

Compression and Conditional Effects: A Product Term Is Essential When Using Logistic Regression to Test for Interaction*

CARLISLE RAINEY

*P*revious research in political methodology argues that researchers do not need to include a product term in a logistic regression model to test for interaction if they suspect interaction due to compression alone. I disagree with this claim and offer analytical arguments and simulation evidence that when researchers incorrectly theorize interaction due to compression, models without a product term bias the researcher, sometimes heavily, toward finding interaction. However, simulation studies also show that models with a product term fit a broad range of non-interactive relationships surprisingly well, enabling analysts to remove most of the bias toward finding interaction by simply including a product term.

Many theories in political science suggest an “interactive” relationship in which the effect of one explanatory variable X depends on the value of another explanatory variable Z (Clark, Gilligan and Golder 2006; Berry, Golder and Milton 2012). Although earlier research lays out a clear and uncontroversial method for testing these claims in the context of *linear* regression (e.g., Friedrich 1982; Brambor, Clark and Golder 2006; Kam and Franzese 2007), political methodologists disagree about the best approach for researchers using *logistic* regression, particularly over whether a product term must be included in the model (Berry and Berry 1991; Frant 1991; Nagler 1991; Berry, DeMeritt and Esarey 2010).¹

The recent debate hinges around the substantive work of Wolfinger and Rosenstone (1980), who theorize that easing registration restrictions should have a smaller positive effect on the probability of voting among individuals with more education. They do not evaluate the presence of interaction by testing the statistical significance of the product term’s coefficient. Indeed, they do not even include a product term in the model. Instead, they simply note that, even without a product term, a logistic regression model can represent the kind of interaction suggested by their theory because the S -shaped response curve creates “compression,” requiring that all variables have smaller effects as the probability of an event approaches 0 or 1:

As a statistical model, [logistic regression] is a more faithful representation of our substantive theory than [linear regression]. As we will see, the impact of most demographic variables is not constant across all types of individuals. Rather, the effect of a variable depends on the probability that the individual would vote. For example, a high-status occupation or a high income has less impact on a

* Carlisle Rainey is Assistant Professor of Political Science in the Texas A&M University, 2010 Allen Building, College Station, TX 77843 (crainey@tamu.edu). The author thanks Kenneth Benoit, Bill Berry, Scott Clifford, Justin Esarey, and two anonymous reviewers for helpful comments on earlier versions of this manuscript. The author also thanks John Oneal and Bruce Russett for making their data available, the Center for Computational Research at the University at Buffalo for providing support for the simulations. Code and data necessary to replicate the simulations and empirical analysis is available at <http://www.carlislerainey.com/research> and at <http://thedata.harvard.edu/dvn/dv/PSRM>. To view supplementary material for this article, please visit <http://10.1017/psrm.2015.59>

¹ Throughout the manuscript, I focus on logistic regression models in particular, but the conclusions generalize to a wide range of models of binary outcomes that rely on S -shaped link functions, including probit models.

college graduate who is 90 percent likely to vote, than it has on a high school dropout, who is only 55 percent likely to vote... With [logistic regression] a variable has very little impact on those who are either very unlikely or nearly certain to vote. It has the greatest impact in the middle of the distribution, on those who are between 40 and 60 percent likely to vote and are the most susceptible to the forces pushing them to vote or not to vote (Wolfinger and Rosenstone 1980, 11).

Instead of including a product term and testing the statistical significance of its coefficient, Wolfinger and Rosenstone (1980) argue for interaction by calculating a second-difference, which is the difference in the effect of easing registration restrictions between individuals with a high and low level of education. Formally, a second-difference $\Delta\Delta$ is the difference in the effect on $\Pr(Y)$ of changing X from a low value to a high value as Z changes from a low value to a high value, which can be defined mathematically as follows:

$$\Delta\Delta = [\Pr(Y|X=x_{hi}, Z=z_{hi}) - \Pr(Y|X=x_{lo}, Z=z_{hi})] \\ - [\Pr(Y|X=x_{hi}, Z=z_{lo}) - \Pr(Y|X=x_{lo}, Z=z_{lo})],$$

where x_{lo} , x_{hi} , z_{lo} , and z_{hi} take on substantively interesting values of the explanatory variables X and Z .² Using this approach, they find, as expected, that easing registration restrictions has a larger effect on the probability of turning out among less-educated individuals.

However, Nagler (1991, see also 1994) rejects this finding because the S-shaped curve automatically builds interaction into the model. He claims that “examining predicted probabilities generated by non-linear models such as [logistic regression] may produce spurious results when used to determine interactive effects between two [explanatory] variables” (Nagler 1991, 1393). Instead of just calculating a second-difference to quantify the amount of interaction, Nagler prefers to argue for interaction by including a product term and testing the statistical significance of its coefficient. When he does this, the coefficient for the product of education and registration restrictions is incorrectly signed (suggesting that easing restrictions has a larger effect on the propensity to vote among the *most* educated), leading Nagler to conclude that the data do not support Wolfinger and Rosenstone’s theory.

But Berry, DeMeritt and Esarey (2010) defend Wolfinger and Rosenstone’s approach against Nagler’s challenge, arguing that “compression should not be viewed as theoretically irrelevant.” Instead, they suggest that “it can be a strong theoretical rationale for expecting interaction between variables in their influence on [the probability of an event]” (Berry, DeMeritt and Esarey 2010, 255). Further, Berry, DeMeritt and Esarey suggest that if researchers expect interaction due to compression alone, then “there is no need to include a product term in the model” (Berry, DeMeritt and Esarey 2010, 261). In this situation, Berry, DeMeritt and Esarey recommend following Wolfinger and Rosenstone’s strategy of (1) excluding a product term, (2) computing a second-difference, and (3) calculating a confidence interval and verifying that it contains only values consistent with the research hypothesis (i.e., is statistically significant).³

² For more details, see Berry, DeMeritt and Esarey (2010). For a discussion of using simulation to calculate confidence intervals for second-differences, see King, Tomz and Wittenberg (2000). See Hanmer and Kalkan (2013) for an alternative.

³ Most of the debate about testing for interaction between two variables in influencing the probability of an event has occurred in political science. An earlier, but less detailed, debate between Berry and Berry (1990, 1991) and Frant (1991) almost exactly parallels the debate among Wolfinger and Rosenstone (1980), Nagler (1991), make an argument similar to Berry, DeMeritt and Esarey (2010). The substantive application focuses on lottery adoption in the US states, but Frant (1991) makes an argument similar to Nagler (1991) and Berry and Berry (1991) make an argument similar to Berry, DeMeritt and Esarey (2010). Also in political science, Tsai and Gill (2013) and Huang and Shields (2000) consider interaction in the context of logistic regression, but focus on

Despite compression's theoretical relevance, I argue that researchers should *always* include a product term, even when they expect interaction on the basis of compression alone. A product term is not necessary to allow the model to represent interaction—it can do that with no product term (Berry, DeMeritt and Esarey 2010). A product term must be included because it allows the model to better represent a *non-interactive* relationship. I present an analytical argument and simulations showing that if no product term is included, a logistic regression model cannot represent a non-interactive relationship and thus overstates the evidence for interaction. Fortunately, I find that simply including a product term eliminates most of this bias.

