Statistical Gestalt: Illustrating Interaction with Indicator Variables

Joseph G. Eisenhauer
University of Detroit Mercy

Regression analyses frequently involve interactions among independent variables. Introductory courses in statistics, however, typically omit any discussion of interaction in the context of regression. The present note suggests a simple pedagogical method, based on indicator variables, for illustrating the value and use of interaction terms in multiple linear regression.

All wholes which are not mere aggregates, i.e. whose totality is something distinct from their parts, must have a principle of unity. …What, then, is it that makes 'man' one; why is he one and not many, e.g. animal + two-footed…? Why do[es] not the sum of these Ideas constitute the ideal man…? (Aristotle, translated 1956, p. 218)

The ancient Greek philosopher, Aristotle, famously pondered how a composite could be grander than the aggregation of its components—why a person, for example, is inevitably more complex than any collection of human features might suggest. As demonstrated below, his philosophy regarding the unity of matter and form, often referred to by the German word Gestalt and loosely abbreviated in common parlance as "the whole is greater than the sum of its parts," aptly describes the impact of two regressors on a response variable when interaction is present.[1]

In empirical research, regression analyses often include interaction terms to capture the joint effects of two (or more) independent variables. Indeed, Dawson (2014) suggests that interaction can be found in nearly any scholarly journal containing empirical work. Because multiple linear regression (MLR) is typically taught toward the end of an introductory college-level statistics course, however, the role of interaction terms in regression is often neglected, despite its importance. Although there are a few exceptions, including Kuiper and Sklar (2013), Moore (2010), and Ramsey and Shafer (2013), introductory textbooks designed for use by undergraduates frequently omit the topic. Similarly, there is remarkably little discussion of this subject among journal articles in the pedagogical literature. Flom and Strauss (2003, p. 1) note that "interpreting interactions is often difficult and sometimes counterintuitive" and offer a graphical approach for understanding interactions, but, in general, the role of interaction in regression has been largely overlooked in work designed to enhance undergraduate teaching.

Aside from mathematics and statistics majors, most college students, including those in liberal arts, education, business, and other disciplines, take only introductory courses in statistics and never formally study more advanced material.[2] Yet interaction effects are commonly found in a wide variety of applied fields, such as political science, sociology, economics, and psychology (Kam & Franzese, 2007), as well as nursing (Stommel & Wills, 2004), ecology (Engqvist, 2005), organization studies (Aguinis, 2002), and management (Dawson, 2014) to name just a few. Thus, at some point in their academic or professional careers, students may encounter a regression having an interaction term and wonder why the product of two or more independent variables is treated as a separate regressor.[3] Alternatively, in conducting their own analyses, they may fail to recognize situations in which interaction terms would be appropriate.

The present study proposes a convenient method for illustrating to undergraduates both why and how an interaction term may be used. In contrast to Aiken and West (1991), Flom and Strauss (2003), and Jaccard and Turrisi (2003), all of whom emphasize interactions involving continuous variables, the approach taken here is based on dummy variables and is designed for expository convenience in the classroom. The lesson is simple to explain, easy to understand, and quickly demonstrated. The following section presents the basic illustration, and a subsequent section provides additional insights, including the application of the example in assessment of learning. A comparison of the indicator approach with effects coding follows, and a short conclusion is given in the final section.

Interaction Between Indicators: An Example

Rather than starting with a didactic lecture on interaction followed by an illustration, the approach taken here begins with an example that challenges students to discover the interaction problem in the regression context and guides them toward a solution. Consider the hypothetical data in Table 1, for a cross-section of 20 individuals, where Salary is a continuous dependent variable measured in dollars (though it could just as well be measured in euros, pounds, or other currency units). Gender and the sector of the economy in which the individual is employed are both shown as indicator, or dummy, variables. Male is coded as 1 for men and as 0 for females. Corporate assumes the value 1 if the individual is employed by a corporation and 0 otherwise. We might think of a position in either the government or the nonprofit sector as the alternative to corporate employment. Naturally, we could employ empirical (rather than hypothetical) data and/or a larger sample size, but the present sample will suffice for illustrating the pedagogical technique. Note that there are equal numbers of non-corporate females, corporate females, non-corporate males, and corporate males. In the language of analysis of variance (ANOVA), these data constitute a balanced, 2×2 factorial experiment with the same number of observations in each cell. For the purpose of MLR, however, Table 1 presents the data as a continuous dependent variable and two binary independent variables. Unequal group sizes would work as well with the indicator approach developed here.

