < p < 1$ onto $-\infty < \eta < \infty$: the logistic, $\eta = \log(p/\{1-p\})$; probit $\eta = \Phi^{-1}(p)$, where $\Phi(\cdot)$ is the inverse cumulative distribution function for the standard normal distribution; and complementary log-log $\eta = \log\{-\log(1-p)\}$. The first two are often very similar and are symmetric about $p = 0.5$; the third does not share this symmetry. For this reason, the complementary log-log link function may prove to be the best of the three, as the data are not expected to be symmetric (see Section 2.1).

To assess the models, a quantity called the deviance can be used (e.g., see Firth 1991). In the case of a binomial distribution, the deviance is

$$\begin{aligned} D(y, \hat{\mu}) &= D(y, \hat{p}) \\ &= 2 \sum_{i=1}^n \left\{ \left(1 - \frac{y_i}{m_i}\right) \log \frac{1 - (y_i/m_i)}{1 - \hat{p}_i} + \frac{y_i}{m_i} \log \frac{(y_i/m_i)}{\hat{p}_i} \right\} \end{aligned}$$

for a sample of size n (in this application, the number of values of r), where y_i are the number of 6's rolled for each r , m_i are the total number of rolls for each r , and \hat{p}_i are the predicted proportions for each r . The deviance plays a role similar to the residual sum of squares for standard regression models; indeed, in the case of a normal distribution the deviance can be shown to be exactly the residual sum of squares. It can also be shown

Table 4. The Analysis of Deviance Table from R After Fitting the Full Model Including the Aspect Ratio r (r), the Group (Group) and the Interaction (r :Group). The first column are the variables, followed by degrees of freedom (df), deviance, residual degrees of freedom, residual deviance, and the p value based on a χ^2 test of the change in deviance.

	Df	Deviance	Resid. Df	Resid. Dev	P(> Chi)
NULL			71	4036.2	
r	1	3894.6	70	141.6	0.0
Group	5	20.4	65	121.2	1.060e-03
r:Group	5	27.3	60	93.9	5.013e-05

that the deviance has approximately a χ^2 -distribution on $n - p$ degrees of freedom (where n is the sample size and $p = \text{rank}(X)$ where X is the design matrix) when ϕ is known, otherwise an F -test is appropriate [details can be found in, e.g., McCullagh and Nelder (1989), sec. 2.3.2]. Two nested models can then be compared by using an analysis of deviance table in a similar way to how the variance is used in an analysis of variance table for normal-based models.

Generalized linear models can be fitted using most popular statistical software packages. The software used here is R.

4.3 Generalized Linear Models: Results

Although each of the students were to analyze their own data, I have the luxury here of using the combined results. This means the possible explanatory variables are the aspect ratio r , the group, and the interaction between them. Note that the dice were not constructed as three separate sets of 12 so any effect due to the set of dice used cannot be definitely measured. However, because the university students rolled far more times than the school students, any group effect can be primarily attributed to the properties of the dice also.

Two results become quickly apparent. First, the complementary log-log link function is the superior link function (on the basis of the residual deviance); see Table 3. Second, the interaction is statistically significant; more formally, an analysis of deviance test indicates that all three possible covariates are statistically significant; see Table 4. The model with the interactions and using the complementary log-log link function will be assumed hereafter unless indicated. In addition, for Group A, the values r_A (rather than r ; see Table 1) were used in all computations.

The mean deviance can be used as an estimate of ϕ ; here it is $\hat{\phi} = 93.9/60 \approx 1.57$, which is close to the nominal value of 1 for the binomial distribution. For this reason, a χ^2 test has been used rather than an F -test; the conclusion are similar in any case.

The fact that Group is significant may be attributed to the set of dice used, or imply that the rolling techniques were different.

To evaluate the model, a Q-Q plot of the quantile residuals (see Dunn and Smyth 1996) can be used to determine if the model appears suitable and, in particular, if the binomial distribution




Figure 5. A Q-Q plot of the quantile residuals for the binomial GLM fitted with a complementary log-log link function and including the interaction terms. A good model would have the points lying (approximately) on the straight line.




Figure 6. The prediction curves for the dice data, using both the complementary log-log and logit link functions. The filled points are from the university students data and are based on more information than the data of the school students (unfilled points) and are weighted proportionally heavier in fitting the model. The triangle identifies the point corresponding to the large negative residual in Figure 7; the cross identifies the point with very high influence as shown in Figure 8.