CURRENT ADVICE TO APPLIED RESEARCHERS

Before the publication of Berry, DeMeritt and Esarey (2010), the literature suggested that when analysts faced a dichotomous outcome and suspected interaction, there was only one path forward: include a product term and conclude that there is interaction between X and Z if and only if the coefficient for the product term was statistically significant (Nagler 1991; Brambor, Clark and Golder 2006, though see Berry and Berry 1991). However, Berry, DeMeritt and Esarey (2010, 2015) have added much needed nuance to this advice. In particular, Berry, DeMeritt and Esarey (2010) lay out an initial framework describing the contexts in which researchers should hypothesize in terms of the latent variable Y^* or $\Pr(Y)$, describe the implications for whether one performs hypothesis tests on the coefficient of the product term or a second-difference, and make suggestions about when to include a product term. Table 1 synthesizes the advice from Nagler (1991), Berry, DeMeritt and Esarey (2010), and Berry, DeMeritt and Esarey (2015), summarizing the current advice to applied researchers.

I find the advice in Table 1 sound, useful, and large improvement over the coarse practice before Berry, DeMeritt and Esarey (2010). However, I disagree with the recommendation to exclude a product term when the analyst puts forward a strong theory predicting interaction on the basis of compression alone (the bold “No” in Table 1). In this paper, I explain why I find Berry, DeMeritt and Esarey's (2010) argument for excluding a product term incomplete.

(Footnote continued)

the calculation of marginal effects on the probability of an event without addressing the need for a product term or discussing contexts in which researchers should care about the probability of an event as opposed to the latent outcome. Although the most detailed discussions of interaction in the context of logistic regression have occurred in political science, especially relating quantities of interest and the need for a product term to substantive theories, similar discussions have emerged in related disciplines. Work in economics focuses on the relationship between the sign and significance of the product term in logit and probit models and the sign and significance of the second derivative $\frac{\partial^2 \Pr(Y)}{\partial X \partial Z}$ (Ai and Norton 2003; Norton, Wang and Ai 2004; Greene 2010). However, this work offers little guidance as to when applied researchers should focus on $\frac{\partial^2 \Pr(Y)}{\partial X \partial Z}$ as opposed to $\frac{\partial^2 Y^*}{\partial X \partial Z}$ and does not discuss the consequences of including or excluding a product term. Research in epidemiology works to define interaction in the context of logistic regression models, usually focusing on differences in risk ratios rather than second-differences or second-derivatives. Again, though, this literature offers little guidance about whether researchers should include or exclude a product term in logistic regression models used to test for interaction. Also outside of political science, Bowen (2012) draws a sharp distinction between the interaction allowed by the inclusion of a product term and the interaction induced by the link function and offers a statistical test to distinguish the two forms of interaction. Bowen does not address whether a product term is necessary to conclude that interaction is present and does not offer advice about the advantages of focusing on the probability of an event instead of the latent outcome variable.

TABLE 1 *Current Advice for Testing Interactive Theories Using Logistic Regression*

| Situation | Description | Include a Product Term? | Quantity of Interest | Source |
|---|---|-------------------------|---|-----------------------------------|
| Interaction in influencing the latent outcome | Guided by a strong theory, the analyst hypothesizes that X and Z interact in influencing the latent outcome variable Y^* . For example, it sometimes makes sense to conceptualize Y^* as utility and derive a probit model using a random utility framework (Train 2009). See Berry, DeMeritt and Esarey (2010, pp. 261–2) for more details and an example. | Yes | $\frac{\partial^2 Y^*}{\partial X \partial Z}$ | Nagler (1991) |
| Interaction due to compression alone | Guided by a strong theory, the analyst hypothesizes that X and Z interact in influencing $\Pr(Y)$ due to compression alone. That is, as the probability of an event approaches 0 or 1, the effect of any explanatory variable (including X and Z) have smaller effects. Although researchers such as Frant (1991) and Nagler (1991) sometimes dismiss compression as an unimportant form of interaction, Berry, DeMeritt and Esarey (2010) make a strong case that this type of interaction is often theoretically meaningful. | No | $\frac{\partial^2 \Pr(Y)}{\partial X \partial Z}$ | Berry, DeMeritt and Esarey (2010) |
| Specification ambiguity | Guided only by weak theoretical intuition, the analyst hypothesizes that X and Z interact in influencing $\Pr(Y)$, but have no strong theoretical rationale for the functional form. In this situation, the analyst lacks the theoretical guidance necessary to theorize about interaction in terms of the latent variable or on the basis of compression alone. | Yes | $\frac{\partial^2 \Pr(Y)}{\partial X \partial Z}$ | Berry, DeMeritt and Esarey (2015) |

Note: this table summarizes three situations in which analysts might theorize interaction in the context of logistic regression. Notice that the quantity of interest varies across these situations, as well as whether the literature recommends including a product term. In this paper, I address the situation in which an analyst suspects interaction due to compression alone. In particular, I agree with Berry, DeMeritt and Esarey (2010) characterization of the theoretical value of compression and the implied quantity of interest, but argue that analysts must still include a product term.

THE ARGUMENT FOR EXCLUDING A PRODUCT TERM

Berry, DeMeritt and Esarey's (2010) (hereafter, BDE) overarching claim, which I wholeheartedly agree with, is that "a statistically significant product term is neither necessary nor sufficient for concluding that there is substantively meaningful interaction among independent variables in their influence on [the probability of an event]" (2010, 261). Given this point though, BDE ask whether it is always necessary to include a product term in the model. To decide whether to include a product term, they suggest that analysts draw a sharp distinction between theories that focus on the unbounded latent variable, denoted as Y^* (perhaps conceptualized as "utility"), and theories that focus on the probability of an event, denoted as $\Pr(Y)$.⁴

BDE Recommendation: When guided by a strong theory, the decision to include a product term must be based on the theorized effects of variables on the unbounded latent dependent variable Y^* , assumed by the model.

Part (a): If the theory assumes that variables interact in influencing the latent unbounded variable Y^* , then the researcher should include a product term (See the situation "Interaction in influencing the latent outcome" in Table 1).

We agree on *Part (a)*. As both BDE and Nagler (1991) note, this situation is entirely analogous to linear regression (e.g., $Y^* = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \dots + \beta_k X_k$), making a product term essential.

Part (b): If the theory assumes that variables should interact in influencing $\Pr(Y)$ strictly on an expectation of compression, then there is no need to include a product term in the model (See the situation "Interaction due to compression alone" in Table 1).

BDE explain:

[I]n any logit or probit model, the marginal effect of X on $\Pr(Y)$ depends on the values of all [explanatory] variables; this marginal effect is greatest when $\Pr(Y)$ is 0.5 and declines when a change in either variable pushes $\Pr(Y)$ toward 0 or 1. We refer to the phenomenon ... as compression, because deviations of $\Pr(Y)$ away from [0.5] compress further possible change in $\Pr(Y)$ to ever-smaller ranges. (Berry, DeMeritt and Esarey 2010, 251)

They continue by suggesting:

[A] researcher... may base his hypothesis that independent variables interact in influencing $\Pr(Y)$ strictly on an expectation of compression. In this case, there is no need to include a product term in the model. Put differently, no product term is required if the analyst's only reason for positing interaction between X and Z in influencing $\Pr(Y)$ is an expectation that when the value of X is extreme, $\Pr(Y)$ is near its limit of 0 or 1 and thus there is little room for $\Pr(Y)$ to change as Z changes. In any event, the analyst should use the parameter estimates for the model (with no product term) to generate estimates of the effects of variables on $\Pr(Y)$ to test his hypothesis; the same tools used for models with a product term—second-differences or marginal effect plots—can be employed. (Berry, DeMeritt and Esarey 2010, 261–2).