Table 1. Raw Data for Salaries

Salary ($)	Male	Corporate
39,000	0	0
40,000	0	0
40,000	0	0
40,000	0	0
41,000	0	0
41,000	1	0
42,000	1	0
43,000	1	0
44,000	1	0
45,000	1	0
50,000	0	1
51,000	0	1
52,000	0	1
53,000	0	1
54,000	0	1
60,000	1	1
61,000	1	1
61,000	1	1
61,000	1	1
62,000	1	1

It is useful to begin by inviting students to develop and test a model for predicting salary with the data at hand. When equipped with the data in Table 1, most students who are being introduced to MLR for the first time will dutifully regress *Salary* on *Male* and *Corporate*, and obtain the following result in an effort to determine the factors that influence salary:

$$\underset{(49.342)}{Salary = 38{,}500} \quad + \underset{(6.659)}{6{,}000Male} \quad + \underset{(16.649)}{15{,}000Corporate} \tag{1}$$

; where t-statistics are given in parentheses.

Students should now be asked to evaluate the results. The intercept and both of the independent variables are statistically significant well below the one percent level and the regression appears to have substantial explanatory power: $R^2 = 0.950$ and adjusted $R^2 = 0.944$. The regression equation correctly indicates that males tend to have significantly higher salaries than females, and that corporate employees enjoy significantly greater incomes than those with jobs in the non-corporate world. As an exercise to reinforce their understanding of the regression coefficients, students may be asked to calculate the grand mean from the regression equation and verify it from the raw data. Such a calculation is always possible with regressions obtained via ordinary least squares (OLS) because the regression line runs through the mean value of each variable. That is, in minimizing the sum of squared errors, the normal equations for the regression ensure that the mean value of the dependent variable is an exact linear function of the mean values of the independent variables.4 The mean value of an indicator variable is just the proportion of observations for which the characteristic of interest is present. Thus, the grand mean of salary, in our example, can be calculated by inserting the proportions of males and corporate employees into equation (1). Because half of the sample is male and half works in the corporate sector, the mean salary for the entire sample can be correctly calculated from equation (1) as $38,500 + $6,000(1/2) + $15,000(1/2) = $49,000.

However, the predictive value of this regression result is limited. To help them see this, students should be directed to use the regression equation to predict salaries for each type of worker in the sample by inserting values for the indicators and to then compare the predictions against the data. Equation (1) indicates that a randomly selected non-corporate female is expected to earn only $38,500 when, in fact, no one in the entire sample earns less than $39,000. The equation further suggests that a non-corporate male should expect a salary of $44,500, though only one such person in the sample actually receives that much. Additionally, the equation predicts a corporate female's salary to be $53,500 despite the fact that only one woman is paid that highly, and it indicates that a corporate male should anticipate an income of $59,500, even though every corporate male in the sample is paid $60,000 or more. Clearly, equation (1) generates embarrassingly poor estimates, especially given that these are in-sample predictions and not extrapolations. This red flag should signal a possible misspecification of the model.

Next, students should construct and evaluate a 2×2 table of mean salaries for each of the four types of individuals in the sample, as shown in Table 2. A careful inspection of the table confirms that although males and those employed by corporations do have higher average incomes than their counterparts, the effects of gender and the employment sector are not simply cumulative. Rather, their interaction creates a type of *Gestalt*: when the attributes Male and Corporate coincide, the increase in average salary is greater than the mere sum of the increases attributable to gender and sector separately. This can be seen in several ways. Compared with non-corporate females, non-corporate males have, on average, a $3,000 advantage. Corporate females have a $12,000 advantage, on average, so one might expect corporate males to have a $15,000 average advantage, but their mean salary is actually $21,000 greater than the mean of non-corporate females. Thus, the whole is indeed greater than the sum of its parts. Viewed from another perspective, the corporate versus non-corporate differential averages $12,000 among females, but $18,000 among males. This lack of uniformity indicates that gender affects the corporate differential. Equivalently, we could say that the male-female gap averages $3,000 in the non-corporate world, but $9,000 in the corporate world, so that the gender gap is magnified in the corporate sector.[5] The lesson to be learned from the juxtaposition of equation (1) and Table 2 is that focusing exclusively on main effects and ignoring interactions, or moderated effects, can be dangerously misleading.

Table 2. Mean Salaries

Attributes:	Female	Male
Non-Corporate	40,000	43,000
Corporate	52,000	61,000

At this stage, students should be encouraged to interpret the differentials in context. It is reasonable to suppose that for-profit corporations can generally afford to pay higher salaries than non-profit organizations or government agencies, which would explain why there is a difference across sectors. If the salary differential between males and females is attributed to discrimination, one might expect the differential to be roughly the same in the corporate and non-corporate sectors. The fact that the gender gap is larger for employees of corporations may indicate the existence of greater discrimination in the corporate world than elsewhere.