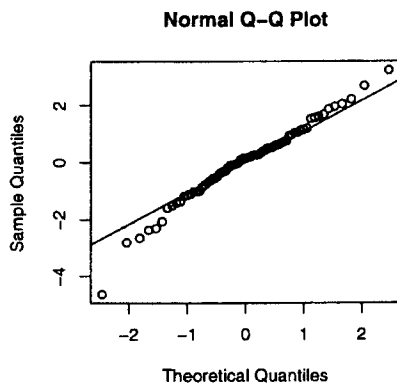


Figure 7. A Q-Q plot of the quantile residuals for the binomial GLM fitted with a complementary log-log link function with only r as a covariate. A good model would have the points lying (approximately) on the straight line.




Figure 8. Cook's distance for the binomial GLM fitted with a complementary log-log link function with only r as a covariate.

appears adequate for modeling the responses. This plot is shown in Figure 5 and indicates the model is adequate, though one point is of concern at the left of the plot. The large negative residual corresponds to Observation 38 (Group A; $r = 0.45$). The complementary log-log link produces a smaller residual deviance than the logit link and a large negative residual; the Q-

Q plot using the logit link function (not shown) indicates some larger positive residuals, and this link also produces a larger residual deviance.

Although the Group is an important factor, a simple relationship between the aspect ratio r and the probability of rolling a 6 is of interest; such a model can be found in R based on the complementary log-log link function as

$$\log(-\log\{1 - \hat{p}_i\}) = \frac{2.04188}{(0.05164)} - \frac{2.92478r}{(0.05800)},$$

where the standard errors are in parentheses below the parameter estimates; both parameters are highly significant. This information could be used to explicitly determine a relationship between r and p ; it is a little messy and is not given here. The prediction curves using the complementary log-log link function and, for comparison, the logit link function are plotted in Figure 6.

The Q-Q plot of the quantile residuals for this model is shown in Figure 7. Observation 28 (Group 3, $r = 0.85$) corresponds to the large negative residual; this is shown in Figure 6 with a triangle over the plotted point. Cook's distance can be used to identify influential observations (see McCullagh and Nelder 1983, sec. 12.7.3); Cook's distance for each observation is shown in Figure 8. Observation 38 (Group A; $r = 0.45$) corresponds to the large value; this is shown as a superimposed cross on the data point in Figure 6. This point is also noted above in the discussion of the full model.

5. CONCLUSIONS AND REFLECTIONS

This article, as mentioned in the Section 1, has two main points of focus, into which this section is divided.

First, the statistical results. While not claiming to be definitive, some results have been presented for estimating the probability of rolling a 6 for $1 \times 1 \times r$ dice (when each 1×1 face is marked with a six). The data are well modeled by a binomial generalized linear model using a complementary log-log link function. There is a suggestion that the dice or the technique of the rollers may be important.

Second, the teaching exercise. I found the exercise to be surprisingly well received by both the school and university students. The experiment was a scaffold for discussing various statistical concepts—such as sampling error, designing experiments, randomization, data recording as well as analysis and graphing. For the university students, concepts such as estimating sample size, modeling, and diagnostic tests were also used. For both groups, there was a great deal of interest in finding estimates of the probabilities. Importantly, the question was simple to pose, easily understood and of interest to the students; the experiment was also simple to perform (once the dice were constructed).

The exercise appeared to be a worthwhile contribution to Maths Day and I believe the students left with a sense of understanding some of the basic statistical concepts. Likewise, the university students found the idea appealing and invested a large amount of time in the assignment, while gaining an understanding of modeling with a binomial GLM and performing some simple diagnostic tests.

[Received March 2002. Revised June 2003.]

REFERENCES

- Dobson, A. J. (1983), *An Introduction to Statistical Modelling*, New York: Chapman and Hall.
- Dunn, P. K., and Smyth, G. K. (1996), "Randomized Quantile Residuals," *Journal of Computational and Graphical Statistics*, 5, 236–244.
- Firth, D. (1991), "Generalized Linear Models," in *Statistical Theory and Modelling: In Honour of Sir David Cox, FRS*, eds. D. V. Hinkley, N. Reid, and E. J. Snell, London: Chapman and Hall, pp. 55–82.
- Hand, D. J., Daly, F., Lunn, A. D., McConway, K. J., and Ostrowski, E. (1996), *A Handbook of Small Data Sets*, London: Chapman and Hall.
- Ihaka, R., and Gentleman, R. (1996), "R: A Language for Data Analysis and Graphics," *Journal of Computational and Graphical Statistics*, 5, 288–314.
- Nelder, J. A., and Wedderburn, R. W. M. (1972), "Generalized Linear Models," *Journal of the Royal Statistical Society, Ser. A*, 135, 370–384.
- McCullagh, P., and Nelder, J. A. (1989), *Generalized Linear Models* (2nd ed.), London: Chapman and Hall.