⁴ In BDE's notation, Y^* refers to the so-called "latent variable" implied in models of binary outcomes. Y^* simply equals the value of the linear predictor, often denoted by $\eta = X\beta = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \dots + \beta_k X_k$ in the notation of generalized linear models, and is transformed into $\Pr(Y)$ by an inverse link function that maps the real line to [0,1]. Sometimes researchers' theories lead to hypotheses about Y^* rather than $\Pr(Y)$. For example, theorists often link formal models of decision making to random utility models, a common framework for deriving models such as logistic regression, particularly the almost identical probit model. Using this approach, researchers can evaluate explanations of utility, represented by the value of $Y^* = \eta = X\beta$.

Although I find BDE's suggestion to exclude a product term problematic, I take two other points from their paper as a starting point.

1. Compression can be substantively interesting and theoretically meaningful. Consider the example from Wolfinger and Rosenstone (1980) discussed above. Recall that their model does not have a product term, the only interaction is due to compression. Yet if their model is correct, then restrictive voting rules not only reduce turnout, but also increase the inequality of political participation because restrictive rules have a larger negative effect on the probability of voting among the least educated. This relationship has enormous normative consequence and policy relevance.
2. If analysts are interested in interaction due to compression, then the quantity of interest is not the product term, but the second-difference (or second derivative). Several authors show that the sign and significance of the second-difference does not depend on the sign and significance of the product term (e.g., Ai and Norton 2003; Tsai and Gill 2013).

Although, I find BDE's argument for these two points compelling, I disagree with the advice to exclude a product term when theorizing interaction on the basis of compression alone.

THE ARGUMENT FOR INCLUDING A PRODUCT TERM

Researchers must employ a statistical model that can represent (1) relationships that are *consistent* with the theory and (2) plausible relationships that are *inconsistent* with the theory. It is obvious that models should be able to represent relationships that are consistent with theory. But if researchers use models that cannot also represent situations *inconsistent* with the theory, then they have assumed their theories are correct. Though the theoretical rationale may be strong, it has not been made vulnerable to the data.

Thus, BDE's analysis is incomplete when they write that "no product term is required if the analyst's only reason for positing interaction between X and Y in influencing $\Pr(Y)$ is an expectation that when the value of X is extreme, $\Pr(Y)$ is near its limit of 0 or 1, and thus there is little room for $\Pr(Y)$ to change as Z changes." Although they correctly note that the theorized relationship can be represented by a model with no product term, they fail to discuss whether this model can also represent relationships that are inconsistent with the theory.

Without a product term, a logistic regression model cannot represent many reasonable situations that are inconsistent with an interactive theory. In particular, it cannot represent a situation in which two explanatory variables have non-zero effects and the probability of an event is greater or less than one half, but there is no interaction. Specifically, for a logistic regression model including X and Z , but excluding their product, if all of the following conditions hold, then X and Z must interact in influencing $\Pr(Y)$.⁵

1. X has a non-zero effect.
2. Z has a non-zero effect.
3. $\Pr(Y) < 0.5$ or $\Pr(Y) > 0.5$.

This can easily be seen by inspecting the second derivative of the logistical regression model with no product term, given by

$$\frac{\partial^2 \Pr(Y)}{\partial X \partial Z} = \Pr(Y)[1 - \Pr(Y)][1 - 2\Pr(Y)]\beta_x \beta_z. \quad (1)$$

⁵ For further discussion, see Nagler (1991) and Berry, DeMeritt and Esarey (2010).

If any of Conditions 1–3 hold, then the second derivative is exactly 0. Otherwise, the second derivative is non-zero, which implies an interactive relationship model. Consequentially, if the researcher incorrectly theorizes compression, but 1–3 hold, then a logistic regression model without a product term will overstate the evidence for interaction.⁶

With a product term, however, a logistic regression model can represent *some* plausible situations inconsistent with an interactive theory. For a logistic regression model including X , Z , and XZ , if all of the conditions (1–3) identified above hold, then X and Z *may or may not* interact in influencing $\text{Pr}(Y)$. The intuition of this claim can be seen by simply inspecting the second derivative of a logistic regression model including a product term, given by

$$\begin{aligned} \frac{\partial^2 \text{Pr}(Y)}{\partial X \partial Z} &= \text{Pr}(Y)[1 - \text{Pr}(Y)]\beta_{xz} \\ &+ \text{Pr}(Y)[1 - \text{Pr}(Y)][1 - 2\text{Pr}(Y)](\beta_x + \beta_{xz}Z)(\beta_z + \beta_{xz}X). \end{aligned} \tag{2}$$

The presence of the product term creates another coefficient that allows the second derivative to be zero, even if Conditions 1–3 hold. Thus, this model is capable of representing situations in which 1–3 hold, but no interaction is present. This means that a product term “uncouples” hypotheses about the effects of X and Z from their interaction.

Although including a product term does not entirely remove the bias toward confirming interactive hypotheses, it does greatly reduce the bias. Thus, including a product term makes the empirical argument more compelling by making the theory more vulnerable to the data.⁷

⁶ Responding to Frant’s (1991) suggestion to include a product term, Berry and Berry write:

Frant claims that this result [that the effect of election proximity on the probability of lottery adoption depends on the state’s fiscal health] “sounds like a *conclusion* (and the authors’ treat it as such); but it is actually an *assumption* of the model.” It is true that interaction among variables in their effects on the probability of adoption is an assumption of our probit model. But all multivariate models make assumptions about the function form of effects of independent variables; yet substantive conclusions are drawn from empirical analyses of these models. Indeed, an interactive regression with a multiplicative term (e.g., $E(Y) = \beta_0 + \beta_1X_1 + \beta_2X_2 + \beta_3X_1X_2$)—similar to the form that Frant proposes we use—also *assumes* interaction among the independent variables (provided $\beta_3 \neq 0$) (1991, 574).

My analytical argument supports Frant’s (1991) position. Although Berry and Berry (1991) are correct that the linear model including a product term assumes a form of interaction (e.g., the marginal effect of X_1 is a *linear* function of X_2), it does not assume that interaction is present (hence their caveat that $\beta_3 \neq 0$). A linear model with a product term can represent situations that are consistent and inconsistent with the interactive theory ($\beta_3 \neq 0$ and $\beta_3 = 0$, respectively). This allows the model to effectively assess the evidence against the null hypothesis and for the research hypothesis. As the analytical argument shows, the logistic regression model with no product term, whereas it can represent an interactive relationship, cannot represent a non-interactive relationship. Both my analytical arguments and simulations strongly suggest that Frant (1991) is correct when he advises researchers to avoid making empirical claims about interaction between variables in influencing the probability of an event when no product term is included in the model.

⁷ This analytical argument provides a formal justification for Huang and Shield’s claim that “[o]f course, adding multiplicative terms to explicitly model this interactive relationship is necessary” (2000, 86).

SIMULATIONS

To illustrate the potential consequences of excluding a product term, I report two simulation studies. The first examines a single relationship in detail. The second looks at a wide range of potential relationships. Because I am particularly interested in what happens if the theory (or researcher's hypothesis) is wrong, I imagine that a researcher incorrectly theorizes interaction due to compression and follows BDE's suggestion to exclude the product term. In almost all cases, excluding the product term biases the researcher toward finding support for her theory.

But why focus specifically on what happens if the researcher is wrong? Why not look at what happens if the researcher correctly theorizes interaction on the basis of compression alone? Careful readers recognize the tradeoff between the size of a hypothesis test (i.e., the probability of finding interaction when there is none) and power (i.e., the probability of finding interaction when it is present). Although analysts usually face a tradeoff between these two, I focus my attention squarely on the size of tests for interaction.