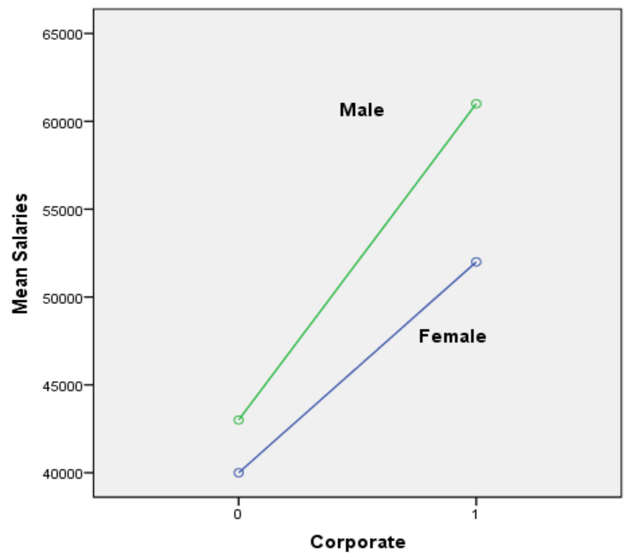

Figure 1. Profile plots for salaries.

As a useful visualization of these salary differentials, students can use the means in Table 2 to construct the profile plots shown in Figure 1. The line segments extend from the mean non-corporate salary to the mean corporate salary for each gender and the positive slope in each case reflects the corporate advantage. If there were no gender gap at all, the two curves in Figure 1 would overlap, but they occupy different positions because the male and female means differ. If no interaction were present between employment sector and gender, the two curves would be parallel to each other. In this case, however, the noticeable difference in slopes yields a wider income gap between males and females within corporations than in the non-corporate sector, and thus, signals the possibility of significant interaction between the employment sector and gender.[6]

Such a discussion highlights the need for an interaction term that captures the joint effects of the two independent variables. The most common method of incorporating interaction in MLR is to employ a product term (Jaccard & Turrisi, 2003). The rationale and mathematics behind this approach can be demonstrated simply, in a manner similar to that of Jaccard and Turrisi, Kam and Franzese (2007), or Aiken and West (1991). For notational consistency with equation (1), let the dependent variable S be a function of two independent variables, M and C. In the absence of an interaction term, the linear-additive model estimated by equation (1) is simply:

$$S = \alpha_0 + \alpha_1 M + \alpha_2 C + \varepsilon \tag{2}$$

where ε denotes random error. Now, since the relationship between *S* and *C* appears to be influenced by *M*, suppose that both α_0 and α_2 are linear functions of *M*, so that

$$\alpha_0 = \gamma_0 + \gamma_1 M\text{; and} \tag{3}$$

$$\alpha_2 = \delta_0 + \delta_1 M\,. \tag{4}$$

Equations (3) and (4) indicate that gender influences the intercept of (2) and the slope of the relationship between *S* and *C*, as depicted in Figure 1. Note that such interaction differs from multicollinearity, which would occur if *C* was a linear function of *M* (or vice versa). With this lead, students can substitute (3) and (4) into (2) and rearrange terms to get an interaction model that includes the multiplicative product of *C* and *M* as an independent variable:

$$S = \gamma_0 + (\alpha_1 + \gamma_1)M + \delta_0 C + \delta_1 (C \times M) + \varepsilon \tag{5}$$

By simply relabeling the coefficients, (5) can be rewritten in more familiar form as follows:[7]

$$S = \beta_0 + \beta_1 M + \beta_2 C + \beta_3 (C \times M) + \varepsilon \tag{6}$$

To estimate equation (6), students can easily construct *CorpMale* as the multiplicative product of the *Male* and *Corporate* dummies, where *CorpMale* = *Corporate* × *Male*. Then, *CorpMale* is itself an indicator variable, the first 15 observations of which are 0 and the final 5 observations of which take the value 1. Because this new independent variable is not a continuous linear function of either *Male* or *Corporate*, multicollinearity should not be a severe problem. While this justification for a multiplicative interaction term is illustrated in the context of indicators, it is important to point out that the same logic applies to continuous variables.[8]

If the regression is now run with the *CorpMale* interaction term included, the result is

$$\underset{(73.030)}{Salary = 40{,}000} + \underset{(3.873)}{3{,}000Male} + \underset{(15.492)}{12{,}000Corporate} + \underset{(5.477)}{6{,}000CorpMale} \tag{7}$$

Each of the independent variables, including the interaction term, is statistically significant. This indicates not only that males are paid more than females while corporate employees are paid more than non-corporate employees, but also that the effects of gender and employment sector are not simply cumulative. Additionally, the fit between the model and the data has actually improved: R^2 = .983 and adjusted R^2 = .979, while the corresponding *F* statistic has likewise increased from 160.76 to 300. Students can also verify that multicollinearity is not an issue by inspecting the variance inflation factors (VIF) for this regression, each of which is 3 or less.[9]