Hypothesis test form compelling statistical arguments because, by design, they rarely lead a researcher to reject the null hypothesis when the null hypothesis is true. This is because hypothesis tests are constructed by fixing $\alpha = \Pr(\text{reject null} \mid \text{null is true})$ at an acceptably low value, typically $\alpha = 0.05$. This is important because researchers and readers interpret a decision to reject the null hypothesis as strong evidence against the null hypothesis (or for the research hypothesis).

Social scientists typically use hypothesis tests because they want to conclude that their research hypothesis is correct only when the data offer strong evidence in favor of their research hypothesis (or equivalently, against their null hypothesis). If a test is improperly sized (e.g., $\alpha = 0.75$), then the decision to reject the null hypothesis is no longer compelling. Making matters worse, readers and reviewers are still likely to interpret the results of this over-sized test as compelling. I focus squarely on the size of tests for interaction because the logic of hypothesis testing requires that the size be fixed at a suitably small value, usually 0.05 in the social sciences, whereas allowing the power of a test to fluctuate with other variables, such as the number of observations. My primary focus on size, or Type I or false-positive errors, should not be taken as meaning that power is unimportant. To address the secondary issue of power, I include a discussion of the statistical power of models that exclude and include a product term in the online Appendix. These simulations show that the ability of a model with a product term to "detect" an interactive relationship (i.e., statistical power) is nearly equal if the amount of interaction is substantial or the sample size is large. However, it is important to point out that when the sample size is small (e.g., <1000) and the amount of interaction is small (e.g., second-difference of $0.04 - 0.09 = -0.05$), then including a product term might lead to a substantial drop in statistical power. However, it is important to note that the gain in statistical power comes from *assuming* an interactive relationship, not *detecting* one.

One might also worry about the impact of including a potentially unnecessary product term XZ in the model (i.e., $\beta_{XZ} = 0$) on the estimated effects of other variables in the model. In the online Appendix, I show that including an unnecessary product term does not lead to bias in the estimated effects of other variables in the model.

STUDY NO. 1: A FIXED RELATIONSHIP

I first consider the non-interactive relationship (i.e., data-generating process) $\Pr(Y) = 0.3 + 0.2X - 0.3Z$ shown in the left panel of Figure 1. I imagine that the analyst correctly suspects that X has a positive effect on $\Pr(Y)$ and Z has a negative effect on $\Pr(Y)$.

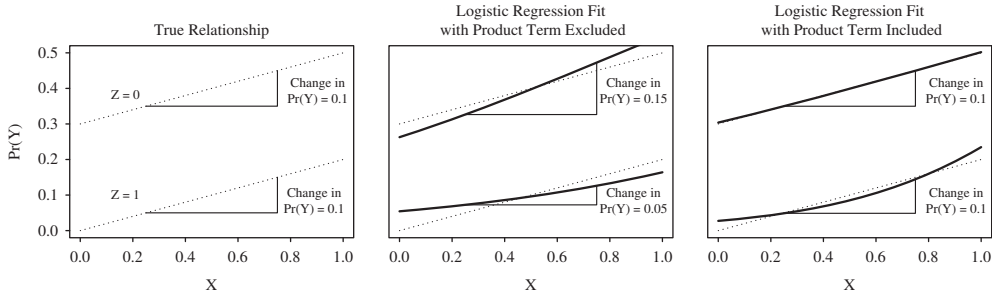


Fig. 1. Logistic regression fits to a non-interactive relationship

Note: this figure shows the true relationship in the first simulation study (dotted lines) as well as how logistic regression models with and without a product term fit the process (solid, bold lines). Notice that X and Z have large effects on $\Pr(Y)$ and $\Pr(Y)$ is always < 0.5 . In this case, excluding a product term forces the model to point toward interaction, despite the fact that the relationship is non-interactive. Including a product term allows the model to fit this non-interactive process surprisingly well.

She also theorizes that because the event is relatively rare and Z pushes $\Pr(Y)$ even closer to 0, there is less room for X to affect $\Pr(Y)$. Therefore, she expects that X should have a smaller effect when $Z = 1$. But this hypothesis is not correct—she *incorrectly* expects that X and Z interact in influencing $\Pr(Y)$ due to compression. Still, she notices that a logistic regression model can represent the kind of interaction suggested by her theory and decides that no product term is necessary (Berry, DeMeritt and Esarey 2010).

Because X is approximately uniformly distributed from 0 to 1 and Z is dichotomous with half of the observations equal to 1 and the other half equal to 0, she decides to test for interaction by estimating a logistic regression model with no product term, calculating the second-difference as X changes from 0.25 to 0.75 and Z changes from 0 to 1. Her hypothesis suggests that this second-difference should be negative, as X should have a smaller effect when $Z = 1$.

Notice the middle panel of Figure 1, which shows how a logistic regression model without a product term might fit the true relationship. Although there is no interaction in the actual data-generating process, the logistic regression model with no product term suggests that there is substantial interaction between X and Z in influencing $\Pr(Y)$. When $Z = 1$, the estimated effect of moving X from 0.25 to 0.75 is 0.05. When $Z = 0$, however, the effect of X triples to 0.15 (the second-difference is $0.05 - 0.15 = -0.1$). Although the hypothesis of interaction is wrong, the logistic regression model with no product term suggests there is strong evidence in favor of it.

Now notice the right panel of Figure 1, which shows the estimates when the model contains a product term. Although this model does not exactly capture the true data-generating process (i.e., notice the non-linearity when $Z = 1$), it estimates the amount of interaction very well. The estimated effect of X when $Z = 0$ is 0.1. When $Z = 1$, the estimate is identical. Therefore, the estimated second-difference is 0 ($0.1 - 0.1 = 0$) and the model correctly suggests that there is no interaction.

Rather than examining how models fit the data-generating process, it is perhaps more informative to evaluate each approach in repeated trials. To do this, I simulate 2,000 data sets containing 1,000 observations each. I use each data set to estimate a second-difference and its confidence interval using a model with and without a product term. Finally, I estimate $\alpha = \Pr(\text{Find Interaction})$, the probability of finding interaction *when none is present* (i.e., a Type I or false-positive error), by calculating the proportion of data sets that pointed toward interaction with each approach (i.e., the confidence interval does not contain 0). By convention, this probability should be close to 0.05. When the model does not include a product term, the probability of finding interaction is 0.98. Thus, finding a statistically significant second-difference is not at all a compelling evidence against the null when the

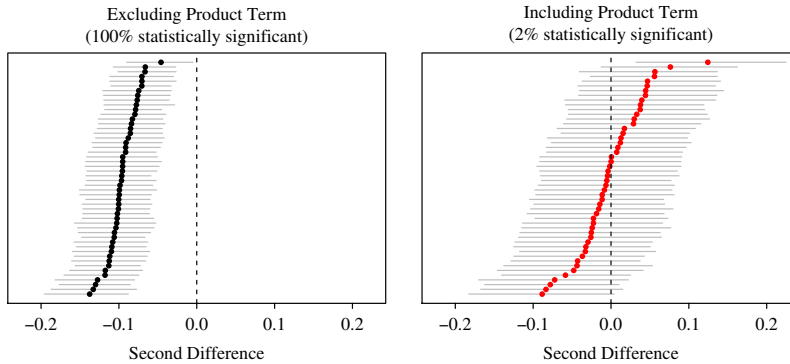


Fig. 2. Fifty estimated second differences and confidence intervals

Note: this figure shows 50 simulations of estimated second-differences and confidence intervals using models with and without a product term. Although second-difference is actually 0 (a non-interactive relationship), the model with no product term consistently finds interaction. Including a product term removes almost all of this bias.

model does not include a product term, as the researcher will likely find this even when there is no interaction. However, the probability of finding interaction drops to 0.06 when the researcher includes a product term in the model. This suggests that finding a statistically significant second-difference is compelling when the model includes a product term, as the researcher would rarely find this when the null is true.