And recognizing from Table 1 that corporate males constitute one-quarter of the sample, the overall mean salary can again be correctly calculated as \$40,000 + \$3,000(1/2) + \$12,000(1/2) + \$6,000(1/4) = \$49,000. But most importantly, the predictive value of the regression has now been dramatically improved by the inclusion of the interaction term: equation (7) correctly predicts the mean salary for each of the four groups of individuals in Table 2. [10]

Of course, the choice of the baseline case, or reference group, for an indicator variable is essentially arbitrary. So, if the indicators are designated differently, the regression equation will have a different form, but it will nevertheless represent the same information. As a homework exercise, students might be asked to consider a reconfiguration of the data in Table 1 such that a *Female* indicator takes the value 1 for a woman and 0 for a man, and the interaction term *CorpFemale* is the multiplicative product of *Corporate* and *Female*. This yields the regression equation

$$\underset{(78.507)}{Salary = 43{,}000} - \underset{(-3.873)}{3{,}000Female} + \underset{(23.238)}{18{,}000Corporate} - \underset{(-5.477)}{6{,}000CorpFemale}. \tag{8}$$

The adjusted and unadjusted coefficients of determination, levels of significance, and other output are identical for equations (7) and (8). Both equations convey precisely the same information regarding mean salaries as Table 2. Thus, equation (7), where the interaction between *Corporate* and *Male* is positive, can readily be converted into equation (8), which portrays the negative interaction between *Corporate* and *Female*. Likewise, replacing *Corporate* with a *Non-corporate* indicator and creating an interaction term between gender and *Non-corporate* yields either:

$$Salary = 61{,}000 \; - \; 9{,}000Female \; - \; 18{,}000Non\text{-}corporate + 6{,}000Non\text{-}corpFemale \qquad (9)$$
$$\quad (111.370) \quad (-11.619) \qquad (-23.238) \qquad\qquad (5.477)$$

or

$$Salary = 52{,}000 \; + \; 9{,}000Male \; - \; 12{,}000Non\text{-}corporate \; - \; 6{,}000Non\text{-}corpMale, \qquad (10)$$
$$\quad (94.939) \quad (11.619) \qquad (-15.492) \qquad\qquad (-5.477)$$

each of which contains the same information as equations (7) and (8).

Having laid out the basic lesson, it is useful to emphasize that, as with any regression study, both theoretical and empirical considerations should guide the development of the model. There may be strong theoretical reasons to anticipate interactions among some subsets of variables and these subsets should be tested for statistically significant effects. In the health care context, for example, we might have an *a priori* expectation that the vital signs of a patient taking multiple medicines will be affected not only by each individual medicine, but also by the interaction of two or more taken together. Empirical techniques, such as forward selection, backward elimination, or stepwise regression, can be used to look for pairs (or other multiples) of variables that exhibit statistically significant interactions. Indeed, interaction will not always be present among the independent variables, and even when the profile plots are not perfectly parallel, this non-uniformity may result from sampling error. Thus, it is important to test the coefficients of interaction terms for statistical significance in the same way that the coefficients of other independent variables are tested. A parsimonious regression model preserves degrees of freedom by including only meaningful interactions. An example without statistically significant interaction is given in the Appendix.

Additional Insights

Instructors who wish (and can devote the time) to advance the discussion a step further can point out that the two indicator variables used in constructing Table 2 allow us to identify four distinct subsets of individuals. Conceptually, this is analogous to introducing a single qualitative characteristic having four categories, which would require the use of three dummy variables for completeness. If we approach the problem from this perspective, and allow non-corporate females to serve as the reference group, then we can reuse the CorpMale, CorpFemale, and Non-corpMale variables that were created for equations (7), (8), and (10), respectively.[11]

The resulting regression equation takes the form:

$$Salary = 40{,}000 \; + \; 12{,}000CorpFemale \; + \; 3{,}000Non\text{-}corpMale \; + \; 21{,}000CorpMale. \qquad (11)$$
$$\quad (73.030) \quad (15.492) \qquad (3.873) \qquad\qquad (27.111)$$

Equation (11) conveys the same fundamental information as equations (7), (8), (9), and (10), and has essentially the same diagnostics (e.g., $R^2 = .983$). The exception is that the VIF values are all reduced to 1.5, indicating even less chance of multicollinearity. The coefficients are statistically significant and again reveal *Gestalt*: non-corporate females would be paid an average of \$3,000 more if they were males and \$12,000 more in corporations, but if they were both corporate employees and male, then the average gain would not be \$15,000, but rather \$21,000.[12] The equivalence of (7) and (11) implies that omitting an interaction term between gender and the employment sector, as was done initially in equation (1), is comparable to omitting a (non-base) category of a qualitative characteristic from the regression, which helps to explain why (1) is a misspecification. Instructors may also take this opportunity to contrast the proper use of the indicator terms with the dummy variable trap—the problem of imposing multicollinearity by creating too many indicator variables and leaving no excluded category to serve as the reference group (see, for example, Hardy, 1993).