To help understand why excluding a product term biases the researcher toward finding interaction, Figure 2 shows the estimates and confidence intervals from only 50 simulated data sets, sorted by the size of the estimated second-difference. Notice the left panel of Figure 2, which shows that the models without a product term always estimate a substantial amount of interaction in the expected direction and that the confidence intervals rarely overlap 0. *This happens because compression serves as the theoretical motivation for the hypothesis and is assumed by the model.* But notice the right panel of Figure 2, which shows that most of this bias disappears when a product term is included in the model.

To get a sense of the robustness of this conclusion, I systematically vary the sample size, the distribution of X , and the change in X used to calculate the second-difference while the true relationship remains the same. For each combination of the simulation parameters, I simulate 2000 data sets. For each data set, I estimate a logistic regression model with and without a product term and compute the confidence interval around the second-difference, as BDE recommend. I then calculate the proportion of simulations in which each model points incorrectly toward interaction to estimate the probability of finding interaction when none is present. The results are presented in Figure 3.

Across all values of the simulation parameters, the models without a product term bias the researcher toward finding interaction and sometimes strongly. The bias is reduced, usually dramatically, when the researcher includes a product term. Indeed, for symmetric distributions of X and centered second-differences, adding a product term to the model eliminates nearly all of the bias, lowering the probability of finding interaction from nearly 1.00 (the worst situation) to about 0.05 (the ideal situation).⁸

⁸ As a robustness check, I repeated this study using a slight different data-generating process. The full results are discussed in the online Appendix, but the basic findings suggest an even more optimistic conclusion about the ability of the product term to remove the bias toward finding interaction. This is as expected as I chose the fixed process discussed above to serve as a worst case scenario.

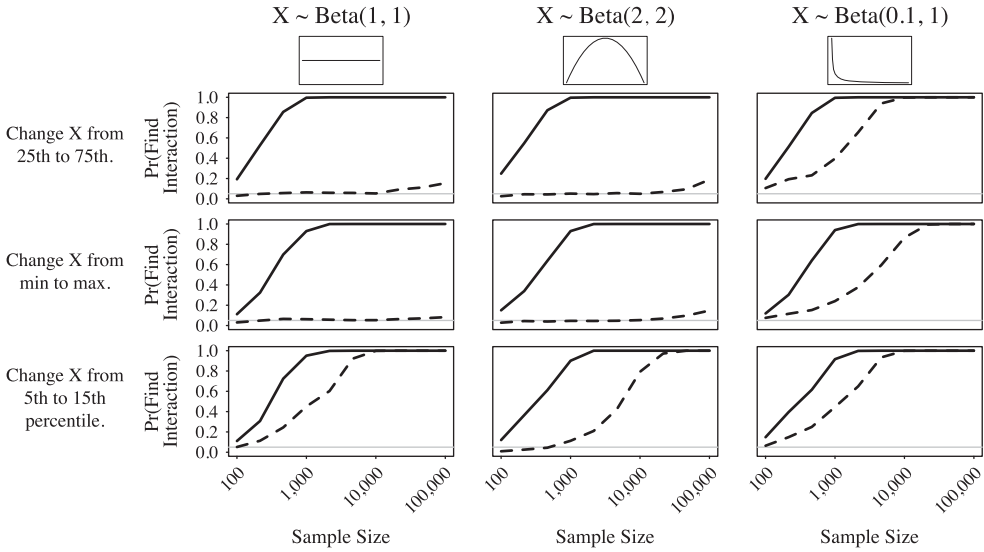


Fig. 3. The size of tests for interaction

Note: this figure shows the probability of concluding interaction when the true relationship is given by $\Pr(Y) = 0.3 + 0.2X - 0.3Z$ (i.e., there is no interaction) using logistic regression models with and without a product term as the sample size, the distribution of X , and the change in X considered in computing the second-difference varies. Notice that the model including a product term performs remarkably well, although the improvement diminishes as the simple size gets large, especially when the distribution of X is skewed and the researcher considers the effect of changing X nears its extremes.

STUDY NO. 2: A DIVERSE SET OF RELATIONSHIPS

To demonstrate that the patterns seen in Figure 3 do not depend on any particular relationship, I now examine a diverse collection of 1000 simulated analyses. Each analysis is generated through a random process (described in the online Appendix) that varies in the following ways:

1. The relationship between $\Pr(Y)$, X , and Z varies, but follows the form of $\Pr(Y) = \beta_0 + \beta_1X + \beta_2Z + \beta_3X^2 + \beta_4Z^2$ and is always monotonic across the range of X and Z . Importantly, X and Z never interact in changing $\Pr(Y)$.⁹
2. The distributions of X and Z vary and might be binary, flat, unimodal, or skewed to the right.
3. The values of x_{lo} , x_{hi} , z_{lo} , and z_{hi} used to construct the second-difference vary from 0 to 1. All combinations are possible, though if a variable is binary, its values are set to 0 and 1.
4. The sample size varies, ranging from 500 to 100,000, though smaller samples are more likely.

Each “analysis” consists of features the researcher totally controls (e.g., values of x_{lo} , x_{hi} , z_{lo} , and z_{hi}), somewhat controls (e.g., sample size), and does not control (e.g., distributions of explanatory variables and the true data-generating process). For each simulated analysis, I use the same procedure

⁹ I use a quadratic functional form here to include a more general set of curves rather than restrict my processes to straight lines. The only condition that I want to place on the processes is that they be monotonic (i.e., the marginal effects never change signs) and there be no interaction. The quadratic form allows a nice mixture of different relationships within these restrictions (though not exhaustive, of course). The restriction disallowing interaction is necessary because I am interested mainly in the probability of finding interaction when there is none present.

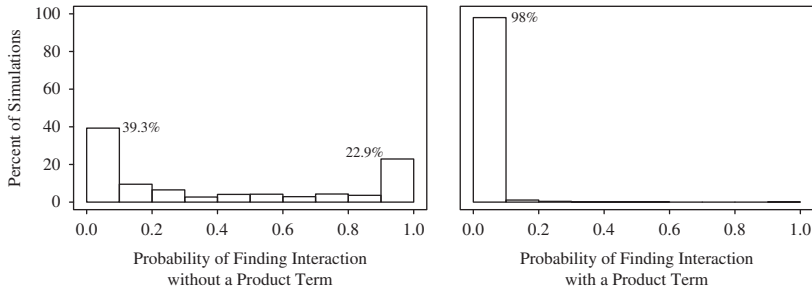


Fig. 4. The distribution of the size of tests for interaction across a diverse set of analyses
 Note: the histogram on the left shows that logistic regression models without a product term tend to find interaction when none exists. The histogram on the right shows researchers can eliminate most of this bias by simply including a product term.