In addition to their pedagogical convenience, interaction examples involving indicator variables are also valuable as assessment instruments. Notice that equation (11) is easily constructed from Table 2 without even viewing the raw data or running a regression. That is, the mean salary for the reference group serves as the intercept, and each category's salary differential from the base case is its regression coefficient. Thus, after learning about the concept of interaction, a student should be able to determine equation (11) on the basis of Table 2 or vice versa.

Similarly, a well-prepared and astute observer of Table 2 should also be able to derive equation (7), or equivalently equations (8), (9), or (10), without having access to the raw data. For the model underlying equation (7), the mean salary of non-corporate females must be the intercept of the regression line; the differential in means when a single indicator assumes the value 1 is that variable's regression coefficient. Further, the coefficient of the *CorpMale* interaction term is the difference between the mean salary for corporate males and the sum of the intercept and the coefficients of *Male* and *Corporate*. To put this in more concrete terms, the figure of 40,000 located in the northwest quadrant of Table 2 is the intercept of the regression equation. Moving to the northeast quadrant, the additional 3,000 is the coefficient of *Male*, and moving from the northwest to the southwest quadrant, the extra 12,000 is the coefficient of *Corporate*. The sum of these three values is 40,000 + 3,000 + 12,000 = 55,000; therefore, to obtain the mean of 61,000 in the southeast quadrant of Table 2, the coefficient of the *CorpMale* interaction term must be the residual 6,000. The end result of this exercise is, of course, equation (7). Thus, by introducing interaction with indicator variables, instructors can readily develop homework assignments or test questions that require students to construct a regression equation from a 2×2 table of means (or vice versa) as a method for evaluating their understanding of regression coefficients.

Indeed, other examples of this type, with different contexts and different sample sizes, can easily be created by working backward from a matrix such as Table 2 to a data set such as that found in Table 1. Given any four subsample means, it is a simple matter to construct hypothetical data distributed around each mean. By establishing sufficiently large and unequal differences in group means, but relatively small deviations of data around each group mean, it is possible to develop statistically significant main effects as well as statistically significant interactions between the independent variables.

The lesson outlined here is designed as a first introduction to interaction and as such, should conclude with a brief comment about more sophisticated issues. A simple generalization of the basic idea behind equation (6) leads to the conclusion that interaction need not be restricted to two variables, but may involve three or more. Because dummy variables are a special case, it is important to note that continuous independent variables can handle some types of interactions that indicators cannot. For example, instructors might use this as a lead-in to a discussion of polynomial models; if we believe that a continuous independent variable interacts with itself (i.e., its effect on the dependent variable changes as its own level changes), then we can create a multiplicative interaction term by simply squaring the variable.[13] This, of course, cannot be done with dummies, since the square of an indicator is itself.

Effects Coding

The approach taken in equation (11), the use of three indicators and the implicit reference group to reflect four categories of a qualitative characteristic, is similar to the technique of effects coding. Effects coding would also designate a reference group, say non-corporate females, and three independent variables, say *CF*, *NCM*, and *CM*, to represent corporate females, non-corporate males, and corporate males, respectively. For the *CF* variable, all corporate females would be coded as 1, all members of the reference group would be coded as –1, and all others would be coded as 0; the variables *NCM* and *CM* would be constructed analogously. Regressing *Salary* on the effects-coded variables, *CF*, *NCM*, and *CM*, would give:

$$\underset{}{Salary} = \underset{(178.923)}{49{,}000} + \underset{(6.325)}{3{,}000CF} - \underset{(-12.649)}{6{,}000NCM} + \underset{(25.298)}{12{,}000CM} \qquad (12)$$

The R^2, adjusted R^2, F, and VIF diagnostics are the same for equations (11) and (12). The principal difference is that the intercept of (11) is the mean salary for the reference group, whereas in (12) the intercept is the un-weighted average of the group means. Note that if there are equal numbers of observations in each group, as in the present example, the un-weighted average of group means is equal to the grand mean. The coefficient of each independent variable measures that group's salary differential from the value taken by the intercept, and the mean for the reference group can be retrieved by setting $NCM = CF = CM = -1$ in equation (12).[14]

One pedagogical advantage of using indicator variables, rather than effects coding, to introduce interaction, is that students are more likely to have already studied the former than the latter by the time they are ready to learn about interaction. Many introductory texts, such as those by Anderson, Sweeney, and Williams (2005); Pearson (2010); and Peck, Olson, and Devore (2008), include sections on indicator variables and regression, but do not cover effects coding. Moreover, in contrast to the example above, empirical investigations often involve unbalanced experiments. Unequal sample sizes among groups will not alter the procedure or interpretation when using indicator variables, but will complicate the interpretation under effects coding. In particular, effects coding in an unbalanced experiment yields a regression intercept that no longer equals the grand mean (Alkharusi, 2012).[15] Therefore, it becomes necessary to discuss the difference between balanced and unbalanced design if one adopts the effects coding approach. Importantly, the use of indicator variables also allows an interaction term, such as the *CorpMale* variable in (7), to be constructed as the multiplicative product of two independent variables (in this case, the *Corporate* and *Male* variables of Table 1) just as product terms are commonly used to capture interaction among continuous variables. In contrast, effects coded variables, such as *CM*, are not product terms, so the insights from their use may not be as readily transferred to frameworks involving interaction among continuous variables. For these reasons, indicator variables may be a simpler way to introduce interaction in entry-level courses. Certainly, however, effects coding can be explored in more advanced courses or those with more time to devote to the topic. Indeed, equations (7), (11), and (12) represent three valid approaches to interaction, and instructors may wish to consider, or even experiment with, each of these methods to discover which one works best for their students.

Conclusion

When independent variables interact with each other, the totality of their influence on a regressand is indeed greater than the sum of their individual impacts taken separately. The example presented above demonstrates how such an interaction effect can be conveniently illustrated with indicator variables. The data are easily replicated for different contexts and other sample sizes, and the analysis can provide a valuable addition to the students' MLR tool kit. Indeed, linking statistical interaction to the broader philosophical concept of *Gestalt* may also help to make the lesson memorable. After the concept of interaction has been mastered using indicator variables, the notion may be more readily transferred to interaction terms constructed from continuous variables or from a combination of continuous and indicator variables. This foundation can promote an appreciation of interaction terms that should be useful to students beyond their initial statistics course—in understanding empirical studies encountered in assigned readings or literature reviews, and in conducting their own regression applications for other courses, as well as in the reading and analyses they may undertake in future careers.

Endnotes

1. *Webster's New Universal Unabridged Dictionary* (Random House, 1996, p. 802) defines gestalt as "a configuration, pattern, or organized field having specific properties that cannot be derived from the summation of its component parts; a unified whole."
2. Garfield, Hogg, Shau, and Wittinghill (2002) explain, "Introductory statistics is often the one and only statistics course taken by students who are not majoring in this discipline." Moore (1998, p.1253) states: "For most of these students, their first course is their one and only formal exposure to statistics." Roiter and Petocz (1996) note, "The introductory statistics course is especially important because it may be the only statistics course taken by future users of statistics."
3. Many popular introductory texts, such as those by Johnson and Kuby (2012) and Johnson and Bhattacharyya (2014), also do not cover interaction in analysis of variance (ANOVA). But even students who do see interaction in the context of ANOVA, may not appreciate the role of a product term in MLR.
4. For a good introductory presentation of the mathematics underlying OLS, see Kenkel (1989).
5. Interaction in which each differential's sign is uniform while its magnitude depends on another variable's level has been called uncrossed interaction; if the sign changes, the interaction is said to be crossed (Cohen & Cohen, 1983).
6. Robinson, Tomek, and Schumacker (2013) discuss testing the difference in slopes as an alternative to testing an interaction term for significance. Note also that with different data in more advanced courses, analysis of covariance (ANCOVA) could be used to isolate main effects by controlling for interaction between a categorical independent variable and a continuous covariate. The least squares (adjusted) group means in ANCOVA are calculated as if the covariate is always at its mean and thus, presupposes the absence of interaction (Huitema, 2011). The present example, however, does not use a continuous covariate and is, therefore, not well suited to ANCOVA.
7. Other assumptions would allow us to arrive at (6) as well. For example, in place of (3) and (4), we could specify $\alpha_0 = \gamma_0 + \gamma_1 C$ and $\alpha_1 = \delta_0 + \delta_1 C$. As emphasized by Kam and Franzese (2007, p. 15),

the end result is the same "because the propositions being modeled are logically symmetric." Notice also that the model in (6) allows for the possibility of significant interaction without significant main effects or vice versa.