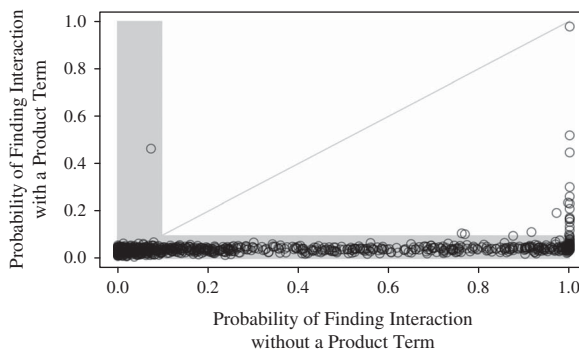


Fig. 5. The relationship between the size of the test for interaction with and without a product term across a diverse set of analyses

Note: this figure shows the probability of concluding interaction when none exists from logistic regression models with and without a product term for many different analyses. Points that fall in the vertical band indicate an analysis in which the model *without* a product term performs well. Points that fall in the horizontal band indicate an analysis in which the model *with* an product term performs well. For many analyses, the model without the product term performs quite poorly, whereas the model with a product term performs well. Finally, notice that almost all points fall below the diagonal line indicating that models without a product term are more likely to mislead researchers in almost every simulated analysis.

as before to estimate the probability of finding interaction using models with and without a product term. Because none of the true relationships are interactive, this probability should be near 0.05.

The left panel of Figure 4 shows that when the researcher uses a logistic regression model without a product term, the probability of finding interaction when none exists is less than 0.1 for only about 40 percent of the simulated analyses. It is >0.9 for about 20 percent of the analyses. This suggests that excluding a product term strongly biases the researcher toward finding interaction for many plausible relationships. On the other hand, the right panel of Figure 4 shows that, when using a model with a product term, the probability of finding interaction is less than 0.1 for almost all of the simulated analyses (over 98 percent). This suggests that including a product term can eliminate most of the bias toward finding interaction for many plausible relationships and analyses.

Figure 5 directly compares the two approaches. Except for a single outlier, notice that every point falls below the diagonal line, except when both approaches fall below 0.1 (i.e., both perform well). This means that in almost every situation in which excluding a product term led to substantial bias, the model with a product term reduced this bias. In only one case did the model without a product

term perform substantially better than the model with a product term. Regardless of the absolute performance of each, the model with a product term performs at least as well as the model without a product term in almost all the analysis generated in my simulation, except when both perform well. In cases in which the model with a product term performs somewhat poorly (e.g., the probability of finding interaction is about 0.3, for example), the model without a product term performs terribly (e.g., the probability of finding interaction is nearly 1.00).

RECOMMENDATIONS FOR APPLIED RESEARCHERS

I suggest that applied researchers take several steps to ensure their claims of interaction are empirically meaningful.

1. Clearly explicate your interactive theory (Clark, Gilligan and Golder 2006; Berry, Golder and Milton 2012) and carefully consider which relationships are consistent and inconsistent with the theory. Only use empirical models that can represent both relationships.
2. As Berry, DeMeritt and Esarey (2010) suggest, specify whether the relevant outcome variable is the latent variable Y^* or the expected value $\Pr(Y)$. If you use the concept Y^* , then test for interaction by testing the statistical significance of the product term. If you rely on $\Pr(Y)$, then test for interaction by simulating confidence intervals around a carefully chosen second-difference or second derivative. See Berry, DeMeritt and Esarey (2010) for the details of calculating a second derivative or second-difference and King, Tomz and Wittenberg (2000) for the details of the simulation procedure. Note that Hanmer and Kalkan (2013) offer an alternative approach in which researchers average the quantity of interest across the observed data rather than choosing a specific comparison of interest.
3. Regardless of whether the theoretically relevant outcome is Y^* or $\Pr(Y)$, and even if you theorize interaction on the basis of compression alone, you must include product terms in order to draw compelling substantive conclusions about interaction from logistic regression models. Otherwise, you are simply pulling assumptions through data and treating them as conclusions.
4. As Nagler (1991) and Frant (1991) note, if you do not include a product term, then avoid making claims about interaction. The interaction is assumed and you should understand and clearly identify it as such.¹⁰

DEMOCRACY, DISTANCE, AND THE PROBABILITY OF CONFLICT

To illustrate how to conduct a compelling test for interaction due to compression, I replicate analyses by Oneal and Russett (2001).¹¹ Berry, DeMeritt and Esarey (2010) explain that the probability of war between two democracies is nearly 0, but much higher if one or both countries is non-democratic. Thus, on the basis of compression alone, Berry, DeMeritt and

¹⁰ Note that while researchers might suspect interaction between two variables (perhaps these are control variables), it is not necessary to include the product term as long as no empirical claim is made about the interaction of these variables. For example, it is not necessary for future studies of conflict to include the product of democracy and distance in their logistic regression models unless the researchers are specifically interested in making compelling empirical arguments that the two interact in influencing the probability of conflict.

¹¹ Although Oneal and Russett estimate several different model using generalized estimating equations (GEE), I estimate only their simplest logistic regression model (2001, 314, table A3.1) using maximum likelihood and the usual standard errors. Using their more complex approach does not substantively alter the results, though it does increase the standard errors slightly.

Esarey (2010) argue that one can meaningfully hypothesize that explanatory variables should have smaller effects on the probability of conflict when both countries are democratic, simply because this probability is so close to 0. Using the example of geographic distance, the authors note that we should expect distance to have a much smaller effect on the probability of conflict when both countries are democracies than when at least one country is not democratic. Importantly, we have no reason to expect that democracy and distance interact in influencing the latent outcome (e.g., the utility of conflict). Thus, we have no reason to theorize about the sign of the product term. Further, the only reason to include a product term is to allow the model to represent a situation in which we are wrong (i.e., there is no interaction between democracy and distance in influencing the probability of conflict).

This leads to three hypotheses about the effects of distance and democracy on the probability of conflict. First, as has been commonly suggested (e.g., Ray 1998; Tomz and Weeks 2013), democracies are much less likely to fight each other.

DEMOCRACY HYPOTHESIS: Regardless of the distance between the two countries in a dyad, as the joint level of democracy in the dyad increases, the probability of conflict in that dyad decreases.

Second, other literature in international relations points out that states that are located geographically further from one another are less likely to engage in conflict (e.g., Diehl 1991; Bremer 1992).

DISTANCE HYPOTHESIS: Regardless of the level of democracy in the dyad, as the distance between the countries in the dyad increases, the probability of conflict in that dyad decreases.

As Berry, DeMeritt and Esarey (2010) point out, these two hypotheses imply a third hypothesis. Because democratic dyads are much less likely to engage in conflict, we should expect the negative effect of distance to be much smaller in these dyads simply due to compression.

INTERACTION HYPOTHESIS: The negative effect of distance on the probability of conflict is larger (more negative) in non-democratic dyads than in democratic dyads.

RESULTS

To test these three hypotheses, I use data from Oneal and Russett (2001). Their logistic regression model of conflict (militarized interstate disputes) includes variables for democracy (smaller of the two Policy IV scores in each dyad), distance (the number of miles between the capitals), and several control variables. I estimated one model using Oneal and Russett's exact specification (with no product terms) and another model that includes the product of democracy and distance.¹² See the online Appendix for the details and the coefficient estimates. The estimates from models with and without a product term are presented in the online Appendix. As each of the hypotheses are stated in terms of the probability of conflict, I compute and plot

¹² Oneal and Russett use GEE to estimate their models, but I rely on the simpler approach of maximum likelihood. The two methods produce substantively similar results, but the standard errors from the GEE approach are slightly larger.

estimates and 90 percent confidence intervals for predicted probabilities, first-differences, and second-differences using Clarify-like simulation (King, Tomz and Whittenburg 2000).