8. Aiken and West (1991) and Kam and Franzese (2007) provide more extensive treatments of the mathematics underlying the use of product terms for interaction. Those interested in background literature may also wish to review Friedrich's (1982) summary of arguments against the use of multiplicative terms and his responses to them.
9. See Robinson and Schumacker (2009) for a discussion of VIF in the presence of interaction.
10. Note from (5) and (6) that $\beta_1 = \alpha_1 + \gamma_1$. In equation (1), α_1 is estimated at 6,000, and in equation (7), β_1 is estimated at 3,000. Thus, interested students should be able to retrieve an estimate of γ_1 as -3,000.
11. Recall that with the data in Table 1, the first 15 observations of *CorpMale* are 0 and the rest are coded 1; for *CorpFemale*, observations 11 through 15 are coded 1 and all others are 0, and for *Non-corpMale*, observations 6 through 10 are coded 1 and all others are coded 0.
12. Proceeding as before, the grand mean can be calculated from (11) as 40,000 + 12,000(1/4) + 3,000(1/4) + 21,000(1/4) = 49,000.
13. As Friedrich (1982, p. 829) observes, "though it is seldom recognized as such…a model with a squared or higher degree term is actually a special case of interaction, the case in which a variable interacts with itself."
14. For more detailed comparisons of dummy and effects coding, see Hardy (1993) and Alkharusi (2012).
15. This can be verified by applying effects coding to the example in the Appendix.

References

Aguinis, H. (2002). Estimation of interaction effects in organization studies. *Organizational Research Methods*, *5*, 207-211.

Aiken, L. S., & West, S. G. (1991). *Multiple regression: Testing and interpreting interactions*. Thousand Oaks, CA: Sage.

Alkharusi, H. (2012). Categorical variables in regression analysis: A comparison of dummy and effect coding. *International Journal of Education*, *4*, 202-210.

Anderson, D. R., Sweeney, D. J., & Williams, T. A. (2005). *Statistics for business and economics* (9th ed.). Mason, OH: Southwestern.

Aristotle. (1956). *Metaphysics*. (J. Warrington, Trans.). London: J.M. Dent & Sons.

Cohen, J., & Cohen, P. (1983). *Applied multiple regression/correlation analysis for the behavioral sciences* (2nd ed.). Hillsdale, NJ: Lawrence Erlbaum Associates.

Dawson, J. F. (2014). Moderation in management research: What, why, when, and how. *Journal of Business and Psychology*, *29*, 1-19.

Engqvist, L. (2005). The mistreatment of covariate interaction terms in linear model analyses of behavioural and evolutionary ecology studies. *Animal Behaviour*, *70*, 967-971.

Flom, P. L., & Strauss, S. M. (2003). Some graphical methods for interpreting interactions in logistic and OLS regressions. *Multiple Linear Regression Viewpoints*, *29*, 1-7.

Friedrich, R. J. (1982). In defense of multiplicative terms in multiple regression equations. *American Journal of Political Science*, *26*, 797-833.

Garfield, J., Hogg, B., Shau, C., & Wittinghill, D. (2002). First courses in statistical science: The status of educational reform efforts. *Journal of Statistics Education, 10(2)*. Retrieved from http://www.amstat.org/publications/jse/v10n2/garfield.html.

Hardy, M. A. (1993). *Regression with dummy variables*. Newbury Park, CA: Sage.

Huitema, B. E. (2011). *The analysis of covariance and alternatives: Statistical methods for experiments, quasi-experiments, and single-case studies* (2nd ed.). Hoboken, NJ: John Wiley & Sons.

Jaccard, J., & Turrisi, R. (2003). *Interaction effects in multiple regression* (2nd ed.). Thousand Oaks, CA: Sage.

Johnson, R. A., & Bhattacharyya, G. K. (2014). *Statistics: Principles and methods* (7th ed.). New York: Wiley.

Johnson, R., & Kuby, P. (2012). *Elementary statistics* (11th ed.). Boston: Brooks/Cole.

Kam, C. D., & Franzese, R. J. (2007). *Modeling and interpreting: Interactive hypotheses in regression analysis*. Ann Arbor, MI: University of Michigan Press.

Kenkel, J. L. (1989). *Introductory statistics for management and economics* (3rd ed.). Boston: PWS-Kent.
Kuiper, S., & Sklar, J. (2013). *Practicing statistics: Guided investigations for the second Course.* New York: Pearson.
Moore, D. S. (1998). Statistics among the liberal arts. *Journal of the American Statistical Association*. *93*, 1253-1259.
Moore, D. S. (2010). *The basic practice of statistics* (5th ed.). New York: W.H. Freeman & Company.
Pearson, R. W. (2010). *Statistical persuasion*. Thousand Oaks, CA: Sage.
Peck, R., Olson, C., & Devore, J. (2008). *Introduction to statistics and data analysis* (3rd ed.). Belmont, CA: Thomson.
Ramsey, F. L., & Shafer, D. W. (2013). *The statistical sleuth: A course in methods of data analysis* (3rd ed.). Boston: Brooks/Cole.
Random House (1996). *Webster's new universal unabridged dictionary*. New York: Author.
Robinson, C., & Schumacker, R. E. (2009). Interaction effects: Centering, variance inflation factor, and interpretation issues. *Multiple Linear Regression Viewpoints*, *35*, 6-11.
Robinson, C., Tomek, S., & Schumacker, R. E. (2013). Tests of moderation effects: Difference in simple slopes versus the interaction term. *Multiple Linear Regression Viewpoints*, *39*, 16-24.
Roiter, K., & Petocz, P. (1996). Introductory statistics courses—a new way of thinking. *Journal of Statistics Education*, 4(2). Retrieved from http://www.amstat.org/publications/jse/v4n2/roiter.html.
Stommel, M., & Wills, C. E. (2004). *Clinical research: Concepts and principles for advanced practice nurses*. Philadelphia: Lippincott, Williams & Wilkins.

Send correspondence to: Joseph G. Eisenhauer
University of Detroit Mercy
Email: joseph.eisenhauer@udmercy.edu

APPENDIX

Illustrating the Absence of Interaction

This appendix briefly illustrates the indicator approach in a setting with unbalanced design and without interaction; it could be used as an extension of the basic lesson or as homework. The hypothetical data in Table A1 show the number of votes garnered by 15 political candidates identified as either Republicans (coded 0) or Democrats (coded 1), and as challengers (0) or incumbents (1). There are fewer Democrats and incumbents than others, and the group means are recorded in Table A2. The mean advantage of Democrats over Republicans is consistent (1,000 votes) regardless of incumbency, and the mean advantage of incumbents over challengers is uniform (4,000 votes) regardless of political party. Thus, the profile plots in Figure A1 are parallel, suggesting a lack of interaction between incumbency and political party.

Table A1. Raw Data for Votes

Votes	Incumbent	Democrat
12,700	0	0
12,800	0	0
13,000	0	0
13,000	0	0
13,200	0	0
13,300	0	0
13,800	0	1
14,000	0	1
14,200	0	1
16,800	1	0
16,900	1	0
17,100	1	0
17,200	1	0
17,900	1	1
18,100	1	1

Table A2. Mean Votes

Attributes:	Democrat	Republican
Challenger	14,000	13,000
Incumbent	18,000	17,000

Regressing *Votes* on the indicators *Incumbent* and *Democrat* yields:

$$\underset{}{Votes} = \underset{(174.705)}{13{,}000} + \underset{(38.763)}{4{,}000}Incumbent + \underset{(9.352)}{1{,}000}Democrat \qquad \text{(A1)}$$

The main effects are statistically significant and R^2 = .993. Even without an interaction term, this regression correctly predicts each of the means in Table A2 and can be used to compute the grand mean as 13,000 + 4,000(6/15) + 1,000(10/15) = 14,933.33 votes. Indeed, as students can easily verify, a multiplicative interaction term (*IncDem* = *Incumbent*×*Democrat*) would now be statistically nonsignificant. This would alter the *t*-statistics, but have essentially no other effect on the coefficients; the result would be:

$$\underset{(155.717)}{Votes = 13{,}000} + \underset{(30.303)}{4{,}000\,Incumbent} + \underset{(6.916)}{1{,}000\,Democrat} + \underset{(0.00)}{0.00\,IncDem} \tag{A2}$$

The three-dummies method described above can also be applied to these data. Letting Republican challengers be the reference group and designating incumbent Republicans and Democratic challengers as *IncRep* and *DemChal*, respectively, generates the regression result:

$$\underset{(155.717)}{Votes = 13{,}000} + \underset{(30.303)}{4{,}000\,IncRep} + \underset{(6.916)}{1{,}000\,DemChal} + \underset{(29.946)}{5{,}000\,IncDem} \tag{A3}$$

where the intercept is once again the mean for the reference group. In this case, Republican challengers would gain, on average, 4,000 votes if they were incumbents in their own party, 1,000 votes if they switched parties, and a total of 5,000 votes if they were incumbents and switched parties. So, for this example, the whole is exactly equal to the sum of its parts—this time, there is no statistical *Gestalt*.

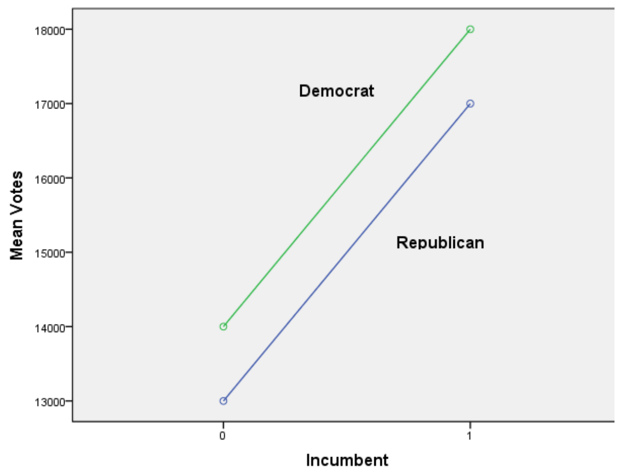

Figure A1. Profile plots for votes.