Figure 6 provides a nice overview of the intuition and the evidence for the hypotheses by plotting the estimated probability of conflict in both democracies and non-democracies from each empirical model.¹³ Notice that the results are consistent with all three hypotheses. First, increasing the distance between the states in a dyad reduces the probability of conflict in that dyad in both democratic (smallest Polity IV score is 10) and non-democratic (smallest Polity IV score is -10) dyads. Second, notice that the probability of conflict is lower in more democratic dyads, regardless of the distance between the states in the dyad. Finally, notice that the pacifying effect of distance is much smaller (i.e., less negative) in democratic dyads.

Although the pattern of the predicted probabilities is consistent with the three hypotheses, I now compute and plot the quantities of direct substantive interest and the confidence intervals. First, Figure 7 plots the effect of moving from a non-democratic dyad (lower Polity IV score is -10) to a democratic dyad (lower Polity IV score is 10) on the probability of conflict as distance varies. The Democracy Hypothesis suggests that this effect should be negative regardless of the distance between the states. Figure 7 shows that the expected negative effect holds for across all distances and the confidence interval never overlaps 0. This pattern holds across the models with and without a product term.

Next, Figure 8 plots the effect of moving from a “nearby dyad” (about 850 miles apart, which is the 25th percentile in the data set) to a “distant dyad” (about 8600 miles apart, which is the 75th percentile) as the democracy of the dyad varies. The Distance Hypothesis predicts that this effect should be negative regardless of the level of democracy in the dyad. Figure 8 shows that the effect is indeed negative across the range of observed levels of democracy and the confidence interval never overlaps 0. Again, this pattern holds across the models with and without a product term.

Finally, and most importantly for my purposes, Figure 9 shows the estimated second-differences—the difference in the effects of distance among democratic and non-democratic states. The negative estimates suggest that distance has a smaller effect among democratic states, which supports the Interaction Hypothesis. Notice that, as before, the results are substantively similar across the model with and without a product term. Overall, then it seems that both models provide strong support for all three hypotheses.

IS PARSIMONY PREFERABLE?

When models with and without a product term provide similar (or nearly identical) conclusions, applied researchers might be tempted to discard the model with a product term for three reasons.

1. The logistic regression model with no product term captures the relationship (democracy and distance have negative effects and interact due to compression) suggested by the strong theory. Berry, DeMeritt and Esarey (2010) recommend excluding the product term in this situation.¹⁴

¹³ Notice that the estimated relationships are similar across both modeling strategies. However, I have shown above that the model with no product term forces an interactive relationship into the results if both variables have an effect and the relevant probabilities are above or below one half.

¹⁴ Berry, DeMeritt and Esarey (2010) write: “[A] researcher may have no reason to believe that the effects of independent variables on the latent variable Y^* are interactive and may base his hypothesis that independent variables interact in influencing $\Pr(Y)$ strictly on an expectation of compression. In this case, there is no need to include a product term in the model.”

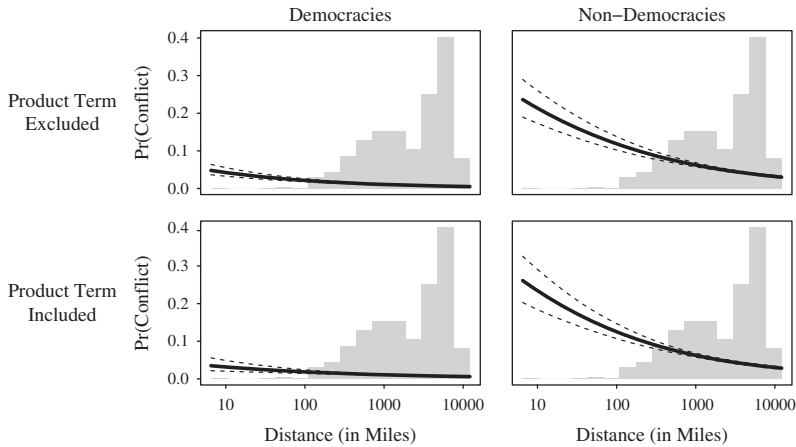


Fig. 6. The estimated probability of conflict as distance varies in democracies and non-democracies using a model with and without a product term

Note: this figure compares the estimated probability of conflict in democracies (Polity IV score of 10) and non-democracies (Polity IV score of -10) from models with and without a product term. Notice that, as expected, there is a much smaller change in the probability of conflict in democracies. However, the model with no product term forces this pattern of compression into the estimates. The model including a product term allows the model to represent non-interactive relationships as well. Because of this, the estimates from the model with a product term are more compelling.

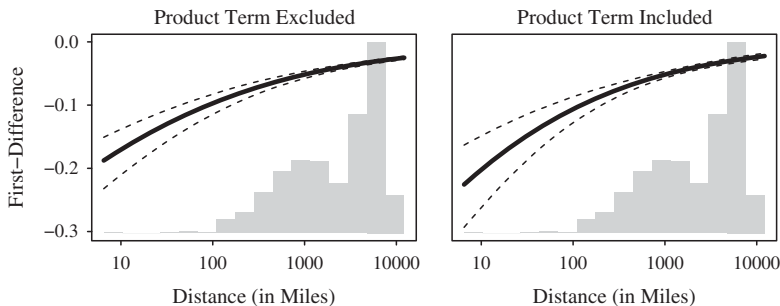


Fig. 7. The effect of democracy on the probability of conflict as distance varies

Note: this figure shows how the negative effect of democracy on the probability of conflict shrinks toward zero as dyads become more distant. This expected pattern is consistent with compression, but the model without a product term forces this pattern into the results while the model with a product term allows other types of relationships to emerge.

2. The AIC suggests the two models fit equally well (AIC = 13,520 for both models) and the BIC points toward the simpler with no product term (BIC = 13,580 for the model without the product term and BIC = 13,589 for the model with the product term).
3. The product term is not statistically significant ($p = 0.23$).

Researchers might be tempted to drop the product term for *any* of these three reasons. In spite of these temptations, *parsimony is not preferable*. Researchers must include the product term XZ when making arguments about X and Z interacting in influencing $\text{Pr}(Y)$. Although the models might suggest similar conclusions, they do not offer similar *evidence* for those

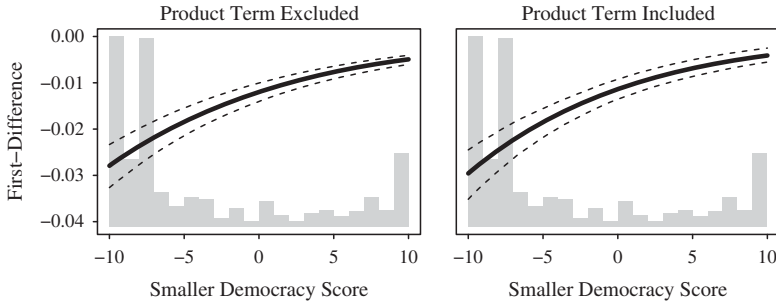


Fig. 8. The effect of distance on the probability of conflict as democracy varies
 Note: this figure shows how the negative effect of distance on the probability of conflict shrinks toward 0 as dyads become more democratic. This expected pattern is consistent with compression, but the model without a product term forces this pattern into the results, whereas the model with a product term allows other types of relationships to emerge.

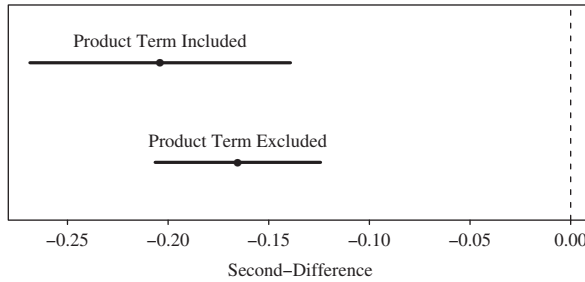


Fig. 9. The change in the estimated effect of distance from democratic to non-democratic dyads
 Note: this figure shows the estimated second-differences from both the models with and without a product term. Although the estimates are similar, the results from the model with a product term are more compelling, because excluding a product term strongly biases the second-differences away from 0 (see the left panel of Figure 2), whereas including removes most of this bias (see the right panel of Figure 2).

conclusions. A model without a product term does not offer compelling evidence for interaction because it suggests interaction *in spite* of the data, whereas a model with a product term suggests interaction *because* of the data.

In the applied example, it is important to include the product term because it “uncouples” the Interaction Hypothesis from the Distance and Democracy Hypotheses. If I do not include a product term in the model, then the large negative effect of democracy and distance guarantee that the two interact in influencing $\text{Pr}(Y)$ (see the left panel of Figure 2 and Equation 1). With no product term in the model, the Distance and Democracy Hypotheses *imply* the Interaction Hypothesis and nothing is learned from examining the second-difference. By including the product term, however, I am able to uncouple the Democracy and Distance Hypotheses from the Interaction Hypotheses (see the right panel of Figure 2 and Equation 2). When the model includes a product term, the Distance and Democracy Hypotheses no longer imply the Interaction Hypothesis and something is learned from examining the second-difference. Although the two estimates of the second-difference presented in Figure 9 are similar, they represent different amounts of evidence for the Interaction Hypothesis. The estimate from the model with no product term is assumed by the structure of the empirical model. The estimate from the model with a product term, on the other hand, represents a meaningful relationship based on the data.

CONCLUSION

When researchers are interested in how explanatory variables interact in influencing the probability of an event, my study shows that (1) models *without* a product term point toward interaction very often when none exists, (2) models *with* a product term exhibit much less bias toward interaction, and (3) models *with* a product term almost always reduce and usually eliminate the bias found in models *without* a product term. Although logistic regression models require “compression” to keep the probability of an event bounded between 0 and 1, researchers can limit the impact of this requirement on tests for interaction by simply including a product term, which allows the models to represent a broad range of non-interactive relationships surprisingly well. Therefore, even when theorizing interaction due to compression alone, researchers should include a product term, making their theories more vulnerable to the data.

REFERENCES

- Ai, Chunrong, and Edward C. Norton. 2003. ‘Interaction Terms in Logit and Probit Models’. *Economics Letters* 80(1):123–29.
- Berry, Frances Stokes, and William D. Berry. 1990. ‘State Lottery Adoptions as Policy Innovations: An Event History Analysis’. *American Political Science Review* 84(2):395–415.
- Berry, Frances Stokes, and William D. Berry. 1991. ‘Specifying a Model of State Policy Innovation’. *American Political Science Review* 85(2):573–9.
- Berry, William, Jacqueline H. R. DeMeritt, and Justin Esarey. 2010. ‘Testing for Interaction in Binary Logit and Probit Models: Is a Product Term Essential?’. *American Journal of Political Science* 54:248–66.
- Berry, William, Jacqueline H. R. DeMeritt, and Justin Esarey. 2015. ‘Bias and Overconfidence in Parametric Models of Interactive Processes’. *American Journal of Political Science*.
- Berry, William, Matt Golder, and Daniel Milton. 2012. ‘Improving Tests of Theories Positing Interaction’. *Journal of Politics* 77(3):653–71.
- Bowen, Harry P. 2012. ‘Testing Moderating Hypotheses in Limited Dependent Variable and Other Nonlinear Models: Secondary Versus Total Interaction’. *Journal of Management* 38(3):860–89.
- Brambor, Thomas, William Roberts Clark, and Matt Golder. 2006. ‘Understanding Interaction Models: Improving Empirical Analyses’. *Political Analysis* 14:63–82.
- Bremer, Stuart A. 1992. ‘Dangerous Dyads: Conditions Affecting the Likelihood of Interstate War, 1816–1965’. *Journal of Conflict Resolution* 36(2):309–41.
- Clark, William Roberts, Michael J. Gilligan, and Matt Golder. 2006. ‘A Simple Multivariate Test for Asymmetric Hypotheses’. *Political Analysis* 14(3):311–31.
- Diehl, Paul F. 1991. ‘Geography and War: A Review and Assessment of the Empirical Literature’. *International Interactions* 17(1):11–27.
- Frant, Howard. 1991. ‘Specifying a Model of State Policy Innovation’. *American Political Science Review* 85(2):571–3.
- Friedrich, Robert. 1982. ‘In Defense of Multiplicative Terms in Multiple Regression Equations’. *American Journal of Political Science* 26:797–833.
- Greene, William. 2010. ‘Testing Hypotheses About Interaction Terms in Nonlinear Models’. *Economics Letters* 107(2):291–96.
- Hanmer, Michael, and Kerem Ozan Kalkan. 2013. ‘Behind the Curve: Clarifying the Best Approach to Calculating Predicted Probabilities and Marginal Effects from Limited Dependent Variable Models’. *American Journal of Political Science* 57(1):263–77.
- Huang, Chi, and Todd G. Shields. 2000. ‘Interpretation of Interaction Effects in Logit and Probit Analyses’. *American Politics Research* 28(1):80–95.
- Kam, Cindy D., and Robert Franzese. 2007. *Modeling and Interpreting Interactive Hypotheses in Regression Analysis*. Ann Arbor: University of Michigan Press.

- King, Gary, Michael Tomz, and Jason Wittenberg. 2000. 'Making the Most of Statistical Analyses: Improving Interpretation and Presentation'. *American Journal of Political Science* 44:341–55.
- Nagler, Jonathan. 1991. 'The Effect of Registration Laws and Education on U.S. Voter Turnout'. *American Journal of Political Science* 85:1393–405.
- Nagler, Jonathan. 1994. 'Scobit: An Alternative Estimator to Logit and Probit'. *American Journal of Political Science* 38(1):230–55.
- Norton, Edward C., Hua Wang, and Chunrong Ai. 2004. 'Computing Interaction Effects and Standard Errors in Logit and Probit Models'. *The Stata Journal* 4(2):154–67.
- Oneal, John R., and Bruce Russett. 2001. *Triangulating Peace: Democracy, Interdependence, and International Organizations*. New York: W. W. Norton.
- Ray, James Lee. 1998. 'Does Democracy Cause Peace?'. *Annual Review of Political Science* 1(1):27–46.
- Tomz, Michael R., and Jessica L. P. Weeks. 2013. 'Public Opinion and the Democratic Peace'. *American Political Science Review* 107(4):849–65.
- Train, Kenneth E. 2009. *Discrete Choice Models with Simulation*. New York: Cambridge University Press.
- Tsai, Tsung-han, and Jeff Gill. 2013. 'Interaction in General Linear Models: Theoretical Issues and An Application to Personal Vote Earning Attributes'. *Social Sciences* 2(2):91–113.
- Wolfinger, Raymond E., and Steven J. Rosenstone. 1980. *Who Votes?* New Haven: Yale University Press.